

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

1-2021

Attribute-aware pedestrian detection in a crowd

Jialiang ZHANG

Lixiang LIN

Jianke ZHU

Yang LI

Yun-chen CHEN

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#), and the [Numerical Analysis and Scientific Computing Commons](#)

Citation

ZHANG, Jialiang; LIN, Lixiang; ZHU, Jianke; LI, Yang; CHEN, Yun-chen; HU, Yao; and HOI, Steven C. H.. Attribute-aware pedestrian detection in a crowd. (2021). *IEEE Transactions on Multimedia*. 23, 3085-3097. Available at: https://ink.library.smu.edu.sg/sis_research/6967

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Jialiang ZHANG, Lixiang LIN, Jianke ZHU, Yang LI, Yun-chen CHEN, Yao HU, and Steven C. H. HOI

Attribute-aware Pedestrian Detection in a Crowd

Jialiang Zhang, Lixiang Lin, Yang Li, Yun-chen Chen, Jianke Zhu *Senior Member, IEEE*,
Yao Hu, Steven C.H. Hoi *Fellow, IEEE*

Pedestrian detection is an initial step to perform outdoor scene analysis, which plays an essential role in many real-world applications. Although having enjoyed the merits of deep learning frameworks from the generic object detectors, pedestrian detection is still a very challenging task due to heavy occlusion and highly crowded group. Generally, the conventional detectors are unable to differentiate individuals from each other effectively under such a dense environment. To tackle this critical problem, we propose an attribute-aware pedestrian detector to explicitly model people’s semantic attributes in a high-level feature detection fashion. Besides the typical semantic features, center position, target’s scale and offset, we introduce a pedestrian-oriented attribute feature to encode the high-level semantic differences among the crowd. Moreover, a novel attribute-feature-based Non-Maximum Suppression (NMS) is proposed to distinguish the person from a highly overlapped group by adaptively rejecting the false-positive results in a very crowd settings. Furthermore, a novel ground truth target is designed to alleviate the difficulties caused by the attribute configuration and extremely class imbalance issue during training. Finally, we evaluate our proposed attribute-aware pedestrian detector on two benchmark datasets including CityPersons and CrowdHuman. The experimental results show that our approach outperforms state-of-the-art methods at a large margin on pedestrian detection.

Index Terms—Pedestrian detection, Attribute-aware, Non-Maximum Suppression (NMS)

I. INTRODUCTION

Pedestrian detection is an initial step to perform outdoor scene analysis, which plays an essential role in many real-world applications, such as autonomous vehicles, security surveillance, robotics and so on.

Being a special case of generic object detection, pedestrian detectors [2], [16], [46] inherit a lot of successful techniques from the conventional anchor-based approaches, such as Faster RCNN [44] and SSD [21]. Most recently, anchor-free detectors like CornerNet [14] and CenterNet [51], are emerging as the promising methods that are able to enjoy the merits of flexibility in network model design. To this end, we present a pedestrian detector under the anchor-free object detection framework.

Comparing to generic object detection, the pedestrian detector has its own characteristics and challenges in the real-world,

Jialiang Zhang, Lixiang Lin, Yang Li, and Jianke Zhu are with the College of Computer Science, Zhejiang University, Hangzhou, China, 310027. E-mail: {zjialiang, lxlin liyang89, jkzhu}@zju.edu.cn.

Yun-chen Chen and Steven C.H. Hoi are with School of Information Systems, Singapore Management University, Singapore. E-mail: {jeanchen, chhoi}@smu.edu.sg

Yao Hu is with company of Alibaba, Youku Cognitive and Intelligent Lab. E-mail: yaohu@alibaba-inc.com

Jianke Zhu is the Corresponding Author.




Fig. 1: Illustration of our attribute-aware pedestrian detector. (a) The predicted bounding boxes of detector [23] without NMS. (b) The bounding boxes using our method colored by a 3-D vector that is mapped from the attribute vector through embedding. (c) The result of detector [23] after greedy-NMS. (d) Our result after our proposed attribute-aware NMS.

such as vastly different scales, poor appearance conditions, and extremely severe occlusion in crowd scenarios.

To localize a visual object, a modern detector usually produce a large amount of overlapped proposals around the target, and then employ the Non-Maximum Suppression (NMS) to select the most possible bounding boxes based on likelihood scores. As the target usually stands without highly overlapped cases, this pipeline works well in the general object detection scenario. However, occlusions commonly occur in pedestrian detection. There are 48.8% annotated pedestrians occluded by each other in the CityPersons dataset [46]. In addition, Wang *et al.* [42] found that the level of crowd occlusion is even higher in the CrowdHuman dataset [35]. Usually, it is hard for the conventional pipeline to chose the bounding boxes in a crowd setting. As shown in Fig. 1, it is challenging to distinguish the bounding boxes generated by a single target or those from multiple people occluded together using a single filtering threshold in the greedy NMS. This is because a lower threshold may lead to missing highly overlapped objects while a higher one often has more false positives.

To deal with pedestrian detection in a crowd, recent methods either design the proper regression losses for more compact boxes generation [42], [48] or develop the post-processing NMS technique. In fact, NMS is greatly affected by the feature extraction, and they should be dealt with in a whole pipeline.

To address the above limitations, we propose a novel attribute-aware pedestrian detector by tackling the crowd occlusion issue from the original feature space side in this paper. Inspired by a human being’s visual perception process, our

algorithm tries to understand the crowd scene as a whole and encode high-level semantic information for each individual rather than alleviating ambiguities from the last separate component like NMS in the detector. To this end, we formulate the pedestrian detection problem as a feature detection task as in [23], and explicitly model each people’s semantic attribute as an embedding vector in a high-level feature space. Combing the basic semantic features including center position, target’s scale and offset, our proposed attribute-aware feature encodes the semantic differences among the persons in a crowd as shown in Fig. 1. Moreover, a density label and a diversity metric are generated as ground truth targets for the attribute-aware training stage. With the very rich semantic information of the proposed attribute feature, a novel attribute-aware NMS is introduced to adaptively reject the false-positive proposals and distinguish the individuals in a very crowd setting.

In summary, the main contributions of this paper are: 1) a novel attribute map that converts each positive instance into a feature vector to encode both density and diversity information simultaneously; 2) an enhanced ground truth target to define attribute loss and address extremely class imbalance issue during training; 3) a novel attribute-aware NMS algorithm by considering the high-level semantic information of each predicted box to effectively refine the detection results; 4) experiments on two standard benchmark datasets [46], [35] show that our proposed approach outperforms state-of-the-art pedestrian detectors at a large margin.

II. RELATED WORK

As pedestrian detection can be viewed as the object localization for the special category, we firstly briefly review the related works on generic object detection. Then, we introduce the recent progress on it. Moreover, we will investigate the techniques related to solving the detection in a crowd, including occlusion handling and Non-Maximum Suppression.

A. Generic object detection

Early generic object detectors are based on the hand-crafted features [4], [6], [28], which aim at finding and classifying the candidates in a sliding window paradigm or region proposal strategy. With the development of deep learning, the modern detectors can be categorized into two classes, anchor-based methods, and anchor-free approaches.

The anchor-based object detectors have dominated in past years. The two-stage detectors including Fast-RCNN [7], Faster-RCNN [32] and their variants [17] generate the object proposals in the first stage and classify them in the second stage. On the other hand, the single-stage methods like SSD [21], YOLOv2 [31] manage to perform detection and classification simultaneously using the feature maps without a separate proposal generation stage.

Recently, anchor-free methods achieve promising results, which directly formulate the detection problem as a regression task having a simple network structure without anchor boxes. CornerNet [14] treats object detection as keypoint detections and their associations, which are free from windows and anchors. CenterNet [51] views an object as a single point,

which is the center point of its bounding box. Then, these center points regress to all other object properties, such as size, 3D location, orientation, and even pose. Our proposed method essentially belongs to the anchor-free detection framework.

B. Pedestrian Detection

Traditional pedestrian detectors [4], [26], [45] rely on integral channel features with sliding window strategy to localize each target. CNN-based detectors are mainly based on Faster RCNN framework. In [44], RPN is employed to generate proposals and provide CNN features followed by a boosted decision forest. Zhang *et al.* [46] applies a plain Faster-RCNN for pedestrian detection. Cai *et al.* [2] employ multi-scale feature maps to match pedestrians of different scales. Mao *et al.* [24] use extra features to further improve performance. Song *et al.* [36] propose to detect the object by predicting the top and bottom vertexes. Liu *et al.* [22] present the asymptotic localization fitting strategy to gradually refine the localization results. Later, they suggest an anchor-free framework to detect center and scale in pedestrian detection [23].

C. Occlusion Handling

Occlusions are quite common for pedestrian detection in a crowd. CityPersons [47] and CrowdHuman datasets [35] are collected especially for crowd scenarios. Part-based models [5], [25], [27], [49], [50], [38], [41] are usually used to deal with occlusions. They learn a series of part detectors, and then design some mechanisms to fuse each result to localize partially occluded pedestrians. In [39], attention mechanism is used to represent occlusion patterns. RepLoss [42] and OR-CNN [48] design two novel regression losses to generate more compact boxes to make them less sensitive to the NMS threshold in crowded scenes.

D. Non-Maximum Suppression

NMS is a critical post-processing step for object detectors. Greedy-NMS suffers from suppressing objects in crowd scenarios. Instead of discarding predicted boxes during suppressing, soft-NMS [1] decreases the scores of neighbors by an increasing function of their overlap with the higher scored bounding box. Learning NMS [9] trains a deep neural network to perform the NMS function using the predicted boxes and their corresponding scores. The relation module is employed to learn the NMS function in [10] as an end-to-end general object detector. Tychsen *et al.* [40] and Jiang *et al.* [12] learn extra localization confidences to guide a better NMS. Idrees *et al.* [11] and Ranjan *et al.* [30] estimate a crowd density map in people counting task. Rujikiatgumjorn *et al.* [33] propose a quadratic unconstrained binary optimization solution to suppress detection boxes, where a hard threshold is still applied to blindly suppress detection boxes as greedy-NMS. Lee *et al.* [15] combine the determinantal point process with individual prediction scores to optimally select the final detection boxes. Adaptive-NMS [20] proposes to estimate the density of predicted boxes to set an adaptive threshold. Despite promising results, its NMS threshold is merely determined by




Fig. 2: Overall architecture of our proposed detector. Detection is performed on Φ_{det} with the down-sampling rate $r = 4$. Our detector consists of four detection heads. Attribute map encodes the density and diversity information of each positive instance into an embedding vector. The length d of embedding vector denotes the density while the angle θ between two of them indicates their similarity. Attribute-aware NMS refines the predicted boxes by taking into account of both density and diversity information from attribute map. The bottom left is our model diagram for (a) ResNet-50 and (b) DLA-34.

the density of the current box with the largest score, which ignores the relationship among boxes. Thus, selecting the bounding boxes is still the most critical issue for pedestrian detection.

III. ATTRIBUTE-AWARE PEDESTRIAN DETECTION

In this section, we first give the overall architecture of the proposed attribute-aware pedestrian detector. To explicitly distinguish a pedestrian from feature representation, four high-level semantic branches are proposed in our network structure. Then, we describe the details of our optimization targets design for conventional detection branches, center, scale, and offset maps, which consists of the definition of ground truth targets as well as the objective functions. Moreover, a novel attribute loss is presented for training our proposed attribute map to encode both density and diversity information. Finally, we introduce an attribute-aware NMS algorithm to fully exploit our proposed attribute-aware structure.

A. Attribute-aware Pedestrian Detector

Under the anchor-free detection framework [23], we introduce four branches into our proposed detector to describe the pedestrian. Each branch produces a feature map to represent a property of the pedestrian in the image. The overall architecture is illustrated in Fig. 2. The backbone is truncated

from a standard network pre-trained on ImageNet [34], such as ResNet-50 [8] and deep layer aggregation network DLA-34 [43].

1) Feature Extraction

There is a total of five stages in ResNet-50 and DLA-34 with the down-sampling rate 2, 4, 8, 16 and 32, respectively. We incorporate semantics from the last four feature maps, in which shallower feature maps provide more precise localization information while deeper feature maps contain more semantic features. Specifically, we augment ResNet-50 with a standard FPN structure [18], where an up-sampling network allows for a higher-resolution output and skip connections for richer semantics. DLA-34 is an image classification network that employs iterative deep hierarchical aggregation to enhance the feature representation. We augment the skip connections with deformable convolution from lower layers to the output as CenterNet [51]. The detailed model diagrams are depicted in Fig. 2. For the last stage, the dilated convolution is adopted to keep the spatial feature size as $1/16$ of the input image. Final detection is performed on an aggregated feature map Φ_{det} with the down-sampling rate $r = 4$.

2) Detection Heads

We introduce four detection heads, which predict four high-level semantic features, including center, scale, offset and attribute map, respectively. All four predicted feature maps are with the same size as Φ_{det} that is $H/r \times W/r$, where




Fig. 3: Illustration for center and offset target. Circle point is the real center $(\frac{x_k}{r}, \frac{y_k}{r})$, while square boxes denote four points next to it. Red indicates positive, while grey is negative. The dashed line denotes the offset, direction from positives to the real center.

input image $I \in \mathcal{R}^{W \times H \times 3}$. Each of them is appended after the aggregated feature maps Φ_{det} . Specifically, we attach one 3×3 convolution layer with 256 channels to Φ_{det} . A final 1×1 convolution is appended to produce the desired output.

The center map aims at localizing the central position of positive instances while the scale map predicts the height and width of them. The offset map is to recover the quantization error caused by the down-sampling rate r for defining the location of positive instances. The attribute map is designed for converting each positive instance into an embedding vector which encodes both diversity and density information simultaneously. We fuse the center map, scale map, and offset map to produce the predicted bounding boxes. These boxes and attribute map are further fed into attribute-aware NMS algorithm in order to effectively produce the refined boxes.

B. Optimization Target

In the following, we present the optimization target for the conventional detection branches, including center, scale, and offset maps, which consists of the definition for ground truth targets as well as the objective functions.

1) Center Loss

The real location of an object's center point is $(\frac{x_k}{r}, \frac{y_k}{r})$. Instead of assigning a single position $(\lfloor \frac{x_k}{r} \rfloor, \lfloor \frac{y_k}{r} \rfloor)$ as positive in [23], we assign all four corners next to the real location as positives while all others as negatives for the center target $\mathcal{C} \in \mathcal{R}^{\frac{W}{r} \times \frac{H}{r} \times 1}$.

As in [23], predicting the center map can be formulated as a classification task with the supervision of center target \mathcal{C} via a penalty-reduced pixel-wise logistic regression with focal loss [19] as below:

$$L_c = -\frac{1}{K} \sum_{i=1}^{W/r} \sum_{j=1}^{H/r} \alpha_{ij} (1 - \hat{p}_{ij})^\gamma \log(\hat{p}_{ij}), \quad (1)$$

where K in Eq. 1 is the number of positive instances in an image. Let $p_{ij} \in [0, 1]$ denote the predicted probability indicating whether the location (i, j) is a positive instance or

not. \hat{p}_{ij} and α_{ij} are defined as follows:

$$\hat{p}_{ij} = \begin{cases} p_{ij} & \text{if } y_{ij} = 1 \\ 1 - p_{ij} & \text{otherwise,} \end{cases} \quad (2)$$

$$\alpha_{ij} = \begin{cases} 1 & \text{if } y_{ij} = 1 \\ (1 - M_{ij})^\beta & \text{otherwise.} \end{cases}$$

where $y_{ij} = 1$ is the ground truth label representing the positive location. To deal with the ambiguity from those negatives surrounding the positives, α_{ij} according to the Gaussian mask M is applied to reduce their contributions to the total loss. As in [14], Hyper-parameters γ and β are set to 2 and 4, respectively.

2) Scale Loss

As for the scale target $\mathcal{S} \in \mathcal{R}^{\frac{W}{r} \times \frac{H}{r} \times s}$, $s = 1$ if only height is predicted. In CrowdHuman dataset, the aspect ratio is not fixed. s is set to 2 so that both height and width can be predicted. For simplicity, we only consider the height prediction in the following illustration. we assign $\log(h_k)$ to the k -th positive regions, which is assigned to the negatives with an extension of one pixel to the positive region with the size of 4×4 .

The prediction of scale map is formulated as a regression task with supervision \mathcal{S} via SmoothL1 loss [7] as follows:

$$L_s = \frac{1}{K} \sum_{k=1}^K \Phi(\hat{s}_k, s_k), \quad (3)$$

where Φ is SmoothL1 loss, \hat{s}_k, s_k represent the predicted height and the ground truth of each positive instance, respectively. The supervision acts only at the positions that have been assigned with scale values. All other locations are ignored.

3) Offset Loss

For offset target $\mathcal{O} \in \mathcal{R}^{\frac{W}{r} \times \frac{H}{r} \times 2}$, we estimate the offset for each corner towards the real center in four directions, which means that offsets are $(\frac{x_k}{r} - \lfloor \frac{x_k}{r} \rfloor, \frac{y_k}{r} - \lfloor \frac{y_k}{r} \rfloor)$, $(\frac{x_k}{r} - \lceil \frac{x_k}{r} \rceil, \frac{y_k}{r} - \lfloor \frac{y_k}{r} \rfloor)$, $(\frac{x_k}{r} - \lfloor \frac{x_k}{r} \rfloor, \frac{y_k}{r} - \lceil \frac{y_k}{r} \rceil)$ and $(\frac{x_k}{r} - \lceil \frac{x_k}{r} \rceil, \frac{y_k}{r} - \lceil \frac{y_k}{r} \rceil)$ for top-left, top-right, bottom-left and bottom-right point, respectively. Fig. 3 illustrates our design.

The branch of offset map is also formulated as a regression task with the supervision of \mathcal{O} via SmoothL1 loss as below:

$$L_o = \frac{1}{K} \sum_{k=1}^K \Phi(\hat{o}_k, o_k), \quad (4)$$

where \hat{o}_k and o_k represent the predicted offset and the ground truth of each positive instance, respectively. Similar to the scale map, the supervision acts only at locations that have been assigned with offset values and all other locations are ignored.

The aforementioned ground truth targets enjoy several key advantages: 1) There is no apparent boundary between four corners. In contrast to the typical design, a detector may predict the other three corners as center as well. However, the offset direction is always bottom-right which makes the localization less precise. Our design regresses offsets in all four directions towards the real center making the localization more precise; 2) By treating a 2×2 region as positives, we increase the positive instances three times larger than the typical design that eases the extremely-class-imbalance issue

during training to some extent; 3) Our design assigns more than one positive location to each ground truth bounding box, which enables us to design an effective attribute loss described in the following.

C. Attribute Loss

Besides the conventional loss for an anchor-free detector, we propose an attribute-aware loss to enforce the neural network to catch a high-level understanding of the crowd. The proposed loss consists of two different types of attribute losses, density loss L_{den} and diversity loss L_{div} , with a weighted sum as below:

$$L_a = \lambda_{den}L_{den} + L_{div} \quad (5)$$

where λ_{den} is set to 5 in our experiment. In this work, we aim to enable the neural networks to capture the high-level semantic information for a crowd. To this end, two attributes are proposed to describe a group of people, the density of people in a crowd and the diversity of a group. Given the predicted bounding box, the density attribute is the maximum IoU between the current bounding box and its surrounding ones. It gives hints to the proposed neural networks that a dense crowd might contain a number of people highly occluded from each other. Moreover, the diversity attribute is a metric for the algorithm to distinguish individuals among a crowd. With less diversity, a crowd should contain fewer people and the proposed detector outputs less bounding boxes.

To achieve this, we generate a density target $\mathcal{D} \in \mathcal{R}^{\frac{W}{r} \times \frac{H}{r} \times 1}$ as the ground truth of density information. Inspired by [20], for each position that the center of a ground truth box b_i located on the \mathcal{D} , the value is defined as the maximum Intersection-over-Union (IoU) between the ground truth bounding box b_i and other nearby boxes in the image as follows

$$d_i := \max_{b_j \in \mathcal{G}, i \neq j} \text{IoU}(b_i, b_j), \quad (6)$$

where \mathcal{G} is the ground truth bounding box set of an image. As for diversity information, we implicitly encode it by proposing an attribute loss instead of designing a diversity ground truth target. More specifically, both density and diversity information are encoded as an m -dimension vector $e \in \mathcal{R}^m$ embedded in each position of the attribute map. With four locations labeled as the positive samples in center target, $N_p = 4$ embedding vectors $e_{k,i}, i = 1, \dots, N_p$ are used to represent the attribute of an object. We employ the norm of embedding vector as the density information. Estimation is formulated as a regression task via SmoothL1 loss as below

$$L_{den} = \frac{1}{N} \sum_{k=1}^N \sum_{i=1}^{N_p} \Phi(\|e_{k,i}\|_2, d_k) \quad (7)$$

where $\|e_{k,i}\|_2$ is the norm of $e_{k,i}$, and d_k is the ground truth density of k -th object. In this loss function, we aim at regressing the norm of $e_{k,i}$ to the density attribute d_k . The supervision happens only at positive locations and all other positions are ignored.

Let $e_{k,i}^*$ and e_k^* denote the normalized embedding vectors and its mean from the k -th object respectively as follows

$$e_{k,i}^* = \frac{e_{k,i}}{\|e_{k,i}\|_2}, \quad (8)$$

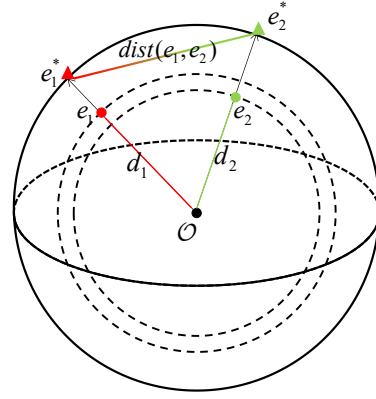


Fig. 4: A sphere \mathcal{O} with radius 1. e_1 and e_2 are embedding vectors and e_1^* and e_2^* are the corresponding normalized embedding vectors. The length of red and green lines d_1 and d_2 denote the density while the changed color line denotes the distance of two embedding vectors. Vectors with similar density can still have large distance.

$$e_k^* = \frac{1}{N_p} \sum_{k=1}^{N_p} e_{k,i}^*. \quad (9)$$

Based on above notations, our proposed diversity loss can be formulated as below:

$$L_{div} = L_1 + L_2 \quad (10)$$

$$L_1 = \sum_{k=1}^N \sum_{j=1, j \neq k}^N \max\left(0, \Delta - \|e_k^* - e_j^*\|_1\right), \quad (11)$$

$$L_2 = \frac{1}{N} \sum_{k=1}^N \sum_{i=1}^{N_p} \|e_{k,i}^* - e_k^*\|_2^2, \quad (12)$$

The hyper-parameter Δ is set to 1. The diversity of an object is defined from two perspectives, dispersion L_1 and gathering L_2 . Dispersion means our embedding vectors of different objects should be apart in the vector space. L_1 loss punishes two vectors from the different identities e_k^* and e_j^* for being close. Meanwhile, gathering L_2 tries to attract embedding vectors that belong to the same identity since we have N_p embedding vectors to represent an object.

Note that a straightforward solution is to design two semantic feature maps. One is for the density map as in [20] trained with L_{den} , and the other is for the diversity map trained with L_{div} , separately. However, we find that the overall network tends to be harder to optimize with more branches attached, which leads to deteriorating detection performance. Besides, a single branch of attribute map is much more elegant. As a matter of fact, we can not train an attribute map with dimension $m = 1$ as in [20] to encode both density and diversity information. The paradox lies in that, there are two predicted bounding boxes b_1, b_2 having the largest IoU with each other, denoted as d . Our objective function targets density

prediction at value d for both b_1 and b_2 , thus with a metric distance 0. However, the diversity prediction requires b_1, b_2 to have a large metric distance.

We show that with embedding vector dimension $m(> 1)$ of an attribute map, these two properties can be held simultaneously. Taking an embedding space of dimension $m = 3$ as an example, the embedding vector will fall in the sphere of radius 1 since the density is a value between 0 and 1. In our design, the density is defined as the length of vector while the diversity is implicitly trained by restricting Euclidean distance between the normalized embedding vectors. Density loss L_{den} is defined on the length of each vector (e_1, e_2) while diversity loss L_{div} is defined on the distance between the paired vectors. As shown in Fig. 4, two embedding vectors with similar density, e_1 and e_2 can still have a large distance, $dist(e_1, e_2)$. In our experiment, we set $m = 4$ if not specified.

Finally, we give a brief proof that larger Euclidean distance between normalized embedding vectors e_1^*, e_2^* is equivalent to larger angle between embedding vectors e_1, e_2 , and vice versa.

$$\begin{aligned} \|e_k^* - e_j^*\|_2^2 &= \|e_k^*\|_2^2 + \|e_j^*\|_2^2 - 2\|e_k^*\|_2\|e_j^*\|_2\cos(\theta) \\ &= 2 - 2\cos(\theta), \end{aligned} \quad (13)$$

where $\theta = \langle e_k, e_j \rangle$ is the angle between e_k and e_j . In a geometrical view, our diversity loss actually restricts the angle between embedding vectors. In summary, training with our attribute loss L_a , the length d of embedding vector can denote the density of each positive instance while the angle θ between two vectors can indicate the similarity of two predicted boxes.

The overall objective function can be derived as follows

$$L = \lambda_c L_c + \lambda_s L_s + \lambda_o L_o + \lambda_a L_a, \quad (14)$$

where $\lambda_c, \lambda_s, \lambda_o$ and λ_a are empirically set to 0.01, 1, 0.03 and 0.01, respectively.

D. Attribute-aware NMS Algorithm

NMS is a critical post-processing step for object detectors. In this section, we first revisit greedy-NMS, and then introduce a density-aware NMS [20]. Moreover, we propose a diversity-aware NMS algorithm to utilize the encoded diversity information from attribute map. To fully exploit our presented attribute map, we propose an attribute-aware NMS algorithm to refine the predicted bounding boxes more effectively by incorporating both density and diversity information.

1) Greedy-NMS

Given a set of predicted bounding boxes \mathcal{B} with their corresponding confidence scores \mathcal{S} , greedy-NMS firstly selects the one $\mathcal{M} \in \mathcal{B}$ with the maximum score and all boxes in \mathcal{B} which have the overlap with \mathcal{M} larger than the threshold N_t are removed. The algorithm continues to select the next bounding box in the remaining set of \mathcal{B} and repeats the process until the end. The suppressing step is as follows:

$$s_i = \begin{cases} s_i, & \text{IoU}(\mathcal{M}, b_i) < N_t \\ 0, & \text{IoU}(\mathcal{M}, b_i) \geq N_t \end{cases},$$

where \mathcal{M} is the bounding box with highest score in \mathcal{B} and s_i is the confidence score of a box b_i . IoU represents the



Fig. 5: Illustrative examples for different NMS algorithm. There are three pedestrians in each picture. The first column is the result of greedy-NMS, the mid column is density-aware NMS, and the last column is our attribute-aware NMS algorithm. The missed bounding box is colored by blue, and the false positive is denoted as red. There are three pedestrians in each picture, greedy

intersection over union score. N_t is a constant threshold that is typically set to 0.5.

2) Density-aware NMS

Similar to [20], the density-aware NMS is designed to incorporate the density information, which is able to adaptively adjust the threshold N_t by the density of each bounding box. The suppressing step of density-aware NMS can be described as below

$$\begin{aligned} N_{\mathcal{M}} &:= \max(N_t, d_{\mathcal{M}}), \\ s_i &= \begin{cases} s_i, & \text{IoU}(\mathcal{M}, b_i) < N_{\mathcal{M}} \\ 0, & \text{IoU}(\mathcal{M}, b_i) \geq N_{\mathcal{M}} \end{cases}, \end{aligned} \quad (15)$$

where $d_{\mathcal{M}}$ is the density of bounding box \mathcal{M} defined in Eq. 6. $N_{\mathcal{M}}$ denotes the threshold for \mathcal{M} .

We argue that the purpose of NMS algorithm is to suppress the predicted bounding boxes that belong to the same identity, which preserves the predicted bounding boxes from different identities. However, neither greedy-NMS nor density-aware NMS addresses these two problems at the same time. While greedy-NMS sets the threshold blindly, density-aware NMS may work well in keeping predicted boxes in crowd scenes, however, it also preserves more boxes belong to the same identity thus leading to more false positives. An illustrative example is shown in Fig. 5.

Algorithm 1: Attribute-aware NMS algorithm

Input:
 $\mathcal{B} = b_1, \dots, b_N, \mathcal{S} = s_1, \dots, s_N$
 $D = d_1, \dots, d_N, \epsilon = \epsilon_1^*, \dots, \epsilon_N^*$
 N_t, δ_t
 \mathcal{B} is the list of initial detection boxes
 \mathcal{S} contains corresponding detection scores
 \mathcal{D} contains corresponding detection densities
 ϵ contains corresponding normalized embedding vectors
 N_t is the NMS threshold
 δ_t is the distance threshold

```

1 begin
2    $\mathcal{F} \leftarrow \{\}$ 
3   while  $\mathcal{B} \neq \text{empty}$  do
4      $m \leftarrow \text{argmax } \mathcal{S}$ 
5      $\mathcal{M} \leftarrow b_m$ 
6      $N_{\mathcal{M}} \leftarrow \text{max}(N_t, d_m)$ 
7      $\mathcal{F} \leftarrow \mathcal{F} \cup \mathcal{M}; \mathcal{B} \leftarrow \mathcal{B} - \mathcal{M}$ 
8     for  $b_i$  in  $\mathcal{B}$  do
9       if  $\|e_m^* - e_i^*\| > \delta_t$  then
10        |  $N_{\mathcal{M},b_i} = \text{max}(N_d, d_m)$ ;
11       else
12        |  $N_{\mathcal{M},b_i} = N_t$ 
13       end
14       if  $\text{iou} \geq N_{\mathcal{M},b_i}$  then
15        |  $\mathcal{B} \leftarrow \mathcal{B} - b_i; \mathcal{S} \leftarrow \mathcal{S} - s_i$ 
16       end
17     end
18   end
19   return  $\mathcal{F}, \mathcal{S}$ 
20 end

```

Fig. 6: The pseudo code for our attribute-aware NMS. The difference from greedy-NMS is highlighted as blue.

3) Attribute-aware NMS and Diversity-aware NMS

To incorporate diversity information, we present a diversity-aware NMS algorithm which takes the similarity between two compared bounding boxes into consideration. The suppressing step of diversity-aware NMS is defined as below

$$N_{\mathcal{M},b_i} := \begin{cases} N_t^{\text{high}}, & \text{dist}(\mathcal{M}, b_i) > \delta_t \\ N_t^{\text{low}}, & \text{dist}(\mathcal{M}, b_i) \leq \delta_t \end{cases}, \quad (16)$$

$$s_i = \begin{cases} s_i, & \text{IoU}(\mathcal{M}, b_i) < N_{\mathcal{M},b_i} \\ 0, & \text{IoU}(\mathcal{M}, b_i) \geq N_{\mathcal{M},b_i} \end{cases},$$

$N_t^{\text{high}}, N_t^{\text{low}}$ are constant thresholds, which are empirically set to 0.6 and 0.5. $\text{dist}(\cdot)$ denotes a distance function to measure the diversity between two bounding boxes. The distance threshold δ_t is set to 0.9. Diversity-aware NMS also has trouble setting distance threshold and constant N_t^{high} in real complicated crowd scenarios.

To fully exploit our attribute map, we propose a novel post-processing algorithm called attribute-aware NMS by taking both density and diversity information of each box into account. The suppressing step of attribute-aware NMS is defined

as follows

$$N_{\mathcal{M},b_i} := f(d_{\mathcal{M}}, \text{dist}(\mathcal{M}, b_i); N_t, \delta_t), \quad (17)$$

$$f(d, \text{dist}; N_t, \delta_t) = \begin{cases} \max(d, N_t), & \text{dist} > \delta_t \\ N_t, & \text{dist} \leq \delta_t \end{cases},$$

$$s_i = \begin{cases} s_i, & \text{IoU}(\mathcal{M}, b_i) < N_{\mathcal{M},b_i} \\ 0, & \text{IoU}(\mathcal{M}, b_i) \geq N_{\mathcal{M},b_i} \end{cases},$$

where $d_{\mathcal{M}}$ represents the density of box \mathcal{M} . $\text{dist}(\cdot)$ is a distance function that accounts for the diversity information between two predicted boxes. Note that the threshold $N_{\mathcal{M},b_i}$ is not only a function of \mathcal{M} but also related to the box b_i . We discuss the properties of our designed threshold function $N_{\mathcal{M},b_i}$ and suppression strategy in the following.

- $N_{\mathcal{M},b_i}$ is high if and only if when $d_{\mathcal{M}}$ is large (\mathcal{M} has a high density) and $\text{dist}(\mathcal{M}, b_i)$ is large (the two compared boxes \mathcal{M} and b_i belong to different identities).
- When \mathcal{M} is not in a crowd scene, $N_{\mathcal{M},b_i}$ will be low since $d_{\mathcal{M}}$ is small.
- When \mathcal{M} locates in the crowded region ($d_{\mathcal{M}}$ is high) and \mathcal{M}, b_i belong to the same identity ($\text{dist}(\mathcal{M}, b_i)$ is small), threshold $N_{\mathcal{M},b_i}$ will be decreased to suppress b_i thus reducing false positives.
- When \mathcal{M} locates in the crowded region and \mathcal{M} and b_i belong to different identities ($\text{dist}(\mathcal{M}, b_i)$ is large), $N_{\mathcal{M},b_i}$ will be high thus neighboring boxes which belong to different identities are preserved.

Thus, our design satisfies the two conditions mentioned before simultaneously. As shown in Fig. 5c, our attribute-aware NMS can address the issue in a crowd. The algorithm is formally described in Fig 6. During inference, the predicted density and $\text{dist}(\cdot)$ function is calculated as follows

$$\text{dist}(b_i, b_j) = \|e_i^* - e_j^*\|_2 \quad (18)$$

$$d_i = \|e_i\|_2 \quad (19)$$

where e_i and e_i^* denotes the embedding vector and the normalized embedding vector of the predicted box b_i , respectively.

IV. EXPERIMENTS

In this section, we first introduce our experimental settings and implementation details. Then, we conduct an ablative analysis for each component of our proposed pedestrian detector on Citypersons dataset [46]. Furthermore, we compare our approach with the state-of-the-art methods on both Citypersons and CrowdHuman datasets [35].

A. Experiment Setup

To evaluate the efficacy of our proposed pedestrian detector, we conduct experiments on the challenging benchmarks in a crowd, including Citypersons and CrowdHuman.

Datasets CityPersons [46] is a newly collected pedestrian detection dataset built on top of the CityScapes benchmark [3], which includes 5,000 images (2,975 for training, 500 for validation and 1,525 for testing). It is a challenging dataset

Method	Backbone	Reasonable	Heavy	Partial	Bare	Test Time
FRCNN [46]	VGG-16	15.4	-	-	-	-
FRCNN+Seg [46]	VGG-16	14.8	-	-	-	-
OR-CNN [48]	VGG-16	12.8	55.7	15.3	6.7	-
RepLoss [42]	ResNet-50	13.2	56.9	16.8	7.6	-
TLL [36]	ResNet-50	15.5	53.6	17.2	10.0	-
TLL+MRF [36]	ResNet-50	14.4	52.0	15.9	9.2	-
ALFNet [22]	ResNet-50	12.0	51.9	11.4	8.4	0.27s/img
WIDERPERSON [17]	VGG-16	11.1	-	-	-	-
CSP [23]	ResNet-50	11.0	49.3	10.4	7.3	0.33s/img
Adaptive-NMS [20]	ResNet-50	10.8	54.0	11.4	6.2	-
APD (ours)	ResNet-50	10.6	49.8	9.5	7.1	0.12s/img
APD (ours)	DLA-34	8.8	46.6	8.3	5.8	0.16s/img

TABLE I: Comparison with state-of-the-art methods on CityPersons validation set. All models are trained on the training set. Our result is tested on the original image size 1024×2048 . The reported testing time is under the same environment with the input size of 1024×2048 on Nvidia GTX 1080Ti GPU.

with various occlusion levels. We train our model on the official training set with 2975 images and test on the validation set with 500 images and test set with 1525 images. Crowd-Human dataset [35] has recently been released to specifically target to the crowd issue in the human detection task, which has 15,000, 4370, and 5,000 images from the Internet for training, validation, and test, respectively. It is of much higher crowdedness than CityPersons. We train on the training set and test on the validation set, where the full body region annotations are used for training and evaluation. In all the two datasets, we employ log-average Miss Rate over False Positive Per Image (FPPI) as the evaluation metric, which ranges in $[10^{-2}, 10^0]$ (denoted as MR^{-2}).

Training Details We implement our proposed method in Pytorch [29]. Adam [13] optimizer is adopted to optimize the network. Also, moving average weights [37] is used to achieve better and more stable training. The input sizes of training images are 640×1280 and 768×1152 for CityPersons and CrowdHumans, respectively. While resizing the inputs size, we keep the original aspect ratio for input images. For CityPersons, we optimize the network on two GPUs (GTX 1080Ti) with four images per GPU for a mini-batch. The learning rate is set to 1×10^{-4} . Training process is stopped after 37.5K iterations. For CrowdHuman, a mini-batch contains 16 images with four GPUs (GTX 1080Ti), and the learning rate is set to 1×10^{-4} . Training process is stopped after 100K iterations.

B. Ablation Study on CityPersons

We conduct an ablative analysis of the proposed method on CityPersons dataset. To this end, each component of our proposed detector is evaluated, including our new ground truth target design, our attribute map as well as attribute-aware NMS algorithm and the sensitivity of our embedding dimension m . The original image size 1024×2048 is used for testing.

1) Ground truth

We compare our ground truth targets design with typical configuration. Our design outperforms conventional method

Backbone	GT Targets	Reasonable	Heavy	Partial	Bare
ResNet-50		11.2	50.0	10.2	7.8
ResNet-50	✓	11.0	49.8	10.1	7.4
DLA-34		10.2	48.1	9.6	6.8
DLA-34	✓	9.5	47.4	8.6	6.6

TABLE II: Comparison to typical ground truth design. GT Targets imply our new design of ground truth targets. Our design performs consistently better than typical design in both ResNet-50 and DLA-34 architecture.

consistently in both ResNet-50 and DLA-34 architecture. It outperforms 0.2% in ResNet-50. Moreover, it increases to 0.7% with a more powerful backbone DLA-34. The performance gain is because our center target design assigns more locations to the positives which eases the extremely class imbalance issue during training, while the offset target design makes the localization more precise. Results are reported in Table II. If not specified, we use our new design in the following experiments.

Backbone	Attribute	Reasonable	Heavy	Partial	Bare
ResNet-50		11.0	49.8	10.1	7.4
ResNet-50	✓	10.6	49.8	9.5	7.1
DLA-34		9.5	47.4	8.6	6.6
DLA-34	✓	8.8	46.6	8.3	5.8

TABLE III: Comparison of our attribute-aware pedestrian detector with non-attribute version. Our detector performs consistently better in both ResNet-50 and DLA-34 architecture.

2) Attribute map v.s. attribute-aware NMS

We first evaluate the effectiveness of our attribute-aware pedestrian detector with both ResNet-50 and DLA-34 architecture. As shown in Fig. III, our detector outperforms 0.4% and 0.7% in ResNet-50 and DLA-34 architecture, respectively.

Backbone	Separate Branches		Attribute	Reasonable	Heavy	Partial	Bare
	Density	Diversity					
DLA-34				9.5	47.4	8.6	6.6
DLA-34	✓			9.2	46.5	8.7	5.9
DLA-34		✓		9.3	47.8	8.6	6.1
DLA-34	✓	✓		9.2	47.7	8.9	5.9
DLA-34			✓	8.8	46.6	8.3	5.8

TABLE IV: Comparison of our attribute-aware pedestrian detector with solely density prediction or diversity prediction or two separate branches for density and diversity prediction. Our attribute-aware pedestrian detector performs better than only density prediction or diversity prediction. Our attribute-aware NMS algorithm still takes effects when predicting density and diversity in two separate branches. However, more branches lead to deteriorate the overall detection performance.

We further conduct an experiment to compare our attribute-aware pedestrian detector with the following three settings: 1) a density map that predicts density information together with density-aware NMS which follows the strategy of adaptive-NMS [20]; 2) a diversity map that predicts diversity information together with diversity-aware NMS, and two separate branches, density map and diversity map, together with attribute-aware NMS.

As shown in Table IV, our attribute-aware pedestrian detector outperforms 0.7% than the default setting with greedy-NMS. Compared with density prediction with density-aware NMS or diversity prediction with diversity-aware NMS, our detector still outperforms 0.4% and 0.5%, respectively. As for predicting density and diversity in two separate branches with attribute-aware NMS, our NMS algorithm still outperforms 0.3% than the default setting with greedy-NMS. However, our attribute map that encodes both density and diversity information simultaneously outperforms 0.4% than two separate branches, which we argue that more detection heads deteriorate the detection performance, and this validates our superior design for attribute map.

Backbone	m	Reasonable	Heavy	Partial	Bare
DLA-34	2	9.2	46.6	8.6	6.1
DLA-34	3	9.0	47.0	8.2	5.9
DLA-34	4	8.8	46.6	8.3	5.8
DLA-34	8	9.0	47.1	8.4	5.9

TABLE V: Sensitivity of attribute-aware pedestrian detector to m . Embedding length $m = 4$ yields the best result.

3) Dimension m of embedding vector

We analyze the sensitivity of our attribute-aware pedestrian detector to the dimension m of embedding vector. We conduct experiments by varying m with 2, 3, 4, and 8. The best result is obtained when $m = 4$. However, the results are comparable for 3, 4 or 8, as shown in Table V.

C. Results on CityPersons

The proposed approach is extensively compared with state-of-the-art methods. Our presented attribute-aware pedestrian detector outperforms all previous methods at a large margin on both validation and the test set. Specifically, we achieve MR^{-2} of 8.8% on the Reasonable setting in validation set I, which is 2% better than the best competitor (10.8% of Adaptive-NMS). Our method also performs consistently better in all other three settings, 46.6%, 8.3% and 5.8% in Heavy, Partial and Bare, respectively. As shown in Table VI, we achieve MR^{-2} of 8.27% on the Reasonable setting in test set, which is 3.05% better than the best competitor (11.32% of OR-CNN). Our detector also outperforms consistently in Small, Heavy and All settings, 3.16%, 15.98%, 4.54%, respectively. Furthermore, the proposed approach runs faster than recent state-of-the-art method, which achieves 0.16s/img compared with 0.27s/img of ALF [22] and 0.33s/img of CSP [23].

Method	Backbone	Reasonable	Small	Heavy	All
FRCNN [46]	VGG-16	12.97	37.24	50.47	43.86
OR-CNN [48]	VGG-16	11.32	14.19	51.43	40.19
RepLoss [42]	ResNet-50	11.48	15.67	52.59	39.17
APD (ours)	DLA-34	8.27	11.03	35.45	35.65

TABLE VI: Comparisons with the state-of-the-art methods on CityPersons testing set. All models are trained on the training set. Our result is tested with the original image size of 1024 × 2048.

D. Results on CrowdHuman

We compared our proposed detector with the most recent state-of-the-art methods, including Adaptive-NMS [20] and FRCNN+FPN. In [20], the result of both single-stage and two-stage anchor-based methods are provided. As shown in Table VII, the proposed pedestrian detector with our new ground truth targets performs better than the competitor at a clear margin. Even compared with two-stage FPN, we still outperform a lot 36.79 compared with 49.73. Using the attribute-aware NMS algorithm, we can further improve our result from 36.79 to 35.76. It demonstrates the efficacy of our proposed detector.

Method	MR^{-2}
FRCNN+FPN [35]	50.42
RetinaNet [35]	63.33
FPN+adaptive-NMS [20]	49.73
RFB+adaptive-NMS [20]	63.03
APD w/o (ours)	36.79
APD (ours)	35.76

TABLE VII: Evaluation of full body detections on the CrowdHuman validation set. w/o denotes that attribute-aware NMS is not used and only our new ground truth targets are adopted.

V. CONCLUSIONS

In this paper, we proposed an attribute-aware pedestrian detector to explicitly capture the semantic attributes in a high-level feature detection manner. A pedestrian-oriented attribute map was employed to encode both density and diversity information of each identity simultaneously. To distinguish the pedestrian from a highly overlapped group, a novel attribute-aware Non-Maximum Suppression (NMS) was proposed to adaptively reject the false-positive results in a very crowd settings. Moreover, we designed the novel ground truth targets to alleviate the training difficulties caused by the attribute configuration and extremely class imbalance. Finally, we conducted the extensive experiments to demonstrate the efficacy of our proposed pedestrian detector. As a result, our approach outperformed the state-of-the-art methods at a large margin on both CityPersons and CrowdHuman datasets. For our future work, we plan to extend our approach for instance segmentation in crowd.

REFERENCES

- [1] N. Bodla, B. Singh, R. Chellappa, and L. S. Davis. Soft-nms—improving object detection with one line of code. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5561–5569, 2017.
- [2] Z. Cai, Q. Fan, R. S. Feris, and N. Vasconcelos. A unified multi-scale deep convolutional neural network for fast object detection. In *European Conference on Computer Vision*, pages 354–370. Springer, 2016.
- [3] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [4] P. Dollár, R. Appel, S. Belongie, and P. Perona. Fast feature pyramids for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(8):1532–1545, 2014.
- [5] M. Enzweiler, A. Eigenstetter, B. Schiele, and D. M. Gavrila. Multi-cue pedestrian classification with partial occlusion handling. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 990–997. IEEE, 2010.
- [6] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2009.
- [7] R. B. Girshick. Fast R-CNN. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 1440–1448, 2015.
- [8] 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*, pages 770–778, 2016.
- [9] J. Hosang, R. Benenson, and B. Schiele. Learning non-maximum suppression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4507–4515, 2017.
- [10] H. Hu, J. Gu, Z. Zhang, J. Dai, and Y. Wei. Relation networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3588–3597, 2018.
- [11] H. Idrees, M. Tayyab, K. Athrey, D. Zhang, S. Al-Maadeed, N. Rajpoot, and M. Shah. Composition loss for counting, density map estimation and localization in dense crowds. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 532–546, 2018.
- [12] B. Jiang, R. Luo, J. Mao, T. Xiao, and Y. Jiang. Acquisition of localization confidence for accurate object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784–799, 2018.
- [13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [14] H. Law and J. Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 734–750, 2018.
- [15] D. Lee, G. Cha, M.-H. Yang, and S. Oh. Individualness and determinantal point processes for pedestrian detection. In *European Conference on Computer Vision*, pages 330–346. Springer, 2016.
- [16] J. Li, X. Liang, S. Shen, T. Xu, J. Feng, and S. Yan. Scale-aware fast r-cnn for pedestrian detection. *IEEE Transactions on Multimedia*, 20(4):985–996, April 2018.
- [17] J. Li, Y. Wei, X. Liang, J. Dong, T. Xu, J. Feng, and S. Yan. Attentive contexts for object detection. *IEEE Transactions on Multimedia*, 19(5):944–954, May 2017.
- [18] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017.
- [19] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017.
- [20] S. Liu, D. Huang, and Y. Wang. Adaptive nms: Refining pedestrian detection in a crowd. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6459–6468, 2019.
- [21] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. E. Reed, C. Fu, and A. C. Berg. SSD: single shot multibox detector. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I*, pages 21–37, 2016.
- [22] W. Liu, S. Liao, W. Hu, X. Liang, and X. Chen. Learning efficient single-stage pedestrian detectors by asymptotic localization fitting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 618–634, 2018.
- [23] W. Liu, S. Liao, W. Ren, W. Hu, and Y. Yu. High-level semantic feature detection: A new perspective for pedestrian detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5187–5196, 2019.
- [24] J. Mao, T. Xiao, Y. Jiang, and Z. Cao. What can help pedestrian detection? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3127–3136, 2017.
- [25] M. Mathias, R. Benenson, R. Timofte, and L. Van Gool. Handling occlusions with franken-classifiers. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1505–1512, 2013.
- [26] W. Nam, P. Dollár, and J. H. Han. Local decorrelation for improved pedestrian detection. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 424–432, 2014.
- [27] W. Ouyang and X. Wang. A discriminative deep model for pedestrian detection with occlusion handling. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3258–3265. IEEE, 2012.
- [28] C. Papageorgiou and T. Poggio. A trainable system for object detection. *International Journal of Computer Vision*, 38(1):15–33, 2000.
- [29] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- [30] V. Ranjan, H. Le, and M. Hoai. Iterative crowd counting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 270–285, 2018.
- [31] J. Redmon and A. Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7263–7271, 2017.
- [32] 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*, pages 91–99, 2015.
- [33] S. Rujikietumjorn and R. T. Collins. Optimized pedestrian detection for multiple and occluded people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3690–3697, 2013.
- [34] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [35] S. Shao, Z. Zhao, B. Li, T. Xiao, G. Yu, X. Zhang, and J. Sun. Crowdhuman: A benchmark for detecting human in a crowd. *CoRR*, abs/1805.00123, 2018.
- [36] T. Song, L. Sun, D. Xie, H. Sun, and S. Pu. Small-scale pedestrian detection based on topological line localization and temporal feature aggregation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 536–551, 2018.
- [37] A. Tarvainen and H. Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep

- learning results. In *Advances in Neural Information Processing Systems*, pages 1195–1204, 2017.
- [38] Y. Tian, P. Luo, X. Wang, and X. Tang. Deep learning strong parts for pedestrian detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1904–1912, 2015.
- [39] Y. Tian, P. Luo, X. Wang, and X. Tang. Pedestrian detection aided by deep learning semantic tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5079–5087, 2015.
- [40] L. Tychsen-Smith and L. Petersson. Improving object localization with fitness nms and bounded iou loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6877–6885, 2018.
- [41] S. Wang, J. Cheng, H. Liu, F. Wang, and H. Zhou. Pedestrian detection via body part semantic and contextual information with dnn. *IEEE Transactions on Multimedia*, 20(11):3148–3159, Nov 2018.
- [42] X. Wang, T. Xiao, Y. Jiang, S. Shao, J. Sun, and C. Shen. Repulsion loss: Detecting pedestrians in a crowd. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7774–7783, 2018.
- [43] F. Yu, D. Wang, E. Shelhamer, and T. Darrell. Deep layer aggregation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2403–2412, 2018.
- [44] L. Zhang, L. Lin, X. Liang, and K. He. Is faster r-cnn doing well for pedestrian detection? In *European Conference on Computer Vision*, pages 443–457. Springer, 2016.
- [45] S. Zhang, R. Benenson, and B. Schiele. Filtered channel features for pedestrian detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, page 4, 2015.
- [46] S. Zhang, R. Benenson, and B. Schiele. Citypersons: A diverse dataset for pedestrian detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4457–4465, 2017.
- [47] S. Zhang, R. Benenson, and B. Schiele. Citypersons: A diverse dataset for pedestrian detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3221, 2017.
- [48] S. Zhang, L. Wen, X. Bian, Z. Lei, and S. Z. Li. Occlusion-aware r-cnn: detecting pedestrians in a crowd. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 637–653, 2018.
- [49] C. Zhou and J. Yuan. Multi-label learning of part detectors for heavily occluded pedestrian detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3486–3495, 2017.
- [50] C. Zhou and J. Yuan. Bi-box regression for pedestrian detection and occlusion estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 135–151, 2018.
- [51] X. Zhou, D. Wang, and P. Krähenbühl. Objects as points. *CoRR*, abs/1904.07850, 2019.