PSC-Net: Learning Part Spatial Co-occurrence for Occluded Pedestrian Detection

Jin Xie*, Yanwei Pang*, Hisham Cholakkal†, Rao Muhammad Anwer†, Fahad Shahbaz Khan†, and Ling Shao†
*Tianjin University †Inception Institute of Artificial Intelligence

Abstract—Detecting pedestrians, especially under heavy occlusions, is a challenging computer vision problem with numerous real-world applications. This paper introduces a novel approach, termed as PSC-Net, for occluded pedestrian detection. The proposed PSC-Net contains a dedicated module that is designed to explicitly capture both inter and intra-part co-occurrence information of different pedestrian body parts through a Graph Convolutional Network (GCN). Both inter and intra-part co-occurrence information contribute towards improving the feature representation for handling varying level of occlusions, ranging from partial to severe occlusions. Our PSC-Net exploits the topological structure of pedestrian and does not require part-based annotations or additional visible bounding-box (VBB) information to learn part spatial co-occurrence. Comprehensive experiments are performed on two challenging datasets: CityPersons and Caltech datasets. The proposed PSC-Net achieves state-of-the-art detection performance on both. On the heavy occluded (HO) set of CityPerosns test set, our PSC-Net obtains an absolute gain of 4.0% in terms of log-average miss rate over the state-of-the-art [34] with same backbone, input scale and without using additional VBB supervision. Further, PSC-Net improves the state-of-the-art [54] from 37.9 to 34.8 in terms of log-average miss rate on Caltech (HO) test set.

I. INTRODUCTION

PEDESTRIAN detection is a challenging problem in computer vision with various real-application applications, *e.g.*, robotics, autonomous driving and visual surveillance. Recent years have witnessed significant progress in the field of pedestrian detection, mainly due to the advances in deep convolutional neural networks (CNNs). Modern pedestrian detection methods can be generally divided into single-stage [22], [30], [3] and two-stage [20], [57], [54], [34], [52], [53], [4], [46], [24]. Single-stage pedestrian detectors typically work by directly regressing the default anchors into pedestrian detection boxes. Different to single-stage pedestrian detectors, two-stage methods first produce a set of candidate pedestrian proposals which is followed by classification and regression of these pedestrian proposals. Most existing two-stage pedestrian detectors [53], [34], [20], [52], [46], [24] are based on the popular Faster R-CNN detection framework [37] that is adapted from generic object detection. Existing pedestrian detectors typically assume entirely visible pedestrians when trained using full body pedestrian annotations.

While promising results have been achieved by existing pedestrian detectors on standard non-occluded pedestrians, their performance on heavily occluded pedestrians is far from satisfactory. This is evident from the fact that the best reported performance [34] on the reasonable (**R**) set (where visibility

ratio is larger than 65%) of CityPersons test set [51] is 9.3 (log-average miss rate) whereas it is 41.0 on the heavy occluded (**HO**) set (where visibility ratio ranges from 20% to 65%) of the same dataset. Handling pedestrian occlusion is an open problem in computer vision and present a great challenge for detecting pedestrians in real-world applications due to its frequent occurrence. Therefore, a pedestrian detector is desired to be accurate with respect to varying level of occlusions, ranging from reasonably occluded to severely occluded pedestrians.

A variety of occluded pedestrian detection approaches exist in literature. A common strategy to address occlusion is based on learning and integrating a set of part detectors [44], [56], [31], [25], [32], [26]. Earlier part-based pedestrian detection approaches utilize body part annotations [26], [27]. Recently, the deployment of pre-trained part models have been investigated [52] to exploit part correlation, typically relying on part detection scores corresponding to the visible regions of the pedestrian. Other methods [44], [56], [31], [25], [55] utilize the bounding-box of the pedestrian and train a large number of independently learned part detectors. Alternatively, the topological structure of the pedestrian has also been exploited [53] to avoid the reliance on body part annotations, leading to promising detection performance. However, it predominantly relies on the detection scores of parts to highlight visible regions of the pedestrian and do not consider spatial co-occurrence relations among different body parts *and* between different sub-regions (*e.g.*, eyes and ears of a head region) within a body part. The part spatial co-occurrence information is expected to enrich the feature representation by exploiting the information about the spatially adjacent parts. Knowledge about the typical configuration of objects (*e.g.*, humans) in a scene and its impact on recognition performance has been extensively studied previously in psychology and computer vision [2], [1], [11]. To the best of our knowledge, modern two-stage CNN-based pedestrian detectors do not explicitly encode the part spatial co-occurrence information. In this work, we introduce a data driven approach that goes beyond part detection scores and explicitly integrates part spatial co-occurrence information among different body parts *and* between different sub-regions within a body part to handle occlusions.

Contributions: We propose a two-stage approach, termed as PSC-Net, to address the problem of occluded pedestrian detection. Our main contribution is the introduction of a novel part spatial co-occurrence (PSC) module designed to explicitly capture both inter and intra-part co-occurrence information of different pedestrian body-parts through a Graph Convolutional
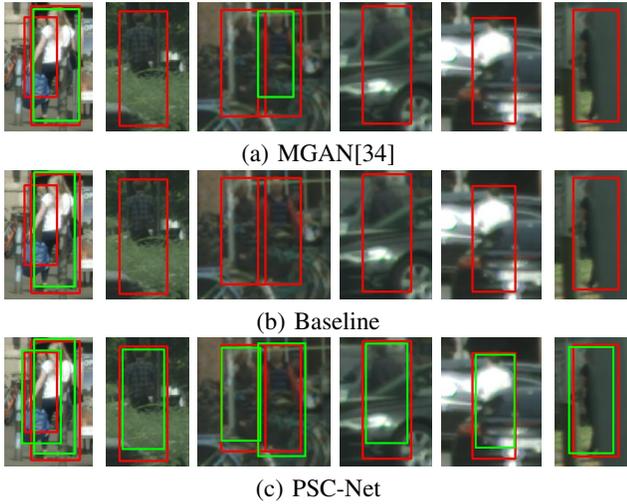
(a) MGAN[34]

(b) Baseline

(c) PSC-Net

Fig. 1: Qualitative detection examples using (a) the state-of-the-art MGAN [34], (b) our baseline and (c) our PSC-Net on CityPersons val. images. In the examples, red boxes denote the ground-truth and detector predictions are indicated by green boxes. The detected regions are cropped from the corresponding images for improved visualization. Note that all detection results are obtained using the same false positive per image (FPPI) criterion. Our PSC-Net accurately detects pedestrians with varying level of occlusions.

Network (GCN) [14]. Our PSC module only uses the standard full body supervision and neither requires part-based annotations nor rely on additional visible bounding-box (VBB) information. To the best of our knowledge, we are the first to propose an approach based on GCN to capture both inter and intra-part spatial co-occurrence for occluded pedestrian detection.

We validate our PSC-Net by performing comprehensive experiments on two standard pedestrian detection datasets: CityPersons [51] and Caltech [10]. Our results clearly demonstrate that the proposed PSC-Net provides a significant improvement over the baseline. On the heavy occluded (**HO**) set of CityPersons, our PSC-Net achieves an absolute gain of 5.5% in terms of log-average miss rate, compared to the baseline. Further, PSC-Net sets a new state-of-the-art on both datasets. On the CityPerson (**HO**) test set, PSC-Net achieves a log-average miss rate of 37.4 with a significant gain of 3.6% over the best reported results in literature [34]. Fig. 1 shows example detections with varying level of occlusions (partial to severe). Our PSC-Net is able to accurately detect pedestrians, compared to both the baseline and the state-of-the-art MGAN [34].

II. RELATED WORK

Convolutional neural networks (CNNs) have significantly advanced the state-of-the-art in numerous computer vision applications, such as image classification [12], [41], [39], [43], object detection [37], [33], [45], [29], [21], object counting [48], [9], [8], [17], image retrieval [35], [36], [23], [49], action recognition [40], [38], [19], [28], and pedestrian detection

[52], [34], [53], [46]. State-of-the-art pedestrian detection methods can generally be divided into single-stage and two-stage methods. Next, we present a brief overview of two-stage pedestrian detection methods.

Two-stage Deep Pedestrian Detection: In recent years, two-stage pedestrian detection approaches [51], [34], [57], [54], [52], [20], [53], [46] have shown superior performance on standard pedestrian benchmarks. Generally, in two-stage pedestrian detectors, a set of candidate pedestrian proposals is first generated. Then, these candidate object proposals are classified and regressed. Zhang *et al.*, [51] propose key adaptations in the popular Faster R-CNN [37] for pedestrian detection. The work of [46] propose an approach based on a bounding-box regression loss designed for crowded scenes. The work of [52] investigate several attention strategies, *e.g.*, channel, part and visible bounding-box, for pedestrian detection. The work of [5] introduce a multi-scale pedestrian detection approach with layers having receptive fields similar to object scales. Zhang *et al.*, [53] propose a loss formulation that enforces candidate proposals to be close to the corresponding objects and integrates structural information with visibility predictions. The work of [3] propose a multi-phase autoregressive pedestrian detection approach that utilizes a stackable decoder-encoder module with convolutional re-sampling layers. In [20], an adaptive NMS strategy is introduced that applies a dynamic suppression threshold to an instance.

Towards Occluded Pedestrian Detection: The problem of occluded pedestrian detection is well studied in literature [44], [56], [57], [52], [46], [54], [53], [34]. Some of these pedestrian detection approaches [44], [56] handle occlusion by exploiting part-based information where typically a set of body part detectors are learned. Each part is designated to handle a specific type (pattern) of occlusion. Other approaches [46], [53], [20] investigate novel loss formulations for detector training to improve pedestrian detection in crowded scenes under heavy occlusions. Multiple pedestrian detectors stacked in a series are also investigated [22], [3] to improve the detection performance.

Most recent approaches [52], [57], [53], [54], [34] investigate the problem of occluded pedestrian detection by utilizing additional visible bounding-box (VBB) annotations together with the standard full body information. Zhang *et al.*, [52] use VBB along with a pre-trained body part prediction model to tackle occluded pedestrian detection. The work of [57] demonstrate that an additional task of visible-region bounding-box prediction can improve the full body pedestrian detection. The work of [53] propose a novel loss that improves the localization, and a part occlusion-aware region of interest pooling, to integrate structure information with visibility predictions. Zhou *et al.*, [54] propose a discriminative feature transformation module that projects the features in to a feature space, where the distance between occluded and non-occluded pedestrians is less. Such a transformation improves the robustness of the pedestrian detector during occlusion. In their approach, the visible bounding-box (VBB) is used to identify the occluded pedestrian. Recently, MGAN [34] propose a mask-guided attention network, using VBB annotation, to emphasize the visible regions while suppressing the
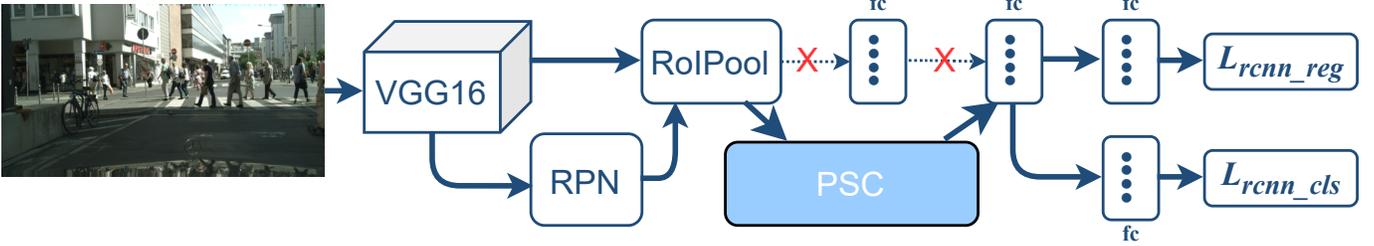
Fig. 2: Overall network architecture of our PSC-Net. It consists of a pedestrian detection (PD) branch and a part spatial co-occurrence (PSC) module. In contrast to the baseline standard PD branch where the RoI features are used for box regression and classification (X), the RoI features in our PSC-Net are fed in to the proposed PSC module to integrate both intra and inter part spatial co-occurrence information. The resulting enriched features are then deployed for final bounding-box regression and classification.

occluded regions, leading to state-of-the-art results on standard benchmarks.

Our Approach: Contrary to above mentioned recent approaches that rely on additional visible bounding-box (VBB) annotations, our proposed PSC-Net only requires the standard full body supervision to handle occluded pedestrian detection. The focus of our design is the introduction of a part spatial co-occurrence (PSC) module that explicitly captures both inter and intra-part co-occurrence information of different body parts through a Graph Convolutional Network (GCN)[14]. To the best of our knowledge, the proposed approach is the first to capture both inter and intra-part co-occurrence information through a GCN to address the problem of occluded pedestrian detection.

III. METHOD

As discussed above, occlusion is one of the most challenging problems in pedestrian detection. In case of occlusion, pedestrian body parts are either partially or fully occluded. Based on the observation that a human body generally has a fixed empirical ratio with limited flexibility of deformation, we propose a two-stage approach that utilizes spatial co-occurrence of different pedestrian body parts as a useful cue for heavily occluded pedestrian detection.

Fig. 2 shows the overall architecture of our proposed PSC-Net. It consists of a standard pedestrian detection (PD) branch (Sec. III-A) and a part spatial co-occurrence (PSC) module (Sec. III-B). The standard pedestrian detection (PD) branch is based on Faster R-CNN [37] typically employed in existing pedestrian detection works [51], [34]. The part spatial co-occurrence (PSC) module encodes both inter and intra-part co-occurrence information of different body parts. The PSC module comprises of two components. In the first component, intra-part co-occurrence of a pedestrian body part is captured by utilizing the corresponding RoI feature. As a result, an enhanced part feature representation is obtained. This representation is used as an input to the second component for capturing the inter-part co-occurrence of spatially adjacent body parts, leading to a final enhanced feature representation that encodes both intra and inter-part information. This final enhanced feature representation of a candidate proposal is then

deployed as an input to the later part of the pedestrian detection (PD) branch which performs final bounding-box regression and classification.

Next, we describe the standard pedestrian detection (PD) branch, followed by a detailed presentation of our part co-occurrence (PSC) module (Sec. III-B).

A. Standard Pedestrian Detector

Here, we describe the standard pedestrian detection (PD) branch that is based on the popular Faster R-CNN framework [37] and typically employed in several pedestrian detection methods [51], [34].

The PD branch consists of a backbone network, a region proposal network (RPN), region-of-interest (RoI) pooling layer and a classification network for final bounding-box regression and classification. In PD branch, the image is first feed into the backbone network and the RPN generates a set of candidate proposals for the input image. For each candidate proposal, a fixed-sized feature representation is obtained through RoI pooling. Finally, this fixed-sized feature representation is passed through a classification network that outputs the classification score and the regressed bounding box locations for the corresponding proposal. The loss function L_f of the standard pedestrian detection (PD) branch is given as follows:

$$L_f = L_{rpn\_cls} + L_{rpn\_reg} + L_{rcnn\_cls} + L_{rcnn\_reg}, \quad (1)$$

where L_{rpn_cls} and L_{rpn_reg} are the classification loss and bounding box regression loss of RPN, respectively, and L_{rcnn_cls} and L_{rcnn_reg} are the classification and bounding box regression loss of the classification network. In PD branch, Cross-Entropy loss is used as classification loss, and Smooth-L1 loss as bounding-box regression loss.

Limitations: To handle heavy occlusions, several recent two-stage pedestrian detection approaches [52], [53], [34] extend the PD branch by exploiting additional visible bounding-box (VBB) annotations along with the standard full body information. However, this reliance on additional VBB information implies that two sets of annotations are required for pedestrian detection training. Further, these approaches obtain a fixed-sized proposal representation by performing a pooling operation (e.g., RoI Pool or RoiAlign) on the high-level features
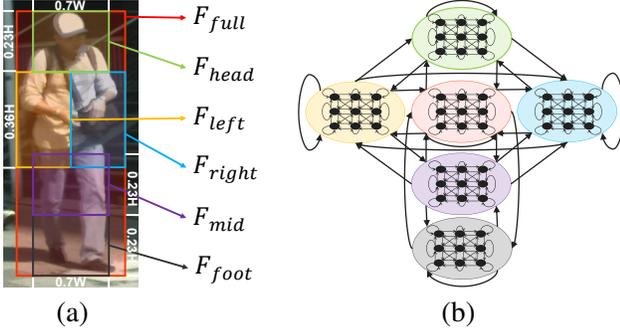
(a)   (b)

Fig. 3: (a) Full body pedestrian bounding-box partitioned into five parts based on empirically fixed ratio of human body. Each body part is shown with a different color and full body is in red. (b) Illustration of intra and inter-part spatial adjacency, used within our PSC module, to capture the spatial co-occurrence information. Our intra-part co-occurrence component employs a graph convolution layer to capture the spatial relation between different sub-regions of each body part. Differently, our inter-part component captures co-occurrence of spatially adjacent body parts using an additional graph convolution layer. Note that node colors in (b) are identical to the corresponding body parts in (a).
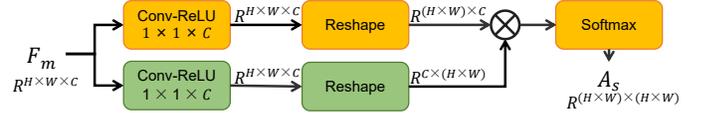

Fig. 4: Computation of intra-part spatial adjacency matrix $\mathcal{A}_s$. Each body part RoI feature F_m is first passed through two parallel convolutional layers that are cascaded by ReLU. The resulting features are re-shaped to perform matrix multiplication, followed by a softmax operation.

from the later layer of the backbone network (*e.g.*, $conv5$ of VGG). Alternatively, several anchor-free methods [47], [42] have investigated the effective utilization of features from multiple ($conv$) layers of the backbone network for pedestrian detection.

In this work, we propose a two-stage pedestrian detection method, PSC-Net, to address heavy occlusions. Our main contribution is the introduction of a *part spatial co-occurrence* (PSC) module that only requires standard full body supervision and explicitly captures both inter and intra-part spatial co-occurrence information of different body parts. Next, we describe the details of our PSC module.

B. Part Spatial Co-occurrence Module

In pedestrian detection, the task is to accurately localize the *full body* of a pedestrian. This task is relatively easier in case of standard non-occluded pedestrians. However, it becomes particularly challenging in case of partial or severe occlusions. Here, we introduce a part spatial co-occurrence (PSC) module that utilizes spatial co-occurrence of different body parts captured through a Graph Convolutional Network (GCN) [14]. In PSC module, the GCN is employed to capture intra-part and inter-part spatial co-occurrence by exploiting the topological structure of pedestrian. The intra-part co-occurrence is expected to improve the feature representation in scenarios where a particular body part is partially occluded whereas the inter-part co-occurrence targets severe occlusion of a particular body part.

Our PSC module neither requires pedestrian body part annotations nor rely on the use of an external pre-trained part model. Instead, it divides the full body bounding-box of a pedestrian into five parts ($F_{head}, F_{left}, F_{right}, F_{mid}, F_{foot}$),

based on empirical fixed ratio of human body (see Fig. 3), as in [53]. The RoI pooling operation is performed on each body part (five) as well as the full body (F_D), resulting in six RoI pooled features for each proposal.

As described above, the RoI pooling operation is performed for five body parts as well as the full body, resulting in an increased feature dimension. Therefore, a direct utilization of all these RoI features will drastically increase the computational complexity of our PSC module. Note that the Faster R-CNN and its pedestrian detection adaptations [51], [34], [53] commonly use a single RoI pooling only on the $conv5$ features of VGG, resulting in 512 channels. To maintain a similar number of channels as in Faster R-CNN and its pedestrian detection adaptations [52], [53], [34], we introduce an additional 1×1 convolution in RoI pooling strategy that significantly reduces the number of channels (572 in total). Consequently, RoI pooled features of each body part and the full body has only 64 and 256 channels, respectively.

1) Intra-part Co-occurrence: Here, we enhance RoI pooled feature representation of each body part by considering their intra-part co-occurrence. For instance, consider a scenario where head part F_{head} is partially occluded, thereby making top-part of the head invisible. Our intra-part co-occurrence component aims to capture spatial relation between different sub-regions (*e.g.* eyes and ears) within a RoI feature F_m of a body part (*e.g.*, F_{head}) through graph convolutional layer,

$$\tilde{F}_m = \sigma(A_s F_m W_s) \tag{2}$$

where σ is the ReLU activation, $W_s \in R^{C \times C}$ is a learnable parameter matrix, C is the number of channels (64 for each body part and 256 for full body) and $A_s \in R^{H \times W \times H \times W}$ is intra-part spatial adjacency matrix of a graph $\mathcal{G}_s = (\mathcal{V}_s, \mathcal{A}_s)$. Here, each pixel within the RoI region is treated as a graph node. In total, there are $H \times W$ nodes $\mathcal{V}_s$ in the graph.

The intra-part spatial adjacency matrix A_s is computed as follow, and is also shown in Fig. 4. We first pass the RoI feature $F_m \in R^{H \times W \times C}$ through two parallel 1×1 convolution layers that are cascaded by ReLU activation. The resulting outputs are re-shaped prior to performing a matrix multiplication, followed by a softmax operation to compute the intra-part spatial adjacency matrix A_s.

The output from the graph convolution layer (Eq. 2) is $\tilde{F}_m \in R^{H \times W \times C}$. This output $\tilde{F}_m$ is first added to its input F_m (original RoI feature), followed by passing it through a fully connected layer to obtain a d dimensional enhanced part feature. The enhanced part features F_e of all six parts (five

parts and full body) are further used to capture inter-part co-occurrence described next.

2) Inter-part Co-occurrence: Our inter-part co-occurrence component is designed to improve feature representation, especially in case of severe occlusion of a particular body part.

The traditional convolution layer only captures information from a small spatial neighborhood, defined by the kernel size (*e.g.* 3×3), and is therefore often ineffective to encode inter-part co-occurrence information of a pedestrian. To address this issue, we introduce an additional graph convolutional layer in our PSC module, that captures inter-part relationship of different body parts and the full body of a pedestrian. We treat each part (including full body) as separate nodes ($\mathcal{V}_p$) of a graph $\mathcal{G}_p = (\mathcal{V}_p, \mathcal{A}_p)$, where $\mathcal{A}_p$ denote the spatial adjacency matrix capturing the neighborhood relationship of different nodes. The graph convolutional operation is used to improve the node features F_e as follows,

$$\tilde{F}_e = \sigma((I - \mathcal{A}_p)F_e W_p) \qquad (3)$$

where $F_e \in R^{n \times d}$ is enhanced features of n parts, A_p is the spatial adjacency matrix, $W_p \in R^{d \times d}$ is learnable parameter matrix and σ is ReLU activation. Here, $n = 6$ is the number of nodes, and $d = 1024$ is channels of each node. $(I - \mathcal{A}_p)$ is used to conduct Laplacian smoothing [16] to propagate the node features over the graph.

The spatial adjacency matrix $\mathcal{A}_p$ captures the relation between different graph nodes. During heavy occlusion, enhanced part features of a particular node may not contain relevant body part information. Therefore, it is desired to assign lesser weights to the edges linking such nodes in $\mathcal{A}_p$. To this end, we introduce a self-attention scheme to each edge that assigns learnable weight a_{ij}. The input of the self attention is the concatenated features of node i and node j. The self attention of each edge a_{ij} is computed by a fully-connected operations followed by a sigmoid activation. The unnormalized spatial adjacency matrix $\hat{\mathcal{A}}_p(i,j)$ is defined as:

$$\hat{\mathcal{A}}_p(i,j) = \begin{cases} a_{ij} & \text{part } i \text{ and part } j \text{ are spatial adjacent} \\ 0 & \text{otherwise} \end{cases} \qquad (4)$$

The spatial adjacency matrix $\mathcal{A}_p(i,j)$ is computed by conducting $L2$ normalize in each row of $\hat{\mathcal{A}}_p(i,j)$.

Afterwards, we employ full connected layer to merge all the features $\tilde{F}_e \in R^{n \times d}$ into a d-dimensional feature vector. The resulting *enriched* features are then utilized as an input to the classification network which predicts the final classification score and bounding-box regression.

IV. EXPERIMENTS

Datasets: We perform experiments on two datasets: CityPersons [51] and Caltech [10]. CityPersons [51] consists of 2975 training, 500 validation, and 1525 test images. Caltech [10] contains 11 sets of videos. First 6 sets (0-5) are used for training, the last 5 sets (6-10) are used for testing. To get a large amount of training data, we sample the videos with 10Hz. The training sets consists of 42782 images in total.

Evaluation Metrics: We report the performance using log-average miss rate (MR) throughout our experiments. Here, MR

Baseline (PD)	Intra-Part Co (Sec. III-B1)	Inter-Part Co (Sec. III-B2)	R	HO
✓			13.8	56.8
✓	✓		11.8	53.1
✓		✓	12.5	52.1
✓	✓	✓	**10.6**	**50.2**

TABLE I: Impact of integrating our intra-part (Intra-Part Co) and inter-part (Inter-Part Co) co-occurrence into the baseline on CityPersons val. set. Performance is reported in terms of log-average miss rate and best results are boldfaced. Our final PSC-Net that integrates both the intra-part and inter-part co-occurrence achieves consistent improvement in performance, with gains of 3.2% and 6.6% on the **R** and **HO** sets, respectively, over the baseline.

is computed over the false positive per image (FPPI) range of $[10^{-2}, 10^0]$ [10]. On CityPersons, we report the results across two different occlusion degrees: Reasonable (**R**) and Heavy Occlusion (**HO**). For both **R** and **HO** sets, the height of pedestrians is larger than 50 pixels. The visibility ratio in **R** set is larger than 65% whereas in **HO** it ranges from 20% to 65%. In addition to these sets, the results are reported on combined (**R + HO**) set on Caltech.

Implementation Details: For both datasets, we train our network on a NVIDIA GPU with the mini-batch consisting of two images per GPU. Adam [13] solver is selected as optimizer. In case of CityPersons, we fine-tune pre-trained ImageNet VGG model [41] on the trainset of the CityPersons. We follow the same experimental protocol as in [51] and employ two fully connected layers with 1024 instead of 4096 output dimensions. The initial learning rate is set to 10^{-4} for the first 8 epochs, and decay it by a factor of 10 for another 3 epochs. For Caltech, we start with a model that is pre-trained on CityPersons. An initial learning rate of 10^{-5} is used for first 3 training epochs and decay it to 10^{-6} for another 1 training epoch. Both the source code and models will be publicly available.

A. CityPersons Dataset

Baseline Comparison: As discussed earlier, our design is the introduction of a part spatial co-occurrence (PSC) module that explicitly captures both intra-part (Sec. III-B1) and inter-part (Sec. III-B2) co-occurrence information of different body parts. Tab. I shows the impact of integrating our intra-part and inter-part co-occurrence components in the baseline. For fair comparison, all results in Tab. I are reported by using the same set of ground-truth pedestrian examples during training. All ground-truth pedestrian examples which are at least 50 pixels tall with visibility ≥ 65 are utilized for training. Further, the input scale of $1.0\times$ is employed during the experiments. The integration of each component into the baseline results in consistent improvement in performance. Further, our final PSC-Net that integrates both the intra-part and inter-part co-occurrence achieves absolute gains of 3.2% and 6.6% on **R** and **HO** sets, respectively, over the baseline. These results demonstrate that both components are required to obtain optimal performance.

State-of-the-art Comparison: Here, we perform a comparison of our PSC-Net with state-of-the-art pedestrian detec-

Methods	Data (Visibility)	Input Scales	R	HO
TLL [42]	-	1×	14.4	52.0
F.RCNN+ATT-part [52]	≥ 65%	1×	16.0	56.7
F.RCNN+ATT-vbb [52]		1×	16.4	57.3
Repulsion Loss [46]		1×	13.2	56.9
Adaptive-NMS [20]		1×	11.9	55.2
MGAN [34]		1×	11.5	51.7
PSC-Net (Ours)		1×	**10.6**	**50.2**
OR-CNN [53]	≥ 50%	1×	12.8	55.7
MGAN [34]		1×	10.8	46.7
PSC-Net (Ours)		1×	**10.3**	**44.9**
ALFNet [22]	≥ 0%	1×	12.0	52.0
CSP [47]		1×	11.0	49.3
MGAN [34]		1×	11.3	42.0
PSC-Net (Ours)		1×	**10.5**	**39.5**
Repulsion Loss [46]	≥ 65%	1.3×	11.5	55.3
Adaptive-NMS [20]		1.3×	10.8	54.0
MGAN [34]		1.3×	10.3	49.6
PSC-Net (Ours)		1.3×	**9.8**	**48.3**
OR-CNN [53]	≥ 50%	1.3×	11.0	51.3
MGAN [34]		1.3×	9.9	45.4
PSC-Net (Ours)		1.3×	**9.6**	**43.6**
Bi-box [57]	≥ 30%	1.3×	11.2	44.2
FRCN +A +DT [54]		1.3×	11.1	44.3
MGAN [34]		1.3×	10.5	39.4
PSC-Net (Ours)		1.3×	**9.9**	**37.2**

TABLE II: Comparison with the state-of-the-art (in terms of log-average miss rate) on the CityPersons validation set. In each case, best results are boldfaced. Our PSC-Net achieves superior performance on both datasets, compared to existing methods. When using same input scale (1.3x), training data visibility (≥ 30%) and backbone (VGG), PSC-Net provides absolute gains of 7.1% and 2.2% over FRCN +A +DT [54] and MGAN [34], respectively on the **HO** set.

Method	R	HO
Adaptive Faster RCNN [51]	13.0	50.5
MS-CNN [7]	13.3	51.9
Rep. Loss [46]	11.5	52.6
OR-CNN [53]	11.3	51.4
Cascade MS-CNN [7]	11.6	47.1
Adaptive-NMS [20]	11.4	-
MGAN [34]	**9.3**	41.0
PSC-Net (Ours)	**9.3**	**37.0**

TABLE III: State-of-the-art comparison (in terms of log-average miss rate) on CityPersons test set. The test set is withheld and results are obtained by sending our PSC-Net detection predictions for evaluation to the authors of CityPersons [51]. Our PSC-Net outperforms existing methods on both **R** and **HO** sets. On heavy occluded **HO** set, PSC-Net achieves an absolute gain of 4.0% over the state-of-the-art [34]. In each case, best results are boldfaced.

tion methods in literature. Tab. II shows the comparison on CityPersons validation set. Note that existing approaches utilize different set of ground-truth pedestrian examples for training. For fair comparison, we therefore select same set of ground-truth pedestrian examples (denoted as data (visibility) in Tab. II) and input scale, when performing a comparison with each state-of-the-art method. Our PSC-Net achieves superior performance on all these settings for both **R** and **HO** sets, compared to the state-of-the-art methods.

When using an input scale of 1× and data visibility (≥ 65%), the attention-based approaches, F.RCNN+ ATT-part [52] and F.RCNN+ATT-vbb [52], obtain log-average miss rates

Detector	Occl.	R	HO	R+HO
CompACT-Deep [6]	×	11.8	65.8	24.6
ATT-vbb [52]	✓	10.3	45.2	18.2
MS-CNN [5]	×	10.0	59.9	21.5
RPN+BF [50]	×	9.6	74.4	24.0
SA-F.RCNN [15]	×	9.7	64.4	21.9
SDS-RCNN [4]	×	7.4	58.6	19.7
F.RCNN [51]	×	9.2	57.6	20.0
GDFL [18]	×	7.9	43.2	15.6
Bi-Box [57]	✓	7.6	44.4	16.1
AR-Ped [3]	×	6.5	48.8	16.1
FRCN +A +DT [54]	✓	8.0	37.9	-
MGAN [34]	✓	6.8	38.2	13.8
PSC-Net (Ours)	✓	**6.4**	**34.8**	**12.7**

TABLE IV: State-of-the-art comparison (in terms of log-average miss rate) on Caltech test set. Best results are bold-faced in each case. Our PSC-Net provides consistent improvements (over) on all sets. On the **HO** set, PSC-Net outperforms the best reported results [54] by reducing the error from 37.9 to 34.8.

of 16.0, 56.7 and 16.4, 57.3 on the **R** and **HO** sets. The work of [46] based on Repulsion Loss obtains log-average miss rates of 13.2 and 56.9 on the **R** and **HO** sets, respectively. The Adaptive-NMS approach [20] that applies a dynamic suppression threshold and learns density scores obtains log-average miss rates of 11.9 and 55.2 on the **R** and **HO** sets, respectively. MGAN [34] learns a spatial attention mask using VBB information to modulate full body features and achieves log-average miss rates of 11.5 and 51.7 on the **R** and **HO** sets, respectively. Our PSC-Net outperforms MGAN, without using VBB supervision, on both sets with log-average miss rates of 10.6 and 50.2. When using same data visibility but 1.3× input scale, Adaptive-NMS [20] and MGAN [34] achieve log-average miss rates of 54.0 and 49.6, respectively on the **HO** set. The same two approaches report 10.8 and 10.3 on the **R** set. PSC-Net achieves superior results with log-average miss rates of 9.8 and 48.3 on the **R** and **HO** sets, respectively. On this dataset, the best existing results of 10.5 and 39.4 are reported [34] on **R** and **HO** sets, respectively when using an input scale of 1.3× and data visibility (≥ 30%). Our PSC-Net outperforms the state-of-the-art [34] with log-average miss rates of 9.9 and 37.2 on the two sets.

Tab. III shows the comparison on CityPersons test set. Among existing methods, the multi-stage Cascade MS-CNN [7] consisting of a sequence of detectors trained with increasing IoU thresholds obtains log average miss rate of 47.1 on the **HO** set. MGAN [34] obtains a log average miss rate of 41.0 on the same set. Our PSC-Net significantly reduces the error by 4.0% over MGAN on the **HO** set. Similarly, PSC-Net also improves the performance on the **R** set.

B. Caltech Dataset

Finally, we evaluate our PSC-Net on Caltech and compare it with state-of-the-art approaches in literature. Tab. II shows the comparison on Caltech test set under three sets: **R**, **H** and **R + HO**. Among existing methods, ATT-vbb [52], Bi-box [57], FRCN +A +DT [54] and MGAN [34], address the problem of occlusions by utilizing VBB information. On
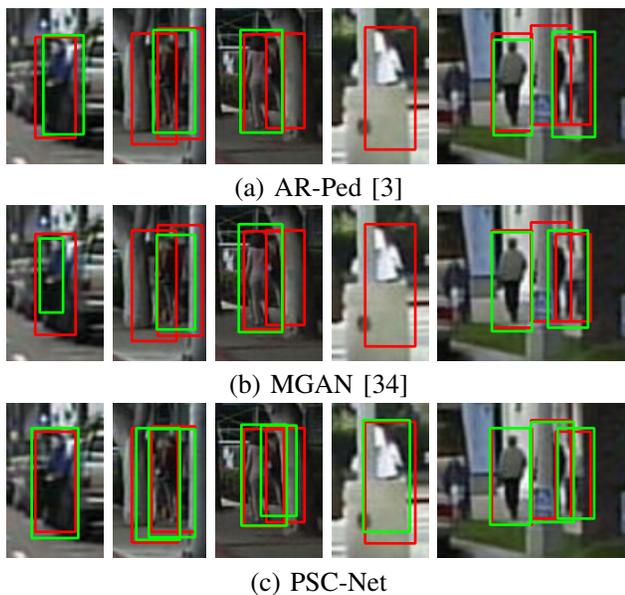
(a) AR-Ped [3]

(b) MGAN [34]

(c) PSC-Net

Fig. 5: Qualitative detection comparison of (c) PSC-Net with (a) AR-Ped [3] and (b) MGAN [34] under occlusions on caltech test images. Here, all detection results are obtained using the same false positive per image criterion. The red boxes denote the ground-truth whereas the detector predictions are in green. For better visualization, the detected regions are cropped from the corresponding images.

the **R**, **H** and **R + HO** subsets, AR-Ped [3], FRCN +A +DT [54] and MGAN [34] report best existing performance, respectively. PSC-Net achieves superior detection performance on all three subsets with log average miss rates of 6.4, 34.8 and 12.7, respectively. Fig. 5 displays example detections showing a visual comparison of our PSC-Net with recently introduced AR-Ped [3] and MGAN [34] under occlusions. All detection results in Fig. 5 are obtained using same false positive per image (FPPI) criterion. Our PSC-Net provides improved detections under these occlusion scenarios compared to the other two approaches.

V. CONCLUSION

We proposed a two-stage approach, PSC-Net, for occluded pedestrian detection. Our PSC-Net consists of a standard pedestrian detection branch and a part spatial co-occurrence (PSC) module. The focus of our design is the PSC module that is designed to capture intra-part and inter-part spatial co-occurrence of different body parts through a Graph Convolutional Network (GCN). Our PSC module only requires standard full body supervision and exploits the topological structure of pedestrian. Experiments are performed on two popular datasets: CityPersons and Caltech. Our results clearly demonstrate that the proposed PSC-Net significantly outperforms the baseline in all cases. Further, the PSC-Net sets a new state-of-the-art on both datasets.

REFERENCES

[1] Moshe Bar and Shimon Ullman. Spatial context in recognition. *Perception*, 25:343–352, 1993.

[2] Irving Biederman, Robert Mezzanotte, and Jan Rabinowitz. Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, 14(2):143–177, 1982.

[3] Garrick Brazil and Xiaoming Liu. Pedestrian detection with autoregressive network phases. In *CVPR*, 2019.

[4] Garrick Brazil, Xi Yin, and Xiaoming Liu. Illuminating pedestrians via simultaneous detection & segmentation. In *ICCV*, 2017.

[5] Zhaowei Cai, Quanfu Fan, Rogerio Schmidt Feris, and Nuno Vasconcelos. A unified multi-scale deep convolutional neural network for fast object detection. In *ECCV*, 2016.

[6] Zhaowei Cai, Mohammad Saberian, and Nuno Vasconcelos. Learning complexity-aware cascades for deep pedestrian detection. In *ICCV*, 2015.

[7] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: High quality object detection and instance segmentation. *arXiv preprint arXiv:1906.09756*, 2019.

[8] Hisham Cholakkal, Guolei Sun, Fahad Shahbaz Khan, and Ling Shao. Object counting and instance segmentation with image-level supervision. In *CVPR*, 2019.

[9] Hisham Cholakkal, Guolei Sun, Salman Khan, Fahad Shahbaz Khan, Ling Shao, and Luc Van Gool. Towards partial supervision for generic object counting in natural scenes. *arXiv preprint arXiv:1912.06448*, 2019.

[10] P. Dollar, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: An evaluation of the state of the art. *T-PAMI*, 2012.

[11] Carolina Galleguillos, Andrew Rabinovich, and Serge Belongie. Object categorization using co-occurrence, location and appearance. In *CVPR*, 2008.

[12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[13] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. ICLR.

[14] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.

[15] Jianan Li, Xiaodan Liang, ShengMei Shen, Tingfa Xu, Jiashi Feng, and Shuicheng Yan. Scale-aware fast R-CNN for pedestrian detection. *T-MM*, 2018.

[16] Q. Li, Z. Han, and X.-M. Wu. Deeper Insights into Graph Convolutional Networks for Semi-Supervised Learning. In *AAAI*, 2018.

[17] Yuhong Li, Xiaofan Zhang, and Deming Chen. Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes. In *CVPR*, 2018.

[18] Chunze Lin, Jiwen Lu, Gang Wang, and Jie Zhou. Graininess-aware deep feature learning for pedestrian detection. In *ECCV*, 2018.

[19] Li Liu, Yi Zhou, and Ling Shao. Deep action parsing in videos with large-scale synthesized data. *IEEE Transactions on Image Processing*, 27(6):2869–2882, 2018.

[20] Songtao Liu, Di Huang, and Yunhong Wang. Adaptive nms: Refining pedestrian detection in a crowd. In *CVPR*, 2019.

[21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016.

[22] Wei Liu, Shengcai Liao, Weidong Hu, Xuezhi Liang, and Xiao Chen. Learning efficient single-stage pedestrian detectors by asymptotic localization fitting. In *ECCV*, 2018.

[23] Yu Liu, Jingkuan Song, Ke Zhou, Lingyu Yan, Li Liu, Fuhao Zou, and Ling Shao. Deep self-taught hashing for image retrieval. *IEEE Transactions on Cybernetics*, 49(6):2229–2241, 2019.

[24] Jiayuan Mao, Tete Xiao, Yuning Jiang, and Zhimin Cao. What can help pedestrian detection? In *CVPR*, 2017.

[25] Markus Mathias, Rodrigo Benenson, Radu Timofte, and Luc Van Gool. Handling occlusions with franken-classifiers. In *ICCV*, 2013.

[26] Krystian Mikolajczyk, Cordelia Schmid, and Andrew Zisserman. Human detection based on a probabilistic assembly of robust part detectors. In *ECCV*, 2004.

[27] Anuj Mohan, Constantine Papageorgiou, and Tomaso A. Poggio. Example-based object detection in images by components. *PAMI*, 2001.

[28] Sanath Narayan, Hisham Cholakkal, Fahad Shahbaz Khan, and Ling Shao. 3c-net: Category count and center loss for weakly-supervised action localization. In *ICCV*, 2019.

[29] Jing Nie, Rao Muhammad Anwer, Hisham Cholakkal, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. Enriched feature guided refinement network for object detection. In *ICCV*, 2019.

[30] Junhyug Noh, Soochan Lee, Beomsu Kim, and Gunhee Kim. Improving occlusion and hard negative handling for single-stage pedestrian detectors. In *CVPR*, 2018.

[31] Wanli Ouyang and Xiaogang Wang. Joint deep learning for pedestrian detection. In *ICCV*, 2013.

[32] Wanli Ouyang, Xingyu Zeng, and Xiaogang Wang. Modeling mutual visibility relationship in pedestrian detection. In *CVPR*, 2013.

[33] Yanwei Pang, Tiancai Wang, Rao Muhammad Anwer, Fahad Shahbaz Khan, and Ling Shao. Efficient featurized image pyramid network for single shot detector. In *CVPR*, 2019.

[34] Yanwei Pang, Jin Xie, Muhammad Haris Khan, Rao Muhammad Anwer, Fahad Shahbaz Khan, and Ling Shao. Mask-Guided attention network for occluded pedestrian detection. In *ICCV*, 2019.

[35] Filip Radenovic, Giorgos Tolias, and Ondrej Chum. Cnn image retrieval learns from bow: Unsupervised fine-tuning with hard examples. In *ECCV*, 2016.

[36] Filip Radenovic, Giorgos Tolias, and Ondrej Chum. Fine-tuning cnn image retrieval with no human annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1655–1668, 2019.

[37] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.

[38] Fahad Shahbaz Khan, Jiaolong Xu, Joost van de Weijer, Andrew Bagdanov, Rao Muhammad Anwer, and Antonio Lopez. Recognizing actions through action-specific person detection. *IEEE Transactions on Image Processing*, 24(11):4422–4432, 2015.

[39] Ling Shao, Ziyun Cai, Li Liu, and Ke Lu. Performance evaluation of deep feature learning for rgb-d image/video classification. *Information Sciences*, 385(1):266–283, 2017.

[40] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.

[41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[42] Tao Song, Leiyu Sun, Di Xie, Haiming Sun, and Shiliang Pu. Small-scale pedestrian detection based on topological line localization and temporal feature aggregation. In *ECCV*, 2018.

[43] Guolei Sun, Hisham Cholakkal, Salman Khan, Fahad Shahbaz Khan, and Ling Shao. Fine-grained recognition: Accounting for subtle differences between similar classes. In *AAAI*, 2020.

[44] Yonglong Tian, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning strong parts for pedestrian detection. In *ICCV*, 2015.

[45] Tiancai Wang, Rao Muhammad Anwer, Hisham Cholakkal, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. Learning rich features at high-speed for single-shot object detection. In *ICCV*, 2019.

[46] Xinlong Wang, Tete Xiao, Yuning Jiang, Shuai Shao, Jian Sun, and Chunhua Shen. Repulsion loss: Detecting pedestrians in a crowd. In *CVPR*, 2018.

[47] Weiqiang Ren Weidong Hu Yinan Yu Wei Liu, Shengcai Liao. High-level semantic feature detection: A new perspective for pedestrian detection. In *CVPR*, 2019.

[48] Cong Zhang, Hongsheng Li, Xiaogang Wang, and Xiaokang Yang. Cross-scene crowd counting via deep convolutional neural networks. In *CVPR*, 2015.

[49] Haofeng Zhang, Li Liu, Yang Long, and Ling Shao. Unsupervised deep hashing with pseudo labels for scalable image retrieval. *IEEE Transactions on Image Processing*, 27(4):1626–1638, 2018.

[50] Liliang Zhang, Liang Lin, Xiaodan Liang, and Kaiming He. Is faster R-CNN doing well for pedestrian detection? In *ECCV*, 2016.

[51] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. In *CVPR*, 2017.

[52] Shanshan Zhang, Jian Yang, and Bernt Schiele. Occluded pedestrian detection through guided attention in cnns. In *CVPR*, 2018.

[53] Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z. Li. Occlusion-aware R-CNN: Detecting pedestrians in a crowd. In *ECCV*, 2018.

[54] Chunluan Zhou, Ming Yang, and Junsong Yuan. Discriminative feature transformation for occluded pedestrian detection. In *ICCV*, 2019.

[55] Chunluan Zhou and Junsong Yuan. Learning to integrate occlusion-specific detectors for heavily occluded pedestrian detection. In *ACCV*, 2016.

[56] Chunluan Zhou and Junsong Yuan. Multi-label learning of part detectors for heavily occluded pedestrian detection. In *ICCV*, 2017.

[57] Chunluan Zhou and Junsong Yuan. Bi-box regression for pedestrian detection and occlusion estimation. In *ECCV*, 2018.