

# Learning to Separate: Detecting Heavily-Occluded Objects in Urban Scenes

Chenhongyi Yang<sup>1</sup>[0000-0003-3895-6895], Vitaly Ablavsky<sup>2</sup>[0000-0003-2703-7666]\*, Kaihong Wang<sup>1</sup>[0000-0002-0637-9862], Qi Feng<sup>1</sup>[0000-0001-6342-3228], and Margrit Betke<sup>1</sup>[0000-0002-4491-6868]

<sup>1</sup> Boston University  
{hongyi,kaiwh,func,betke}@bu.edu  
<sup>2</sup> University of Washington  
vxa@uw.edu

**Abstract.** While visual object detection with deep learning has received much attention in the past decade, cases when heavy intra-class occlusions occur have not been studied thoroughly. In this work, we propose a Non-Maximum-Suppression (NMS) algorithm that dramatically improves the detection recall while maintaining high precision in scenes with heavy occlusions. Our NMS algorithm is derived from a novel embedding mechanism, in which the semantic and geometric features of the detected boxes are jointly exploited. The embedding makes it possible to determine whether two heavily-overlapping boxes belong to the same object in the physical world. Our approach is particularly useful for car detection and pedestrian detection in urban scenes where occlusions often happen. We show the effectiveness of our approach by creating a model called SG-Det (short for Semantics and Geometry Detection) and testing SG-Det on two widely-adopted datasets, KITTI and CityPersons for which it achieves state-of-the-art performance. Our code is available at <https://github.com/ChenhongyiYang/SG-NMS>.

## 1 Introduction

Recent years have witnessed significant progress in object detection using deep convolutional neural networks (CNNs) [1,2,3]. The approach taken by many state-of-the-art object detection methods [4,5,6,7,8] is to predict multiple bounding boxes for an object and then use a heuristic method such as non-maximum suppression (NMS) to remove superfluous bounding boxes that stem from duplicate detected objects.

The Greedy-NMS algorithm is easy to implement and tends to work well in images where objects of the same class do not significantly occlude each other. However, in urban scenes, where the task is to detect potentially heavily occluded cars or pedestrians, Greedy-NMS does not perform adequately. The decrease in accuracy is due to the fundamental limitation of the NMS algorithm, which uses

---

\* Work performed at Boston University.

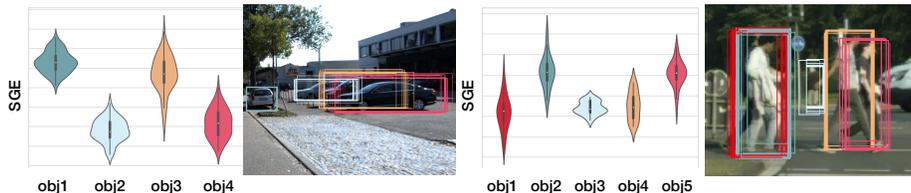


Fig. 1: Learned Semantics-Geometry Embedding (SGE) for bounding boxes predicted by our proposed detector on KITTI and CityPersons images. Heavily overlapped boxes are separated in the SGE space according to the objects they are assigned to. Thus, distance between SGEs can guide Non-Maximum-Suppression to keep correct boxes in heavy intra-class occlusion scenes.

a fixed threshold to determine which bounding boxes to suppress: The algorithm cannot suppress duplicate bounding boxes belonging to the *same* object while preserving boxes belonging to *different* objects, where one object heavily occludes others. Soft-NMS [9] attempts to address this limitation by not removing overlapping boxes but instead lowering their confidence; however, all overlapping boxes are still treated as false positives regardless of how many physical objects are in the image.

The limitation of NMS could be circumvented with an oracle that assigns each bounding box an identifier that corresponds to its physical-world object. Then, a standard NMS algorithm could be applied per set of boxes with the same identifier (but not across identifiers), thus ensuring that false positives from one object do not result in suppression of a true positive from a nearby object.

To approximate such an oracle, we can try to learn a mapping from boxes into a latent space so that the heavily overlapping boxes can be separated in that space. Naively, this mapping can be implemented by learning an embedding for every box based on its region features, e.g., the pooled features after RoIPooling [7]. However, the usefulness of such an embedding would be limited because heavily overlapping boxes tend to yield similar region features, thus would map to nearby points in the embedding space. In this paper, we demonstrate that by considering both the region features and the geometry of each box, we can successfully learn an embedding in a space where heavily overlapping boxes are separated if they belong to different objects. We call the learned embedding Semantics-Geometry Embedding (SGE). We also propose a novel NMS algorithm that takes advantage of the SGE to improve detection recall.

We visualize the concept of a Semantics-Geometry Embedding (SGE) in Fig. 1, where boxes belonging to the same object are mapped to a similar SGE and boxes belonging to different but occluded objects are mapped to SGEs that are far away. Although the embedding algorithm may assign boxes in disparate parts of an image to similar SGEs, this does not negatively impact our SG-NMS algorithm because these boxes can be easily separated based on their intersection-over-union (IoU) score. The SGE is implemented as an associative embedding [10] and learned using two loss functions, *separation* and *group loss*.

To train the SGE with the object detector end-to-end, we propose a novel Serial Region-based Fully Convolutional Network (Serial R-FCN), where the geometric feature of each detected box is precisely aligned with its semantic features. We combine this network with the Semantics-Geometry Non-Maximum-Suppression (SG-NMS) algorithm in a model we call SG-Det (short for Semantics and Geometry Detection). In summary, we make three main contributions:

1. A bounding-box-level *Semantics-Geometry Embedding (SGE)* is proposed, and a Non-Maximum Suppression algorithm, called *Semantics-Geometry Non-Maximum-Suppression (SG-NMS)*, based on this embedding, is derived. The algorithm markedly improves object detection in scenarios with heavy intra-class occlusions.
2. A *serial R-FCN* with self-attention in each head is presented that not only provides the ability to learn the above-mentioned SGE end-to-end, but also improves object detection accuracy.
3. The model *SG-Det* is proposed, which combines the serial R-FCN and the SG-NMS algorithm. *SG-Det* achieves state-of-the-art performance on the tasks of car detection for the KITTI [11] dataset and pedestrian detection for the CityPersons [12] dataset by dramatically improving the detection recall in heavily-occluded scenes.

## 2 Related Works

**Object Detection.** CNN-based object detectors can be divided into one-stage and two-stage approaches. One-stage detectors [4,5,13] directly predict the object class and the bounding box regressor by sliding windows on the feature maps. Two-stage object detectors [7,8,14,15], first compute regions of interest (RoIs) [8,16,17,18,19] and then estimate the class label and bounding box coordinates for each RoI. Although the two-stage approaches often achieve higher accuracy, they suffer from low computational efficiency. R-FCN [14] addresses this problem by replacing the computation in fully-connected layers with nearly cost-free pooling operations.

**Non Maximum Suppression.** NMS is widely used in modern object detectors to remove duplicate bounding boxes, but it may mistakenly remove boxes belonging to different objects. Soft-NMS [9] was proposed to address this problem by replacing the fixed NMS threshold with a score-lowering mechanism. However, highly-overlapping boxes are still treated as false positives regardless of the semantic information. In Learning-NMS [20], a neural network is used to perform NMS, but the appearance information is still not considered. The Adaptive-NMS approach [21] learns a threshold with the object detector, but when the threshold is set too high, false positives may be kept. The relation of bounding boxes can also be used to perform NMS by considering their appearance and geometric features [22], but this does not handle intra-class occlusion. The localization quality of each box can be learned to help NMS with keeping accurate boxes [23,24,25,26].

**Other Occlusion Handling Approaches.** There are many other methods designed to handle occlusion, including both intra-class or inter-class occlusion.

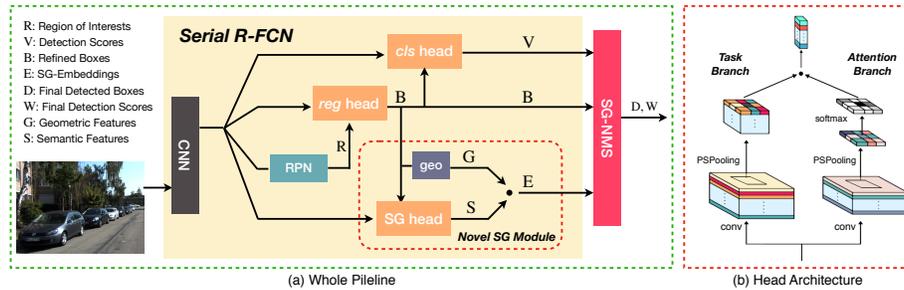


Fig. 2: (a) Overview of our proposed model SG-Det. An input image is first processed by a backbone CNN to yield feature maps. A Region Proposal Network (RPN) [8] is used to extract regions of interests (ROIs). The ROIs will first be refined by the regression head and then fed into the classification head to produce detection scores, making the whole pipeline serial. A novel **Semantics-Geometry Module**, parallel to the classification head, is added to learn the SG embedding for each refined box. Finally, the detected box, detection scores, and SG embeddings are fed into the SG-NMS algorithm to produce final detections. (b) All heads (orange boxes in Fig. 2(a)) share a similar architecture. The feature map computed by the backbone network is processed by two branches to yield two score maps. Then a Position Sensitive RoI-Pooling [14] is applied to produce two grids of  $k^2$  position-sensitive scores – a task-specific score and an attention score. A softmax operation transforms the attention score into a discrete distribution over the  $k^2$  grids. Finally the  $k^2$  task scores are aggregated by the attention distribution to yield the final output scores.

Most of them focus on detecting pedestrians in crowd scenes. Repulsion loss [27] was proposed to prevent boxes from shifting to adjacent objects. The occluded person is detected by considering different body parts separately [17,28,29,30,31]. A novel expectation-maximization merging unit was proposed to resolve overlap ambiguities [32]. Additional annotations such as head position or visible regions have been used [33,34,35] to create robust person detectors. Although these approaches have been shown to be effective in detecting occluded persons, it is difficult to generalize them to other tasks like car detection.

### 3 Methodology

In this section, we first introduce the proposed Semantics-Geometry Embedding (Section 3.1), then the Semantics-Geometry NMS algorithm (Section 3.2), and finally the proposed Serial R-FCN (Section 3.3). The overview of the combined proposed model *SG-Det* is shown in Figure 2.

See also <http://www.cs.bu.edu/fac/betke/research/SG-Det>.

### 3.1 Semantics-Geometry Embedding

Our key idea for separating occluded objects in an image is to map each putative detection to a point in a latent space. In this latent space, detections belonging to the same physical object form a tight cluster; detections that are nearby in the image plane, but belong to different physical objects, are pushed far apart.

To implement this idea, we design an embedding for each bounding box that takes the form of a dot-product

$$e = \mathbf{s}^T \cdot \mathbf{g}, \quad (1)$$

where  $\mathbf{g}$  is the *geometric feature* and  $\mathbf{s}$  is the *semantic feature*. The geometric feature has a fixed form  $\mathbf{g} = (x, y, w, h)^T$  with center coordinates  $(x, y)$  and width and height  $(w, h)$  of the bounding box. We tried different kinds of geometric features (e.g., [22]) and feature vectors with higher dimensions produced by a fully-connected layer, but found that such complexity did not provide further significant improvement.

Unlike the geometric feature, the semantic feature  $\mathbf{s}$  is a weight output by a function that yields a vector compatible with  $\mathbf{g}$ ; the function is implemented as a neural network, as shown in the Semantics-Geometry Head in Fig. 2. Note that the SGE is computed by the linear transformation of the geometric feature taking the learned semantic feature as a weight. An interpretation is that the neural network automatically learns how to distinguish the bounding boxes belonging to different objects. Note that a similar idea was proposed that combined geometric and semantic features in a Relation Network [22], but our approach is much simpler and can handle intra-class occlusion effectively.

We train the SGE using the loss function defined in Eq. 1. The training is carried out end-to-end, jointly with the object-detection branch using the loss function defined later (Eq. 6).

The loss function is derived for the SGE by extending the notion of an *associative embedding* [10,36]. Specifically, we use a *group loss* to group the SGEs of boxes belonging to the same object, and use a *separation loss* to distinguish SGEs of boxes belonging to different objects. For one image, the ground-truth boxes are denoted by  $B^* = \{b_1^*, b_2^*, \dots, b_M^*\}$ . For each refined box  $b_i$  in the refined box set  $B = \{b_i\}$ , let  $b_j^*$  be the ground truth box with the largest IoU. If  $\text{IoU}(b_i, b_j^*) > \theta$ , box  $b_i$  would be “assigned” to  $b_j^*$ . Thus the refined bounding boxes are divided into  $M + 1$  sets:  $B = B_1 \cup B_2 \cup \dots \cup B_{M+1}$ , where  $B_{M+1}$  is the set of refined boxes that are not assigned to any ground truth box. Then the *group* and *separation losses* are defined as:

$$L_{\text{group}}(\{e_i\}) = \sum_{j=1}^M \sum_{b_i \in B_j} |e_i - e_j^*|, \quad (2)$$

$$L_{\text{sep}}(\{e_i\}) = \sum_i p_i^* \max(0, \sigma - |e_i - \tilde{e}_i|), \quad (3)$$

where  $e_j^*$  is the SGE of ground truth box  $b_j^*$ ,  $\tilde{e}_i$  is the ground truth box with the second largest IoU with respect to  $b_i$ , and its SGE is  $\tilde{e}_i$ . We use  $\sigma$  to stabilize the

---

**Algorithm 1:** The proposed Semantics-Geometry NMS.

---

**Input:**  $B = \{b_i\}$ : List of detection locations (boxes),  
 $V = \{v_i\}$ : List of detection scores,  
 $E = \{e_i\}$ : List of SGEs,  
 $N_t$ : IOU threshold,  
 $\Phi(\cdot): \mathbb{R} \rightarrow \mathbb{R}$ , a monotonically increasing function

**begin**  
 $D \leftarrow \{\}$ ;  $W \leftarrow \{\}$  // set D as detected boxes and W as their scores  
**while**  $B \neq \emptyset$  **do**  
 $m \leftarrow \arg \max_{i \in \{1 \dots N\}} V$   
 $D \leftarrow D \cup \{b_m\}$ ;  $W \leftarrow W \cup \{s_m\}$   
 $B \leftarrow B \setminus \{b_m\}$ ;  $V \leftarrow V \setminus \{v_m\}$   
**for**  $b_i \text{ IN } B$  **do**  
 $\tau \leftarrow \text{IoU}(b_m, b_i)$   
// compare not only IoU but also the embedding distance  
**if**  $\tau \geq N_t$  AND  $d(e_m, e_i) \leq \Phi(\tau)$  **then**  
 $B \leftarrow B \setminus b_i$ ;  $V \leftarrow V \setminus v_i$   
**return** D, W

---

training process by preventing the distances between embeddings to be infinite. We found that the model performance is not sensitive to the actual value of  $\sigma$ . In the definition of the *separation loss*,  $p_i^*$  is a indicator variable which is 1 only if  $b_i \notin B_{M+1}$  and  $\text{IoU}(b_i, \tilde{b}_i^*) > \rho$ .

Some readers may confuse our loss functions with the Repulsion Loss (RL) [27], which is completely different. The RL was proposed to improve bounding box regression so that the detected bounding boxes better fit ground-truth objects. In contrast, our method does not affect the bounding box regression. The embedding trained through the two loss functions is used to determine if two overlapping boxes belong to the same object. Another difference is that the RL is performed in the box-coordinate space, while our group and separation losses are performed in the latent embedding space.

### 3.2 Semantics-Geometry Non-Maximum Suppression

We now derive our simple, yet effective NMS algorithm, SG-NMS, which takes advantage of the Semantics-Geometry Embedding. Its pseudo code is given in Algorithm 1.

SG-NMS first selects the box with the highest detection score as the *pivot box*. For each of the remaining boxes, its IoU with the pivot box is denoted by  $\tau$ , and the box will be kept if the  $\tau < N_t$ . When  $\tau > N_t$ , SG-NMS checks the distance between its SGE and the SGE of the pivot box. If the distance is larger than  $\Phi(\tau)$ , the box will also be kept. Here  $\Phi(\cdot)$  is a monotonically increasing function, which means that, as  $\tau$  increases, a larger distance is required to keep it. In this work, we consider three kinds of SG-NMS algorithms: SG-NMS-Constant,

SG-NMS-Linear and SG-NMS-Square, which respectively correspond to:

$$\Phi(\tau) = t_c, \quad \Phi(\tau) = t_l \cdot \tau, \quad \text{and} \quad \Phi(\tau) = t_s \cdot \tau^2, \quad (4)$$

where  $t_c$ ,  $t_l$ , and  $t_s$  are hyper-parameters.

### 3.3 The proposed Serial R-FCN

In order to compute SGEs that can capture the difference between geometric features of boxes belonging to different objects, we need to align extracted semantic features strictly with the refined boxes after bounding box regression. However, this cannot be achieved by normal two-stage CNN-based object detectors where the pooled feature is aligned with the RoI instead of the refined box because of the bounding-box regression.

To address this problem, we propose *Serial R-FCN*, see Fig. 2 (a). In Serial R-FCN, the classification head along with the SG module is placed *after* the classification-agnostic bounding box regression head [7]; thus, the whole pipeline becomes a serial structure. The classification head and the SG module use the refined boxes for feature extraction rather than the RoIs. Thus, the pooled features are strictly aligned with the refined boxes.

A light-weight self-attention branch is added into each head, as in Fig. 2 (b). The output of the attention head is a discrete distribution over the  $k^2$  position-sensitive grid. The position-sensitive scores are then aggregated through a weighted sum based on that distribution. There are two reasons why we introduced the self-attention in each head: 1). The self-attention helps the network to capture the semantic difference between heavily overlapping boxes and hence the SGE can be learned effectively. 2) we suggest that merging the position-sensitive scores by averaging (as done previous work [14]) could be sub-optimal, while adding the self-attention module helps the model to learn how to merge the score better. The idea of our Serial R-FCN is similar to a Cascade R-CNN [15]. However, while Cascade R-CNN stacks multiple classification and regression heads, we here only use one regression head and one classification head, thus do not introduce an extra parameter. Although the serial structure can be used by any two-stage detector, it suits the R-FCN best since no extra operation is added, and so the computation of the refined box is nearly cost free.

Placing the classification head after the regression head can bring us another benefit: It enables us to train the classification head using a higher IoU threshold. This yields more accurate bounding boxes. Without the serial structure, setting the IoU threshold to a very high value would result in the shortage of positive samples. However, in practice, we find that simply adopting the serial structure could easily yield a network that overfits on the training data. The reason is that as training progresses, the regression head becomes more and more powerful so that the classification head cannot receive enough hard negative examples (i.e., boxes whose IoU with the ground truth box is slightly smaller than the training threshold). The result is that the model cannot distinguish these examples and true positives when the model is tested. To alleviate the overfitting

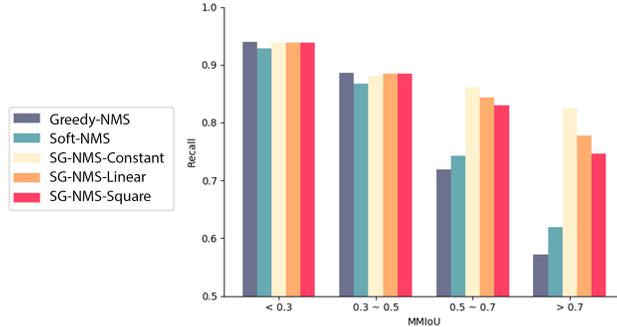


Fig. 3: Detection recall of the proposed SG-NMS and competing NMS algorithms on the KITTI validation set for different levels of occlusion, denoted by the maximal-IoU (MMIOU) among ground-truth boxes.

problem, we propose the simple but effective approach to add some noise to the refined bounding box so that the classification head continues to obtain hard false examples. Formally, during training, a box  $b = (x, y, w, h)$  is transformed to  $b' = (x', y', w', h')$  to train the classification head and the SG module:

$$\begin{aligned} x' &= \sigma_x w + x, & y' &= \sigma_y h + y, \\ w' &= w \cdot \exp(\sigma_w), & h' &= h \cdot \exp(\sigma_h), \end{aligned} \quad (5)$$

where  $\sigma_x, \sigma_y, \sigma_w, \sigma_h$  are noise coefficient drawn from a uniform distribution  $\prod_{j=1}^{j=k} (-\zeta_k, \zeta_k)$  where the four dimensions correspond to  $x, y, w, h$  respectively. In practice we set  $\zeta_x = \zeta_y = 0.05$  and  $\zeta_w = \zeta_h = 0.2$ .

The whole pipeline is trained end-to-end with the loss functions

$$L_{\text{total}} = L_{\text{rpn}} + \alpha L_{\text{det}} + \beta L_{\text{SGE}}, \quad (6)$$

$$L_{\text{rpn}} = L_{\text{cls-anchor}} + L_{\text{reg-anchor}}, \quad (7)$$

$$L_{\text{det}} = L_{\text{cls-rbox}} + L_{\text{reg-roi}}, \quad (8)$$

$$L_{\text{SGE}} = L_{\text{group-rbox}} + L_{\text{sep-rbox}}, \quad (9)$$

where the  $L_{\text{rpn}}$  is the commonly used loss to train the Region Proposal Network (RPN) [8],  $L_{\text{det}}$  is object detection loss [7] and  $L_{\text{SGE}}$  is the loss to train SGE as described in Sec. 3.1. We use two hyper-parameter  $\alpha$  and  $\beta$  to balance between losses (Eq. 6). The RPN classification and regression losses are applied to the anchor boxes (Eq. 7), the regression loss to RoIs (Eq. 8), and the classification, group, and separation losses to the refined boxes (Eq. 9).

## 4 Experiments

We conducted quantitative experiments on two commonly used urban scenes datasets: KITTI [11] and CityPersons [12]. To show the advantage of our *SG-Det model* and also to give deep insights into our approach, we first conducted

several experiments using different settings on the KITTI validation set. We then compared our approach with other state-of-the-art methods on the KITTI test set, and finally we show the performance on the CityPersons datasets. Results demonstrate the effectiveness of SG-Det to detect heavily-occluded cars and pedestrians in urban scenes. See also [www.cs.bu.edu/~betke/research/SG-Det](http://www.cs.bu.edu/~betke/research/SG-Det).

#### 4.1 Datasets

**KITTI** contains 7,481 images for training and validation, and another 7,518 images for testing. We evaluated our methods on the car detection task where intra-class occlusions tend to happen the most. The dataset has a standard split into three levels of difficulty: Easy, Moderate, and Hard, according to the object scale, occlusion level, and maximum truncation. To further demonstrate how our methods handles intra-class occlusions, we proposed a new difficulty split that divide the dataset into disjoint subsets based on the max-mutual-IoU (MMIoU), denoted by  $MMIoU$ , between ground-truth boxes. The max-mutual-IoU of a ground-truth box is defined by its maximum IoU with other ground-truth boxes in the same category. We separate the validation set into three levels: Bare ( $0 < MMIoU \leq 0.2$ ), Partial ( $0.2 < MMIoU \leq 0.5$ ) and Heavy ( $0.5 < MMIoU$ ). Average Precision (AP) is used to evaluate performance [11]. Following prior work [37], we randomly held out 3,722 images for validation and use the remaining 3,759 images for training, in which a simple image  $L_2$  similarity metric was adopted to differentiate training and validation images.

**CityPersons** contains 5,000 images (2,975 for training, 500 for validation, and 1,525 for testing). The log-average Miss Rate (MR) is used to evaluate performance. Following [27], we compare the detection log-average Miss Rate (MR) in different occlusion degrees. Following prior work [27], we separated the data into four subsets according to occlusion degree.

#### 4.2 Implementation Details

We implemented our Serial R-FCN in TensorFlow [38] and trained it on a Nvidia Titan V GPU. For KITTI, we chose a ResNet-101 [2] based on a Feature Pyramid Network (FPN) as the backbone and set the batch size to 4. The model was trained for 100,000 iterations using the Adam [39] optimizer with learning rate of 0.0001. For CityPersons, we chose a ResNet-50 [2] as the backbone network and trained the model for 240,000 iterations with batch size of 4, and the initial learning rate was set to 0.0001 and decreased by a factor of 10 after 120,000 iterations. In all experiments, OHEM is adopted to accelerate convergence [14]. For both datasets, we set  $\theta$ ,  $\sigma$ ,  $\rho$  to 0.7, 0.3, and 1.0, respectively, and set  $\alpha$  and  $\beta$  to 1. Our code is available <https://github.com/ChenhongyiYang/SG-NMS>.

#### 4.3 Effectiveness of SG-NMS

We report the performance of different NMS algorithms on the KITTI validation set applied to the same initial boxes so that a fair comparison is ensured (Table 1). For Soft-NMS, we only report the results of the linear version because we

Algorithm	Easy	Moderate	Hard	↑
Greedy-NMS	<b>97.98</b>	95.16	90.21	
Soft-NMS	97.72	95.13	91.15	
SG-Constant	97.56	95.35	92.31	
SG-Linear	97.69	<b>95.41</b>	<b>92.54</b>	
SG-Square	97.52	95.14	92.38	

Table 1: Average precision (AP) in % of the proposed SG-NMS algorithm and other commonly-used NMS algorithms on the KITTI validation set.

$N_t$	Greedy			SG-Constant			SG-Linear			SG-Square						
	Bare	Partial	Heavy	$c$	Bare	Partial	Heavy	$t_l$	Bare	Partial	Heavy	$t_s$	Bare	Partial	Heavy	↑
soft	<b>94.63</b>	84.62	54.62	1.2	<b>93.77</b>	84.49	57.21	2.0	<b>93.74</b>	85.14	58.72	3.0	<b>93.65</b>	<b>85.23</b>	60.36	
0.3	94.33	76.56	35.10	1.1	93.73	<b>84.65</b>	58.74	1.9	93.72	<b>85.19</b>	59.17	2.9	93.62	85.16	60.24	
0.4	94.03	83.38	40.58	1.0	93.46	84.33	60.02	1.8	93.70	85.11	59.81	2.8	93.61	85.14	60.83	
0.5	93.63	<b>85.15</b>	50.63	0.9	93.58	84.52	<b>60.05</b>	1.7	93.68	85.10	60.19	2.7	93.59	85.17	61.48	
0.6	91.56	82.85	55.49	0.8	93.46	84.33	60.02	1.6	93.64	85.06	60.24	2.6	93.58	85.09	<b>62.08</b>	
0.7	46.25	27.24	<b>57.21</b>	0.7	93.31	83.83	59.05	1.5	93.60	84.97	<b>61.08</b>	2.5	93.55	84.98	62.03	

Table 2: AP (in %) of NMS algorithms with different thresholds and occlusion levels (highest AP per level in red).

find its performance is consistently better than the Gaussian version. All three SG-NMS algorithms outperform the Greedy-NMS and Soft-NMS on the Moderate and Hard levels. In particular, SG-NMS-Linear outperforms Greedy-NMS and Soft-NMS by 2.33 pp and 1.39 pp, respectively, on the Hard level where heavy intra-class occlusions occur. We also explored the efficacy of the Relation Network [22] in occlusion situations, but found that it did not work well due to generating numerous false positive detections in crowded scenes.

We report the detection recall on different MMIOU intervals and show the results in Fig. 3. When MMIOU is less than 0.5, the tested NMS algorithms achieve similar recall scores. When there is severe intra-class occlusion, i.e.,  $\text{MMIOU} > 0.5$ , the recall of Greedy-NMS and Soft-NMS drops significantly. However, all three SG-NMS keep a relatively high recall. When  $\text{MMIOU} > 0.5$ , the difference in recall among the three SG-NMS algorithms is caused by the different slope of their  $\Phi(\cdot)$  function. This result indicates that our SG-NMS improves the detection by promoting detection recall for objects in crowded scenes.

We report how the hyper-parameter  $t$ , introduced by our SG-NMS, affects detection performance (Table 2). Overall, the variants of SG-NMS outperform Greedy-NMS and Soft-NMS for the Heavy and Partial occlusion levels, while maintaining high performance for the Bare level. For the Heavy level, the best result, 62.08%, is achieved by SG-NMS-Square, which is 4.87 percent points (pp) higher than the best result of Greedy-NMS and Soft-NMS. Although Greedy-NMS can achieve an AP of 55.49% for the Heavy level (when  $N_t = 0.6$ ), the AP

Model	SG	Noise	Attention	Easy	Moderate	Hard	Bare	Partial	Heavy $\uparrow$
R-FCN (FPN)	-	-	-	95.57	95.08	88.66	92.41	81.83	45.96
Our SG-Det			✓	94.77	94.44	89.30	92.35	80.94	44.03
		✓		95.84	94.55	90.10	93.27	84.87	47.79
			✓	95.04	95.12	90.01	93.11	83.28	43.98
	✓			94.62	94.50	89.54	92.45	81.56	52.62
		✓	✓	<b>97.98</b>	95.16	90.21	<b>93.63</b>	85.15	50.63
	✓	✓		97.80	<b>95.24</b>	91.86	93.15	82.71	51.58
	✓		95.25	94.50	92.30	93.21	84.62	58.43	
	✓	✓	97.52	95.14	<b>92.38</b>	93.59	<b>85.17</b>	<b>61.48</b>	

Table 3: AP for different settings for the proposed SG-Det model and a baseline R-FCN model on car detection on the KITTI validation set. SG stands for SG-NMS; Noise stands for box noise, Attention stands for the self-attention branch used in each head.

in the Bare and Partial levels drops significantly due to the false-positive boxes it generates.

#### 4.4 Ablation Study

We conducted an ablation study that demonstrates how the different model components affect the overall detection performance (Table 3). Our SG-Det model is proposed for detecting occluded objects, thus the analysis is focused on the detection of objects at the Hard difficulty (in the official split) and the Heavy occlusion level.

When the self-attention and bounding box noise are removed from our Serial R-FCN, we obtain a baseline Serial R-FCN that achieves an AP of 89.30% on the Hard and 44.03% on the Heavy occlusion level. When SG-NMS is included, the detection AP on the Heavy level is improved by 8.59 pp. When the self-attention branch is added into each head, the detection AP in the Hard and Heavy levels is lifted by 0.8 pp and 3.76 pp, respectively, compared to the baseline Serial R-FCN. This verifies our assumption that the learnable score aggregation enabled by the self-attention is superior to the naive average aggregation. By adding SG-NMS, the APs are further improved to 92.3% and 58.43%, which indicates that the self-attention head is important in capturing the semantic difference between heavily overlapping boxes. By adding box noise during training, the detection APs for all settings are improved, except for the heavy occlusion level. This means that the box noise can improve the detection precision by alleviating the overfitting problem in the Serial R-FCN, but it cannot help with improving the detection recall for heavily occluded objects. By combining self-attention, box noise, and SG-NMS, the full SG-Det model achieves APs of 92.38% and 61.48% on the Hard difficulty and Heavy occlusion level, respectively.

To conclude, we note that self-attention is useful to capture the semantic difference between heavily overlapping boxes. The box noise can alleviate the

Embedding	Bare	Partial	Heavy $\uparrow$
SE	93.31	83.95	55.60
GE	<b>94.11</b>	78.53	38.12
SGE	93.59	<b>85.14</b>	<b>61.48</b>

Table 4: Comparison of AP between the proposed Semantics-Geometry Embedding (SGE) , the pure Semantic Embedding (SE) and the pure Geometric Embedding (GE) on the KITTI validation set.

$\rho$	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40	0.45	0.50
Bare	93.92	94.03	92.50	<b>94.24</b>	93.28	93.68	92.98	93.56	93.17	93.07
Partial	81.73	82.89	83.75	84.06	<b>85.70</b>	85.10	84.89	82.38	83.24	83.52
Heavy	53.42	54.73	58.37	57.08	60.05	<b>60.19</b>	58.93	56.05	55.11	53.25

Table 5: Comparison of AP using different  $\rho$  during training.

overfitting problem so that the detection precision is improved and the SG-NMS algorithm can improve the detection performance for heavily occluded objects.

#### 4.5 Discussion

**The importance of Semantics and Geometry.** We explored the importance of the semantic and geometric features by removing them from the embedding calculation. We first removed the semantic features by computing a Geometric Embedding (GE) for each box, where the GE is computed using a fixed  $\hat{s}$  that is the mean of all the  $s$  vectors in the validation set. The performance of GE, shown in Table 4, is inferior than our SGE in occlusion situations, demonstrating the benefit of computing semantic features adaptively. Then we tested the purely-semantic model: for every box, a 1D Semantic Embedding (SE) is computed directly from its pooled region feature (Table 4). Our SGE performs better than the SE for all three occlusion levels. In fact, the two loss functions, defined in Sec. 3.1 for the SE, produce very unstable results during training. This means it is difficult for the neural network to learn such an embedding based on semantic features only, and it reveals the benefit of including geometric features.

**How to set  $\rho$  when training SGE?** We use a hyper-parameter  $\rho$  to determine occlusion during training (Sec. 3.1): For a detected box  $b$ , if its second largest IoU with any ground-truth box is larger than  $\rho$ , we assert  $b$  is occluded or occludes another object. Thus, the value of  $\rho$  becomes critical to the performance. In Table 5, we report the AP on different  $\rho$  using SG-NMS-Linear with  $t_l = 1.7$ . The results show that the performance on the bare difficulty level does not depend on  $\rho$ , which is reasonable because our SGE and SG-NMS do not affect objects without occlusion. The best  $\rho$  for the partial and heavy difficulty levels

Model	Runtime ↓	Easy	Moderate	Hard ↑
RRC [37]	3.60 s	95.68	93.40	87.37
SenseKITTI [40]	4.50 s	94.79	93.17	84.38
SDP+RPN [41]	0.40 s	95.16	92.03	79.16
ITVD [42]	0.30 s	95.85	91.73	79.31
SINet+ [43]	0.30 s	94.17	91.67	78.6
Cascade MS-CNN [15]	0.25 s	94.26	91.60	78.84
LTN [44]	0.40 s	94.68	91.18	81.51
Aston-EAS [45]	0.24 s	93.91	91.02	77.93
Deep3DBox [46]	1.50 s	94.71	90.19	76.82
R-FCN(FPN) [14,47]	0.20 s	93.53	89.35	79.35
Ours	0.20 s	95.81	93.03	83.00

Table 6: Runtime and AP (%) on the KITTI test set as reported on the KITTI leaderboard. All methods are ranked based on Moderate difficulty.

are 0.25 and 0.3, so we suggest to use a  $\rho$  of 0.27. A different value for  $\rho$  leads to a decrease in performance. To explain this, we suggest that a low value of  $\rho$  brings too much noise into the computing of the *group loss*, while a high value of  $\rho$  results in the model failing to capture the semantic difference of overlapping boxes that belong to different objects.

#### 4.6 Comparison with Prior Methods

We compared our model with other state-of-the-art models on the KITTI car detection leaderboard (Table 6). Our Serial R-FCN and SG-NMS are ranked at the third place among the existing methods. The respective APs on the Moderate and Hard level are 1.00 pp and 3.84 pp higher than the fourth-place values [41]. Although RRC [37] and sensekitti [40] are ranked higher than ours, the speed of our method is more than ten times faster than theirs. A reason is that our main contribution focuses on the post-processing step rather than the detection pipeline.

#### 4.7 Experiments on CityPersons

We compare miss rates of NMS algorithms on the CityPersons validation set for different occlusion degrees in Table 7. We also compare our model with existing methods. The NMS hyper-parameters are obtained from a grid search, and we report the best result for each algorithm. With Greedy-NMS, our Serial R-FCN achieves miss rates of 11.7% (reasonable difficulty level) and 52.4% (heavy). Using Soft-NMS yields a slight improvement. SG-NMS-Linear and SG-NMS-Square yield 0.2 and 0.7 pp improvements (reasonable difficulty), but using SG-NMS-Constant harms the performance for this level because a single threshold cannot handle the various complex occlusion situations. All three SG-NMS improve performance on the heavy and partial occlusion levels. Especially, the

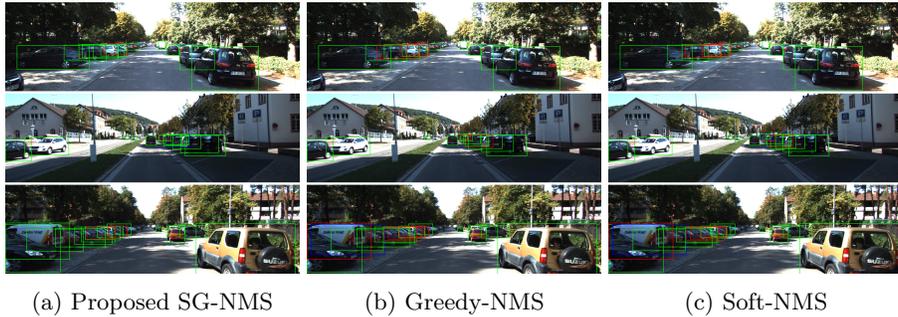


Fig. 4: Visualization of results with true positive (green), false positive (blue) and missed (red) detections. Two failure cases of SG-NMS are shown with a false positive detection ((a) middle), and a missed detection ((a) bottom).

Model	Reasonable	Bare	Partial	Heavy ↓
Adapted FasterRCNN [12]	12.8	-	-	-
Repulsion Loss [27]	11.6	7.0	14.8	55.3
OR-CNN [17]	11.0	<b>5.9</b>	13.7	51.3
Adaptive-NMS [20]	<b>10.8</b>	6.2	11.4	54.0
SR-FCN+Greedy-NMS	11.7	7.5	11.0	52.4
SerialR-FCN+Soft-NMS	11.4	7.1	10.9	51.8
SR-FCN+SG-Constant	11.5	7.4	11.2	52.3
SR-FCN+SG-Linear	11.3	7.3	10.8	51.6
SR-FCN+SG-Square	11.0	7.2	<b>10.7</b>	<b>51.1</b>

Table 7: The miss rate (%) on the CityPersons validation set.

SG-NMS-Square improves the respective miss rate to 10.7% and 51.1% on the partial and heavy occlusion levels, making our methods superior to the state of the art on those two levels. This means our method excels at handling occlusions.

## 5 Conclusion

In this paper, we presented two contributions, a novel Semantics-Geometry Embedding mechanism that operates on detected bounding boxes and an effective Semantics-Geometry Non-Maximum-Suppression algorithm that improves detection recall for heavily-occluded objects. Our combined model SG-Det achieves state-of-the-art performance on KITTI and CityPersons datasets by dramatically improving the detection recall and excelling in a low run time.

**Acknowledgements** We acknowledge partial support of this work by the MURI Program, N00014-19-1-2571 associated with AUSMURIB000001, and the National Science Foundation under Grant No. 1928477.

## References

1. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, Ieee, 2009.
2. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
3. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014.
4. J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788, 2016.
5. W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single shot multibox detector," in *European Conference on Computer Vision*, pp. 21–37, Springer, 2016.
6. R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 580–587, 2014.
7. R. Girshick, "Fast R-CNN," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1440–1448, 2015.
8. S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems*, pp. 91–99, 2015.
9. 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*, pp. 5561–5569, 2017.
10. A. Newell, Z. Huang, and J. Deng, "Associative embedding: End-to-end learning for joint detection and grouping," in *Advances in Neural Information Processing Systems*, pp. 2277–2287, 2017.
11. A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3354–3361, IEEE, 2012.
12. 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*, pp. 3213–3221, 2017.
13. 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*, pp. 2980–2988, 2017.
14. J. Dai, Y. Li, K. He, and J. Sun, "R-FCN: Object detection via region-based fully convolutional networks," in *Advances in Neural Information Processing Systems*, pp. 379–387, 2016.
15. Z. Cai and N. Vasconcelos, "Cascade R-CNN: Delving into high quality object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6154–6162, 2018.
16. J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *International Journal of Computer Vision*, vol. 104, no. 2, pp. 154–171, 2013.
17. 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)*, pp. 637–653, 2018.

18. C. L. Zitnick and P. Dollár, “Edge boxes: Locating object proposals from edges,” in *European Conference on Computer Vision*, pp. 391–405, Springer, 2014.
19. K. He, X. Zhang, S. Ren, and J. Sun, “Spatial pyramid pooling in deep convolutional networks for visual recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 9, pp. 1904–1916, 2015.
20. J. Hosang, R. Benenson, and B. Schiele, “Learning non-maximum suppression,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4507–4515, 2017.
21. 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*, pp. 6459–6468, 2019.
22. 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*, pp. 3588–3597, 2018.
23. L. Tychsen-Smith and L. Petersson, “Improving object localization with fitness nms and bounded IoU loss,” *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun 2018.
24. Y. He, C. Zhu, J. Wang, M. Savvides, and X. Zhang, “Bounding box regression with uncertainty for accurate object detection,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2888–2897, 2019.
25. 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)*, pp. 784–799, 2018.
26. Z. Tan, X. Nie, Q. Qian, N. Li, and H. Li, “Learning to rank proposals for object detection,” in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
27. 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*, pp. 7774–7783, 2018.
28. S. Zhang, J. Yang, and B. Schiele, “Occluded pedestrian detection through guided attention in CNNs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6995–7003, 2018.
29. J. Noh, S. Lee, B. Kim, and G. Kim, “Improving occlusion and hard negative handling for single-stage pedestrian detectors,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 966–974, 2018.
30. 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*, pp. 1904–1912, 2015.
31. 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*, pp. 3486–3495, 2017.
32. E. Goldman, R. Herzig, A. Eisenschat, J. Goldberger, and T. Hassner, “Precise detection in densely packed scenes,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
33. C. Zhou and J. Yuan, “Bi-box regression for pedestrian detection and occlusion estimation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 135–151, 2018.
34. Y. Pang, J. Xie, M. H. Khan, R. M. Anwer, F. S. Khan, and L. Shao, “Mask-guided attention network for occluded pedestrian detection,” in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4967–4975, 2019.

35. K. Zhang, F. Xiong, P. Sun, L. Hu, B. Li, and G. Yu, "Double anchor R-CNN for human detection in a crowd," *arXiv preprint arXiv:1909.09998*, 2019.
36. H. Law and J. Deng, "Cornersnet: Detecting objects as paired keypoints," in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 734–750, 2018.
37. J. Ren, X. Chen, J. Liu, W. Sun, J. Pang, Q. Yan, Y.-W. Tai, and L. Xu, "Accurate single stage detector using recurrent rolling convolution," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5420–5428, 2017.
38. M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.
39. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
40. B. Yang, J. Yan, Z. Lei, and S. Z. Li, "Craft objects from images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6043–6051, 2016.
41. F. Yang, W. Choi, and Y. Lin, "Exploit all the layers: Fast and accurate cnn object detector with scale dependent pooling and cascaded rejection classifiers," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2129–2137, 2016.
42. W. Liu, S. Liao, W. Hu, X. Liang, and Y. Zhang, "Improving tiny vehicle detection in complex scenes," in *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6, IEEE, 2018.
43. X. Hu, X. Xu, Y. Xiao, H. Chen, S. He, J. Qin, and P.-A. Heng, "Sinet: A scale-insensitive convolutional neural network for fast vehicle detection," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 3, pp. 1010–1019, 2018.
44. T. Wang, X. He, Y. Cai, and G. Xiao, "Learning a layout transfer network for context aware object detection," *IEEE Transactions on Intelligent Transportation Systems*, 2019.
45. J. Wei, J. He, Y. Zhou, K. Chen, Z. Tang, and Z. Xiong, "Enhanced object detection with deep convolutional neural networks for advanced driving assistance," *IEEE Transactions on Intelligent Transportation Systems*, 2019.
46. A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka, "3d bounding box estimation using deep learning and geometry," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7074–7082, 2017.
47. 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*, pp. 2117–2125, 2017.