

Orientation-aware Vehicle Re-identification with Semantics-guided Part Attention Network

Tsai-Shien Chen^{1,2}, Chih-Ting Liu^{1,2}, Chih-Wei Wu^{1,2}, and Shao-Yi Chien^{1,2}

¹ Graduate Institute of Electronic Engineering, National Taiwan University

² NTU IoX Center, National Taiwan University

{tschen, jackieliu, cwwu}@media.ee.ntu.edu.tw, sychien@ntu.edu.tw

Abstract. Vehicle re-identification (re-ID) focuses on matching images of the same vehicle across different cameras. It is fundamentally challenging because differences between vehicles are sometimes subtle. While several studies incorporate spatial-attention mechanisms to help vehicle re-ID, they often require expensive keypoint labels or suffer from noisy attention mask if not trained with expensive labels. In this work, we propose a dedicated Semantics-guided Part Attention Network (SPAN) to robustly predict part attention masks for different views of vehicles given only image-level semantic labels during training. With the help of part attention masks, we can extract discriminative features in each part separately. Then we introduce Co-occurrence Part-attentive Distance Metric (CPDM) which places greater emphasis on co-occurrence vehicle parts when evaluating the feature distance of two images. Extensive experiments validate the effectiveness of the proposed method and show that our framework outperforms the state-of-the-art approaches.

Keywords: Vehicle re-identification, spatial attention, semantics-guided learning, visibility-aware features

1 Introduction

Vehicle re-identification (re-ID) aims to match vehicle images in a camera network. Recently, this task has drawn increasing attention due to practical applications such as urban surveillance and traffic flow analysis. While deep Convolutional Neural Networks (CNN) have shown remarkable performance in vehicle re-ID over the years [22, 23, 33], various challenges still hinder the performance of vehicle re-ID. One of them is that a vehicle captured from different viewpoints usually has dramatically different visual appearances. On the other hand, two different vehicles of the same color and car model are likely to have very similar appearances. As illustrated in the left part of Fig. 1, it is challenging to distinguish vehicles by comparing the features extracted from the whole vehicle images. In such case, the minor differences in specific parts of vehicle such as decorations or license plates would be a great benefit to identifying two vehicles. Furthermore, when two vehicles are presented in different orientations, a desired vehicle re-ID algorithm should be able to focus on the parts (views) that both

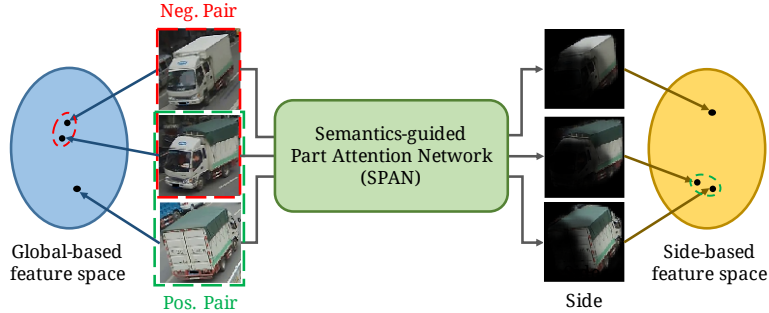


Fig. 1: **Concept illustration of Semantics-guided Part Attention Network.** The example images show intra-class difference and inter-class similarity in the vehicle re-ID problem. It is challenging to separate the negative images merely based on global feature due to the similar car model and viewpoint. In this example, it is easier to distinguish two vehicles by the side-based feature. This motivates us to generate the part (view) attention maps and then emphasize the feature of the co-occurrence vehicle parts for better re-ID matching.

appear in the two vehicle images. For example, in the right part of Fig. 1, it is easier to distinguish the vehicles by comparing their side views. To reach this idea, we divide it into two steps.

The first step is to extract the feature from specific parts of vehicle images. A number of work has been proposed to achieve this purpose by learning orientation-aware features. Nonetheless, existing methods either rely on expensive vehicle keypoints as guidance to learn an attention mechanism for each part of a vehicle [34, 11] or use only viewpoint labels but produce noisy and unsteady attention outcome which will thus hinder the network to learn subtle differences between vehicles [42]. In this paper, we introduce the *Semantics-guided Part Attention Network (SPAN)* to generate attention masks for different parts (front, side and rear views) of a vehicle. As shown in Fig. 1, our SPAN learns to produce meaningful attention masks. The masks not only help disentangle features of different viewpoints but also improve the interpretability of our learning framework. It is also worth noting that, instead of expensive keypoints or pixel-level labels for training, our SPAN requires only *image-level viewpoint labels* which are much easier to be derived from known camera pose and traffic direction.

For the second step, we design a *Co-occurrence Part-attentive Distance Metric (CPDM)* to better utilize the part features when measuring the distance of images. The intuition of this metric is that the network should focus on the parts (views) that both appear in the compared vehicle images. Therefore, the proposed metric allows us to automatically adjust the importance of each part feature distance according to the part visibility in two compared vehicle images.

We conduct experiments on two large-scale vehicle re-ID benchmarks and demonstrate that our method outperforms current state-of-the-arts. Ablation

studies prove that the attention masks generated by SPAN extract helpful part features and our CPDM can better utilize the global and part features to improve the re-ID performance. Moreover, qualitative results show that our SPAN can robustly generate meaningful attention maps on vehicles of different types, colors, and orientations. We now highlight our contributions: (1) We propose a Semantics-guided Part Attention Network (SPAN) to generate robust part attention masks which can be used to extract more discriminative features. (2) Our SPAN only needs image-level viewpoint labels instead of expensive keypoints or pixel-level annotations for training. (3) We introduce the Co-occurrence Part-attentive Distance Metric (CPDM) to facilitate vehicle re-ID by focusing on the parts that jointly appear in the compared images. (4) Extensive experiments on public datasets validate the effectiveness of each component and demonstrate that our method performs favorably against state-of-the-art approaches.

2 Related Work

Re-Identification (*re-ID*). Re-identification studies the problem of identifying identities in different camera views. There are large numbers of studies that focus on re-identifying human [39, 38, 3, 15] and vehicles [29, 34, 42, 20]. Most re-ID methods can be categorized into two types: feature learning and distance metric learning. Feature learning methods [38, 15, 37, 3, 30, 39] aim to learn a more discriminative embedding space. Distance metric learning methods [35, 41, 4, 12, 2] design distance functions for comparing features of two images. In this work, we design an orientation-aware feature extraction network as well as an orientation-aware distance metric for solving the vehicle re-ID problem.

Vehicle Re-Identification. Vehicle re-ID has received more attention for the past few years due to the releases of large-scale annotated vehicle re-ID datasets. Liu *et al.* [22, 23] released a high-quality multi-viewed VeRi-776 dataset. Tang *et al.* [33] proposed a city-scale traffic camera CityFlow dataset. With several datasets, numerous vehicle re-ID methods have been proposed recently. Some methods use CNN model to tackle the vehicle re-ID problem [29, 20, 33]. However, those methods lack spatial guidance and could be hard to distinguish two similar vehicles with only subtle difference. In contrast, the others adopt the extra information, such as viewpoint or keypoint labels, to generate spatial attentive features. Wang *et al.* [34] and Khorramshahi *et al.* [11] used 20 vehicle keypoints to generate attention maps by categorizing keypoints into four groups which respectively represent front, rear, left or right view of vehicle. Yet, the keypoint information is hard to acquire in real-world scenarios. Also, the keypoint is insufficient to cover all crucial features. Zhou *et al.* [42] proposed a viewpoint-aware attention model to produce attention map for different viewpoints and further generate multi-view features from single view input image. However, due to the lack of direct supervision on the generated attention maps, the attention outcomes are noisy and would unfavorably affect the learning of network. In contrast, we design a dedicated network and adopt specific loss functions to

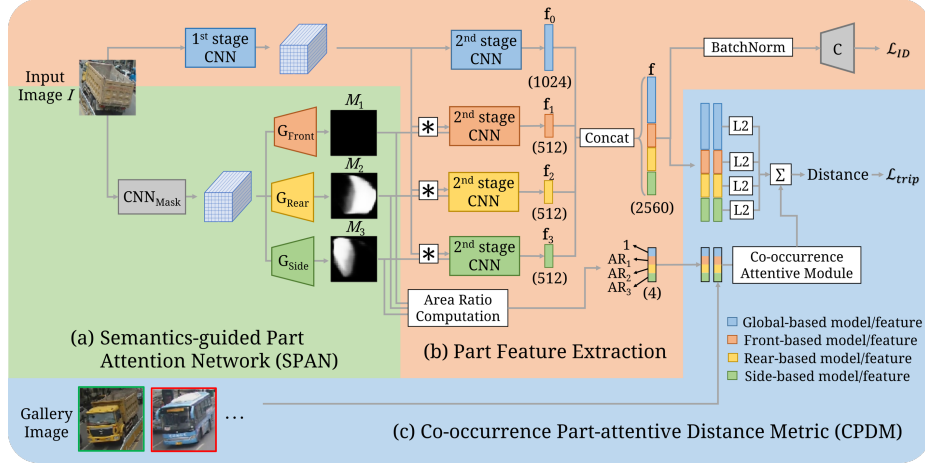


Fig. 2: **Architecture of our proposed framework.** (a) Semantics-guided Part Attention Network (SPAN) generates the attention masks for each part (view) of a vehicle image. (b) With the attention masks generated by SPAN, Part Feature Extraction produces one global and three part attentive features which are then concatenated into a representative feature. (c) Co-occurrence Part-attentive Distance Metric (CPDM) calculates a weighted feature distance with emphasis on the vehicle parts that appear in both compared images.

supervise the generation of attention maps. Moreover, our network only requires image-level viewpoint labels rather than keypoint labels during training.

Visibility-Aware Features. Utilizing visibility-aware features has gained growing interest considering that there are lots of occluded images in real-world scenarios. Sun *et al.* [31] and Miao *et al.* [5] pre-define several regions among whole images by horizontally or vertically partitioning the images and then produce the confidence score for each region to represent their visibility. However, the visibility of pre-defined region is hard to represent its importance for re-ID matching. For example, the highly visible regions but containing mostly background would be overemphasized while the smaller regions but containing some critical appearances would be neglected. To avoid the issue mentioned above, in this work, we directly use the visibility of specific parts of vehicle to represent its importance. Note that it is only possible when the specific parts are accurately located.

3 Proposed Method

The proposed learning framework for vehicle re-ID consists of three sub-modules as depicted in Fig. 2. First, we learn a Semantics-guided Part Attention Network (SPAN) to predict the attention masks for each part (view) of a vehicle in Sec. 3.1. Then, in Sec. 3.2, we apply the attention masks to our main feature

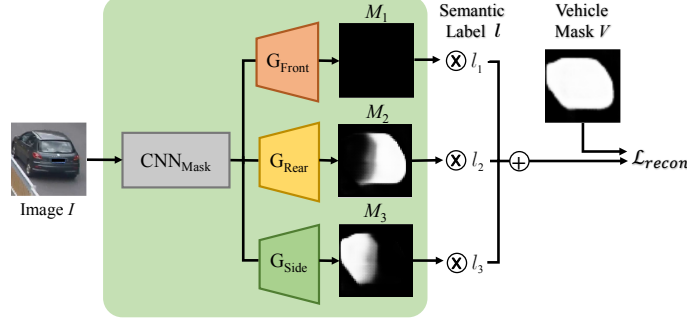


Fig. 3: **Mask Reconstruction Loss.** The part masks selected by semantic label should jointly reconstruct the whole foreground vehicle mask.

extraction network to generate part features in addition to the global features. During both training and inference, the global and part features are combined to evaluate feature distance between two vehicle images with our proposed Co-occurrence Part-attentive Distance Metrics (CPDM) in Sec. 3.3. Last, the overall model learning scheme of our framework is introduced in Sec. 3.4.

3.1 Semantics-guided Part Attention Network

The goal of our Semantics-guided Part Attention Network (SPAN) is to generate a set of attention masks for different parts (e.g. front, side, and rear view) of a vehicle image. An intuitive approach would be to train a segmentation network with pixel-wise view labels to predict segmented part masks. However, pixel-level annotation is expensive to obtain in real-world data. Instead, we turn to the image-level semantic labels, such as the viewpoint of vehicles which are much easier to be derived from known camera pose and the traffic direction, to learn our attention network. Given a vehicle image I , we define its corresponding semantic label vector as $l \in \mathcal{R}^3$. The semantic label l is encoded from its viewpoint. Its elements represent whether the front, rear or side view of image I are visible or not, respectively. To be more specific, $l_i = 1$ if the i^{th} view is visible, while $l_i = 0$ if it is not. For example, for a vehicle image with the front-side viewpoint, its semantic label vector l will be assigned with $[1, 0, 1]$.

As shown in Fig. 2 (a), our network predicts the attention masks of front, rear and side views M_1, M_2, M_3 with a shared feature extractor CNN_{Mask} and three mask generators G_{Front} , G_{Rear} , and G_{Side} . To ensure our SPAN generating ideal masks, we meticulously design a novel loss function, named mask reconstruction loss, with two auxiliary losses to supervise the learning of network.

Mask Reconstruction Loss. As illustrated in Fig. 3, the main idea of mask reconstruction loss is that the attention masks selected by corresponding semantic labels should jointly reconstruct the foreground mask of a vehicle. For instance, if the image is with rear-side viewpoint, the rear and side masks should jointly

reconstruct the whole vehicle foreground mask to the greatest extent possible. To this end, we first need the foreground mask of each vehicle image, which is also automatically generated by our deep segmentation network trained with the preliminary results by GrabCut [28] as the target. The detail of generating foreground masks is shown in the supplementary material; notes that any manually annotated pixel-level label is **not** required here. Thus, with the foreground masks (denoted as V), our mask reconstruction loss can be written as:

$$\mathcal{L}_{recon} = \|V - \sum_{i=1}^3 (l_i \times M_i)\|_2, \quad (1)$$

which represents the mean square error (MSE) between the foreground mask and the generated mask gated by the semantic label.

Area Constraint Loss. While imposing the mask reconstruction loss, we note that the training is unstable and often leads to undesired results. Take the qualitative result in Fig. 6 “w/o \mathcal{L}_{area} ” as example, we observe that, for a vehicle image with two visible views, the network only uses single representative mask generator to predict the whole vehicle mask. To prevent network from cheating, we design the area constraint loss to limit the maximum area of each predicted attention mask. Here, we define the area of mask as its L1-norm (sum of all elements) and also define the maximum area ratio of i^{th} view for a semantic label \mathbf{l} as $a_{\mathbf{l},i}$. Our area constraint loss can be formulated as:

$$\mathcal{L}_{area} = \sum_{i=1}^3 \left[\frac{\|M_i\|_1}{\|V\|_1} - a_{\mathbf{l},i} \right]_+, \quad (2)$$

where $\|\cdot\|_1$ represents L1-norm of a given mask. $[\cdot]_+$ is the hinge function since we only penalize the mask with the area ratio (over the whole foreground mask) larger than our expected ratio. For the setting of max area ratio a , the ratio of invisible parts should be 0 intuitively while the ratio of visible part should be 1 for images with merely one visible views. For images with two visible views, the ratio of each view should be set within the range from 0.5 to 1.

Spatial Diversity Loss. In addition to the situation mentioned above, we observe other unfavorable results. Such as the qualitative result in Fig. 6 “w/o \mathcal{L}_{div} ”, for a vehicle image with two visible views, the two corresponding mask generators may predict whole vehicle masks with values of 0.5. Therefore, similar to Li *et al.* [16], we introduce a spatial diversity loss to restrict the overlapped area between masks of different views with the following formulation:

$$\mathcal{L}_{div} = \sum_{(i,j) \in P} [(M_i \cdot M_j) - m_{i,j}]_+, \quad (3)$$

where $m_{i,j}$ is the margin representing the tolerable overlapped area between i^{th} and j^{th} view and P is the set of all view index pairs. For two mutually exclusive




Fig. 4: **Illustration of general-purposed SPAN.** (a)(b) show the output masks of the vehicle and multi-digit images with different semantic labels respectively. Our proposed SPAN is able to learn to generate the part attention map or localization only given the image-level semantic labels.

views, such as front and rear, the margin is set to 0 intuitively. For two adjacent views, such as front and side, the margin parameter is set to a positive value to tolerate the overlapped situation (e.g. front-side view mirror and headlight could be hard to uniquely assign to either front or side view).

Discussion. SPAN is general-purposed and can be extended to weakly-supervised segmentation which has much weaker supervision setting than regular segmentation because it only requires image-level label for training. It can also well perform on other datasets besides to vehicle images. Fig. 4 shows the example results on the multi-digit dataset based on MNIST [14] created by ourselves. For the multi-digit dataset, the semantic label represents which digit is visible in the image. Hence, the network can learn to generate the localization of each digit.

3.2 Part Feature Extraction

With the attention masks generated by our SPAN, we design a part feature extraction module to learn orientation-aware features for vehicle re-ID. As shown in Fig. 2 (b), the module includes two convolution stages. The 1st-stage CNN transforms input images into 1st-stage feature maps. Then, four distinct 2nd-stage CNNs respectively dedicated for extracting global-based, front-based, rear-based and side-based features follow the previous stage. The global-based model simply takes 1st-stage feature map as input and generates the global feature \mathbf{f}_0 . The other three branches apply the part attention masks to the 1st-stage feature map by element-wise matrix multiplication and then extract part features \mathbf{f}_1 , \mathbf{f}_2 and \mathbf{f}_3 by corresponding 2nd-stage CNNs. With one global feature and three part features, unlike previous methods [34, 42, 11] which embed all part features into one unified vector by additional network, our network simply concatenates them into one representative feature \mathbf{f} to best utilize all possible features for vehicles.

3.3 Co-occurrence Part-attentive Distance Metric

To fully utilize the part features extracted by our SPAN and part feature extraction module, we design the Co-occurrence Part-attentive Distance Metric

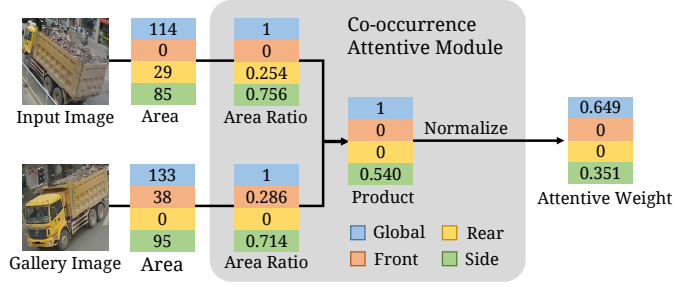


Fig. 5: **Co-occurrence Attentive Module.** To correctly recognize these two positive images, the feature of side view (co-occurrence view) should be emphasized, while front and rear feature should be relatively neglected. Co-occurrence attentive module is able to re-weigh the importance of each view accordingly.

(CPDM) for both training the CNN and matching images during inference. We note that, in addition to the global feature, the features of the same visible parts on different vehicles are also critical for re-ID. Moreover, the co-occurrence part with greater area ratio often represents higher clarity or is likely to include more key features in the original image. Therefore, we develop the Co-occurrence Attentive Module to re-weigh the importance of different feature distances by comprehensively considering the area ratio of each view in both images. Fig. 5 illustrates an example of Co-occurrence Attentive Module. Given a vehicle image, we first compute the area of global, front, rear and side view by calculating the L1-norm of the attention masks generated by SPAN (the area of global view is defined as the summation of the ones of front, rear and side view). The area ratios of each view are then normalized by the global area. We denote the area ratio of i^{th} view in image I_a as $AR_{a,i}$. For arbitrary two images I_a and I_b , the attentive weight of i^{th} view $w_{(a,b),i}$ can be written as:

$$w_{(a,b),i} = \frac{AR_{a,i} \times AR_{b,i}}{\sum_{i=0}^3 AR_{a,i} \times AR_{b,i}}. \quad (4)$$

Finally, we use the attentive weights to adjust the weighting for combining feature distances of all global and part features. The final distance $Dist_{(a,b)}$ between two vehicle images I_a and I_b is calculated by:

$$Dist_{(a,b)} = \sum_{i=0}^3 w_{(a,b),i} \times \|\mathbf{f}_{a,i} - \mathbf{f}_{b,i}\|_2, \quad (5)$$

which is the weighted summation of feature euclidean distances in each view.

Discussion. For two images with completely disjoint views, the attentive weights are all 0 for front, rear, and side views. Hence, the distance between I_a and I_b will be fully determined by their global features $\mathbf{f}_{a,0}$ and $\mathbf{f}_{b,0}$.

3.4 Model Learning Scheme

The learning scheme for our feature learning framework consists of two steps. In the first step (Fig. 2 (a)), we optimize our SPAN with the following loss:

$$\mathcal{L}_{step1} = \lambda_{recon}\mathcal{L}_{recon} + \lambda_{area}\mathcal{L}_{area} + \lambda_{div}\mathcal{L}_{div}. \quad (6)$$

Instead of training SPAN end-to-end with the re-ID feature extractor network [26], we train this network in advance because SPAN relies on clean viewpoint labels, which is not the case of our experimenting datasets. As a result, we train SPAN with a smaller dataset than the original one but with cleaner viewpoint labels.

In the second stage, we optimize the rest of our network (Fig. 2 (b)(c)) with two common re-ID losses while SPAN is fixed. The first one for metric learning is the triplet loss (\mathcal{L}_{trip}) [27], which is calculated based on the weighted distance introduced in Sec. 3.3. The other loss for the discriminative learning is the identity classification loss (\mathcal{L}_{ID}) [40]. The overall loss is computed as follows:

$$\mathcal{L}_{step2} = \lambda_{trip}\mathcal{L}_{trip} + \lambda_{ID}\mathcal{L}_{ID}. \quad (7)$$

During inference, given a query and a gallery image, we extract their features separately by SPAN and the part feature extraction module. The distance of the query and gallery images are then computed by our CPDM for re-ID matching.

4 Experiments

4.1 Datasets and Evaluation Metrics

Our framework is evaluated on two benchmarks, VeRi-776 [22, 23] and CityFlow-ReID [33], which are two large-scale vehicle re-ID datasets with multiple viewpoints. VeRi-776 dataset contains 776 different vehicles captured, which is split into 576 vehicles with 37,778 images for training and 200 vehicles with 11,579 images for testing. Wang *et al.* [34] released the annotated keypoints and viewpoint information for VeRi-776 dataset, which has been widely adopted by other work. In this paper, we only use the viewpoint labels to train our proposed SPAN. CityFlow-ReID is a subset of images sampled from the CityFlow dataset [33]. It consists of 36,935 images of 333 identities in the training set and 18,290 images of another 333 identities in the testing set. However, the viewpoint information of CityFlow-ReID is not available. Thus, we utilize the SPAN pre-trained on VeRi-776 to generate corresponding attention masks. Note that, though VehicleID [19] dataset is also a widely adopted benchmark, it only covers the images with front or rear viewpoint and cannot validate the effectiveness of our method. Hence, we would not use VehicleID in the following experiments.

As in previous vehicle re-ID works, we employ the standard metrics, namely the cumulative matching curve (CMC) and the mean average precision (mAP) [39] to evaluate the results. We report the rank-1 accuracy (R-1) in CMC and the mAP for the testing set in both datasets.

4.2 Implementation Details

For our SPAN (Fig. 2 (a)), we adopt the former four blocks in ResNet-34 [7] (*conv1* to *conv4*) as the feature extractor (CNN_{mask}) to extract the mid-level features which retain more spatial information than those after the last block (*conv5*). Afterwards, the feature map is fed into three generative blocks to generate the part masks. The detailed architecture of SPAN is shown in the supplementary material. This network is trained in advance on a subset of VeRi-776 dataset with balanced images in each viewpoint. For optimizing SPAN with \mathcal{L}_{step1} , the coefficients λ_{recon} and λ_{div} are set to 1 and λ_{area} is 0.5.

For our part feature extraction (Fig. 2 (b)), we adopt ResNet-50 [7] as our backbone which is split into two stages. The first four blocks (*conv1* to *conv4*) are in the first stage and the last block (*conv5*) with one fully-connected layer are in the second stage to generate a 1024-d or 512-d feature vector. For optimizing with triplet loss (\mathcal{L}_{trip}), we adopt the *PK* training strategy [8], where we sample $P = 8$ different vehicles and $K = 4$ images for each vehicle in a batch of size 32. In addition, for training identity classification loss (\mathcal{L}_{ID}), we adopt a BatchNorm [25] and a fully-connected layer as the classifier [25, 18]. The training process lasts for 30,000 iterations with λ_{trip} and λ_{ID} all set to 1 in \mathcal{L}_{step2} .

4.3 Ablation Studies and Visualization

In this section, to assess the effectiveness of our Semantics-guided Part Attention Network (SPAN) and Co-occurrence Part-attentive Distance Metric (CPDM), we conduct ablation studies quantitatively on VeRi-776 dataset and visualize the qualitative results of our attention masks compared with the existing methods.

Loss Functions of Our SPAN. We adopt three loss functions to help generating steady and clear attention masks when training SPAN. To evaluate the influence of each loss function, we conduct experiments with multiple combinations of losses and report the re-ID results on VeRi-776 in Table 1 and the corresponding qualitative results of our part attention masks in Fig. 6.

As listed in the first row in Table 1, we show the baseline method which simply transferred the whole vehicle image into a 1024-dim global feature and adopted euclidean distance as the feature distance metric. Except for the baseline method, all other methods in Table 1 adopt CPDM and utilize same architecture in SPAN but trained with different combinations of proposed loss functions. As shown in the second to fourth rows in Table 1 and the corresponding visualized attention masks in Fig. 6, the re-ID performance of those methods are almost the same as the baseline owing to the unfavorable generated attention masks, which cannot benefit the part feature extraction and the following CPDM. Only when simultaneously supervised by proposed three loss functions, our SPAN can generate clear and meaningful attention masks which can further improve the re-ID performance by a large margin as shown in the last row in Table. 1.

Method	Training Loss			VeRi-776	
	\mathcal{L}_{recon}	\mathcal{L}_{area}	\mathcal{L}_{div}	R-1	mAP
Baseline	-	-	-	92.0	59.1
only \mathcal{L}_{recon}	✓	✗	✗	91.8	58.9
w/o \mathcal{L}_{area}	✓	✗	✓	92.1	59.2
w/o \mathcal{L}_{div}	✓	✓	✗	92.5	59.7
SPAN(Ours)	✓	✓	✓	93.9	68.6

Table 1: **Ablation study of the loss** Fig. 6: **Results of SPAN training w/ functions for training SPAN (%)**. **different combinations of losses**.

Selection of Hyper-parameters in Loss Functions. There are two hyper-parameters which should be selected for loss functions, including max area ratio a in \mathcal{L}_{area} and margin m in \mathcal{L}_{div} . The physical meanings of selection have been discussed in Sec. 3.1. We finally choose $a = 0.7$ for the visible views of two-view images and $m = 0.04$ for two adjacent views based on the experimental results shown in the supplementary material.

Qualitative Results of Part Attention Masks. To verify the robustness of our proposed SPAN, we show some qualitative results of our part attention masks in Fig. 7 (a) and show the comparisons with Wang *et al.* [34] and Zhou *et al.* [42] (VAMI) in Fig. 7 (b)(c), respectively. In Fig. 7 (a), the produced masks from our SPAN can correctly cover all regional features which are belonging to their views while eliminates all redundant information such as features from background or other views. For example, headlights and front bumper are all covered in front mask, while door or background are not. In Fig. 7 (a), the demonstrated vehicles are all in different colors, types (sedan, SUV, pickup, truck, bus, etc.) and orientations, proving that SPAN is robust for various vehicles.

In contrast, the attention masks generated by the previous work [34, 42] are more noisy and unsteady. As shown in the left half of Fig. 7 (b), the front mask generated by Wang *et al.* [34] cannot cover the front windshield which possibly contains crucial features such as stickers or patterns. Also, in the right half of Fig. 7 (b), the front face of given vehicle image is not visible but the generated front attention mask fails to shield all features and instead activates on the background. The other example of generating unsteady masks is shown in Fig. 7 (c). Both rear and front masks generated by VAMI [42] fail to consistently embed the rear or front windshield among different vehicle images, which will make the network hard to distinguish two images based on those part features.

Component Analysis of the Proposed Model. Here, we report the re-ID performances to evaluate the effectiveness of each sub-module in our proposed framework in Table 2. The first row demonstrates the baseline model which

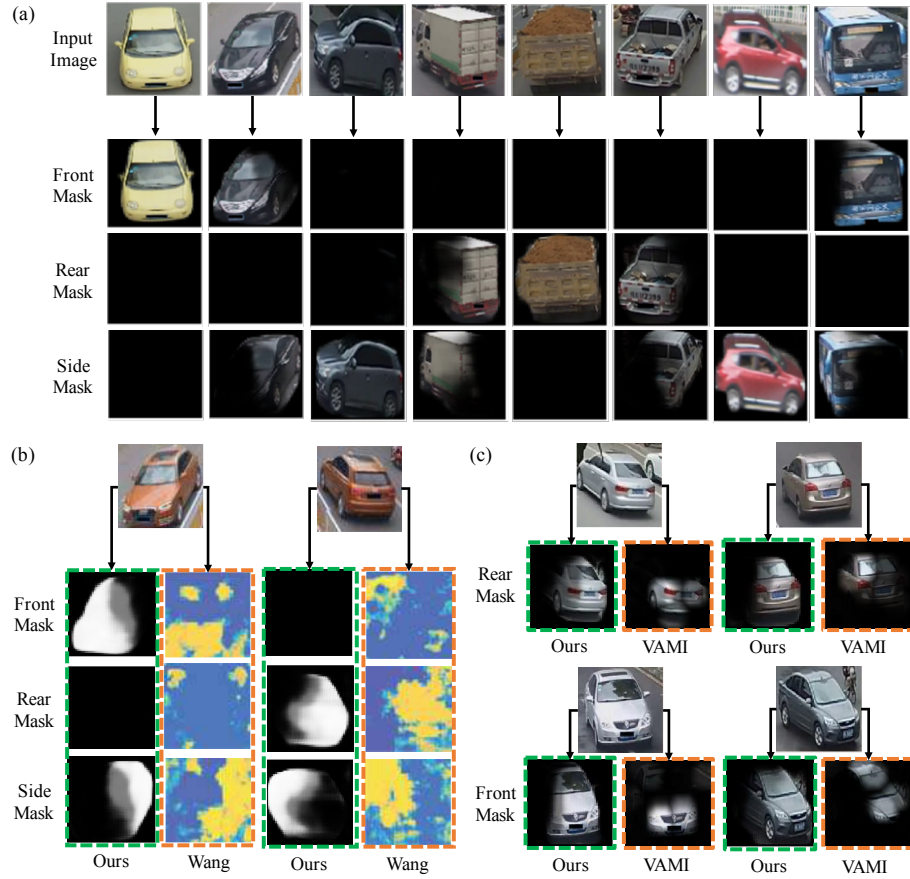


Fig. 7: **Qualitative Part Masks.** (a) shows some examples of the part masks generated by SPAN. Note that the demonstrated vehicles are all in different colors, types and orientations to verify the robustness of SPAN. (b)(c) show the comparison with Wang *et al.* [34] and Zhou *et al.* [42] (VAMI) respectively. The attention maps generated by their methods are directly from their papers.

simply transfers the whole vehicle image into a global feature and uses standard euclidean distance to evaluate the distance between two vehicles. Next, based on our SPAN, we conduct two experiments with different aggregation techniques to combine the global and part features into one vector. The first one utilizes an additional fully-connected layer (FC) to embed the whole features, as shown in the second row in Table 2 (SPAN w/ FC). The other directly concatenates all global and part features, as shown in the third row in Table 2 (SPAN w/ Cat). It shows that compared to the baseline method, the performances are all boosted with the global and part features jointly be utilized. However, concatenating all the features can retain more part information, which achieves better performance

Table 2: **Ablation studies of the proposed method in terms of R-1 and mAP (%)**. The effectiveness analysis for each component including the usage of SPAN, feature aggregation methods (Agg.) and distance metrics (Dist.).

Method	Sub-modules			VeRi-776	
	SPAN	Agg.	Dist.	R-1	mAP
Baseline	✗	-	Euc.	92.0	59.1
SPAN w/ FC	✓	FC	Euc.	92.6	60.3
SPAN w/ Cat	✓	Concat.	Euc.	93.0	63.1
SPAN w/ CPDM (Ours)	✓	Concat.	CPDM	94.0	68.9

for re-ID (from 59.1% to 63.1% than from 59.1% to 60.3% in mAP). Last, we report the results in the last row with the concatenated features and the usage of our CPDM, which is also our final proposed method (**SPAN w/ CPDM**). It shows that with our proposed method, the re-ID performance can outperform the baseline method by a large margin (**9.5%** in mAP), proving that CPDM can better utilize the global and part features to measure the distance between two vehicles by enhancing the importance of the co-occurrence part.

4.4 Comparison with the State-of-the-Arts

We compare our proposed framework with the state-of-the-art vehicle re-ID methods and report the results on VeRi-776 and CityFlow-ReID datasets in Table 3. Note that there are a few of recent works which cannot be fairly compared with ours due to different setting such as the usage of external vehicle re-ID dataset [17], manually annotated bounding boxes for crucial features [6] and large-scale synthetic dataset with various kinds of pixel-level annotations [32]. Therefore those works are not shown in our comparison in Table 3.

Previous vehicle re-ID methods can be mainly summarized into three categories: spatial-attentive feature learning [34, 42, 20, 11, 21, 9], distance metric learning [1] and embedding learning [24, 43]. For spatial-attentive feature learning, proposed methods attempted to guide the network focusing on the regional features which may be useful to distinguish from two vehicles. RAM [20], GRFGLL [21] and DFFMG [9] simply partitioned the images horizontally and vertically into several regions and extract the corresponding regional features; however, when the given images are in different orientations, the features would fail to consistently attends on same parts of vehicle. To extract orientation-aware features, OIFE [34] and AAVER [11] used extra expensive keypoints information to train their orientation-based region proposal network. Yet, they usually lose some informative information like the sticker on the windshield which is not covered by annotated keypoints. Instead, VAMI [42] used the viewpoint information to generate representative features of each viewpoint and used them to guide the network producing the viewpoint-aware attention maps and features, but the attention outcomes are not steady. To sum up, the unfavorable attention masks generated by existing work would hinder the re-ID performance on the

Table 3: **Comparison with state-of-the-arts re-ID methods on VeRi-776 and CityFlow-ReID dataset(%)**. Upper/Lower Group: methods **without/with** spatial-attentive mechanism. All listed scores are from the methods **without** adopting spatial-temporal information [23] or re-ranking [36].

Method	VeRi-776			CityFlow-ReID		
	R-1	R-5	mAP	R-1	R-5	mAP
EALN [24]	84.4	94.1	57.4	-	-	-
MoV1+BS [13]	90.2	96.4	67.6	49.0	63.1	31.3
MTML [10]	92.3	95.7	64.6	48.9	59.7	23.6
OIFE [34]	68.3	89.7	48.0	-	-	-
VAMI [42]	77.0	90.8	50.1	-	-	-
RAM [20]	88.6	94.0	61.5	-	-	-
AAVER [11]	89.0	94.7	61.2	-	-	-
GRF-GGL [21]	89.4	95.0	61.7	-	-	-
DFFMG [9]	-	-	-	48.0	60.0	25.3
SPAN w/ CPDM (Ours)	94.0	97.6	68.9	59.5	61.9	42.0

benchmarks. In contrast, our method (**SPAN w/ CPDM**) achieves clear gains of **7.2%** and **16.7%** for mAP in VeRi-776 and CityFlow-ReID datasets compared to [21] and [9] respectively, indicating that we can benefit from more meaningful attention masks and better utility of global and part features. Also, our method outperforms other state-of-the-arts in both datasets.

5 Conclusion

In this paper, we present a novel vehicle re-ID feature learning framework including Semantics-guided Part Attention Network (SPAN) and Co-occurrence Part-attentive Distance Metric (CPDM). Our newly-designed SPAN can generate robust and meaningful attention masks on vehicle parts given only the image-level semantic labels for training. This is attributed to the direct supervision by our proposed mask reconstruction loss and two auxiliary losses. With the help of robust attention masks, the part feature extraction network is able to learn a more discriminative representation. Finally, our proposed CPDM can place emphasis on the vehicle parts that co-occurs in two images to better measure the distance between two vehicles. Both qualitative and quantitative results confirm the quality of generated attention masks and the benefit of dedicated part feature extraction and distance metric. Experiments also show that our proposed framework performs favorably against existing vehicle re-ID methods.

Acknowledgment

This research was supported in part by the Ministry of Science and Technology of Taiwan (MOST 108-2633-E-002-001), National Taiwan University (NTU-108L104039), Intel Corporation, Delta Electronics and Compal Electronics.

References

1. Bai, Y., Lou, Y., Gao, F., Wang, S., Wu, Y., Duan, L.Y.: Group-sensitive triplet embedding for vehicle reidentification. *IEEE Transactions on Multimedia* **20**(9), 2385–2399 (2018)
2. Bak, S., Carr, P.: One-shot metric learning for person re-identification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2990–2999 (2017)
3. Chen, D., Yuan, Z., Chen, B., Zheng, N.: Similarity learning with spatial constraints for person re-identification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1268–1277 (2016)
4. Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: a deep quadruplet network for person re-identification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 403–412 (2017)
5. Ge, Y., Li, Z., Zhao, H., Yin, G., Yi, S., Wang, X., et al.: Fd-gan: Pose-guided feature distilling gan for robust person re-identification. In: *Advances in neural information processing systems*. pp. 1222–1233 (2018)
6. He, B., Li, J., Zhao, Y., Tian, Y.: Part-regularized near-duplicate vehicle re-identification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3997–4005 (2019)
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
8. Hermans, A., Beyer, L., Leibe, B.: In defense of the triplet loss for person re-identification. *arXiv* **1703.07737** (2017)
9. Huang, P., Huang, R., Huang, J., Yangchen, R., He, Z., Li, X., Chen, J.: Deep feature fusion with multiple granularity for vehicle re-identification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*. pp. 80–88 (2019)
10. Kanaci, A., Li, M., Gong, S., Rajamanoharan, G.: Multi-task mutual learning for vehicle re-identification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*. pp. 62–70 (2019)
11. Khorramshahi, P., Kumar, A., Peri, N., Rambhatla, S.S., Chen, J.C., Chellappa, R.: A dual path model with adaptive attention for vehicle re-identification. *arXiv* **1905.03397** (2019)
12. Koestinger, M., Hirzer, M., Wohlhart, P., Roth, P.M., Bischof, H.: Large scale metric learning from equivalence constraints. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2288–2295 (2012)
13. Kuma, R., Weill, E., Aghdasi, F., Sriram, P.: Vehicle re-identification: an efficient baseline using triplet embedding. In: *2019 International Joint Conference on Neural Networks (IJCNN)*. pp. 1–9. *IEEE* (2019)
14. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
15. Li, D., Chen, X., Zhang, Z., Huang, K.: Learning deep context-aware features over body and latent parts for person re-identification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 384–393 (2017)
16. Li, S., Bak, S., Carr, P., Wang, X.: Diversity regularized spatiotemporal attention for video-based person re-identification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 369–378 (2018)

17. Liu, C.T., Lee, M.Y., Wu, C.W., Chen, B.Y., Chen, T.S., Hsu, Y.T., Chien, S.Y., Center, N.I.: Supervised joint domain learning for vehicle re-identification. In: Proc. CVPR Workshops. pp. 45–52 (2019)
18. Liu, C.T., Wu, C.W., Wang, Y.C.F., Chien, S.Y.: Spatially and temporally efficient non-local attention network for video-based person re-identification (2019)
19. Liu, H., Tian, Y., Yang, Y., Pang, L., Huang, T.: Deep relative distance learning: Tell the difference between similar vehicles. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2167–2175 (2016)
20. Liu, X., Zhang, S., Huang, Q., Gao, W.: Ram: a region-aware deep model for vehicle re-identification. In: IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6 (2018)
21. Liu, X., Zhang, S., Wang, X., Hong, R., Tian, Q.: Group-group loss-based global-regional feature learning for vehicle re-identification. *IEEE Transactions on Image Processing* **29**, 2638–2652 (2019)
22. Liu, X., Liu, W., Ma, H., Fu, H.: Large-scale vehicle re-identification in urban surveillance videos. In: IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6 (2016)
23. Liu, X., Liu, W., Mei, T., Ma, H.: A deep learning-based approach to progressive vehicle re-identification for urban surveillance. In: European Conference on Computer Vision (ECCV). pp. 869–884. Springer (2016)
24. Lou, Y., Bai, Y., Liu, J., Wang, S., Duan, L.Y.: Embedding adversarial learning for vehicle re-identification. *IEEE Transactions on Image Processing* **28**(8), 3794–3807 (2019)
25. Luo, H., Gu, Y., Liao, X., Lai, S., Jiang, W.: Bag of tricks and a strong baseline for deep person re-identification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop (2019)
26. Miao, Y., Gowayyed, M., Metze, F.: Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding. In: IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU). pp. 167–174 (2015)
27. Ristani, E., Tomasi, C.: Features for multi-target multi-camera tracking and re-identification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6036–6046 (2018)
28. Rother, C., Kolmogorov, V., Blake, A.: Grabcut: Interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics (TOG)* **23**(3), 309–314 (2004)
29. Shen, Y., Xiao, T., Li, H., Yi, S., Wang, X.: Learning deep neural networks for vehicle re-id with visual-spatio-temporal path proposals. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1900–1909 (2017)
30. Shi, Z., Hospedales, T.M., Xiang, T.: Transferring a semantic representation for person re-identification and search. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4184–4193 (2015)
31. Sun, Y., Xu, Q., Li, Y., Zhang, C., Li, Y., Wang, S., Sun, J.: Perceive where to focus: Learning visibility-aware part-level features for partial person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 393–402 (2019)
32. Tang, Z., Naphade, M., Birchfield, S., Tremblay, J., Hodge, W., Kumar, R., Wang, S., Yang, X.: Pamtri: Pose-aware multi-task learning for vehicle re-identification using highly randomized synthetic data. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 211–220 (2019)

33. Tang, Z., Naphade, M., Liu, M.Y., Yang, X., Birchfield, S., Wang, S., Kumar, R., Anastasiu, D., Hwang, J.N.: Cityflow: A city-scale benchmark for multi-target multi-camera vehicle tracking and re-identification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8797–8806 (2019)
34. Wang, Z., Tang, L., Liu, X., Yao, Z., Yi, S., Shao, J., Yan, J., Wang, S., Li, H., Wang, X.: Orientation invariant feature embedding and spatial temporal regularization for vehicle re-identification. In: IEEE International Conference on Computer Vision (ICCV). pp. 379–387 (2017)
35. Yu, H.X., Wu, A., Zheng, W.S.: Cross-view asymmetric metric learning for unsupervised person re-identification. In: IEEE International Conference on Computer Vision (ICCV). pp. 994–1002 (2017)
36. Yu, R., Zhou, Z., Bai, S., Bai, X.: Divide and fuse: A re-ranking approach for person re-identification. arXiv preprint arXiv:1708.04169 (2017)
37. Zhao, H., Tian, M., Sun, S., Shao, J., Yan, J., Yi, S., Wang, X., Tang, X.: Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1077–1085 (2017)
38. Zhao, L., Li, X., Zhuang, Y., Wang, J.: Deeply-learned part-aligned representations for person re-identification. In: IEEE International Conference on Computer Vision (ICCV). pp. 3219–3228 (2017)
39. Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., Tian, Q.: Scalable person re-identification: A benchmark. In: IEEE International Conference on Computer Vision (ICCV). pp. 1116–1124 (2015)
40. Zheng, Z., Zheng, L., Yang, Y.: A discriminatively learned cnn embedding for person reidentification. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **14**(1), 13 (2018)
41. Zhou, J., Yu, P., Tang, W., Wu, Y.: Efficient online local metric adaptation via negative samples for person re-identification. In: IEEE International Conference on Computer Vision (ICCV). pp. 2420–2428 (2017)
42. Zhou, Y., Shao, L.: Viewpoint aware attentive multi-view inference for vehicle re-identification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6489–6498 (2018)
43. Zhu, J., Zeng, H., Huang, J., Liao, S., Lei, Z., Cai, C., Zheng, L.: Vehicle re-identification using quadruple directional deep learning features. *IEEE Transactions on Intelligent Transportation Systems* (2019)