

LR-CNN: LOCAL-AWARE REGION CNN FOR VEHICLE DETECTION IN AERIAL IMAGERY

Liao, Wentong^{1,*†}, Xiang Chen^{2,*}, Jingfeng Yang³, Stefan Roth², Michael Goesele², Michael Ying Yang⁴, Bodo Rosenhahn¹

¹ Leibniz Universität Hannover, Germany - (liao, rosenhahn)tnt.uni-hannover.de

² Technische Universität Darmstadt, Germany - (xiang.chen, stefan.roth)@visinf.tu-darmstadt.de, research@goesele.org

³ Chinese Academy of Sciences, China - ioaniu@163.com

⁴ Faculty ITC, University of Twente, The Netherlands - michael.yang@utwente.nl

Commission II, WG II/4

KEY WORDS: Deep Learning, Object Detection, Vehicle Detection, Twin Region Proposal, Feature Enhancement

ABSTRACT:

State-of-the-art object detection approaches such as Fast/Faster R-CNN, SSD, or YOLO have difficulties detecting dense, small targets with arbitrary orientation in large aerial images. The main reason is that using interpolation to align RoI features can result in a lack of accuracy or even loss of location information. We present the Local-aware Region Convolutional Neural Network (LR-CNN), a novel two-stage approach for vehicle detection in aerial imagery. We enhance translation invariance to detect dense vehicles and address the boundary quantization issue amongst dense vehicles by aggregating the high-precision RoIs' features. Moreover, we resample high-level semantic pooled features, making them regain location information from the features of a shallower convolutional block. This strengthens the local feature invariance for the resampled features and enables detecting vehicles in an arbitrary orientation. The local feature invariance enhances the learning ability of the focal loss function, and the focal loss further helps to focus on the hard examples. Taken together, our method better addresses the challenges of aerial imagery. We evaluate our approach on several challenging datasets (VEDAI, DOTA), demonstrating a significant improvement over state-of-the-art methods. We demonstrate the good generalization ability of our approach on the DLR 3K dataset.

1. INTRODUCTION

Vehicle detection in aerial photography is challenging but widely used in different scenarios, *e.g.*, traffic surveillance, urban planning, satellite reconnaissance, or UAV detection. Since the introduction of Region-CNN (Girshick et al., 2014), which uses region proposals and learns possible region features using a convolutional neural network instead of traditional manual features, many excellent object detection frameworks based on this structure were proposed, *e.g.*, Light-head R-CNN (Li et al., 2017), Fast/Faster R-CNN (Girshick, 2015, Ren et al., 2015), YOLO (Redmon, Farhadi, 2017, Redmon, Farhadi, 2018), and SSD (Liu et al., 2016). These frameworks do, however, not work well for aerial imagery due to the challenges specific to this setting.

In particular, the camera's bird's eye view and the high-resolution images make target recognition hard for the following reasons: (1) Features describing small vehicles with arbitrary orientation are difficult to extract in high-resolution images. (2) The large number of visually similar targets from different categories (*e.g.*, building roofs, containers, water tanks) interfere with the detection. (3) There are many, densely packed target vehicles with typically monotonous appearance. (4) Occlusions and shadows increase the difficulty of feature extraction. Fig. 1 illustrates some challenging examples in aerial imagery.

(Xia et al., 2018) evaluate recent frameworks on the DOTA dataset. Their results indicate that two-stage object detection frameworks (Dai et al., 2016, Ren et al., 2015) do not work well for finding objects in dense scenarios, whereas one-stage

* These authors contributed equally to this work

† Corresponding author

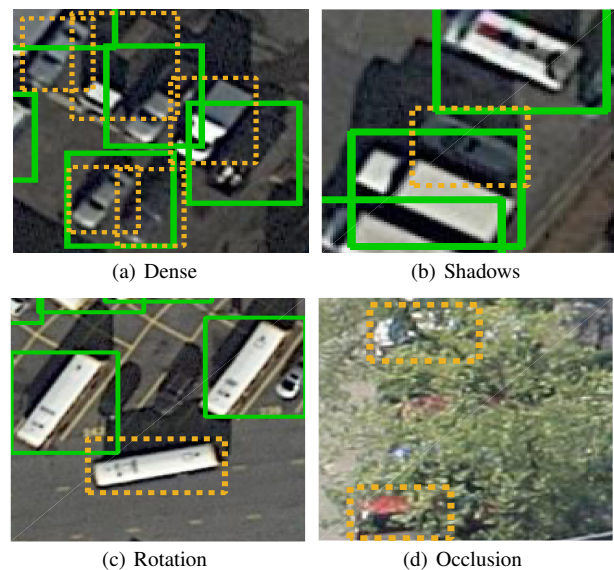


Figure 1. Dense, arbitrary orientation, shadows, and occlusion are typical challenges for vehicle detection in aerial imagery. Green boxes indicate detection results of Faster R-CNN. Orange dashed boxes mark undetected vehicles.

object detection frameworks (Liu et al., 2016, Redmon, Farhadi, 2017) cannot detect dense and small targets. Moreover, all frameworks have problems detecting vehicles with arbitrary orientation. We argue that one of the important reasons is that RoI pooling uses interpolation to align region proposals of all sizes, which leads to a reduced accuracy or even loss of spatial information of the feature.




Figure 2. Architecture: The backbone is a ResNet-101. Blue components represent subnetworks, gray color denotes feature maps, and yellow color indicates fully connected layers. The Region Proposal Network (RPN) proposes candidate RoIs, which are then applied to the feature maps from the third and the fourth convolutional blocks, respectively. Afterwards, RoIs from the third convolutional block are fed into the Localization Network to find the transformation parameters of local invariant features, and the Grid Generator matches the correspondence of pixel coordinates between RoIs from the third and the fourth convolutional blocks. Next, the Sampler determines which pixels are sampled. Finally, the regression and classifier output the vehicle detection results.

To address these problems, we propose the Local-aware Region Convolutional Neural Network (LR-CNN) for vehicle detection in aerial imagery. The goal of LR-CNN is to make the deeper high-level semantic representation regain high-precision location information. We, therefore, predict affine transformation parameters from the shallower layer feature maps, containing a wealth of location information. After spatial transformation processing the pixels of the shallower layer feature maps are projected based on these transformation parameters onto the corresponding pixels of deeper feature maps containing higher-level semantic information. Finally, the resampled features, guided by the loss function, possess local invariance and contain location and high-level semantic information. To summarize, our contributions are the following:

- A novel network framework for vehicle detection in aerial imagery.
- Preserving the aggregate RoIs' feature translation invariance and addressing the boundary quantization issue for dense vehicles.
- Proposing a resampled pooled feature, which allows higher-level semantic features to regain location information and have local feature invariance. This allows detecting vehicles at an arbitrary orientation.
- An analysis of our results showing that we can detect vehicles in aerial imagery accurately and with tighter bounding boxes even in front of complex backgrounds.

2. RELATED WORK

Object detection. Recent object detection techniques can be roughly summarized in two ways. Two-step strategies first generate many candidate regions, which likely contain objects of interest. Then a separate sub-network determines the categories of each of these candidates and regresses the location. The most representative work is Faster R-CNN (Ren et al., 2015), which introduced the Region Proposal Network (RPN) for candidate

generation. It is derived from R-CNN (Girshick et al., 2014), which uses Selective Search (Uijlings et al., 2013) to generate candidate regions. SPPnet (He et al., 2014) proposed a Spatial Pyramid Pooling layer to obtain multi-scale features at a fixed feature size. Lastly, Fast R-CNN (Girshick, 2015) introduced the ROIpooling layer and enabled the network to be trained in an end-to-end fashion. Because of its high precision and good performance on small objects and dense objects, Faster R-CNN is currently the most popular pipeline for object detection. In contrast, one-step approaches predict the location of objects and their category labels simultaneously. Representative works are YOLO (Redmon et al., 2016, Redmon, Farhadi, 2017, Redmon, Farhadi, 2018) and SSD (Liu et al., 2016). Because there is no separate region proposal step this strategy is fast but achieves lower detection accuracy.

Vehicle detection. Vehicle detection is a special case of object detection, *i.e.* the aforementioned methods can be directly applied (Shi et al., 2017, Wu et al., 2018). These methods are, however, carefully designed to work on images collected from the ground, in which the objects have rich appearance characteristics. In contrast, visual information is very limited and monotonous when seen from an aerial perspective. Moreover, aerial images have much higher resolution (*e.g.*, 5616×3744 in ITCVD (Yang et al., 2019) compared to 375×500 in ImageNet (Deng et al., 2009)) and cover a wider area. The objects of interest (vehicles in this work) are much smaller, and their scale, size, and orientation vary strongly. An important prior for object detection on ground-view images is that the main or large objects within an image are mostly at the image center (Redmon, Farhadi, 2017). In contrast, an object's location is unpredictable in an aerial image. Selective search, RPN, or YOLO are therefore likely not ideal to handle these challenges. Given inaccurate region proposals, the following classifier cannot work well to make a final decision. More challenges include that vehicles can be in dark shadow, occluded by buildings, or packed densely on parking lots. All these challenges make the existing sophisticated object detection algorithms not well suited for aerial images.

Vehicle detection in aerial images has been investigated by many

recent studies, *e.g.* (Azimi et al., 2018, Hinz, 2004, Liu et al., 2017, Qu et al., 2017, Razakarivony, Jurie, 2015, Tang et al., 2017, Yang et al., 2018). (Tang et al., 2017, Yang et al., 2018) extract features from shallower convolution layers (*conv3* and *conv4*) through skip connections and fuse with the final features (output of *conv5*). Then a standard RPN is used on multi-scale feature maps to obtain proposals at different scales. (Tang et al., 2017) train a set of boosted classifiers to improve the final prediction accuracy. (Yang et al., 2018) use the focal loss (Lin et al., 2020) instead of the cross entropy as loss function for the RPN and the classification layer during training to overcome the easy/hard examples challenge. They report a significant improvement in this task. (Azimi et al., 2018) propose to extract features hierarchically at different scales so that the network is able to detect objects in different sizes. To address the arbitrary orientation problem, they rotate the anchors of the proposals to some predefined angles (Ma et al., 2018), similar to (Liu et al., 2017). The number of anchors increases, however, dramatically to $N_{\text{scales}} \times N_{\text{ratios}} \times N_{\text{angles}}$ and computation is costly.

3. OUR APPROACH

Motivated by DFL-CNN (Yang et al., 2018), our approach uses a two-stage object detection strategy, as shown in Fig. 2. In this section, we will give details for each of the sub-networks and discuss how our approach improves the accuracy for detecting vehicles in aerial images.

3.1 Base feature extractor

Excessive downsampling can lead to a loss of feature information for small target vehicles. In contrast, low-level features from shallower layers can retain not only rich feature details of small targets, but also rich spatial information. We adopt ResNet-101 (He et al., 2016) and extract the base features from the shallow layers. As shown in Fig. 2, we use feature maps from the third and fourth convolutional block, which have the same resolution. Since there is a 69 convolutional layer gap between the output of the third and fourth convolutional blocks, the latter contains deeper features, whereas the third convolutional block is relatively shallow and its output retains better spatial information of the pooled objects' features.

3.2 Region proposal network

Twin region proposals. We model the region proposal network (RPN) as in (Ren et al., 2015). For each input image, the RPN outputs 128 potential RoIs, which are mapped to the feature maps from the third $F_{\text{RoI}_{\text{conv3},x}}$ and fourth $F_{\text{RoI}_{\text{conv4},x}}$ convolutional block. (He et al., 2017) argue that the RoI pooling's nearest neighbor interpolation leads to a loss in translation invariance of the aligned RoI features. Low RoI alignment accuracy is, however, counterproductive for region proposal features that represent small target vehicles. We, therefore, use RoIAlign (He et al., 2017) instead of RoI pooling to aggregate high-precision RoIs.

RoI feature processing. As Fig. 3 illustrates, the $N \times 512 \times 128 \times 128$ input from the third convolutional block will be sent into a large separable convolution (LSC) module containing two separate branches. Afterwards, the $N \times 512 \times 128 \times 128$ feature is compressed to $N \times 147 \times 128 \times 128$ position-sensitive score maps, which have 49 3-channel feature map blocks. This will greatly reduce the computational expense of generating

position-sensitive score maps since the feature is now much thinner than it used to be (Li et al., 2017).

In the LSC module, each branch uses a large kernel size to enlarge the receptive field to preserve large local features. Large local features, while not accurate enough, retain more spatial information than local features extracted with small convolution kernels. This means that the larger local features facilitate further affine transformation parameterization, which effectively preserves the spatial information.

Position-sensitive RoIAlign. As discussed above, RoI pooling increases noise in the feature representation when RoIs are aggregated. Additionally, (Dai et al., 2016) demonstrates that the translation invariance of the feature is lost after the RoI pooling operation. Inspired by both and following the structure of (Dai et al., 2016) we build the position-sensitive RoIAlign by replacing RoI pooling with RoIAlign. As the structure of position-sensitive RoIAlign indicates in Fig. 3, after aggregating by RoIAlign the precision of the RoIs' alignment strongly improves the sensitive position scoring and significantly reduces the noise of the small target feature.

RPN loss. Since the distribution of large and small vehicle samples in aerial images is sparse, the ratio of positive and negative examples for training is very unbalanced. Hence, we use the focal loss (Lin et al., 2020), which reduces the weight for easy to classify examples, in order to improve the learnability of dense vehicle detection. The loss function of the RPN is defined as

$$L_{\text{RPN}}(\{p_i\}, \{t_i\}) = \frac{\alpha}{N_{\text{cls}}} \sum_i (p_{t,i} - 1)^\gamma \log(p_{t,i}) + \frac{\lambda}{N_{\text{reg}}} \sum_i p_i^* f_{\text{smooth L1},i} \quad (1)$$

with

$$p_{t,i} = \begin{cases} p_i, & p_i^* = 1 \\ 1 - p_i, & \text{otherwise} \end{cases} \quad (2)$$

$$f_{\text{smooth L1},i} = \begin{cases} 0.5(t_i - t_i^*)^2, & |t_i - t_i^*| < 1 \\ |t_i - t_i^*| - 0.5, & \text{otherwise.} \end{cases} \quad (3)$$

Here, i denotes the index of the proposal, p_i is the predicted probability of the corresponding proposal, p_i^* represents the ground truth label (positive = 1, negative = 0). t_i describes the predicted bounding box vector and t_i^* indicates the ground truth box vector if $p_i^* = 1$. We set the balance parameters $\alpha = 1$ and $\lambda = 1$. The focusing parameter of the modulating factor $(p_{t,i} - 1)^\gamma$ is $\gamma = 2$ as in (Lin et al., 2020).

3.3 Resampled pooled feature

(Dai et al., 2016, He et al., 2017, Jiang et al., 2018) argue that RoI pooling uses interpolation to align the region proposal, which causes the pooled feature to lose location information. Due to this, they propose higher precision interpolations to improve the precision of RoI pooling. We instead assume that the region proposal undergoes an affine transformation after interpolation alignment, such as stretching, rotation, shifting, *etc.* We thus exploit spatial transformer networks (STNs) (Jaderberg et al., 2015) to let the deep high-level semantic representation regain location information from the shallower features that retain the spatial information. Thereby, we strengthen the local feature invariance of the target vehicle in the RoI.




Figure 3. The specific architecture of the Large Separable Convolution and Position-Sensitive RoIAlign blocks in Fig. 2. This subnet consists of three modules: Large separable convolutions (LSC), position-sensitive score maps, and RoIAlign. Each color of the output stands for the pooled results from each corresponding 3-channel position-sensitive score maps. Combined with region proposals from RPN, the position-sensitive RoIAlign creates a $128 \times 3 \times 7 \times 7$ output for the localization network.

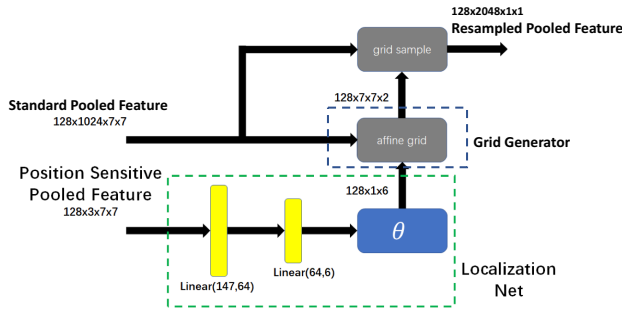


Figure 4. The specific architecture of the resampled pooled feature subnetwork in Fig. 2, which consists of Localization Network, Grid Generator, and Sampler.

The STN trains a model to predict the spatial variation and alignment of features (including translation, scaling, rotation, and other geometric transformations) by adaptively predicting the parameters of an affine transformation. Fig. 4 depicts the architecture of a resampled pooled feature subnetwork. Six parameters are sufficient to describe the affine transformation (Jaderberg et al., 2015). We feed the position-sensitive pooled feature F_{ps} from $F_RoI_{conv3-x}$ into the localization network and then parameterize the location information in the RoI as θ , which are regressed 2×3 parameters for describing the affine transformation. Next, standard pooled features F_{st} from $F_RoI_{conv4-x}$ are converted to a parameterised sampling grid to model the correspondence coordinate matrix M_t with transformation $T(\theta)$. It is placed at the pixel level between the resampled pooled feature F_{rp} and F_{st} by the grid generator. Once M_t has been modeled, F_{rp} will be pixel-wise resampled from F_{st} , and thus the spatial information is re-added to F_{rp} .

The feature map visualization in Fig. 9 shows that our resampled pooled features have enhanced the local feature invariance, and the feature representation of the vehicle placed at any direction is also very strong.

3.4 Loss of classifier and regressor

For the final classifier and regression, we continue using the focal loss and the smooth L_1 loss function, respectively:

$$L_{LR-CNN}(\{p_j\}, \{t_j\}) = \frac{\alpha}{N_{cls}} \sum_j (p_{t,j} - 1)^\gamma \log(p_{t,j}) + \frac{\lambda}{N_{regr}} \sum_j p_j^* f_{smooth L1,j}, \quad (4)$$

where j represents the index of the proposal. All other definitions are as in Eq. (1). The parameters remain as $\alpha = 1$, $\lambda = 1$ and $\gamma = 2$. The total loss function can then be represented as

$$L = L_{RPN}(\{p_i\}, \{t_i\}) + L_{LR-CNN}(\{p_j\}, \{t_j\}). \quad (5)$$

4. EXPERIMENTS

4.1 Datasets

We evaluate the proposed method on three datasets with different characteristics, testing different aspects of the accuracy of our method.

The VEDAI (Razakarivony, Jurie, 2015) dataset consists of satellite imagery taken over Utah in 2012. It contains 1210 RGB images with a resolution of 1024×1024 pixels. VEDAI contains sparse vehicles and is challenging due to strong occlusions and shadows.

DOTA (Xia et al., 2018) has 2806 aerial images, which are collected with different sensors and platforms. Their resolutions range from 800×800 to about $4k \times 4k$ pixels. The dataset is randomly split into three sets: Half of the original images form the training set, 1/6 are used as validation set, and the remaining 1/3 form the testing set. Annotations are publicly accessible for all images not in the testing set. The experimental results on DOTA reported in this paper are therefore from the validation set. Furthermore, we evaluate the accuracy of detecting large and small vehicles separately for comparison purposes.

The DLR 3K dataset (Liu, Mattyus, 2015) consists of 20 images (10 images for training and the other 10 for testing), which are captured at the height of about 1000 feet over Munich with a resolution of 5616×3744 pixels. This dataset is used to evaluate the generalization ability of our method.

DOTA and VEDAI provide annotations of different kinds of object categories. Given the goal of this paper, we only use the vehicle annotations. Our method can, however, likely be generalized to detect arbitrary categories of interest.

Because of the very high resolution of the images and limited GPU memory, we process images larger than 1024×1024 pixels in tiles. *I.e.*, we crop them into 1024×1024 pixel patches with an overlap of 100 pixels. This truncates some targets. We only keep targets with more than 50% remaining as positive samples.

In order to assess the accuracy of our framework, we adopt the standard VOC 2010 object detection evaluation metric (Everingham et al., 2015) for quantitative results of *precision*, *recall*, and *average precision*.

	VEDAI	DOTA	DOTA		
	AP	AP	SV AP	LV AP	mAP
F R-CNN	87.24	42.92	33.79	45.50	39.65
DFL	90.54	62.62	45.56	61.63	53.60
Ours	92.54	70.33	56.09	77.86	66.97

Table 1. Experimental results showing average precision (AP) and mean AP (mAP) when detecting small (SV) and large vehicles (LV) separately in percent.




Figure 5. Precision-Recall curves given by different methods on the DOTA dataset. The color denotes the method while the line type denotes different tasks.

4.1.1 Implementation details We use ResNet-101 as backbone network to learn features and initialize its parameters with a model pretrained on ImageNet (Deng et al., 2009). The remaining layers are initialized randomly. During training, stochastic gradient descent (SGD) is used to optimize the parameters. The base learning rate is 0.05 with a 10% decay every 3 epochs. The IoU thresholds for NMS are 0.7 for training and 0.5 for inference. The RPN part is trained first before the whole framework is trained jointly. All experiments were conducted with NVIDIA Titan XP GPUs. A single image with size 1024×1024 keeps a maximum of 600 RoIs after NMS, and takes ca. 1.4s during training and ca. 0.33s for testing.

4.2 Results and comparison

We compare our method with the state-of-the-art detection methods DFL (Yang et al., 2018) and the standard Faster R-CNN (Ren et al., 2015) as baseline. We evaluated these methods with their own settings on all datasets.

4.2.1 Quantitative results Tab. 1 summarizes the experimental results. Note that our method outperforms all methods on all datasets. Furthermore, small vehicle and large vehicle on the DOTA Evaluation Server get **68.56%** and **69.87%** of AP respectively, and the mAP is **69.22%**. Particularly, compared to the baseline method and the state-of-the-art, our model increases the AP by 27.41% and 7.71% on the most challenging dataset DOTA, respectively, corresponding to 63.9% and 12.3% relative gains. When small and large vehicles are considered as two classes, our model achieves 55.1% and 71.1% relative gains, respectively, against the baseline. The significant gains prove that our Large Separable Convolution, Position-Sensitive RoIAlign and Spatial Transform Network modules work efficiently.

Fig. 5 depicts the *precision-recall* curves of different methods on DOTA. We can see that for vehicle detection our method (blue solid line) has a wider smooth region (until a recall of

	Training data		HRPN
	VEDAI	DOTA	
Faster R-CNN	60.25%	68.51%	79.54%
DFL	61.69%	83.04%	
Ours	69.19%	89.21%	

Table 2. Experimental results (AP) on the training set of DLR 3K with the models trained on different other datasets. These experiments evaluate the generalization ability. For comparison, we cite the results of HRPN trained and evaluated on DLR 3K.

Features from	DOTA Dataset		
	SV AP	LV AP	mAP
conv3_x	56.09%	77.86%	66.97%
conv4_x	55.81%	75.39%	65.60%

Table 3. Ablation study. STN is fed with features from different convolution blocks of the backbone network for small (SV) and large (LV) vehicles.

0.65) and smoother tendency, which means our method is more robust and has higher object classification precision than others. In contrast, both Faster R-CNN and DFL (red and green solid lines, respectively) have a rapid drop at the high-precision end of the plot. In other words, our method achieves higher recall without the cost of obviously sacrificing precision. We also can see that small vehicle detection is more difficult for all methods: The curves (pointed lines) begin to obviously drop much earlier (for LR-CNN at a recall of 0.4) than the general or large-vehicle detection (at a recall of 0.65), and the transition region is also wide (until a recall of 0.67 for LR-CNN). It is worth mentioning that DFL and LR-CNN have very good curves for large vehicle detection (dashed lines) with long smooth regions and a rapid drop.

4.2.2 Qualitative results Fig. 6 gives a qualitative comparison between different methods on DOTA. It shows a typical complex scene: vehicles are in arbitrary places, dense or sparse, and the background is complex. As shown in the first row, Faster R-CNN fails to detect many vehicles, especially when they are dense (Regions 2, 3) or in shadow (Regions 5, 6). DFL detects more small vehicles. In particular, it is sensitive to the dark small vehicles, e.g., an unclear car on the road (Region 1) is detected. However, this has side effects: DFL cannot distinguish small dark vehicles from shadow well. E.g., the shadow of the white vehicle in Region 4 is detected as a small vehicle but the vehicles in Regions 5 and 6 are not detected. Furthermore, its accuracy for detecting vehicles in dense cases and classifying the vehicles' type is not good enough (Regions 2, 3). Fig. 6(c) shows that our method distinguishes large and small vehicles well. It can also detect individual vehicles in dense parts of the scene. The advantages of detecting vehicles in dense situations and distinguishing the vehicles from the similar background objects are further showcased in the second row.

4.2.3 Generalization ability To evaluate the generalization ability of our approach, we test it on the DLR 3K dataset with models trained on different datasets. Because the ground truth of the test set of DLR 3K is not publicly accessible, we test the models on the training and validation set whose annotations are available. We also compare the results with the ones reported in HRPN (Tang et al., 2017), which was trained on DLR 3K. Experimental results are listed in Tab. 2. We can see that, for each method, the model trained on DOTA reports higher AP than that trained on VEDAI. The main reason is that DOTA has more and more diverse training samples. DFL and our method trained on DOTA outperform HRPN with our method reporting




Figure 6. Qualitative comparison. Green boxes indicate detected large vehicles; blue boxes show detected small vehicles. The number of detected vehicles is shown at the bottom right. We use dashed red and yellow boxes to highlight challenging image parts, which can be handled correctly by our method.

about 10% better results than HRPN. These results show that our model has good generalization abilities as well as transferability. For better understanding, we show some examples in Fig. 7. When comparing the dashed purple boxes (results of models trained on VEDAI) with the green boxes (results of models trained on DOTA) from the same method, we can see that the models trained on DOTA detect more vehicles. When comparing the results of different methods trained on DOTA, we can see that LR-CNN successfully detects more vehicles. Within the region highlighted by the dashed yellow box where vehicles are dense, LR-CNN successfully detects almost all individual vehicles.

4.2.4 Ablation study To evaluate the impact of the STN placed at different locations in the network, we conduct an ablation study. We do not provide separate experiments to evaluate the impact of focal loss and RoIAlign pooling because these have been provided in (Lin et al., 2020, Tang et al., 2017) and (He et al., 2017), respectively. Tab. 3 reports our results. When the STN is placed at the output of the *conv3_x* block, the model achieves better results, especially for large vehicle detection. The reason is that the STN mainly processes spatial information, which is much richer in the output features of *conv3_x* than in those of *conv4_x*.

For better understanding, we visualize some feature maps in Fig. 9. The features extracted from *conv3_x* (second row) contain more spatial and detailed information than those from *conv4_x* (fourth row): The edges are clearer and the locations corresponding to the vehicle show stronger activations. Comparing the feature maps before and after the STN (2nd row vs. 3rd row and 4th row vs. 5th row) shows that the activations of the background regions are weaker after the STN. Active regions corresponding to the foreground are closer to the vehicle's shape and orientation than before applying the STN since the features are transformed and regularized by the STN module. Furthermore, after STN processing, in addition to being accurate in position, the feature representation is also slimmer. This is why

our bounding boxes are tighter than other detectors. From these observation, we can intuitively conclude that the STN module is better able to find the transformation parameters on *conv3_x* to regularize the features used to regress the location and classify the RoIs.

Fig. 8 illustrates how the quality of proposals from RPN affects the final localization and classification. When comparing the final detection results (green boxes) with the RPN proposals (dashed purple boxes) of different methods, we can make the following observations: LR-CNN correctly detects more vehicles. In addition, the green bounding boxes given by LR-CNN are tighter, which means that LR-CNN gives more precise localization. To analyze the reasons for this, we compare the proposals (dashed purple boxes) of different methods. We can see that the proposals given by DFL and our method are closer to the targets than the ones of Faster R-CNN. Even though each vehicle is detected by its own RPN, the final classifier removes these proposals (Proposals 2 and 4) since they deviate from the ground truth location too much and contain too much background. Thus, the features pooled from these RoIs are not precise enough to represent the targets. Consequently, the final classifier cannot determine well based on these features whether they are an object of interest, especially in dense cases. To analyze why LR-CNN localizes the objects better, we look at the mathematical definition of target regression. The regression target for width is

$$t_w = \log \frac{G_w}{P_w} = \log \left(1 + \frac{G_w - P_w}{P_w} \right). \quad (6)$$

G_w denotes the ground truth width and P_w is the prediction. The target height t_h is handled equivalently. Only when the prediction is close to the target, the equation can approximate a linear relationship: $\lim_{x \rightarrow 0} \log(1 + x) = 1 + x$ (because the regression targets of center shift (x, y) are already defined as a linear function and all these four parameters are predicted simultaneously. The regression layer is easier to be trained and



Figure 7. Qualitative comparison of the generalization ability. The green boxes denote detection results given by the model that was trained on DOTA. The dashed purple boxes denote detection results of a model trained on VEDAI. We use a dashed yellow box to highlight a challenging image region that can be handled correctly by our method.




Figure 8. Example images showing the differences between the bounding box predicted by RPN (dashed purple box) and the finally predicted location (solid green box) regressed by the classification layer.




Figure 9. Feature map (one example per column). Colors show activation strength.

works better when all the four target equation are linear). For all these reasons, our framework obtains better proposals in our RPN and yields better final classification and localization.

4.2.5 Discussion Compared to Faster R-CNN and DFL, our approach performs much better on detecting small targets. This improvement benefits from the skip connection structure that fuses the richer detail information from the shallower layers with the features from deeper layers, which contain higher-level semantic information. This is important for detecting small objects in high-resolution aerial images. In our method, the

position-sensitive RoIAlign pooling is adopted to extract more accurate information compared with the traditional RoI pooling. An accurate representation is important for precisely locating and classifying small objects. Then our final classifier works better to determine the targets and further refine their location. Most importantly, the STN module in our framework regularizes the learned features after RoIAlign pooling well, which reduces the burden of the following layers that are expected to learn powerful enough feature representations for classification and further regression. That is the reason why LR-RCNN distinguishes small and large vehicles better and has more precise detection. All the above elements enable our method to have a good generalization ability and to reach a new state-of-the-art in vehicle detection in high resolution aerial images.

5. CONCLUSION

We present an accurate local-aware region-based framework for vehicle detection in aerial imagery. Our method improves not only the boundary quantization issue for dense vehicles by aggregating the RoIs' features with higher precision, but also the detection accuracy of vehicles placed at arbitrary orientations by the high-level semantic pooled feature regaining location information via learning. In addition, we develop a training strategy to allow the pooled feature of location information lacking the precision to reacquire the accurate spatial information from shallower layer features via learning. Our approach achieves state-of-the-art accuracy for detecting vehicles in aerial imagery and has good generalization ability. Given these properties, we believe that it should also be easy to generalize by detecting additional object classes under similar circumstances.

ACKNOWLEDGMENT

This work was supported by German Research Foundation (DFG) grants COVMAP (RO 2497/12-2) and PhoenixD (EXC 2122, Project ID 390833453).

REFERENCES

- Azimi, S. M., Vig, E., Bahmanyar, R., Körner, M., Reinartz, P., 2018. Towards multi-class object detection in unconstrained remote sensing imagery. *arXiv preprint arXiv:1807.02700*. <https://arxiv.org/abs/1807.02700>.
- Dai, J., Li, Y., He, K., Sun, J., 2016. R-fcn: Object detection via region-based fully convolutional networks. *Neural Information Processing Systems*, 379–387.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. *Computer Vision and Pattern Recognition*.
- Everingham, M., Eslami, S. A., Van Gool, L., Williams, C. K., Winn, J., Zisserman, A., 2015. The PASCAL visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1), 98–136.
- Girshick, R., 2015. Fast R-CNN. *International Conference on Computer Vision*, 1440–1448.
- Girshick, R., Donahue, J., Darrell, T., Malik, J., 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. *Computer Vision and Pattern Recognition*, 580–587.
- He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask R-CNN. *International Conference on Computer Vision*, 2980–2988.
- He, K., Zhang, X., Ren, S., Sun, J., 2014. Spatial pyramid pooling in deep convolutional networks for visual recognition. *European Conference on Computer Vision*, Springer, 346–361.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. *Computer Vision and Pattern Recognition*, 770–778.
- Hinz, S., 2004. Detection of vehicles and vehicle queues in high resolution aerial images. *Photogrammetrie-Fernerkundung-Geoinformation*.
- Jaderberg, M., Simonyan, K., Zisserman, A. et al., 2015. Spatial transformer networks. *Neural Information Processing Systems*, 2017–2025.
- Jiang, B., Luo, R., Mao, J., Xiao, T., Jiang, Y., 2018. Acquisition of Localization Confidence for Accurate Object Detection. *arXiv preprint arxiv:1807.11590*. <https://arxiv.org/abs/1807.11590>.
- Li, Z., Peng, C., Yu, G., Zhang, X., Deng, Y., Sun, J., 2017. Light-head R-CNN: In defense of two-stage object detector. *arXiv preprint arxiv:1711.07264*. <https://arxiv.org/abs/1711.07264>.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2020. Focal loss for dense object detection. *Transactions on Pattern Analysis and Machine Intelligence*, 42(1), 318–327.
- Liu, K., Mattyus, G., 2015. Fast multiclass vehicle detection on aerial images. *IEEE Geosci. Remote Sensing Lett.*, 12(9), 1938–1942.
- Liu, L., Pan, Z., Lei, B., 2017. Learning a rotation invariant Detector with rotatable bounding box. *arXiv preprint arXiv:1711.09405*. <https://arxiv.org/abs/1711.09405>.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A. C., 2016. SSD: Single shot multibox detector. *European Conference on Computer Vision*, Springer, 21–37.
- Ma, J., Shao, W., Ye, H., Wang, L., Wang, H., Zheng, Y., Xue, X., 2018. Arbitrary-oriented scene text detection via rotation proposals. *IEEE Transactions on Multimedia*, 20(11), 3111–3122.
- Qu, T., Zhang, Q., Sun, S., 2017. Vehicle detection from high-resolution aerial images using spatial pyramid pooling-based deep convolutional neural networks. *Multimedia Tools and Applications*, 76(20), 21651–21663.
- Razakarivony, S., Jurie, F., 2015. Vehicle detection in aerial imagery: A small target detection benchmark. *Journal of Visual Communication and Image Representation*, 34.
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection. *International Conference on Computer Vision*, 779–788.
- Redmon, J., Farhadi, A., 2017. YOLO9000: better, faster, stronger. *arXiv preprint arxiv:1612.08242*. <https://arxiv.org/abs/1612.08242>.
- Redmon, J., Farhadi, A., 2018. YOLOv3: An incremental improvement. *arXiv preprint arxiv:1804.02767*. <https://arxiv.org/abs/1804.02767>.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. *Neural Information Processing Systems*, 91–99.
- Shi, K., Bao, H., Ma, N., 2017. Forward vehicle detection based on incremental learning and Fast R-CNN. *CIS*, 73–76.
- Tang, T., Zhou, S., Deng, Z., Zou, H., Lei, L., 2017. Vehicle detection in aerial images based on region convolutional neural networks and hard negative example mining. *Sensors*, 17(2), 336.
- Uijlings, J. R., Van De Sande, K. E., Gevers, T., Smeulders, A. W., 2013. Selective search for object recognition. *International Journal of Computer Vision*, 104(2), 154–171.
- Wu, C.-W., Liu, C.-T., Chiang, C.-E., Tu, W.-C., Chien, S.-Y., Center, N. I., 2018. Vehicle re-identification with the space-time prior. *CVPR Workshop (CVPRW) on the AI City Challenge*.
- Xia, G.-S., Bai, X., Ding, J., Zhu, Z., Belongie, S., Luo, J., Datcu, M., Pelillo, M., Zhang, L., 2018. Dota: A large-scale dataset for object detection in aerial images. *Computer Vision and Pattern Recognition*.
- Yang, M., Liao, W., Li, X., Cao, Y., Rosenhahn, B., 2019. Vehicle Detection in Aerial Images. *Photogrammetric Engineering & Remote Sensing (PE&RS)*, 85(4), 297–304.
- Yang, M. Y., Liao, W., Li, X., Rosenhahn, B., 2018. Deep learning for vehicle detection in aerial images. *International Conference on Image Processing (ICIP)*, 3079–3083.