

On Learning Vehicle Detection in Satellite Video

Roman Pflugfelder^{1,2}, Axel Weissenfeld¹, Julian Wagner²

¹AIT Austrian Institute of Technology, Center for Digital Safety & Security

²TU Wien, Institute of Visual Computing & Human-Centered Technology

{roman.pflugfelder|axel.weissenfeld}@ait.ac.at, e1326108@student.tuwien.ac.at

Abstract. *Vehicle detection in aerial and satellite images is still challenging due to their tiny appearance in pixels compared to the overall size of remote sensing imagery. Classical methods of object detection very often fail in this scenario due to violation of implicit assumptions made such as rich texture, small to moderate ratios between image size and object size. Satellite video is a very new modality which introduces temporal consistency as inductive bias. Approaches for vehicle detection in satellite video use either background subtraction, frame differencing or subspace methods showing moderate performance (0.26 - 0.82 F_1 score). This work proposes to apply recent work on deep learning for wide-area motion imagery (WAMI) on satellite video. We show in a first approach comparable results (0.84 F_1) on Planet's SkySat-1 LasVegas video with room for further improvement.*

1. Introduction

Object detection, i.e. the recognition and localisation of objects, in visual data is a very important and still unsolved problem. For example, the problem becomes challenging in aerial imaging and remote sensing as the data and scenes differ significantly from the case considered usually in computer vision [6, 25].

Such remote detection is important in surveillance, as demanding applications let surveillance currently undergo a transition from near to mid distances (as with security cameras) to sceneries such as whole cities, traffic networks, forests, and green borders. Beside coverage new, low orbit satellite constellations¹ will allow multiple daily revisits and constantly falling costs per image. Such applications can be found e.g. in urban planning, traffic monitoring,

¹<https://earthispace.com>, 11/03/2019




Figure 1. Results of the proposed method. Top, left: video frame of the SkySat-1 LasVegas video showing a city highway with multiple cars. Top, right: vehicle labelling provided by Zhang et al. [34, 33]. Bottom, left: the method's response (heat map). Bottom, right: the final segmentation result. The network detects all labelled cars and even a bus or truck at the right image border.

driver behaviour analysis, and road verification for assisting both scene understanding and land use classification. Civilian and military security is another area to benefit with applications including military reconnaissance, detection of abnormal or dangerous behaviour, border protection, and surveillance of restricted areas.

Although remotely acquired data shows great reduction of occlusion and perspective distortion due to the overhead view, new difficulties arise. Typical aerial and satellite images are very large in resolution and data size. For example, wide-area motion imagery (WAMI) provides instead of a few megapixel (MP) typical for security cameras up to 400 MP per image frame and three image frames per second (2.2 TB/s for 16 bit per px). Satellite video gives today 4K RGB video with 30 frames per second (759.3 MB/s). Satellite images capture large sceneries, usually dozens of square kilometers which introduce instead of a few visually large objects, thou-

sands of tiny objects coming from hundreds of categories in a single image. At the same time these objects reduce in pixel size by orders of magnitude from 10^4 px to 10^2 px, to even 10 px for satellite video [34], depending on the camera’s ground sample distance (GSD)².

This severe magnification of scenery and reduction of object size to very tiny appearances have consequences. Object detection becomes very ambiguous and sensitive to noise and nuisances and the search space dramatically increases and becomes very sparse. Inferred labels of data usually capture instead of the bounding box or contour sole positions, as the extent of objects is even for humans, e.g. in WAMI or satellite video, unrecognizable. All this leads to major difficulties if not inapplicability of vanilla methods [17]. Manual labelling of data is furthermore very tedious, for many cases impossible, hence, research on object detection in satellite video relies currently on background subtraction and frame differencing [16, 28, 32, 18, 4, 3].

Recent literature [22, 23, 35, 15, 31, 8, 27, 17, 34, 13, 2, 29, 30] also suggests to apply deep learning on aerial and satellite high resolution RGB single images, however, the work shows moderate performance for GSD larger than 15 cm [23]. All work is also tested with rather narrow datasets of very different sceneries which makes the validity of the results questionable and the comparison of methods difficult. It is therefore unclear, if deep learning on high resolution images will further improve, given the limitations of the data.

Another problem of still images is the impossibility to capture the dynamic behaviour of vehicles which is essential for many applications. For example, vehicle heading and speed are important indicators in traffic models. Although rapid retargeting for multi-angular image sequences with Worldview-2 is possible [21], the time interval of around one minute between consecutive images is too large for reasonable analysis.

For these reasons the paper addresses the problem of vehicle detection in satellite video. Such video was introduced 1999 by DLR-TubSat, since 2013 Planet’s SkySat-1 delivers up to 120 s, 30 Hz, 2K panchromatic video covering two areas of 1.1 km^2 with up to 80 cm GSD. China’s Jilin programme launched 2015, now provides even 4 MP color video.

²GSD is the spatial distance of two adjacent pixels on the image measured on the ground.

To the best of our knowledge this is the first work on using neural networks and deep learning to directly regress positions of vehicles *in satellite video*. Inspired by recent work on WAMI [17] this paper proposes to exploit the temporal consistency in satellite video by using a neural network and deep learning instead of using background subtraction or frame differencing, by this improving over the state-of-the-art in vehicle detection with satellite video. To overcome shortage of labelled video, this work follows in this context the novel idea of transfer learning by recognising similarity of WAMI and satellite video data.

To summarise, the contributions of this work are

- the confirmation of results in LaLonde et al. [17] which shows clearly improvement in vehicle detection (from 0.79 to 0.93 in F_1 score) when using a spatiotemporal convolutional network,
- empirical results showing the applicability of FoveaNet [17] to reduced resolution (0.91 F_1 score for 40% of the original image resolution and 0.79 F_1 score for 20%), yielding sizes of up to 3.6×1.8 px for vehicles which simulates satellite video and finally,
- a transfer learning approach that uses labelled WAMI data to train a detector for satellite video with 0.84 F_1 score which is comparable to the currently best (subspace) method E-LSD[33] with 0.83 F_1 score on the same data.

2. Related Work

Deep learning significantly improved previously handcrafted methods of object recognition [6]. Neural networks and back-propagation allow a learning formalism, where features and inference are jointly learnt from data in a neat end-to-end framework. Object detection is designed either as direct regression of bounding box image coordinates [24] or by using the idea of object proposals as intermediate step [25].

These developments triggered also work on deep learning for object detection in remote sensing [22, 23, 35, 15, 31, 8, 27, 17, 34, 13, 2, 29, 30]. Applying deep learning for remote sensing is challenging, as labels are very expensive for satellite data and good augmentation, transfer learning or even unsupervised methods circumventing this problem are currently unknown [38, 20]. Besides deep learning, object detection in remote sensing can be categorised

according to the approach taken as well as the sensor modality, i.e. satellite image, sequence of multi-angular satellite images, satellite video, aerial image and WAMI.

Applying a classifier on top of a sliding window is one possible approach. Using a convolutional neural network in combination with hard negative mining showed by a F_1 score of 0.7 reasonable results with 15 cm GSD on aerial images [15]. Following the golden standard [25], adapted variants of the base feature, region proposal and Fast R-CNN network have been proposed such as using skip connections in the base and focal loss [31], or using a dilated, multi-scale VGG16 as base in combination with hard negative mining [8] which gives AP and Recall larger than 0.8 in their experiments. Guo et al. [27] introduces proprietary base, region proposal and detection networks, but did not show results on vehicles. This approach is useful with aerial images, but fails entirely for 1m GSD video as shown by [34] (F_1 score of 0.5). Results on high resolution satellite images are still unknown in literature.

Another idea is to pixel-wise classify vehicle vs. background (semantic segmentation), e.g. by combining Inception and ResNet to give a heatmap. Assuming a fixed vehicle size and using non-maxima suppression gives excellent results [23] (F_1 score larger than 0.9). Imbert proposes a generative U-Net in combination with hard negative mining for satellite images but kept unfortunately results in absolute F_1 scores confidential.

Spatiotemporal information is a further cue important in object detection, especially with WAMI and satellite video. The standard is to use background subtraction (BGS) [35, 16, 28, 32, 1] and frame differencing (FD) [18, 4, 3], except Al-Shakarji et al. [2] who combined YOLO with spatiotemporal filtering on WAMI (F_1 score of 0.7), and Mou and Zhu [22] who use KLT tracking on video with a SegNet on overlapping multispectral data, however, they did not show results for vehicles. Zhang and Xiang [35] apply a ResNet classifier trained on CIFAR on proposals from a mixture of Gaussians foreground model, but did not show a proper evaluation.

The standard here is to apply connected component analysis [16, 28], saliency analysis, segmentation [32, 18], distribution fitting [4, 3] followed by morphology. F_1 scores of larger than 0.9 for ships and scores between 0.6 and 0.8 for vehicles on the Burji Khalifa [32], Valencia [4, 3] and Las Vegas [16]

videos suggest BGS, FD for larger objects. Both BGS and FD depend heavily on registration and parallax correction, hence, these methods introduce various nuisances for vehicles which are difficult to handle. Evaluation on single, selective scenes is further too narrow to draw a final conclusion.

Very recent work [33] suggests a subspace approach for discriminating vehicles and background. The idea shows potential with F_1 score results of larger than 0.8 on the simple Las Vegas video, which therefore needs further evaluation with more complex traffic patterns.

Another problem is the sparsity of vehicle occurrences in very large images as in WAMI which has been tackled by clustering the large images to draw attention to certain parts of the image and then to apply convolutional neural networks on single images [29][30] or multiple video frames [17] for final detection. Such clustering combined with deep spatiotemporal analysis shows excellent results on WAMI (F_1 score larger than 0.9) [17].

Also very recently tracking of airplanes, trains and vehicles has been considered for satellite video [10, 9, 26, 12], either by using optical flow [10, 9], correlation trackers (KLT) [26] or a combination of correlation and Kalman filters [12].

3. Methodology

With our goal of detecting moving vehicles in satellite videos, we were inspired by the work of Lalonde et al. [17], who designed two neural networks, denoted as ClusterNet and FoveaNet, to detect vehicles in WAMI. The ClusterNet proposes regions of objects (ROOBI) based on areas of interest (AOI), which are input to the FoveaNet. Instead of using the ClusterNet to determine ROOBIs we split the AOI into square tiles (ROOBIs) with size $N \times N$; e.g. $N=128$ px. The object detection based on the FoveaNet consists of two steps as depicted in Fig. 2.

3.1. FoveaNet and thresholding

The FoveaNet is a fully convolutional neural network (CNN) and consists of eight convolutional layers. The number of filters per convolution are 32, 32, 32, 256, 512, 256, 256 and 1. Their filter sizes are summarized in Tab. 3. After the first convolution a 2×2 max pooling is carried out. Moreover, during training the 6th and 7th convolutional layers have a 50% dropout. The heatmap is generated by the final 1×1 convolutional layer where each neuron gives a




Figure 2. The object detection process consists of two steps [17]: The FoveaNet predicts a heatmap, which indicates the likelihood that an object is at a given image coordinate. Vehicles are detected by thresholding the heatmap.

vote of the likelihood of a moving vehicle at pixel level.

The input to the network is a stack of frames with size $N \times N \times c$, where $N \times N$ is the ROOBI size and c depicts the number of consecutive adjoining frames in a stack. Hereinafter we refer to c as channels. Thereby, the CNN shall learn to predict the positions of the objects of the central frame. We believe the FoveaNet is capable to learn spatiotemporal features by feeding the network with stacks of multiple frames (e.g. $c=5$), which are especially important in lower resolution images as existing in satellite videos.

The ground truth is based on heatmaps H , which are created by superimposing Gaussian distributions, where the center of each distribution is the pixel position (x, y) of the vehicle in the image:

$$H(x, y) = \sum_{n=1}^N \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (1)$$

where n are the downsampled ground-truth coordinates provided in pixel positions and σ is the variance of the Gaussian blur. During training the network learns to minimize the Euclidean distance between the network output and the generated ground truth heatmaps.

The original FoveaNet uses ReLUs as activation functions. We discovered, however, the problem known as the ‘‘Dying ReLU’’ problem³. During training, a weight update triggered by a large gradient flowing through a ReLU can make the neuron inactive. If this happens, the gradient flowing through this ReLU will always be zero and the network continues to give the same output. In our trainings we frequently discovered this phenomenon ($\sim 71\%$ of the cases) using the Xavier initialization [11]. Hence, we replaced the ReLUs with either ELUs

(Exponential Linear Unit) or Leaky ReLUs.

The second step processes the predicted heatmap to determine the objects’ positions. For this, the heatmaps are converted into segmentation maps via OTSU thresholding [17]. If the segmented area is larger than a threshold α , then the center of the area is defined as the object position.

3.2. Transfer learning

To the best of our knowledge there are currently no annotated datasets of satellite videos publicly available. In contrast, there are some labeled WAMI datasets accessible; e.g. the WPAFB dataset⁴ contains over 160.000 annotated moving vehicles. WAMI and satellite images, however, differ considerably, among other things due to the different GSD. For instance, the WPAFB images have about four times higher GSD than the LasVegas video. Our core idea is to use transfer learning for a domain transfer from WAMI to satellite images. For this, we train our CNN based on the WPAFB dataset. Afterwards we fine-tune the CNN on satellite video data.

4. Experimental Evaluation and Results

Our network was trained from scratch using PyTorch - we used Adam with a learning rate of $1e-5$ and a batch size of 32. Data preparation includes frame registration to compensate camera motion.

We conducted three experiments. In the first experiment we carry out a baseline evaluation to reproduce the results of [17]. For the second experiment, we reduce the image resolution (GSD) of the WPAFB dataset. Thereby, the vehicle size in these low-resolution images is in the same order as in satellite videos. In the third experiment, we carry out a fine-tuning and evaluate the FoveaNet on satellite data.

³<http://cs231n.github.io/neural-networks-1>, 11/03/2019

⁴<https://www.sdms.afrl.af.mil/index.php?collection=wpafb2009>, 11/03/2019



Figure 3. AOI 40 contains a lot of dense traffic passing the intersection. On the contrary, AOI 41 contains mostly single vehicles driving on the road. Traffic patterns of AOI 34 are a combination of AOI 40 and AOI 41.

Detections are considered true positives if they are within a certain distance θ of a ground truth coordinate. If multiple detections are within this radius, the closest one is taken and the rest, if they do not have any other ground truth coordinates within the distance θ , are marked as false positives (FP). Any detections that are not within θ of a ground truth coordinate are also marked as FP. Ground truth coordinates which have no detections within θ are marked as false negatives. Quantitative results are compared in terms of precision, recall, and F_1 measure.

To compare our results with LaLonde et al. [17] we selected three of their AOIs (area of interest) - 34, 40 and 41. The contents of the AOIs 40 and 41 with respect to traffic patterns widely differ as displayed in Fig. 3. Whereas AOI 40 contains a lot of dense traffic at a main intersection, AOI 41 mainly consists of single vehicles on the road. AOI 34 is a combination of both traffic patterns. Data was split into training and testing in the following manner: AOI 34 was trained on AOIs 40 and 41. AOI 40 was trained on AOIs 34 and 41 and AOI 41 was trained on 34 and 40. In contrast to [17], we omitted AOI 42 for training as it is a sub-region of AOI 41.

For training and evaluation based on the WPAFB dataset, only frames with moving vehicles were included. We excluded frames without moving vehicles as our approach focuses solely on the detection and omits the region proposal part (ClusterNet) of [17]. A vehicle is defined as moving if it moves at least ω pixel within 5 frames.

4.1. Experiment 1: Baseline evaluation

In the first experiment we reproduced the results in [17]. For this, we set the following parameters: $N=100$ px (ROOBI edge length), $\sigma=2$ (variance of Gaussian blur), $\theta=40$ px (evaluation threshold),

$\omega=15$ px (threshold for removing stationary cars) and $\alpha=15$ px (threshold to disregard small segments). Tab. 1 indicates that our results are in the same order of magnitude than [17]. For instance, we achieve a F_1 score of 0.90 in AOI 34 ($c=5$), whereas Lalonde et al. have a F_1^* score of 0.93. The difference in the results is most likely due to the implementation differences of the second step, where we i.a. do not split connected regions into multiple detections. This presumption is confirmed looking at the evaluation results of AOI 40, where the differences of the F_1 score are greatest. AOI 40 contains a lot of dense traffic at the intersection resulting in connected regions, which cause false negative detections (Fig. 4). Furthermore, the results in Tab. 1 confirm that the network is learning spatiotemporal features which improve the overall performance comparing single versus multi-channels. For instance, the precision of AOI 34 increases from 0.73 ($c=1$) to 0.87 ($c=5$).

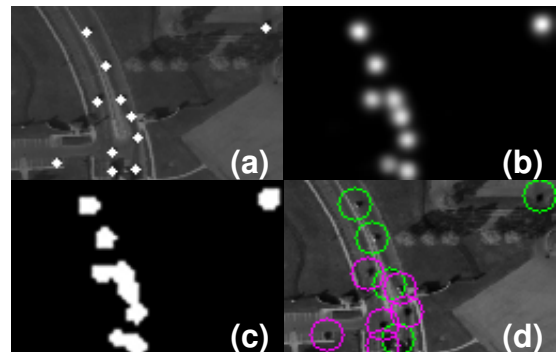


Figure 4. The detection of vehicles in crowded scenes is error-prone. Detection results of a ROOBI with reduced resolution (SF=0.2): (a) ground truth, (b) predicted heatmap, (c) after thresholding, (d) detected vehicles (green: true positives, pink: false negatives)

Table 1. Results are based on three AOIs of the WPAFB dataset with various channel sizes (c). For comparison, F_1^* scores of [17] are provided. Results of the second experiment include two scaling factors - 0.4 and 0.2.

		Experiment 1				Experiment 2					
		Full Resolution				Scaling factor 0.4			Scaling factor 0.2		
AOI	c	Prec.	Rec.	F_1	F_1^*	Prec.	Rec.	F_1	Prec.	Rec.	F_1
34	1	0.73	0.88	0.79		0.55	0.55	0.55	0.39	0.35	0.37
40	1	0.73	0.82	0.77		0.55	0.48	0.51	0.20	0.21	0.20
41	1	0.76	0.90	0.82		0.60	0.72	0.65	0.28	0.42	0.34
34	3	0.86	0.94	0.90		0.93	0.77	0.84	0.80	0.61	0.69
40	3	0.92	0.89	0.90		0.95	0.69	0.80	0.93	0.56	0.70
41	3	0.93	0.93	0.93		0.97	0.84	0.90	0.89	0.69	0.77
34	5	0.87	0.93	0.90	0.93	0.94	0.78	0.85	0.91	0.63	0.74
40	5	0.92	0.89	0.90	0.98	0.96	0.70	0.81	0.90	0.57	0.70
41	5	0.93	0.92	0.93	0.93	0.97	0.85	0.91	0.90	0.70	0.79

4.2. Experiment 2: Downscaled WPAFB dataset

For the second experiment we reduced the images by a scaling factor (SF) of 0.4 and 0.2 resulting in 40% and 20% of the original image resolution, respectively. We selected a SF of 0.2, because this factor reduces the typical vehicle object size in the WPAFB dataset from the order of 18×9 px to 3.6×1.8 px, which is like the vehicle size in satellite videos. The following parameters were set for the experiments: SF=0.4 with $N=100$ px, $\sigma=2$, $\theta=16$ px, $\omega=6$ px, $\alpha=15$ px and SF=0.2 with $N=100$ px, $\sigma=1$, $\theta=8$ px, $\omega=3$ px, $\alpha=3.5$ px. Comparing results of detections based on $c=1$ (Tab. 1) indicate that the performance significantly decreases with lower image resolutions; e.g. the F_1 score of AOI 40 decreases from 0.77 to 0.20 (SF=0.2). In contrast, the detection results significantly improve if the number of channels is increased. These results confirm our hypothesis that the learned spatiotemporal features are of great importance for detecting tiny objects such as vehicles under low resolution.

One of the main problems with low resolution images is the small distance between neighboring vehi-




Figure 5. Example of the SkySat-1 LasVegas video in which both AOIs are shown. AOI 1 (400×400 px) is used for evaluation and AOI 2 (600×400 px) for training. Two ROOBIs are sketched as yellow dashed rectangles.

cles as displayed in Fig. 4. In this case the FoveaNet creates a heatmap with a large number of connected regions, which result in a large number of false negative detections. To deal with small distances between neighboring cars we reduced the variance σ of Eq. 1, which improved the detection results. Otherwise, this issue has not been addressed in this work, although enhancing step 2 of the object detection will most likely improve the results.

4.3. Experiment 3: Satellite video

The third experiment is conducted to evaluate the detection performance of the FoveaNet on the panchromatic satellite SkySat-1 LasVegas video⁵ consisting of 700 frames, whose GSD is ~ 1.0 m and its frame rate is 30 fps. We defined two AOIs as illustrated in Fig. 5. While AOI 2 is mainly composed of straight parallel roads, AOI 1 contains additionally a bridge which results in more complex traffic patterns. The ground truth which was kindly shared by [33] consists of bounding boxes for moving vehicles. We used the center points of those bounding boxes as ground truth analogous to the WPAFB ground truth.

For training and evaluation we set $\theta=8$ px, $\alpha=4$ px, $\sigma=1$, $c=5$ and $N=128$ px. Additionally, we set SF=0.2 and $\omega=3$ px for training the WPAFB dataset. We observed in this experiment higher efficiency in training by replacing the ELUs with Leaky ReLUs.

Tab. 2 shows the results of nine individual experiments using FoveaNet with different filter sizes in the respective convolutional layers (Tab. 3). FoveaNet is trained on the 80% reduced WPAFB and directly applied to the LasVegas video. We observe high recall (>0.8) but average precision which proves applicability of transfer learning.

⁵<https://www.youtube.com/watch?v=IKNAY5ELUZY>

In contrast to LaLonde et al. [17], we do not observe large influence of the filter size to the final performance of the network. The argument that large filter sizes in the first layer are needed for spatial contextual information seems to be misleading, as context is introduced in higher layers of a deep network by the network’s receptive field. We argue that the filter size depends on the pixel distance of vehicles in consecutive frames so that the spatiotemporal network can exploit temporal information which is empirically confirmed by our experiments.

We then choose slightly smaller filter sizes (13-11-9-7-5-3-3-1) for the convolutional layers in FoveaNet, as this configuration shows best final results. We fine-tuned the network on AOI 2 which improved F_1 score from 0.55 to 0.84. A qualitative result of this experiment is shown in Fig. 1. The heat map of the network reconstructs amazingly well the ground truth. It detects not only cars but also buses and trucks which the network never saw before. Three experiments with varying filter sizes show further that filter sizes have minor influence on the result. We clearly see that our proposed method outperforms most methods for vehicle detection in satellite video except E-LSD[33] which is comparable to our results.

We then performed an experiment where we directly trained all layers of FoveaNet on AOI 2. Surprisingly, the overall results are only slightly worse which indicates that the learning problem is not as complex as for the WPAFB dataset. We conclude from all observations that FoveaNet learns to detect moving spots by characterising the slope of linear movement in spacetime which is a much simpler learning problem as learning spatiotemporal changes of visual appearance. However, pre-training on WPAFB is important for the network to generalise as can be seen in Fig. 6. Without pre-training the network is in this example not able to detect more complex motion patterns such as the moving vehicle on the bridge. It is an open question if such patterns could be learned by sole data augmentation.

Finally, we performed an experiment where we studied the effect of the frame rate of videos. Beside our baseline of considering every 10th image frame of the satellite video, we experimented with every 5th, 15th and 30th (1 fps) video frame. The results indicate less influence of higher frame rates on performance. This again supports our hypothesis that very simple features such as typical slopes of vehicle




Figure 6. From left to right. Top: input image and ground truth. Middle: estimated and thresholded heatmap, FoveaNet trained with AOI 2. Bottom: estimated and thresholded heatmap, FoveaNet after fine-tuning.

trajectories in spacetime are learned by the network. This presumption needs however further experiments and insight.

5. Conclusion

This paper considers vehicle detection in satellite video. Vehicle detection in remote sensing is challenging as the objects usually appear tiny compared to the size of typical aerial and satellite images and discrimination of objects of interest from background is frequently ambiguous. Satellite video is a very new modality introduced 2013 by Skybox (now Planet) which might overcome the problem by introducing high temporal resolution. This allows to exploit temporal consistency of moving vehicles as inductive bias. Current state-of-the-art methods use either background subtraction, frame differencing or subspace learning in video, however, performance is currently limited (0.26 - 0.82 F_1 score).

The method in this paper is motivated by recent work in WAMI which exploits video in spatiotemporal convolutional networks[17]. We apply FoveaNet to the domain of satellite video by transfer learning the network with WPAFB and a small amount of available labelled video frames of the SkySat-1 LasVegas video which yields comparable results (0.84 F_1 score). Several ablation studies show minor influence of the filter sizes in the convolutional layers and minor influence of the frame rate (tempo-

WPAFB				LasVegas AOI 1				SOTA			
Conf.	Prec.	Rec.	F_1		Prec.	Rec.	F_1		Prec.	Rec.	F_1
1	0.56	0.67	0.61	scratch	0.87	0.80	0.83	ViBe[5]	0.58	0.17	0.26
2	0.46	0.76	0.57	fine-tuning				GMMv2[39]	0.65	0.27	0.38
3	0.40	0.79	0.53	1	0.84	0.82	0.83	GMM[14]	0.46	0.50	0.48
4	0.42	0.81	0.55	4	0.86	0.82	0.84	Fast-RCNN-LRP[34]	0.58	0.44	0.50
5	0.43	0.85	0.58	9	0.76	0.85	0.80	GoDec[36]	0.95	0.36	0.52
6	0.47	0.80	0.60	skip 5	0.84	0.83	0.84	RPCA-PCP[7]	0.94	0.41	0.57
7	0.46	0.82	0.59	skip 10	0.86	0.82	0.84	Decolor[37]	0.77	0.59	0.67
8	0.46	0.83	0.59	skip 15	0.85	0.81	0.83	LSD[19]	0.87	0.71	0.78
9	0.45	0.70	0.55	skip 30	0.83	0.82	0.83	E-LSD[33]	0.85	0.79	0.82

Table 2. Left: Evaluation results of nine different filter size configurations (see Tab. 3) of the FoveaNet. Middle: Results of the FoveaNet trained from scratch, fine-tuned with different filter sizes and different fps (conf. 4). Right: Evaluation results of state-of-the-art (SOTA) methods are presented.

conf.	filter size	conf.	filter size
1	19-17-15-13-11-9-7-1	6	9-7-5-3-3-3-3-1
2	17-15-13-11-9-7-5-1	7	7-5-3-3-3-3-3-1
3	15-13-11-9-7-5-3-1	8	5-3-3-3-3-3-3-1
4	13-11-9-7-5-3-3-1	9	3-3-3-3-3-3-3-1
5	11-9-7-5-3-3-3-1		

Table 3. Filter size configurations of the various experiments. Conf. 3 corresponds to the filter sizes suggested by LaLonde et al. [17].

ral resolution) on the overall result. This indicates a much simpler learning problem than for the original high-resolution WAMI data, however, we show that temporal information is essential for a good detection performance. Improvements of FoveaNet, e.g. including the final segmentation of the heat map into the network, are left for future work.

Acknowledgment

This research was supported by the Austrian Research Promotion Agency (FFG) under grant MiTrAs-867030 and by the European Union’s H2020 programme under grant FOLDOUT-787021. The imagery and video used in this work is in courtesy of the U.S. air force research laboratory sensors directorate layered sensing exploitation division and Planet Inc. under creative common CC-BY-NC-SA⁶.

References

- [1] S. A. Ahmadi, A. Ghorbanian, and A. Mohamadzadeh. Moving vehicle detection, tracking and traffic parameter estimation from a satellite video: a perspective on a smarter city. *IJRS*, 40(22):8379–8394, 2019. 3

- [2] N. Al-Shakarji, F. Bunyak, H. Aliakbarpour, G. Seetharaman, and K. Palaniappan. Multi-cue vehicle detection for semantic video compression in georegistered aerial videos. In *CVPR*, June 2019. 2, 3
- [3] W. Ao, Y. Fu, X. Hou, and F. Xu. Needles in a haystack: Tracking city-scale moving vehicles from continuously moving satellite. *IEEE TIP*, preprint:1–1, 2019. 2, 3
- [4] W. Ao, Y. Fu, and F. Xu. Detecting tiny moving vehicles in satellite videos. *CoRR*, abs/1807.01864, 2018. 2, 3
- [5] O. Barnich and M. Van Droogenbroeck. Vibe: A universal background subtraction algorithm for video sequences. *IEEE TIP*, 20(6):1709–1724, June 2011. 8
- [6] R. Benenson, M. Mathias, T. Tuytelaars, and L. Van Gool. Seeking the strongest rigid detector. In *CVPR*, June 2013. 1, 2
- [7] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, June 2011. 8
- [8] P. Ding, Y. Zhang, W.-J. Deng, P. Jia, and A. Kuijper. A light and faster regional convolutional neural network for object detection in optical remote sensing images. *ISPRS J. P&RS*, 141:208 – 218, 2018. 2, 3
- [9] B. Du, S. Cai, C. Wu, L. Zhang, and D. Tao. Object tracking in satellite videos based on a multi-frame optical flow tracker. *CoRR*, abs/1804.09323, 2018. 3
- [10] B. Du, Y. Sun, S. Cai, C. Wu, and Q. Du. Object tracking in satellite videos by fusing the kernel correlation filter and the three-frame-difference algorithm. *IEEE Geoscience and Remote Sensing Letters*, 15(2):168–172, Feb 2018. 3
- [11] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010. 4

⁶<https://creativecommons.org/licenses/by-nc-sa/4.0/legalcode>

- [12] Y. Guo, D. Yang, and Z. Chen. Object tracking on satellite videos: A correlation filter-based tracking method with trajectory correction by kalman filter. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(9):3538–3551, Sep. 2019. 3
- [13] J. Imbert. Fine-tuning of fully convolutional networks for vehicle detection in satellite images: Data augmentation and hard examples mining. Master’s thesis, KTH, 2019. 2
- [14] P. KaewTraKulPong and R. Bowden. *Video-Based Surveillance Systems*, chapter An Improved Adaptive Background Mixture Model for Real-time Tracking with Shadow Detection, pages 135–144. Springer, 2002. 8
- [15] Y. Koga, H. Miyazaki, and R. Shibasaki. A cnn-based method of vehicle detection from aerial images using hard example mining. *Remote Sensing*, 10, Jan. 2018. 2, 3
- [16] G. Kopsiaftis and K. Karantzas. Vehicle detection and traffic density monitoring from very high resolution satellite video data. In *IGARSS*, July 2015. 2, 3
- [17] R. LaLonde, D. Zhang, and M. Shah. Clusternet: Detecting small objects in large scenes by exploiting spatio-temporal information. In *CVPR*, June 2018. 2, 3, 4, 5, 6, 7, 8
- [18] H. Li, L. Chen, F. Li, and M. Huang. Ship detection and tracking method for satellite video based on multiscale saliency and surrounding contrast analysis. *Journal of Applied Remote Sensing*, 13(2):1 – 17, 2019. 2, 3
- [19] X. Liu, G. Zhao, J. Yao, and C. Qi. Background subtraction based on low-rank and structured sparse decomposition. *IEEE TIP*, 24(8):2502–2514, Aug 2015. 8
- [20] L. Ma, Y. Liu, X. Zhang, Y. Ye, G. Yin, and B. A. Johnson. Deep learning in remote sensing applications: A meta-analysis and review. *ISPRS J. P&RS*, 152:166 – 177, 2019. 2
- [21] L. Meng and J. P. Kerekes. Object tracking using high resolution satellite imagery. *IEEE J. of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(1):146–152, 2 2012. 2
- [22] L. Mou and X. X. Zhu. Spatiotemporal scene interpretation of space videos via deep neural network and tracklet analysis. In *IGARSS*, July 2016. 2, 3
- [23] T. Mundhenk, G. Konjevod, W. Sakla, and K. Boakye. A large contextual dataset for classification, detection and counting of cars with deep learning. In *ECCV*, Oct. 2016. 2, 3
- [24] J. Redmon and A. Farhadi. Yolov3: An incremental improvement. *CoRR*, abs/1804.02767, 2018. 2
- [25] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 1, 2, 3
- [26] J. Shao, B. Du, C. Wu, J. Wu, R. Hu, and X. Li. Vcf: Velocity correlation filter, towards space-borne satellite video tracking. In *ICME*, July 2018. 3
- [27] G. Wei, Y. Wen, Z. Haijian, and H. Guang. Geospatial object detection in high resolution satellite images based on multi-scale convolutional neural network. *Remote Sensing*, 10(1), 2018. 2, 3
- [28] A. Xu, J. Wu, G. Zhang, S. Pan, T. Wang, Y. Jang, and X. Shen. Motion detection in satellite video. *Journal of Remote Sensing and GIS*, 6(2):1–9, 2017. 2, 3
- [29] F. Yang, H. Fan, P. Chu, E. Blasch, and H. Ling. Clustered object detection in aerial images. *CoRR*, abs/1904.08008, 2019. 2, 3
- [30] F. Yang, H. Fan, P. Chu, E. Blasch, and H. Ling. Clustered object detection in aerial images. In *ICCV*, 2019. 2, 3
- [31] M. Y. Yang, W. Liao, X. Li, and B. Rosenhahn. Vehicle detection in aerial images. *CoRR*, abs/1801.07339, 2018. 2, 3
- [32] T. Yang, X. Wang, B. Yao, J. Li, Y. Zhang, Z. He, and W. Duan. Small moving vehicle detection in a satellite video of an urban area. *Sensors*, 16(9):1528, Sept. 2016. 2, 3
- [33] J. Zhang, X. Jia, and J. Hu. Error bounded foreground and background modeling for moving object detection in satellite videos. *CoRR*, 2019. 1, 2, 3, 6, 7, 8
- [34] J. Zhang, X. Jia, and J. Hu. Local region proposing for frame-based vehicle detection in satellite videos. *Remote Sensing*, 11(20):2372, Oct. 2019. 1, 2, 3, 8
- [35] X. Zhang and J. Xiang. Moving object detection in video satellite image based on deep learning. In *LIDAR Imaging Detection and Target Recognition*, volume 10605, pages 1149 – 1156. SPIE, 2017. 2, 3
- [36] T. Zhou and D. Tao. Godec: Randomized low-rank & sparse matrix decomposition in noisy case. In *ICML*, 2011. 8
- [37] X. Zhou, C. Yang, and W. Yu. Moving object detection by detecting contiguous outliers in the low-rank representation. *IEEE TPAMI*, 35(3):597–610, March 2013. 8
- [38] X. X. Zhu, D. Tuia, L. Mou, G. Xia, L. Zhang, F. Xu, and F. Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):8–36, Dec 2017. 2
- [39] Z. Zivkovic. Improved adaptive gaussian mixture model for background subtraction. In *CVPR*, Aug. 2004. 8