SpotNet: Self-Attention Multi-Task Network for Object Detection

Hughes Perreault, Guillaume-Alexandre Bilodeau, Nicolas Saunier Polytechnique Montral Montral, Canada Email: {hughes.perreault, gabilodeau, nicolas.saunier}@polymtl.ca Maguelonne Hritier Genetec Montral, Canada Email: mheritier@genetec.com

Abstract-Humans are very good at directing their visual attention toward relevant areas when they search for different types of objects. For instance, when we search for cars, we will look at the streets, not at the top of buildings. The motivation of this paper is to train a network to do the same via a multi-task learning approach. To train visual attention, we produce foreground/background segmentation labels in a semi-supervised way, using background subtraction or optical flow. Using these labels, we train an object detection model to produce foreground/background segmentation maps as well as bounding boxes while sharing most model parameters. We use those segmentation maps inside the network as a self-attention mechanism to weight the feature map used to produce the bounding boxes, decreasing the signal of non-relevant areas. We show that by using this method, we obtain a significant mAP improvement on two traffic surveillance datasets, with state-of-the-art results on both UA-DETRAC and UAVDT.

Keywords-Object Detection; Segmentation; Self-Attention; Multi-Task Learning; Traffic Scenes;

I. INTRODUCTION

There is increasing interest in automatic road user detection for intelligent transportation systems, advanced driver assistance systems, traffic surveillance, etc. Road user detection has its own set of challenges and difficulties, such as the high speed of some road users, the frequent occlusion between them and the small size of road users appearing afar. Despite huge improvements in the last years thanks to advancements in deep learning-based methods, results still need to be improved for reliable practical applications.

Recently, a new family of object detectors were proposed based on keypoint detection rather than based on bounding box classification [1]–[3]. This approach presents several advantages, including not having to manually design anchor boxes and having to process fewer candidate boxes. Detecting objects in this way is deceptively simple and elegant, and quite fast. It yields state-of-the-art accuracy results on several datasets. Therefore, in this work, we build upon CenterNet [3] by designing a novel convolutional neural network (CNN) model that directs its attention towards the areas of interest and thus decreases the probability of having false detections in incongruous areas.

Our contributions are: 1) a self-attention mechanism based on multi-task learning (object detection and segmentation)



Figure 1. A visualisation of the attention map produced by SpotNet on top of its corresponding image, from the UAVDT [6] dataset.

and 2) a semi-supervised training method that capitalizes on automatic foreground/background segmentation annotations.

The idea of attention and self-attention has been around for some time now, most notably in image captioning [4] and natural language processing (NLP) [5]. In those works, neural networks are trained to learn which parts of the input are the most important to solve the task. But they do so progressively, using recurrent neural networks. Can a simple CNN learn which areas it should use to increase its visual attention? In this work, we show that it is indeed possible and beneficial for object detection by using a semi-supervised training approach and multi-task learning. The network is trained for both object detection and foreground/background segmentation, the latter being also used to weight object detection feature maps. Indeed, the foreground/background segmentation is used in an internal attention mechanism that gives more weight to areas useful for detection. In figure 1, we can visualize what the network learns from this approach, that is to concentrate the keypoint search on areas where there are indeed road users, and therefore reducing the response of any other neuron. One can see this process as shining a spotlight on relevant areas and dimming the lights everywhere else. Therefore, we named our method, SpotNet.

This attention approach is particularly beneficial for keypoint-based methods since we are globally looking for keypoints on the whole image at the same time, and not just classifying the object in a cropped bounding box. However, a question remains. How can we train such a self-attention process? Typically, object detection datasets do not provide the segmentation ground-truth since it is very costly and time-consuming to produce. Instead, we rely on classical computer vision techniques to generate automatic pixel-wise annotation labels and on datasets providing video sequences instead of single frames to train the network. In the case of fixed camera video sequences, we successfully employ a background subtraction method to obtain the automatic annotations, while in the case of moving camera video sequences, we rely on dense optical flow for the same purpose.

Although we use imperfect foreground/background segmentation annotations, we can train a network to produce quality segmentation maps by using multi-task learning. The detection and segmentation tasks are trained jointly by sharing all the parameters of the backbone network. Both tasks are mutually beneficial. Indeed, by producing a better segmentation, the object detection task benefits from a better attention mechanism. And by producing better object detection, the parameters of the backbone network get better at recognizing the features of interest from the images to improve the segmentation maps. We validated our method on two popular traffic scene datasets, and we show that our method is the state-of-art on these datasets by improving significantly the performance of the base network (CenterNet).

II. RELATED WORK

Object detection as meant in this paper is the task of drawing a rectangular bounding box around objects of interest in an image, as well as producing a class label for each box. All state-of-the-art object detection methods have been based on deep learning since its rise. They can broadly be split up into two main categories, two-stage and one-stage methods.

Two-stage methods divide the task of object detection into two steps, producing a set of object candidates, and then computing a score, a label and a coordinate offset for each box. The first deep learning-based method was R-CNN [7], which used an external method to produce box candidates, namely selective search [8]. It then passed each candidate in a CNN to compute features for each box. A classification is done on those features by a SVM afterwards. Fast R-CNN [9] aimed to increase the speed of R-CNN by passing the whole image through a CNN once, and afterwards just cropped the relevant parts of the feature map for each box candidate for classification. Faster R-CNN [10] is a further improvement that introduced the RPN, a region proposal network that shares most of its parameters with the classification and regression parts, making it even faster and more efficient than its predecessors. RFCN [11] further builds upon Faster R-CNN by learning to detect and classify parts of objects and then using a grid of parts to vote on each object. Cascade R-CNN [12] addresses the problems of the mismatch between the minimum IOU (Intersection over union) used to evaluate during inference, and the minimum IOU used to select a positive sample during training. They also address overfitting during training by training while progressively increasing the IOU thresholds.

One-stage methods aim to reduce the processing time of two-stage methods by removing the candidate proposal phase and by detecting objects directly from the feature map. The first one-stage method was YOLO [13] which divides the input image into a regular grid and makes each cell predict two bounding boxes. Further iterations of the method, YOLOv2 [14] and YOLOv3 [15] built upon it by using anchor boxes, a better backbone network and several other tweaks. SSD [16] addresses the multi-scale detection problem by combining feature maps at multiple spatial resolutions and then applying anchor boxes to look for objects. RetinaNet [17] uses an FPN (Feature pyramid network) [18] to produce a multi-scale pyramid of features and applies a set of anchor boxes followed by non-maximal suppression to find objects. CornerNet [1] uses the Hourglass network [19] paired with corner pooling layers to detect a set of top-left corners and bottom-right corners, and combines them with a learned embedding. Keypoint Triplets [2] builds upon CornerNet by improving the corner pooling layers, and by also detecting a center keypoint to validate each object. Objects as Points [3] detects an object as a center keypoint and regresses the size of the object to find the bounding box.

Attention mechanisms in object detection have been around for a while. In Geometric Proposal for Faster R-CNN [20], the authors re-rank the proposals of the region proposal network depending on a geometric estimation of the scene, outperforming the standard Faster R-CNN by a large margin. Their geometric estimation of the scene is mostly based on vehicle scale. The HAT [21] method uses a hierarchical attention mechanism that first trains a partspecific attention model. Then an LSTM models the relations between those parts, making it a part-aware detector. FG-BR Net [22] uses background subtraction methods to produce a foreground image that is fed as another input to the network. They also introduce a feedback process from the detection outputs to the background subtraction to keep static objects in the foreground image. Compared to these models, our attention mechanism is simple, elegant and fast. Furthermore, we do not need background subtraction at inference, only during the training phase.

III. PROPOSED METHOD

Figure 2 shows a detailed overview of our complete model.



Figure 2. Overview of SpotNet: the input image first passes through a double-stacked hourglass network; the segmentation head then produces an attention map that multiplies the final feature map of the backbone network; the final center keypoint heatmap is then produced as well as the size and coordinate offset regressions for each object.

A. Base network

Our method is based upon CenterNet [3], not to be confused with the homonym method CenterNet, or keypoint triplets [2]. This method trains a backbone network to recognize the center point of objects by assigning the center pixel of a box to be the ground-truth center and gives a reduced loss for other close points. The width and height of the bounding box are regressed, as well as the coordinate offset of the box (to compensate from the error caused by the smaller spatial resolution of the output). The final output is thus a center point heatmap for each possible label, an object size for each point, and an offset for each point, the size and offset being label agnostic.

B. Multi-task Learning

Our main idea is to train a network to perform multiple tasks, to make it better for at least one of the tasks. In our case, we train a network to perform segmentation of objects of interest while performing bounding box detection, thus making the shared parameters more generic and less prone to overfitting. To do this, we add a two-class (foreground/background) segmentation head to the network and train this head with semi-supervised annotations from training datasets (more details in subsection III-D).

The added segmentation head takes as input a feature map that has been reduced by a factor of four in terms of spatial dimension when compared to the input. It consists of three 3×3 convolutions, with upsampling layers in between. The channel dimension is reduced to 1 in the last convolution, thus resulting in a segmentation map that is the same width and height as the input, with a single channel. The loss, L_{seg} , used to train this head is the binary cross-entropy, given by

$$L_{seg} = -\frac{1}{N} \sum_{i=1}^{N} y_i * \log(x_i) + (1 - y_i) * \log(1 - x_i),$$
(1)

where y_i is the annotation label for sample *i*, x_i its predicted label by the network and N the number of samples. We found out during our experiments that it works better than the initial mean squared error loss that we had tried initially.

C. Self-Attention Mechanism

To further benefit from our learned segmentation map, we implement a simple yet effective self-attention mechanism within the network. Once we obtain our segmentation map, we downsample it by a factor of 4 to reduce it to the spatial dimension of the original feature map. To attenuate the response at locations unlikely to contain an object of interest, we multiply every channel of the feature map with our segmentation map, thus reducing the probability of false positives in irrelevant areas.

D. Semi-Supervised annotations

To train our model to produce foreground/background segmentation maps, we had to produce semi-supervised

pixel-wise segmentation annotations. To do that, we took advantage of having access to full video sequences, despite training and evaluating on a single frame at a time. For the fixed camera video sequences, we used the background subtraction method PAWCS [23]. Since background subtraction is not designed to work with a moving background, for the moving camera video sequences, we used Farneback optical flow [24] followed by some basic image processing and a threshold on motion magnitude. For both automatic two-class segmentation results, we then do an intersection with the ground-truth bounding boxes for each frame to reduce noise and to obtain pixel-wise segmentation annotations only for the object categories to detect. All other object categories, not inside ground-truth training bounding boxes, are therefore labelled as background. This results in fairly good foreground/background segmentation maps, with sometimes squared corners at one or more sides, due to the intersection with bounding boxes, as can be seen in Figure 3. In our experiments, we find that not only are these non-perfect segmentation annotations good enough to train good attention maps, it also allows our segmentation head to produce segmentation maps comparable to good unsupervised foreground/background segmentation methods.

It should be noted that although our method requires videos for training to obtain the semi-supervised segmentation annotations, once trained, it can be applied to single images.



Figure 3. Example of semi-supervised annotations on UA-DETRAC [25] produced by PAWCS [23] and the intersection with the ground-truth bounding boxes.

E. Training for multiple tasks

To adapt the training loss of the whole network, we added the binary cross-entropy loss of our segmentation head (equation 1) to the original CenterNet loss. The center point heatmap loss L_{heat} is calculated with the focal loss [17], and the losses for the regressions for the offset L_{off} and width/heigth L_{WH} are formulated as L1 losses as in the original paper [3]. The total loss L_{tot} is given by

$$L_{tot} = L_{heat} + L_{off} + L_{seq} + 0.1 * L_{WH}.$$
 (2)

The total loss is thus the sum of all losses, with the width and height regression having less weight than the others, 0.1 compared to 1.

IV. EXPERIMENTS

A. Datasets

To validate the effectiveness of our method, we trained and evaluated it against other state-of-the-art methods on two datasets of traffic scenes, namely UA-DETRAC [25] and UAVDT [6]. Figure 4 and Figure 5 show example frames of UA-DETRAC and UAVDT respectively, with their groundtruth. These two datasets were captured with very different settings, UA-DETRAC being filmed with a fixed camera for every scene, and UAVDT with a moving camera. Both datasets have pre-determined test sets, and we used a subset of the training data to do the validation.



Figure 4. Sample from UA-DETRAC with the ground-truth bounding boxes in yellow.



Figure 5. Sample from UAVDT with the ground-truth bounding boxes in yellow.

Evaluation is done using the Matlab code provided by the authors of both datasets. A strict training and validation protocol was followed and the testing data was never seen by the network before the final evaluation. The performance measure used for evaluation is the mAP, the mean Average Precision, with a minimum IOU of 0.7 between inferred and ground-truth bounding boxes. The minimum IOU is the minimum overlap of a bounding box with the ground-truth to be considered a true detection. The IOU is computed as the intersection of the boxes divided by the union of the boxes. The mean average precision is the mean of the average precisions for all classes for multiple values of recall, ranging from 0 to 1 with small steps, typically of 0.1.

B. Implementation Details

We used the stacked hourglass network as our backbone because it shows the best performance for keypoint estimation. This network is composed of modules of downsampling and convolutions followed by upsampling and convolutions with skip connections in an encoder-decoder fashion. For our experiments, we use the Hourglass-104 version as in [1] which stacks two encoder-decoder modules. We implemented the model in PyTorch 0.4.1 using Cuda 10.0. Experiments were run on a workstation with 32 GB of RAM and a NVIDIA GTX 1080Ti GPU. The Github repository for this project is https://github.com/hu64/SpotNet.

C. Object detection results

The experimental results are shown in Table I for UA-DETRAC and in Table II for UAVDT. We outperform our baseline, CenterNet, by a very significant margin on both datasets while being the state-of-art results on both datasets as well. The results are very coherent, showing approximately the same percentage of improvement over CenterNet on both datasets, the absolute value on UAVDT being smaller.

For UA-DETRAC, not only do we outperform all previously published results, we do so in every category, showing the benefit of our self-attention mechanism based on multitask learning. Moreover, the improvements are particularly impressive for the category hard and cloudy, meaning that our model is particularly good for hard examples. It is interesting to note that the improvement for the easy category is very small, due to the mAP values being already very high. Nonetheless, improvement is consistent across all categories. At the moment of writing, our model outperforms every published result on this dataset, including ensemble models from challenges [26].

The UAVDT dataset is more difficult than UA-DETRAC due to its high density of small vehicles and aerial point of view, but the percentage of improvement remains consistent. Our model also outperforms every published result on this dataset by a very significant margin.

D. Foreground/Background segmentation results

Although that was not our principal objective, it is nonetheless interesting to see how we do on specialized foreground/background benchmarks. To produce results, we used our best model trained on UA-DETRAC and ran it on three sequences of the changedetection.net dataset [32] containing only vehicles (because UA-DETRAC includes only annotations for vehicles). To obtain the foreground, we took the attention maps produced by our network, applied a binary threshold and then masked the resulting image with the bounding boxes detected by our network to remove noise. We can see in table III that our method produces competitive results, although we do not quite reach state-ofthe-art foreground/background performance. Figure 6 shows qualitative results on a few frames. Our method does not always fit the object boundaries very well. This is expected since the training annotations are imperfect. Nevertheless, we outperform several classical methods, at no additional cost when producing bounding boxes. It is important to note that a limitation of our model is that it must be trained on the objects we want to segment.

V. DISCUSSION

A. Ablation study

To detail the contribution of each part of our model, we conducted an ablation study on UA-DETRAC. Table IV shows that even though multi-task learning helps, the biggest contribution comes from combining our attention process with it. To further understand the contribution of each part, we draw the precision/recall curve (Figure 7) compared to several other methods on UA-DETRAC. On this curve, we can note that the multi-task learning by itself (SpotNet No Attention) helps to be more precise, but does not help to detect more objects, i.e. to reach improved values of recall. On the other hand, the attention mechanism does both, it helps to be even more precise for the same values of recall (fewer false positives), and it also allows the model to detect more and reach significantly higher values of recall.

Since the network is looking for keypoints on the whole image, it is natural that concentrating the search on learned foreground pixels will increase the probability that the keypoints found belong to the objects of interest, thus reducing the rate of false positives. Furthermore, the experiments show that this increases recall because the network can concentrate on useful information.

It is expected that learning the segmentation task jointly with the object detection task can be mutually beneficial since both tasks have a large overlap in what needs to be learned. The main difference is that object detection needs to separate instances, while segmentation needs a more precise border around the objects. We show that semi-supervised annotations are good enough for our purpose, and multitask learning by itself, based on those annotations, improves precision.

 Table I

 Results on the UA-DETRAC dataset [25]. 3D-DETNET RESULTS ARE FROM [27], AND OTHERS RESULTS ARE REPORTED AS IN THE RESULTS SECTION OF THE UA-DETRAC WEBSITE (BOLDFACE: BEST RESULT, Italic: INDICATES OUR BASELINE).

Model	Overall	Easy	Medium	Hard	Cloudy	Night	Rainy	Sunny
SpotNet (ours)	86.80%	97.58 %	92.57%	76.58%	89.38%	89.53%	80.93%	91.42%
CenterNet [2]	83.48%	96.50%	90.15%	71.46%	85.01%	88.82%	77.78%	88.73%
FG-BR_Net [22]	79.96%	93.49%	83.60%	70.78%	87.36%	78.42%	70.50%	89.8%
HAT [21]	78.64%	93.44%	83.09%	68.04%	86.27%	78.00%	67.97%	88.78%
GP-FRCNNm [20]	77.96%	92.74%	82.39%	67.22%	83.23%	77.75%	70.17%	86.56%
R-FCN [11]	69.87%	93.32%	75.67%	54.31%	74.38%	75.09%	56.21%	84.08%
EB [28]	67.96%	89.65%	73.12%	53.64%	72.42%	73.93%	53.40%	83.73%
Faster R-CNN [10]	58.45%	82.75%	63.05%	44.25%	66.29%	69.85%	45.16%	62.34%
YOLOv2 [14]	57.72%	83.28%	62.25%	42.44%	57.97%	64.53%	47.84%	69.75%
RN-D [29]	54.69%	80.98%	59.13%	39.23%	59.88%	54.62%	41.11%	77.53%
3D-DETnet [27]	53.30%	66.66%	59.26%	43.22%	63.30%	52.90%	44.27%	71.26%



Figure 6. Example of foreground/background segmentation maps obtained with several segmentation methods. First row: frame 1015 of "highway", second row: frame 967 of "traffic", third row: frame 883 of "boulevard".

 Table II

 RESULTS ON THE UAVDT [6] DATASET (BOLDFACE: BEST RESULT, Italic: INDICATES OUR BASELINE).

Model	Overall		
SpotNet (Ours)	52.80%		
CenterNet [2]	51.18%		
Wang et al. [30]	37.81%		
R-FCN [11]	34.35%		
SSD [16]	33.62%		
Faster-RCNN [10]	22.32%		
RON [31]	21.59%		

Table III Results on the changedetection.net [32] dataset. Results are averaged for sequences "highway", "traffic" and "boulevard" (**Boldface**: best result).

Model	Average F-Measure
PAWCS [23]	0.872
SuBSENSE [33]	0.831
SpotNet (Ours)	0.806
SGMM [34]	0.766
KNN [35]	0.731
GMM [36]	0.709

B. Limitations of our Model

One of the limitations of our model is the fact that it needs semi-supervised annotations to be trained properly. However, we believe that in most real-world applications, video sequences are available and we can thus run background subtraction or optical flow to generate them. In other cases, pre-trained semantic segmentation methods could be used to obtain the desired annotations.

Attention Multi-Task Overall Medium Hard Cloudy Night Sunny Easy Rainy 86.80% 97.58% 92.57% 76.58% 89.38% 89.53% 80.93% 91.42% 84.57% 96.72% 90.85% 73.16% 86.53% 88.76% 78.84% 90.10%

90.15%

71.46%

85.01%

88.82%

83.48%

96.50%



Figure 7. Precision/Recall curve of our model compared with a variant and other methods.

VI. CONCLUSION

In this paper, we presented a novel multi-task model equipped with a self-attention process, and we trained it with semi-supervised annotations. We show that these improvements allow us to reach state-of-the-art performance on two traffic scenes datasets with different settings. We argue that not only does this improve accuracy by a large margin, it also provides instance segmentations of the road users almost at no cost.

ACKNOWLEDGMENT

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), [RDCPJ 508883 - 17], and the support of Genetec.

REFERENCES

- H. Law and J. Deng, "Cornernet: Detecting objects as paired keypoints," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 734–750.
- [2] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "Centernet: Keypoint triplets for object detection," in *Proceedings* of the IEEE International Conference on Computer Vision, 2019, pp. 6569–6578.
- [3] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," arXiv preprint arXiv:1904.07850, 2019.

[4] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *International conference on machine learning*, 2015, pp. 2048–2057.

88.73%

77.78%

- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing* systems, 2017, pp. 5998–6008.
- [6] D. Du, Y. Qi, H. Yu, Y. Yang, K. Duan, G. Li, W. Zhang, Q. Huang, and Q. Tian, "The unmanned aerial vehicle benchmark: Object detection and tracking," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 370–386.
- [7] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *The IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2014.
- [8] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *International journal of computer vision*, vol. 104, no. 2, pp. 154– 171, 2013.
- [9] R. Girshick, "Fast r-cnn," in *The IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [10] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [11] J. Dai, Y. Li, K. He, and J. Sun, "R-fcn: Object detection via region-based fully convolutional networks," in Advances in Neural Information Processing Systems 29. Curran Associates, Inc., 2016, pp. 379–387.
- [12] Z. Cai and N. Vasconcelos, "Cascade r-cnn: Delving into high quality object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6154–6162.
- [13] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [14] J. Redmon and A. Farhadi, "Yolo9000: Better, faster, stronger," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [15] —, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.

 Table IV

 ABLATION STUDY ON THE UA-DETRAC [25] DATASET.

- [16] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.
- [17] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [18] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [19] A. Newell, K. Yang, and J. Deng, "Stacked hourglass networks for human pose estimation," in *European conference* on computer vision. Springer, 2016, pp. 483–499.
- [20] S. Amin and F. Galasso, "Geometric proposals for faster rcnn," in 2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS). IEEE, 2017, pp. 1–6.
- [21] S. Wu, M. Kan, S. Shan, and X. Chen, "Hierarchical attention for part-aware face detection," *International Journal of Computer Vision*, vol. 127, no. 6-7, pp. 560–578, 2019.
- [22] Z. Fu, Y. Chen, H. Yong, R. Jiang, L. Zhang, and X.-S. Hua, "Foreground gating and background refining network for surveillance object detection," *IEEE Transactions on Image Processing*, vol. 28, no. 12, pp. 6077–6090, 2019.
- [23] P.-L. St-Charles, G.-A. Bilodeau, and R. Bergevin, "A selfadjusting approach to change detection based on background word consensus," in 2015 IEEE winter conference on applications of computer vision. IEEE, 2015, pp. 990–997.
- [24] G. Farnebäck, "Two-frame motion estimation based on polynomial expansion," in *Scandinavian conference on Image analysis*. Springer, 2003, pp. 363–370.
- [25] L. Wen, D. Du, Z. Cai, Z. Lei, M.-C. Chang, H. Qi, J. Lim, M.-H. Yang, and S. Lyu, "UA-DETRAC: A New Benchmark and Protocol for Multi-Object Detection and Tracking," *arXiv CoRR*, vol. abs/1511.04136, 2015.
- [26] S. Lyu, M.-C. Chang, D. Du, W. Li, Y. Wei, M. Del Coco, P. Carcagnì, A. Schumann, B. Munjal, D.-H. Choi et al., "Ua-detrac 2018: Report of avss2018 & iwt4s challenge on advanced traffic monitoring," in 2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS). IEEE, 2018, pp. 1–6.
- [27] S. Li and F. Chen, "3d-detnet: a single stage video-based vehicle detector," in *Third International Workshop on Pattern Recognition*, vol. 10828. International Society for Optics and Photonics, 2018, p. 108280A.
- [28] L. Wang, Y. Lu, H. Wang, Y. Zheng, H. Ye, and X. Xue, "Evolving boxes for fast vehicle detection," in 2017 IEEE International Conference on Multimedia and Expo (ICME). IEEE, 2017, pp. 1135–1140.
- [29] H. Perreault, G.-A. Bilodeau, N. Saunier, and P. Gravel, "Road user detection in videos," *arXiv preprint arXiv:1903.12049*, 2019.

- [30] T. Wang, R. M. Anwer, H. Cholakkal, F. S. Khan, Y. Pang, and L. Shao, "Learning rich features at high-speed for singleshot object detection," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1971–1980.
- [31] T. Kong, F. Sun, A. Yao, H. Liu, M. Lu, and Y. Chen, "Ron: Reverse connection with objectness prior networks for object detection," in *IEEE Conference on Computer Vision* and Pattern Recognition, vol. 1, 2017, p. 2.
- [32] N. Goyette, P.-M. Jodoin, F. Porikli, J. Konrad, and P. Ishwar, "Changedetection. net: A new change detection benchmark dataset," in 2012 IEEE computer society conference on computer vision and pattern recognition workshops. IEEE, 2012, pp. 1–8.
- [33] P.-L. St-Charles, G.-A. Bilodeau, and R. Bergevin, "Subsense: A universal change detection method with local adaptive sensitivity," *IEEE Transactions on Image Processing*, vol. 24, no. 1, pp. 359–373, 2014.
- [34] R. H. Evangelio, M. Pätzold, and T. Sikora, "Splitting gaussians in mixture models," in 2012 IEEE Ninth international conference on advanced video and signal-based surveillance. IEEE, 2012, pp. 300–305.
- [35] Z. Zivkovic and F. Van Der Heijden, "Efficient adaptive density estimation per image pixel for the task of background subtraction," *Pattern recognition letters*, vol. 27, no. 7, pp. 773–780, 2006.
- [36] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," in *Proceedings*. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149), vol. 2. IEEE, 1999, pp. 246–252.