



Heriot-Watt University
Research Gateway

POL-LWIR Vehicle Detection: Convolutional Neural Networks Meet Polarised Infrared Sensors

Citation for published version:

Sheeny, M, Wallace, A, Emambakhsh, M, Wang, S & Connor, B 2018, POL-LWIR Vehicle Detection: Convolutional Neural Networks Meet Polarised Infrared Sensors. in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, pp. 1328-1334, 14th IEEE Workshop on Perception Beyond the Visible Spectrum in conjunction with CVPR 2018, Salt Lake City, Utah, United States, 18/06/18. <https://doi.org/10.1109/CVPRW.2018.00171>

Digital Object Identifier (DOI):

[10.1109/CVPRW.2018.00171](https://doi.org/10.1109/CVPRW.2018.00171)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Peer reviewed version

Published In:

2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)

Publisher Rights Statement:

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

POL-LWIR Vehicle Detection: Convolutional Neural Networks Meet Polarised Infrared Sensors

Marcel Sheeny¹, Andrew Wallace¹, Mehryar Emambakhsh², Sen Wang¹ and Barry Connor³

¹Heriot-Watt University, ²Cortexica Vision Systems, ³Thales UK

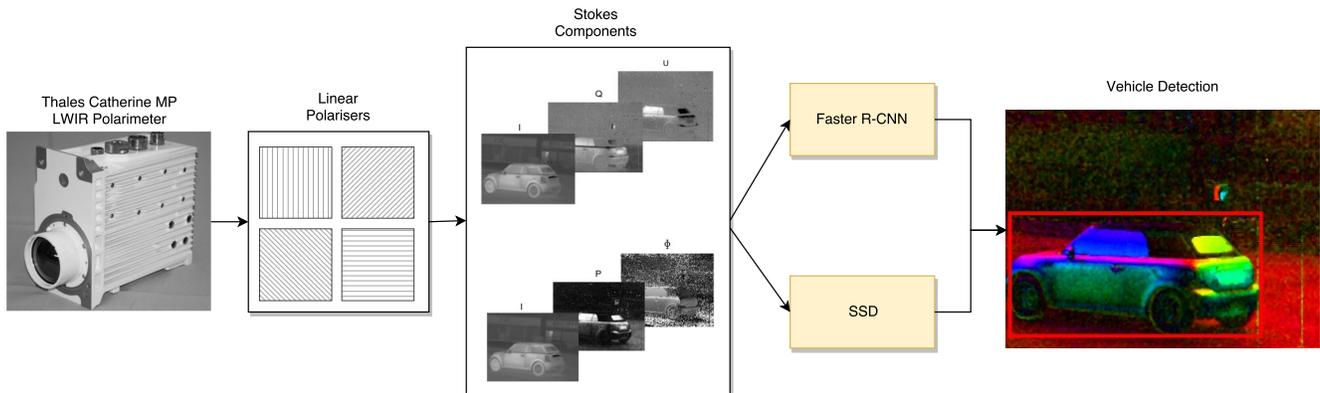


Figure 1: POL-LWIR Vehicle detection. This paper uses the Thales Catherine MP LWIR sensor, which is based on long wave polarised infrared technology. It contains 4 linear polarisers (0° , 45° , 90° , 135°). From the linear polarisers we can compute the Stokes components I, Q, U, P and ϕ . Two configurations are created (I, Q, U and I, P, ϕ), which are passed to 2 types of neural networks: (Faster R-CNN [24] and SSD [22]). The networks are trained to detect vehicles in both day and night conditions.

Abstract

For vehicle autonomy, driver assistance and situational awareness, it is necessary to operate at day and night, and in all weather conditions. In particular, long wave infrared (LWIR) sensors that receive predominantly emitted radiation have the capability to operate at night as well as during the day. In this work, we employ a polarised LWIR (POL-LWIR) camera to acquire data from a mobile vehicle, to compare and contrast four different convolutional neural network (CNN) configurations to detect other vehicles in video sequences. We evaluate two distinct and promising approaches, two-stage detection (Faster-RCNN) and one-stage detection (SSD), in four different configurations. We also employ two different image decompositions, the first based on the polarisation ellipse, the second on the Stokes parameters themselves. To evaluate our approach, the experimental trials were quantified by mean average precision (mAP) and processing time, showing a clear trade-off be-

tween the two factors. For example, the best mAP result of 80.94 % was achieved using Faster-RCNN, but at a frame rate of 6.4 fps. In contrast, MobileNet SSD achieved only 64.51 % mAP, but at 53.4 fps.

1. Introduction

The future of autonomous cars is still uncertain, but impressive new results are being achieved with most car manufacturers promising level 4 autonomy by 2020. A necessary capability for autonomy is sensory perception, but the vast majority of research is based on publicly available video benchmarks like KITTI [9] and CityScapes [6]. These datasets are acquired during daytime, in good weather conditions, using video cameras. For full autonomy and situational awareness in all weather conditions, sensors and perceptual algorithms workable continuously in 24 hours capability are required. Infrared sensors are capable of sensing beyond the visible spectrum and are robust to falling illumi-

nation.

The majority of previous works, e.g. [19] [16], have used IR sensors in a military or surveillance context to detect *hot* objects, especially during night. From the perspective of a commercial road vehicle, IR sensors have also been used at night to detect other actors such as vehicles and pedestrians [17]. However, it has long been recognised that additional analysis of polarisation state, governed by the material refractive index, surface orientation and angle of observation, can lead to better discrimination. False colour representations of the polarimetric data can visually reveal hidden targets [31] and to some extent, 3D structure [18]. By exploiting knowledge about the dependence of polarisation state on surface and viewing angle, and the fixed vehicle to road geometry, Connor *et al.* [5] constructed a road segmentation algorithm. Bartlett *et al.* [4] used polarisation to improve anomaly detection. Using a nearest neighbour detection based on Euclidean distance, foreground vs background clustering is performed. The points whose distance from background components are higher than a threshold are labeled anomalous. Romano *et al.* [26] also used polarisation for anomaly detection. A data cube was formed from the primary (i_0, i_{90}) imagery and the sample covariance within a local (sliding) window is compared to the sample covariance of the entire image. They found this method is better at discriminating between target and background pixels at different times of the day than using Stokes components.

With its dramatic increase in popularity, there have been a number of recent studies based on deep neural networks to recognise objects in IR images. For example, Rodger *et al.* [25] used CNNs to classify pedestrians, vehicles, helicopters, airplanes and drones using a LWIR sensor. Abbott *et al.* [2] used the YOLO method [23] and transfer learning from a high resolution IR to a lower resolution IR (LWIR) dataset to detect vehicles and pedestrians. Lie *et al.* [21] used the KAIST dataset [15] to fuse RGB and thermal information in a Faster R-CNN architecture, again to detect pedestrians. Gundogu *et al.* [12] combined a CNN detection stage with a long term correlation tracker to detect tanks in cluttered backgrounds. However, all of these studies only used intensity data.

Dickson *et al.* [7] exploited a polarised LWIR sensor to detect vehicles in both still images and to a lesser extent, video sequences. In rural settings, LWIR emissions from man-made objects appear more strongly polarised than vegetation. However in urban settings, most of the environment is man-made. Therefore, in their work on vehicle recognition [7], the key observation was that although there are many other man-made structures in urban scenes, there is a distinct, differential spatial arrangement of surface signatures in LWIR polarimetric images of vehicles due to their regular structure and size. This leads to a regular pattern of

pixel clusters in a 2D space encoding the degree and angle of polarisation.

In this paper, our contribution is to evaluate the effectiveness of CNNs to detect vehicles in polarised LWIR data. We evaluate the two main research directions in deep learning for object detection: two-stage detection, which first proposes the bounding boxes then performs classification in each bounding box based on Faster R-CNN [24], and one-stage detection which detects and classifies in a single network, based on Single-Shot Multibox Detection (SSD) [22]. To the best of our knowledge, this is the first paper to exploit the use of polarised infrared together with neural networks for object detection.

2. Methodology

2.1. Sensing and the Polarisation Parameters

Figure 1 illustrates a schematic diagram of our approach. We use the Thales Catherine MP LWIR Polarimeter [7], operating in the range of $8\mu m$ to $12\mu m$ to record video images. Each pixel of a 320×256 image frame has 2×2 sensing sites that contains linear polarisers oriented at 0, 45, 90 and 135 degrees. The data capture rate is 100 frames per second (fps). Our dataset was collected in Glasgow, UK on 14th and 15th of March, 2013. Seven sequences are recorded and the bounding boxes of the vehicles are annotated, from which we use 4 sequences for training and 3 for testing. In total, we have 10659 annotated frames for training and 4453 annotated frames for testing.

The polarisation state of the emitted LWIR radiation can be expressed in terms of the Stokes vector, $\{I, Q, U, V\}$ [3]: The I component measures the total intensity; the Q and U components describe the radiation polarised in the horizontal direction and in a plane rotated 45° from the horizontal direction, respectively; the V component describes the amount of right-circularly polarised radiation. To measure the V component we require an additional quarter waveplate. As a result we can only measure I, Q and U . Therefore, with respect to the measured intensities at each pixel site, we deduce that

$$I = \frac{1}{2}(i_0 + i_{45} + i_{90} + i_{135}) \quad (1)$$

$$Q = i_0 - i_{90} \quad (2)$$

$$U = i_{45} - i_{135} \quad (3)$$

The degree of linear polarisation, P , and the angle of polarisation, ϕ , can also be calculated as follows,

$$P = \frac{\sqrt{Q^2 + U^2}}{I} \quad (4)$$

$$\phi = \frac{1}{2} \tan^{-1} \left(\frac{U}{Q} \right) \quad (5)$$

2.2. Faster-RCNN

The Faster R-CNN method [24] relies on a two-stage object detection procedure. First, a sub-network is used to propose the bounding boxes; Second, a separate sub-network is used to classify objects within each bounding box. The idea of Faster R-CNN evolved from R-CNN [11], which proposes several bounding boxes based on the selective search algorithm [30]. Selective search applies a segmentation algorithm in many stages to under and over segment the image. Bounding boxes are proposed at each region segment; these are input to a CNN to classify the type of object. This CNN can be chosen from popular successful architectures, such as VGG [28], InceptionNet [29] or ResNet [13].

Since selective search usually outputs large number of regions (~ 2000 regions), it is computationally expensive. Fast R-CNN [10] reduces this complexity by running a CNN over the whole image. Proposed regions are transformed to the last feature map before the fully connected layers and the regions in the feature map are classified in a simple neural network, resulting in a significant computational complexity reduction. Despite such improvement, the authors realised that the selective search is indeed a bottleneck preventing faster execution of the overall algorithm. As a solution, Faster R-CNN [24] created a network to learn how to generate bounding boxes. This pipeline is also known as region proposal network (RPN). The RPN creates a grid in the original image which anchors the bounding box annotations to the map. Using the previous annotations of the bounding box in the original image, the RPN learns how to propose bounding boxes. RPN can also reduce the number of proposals compared to selective search. Replacing selective search with the RPN improved both speed and accuracy (using the PASCAL VOC 2012 dataset).

Since it is easy to plug any CNN into the Faster R-CNN method, we decided to use ResNet-50 and ResNet-101 [13], where the number attached to the name of the network relates the number of layers used. ResNet uses residual layers that are CNNs with shortcut connections. Those connections skip the current layer and the skipped output is added to the output after the convolution is applied. ResNet has a trade-off between accuracy and depth of the network: the smaller the network, the faster it performs.

2.3. Single Shot Multi-box Detection

Single Shot Multi-box Detection (SSD) [22] uses one-stage detection, in which the output of a single network is

a set of bounding boxes with the respective classes. This is different from Faster R-CNN which has two stages, the region proposal and the classification stages.

The use of one-stage detection attracted the attention of many researchers in the field. Sermane *et al.* [27] used a CNN over the whole image, where each cell of the last feature map corresponds to a region in the original image. The regions are always uniformly distributed over the image, which constitutes a major disadvantage as there are no *a priori* reason why this should be the case. The YOLO network [23] created a CNN which is trained using anchor boxes from the annotation (similar to RPN from Faster R-CNN) to output the region location plus its classification. SSD is similar to YOLO in the sense that it uses a CNN to output the region's location and its classification result. However, SSD generates the output at several scales of the feature map produced by the convolution. The output maps are based on a grid of anchor boxes, as with Faster R-CNN. All the results at several scales are then combined, followed by a non-maximum suppression step to remove multiple detections of the same object.

The architecture of SSD is based on convolutional stages of other networks, such as VGG [28] and InceptionNet [29] and MobileNet [14]. InceptionNet (GoogleNet) was the winner of the ImageNet 2014 competition for image classification. This network applies convolutions of different sizes (1×1 , 3×3 , 5×5) in the same layer and concatenates them into a single feature map. These convolutions are called "Inception modules", and when done several times create a deep network of inception modules. This showed that applying several convolutions to the same feature map can retrieve more robust features. MobileNet [14] was designed to be a fast and small CNN to run on low powered devices. It is a convolutional layer approximator; Instead of applying $N \times M \times C$ convolutions, it first applies a $1 \times 1 \times C$ convolution to reduce the size of the input to $W \times H \times 1$, then applies a $N \times M \times 1$ convolution (where N and M are the convolution mask sizes, W and H are width and height of the image and C is the number of channels). This strategy reduces the number of weights to be learned by the network and the complexity of the convolution. SSD with MobileNet is used in our experiments to evaluate how a small network can learn the polarised infrared features.

2.4. Experiments, Training and Evaluation

We have evaluated two configurations of the measured polarised image data to train and test our several CNN architectures. The first configuration uses the I, Q, U parameters as the input image planes. In Fig. 2a, we can visualise the image of each component I, Q, U . Again, in Fig. 2b we can visualise the image of each component I, P, ϕ . We use four configurations of neural network for our experiments.

1. The SSD network using the InceptionV2 network [29]

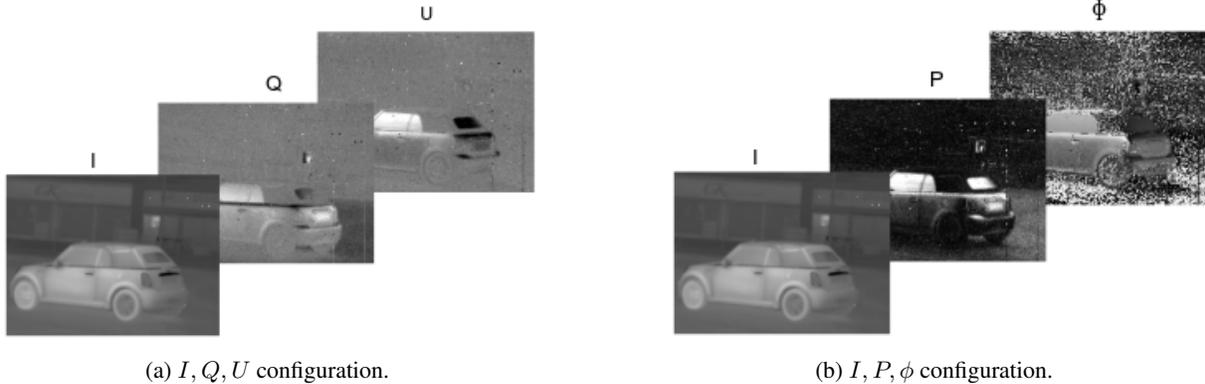


Figure 2: Visualisation of each configuration based on the Stokes components.

to extract features.

2. The SSD network using the MobileNet network [14] to extract features.
3. Faster R-CNN using ResNet-50 [13] to extract features.
4. Faster R-CNN using ResNet-101 [13] to extract features.

We trained our 4 different configurations on both (I, Q, U) and (I, P, ϕ) data, and for comparison with previous work that has applied CNNs to intensity data alone, on the I data in isolation. The networks are trained using a i7-7700HQ, 32 GB ram, NVIDIA Titan X and developed using the Tensorflow Object Detection API [1]. The network weights for both SSD and Faster R-CNN are initialised from the MS-COCO object detection dataset [20]. The parameters for the Faster R-CNN networks are: batch size 1, learning rate 0.0003, momentum 0.9. The parameters for the SSD networks were: batch size 24, learning rate 0.004, momentum 0.9. Eq. 6 shows the gradient descent formula.

$$W_{t+1} = W_t - \alpha \nabla f(x; W) + \eta \Delta W \quad (6)$$

where η is the momentum, α is the learning rate, t is the time current time step, W is all weights of the network and $\nabla f(x; W)$ is the derivative of the function that represents the network and x is our dataset.

The evaluation metrics used are mean average precision (mAP) and processing time in frames per second (fps). The mAP classifies correct detection when Intersection over Union (IoU), > 0.5 , which follows the PASCAL VOC protocol [8]. (The KITTI protocol [9] uses $IoU > 0.7$ for vehicles. However, since unlike KITTI, our annotations are not pixel-level, we followed the PASCAL criteria.) To compute the fps we compute the average time over 100 frames. Tables 1 and 2 show the results for each configuration.

Table 1: Results for each configuration

	mAP [I]	mAP [I,Q,U]	mAP [I,P,phi]
MobileNet SSD	48.50 %	64.51 %	58.56 %
InceptionV2 SSD	59.79 %	72.17 %	73.24 %
Faster R-CNN Resnet-50	75.63 %	72.82 %	76.43 %
Faster R-CNN Resnet-101	75.21 %	73.67 %	80.94 %

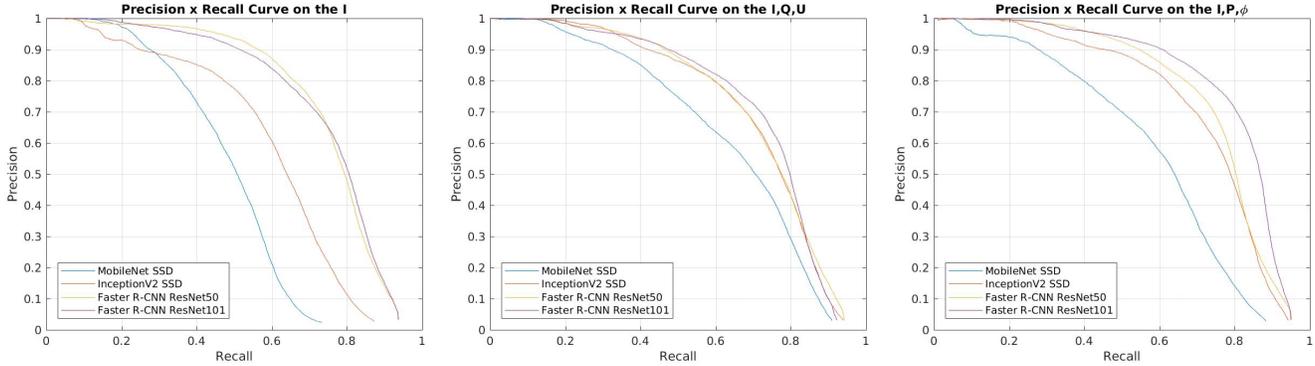
Table 2: Computational speed (fps) for each network configuration.

	fps
MobileNet SSD	53.4
InceptionV2 SSD	37.2
Faster R-CNN Resnet-50	7.8
Faster R-CNN Resnet-101	6.4

Precision-recall curves are also generated for each results and can be visualised in Fig. 3a for I alone, Fig. 3b for I, Q, U and Fig. 3c for I, P, ϕ .

Qualitative examples of images can be seen in Fig. 4. This uses a pseudo-colour display, converting P and ϕ to HSV colour space. Based on qualitative results we can see that the main problem of the SSD lies with small objects. The network needs to learn the location and features at the same time, which affects the detection of small objects.

Considering the results of Table 1 and the precision-recall curves, a key question is whether the use of polarised data improves detection when compared to the intensity data alone, as used by most previous authors. We believe that it is indeed the case, particularly for the SSD examples, although we qualify that statement by noting that the difference in the R-CNN results is less definitive, and that the dataset is limited and many more trials are needed for full statistical verification. In general, for all datasets, Faster R-CNN performs better for this limited trial, as measured by mAP, although this comes with the penalty of a much



(a) Precision-Recall curve for the I configuration. (b) Precision-Recall curve for the I, Q, U configuration. (c) Precision-Recall curve for the I, P, ϕ configuration.

Figure 3: Precision-Recall curves

slower frame rate. However, this latter result is consistent with the published results on video sequences, where splitting the tasks of region proposal and classification arguably makes the network more robust in learning directly the object features. In contrast, SSD needs to learn the localisation together with classification, and hence both location and object characteristics influence the weights of the network which can degrade performance. Nevertheless, in our trial, SSD-InceptionV2 achieves similar performance to Faster R-CNN ResNet-50, at a much increased speed, since it just needs one CNN for both region proposal and classification.

Comparing I, Q, U and I, P, ϕ , the best result is obtained with the I, P, ϕ parametrisation. Although the differences are not proven as significant, such that much more extensive characterisation is required. At this stage, given the complexity of these neural networks, it is hard to define what type of feature is being learned in each case, although from Dickson *et al.* [7] the authors claim that material, shape and surface and viewing angles influence the underlying polarisation patterns. As a specific example, one can see that the I, P, ϕ configuration does detect an occluded car that the I, Q, U does not in the Faster R-CNN example. However, although this occurs more often than the converse, much better understanding of the network and more extensive trials are necessary to draw reliable conclusions. For a necessary perception by an autonomous car, computational time is clearly quite a crucial factor. As expected, the one-stage SSD architectures shows higher frame rates. MobileNet SSD is the fastest and can process on average at 53.4 fps, but it has the lowest mAP.

3. Conclusions

In this paper, we evaluated and compared a series of different CNN architectures for vehicle detection in polarised

long wave infrared image sequences, using two different image decompositions. We showed that the use of polarised infrared data was effective for vehicle detection, and appeared to perform better when CNNs are used for detection in infrared intensity data alone, confirmed also by previous researches. Faster R-CNN based networks achieved better results in terms of detection accuracy, splitting the tasks of region proposal and classification to make the network more robust. However, it should be mentioned that improving the accuracy of one-stage detection network is quite an active field of research, providing much higher frame rates. We could reach no firm conclusion on which image decomposition was preferable, although anecdotally the $\{I, P, \phi\}$ parametrisation is both more intuitive in describing the polarisation ellipse and achieved the best overall result with Faster R-CNN ResNet 101. Our detection rates are similar to those of KITTI dataset for vehicle detection from simple video data in daylight using the same networks. Overall, our work shows that polarising LWIR data is a relatively robust option for day and night operation.

Acknowledgements

We acknowledge the support of the Engineering and Physical Research Council, grant references EP/N012402/1 (TASCC: Pervasive low-TeraHz and Video Sensing for Car Autonomy and Driver Assistance (PATH CAD)) and EP/G037523/1 (Optics and Photonics Technologies) conducted with and supported by Thales UK.

Thanks to NVIDIA for the donation of the TITAN X GPU.

References

- [1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. Tensor-

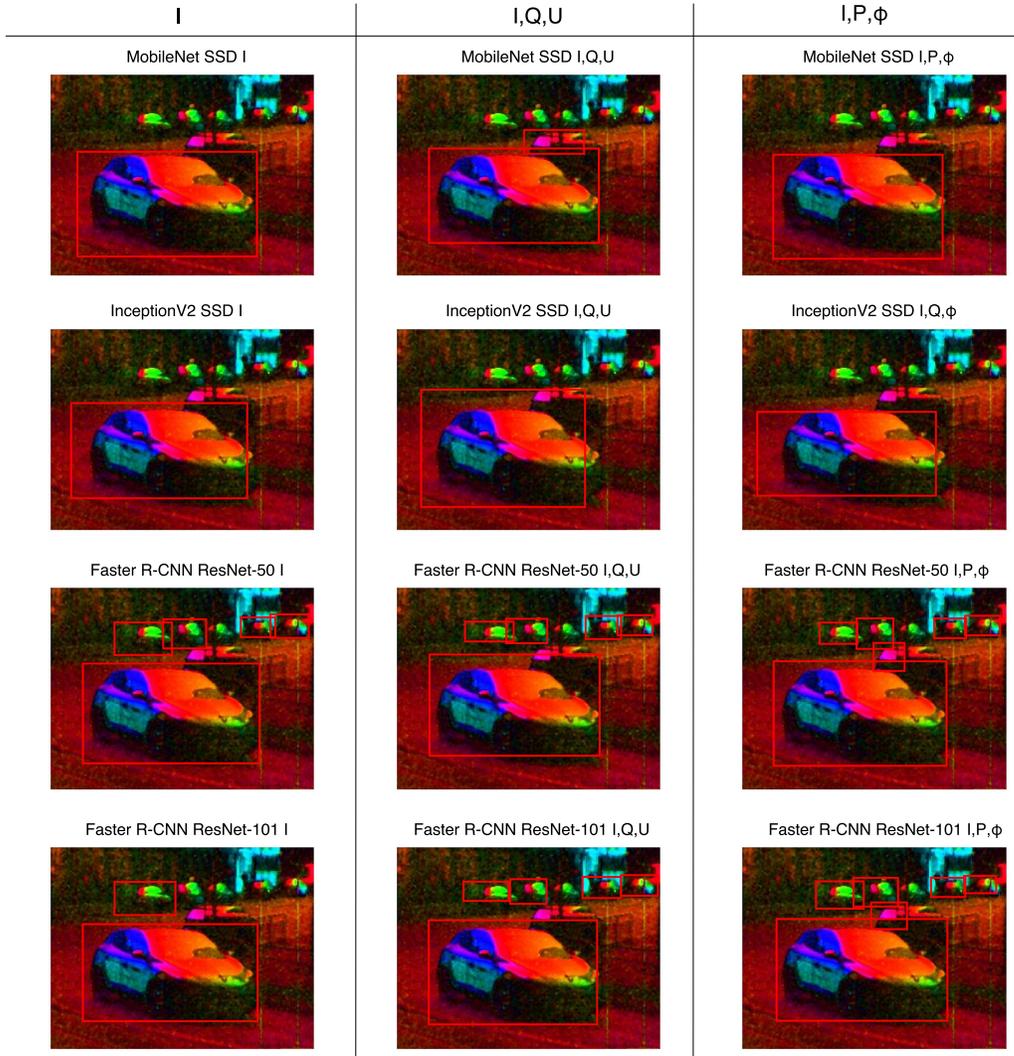


Figure 4: Qualitative results for each configuration with threshold at 0.7.

flow: A system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016. 4

- [2] R. Abbott, J. M. Del Rincon, B. Connor, and N. Robertson. Deep object classification in low resolution lwir imagery via transfer learning. In *Proceedings of 5th IMA Conference on Mathematics in Defence.*, 2017. 2
- [3] R. Azzam and N. Bashara. *Ellipsometry and polarized light*. North-Holland personal library. North-Holland Pub. Co., 1977. 2
- [4] B. D. Bartlett, A. Schlamm, C. Salvaggio, and D. W. Messinger. Anomaly detection of man-made objects using spectropolarimetric imagery. In *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XVII*, volume 8048, page 80480B. International Society for Optics and Photonics, 2011. 2
- [5] B. Connor, J. Letham, N. Robertson, and I. Carrie. Scene

understanding and task optimisation using multimodal imaging sensors and context: a real-time implementation. In *Infrared Technology and Applications XXXVII*, volume 8012, page 80120A. International Society for Optics and Photonics, 2011. 2

- [6] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016. 1
- [7] C. N. Dickson, A. M. Wallace, M. Kitchin, and B. Connor. Long-wave infrared polarimetric cluster-based vehicle detection. *JOSA A*, 32(12):2307–2315, 2015. 2, 5
- [8] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of com-*

- puter vision*, 111(1):98–136, 2015. 4
- [9] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 1, 4
- [10] R. Girshick. Fast r-cnn. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 1440–1448. IEEE, 2015. 3
- [11] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 3
- [12] E. Gundogdu, A. Koc, and A. AA. Automatic target recognition and detection in infrared imagery under cluttered background. In *Proceedings of SPIE Conferenc on Security and Defence*, volume 10432, Target and Background Signatures, 2017. 2
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 4
- [14] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 3, 4
- [15] S. Hwang, J. Park, N. Kim, Y. Choi, and I. S. Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*, pages 1037–1045. IEEE, 2015. 2
- [16] M. Khan, G. Fan, D. Heisterkamp, and L. Yu. Automatic target recognition in infrared imagery using dense hog features and relevance grouping of vocabulary. In *Proceedings of CVPR Workshop*, pages 293–298, 2014. 2
- [17] J.-Y. Kwak, B. Ko, and J. Nam. Pedestrian tracking using online boosted random ferns learning in far-infrared imagery for safe driving at night. *IEEE Transactions on Intelligent Transportation Systems*, 18(1):69–81, 2017. 2
- [18] D. A. Lavigne, M. Breton, G. Fournier, J.-F. Charette, M. Pichette, V. Rivet, and A.-P. Bernier. Target discrimination of man-made objects using passive polarimetric signatures acquired in the visible and infrared spectral bands. In *Polarization Science and Remote Sensing V*, volume 8160, page 816007. International Society for Optics and Photonics, 2011. 2
- [19] G. Lin, G. Fan, L. Yu, X. Kang, and E. Zhang. Heterogeneous structure fusion for target recognition in infrared imagery. In *Proceeding of CVPR Workshop*, volume 8160, pages 118–125, 2015. 2
- [20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 4
- [21] J. Liu, S. Zhang, S. Wang, and D. N. Metaxas. Multispectral deep neural networks for pedestrian detection. *arXiv preprint arXiv:1611.02644*, 2016. 2
- [22] 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*, pages 21–37. Springer, 2016. 1, 2, 3
- [23] 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*, pages 779–788, 2016. 2, 3
- [24] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2017. 1, 2, 3
- [25] I. Rodger, B. Connor, and N. Robertson. Classifying objects in lwir imagery via cnns. In *Proc. SPIE: Electro-Optical and Infrared Systems: Technology and Applications XIII*, 2016. 2
- [26] J. M. Romano, D. Rosario, and J. McCarthy. Day/night polarimetric anomaly detection using spice imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 50(12):5014–5023, 2012. 2
- [27] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. Lecun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *International Conference on Learning Representations (ICLR2014), CBLS, April 2014*, 2014. 3
- [28] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 3
- [29] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [30] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013. 3
- [31] G. Wong, C. McMaster, R. Struthers, A. Gorman, P. Sinclair, R. Lamb, and A. R. Harvey. A novel snapshot polarimetric imager. In *Electro-Optical and Infrared Systems: Technology and Applications IX*, volume 8541, page 85410G. International Society for Optics and Photonics, 2012. 2