

Optimized Vision-Directed Deployment of UAVs for Rapid Traffic Monitoring

Christos Kyrkou, *Member, IEEE*, Stelios Timotheou, *Senior Member, IEEE*, Panayiotis Kolios, *Member, IEEE*, Theocharis Theocharides, *Senior Member, IEEE*, and Christos G. Panayiotou, *Senior Member, IEEE*

KIOS Research and Innovation Center of Excellence
Department of Electrical and Computer Engineering
University of Cyprus, Nicosia, Cyprus

{kyrkou.christos, timotheou.stelios, pkolios, ttheocharides, christosp}@ucy.ac.cy

Abstract—The flexibility and cost efficiency of traffic monitoring using Unmanned Aerial Vehicles (UAVs) has made such a proposition an attractive topic of research. To date, the main focus was placed on the types of sensors used to capture the data, and the alternative data processing options to achieve good monitoring performance.

In this work we move a step further, and explore the deployment strategies that can be realized for rapid traffic monitoring over particular regions of the transportation network by considering a monitoring scheme that captures data from a visual sensor on-board the UAV, and subsequently analyzes it through a specific vision processing pipeline to extract network state information. These innovative deployment strategies can be used in real-time to assess traffic conditions, while for longer periods, to validate the underlying mobility models that characterise traffic patterns.

I. INTRODUCTION

UAVs are considered as a promising technology for both environmental and infrastructure monitoring with broad use in a plethora of applications [1]. This is due to the fact that UAVs have become both affordable and highly capable platforms mainly due to the Moores-law-style pace of technological advancements governing the electronic components that make up much of the aircraft and its payload.

Traffic monitoring in transportation networks is one such application for which the use of UAVs provides clear benefits both in terms of flexibility in deploying the sensors where there is a need (in an easy and fast way), as well as in terms of cost efficiency since on-board sensors can capture a significantly larger area compared to ground-mounted alternatives with confined view.

Research on the topic has focused mainly on addressing the challenges of remotely capturing information above the ground using a variety of sensors including visual and infra-red cameras, and radar, and processing that data to detect and track vehicles [2] - [8]. While a single UAV with fixed trajectory is considered in most of the literature, solutions with multiple UAVs with adaptive trajectories have also been studied to improve detection and tracking performance [9], [10].

The proposed solution extends previous work by incorporating data acquisition and processing from UAVs into a holistic framework that considers not just the UAV capabilities but also the limitations posed by the underlying road traffic network to

develop optimization strategies for the deployment of UAVs for traffic monitoring. Specifically the contributions of this paper are twofold. First, we introduce a vision processing pipeline that can provide traffic monitoring information and is suitable for deployment on UAV platforms. Second, we consider where to place a number of UAVs above particular locations of the transportation network to effectively capture the traffic patterns of the region of interest. To simplify the problem, no mobility is considered in this work. As shown in the sequel, the proposed UAV deployment strategies provide a valuable tool to transport officers that would like to assess traffic conditions in a particular area in real time or capture footage that will help assess the assumed mobility models used for a particular region of the network.

To do so, an arbitrary transportation network topology is used together with data associated with the constituent areas. The maximum allowed height of buildings in certain areas is used as the main parameter affecting the line-of-sight between the roadways and the drone hovering locations. The aim is then to find the particular locations for UAV deployment that will allow for the desired traffic monitoring level across the region of interest. The problem is mathematically formulated as an optimization program that finds the least number of UAV hovering locations considering the data acquisition and processing constraints jointly with the Field-of-View (FoV) obstructions caused by building heights at the different areas. The solution to this problem provides a deployment map that transport officers can use to capture footage of road traffic that will help assess congestion levels and mobility patterns.

The rest of the paper is organized as follows. Section II first reviews the state-of-the-art practises for capturing and processing road traffic footage. It then describes an on-board solution that can be used to efficiently analyse the captured footage for vehicle detection. As shown later, the proposed solution has been developed and tested using commercial hardware under real road traffic conditions with good performance results. Section III then goes on to model the underlying transportation network based on the FoV constraints and then provides a formulation for solving this problem using Integer Linear Programming. Finally, Section IV provides experimental results for a case study conducted with real data in the capital city of Cyprus, Nicosia where deployment strategies

are computed for a number of road network levels (i.e., considering only the primary road, the primary and secondary road, and the complete network including residential roads). Section V concludes this paper with the key contributions and findings of the proposed work.

II. COMPUTER VISION PIPELINE FOR DATA CAPTURING AND PROCESSING

Vehicle detection has been extensively studied in the literature while the latest state-of-the-art techniques use deep convolutional neural networks (CNN) [12]. A more recent example is shown in [11] where the authors utilize a deep learning framework that performs scene analysis of aerial images to first segment the image into various regions, and then extract the regions that correspond to vehicles and classify them into subcategories. As the pipeline uses deep neural networks the segmentation and classification tasks run off-line on an NVIDIA Tesla GPU, and not on the UAV. On the contrary in this work we target on-board processing on a UAV platform which may not have such high-end hardware. Hence, in this work we propose a lightweight pipeline capable of running even on low-end devices such as the Raspberry Pi. The pipeline first identifies road regions and then uses a combination of motion detection and image classifier to detect and count the vehicles over time to estimate the traffic density. For low-end platforms the Haar cascade detector [13] is used, while for more capable platforms it can be substituted with a convolutional neural network [17]. We provide results for both classifiers.

The overall computer vision pipeline is shown in Fig. 1. There are four main stages in the pipeline, the processing to extract the road mask, the detection of moving objects, the classification of vehicles, and the final post-processing for vehicle count and accumulation over time.

A. Extraction of Road Segments

The extraction of road segments is important as we target road urban road traffic monitoring and this optimization will the classification to focus on promising areas which contain vehicles of interest, thus maximizing both the performance as well as the accuracy. The extraction happens in HSV space by thresholding the input image for different value ranges. In addition, in order to guarantee the smoothness of the road mask and that no road pixels will be discarded either due to illumination changes or the presence of vehicles on the road, a running average algorithm is employed, due to its low computational complexity, to isolate the background and use that to find the road segments.

B. Motion Detection

The detection of moving objects on the extracted road segment acts as an efficient way of identifying the vehicles from the road segment. The motion detection relies on identifying differences between successive frames. The difference image is then analyzed to identify the different contours that correspond to moving vehicles. Some false detections can be

discarded based on size and aspect ratio. In cases where the vehicles may stop due to traffic or at junctions, the motion detection may fail to identify the vehicles. For this reason, we also include in the pipeline a classifier that analyzes the image to detect vehicles through machine learning. By incorporating both approaches we are able to compensate for failures of one or the other approach thus resulting in a more robust detection process.

C. Vehicle Classification

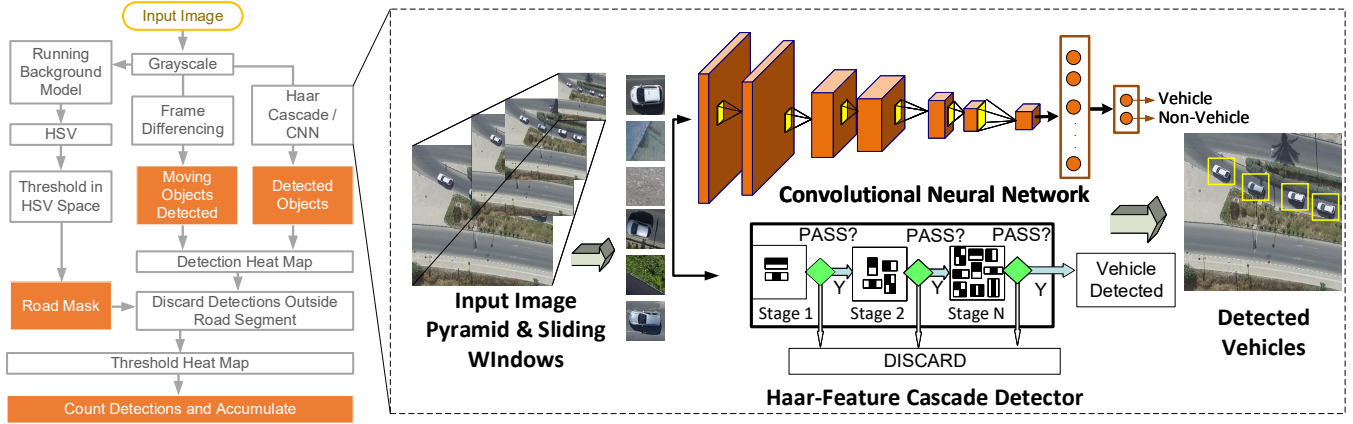
The vehicle classification algorithm can be based either on a Haar Cascade or Convolutional Neural Network. The latter is faster and does not require a specialized GPU so it can be executed on on-board processing platform of the UAV, while the former can offer higher detection accuracy and can be used in cases where the image is offloaded to a dedicated computing infrastructure or when the on-board hardware is powerful enough to handle it. Following we provide more details for the two classification approaches. The architecture of the two classifiers are shown in Fig. 1-(iii).

1) *Haar Feature-based Cascade Classifier for Object Detection*: The Haar cascade classification approach would be particularly attractive for on-board processing of the camera images as it does not require specialized hardware and is capable of operating in real-time. The approach relies training boosted classifiers of increasing complexity using Haar-like features and arranging them in a cascade structure in order to rapidly discard background regions. The cascade is trained to distinguish between car and non-car images of 10×10 pixels, with a total of 2,884 positive and 83,000 negative samples, resulting in a 20 stage cascade.

2) *Convolutional Neural Network*: Deep learning approaches have demonstrated state-of-the-art performance in various computer vision tasks such as object detection and recognition [17]. However, they often require a dedicated GPU to run which makes it difficult to deploy for on-board processing in UAVs. Nevertheless, it should be the go-to approach when the visual processing can be off-loaded from the UAV. As such, we also develop a CNN architecture to detect vehicles and offer this as a possible alternative to the Haar Cascade. The network architecture follows the VGG style of 3×3 convolutional layers followed by max-pooling, and has a total of 960,000 parameters which are much less than pre-trained models such as VGG [11]. We train the network on a much larger dataset than the cascade classifier with image augmentation for a total of 95,000 positive and 245,000 negative examples.

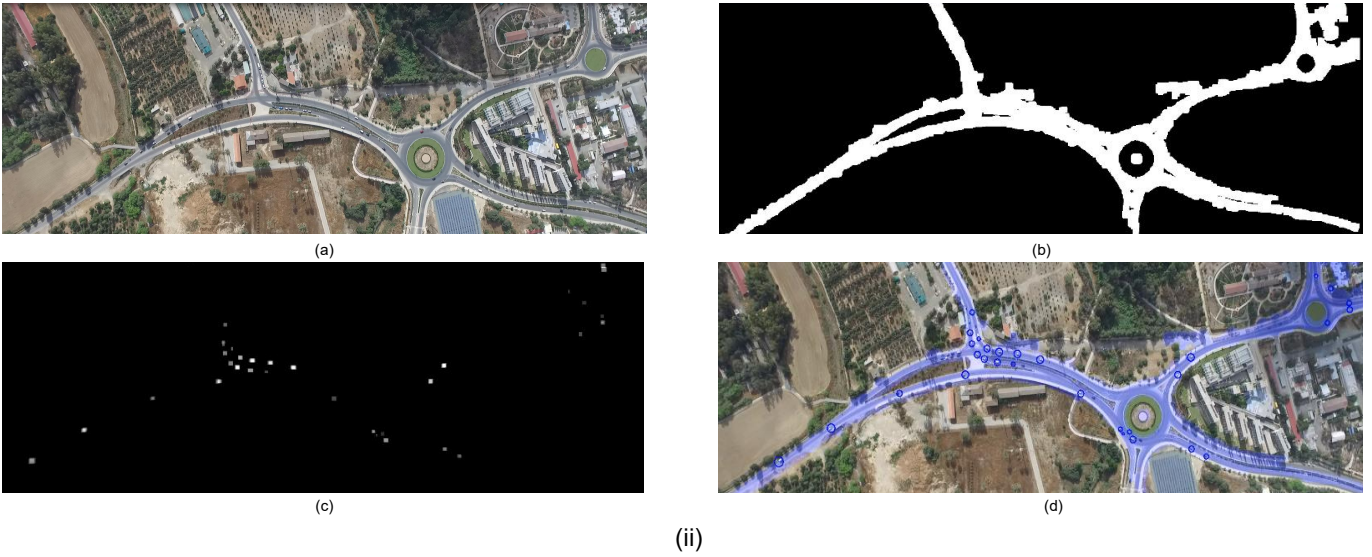
D. Result Aggregation

The detected moving vehicles and the classifier detected vehicles are aggregate to produce a detection map of vehicle locations. Areas where multiple detections overlap indicate a high confidence from both the motion detection and the classifier. Hence, by appropriately thresholding the detection map and by keeping only highly detected regions it is possible to remove any false positives. The preserved regions are



(i)

(iii)



(ii)

Fig. 1. (i) Vision Pipeline Main Stages: a) Input Image b) Road Mask c) Moving Vehicles d) Road Segmentation and Vehicle Detection. (ii) Computer Vision Pipeline for vehicle density estimation. (iii) Classification Process using either a haar-cascade classifier or a convolutional neural network.

counted and accumulated to produce a density estimation of the region for a given time period.

As shown in Fig. 1 the whole process has been developed and tested using actual road footage captured by a commercial drone providing results that demonstrate the applicability of such an approach under real settings. The computer vision pipeline has been evaluated under different conditions and height levels as shown in Section IV-A. Such an analysis can provide useful information that can be used to construct a feasible deployment strategy depending for the UAVs in order to ensure maximum coverage and detection performance.

III. MODELLING AND PROBLEM FORMULATION

As introduced in Section I, the proposed UAV deployment strategy is based on finding appropriate UAV hovering locations in order monitor a certain part of the road traffic network. Towards this direction, space is discretized so that UAVs need to be able to monitor a finite set of points that ensure coverage

of the road infrastructure of interest, hereafter termed Points of Interest (PoIs). Discretization of space also results in a set of candidate UAV hovering locations, so that by deploying the UAVs at some of the locations all PoIs will be in line-of-sight.

To construct the set of UAV hovering locations, we exploit the fact that certain subareas of the monitored area usually have a maximum building height h_M , such that no building structure in the specific subarea can be taller. This ensures line-of-sight for all vehicles located in a circular region of radius R around a UAV, provided that the particular region falls within the field-of-view of the UAV camera. Knowing R_l for each subarea l , means that we can construct a non-uniform grid of UAV hovering locations by considering points on a square lattice with horizontal/vertical distance R_l/ρ , where ρ is a granularity constant. The higher the value of ρ , the denser is the constructed PoI grid and hence the area that may be left uncovered is smaller.




Fig. 2. Geometric representation of the line-of-sight between a UAV and a vehicle inside a transportation network.

To determine R , let us consider Fig. 2 depicting a UAV aiming to detect a specific vehicle. In the figure, let w_{UB} and w_{BV} denote the horizontal distance between the UAV and the building possibly obstructing the line-of-sight with the vehicle, and the distance between that building and the vehicle, respectively. Let also h_U , denote the height of the UAV above the ground level, and h_B and h_V the height of the building and the vehicle, respectively. In this case, line-of-sight is maintained between the UAV and the vehicle when $\theta \geq \phi$, where

$$\tan(\theta) = \frac{w_{UB} + w_{BV}}{h_U - h_V}, \quad (1)$$

$$\tan(\phi) = \frac{w_{BV}}{h_B - h_V}. \quad (2)$$

Considering the maximum building height in a given subarea, h_M , the maximum range at which the UAV can detect vehicles (maintains line-of-sight) is achieved when $\theta = \phi$ such that

$$R = \frac{h_U - h_V}{h_M - h_V} w_{BV}. \quad (3)$$

Let \mathcal{N} and \mathcal{M} denote the sets of UAV candidate hovering locations and PoIs of the road network, respectively. Then, a UAV positioned at $i \in \mathcal{N}$ can monitor vehicles located at PoI $j \in \mathcal{M}$ if $d_{ij} \leq R_j$, where $d_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$, R_j is the maximum detection range of PoI j from a UAV, while (x_i, y_i) and (x_j, y_j) are the coordinates of points i and j , respectively.

In this way, we can form matrix $\mathbf{C} \in \{0, 1\}^{|\mathcal{N}| \times |\mathcal{M}|}$ with elements C_{ij} denoting whether PoI $j \in \mathcal{M}$ can be monitored from a UAV located at point $i \in \mathcal{N}$ ($C_{ij} = 1$) or not ($C_{ij} = 0$). Let also decision variables x_i , $i \in \mathcal{N}$ denote whether a UAV is deployed at i ($x_i = 1$) or not ($x_i = 0$). Then, the deployment

problem can be addressed using the following formulation

$$\min_{\mathbf{x}} \sum_{i \in \mathcal{N}} x_i \quad (4a)$$

$$\text{s.t.} \sum_{i \in \mathcal{N}} C_{ij} x_i \geq 1, \quad j \in \mathcal{M}, \quad (4b)$$

$$x_i \in \{0, 1\}, \quad i \in \mathcal{N}. \quad (4c)$$

Formulation (4) aims to minimize the number of UAVs that are needed to cover all PoIs. This formulation is known in the literature as the *set covering problem* which is known to be NP-hard [15]. Nonetheless, as the problem is a Mixed Integer Linear Program (MILP), we can employ standard MILP solvers for its solution, such as the Gurobi Solver [16].

IV. EVALUATION RESULTS

A. Vehicle Detection Performance

The computer vision pipeline presented in Section II was evaluated using footage captured by a DJI Matrice 100 UAV. Different scenes were captured corresponding to different traffic conditions as well as UAV hovering heights (as shown in Fig. 3) in order to test the sensitivity of the proposed vehicle detection pipeline. Overall, the proposed approach was able to reliably estimate the number of vehicles within an accuracy between 86 – 90% which indicates that it can be reliably used to estimate the traffic conditions in an area. The maximum height which the UAV was able to reliably detect the vehicles given a direct line of sight was found at 500m. This height was then selected as the maximum one for the evaluation of the deployment strategy presented next.

B. Deployment Strategy Evaluation

For the evaluation of the proposed UAV deployment strategy we have used the metropolitan area of Nicosia, the capital city of Cyprus, and examined the number of static UAVs required to achieve full coverage of specific parts of Nicosia's road network. The road network is comprised of all primary, secondary and residential links, and spans an area of around 100 km². For the evaluation we have used the following parameter values: $h_V = 2$ m, $h_U = 500$ m, $w_{BV} = 4$ m. The maximum height of buildings in Nicosia is determined from the postcode (areas of approximately 0.5-2 km²). A total of eleven unique maximum building heights arise in the considered area $h_M = \{5.0, 5.5, 8.0, 10.0, 11.5, 13.5, 17.0, 24.0, 38.0, 45.0, 52.0\}$ m which result in the maximum coverage ranges $R = \{664.0, 569.1, 332.0, 249.0, 209.7, 173.2, 132.8, 90.5, 55.3, 46.3, 39.8\}$ m, respectively. The maximum building heights in the metropolitan area of Nicosia are shown in Fig. 4.

We have solved the UAV deployment optimization problem in Nicosia for three different networks:

- The *primary road network* which includes Nicosia's arterial roads, comprising a set of 3507 PoIs.
- The *primary and secondary road network* which include Nicosia's main roads in addition to the arterial roads, comprising a set of 6164 PoIs.



Fig. 3. Detection Results for different heights: (a) 100m (b) 250m (c) 500m.

- The *entire road network* of Nicosia which includes all roads (primary, secondary and residential), comprising a set of 32725 PoIs.

For the above road networks we have defined matrix C in problem (4), by considering a non-uniform grid comprised of 86191 candidate UAV hovering locations, obtained by setting $\rho = 5$, as shown in Fig. 5. All zero columns have been removed from C to reduce the dimensionality of matrix C .

Figs. 6 - 8 depict the positions of the deployed UAVs to monitor the primary, primary and secondary, and entire road network of Nicosia, respectively. As can be seen from the figures, most UAVs are deployed across the road network to improve coverage. Nonetheless, in certain cases UAVs are deployed between roads or at intersections to maximize coverage. The figures also highlights the importance of the maximum building high, as in postcode areas with tall buildings the concentration of UAVs is significantly higher than areas of shorter building developments.

According to the solution of (4), a total of 198, 305 and 709 UAVs need to be deployed to cover the primary, primary and secondary, and entire road network of Nicosia, respectively. These numbers indicate that transportation operators can reduce the number of UAVs needed to monitored




Fig. 4. Illustration of the Nicosia postcode areas. The area colours indicate the maximum building height.

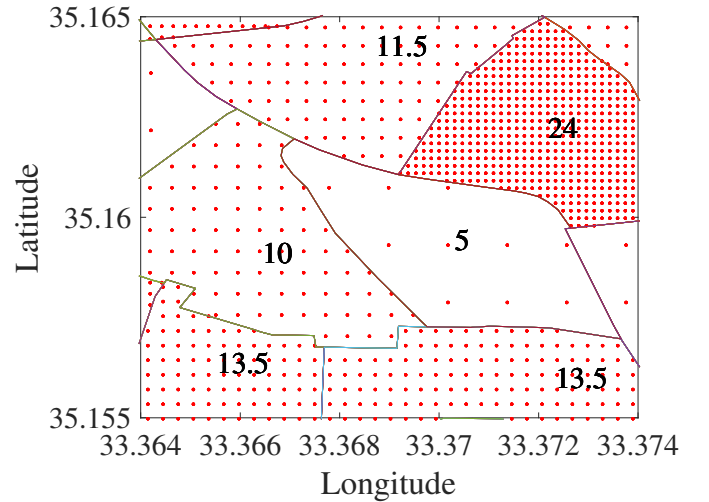


Fig. 5. Illustration of the non-uniform grid produced for the UAV hovering locations. Different regions in the figure represent postcode areas, while the numbers indicate maximum building heights (m) in different postcode areas.

the infrastructure, if a smaller set of roads is monitored. In addition, operators can utilize the particular deployment strategies to deploy UAVs in specific areas of the network to achieve rapid and detailed monitoring and to collect valuable footage that can be used to verify their road traffic mobility models.

V. CONCLUSION AND FUTURE WORK

This work serves as a practical case study on the use of UAVs for road traffic monitoring. The deployment architecture presented considers parameters related with capturing and analysing footage from above as well as the underlying road network topology and the obstacles obstructing collection of that footage. To solve the deployment challenge that arises, an optimization problem is formulated and solved to investigate the potentials of the proposed solution.




Fig. 6. UAV deployment for covering the primary road network of Nicosia. The square markers indicate UAV positions, while the bullet markers points-of-interest that need to be covered. PoIs have the same colour with the UAV that covers them. PoIs covered by multiple UAVs are shown with the colour of the nearest UAV.




Fig. 7. UAV deployment for covering the primary and secondary road network of Nicosia.

As future work, we will be exploring solutions based on light-weight deep learning algorithms (e.g., [14]) that will be able to run in real-time on the limited on-board hardware of UAVs.

ACKNOWLEDGMENT

This work has been supported by the European Unions Horizon 2020 research and innovation programme under grant agreement No 739551 (KIOS CoE).

REFERENCES

- [1] European Commission, *Towards a European strategy for the development of civil applications of Remotely Piloted Aircraft Systems (RPAS)*, Commission Staff Working Document (SWD(2012)259), Sept. 2012.
- [2] K. Kanistras, G. Martins, M.J. Rutherford, and K.P. Valavanis, *A survey of unmanned aerial vehicles (UAVs) for traffic monitoring*, International Conference on Unmanned Aircraft Systems, Atlanta, GA, 2013, pp. 221-234.




Fig. 8. UAV deployment for covering the entire road network of Nicosia.

- [3] J. Apeltauer, A. Babinec, D. Herman, T. Apeltauer, *Automatic Vehicle Trajectory Extraction for Traffic Analysis from Aerial Video Data*, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XL-3/W2, 2015, pp. 9-15.
- [4] A. De Bruin, M. Booysen, *Drone-based Traffic Flow Estimation and Tracking Using Computer Vision*, Annual Southern African Transport Conference, July 2015.
- [5] G. Maria, E. Baccaglioni, D. Brevi, M. Gavelli and R. Scopigno, *A drone-based image processing system for car detection in a smart transport infrastructure*, Mediterranean Electrotechnical Conference, Lemesos, 2016, pp. 1-5.
- [6] N. Audebert, B. Le Saux, and S. Lefvre, *Segment-before-Detect: Vehicle Detection and Classification through Semantic Segmentation of Aerial Images*, Remote Sensing, 9(4), 2017.
- [7] E.N. Barmounakis, E.I. Vlahogianni, and J.C. Golias, *Unmanned Aerial Aircraft Systems for transportation engineering: Current practice and future challenges*, International Journal of Transportation Science and Technology, 5(3), 2016, pp. 111-122.
- [8] M. Elloumi, B. Escrig, R. Dhaou, H. Idoudi and L. A. Saidane, *Designing an energy efficient UAV tracking algorithm*, International Wireless Communications and Mobile Computing Conference, Valencia, 2017, pp. 127-132.
- [9] E. Yanmaz, R. Kuschnig, M. Quaritsch, C. Bettstetter, and B. Rinner, *On path planning strategies for networked unmanned aerial vehicles*, IEEE Conference on Computer Communications Workshops, Shanghai, 2011, pp. 212-216.
- [10] S. Rasmussen, K. Kalyanam, and D. Kingston, *Field experiment of a fully autonomous multiple UAV/UGS intruder detection and monitoring system*, International Conference on Unmanned Aircraft Systems, Arlington VA, 2016, pp. 1293-1302.
- [11] N. Audebert, B. Le Saux, and S. Lefvre. *Segment-before-detect: Vehicle detection and classification through semantic segmentation of aerial images*. Remote Sensing, 9(4), 2017.
- [12] X. Chen, S. Xiang, C. L. Liu, and C. H. Pan. *Vehicle detection in satellite images by parallel deep convolutional neural networks*. In 2013 2nd IAPR Asian Conference on Pattern Recognition, pages 181-185, Nov 2013.
- [13] C. Kyrkou and T. Theocharides. *A flexible parallel hardware architecture for adaboost-based real-time object detection*. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 19(6):1034-1047, June 2011.
- [14] V. Sze, Y. Chen, T. Yang, and J. S. Emer. *Efficient processing of deep neural networks: A tutorial and survey*. CoRR, abs/1703.09039, 2017.
- [15] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*. MIT press, 2009.
- [16] I. Gurobi Optimization, "Gurobi optimizer reference manual," 2016.
- [17] C. Szegedy, A. Toshev, and D. Erhan. *Deep neural networks for object detection*. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2553-2561. Curran Associates, Inc., 2013.