Transportation Engineering 4 (2021) 100068

Contents lists available at ScienceDirect

Transportation Engineering

journal homepage: www.elsevier.com/locate/treng

Real-time traffic quantization using a mini edge artificial intelligence platform

André Broekman*, Petrus Johannes Gräbe, Wynand J.vdM. Steyn

Department of Civil Engineering, Engineering 4.0, Hillcrest campus, University of Pretoria, South Africa

ARTICLE INFO

Keywords:
Civiltronics
Traffic analysis
Mini edge computing
Object detection
Artificial intelligence
Digital twin
IoT

ABSTRACT

Traffic analysis is dependent on reliable and accurate datasets that quantify the vehicle composition, speed and traffic density over a long period of time. The utilisation of big data is required if equitable and efficient transportation networks are to be realised for smart, interconnected cities of the future. The rapid and widespread adoption of digital twins, IoT (Internet of Things), artificial intelligence and mini edge computing technologies serve as the catalyst to rapidly develop and deploy smart systems for real-time data acquisition of traffic in and around urban and metropolitan areas. This paper presents a proof of concept of a mini edge computing platform for real-time edge processing, which serves as a digital twin of a multi-lane freeway located in Pretoria, South Africa. Video data acquired from an Unmanned Aerial Vehicle (UAV) is processed using a neural network architecture designed for real-time object detection tracking of vehicles. The implementation successfully counted vehicles (cars and trucks) together with an estimation of the speed of each detected vehicle. These results compare favourably to the ground truth data with vehicle counting accuracies of 5% realised. Detection of sparse motorcycles and pedestrians were less than optimal. This proof of concept can be easily scaled and deployed over a wide geographic area. Integration of these cyber-physical assets can be incorporated into existing video monitoring systems or fused with optical sensors as a single data acquisition system.

1. Introduction

The Fourth Industrial Revolution (4IR) will see the integration of the physical world with digital and intelligent (biological) engineering. Tao et al. [1] defines a digital twin as "an integrated multi-physics, multi-scale, and probabilistic simulation of a complex product and uses the best available physical models, sensor updates, etc., to mirror the life of its corresponding twin". The convergence of physical infrastructure and virtual spaces is a combination of both elements tied through interconnected data of the physical asset. For engineering practitioners, cyber-physical data better serves the lifecycle management of infrastructure and assets. Internet of Things (IoT), described as a computer that can make sense of information without the aid of human intervention [2], remains the primary catalyst in driving connectivity of physical assets. Even though processing capabilities have increased exponentially in recent years, network bandwidth cannot keep pace with the ever-increasing demands induced by the number of devices generating data [3]. The number of IoT devices in the world are expected to reach 50 billion by the end of 2020 [4]. With data increasingly produced at the edge of a connected network, it stands to reason that the processing of said data should be executed at the edge sensor platform itself, intelligently reacting to the surrounding environment. AI stands to directly benefit the 20 sustainable development goals (SDGs) [5] as set out by the 2030 Agenda for Sustainable Development [6] for up to 71% of studies conducted – in particular, affordable and clean energy alongside sustainable cities and communities stand to benefit the most from the advances and deployment of these in intelligent technologies. This is directly tied with the provision of functional and well-maintained infrastructure to support growing economies and communities. Improved infrastructure management will form a direct by-product of hyper-connected, real-time pavement structures, materials, environment and traffic [7]. Despite the benefits posed by various studies, many AI models remain confined to controlled settings in laboratory environments using limited datasets. Extrapolation of this data is of limited practical value.

Civiltronics is the fusion of traditional civil engineering knowledge spheres with facets from electronic engineering, computer science, information technology and materials science [8]. The combination of state-of-the-art technological advancements such as miniaturized sensor platforms, artificial intelligence (AI) edge computing, Unmanned Aerial Vehicles (UAVs) featuring payload capabilities and additive manufac-

* Corresponding author.
E-mail addresses: u13025059@tuks.co.za (A. Broekman), hannes.grabe@up.ac.za (P.J. Gräbe), wynand.steyn@up.ac.za (W.J.vdM. Steyn).

https://doi.org/10.1016/j.treng.2021.100068
Received 19 August 2020; Received in revised form 20 February 2021; Accepted 15 March 2021

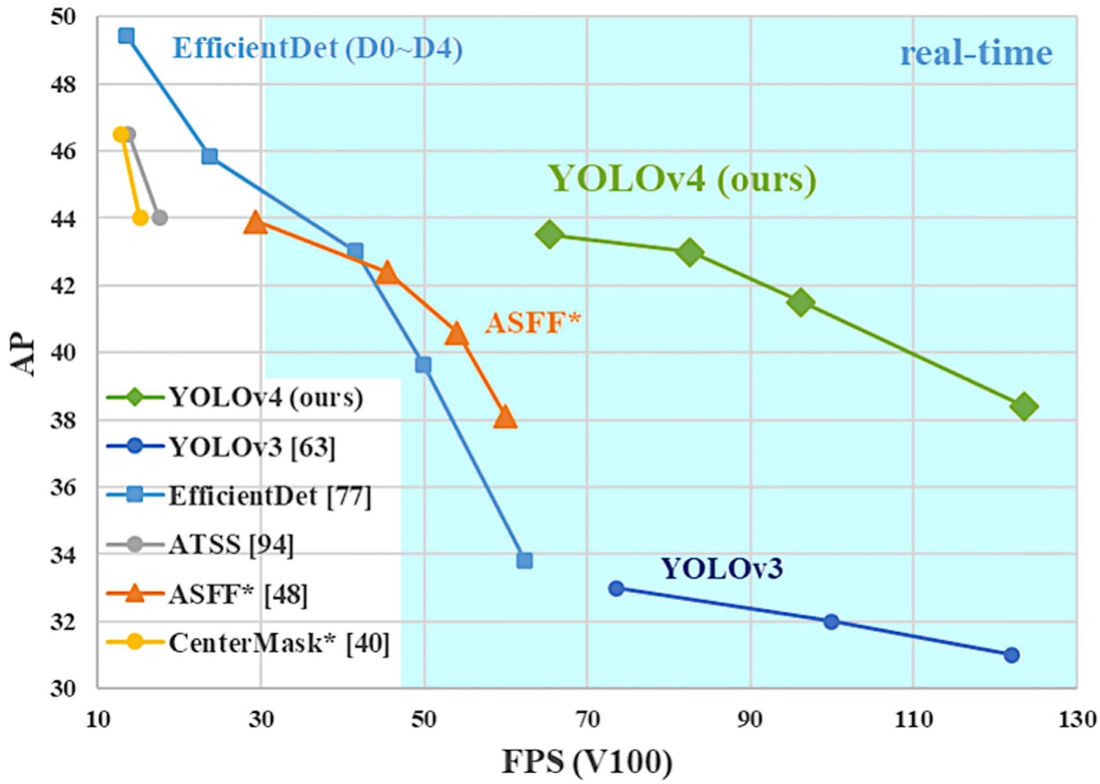

Fig. 1. The performance of YOLO4 compared to other state-of-the-art object detectors (from [26]).

turing techniques of inorganics and engineering material are opening new avenues for transdisciplinary research. The nexus of digital twins, IoT, Civiltronics and mini edge devices adds a new dimension to the capabilities available in the pursuit of addressing challenging engineering questions amid a new industrial revolution, supported by seamless integration of cyber-physical systems. These opportunities drive auxiliary transportation research projects that consider the impact of autonomous vehicles on existing methods of testing [9,10] and the digitization of railway infrastructure [11] and operations [12,13] for the realization of interconnected smart cities [14] and smart transportation [15]. The relevance thereof is clear in the context of South Africa's unequal spatial and urban development, where improved transportation infrastructure promotes and secures upward social mobility [16].

This paper presents a proof-of-concept to determine the efficacy and accuracy attained by a real-time, open-source, low-cost vehicle identification and classification artificial intelligence platform. This forms a primary component in the establishment of an integrated, wireless, digital twin of a pavement structure [17,18], where the performance of the asset itself is instrumented with temperature sensors [19] alongside environmental parameters affecting the structural performance (using LoRaWAN technology [20]). Accurate vehicle detection and classification performance is demonstrated which could be expanded for rural environments, which contribute the majority of South Africa's road network [21].

2. Materials and methods

Considering these emerging trends, a quantitative proof of concept is presented with the successful implementation of a real-time traffic analysis platform using mini edge computing technology [22]. Traffic analysis quantifies a road's performance regarding specified traffic volumes [23]. A multitude of neural network architectures have been developed during the last few years, including Faster R-CNN [24], YOLO9000 [25] and YOLOv4 [26]. Fig. 1 highlights the continued improvements realized with optimized object detection neural network architectures that provide increased average precision (AP) for predictions whilst retaining the same level of performance for real-time applications.

Accurate estimations of traffic volumes over the life of pavement infrastructure are paramount to correctly design and construct a roadway without significant deterioration. Even though design criteria are specified for different types of road capacities in South Africa [27–29], accurately monitoring the true distribution of vehicle traffic associated with these assets is not typically mandated nor considered. Approximated results using neural networks with known performance and inference accuracy is deemed a significant improvement over the sparse or non-existent availability of traffic data currently available. A "deemed to satisfy" approach provide the ability to rapidly develop and deploy edge computing solutions at scale, in contrast to the lengthy lifecycles typically associated with other comparable engineered solutions.

OpenDataCam [30] is an open-source software utility used for processing visual information from urban environments into valuable information, thereby creating a digital twin. OpenDataCam is designed specifically for Linux-based, Nvidia GPU (Graphics Processing Unit) CUDA (Compute Unified Device Architecture) enabled hardware for real-time inference applications. Mini edge devices are designed to process the data at the edge prior to uploading the processed information and relevant statistics to an integrated backend service. This approach indirectly addresses privacy and security concerns associated with the protection of personal privacy, by avoiding the storage of any media footage on the device or secondary service. Only the vehicle statistics that is composed of abstract, unidentifiable metadata is stored and transmitted. OpenDataCam is designed for compatibility to run efficiently on mini edge devices such as the low-cost Nvidia Jetson Nano [31] and the Nvidia TX2 [32] development kits. The Nvidia TX2 (Fig. 2) provides 1.33 TFLOPS (trillion floating point operations per second) in a small form factor. This unprecedented power efficiency enables real-time in-

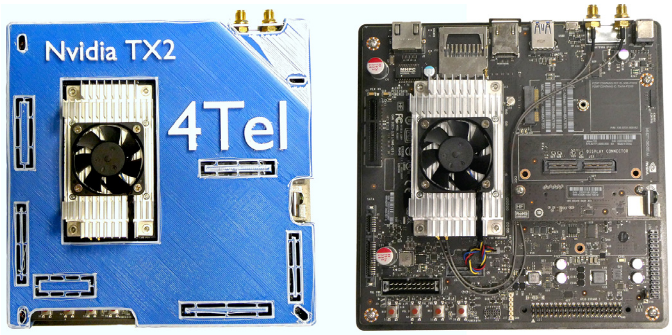

Fig. 2. Nvidia TX 2 development kit with (left) and without (right) the cover.

ference at the edge of the application, featuring a discrete GPU alongside a multi-core processor and ample memory for accelerated performance.

3. Experimental work

A Mavic Air UAV was used to capture high-resolution video of the N4 freeway (Fig. 3), adjacent to the Engineering 4.0 campus, located on the Hillcrest campus of the University of Pretoria in South Africa. This section of freeway serves as the primary arterial linking the Hatfield metropolitan area to the national freeway system. Referring to Fig. 3, Hatfield is located toward the left-hand side (driving West) with the N1/N4 flying saucer interchange located nearby toward the right-hand side (driving East). The video spans a total of 18 min and 45 s in length and was captured using a resolution of 2160 p (4 K) at a framerate of 29.97 fps (frames per second). The freeway was comparatively quiet with the video footage acquired 11:30 local time on a weekday. The UAV remained stationary throughout the data acquisition process with small yaw rotations throughout the video resulting from drift during the flight. The UAV was piloted by a certified operator, conforming to regulations as required by the South African Civil Aviation Authority (SACAA). Only the traffic driving toward Hatfield (Fig. 3, right-to-left) is considered due to the unobstructed geometry, lack of an off ramp which reduces the average speed and lack of sporadic object detection caused by the tree cover. The 4 K (8.3 MP) video was cropped using Blender [33], an open-source animation and graphics suite, reducing the resolution to 1220 × 560 px (0.68 MP). This decision followed from preliminary, qualitative analysis of the video considering the performance of different resolutions and aspect ratios. The cropped video file is available from the linked data repository [41].

4. Calculations

4.1. Ground truth

After reviewing the footage, a classification scheme was implemented using only four class labels, namely *car, truck, motorbike* and *person* to generate the ground truth in the form of a CSV file. This classification follows from the implementation of the MS COCO (Microsoft Common Objects in Context) [34] dataset that serves as the primary training dataset for the neural network implementation. There is no quantitative class distinction or definition for vehicles given the variety of sizes, color, axle and canopy configurations observed for the sample population. The resulting interpretation is highly subjective. Fig.4 illustrates a variety of vehicles that was considered to belong to the car class. Similarly, Fig. 5 illustrates a collection of the truck class. Minibus taxis, which serve as the primary mode of transportation for the majority of the population, does not have a distinct classification scheme. These were collectively labeled as a truck. Edge cases that could potentially be mislabeled or missed entirely for classification, such as the trailer ferrying vehicles to a car dealership, were rarely encountered and did not affect the inference accuracy significantly. A total of 592 vehicles were counted for the ground truth, of which 455 were cars and 137 trucks, in addition to 4 motorcycles and 1 pedestrian walking along the shoulder of the freeway.

4.2. OpenDataCam implementation

OpenDataCam features a flexible front- and back-end configuration and interface. The detailed installation procedure was followed as provided on the GitHub page [30]. The video file path used for inference is specified alongside the inference parameters and thresholds in the JSON (JavaScript Object Notation) configuration file. The preconfigured tracker settings were modified slightly to more closely resemble the ground truth dataset. The intersection over union (IoU) was lowered to 18%, improving the performance of adjacent, occluded vehicles moving in parallel. A relatively low confidence threshold of 43% for object registration compensates for the relatively poor video quality and the presence of slight motion blur. Trajectories are calculated on the previous 4 frames with an unmatched frame tolerance of 5 frames ensuring consistent tracking results of unique vehicles. A minimum angle of 15° relative to the counting lines provided ample tolerance for perspective

Fig. 3. Photograph of the N4 highway located adjacent to the engineering 4.0 campus.

Fig. 4. A collection of vehicles that were labelled as a single "car" for the ground truth.

Fig. 5. A collection of vehicles that were labelled as a single "truck" for the ground truth.

distortions caused by the drone rotation and relative angles between the camera and the vehicles. The distance between the markers was determined by comparing corresponding road features to those captured by high-resolution satellite imagery. The center-to-center distance of the lane markings are 12.7 m with each of the lanes measuring 3.4 m wide.

Traffic analysis applications require a small image frame. Traffic flow is concentrated along a narrow vertical section of the video whereas the width should produce a suitable aspect ratio together with an accurate estimation speed over the limited distance. Additionally, the proportional area of the image frame occupied by the vehicles is proportionally larger, benefitting the prediction and refinement of bounding boxes. This configuration is beneficial for the YOLO network, which more closely resembles the input distribution of the COCO training dataset. The YOLOv4-tiny neural network architecture was selected for. The *tiny* version of the YOLOv4 [28] neural network is a smaller version of the original network, comparable to that of the original YOLO implementation [35]. The reduced capacity and corresponding computational complexity improve inference performance by more than 400% compared to the larger network implementation with a negligible accuracy loss in practice (38.1 and 64.9% mAP for YOLOv4 and YOLOv4-tiny).

Fig. 6 illustrates the OpenDataCam GUI (graphical user interface). The interactive display allows the user to either draw, save or load line counters (Fig. 6, top). The total vehicle count is illustrated for every line counter. The individual class statistics are also accessible for each of the line counters on the dashboard. These statistics were downloaded in CSV format for further analysis (Fig. 6, bottom). Three parallel line counters were added perpendicular to the flow direction of the traffic. These line counters were used to both determine the count accuracy as a function of the location on the image frame and to estimate the velocity of the vehicles. Every object classification instance is assigned a unique ID that is stored alongside the frame number when the object intersects one of the line counters. With a centre-to-centre distance between the lines of approximately 13.5 m, the difference in the number of frames - corresponding to a known time difference – can be used to determine

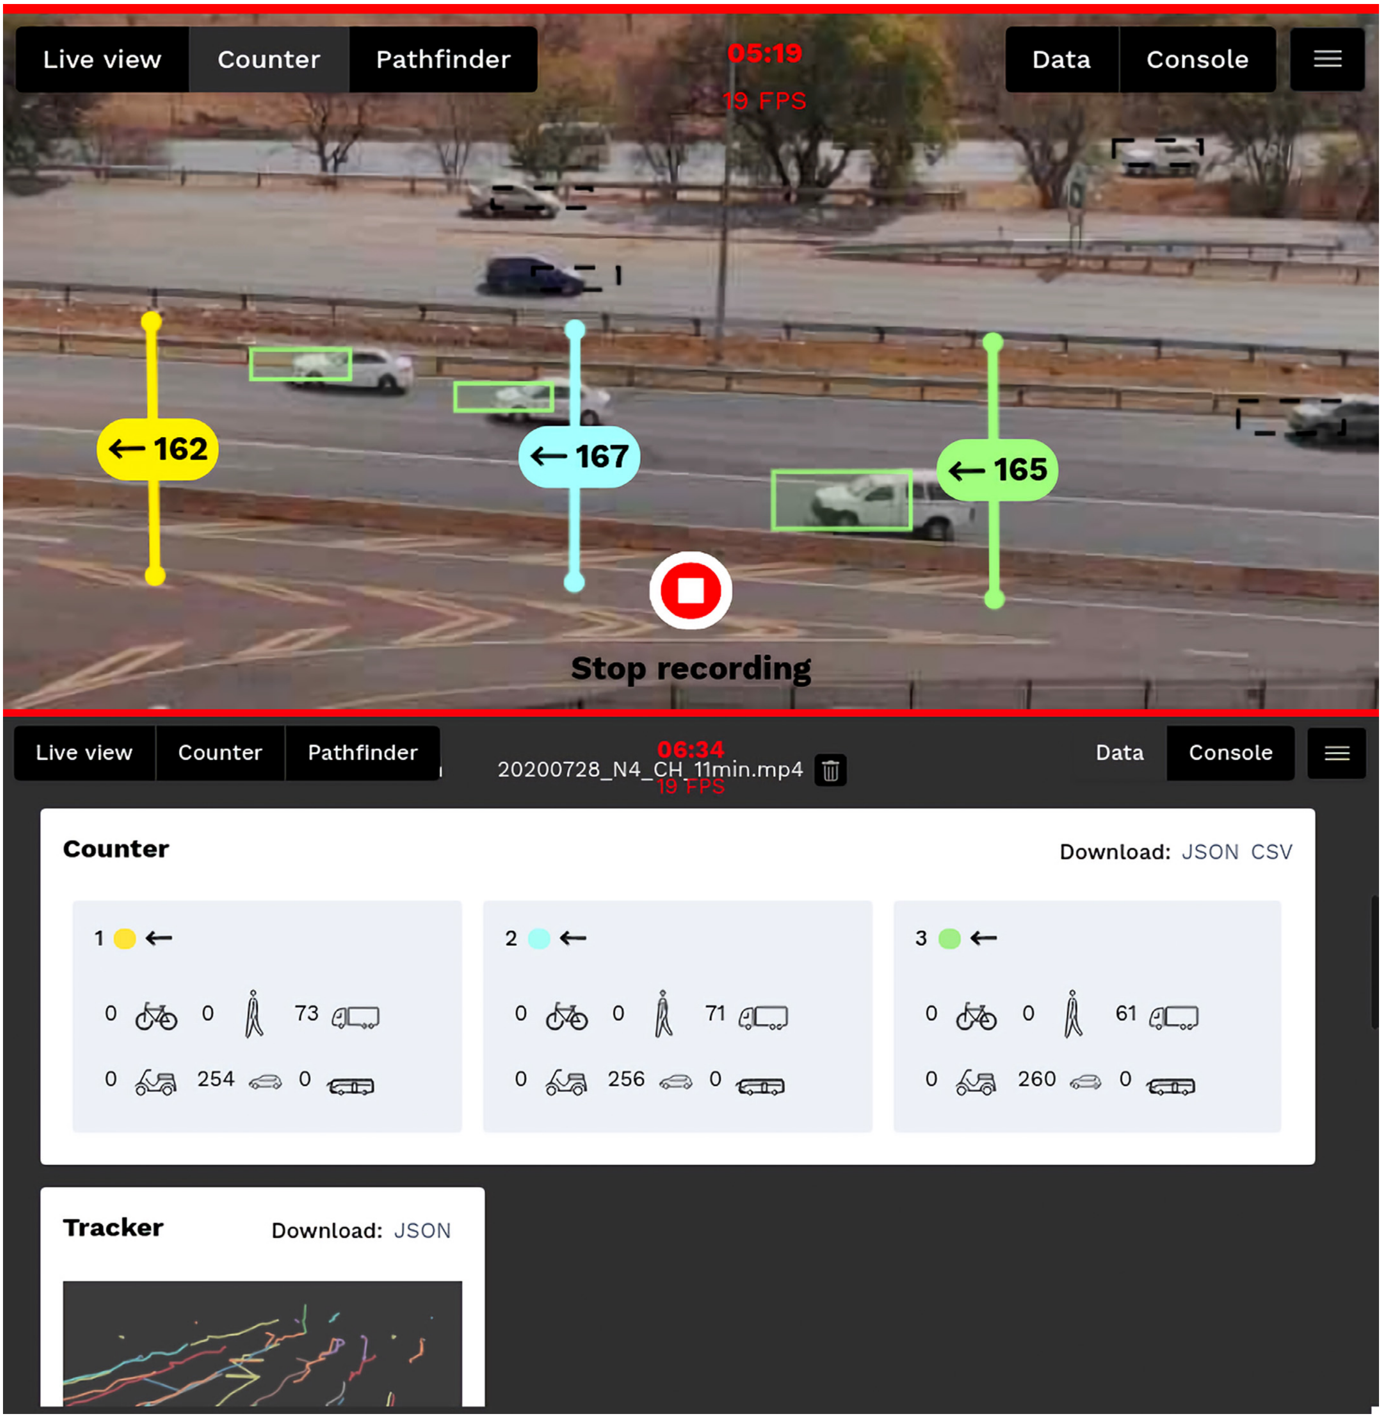

Fig. 6. Line counter configuration (top) and real-time data dashboard (bottom) for different object classes.

the speed of every tracked object. The OpenDataCam line counter configuration file is available from the linked data repository [41].

4.3. Data processing

With the counter lines defined, the "Start processing" button is pressed to initiate processing of the specified video file stored on the computer. The bounding boxes are illustrated for the detected objects in real-time during playback of the video together with the class and confidence score of every instance. The line count information was downloaded as an aggregated CSV file following the successful processing of the video file. Python was the programming language of choice to process the data. The Python source code and CSV files are available from the corresponding data repository [41]. The CSV file contains the frame number when an object (with its corresponding unique ID) crossed a specific line counter. The vehicle velocity could be determined provided that the same unique object crossed at least two of the three lines. The cumulative number of objects detected for each class can be compared to the ground truth as a measure of accuracy of the neural network. Furthermore, the data generated by OpenDataCam is stored in a MongoDB database on the local machine. The MongoDB Python package features a full API interface with the ability to view any of the databases (video files), including those of active tasks being processed such as a live video feed. This implementation of extracting real-time vehicle counts was

Table 1
Summary of vehicle counts for all line counters and the ground truth.

Vehicle class	Line counter 1	Line counter 2	Line counter 3	Average line counter	Ground truth
Car	461 (73%)	437 (72%)	408 (71%)	435 (72%)	455 (75%)
Truck	167 (27%)	161 (27%)	165 (29%)	164 (27%)	134 (23%)
Motorbike	0 (0%)	1 (1%)	0 (0%)	1 (1%)	4 (1%)
Pedestrian	0 (0%)	0 (0%)	0 (0%)	0 (0%)	1 (1%)
Total	628 (100%)	599 (100%)	573 (100%)	600 (100%)	594 (100%)

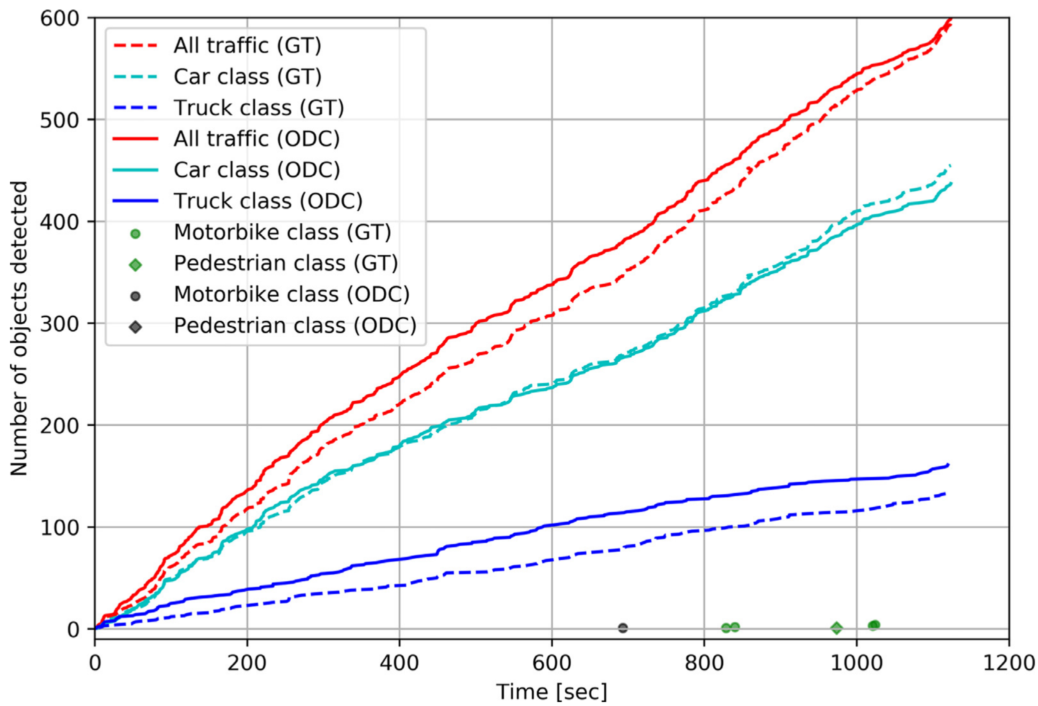

Fig. 7. Comparison between aggregated inference results obtained by line counter 2 from OpenDataCam (ODC) and the corresponding ground truth.

successfully demonstrated, marking the inception of a digital twin for this section of freeway.

5. Results and discussion

5.1. Line counter statistics

Table 1 summarises the class counts as derived by OpenDataCam, alongside the ground truth data, for the settings specified in Section 4.2. These settings strike a balance between both undercounting due to sub-optimal video quality and over counting, where false positives are more numerous if the thresholds are set too low. On average, the line counters undercounted cars by 4.4% and the trucks were overcounted by 22.4% respectively. However, if all the detected objects are considered as a single aggregated value, thereby effectively eliminating the bias of the training, the average line counter overestimates the number of objects detected by only 1%. There exists some measure of variation among the line counters. Line counter 2 provided the most accurate counting data, with line counters 1 and 3 on either side over- and underestimating the total number of objects detected by 5.7% and 3.5% respectively. These results could be attributed to the variation of the relative viewing angles and apparent size of the vehicles with respect to the camera producing slight variations in the object detection accuracy. The line counters performed poorly in detecting passing motorcycles and a walking pedestrian. This is likely due to the small relative size reducing the inference accuracy. Reducing the minimum threshold produced more favourable results for these two classes at the expense of significantly overestimating the number of vehicles.

5.2. Cumulative vehicle traffic

Fig. 7 illustrates the cumulative traffic over time graphically for both the OpenDataCam (left) implementation (line counter 2) and the ground truth (right). The traffic density of the freeway is approximated as the average gradient of the traffic count-time curve, divided by the total number of lanes, which is equal to approximately 479 vehicles-per-hour per-lane. This comparison illustrates subtle similarities, most notably for high-frequency events occurring over short periods, where the arrival of vehicles is highly stochastic.

5.3. Vehicle speed

Fig. 8 illustrates the approximated speed of every valid object pair (the same object identification number of a vehicle detected by two consecutive line counters). The speed is determined by dividing the fixed distance (13.5 m) between the counter lines by the time difference when the object was identified by the respective counter lines. The average vehicle speed as determined by OpenDataCam's results was 91 km/h with a standard deviation of 15 km/h (Fig. 8). The average speed of the trucks (90 km/h) was nearly equal to that of the cars (91 km/h).

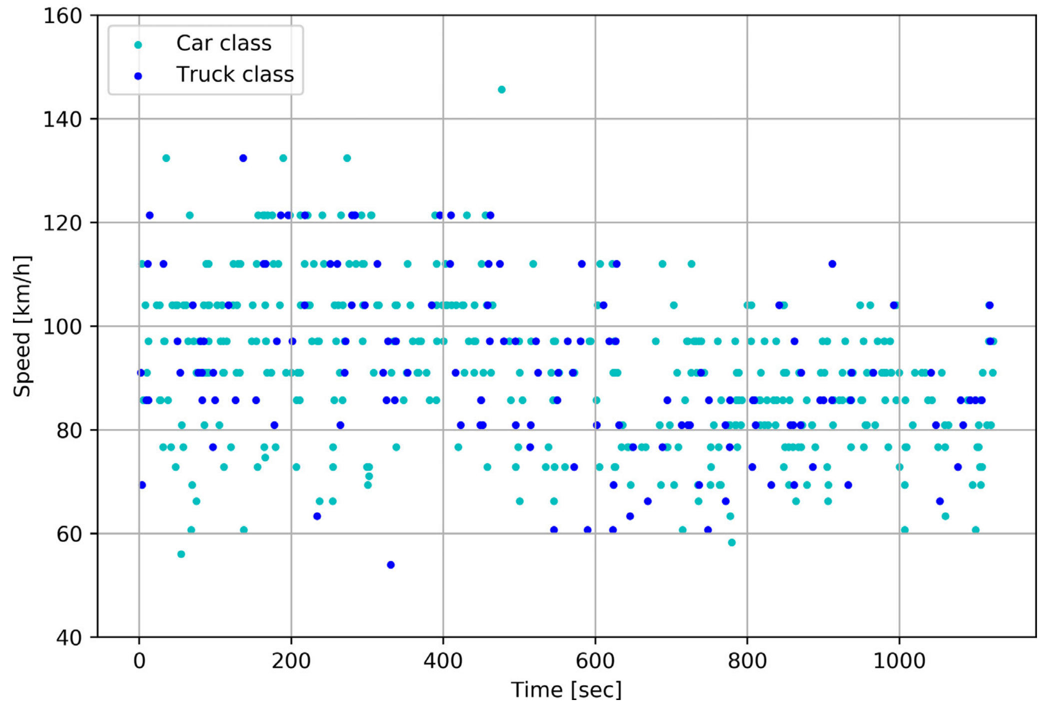

Fig. 8. Approximated velocity for every valid vehicle detected by OpenDataCam as a function of time.

The resolution of the speed calculation is limited by the short separation distance of the line counters and the frame rate of the video. Using a higher video framerate or increasing the distance between the successive line counters will increase the resolution and resulting accuracy of the speed measurements. The relationship between the average vehicles speeds and environmental parameters are of particular value; the stiffness and resulting stress distribution throughout the granular layers of the pavement structure is particularly sensitive to both the air temperature and the loading characteristics (static slow-moving or dynamic) imposed by the vehicles.

5.4. General observations

Fig. 9 illustrates a sample of the neural network processing the video file in real-time. Performance of the YOLOv4-tiny neural network varies between 20 and 28 fps for a video resolution of 1220 × 560 px. A summary of observations and edge cases are summarised:

- The truck ferrying vehicles (Fig. 9, top-right) was correctly identified as a truck, but with only two out of the six vehicles detected at all;
- A flat-bed truck towing a trailer was identified as two independent trucks (Fig. 9, bottom-right);
- A truck carrying what appears to be petroleum products was incorrectly identified as two independent, superimposed trucks (Fig. 9, bottom-left);
- Minibus taxis were identified as either a truck or car with approximately an equal distribution;
- Tarps were used by a few truck operators to cover cargo. The trucks were subsequently mislabelled, with one classified as a “boat” and another a “bench”, and
- A transgressing pedestrian walking illegally along the median was not detected due to the pedestrian’s relatively small size.

The limitation of only two vehicle classes provides limited granularity considering the variety of vehicles observed over the short time span. This could be easily addressed with the introduction of transfer learning, whereby the network continues training from its current state using more domain specific data and a larger variety of customized class labels. The adverse effect of noise and low video resolution on the network’s performance could also be improved using Super-Resolution Convolutional Neural Networks (SRCNN) [36], for example, waifu2x [37]. The Highway Capacity Manual (HCM), produced by the Transportation Research Board [38], defines the level of service (LOS). LOS defines the qualitative ranking of the traffic operational performance and capacity utilization (congestion) for a variety of transportation facilities [23]. Based on a visual assessment, a LOS B was realized during the period of data acquisition, with vehicle speeds at or near free-flow speeds observed. This observation is also reflected in Fig. 8. The actual LOS is not calculated owing to a lack of information concerning peak traffic flow and the free flow speed (FFS) and falls outside the scope of demonstrating the proof of concept.

5.5. Future improvements and research

In the South African context, minibus taxis and busses represent the primary mode of transportation. Transfer learning can easily address the existing shortcomings of the demonstrated YOLOv4 and COCO dataset implementation to identify and track these specific classes of vehicles with increased accuracy and granularity. Tethered drones sporting superior optical zoom capabilities, whilst significantly more costly, could serve as cost-effective observation platforms where existing infrastructure and internet services are more limited in certain areas. The application of OpenDataCam’s framework is not restricted to only traffic applications. Additional transportation applications include the potential of improving passenger safety on train platforms [39] and the protection of infrastructure from theft and vandalism. For agricultural applications [40], counting and classification of individual fruits in automated packhouses could be accomplished at a significantly reduced cost compared to existing, high-end computer vision solutions available.

5.6. Limitations

The media utilized for the proof of concept only spans a small window of time subject to ideal weather conditions. Note that only the vehicles closest to the camera (travelling East to West) were considered

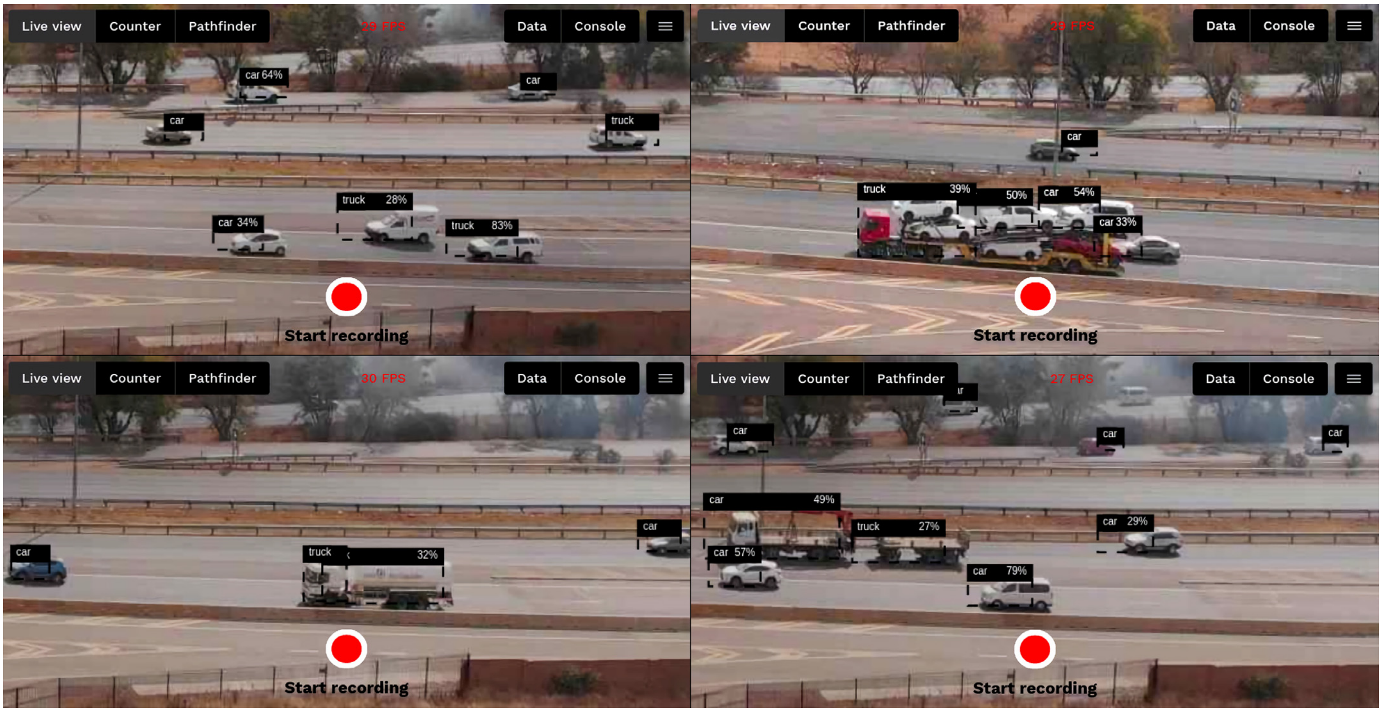

Fig. 9. A collection of real-time inference from the OpenDataCam browser-based interface.

as the obstructions (natural vegetation, barriers and lighting masts) resulted in limited detection accuracy of the traffic in the opposing travel direction. Whilst the location under discussion is not prone to fog or prolonged periods of rain or snow, these climatic effects would have to be considered for implementation on a national scale. Cameras designed for video conferencing (web cameras) are ideally positioned as a low-cost solution during daylight hours when most vehicles are travelling along the freeway. With freight and cargo transported during the evenings, more sophisticated optical sensors would have to be considered, for example, infrared capabilities. To address shadows associated with the movement of the sun, transfer learning could prove as an invaluable method to improve the overall accuracy and robustness of the neural network should this natural phenomena adversely affect the performance.

6. Conclusions

Digital twins, IoT, Civiltronics and mini edge computing are rapidly diverting traditional engineering investigation, analysis and design methodologies to a cyber-physical format. This proof of concept demonstrated the successful implementation of a real-time neural network deployed on mini edge hardware, designed to detect and track vehicles in real-time for the purpose of improved traffic analysis. Accurate detection and classification of vehicles (cars, trucks) illustrated the high level of detail and fidelity provided by the implementation. The integration of line counters included within the OpenDataCam software framework was able to accurately count and classify vehicles, with accuracies of 5% realized for vehicle classification. Detection of motorcycles and pedestrians were less than optimal, although easily resolved using improved video. Considering the suboptimal video quality and resolution and stability of the UAV, the distribution of vehicle speeds could be accurately quantified. The benefits of implementing low-cost hardware were highlighted alongside the potential to duplicate the functionality at scale without the bandwidth requirements traditionally associated with large-scale video camera networks. Development of a permanent monitoring station is underway for long-term traffic monitoring applications at the Engineering 4.0 campus.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships which have, or could be perceived to have, influenced the work reported in this article.

CRediT authorship contribution statement

André Broekman: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing - original draft. **Petrus Johannes Gräbe:** Funding acquisition, Project administration, Resources, Supervision, Writing - review & editing. **Wynand J.vdM. Steyn:** Conceptualization, Funding acquisition, Investigation, Project administration, Resources, Supervision, Writing - review & editing.

Acknowledgments

4Tel Pty is gratefully acknowledged for sponsoring the Chair in Railway Engineering in the Department of Civil Engineering at the University of Pretoria.

References

[1] F. Tao, K. Cheng, Q. Qi, M. Zhang, H. Zhang, F. Sui, Digital twin-driven product design, manufacturing and service with big data, Int. J. Adv. Manuf. Technol. 94 (2018) 3563–3576, doi:10.1007/s00170-017-0233-1.

[2] W. Shi, J. Cao, Q. Zhang, Y. Li, L. Xu, Edge computing: vision and challenges, IEEE Internet Things J. 3 (5) (2016) 637–646, doi:10.1109/JIOT.2016.2579198.

[3] S. Deng, H. Zhao, W. Fang, J. Yin, S. Dustdar, A.Y. Zomaya, Edge intelligence: the confluence of edge computing and artificial intelligence, IEEE Internet Things J 7 (8) (2020) 7457–7469, doi:10.1109/JIOT.2020.2984887.

[4] D. Evans, The Internet of Things: How the Next Evolution of the Internet is Changing Everything, CISCO White Paper, 2011 https://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf (accessed 19 February 2021).

[5] R. Vinuesa, H. Azizpour, I. Leite, M. Balaam, V. Dignum, S. Domisch, A. Felländer, S.D. Langhans, M. Tegmark, F.F Nerini, The role of artificial intelligence in achieving the sustainable development goals, Nat. Commun. 11 (233) (2020), doi:10.1038/s41467-019-14108-y.

[6] UN General Assembly (UNGA). A/RES/70/1 Transforming our world: the 2030 agenda for sustainable development. Resolut 25 (2015). https://www.un.org/en/development/desa/population/migration/generalassembly/docs/globalcompact/A_RES_70_1_E.pdf (accessed 22 March 2021).
[7] W.J. Steyn, Selected implication of a hyper-connected world on pavement engineering, Int. J. Pavement Res. Technol. (13) (2020) 673–678, doi:10.1007/s42947-020-6012-7.
[8] W.J.vdM. Steyn, A. Broekman, in: Civiltronics: Fusing civil and Electronics Engineering in the 4IR Era, Edition, SAICE Magazine, 2020, pp. 24–28. https://saice.org.za/civil-engineering-magazine/.
[9] W.J.vdM. Steyn, J. Maina, Guidelines for the use of accelerated pavement testing data in autonomous vehicle infrastructure research, J. Traffic Transp. Eng. 6 (3) (2019) 273–281, doi:10.1016/j.jtte.2019.05.001.
[10] W.J.vdM. Steyn, G. Jordaan, A. Broekman, A. Marais, Evaluation of novel chip seals applications during periods of low temperatures, in: Proceedings of the 12th Conference on Asphalt Pavements for Southern Africa (CAPSA), Sun City, South Africa, 2019 13 –16 October.
[11] A. Broekman, P.J. Gräbe, in: Keeping research on track: applications of machine learning and virtual reality, Edition, SAICE Magazine, 2020, pp. 41–45. https://saice.org.za/civil-engineering-magazine/.
[12] P.J. Gräbe, A. Broekman, Condition monitoring on the digital railway, UP Innovate Issue 14, (2019) 32–34. https://www.up.ac.za/media/shared/404/ZP_Files/innovate-14_web.zp183250.pdf (accessed 1 August 2020).
[13] D. Wust, M. Wood, Next-generation rail systems using artificial intelligence and machine learning, in: Proceedings of the International Heavy Haul STS Conference (IHHA2019), Narvik , Norway, 2019, pp. 617–624.
[14] A. Gharaibeh, M.A. Salahuddin, S.J. Hussini, B. Khreishah, I. Khalil, M. Guizini, A. Al-Fuqaha, Smart cities: a survey on data management, security, and enabling technologies, IEEE Comm. Surv. Tutor. 19 (4) (2017) 2456–2501, doi:10.1109/COMST.2017.2736886.
[15] Z. Karami, R. Kashef, Smart transportation planning: data, models and algorithms, Transp. Eng. 2 (2020), doi:10.1016/j.treng.2020.100013.
[16] N. Lionjanga, C. Venter, Does public transport accessibility enhance subjective well-being? A study of the city of Johannesburg, Res. Transp. Econ. 69 (2015) 523–535, doi:10.1016/j.retrec.2018.07.011.
[17] W.Jvd.M. Steyn, A. Broekman, Process for the development of a digital twin of a local road, in: Proceedings of the GeoChina 2021 Conference: Civil and Transportation Infrastructures: From Engineering to Smart and Green Life Cycle Solutions, 19 –21 Jul, 2021 (provisionally accepted for publication).
[18] W.JvdM. Steyn, A. Broekman, G. Jordaan, Digital twinning of asphalt pavement surfacings using visual simultaneous localization and mapping, in: Proceedings of the 2nd International Conference on Advances in Materials and Pavement Performance Prediction (AM3P), San Antonio, Texas, USA, 2020, pp. 97–100, doi:10.1201/9781003027362-24. 27-19 May.
[19] W.JvdM. Steyn, A. Broekman, Development of a Digital Twin of a local road network - a case study, J. Test. Eval. (2021) (provisionally accepted for publication).
[20] Scoop. Internet of Things for Engineering 4.0 – The LoRaWAN Catalyst Powered by Scoop. https://scoop.co.za/blog/post/internet-of-things-mikrotik-lorawan, 2020 (accessed 19 February 2021).
[21] M. Henderson, G. Van Zyl, Management of unpaved roads: developing a strategy and refining models, in: Proceedings of the 36th Southern African Transport Conference (SATC), CSIR International Convention Centre, Pretoria, South Africa, 2017, pp. 61–78. 10 - 13 July.
[22] Forbes, Classifying the modern edge computing platform. https://www.forbes.com/sites/janakirammsv/2020/07/14/classifying-the-modern-edge-computing-platforms, 1010 (accessed 31 July 2020.
[23] F.L. Mannering, S.S. Washburn, Principles of Highway Traffic Engineering and Traffic Analysis, 5th ed., John Wiley and Sons, Singapore, 2013 SI Version.
[24] S. Ren, K. He, R. Girshick, J. Sun, R.C.N.N. Faster, Towards real-time object detection with region proposal networks, IEEE Trans. Pattern Anal. Mach. Intell. 39 (6) (2015) 1137–1149, doi:10.1109/TPAMI.2016.2577031.
[25] J. Redmon, A. Farhadi, YOLO9000: better, faster, stronger, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 2017, pp. 7263–7271, doi:10.1109/CVPR.2017.690. 21 –26 July.
[26] A. Bochkovskiy, C.Y. Wang, H.Y.M. Liao, YOLOv4: optimal speed and accuracy of object detection (2020). https://arxiv.org/abs/2004.10934 (accessed 22 March 2021).
[27] Council for Scientific and Industrial Research (CSIR), Guidelines for human settlement planning and design, Chapter 7: roads: geometric design and layout planning (2000), https://www.csir.co.za/sites/default/files/Documents/Chapter_07_Vol_II_Roads.pdf
[28] South African National Roads Agency Limited (SANRAL), COTO standards specification, TMH16, Volume 2: South African traffic impact and site traffic assessment standards and requirements manual (2014), https://www.accessmanagement.info/sites/default/files/TMH16-Traffic-Impact-Study-Manual-Vol-2.pdf.
[29] South African National Roads Agency Limited (SANRAL), South African guidelines for single carriageway National roads, First Draft (2009), https://www.nra.co.za/download/141/may-2009/4259/policy-design-guideline-standards-for-single-carriageway-national-roads-draft-may-2009.pdf
[30] OpenDataCam, OpenDataCam 3.0.1 – an open source tool to quantify the world. https://github.com/opendatacam/opendatacam, 2021 (accessed 19 February 2021)
[31] Nvidia, Jetson Nano Developer Kit. https://developer.nvidia.com/embedded/jetson-nano-developer-kit, 2021 (accessed 19 February 2021).
[32] Nvidia, Harness AI at the Edge with the Jetson TX2 Developer Kit. https://developer.nvidia.com/embedded/jetson-tx2-developer-kit, 2021 (accessed 19 February 2021).
[33] Blender, Open source 3D creation. Free to use for any purpose, forever. https://www.blender.org/, 2021 (accessed 19 February 2021).
[34] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, Microsoft COCO: common objects in context, in: Proceedings of the 13th European Conference Computer Vision, Zurich, Switzerland, 2014, pp. 740–755. arXiv:1405.0312.
[35] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: unified, real-time object detection, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 2016, pp. 779–788, doi:10.1109/CVPR.2016.91.
[36] C. Dong, C.C. Loy, K. He, X. Teng, Image super-resolution using deep convolutional networks, in: Proceedings of the 13th European Conference on Computer Vision Part IV, Zurich, Switzerland, 2014, pp. 184–199, doi:10.1007/978-3-319-10593-2_13. September 6-12.
[37] waifu2x, Image Super-resolution for anime-style art using deep convolutional neural networks. https://github.com/nagadomi/waifu2x, 2021 (accessed 19 February 2021).
[38] Transportation Research BoardHighway Capacity Manual, National Research Council, Washington, DC, 2010.
[39] R.L. Myoya, P.J. Gräbe, C.J. Venter, A case study on platform-train interface occurrences at a metrorail passenger station, in: Proceedings of the 37th Annual South African Transport Conference, Pretoria, South Africa, 2018 9-12 July.
[40] A. Broekman, Steyn W.Jvd.M, SMAVO: packhouse optimization using smart avocadoes, in: Proceedings of the SAAGA Research Symposium, Tzaneen, Limpopo, South Africa, 2020.
[41] A. Broekman, P.J. Gräbe, W.Jvd.M. Steyn, DroneTrafficZA2020: a sample dataset of N4 freeway traffic adjacent to Engineering 4.0 in South Africa, Mendeley Data, v 1, 2020. 10.17632/wdfkxsnpzc.1.