Supplementary material for: Estimating road vehicle speed from high-resolution satellite imagery for environmental applications: a case study of Barcelona.

Annalisa Sheehan[1]*, Andrew Beddows[2], John Gulliver[1], David C. Green[2,3] and Sean Beevers[2,3]

[1] Population Health Research Institute, School of Health and Medical Sciences, City St George's, University of London, London, United Kingdom.

[2] MRC Centre for Environment and Health, Imperial College, London W12 0BZ, UK

[3] NIHR HPRU in Environmental Exposures and Health, Imperial College, London W12 0BZ, UK

* Corresponding author: asheehan@sgul.ac.uk

Supplementary A: Extracting pixels relating to vehicle movement.

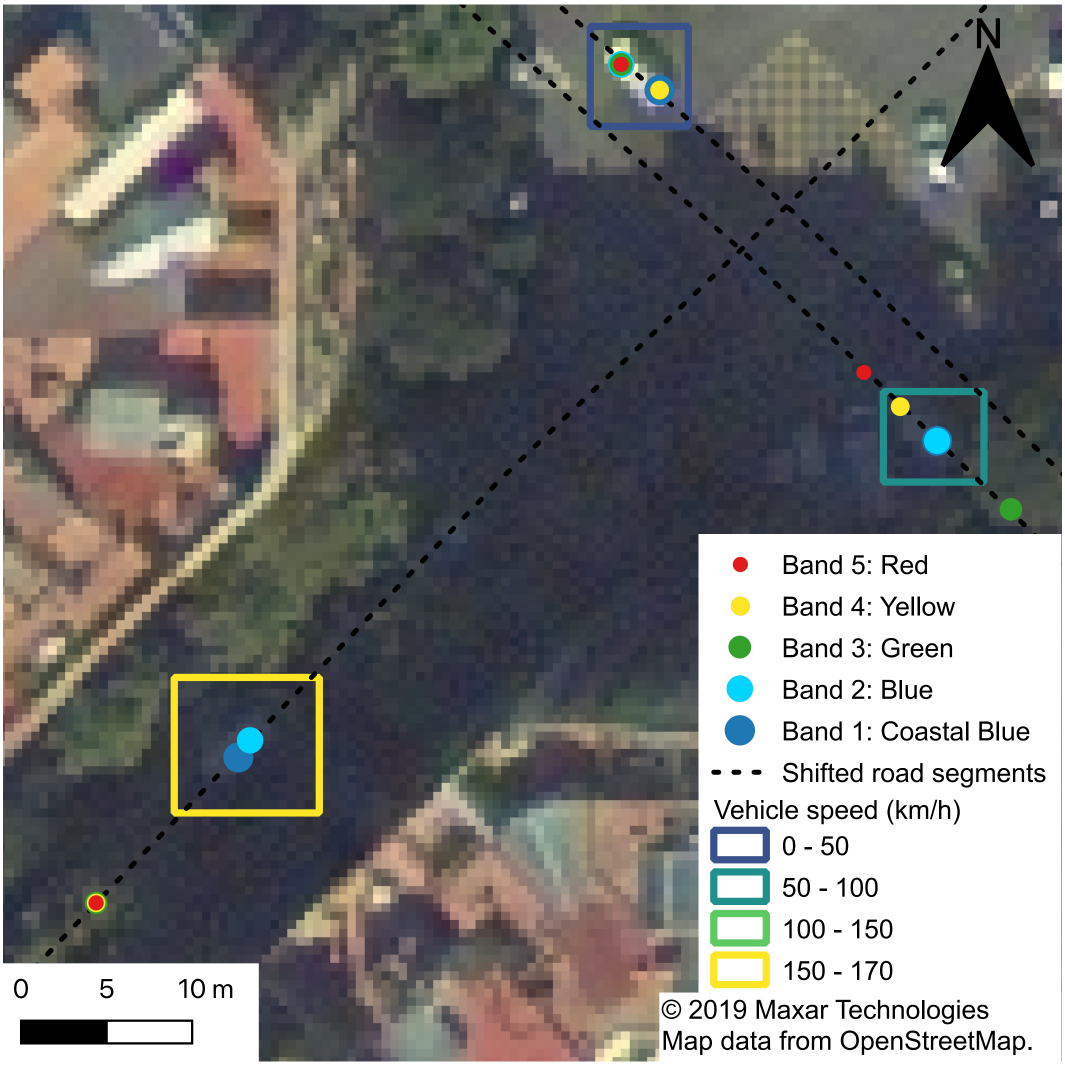

Figure A1. Map showing the locations of the extracted spectral band pixels using the final vehicle speed method (Coastal Blue, Blue, Green, Yellow and Red) and the corresponding estimated vehicle speed (km/h). Example of how vegetation and shadow influence the vehicle speed method and the extracted pixels.

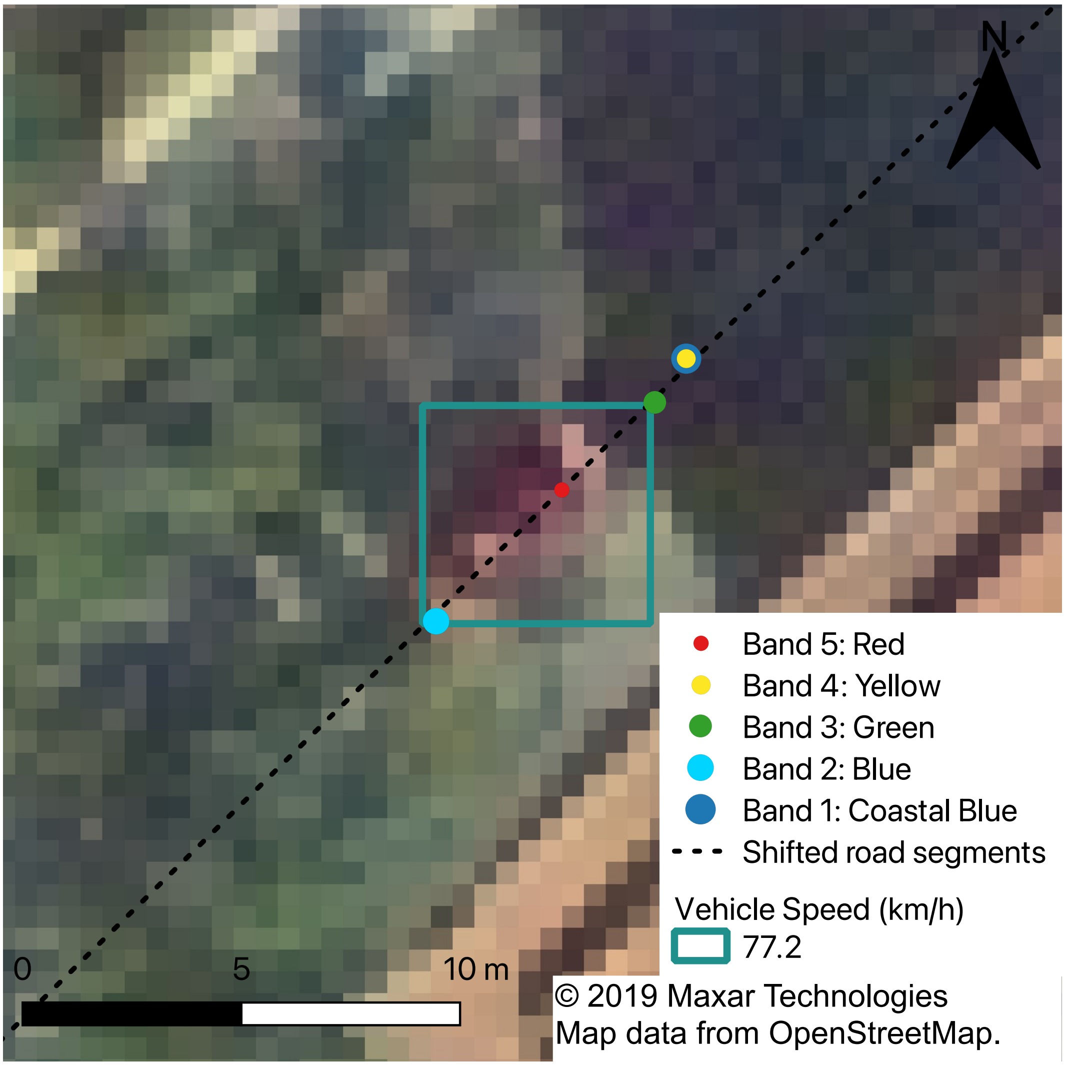

Figure A2: Map showing the locations of the extracted spectral band pixels using the final vehicle speed method (Coastal Blue, Blue, Green, Yellow and Red) and the corresponding estimated vehicle speed (km/h). Example of the vehicle speed method applied to a red vehicle.

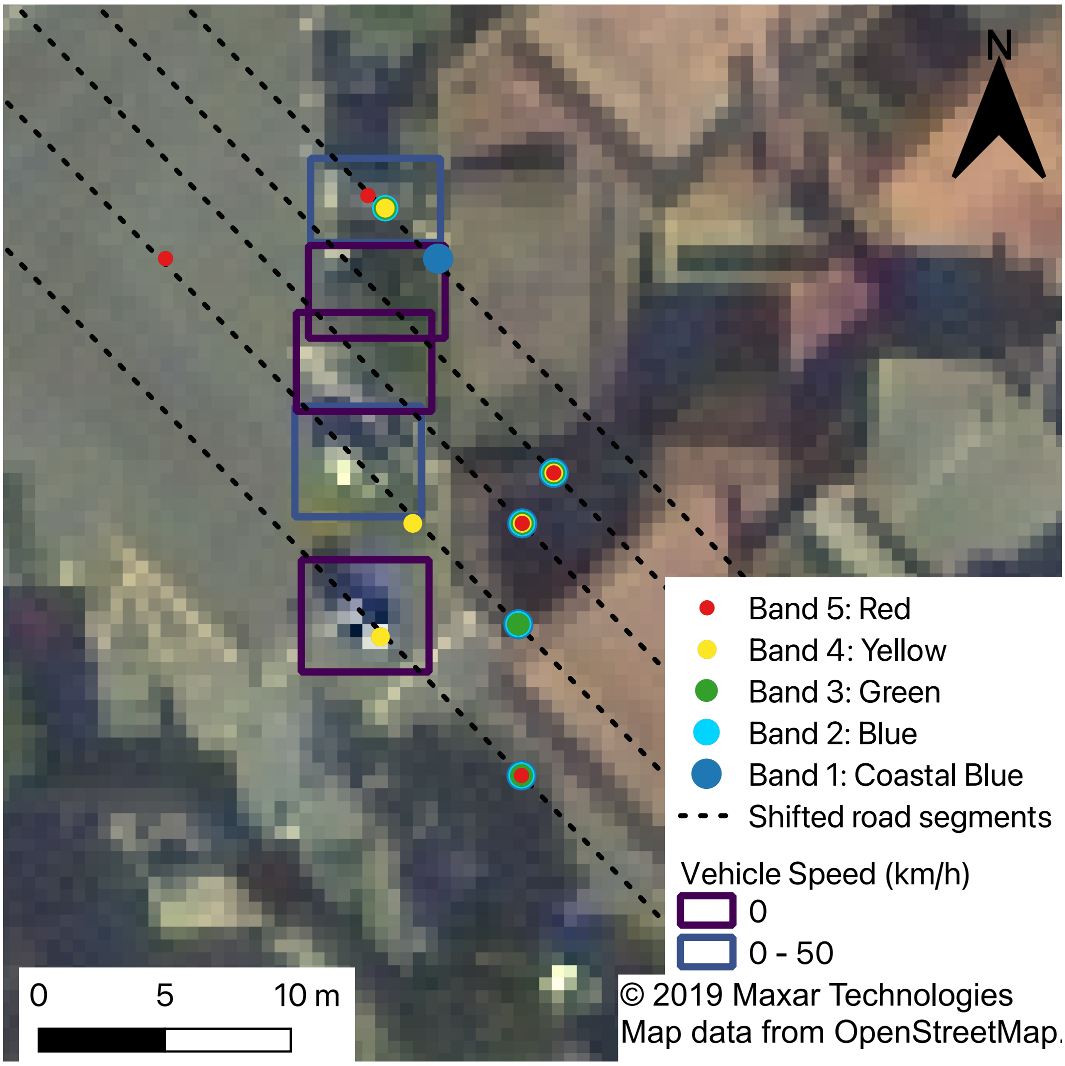

Figure A3: Map showing the locations of the extracted spectral band pixels using the final vehicle speed method (Coastal Blue, Blue, Green, Yellow and Red) and the corresponding estimated vehicle speed (km/h). Example of how buildings influence the vehicle speed method and the extracted pixels.

Supplementary B: Removing parked and erroneous vehicle detections

B.1: Method to remove parked and erroneous vehicle detections

The DL OD vehicle detections data set (section 2.1.2.) included vehicles which were parked or erroneously detected. To remove detections not related to moving vehicles, whilst ensuring moving vehicles which were static at the time of image capture (e.g. waiting at traffic lights) were kept, four cleaning steps were completed.

The first cleaning step, `*distance to road threshold',* was to remove any vehicle detections in tunnels, and further than 15 m from a major road and 10 m from a minor road. This step was aimed at removing erroneous detections within buildings or in vegetated areas. The distances 15 m and 10 m were chosen after a subset of major and minor roads in Barcelona were measured to determine the average road width.

The second cleaning step, *`maximum speed threshold'*, was to remove any vehicles with unrealistically high vehicle speeds. A conservative threshold value of 170 km/h was used based on the 5 m buffer around each detection used in the vehicle speed method and the maximum speed a vehicle can be moving to cover the distance.

The third cleaning step, *`road edge threshold',* used a road edge data set to remove parked vehicles from minor roads. A road edge data set for Barcelona was created using the x and y offsets for the shifted roads (produced for all detections in the vehicle speed method) and the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) clustering method (*Ester et al., 1996*). The DBSCAN method requires two criteria: a threshold for the minimum number of points considered and an arbitrary distance measure (*Schubert et al., 2017; Khan et al., 2014*). In this work, the DBSCAN clustering method used a minimum of two neighbours and a 1 m distance between points. The minimum of two neighbours were chosen to limit the impact of a single erroneous detection being classified as a road edge or lane. A 1 m distance between points was used to ensure each road lane and edge was clustered into separate clusters. The outer road clusters were selected on roads where there were more than two clustered sets of roads, to ensure both sides of the road edge were extracted and to account for roads that only allowed parking on a single side of the road. The road edge data set was then buffered by 1 m and intersected with the vehicle detections to remove any detections assumed to be parked vehicles on the sides of roads.

The final cleaning step, *`removing SLR roads',* removed all service, living streets and residential roads from the data set. The method struggled to identify vehicles on these road types as they tend to be narrow streets where buildings and vegetation often obstruct the road from view, with a large proportion of parked vehicles.

Figure B1 show the detections removed at each data set cleaning step, with black BBOX's representing the original model detections (and the detections removed in the first cleaning step) and yellow BBOX's representing the final cleaned vehicle speed data set. In **Figure B1,** the detections mainly removed in the first cleaning step (black) are over car parks and buildings. The red BBOX's (detections removed during the road edge thresholding step) mainly cluster on either side of minor roads. Finally, the orange BBOX's show the detections on SLR, which are the smaller side roads. The final vehicle speed data set in yellow, shows the final cleaned vehicle speed data set with the BBOX's on the main roads in the image.

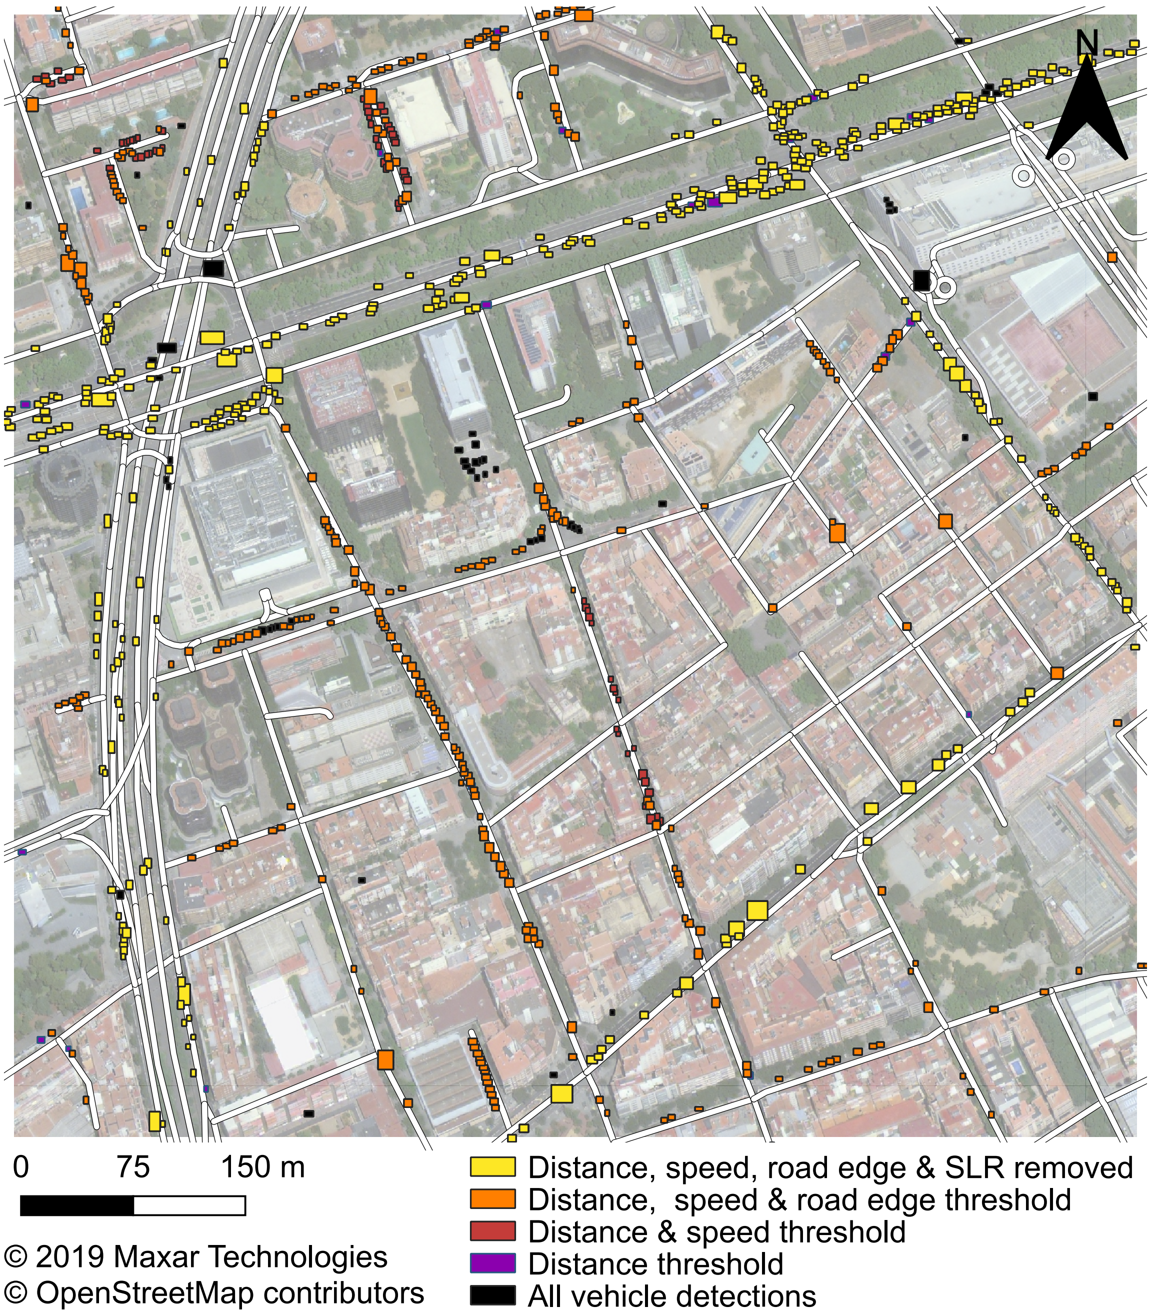

Figure B1. July subset vehicle speed validation tile highlighting the detections removed at each cleaning step. The black BBOXs represent the original model detections and the detections removed using the distance threshold (1st cleaning step), purple BBOXs representing further detections removed after using the speed threshold (1st and 2nd cleaning steps), red BBOXs representing further detections removed after using the road edge threshold (1st, 2nd and 3rd cleaning steps), orange BBOXs representing further detections removed after removing all SLR roads (1st, 2nd, 3rd and 4th cleaning steps) and yellow BBOXs representing the final cleaned vehicle speed data set after all four cleaning steps have been applied.

Figure B2 shows the detections removed at each data set cleaning step, with black BBOX's representing the original model detections (and the detections removed in the first cleaning step) and yellow BBOX's representing the final cleaned vehicle speed data set. In **Figure B2** the detections mainly removed in the first cleaning step are over car parks and in buildings, with detections on minor roads and other erroneous detections removed over the data cleaning steps, and the final data set in yellow showing detections on main roads.

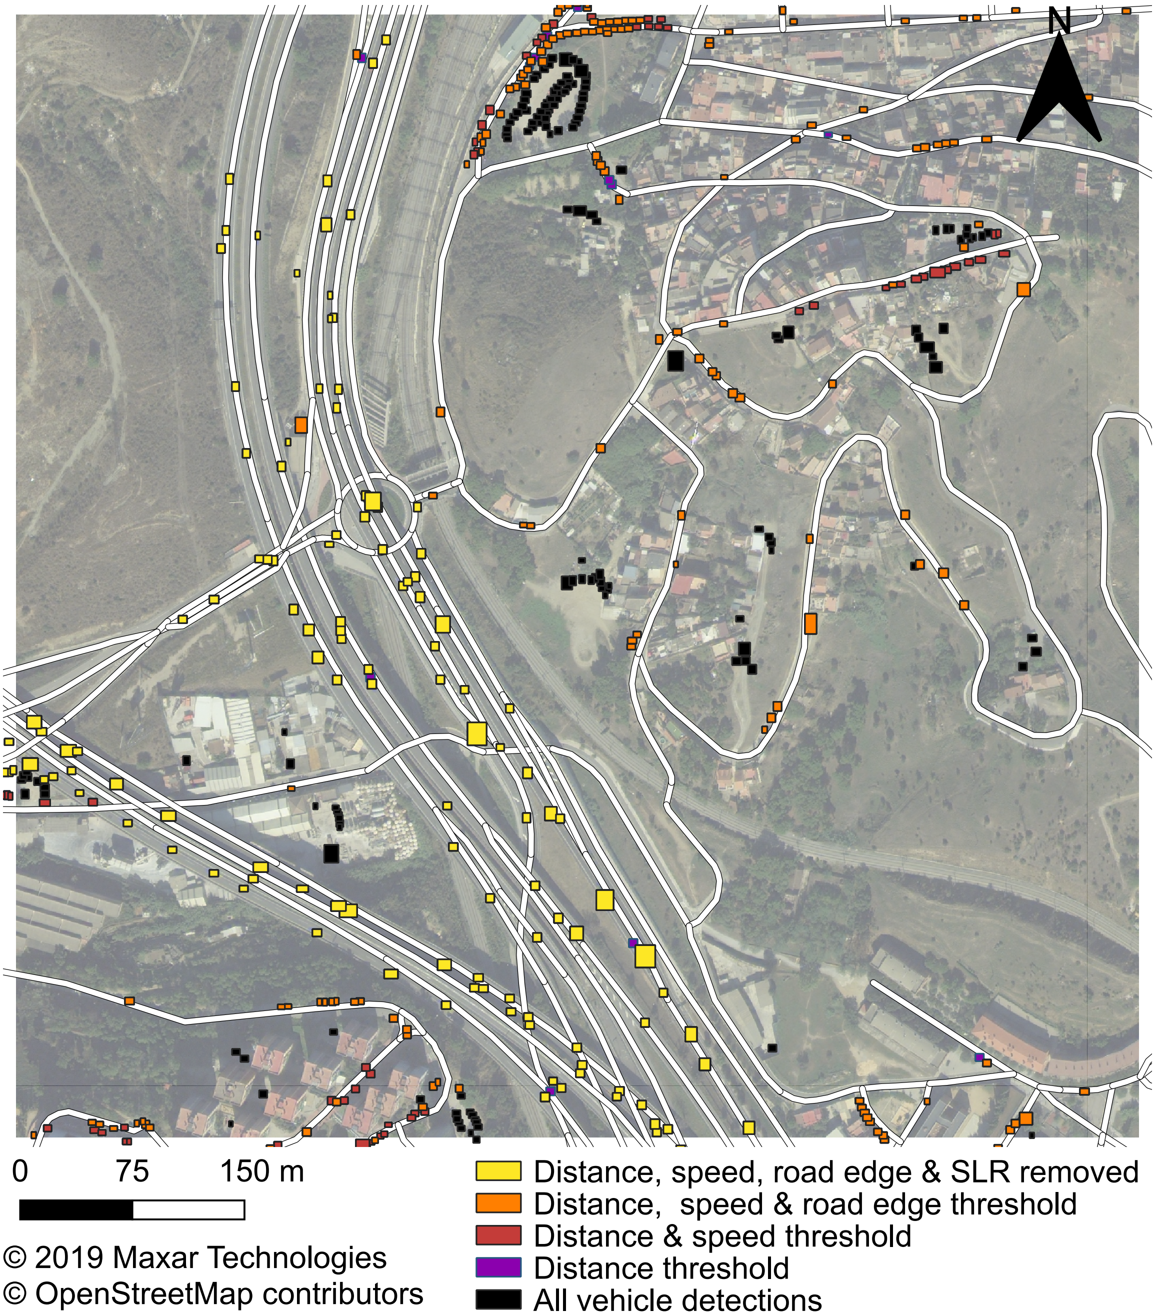

Figure B2. October subset vehicle speed validation tile highlighting the detections removed at each cleaning step. The black BBOX's represent the original model detections and the detections removed using the distance threshold (1st cleaning step), purple BBOX's representing further detections removed after using the speed threshold (1st and 2nd cleaning steps), red BBOX's representing further detections removed after using the road edge threshold (1st, 2nd and 3rd cleaning steps), orange BBOX's representing further detections removed after removing all SLR roads (1st, 2nd, 3rd and 4th cleaning steps) and yellow BBOX's representing the final cleaned vehicle speed data set after all four cleaning steps have been applied.

B.2: Evaluating the vehicle detection cleaning steps

To understand whether the erroneous and parked vehicle detections were successfully removed, the speed data set, at each cleaning step, was evaluated with a manually digitised validation data set of static and moving vehicles in a subset of the Barcelona imagery (*Sheehan et al., 2023*).

For each cleaning step the number of ground truth BBOXs (validation data set) and cleaned model detections were counted, the number of True Positives (TP's), False Positives (FP's) and False Negatives (FN's) were calculated. Validation metrics including precision, recall and F1 score were produced, and equations shown below. The precision metric (Equation 1) is the percentage correct detections out of the total number of model detections.

$$\mathbf{Precision} = \frac{\mathbf{True\ Positives}}{(\mathbf{True\ Positives} + \mathbf{False\ Positives})} \tag{1}$$

range = 0 and 1 (all the model's identified objects are correct).

The recall metric (Equation 2) measures the percentage correct detections out of the total number of validation data set detections.

$$\mathbf{Recall} = \frac{\mathbf{True\ Positives}}{(\mathbf{True\ Positives} + \mathbf{False\ Negatives})} \tag{2}$$

range = 0 and 1 (model correctly identified all the objects).

The F1 score (Equation 3) is a performance metric that provides an overall evaluation of a model through combining the precision and recall metrics into a singular value.

$$\mathbf{F1\ score} = \mathbf{2} \times \frac{\mathbf{Precision} \times \mathbf{Recall}}{\mathbf{Precision} + \mathbf{Recall}} \tag{3}$$

range = 0 (both precision and recall are low: the model did not correctly identify objects and had many false detections) and 1 (both the precision and recall are high: the model identified all objects in an image correctly).

B.3: Moving and static vehicle detections evaluation

To understand whether the cleaning method was successful in removing parked and erroneous detections from the DL OD vehicle detection data set, an evaluation using the Static and Moving (SM) validation data set was completed on each of the different cleaning steps and is shown in **Table B1**. The F1 score greatly improved from 0.27 with no cleaning steps to 0.73 after all four cleaning steps were completed (**Table B1**, final row). The recall shows that the vehicle speed method and cleaning steps have successfully identified 76 % of SM vehicles on motorway, trunk, primary, secondary and tertiary roads in Barcelona. The results show that many erroneous detections and parked vehicles have been removed as the precision has increased from 0.16 to 0.71. Overall, removing the SLR roads from the data set (cleaning step four) had the greatest impact on precision and F1 score.

Table B1. Object detection SM vehicle validation for the four vehicle speed cleaning steps and different parameters for the road edge data set.

Data cleaning step	Precision	Recall	F1 score
SM Validation	0.16	0.79	0.27
Distance and speed threshold	0.24	0.74	0.36
Distance, speed and road edge	0.28	0.72	0.40
Distance, speed, road edge and SLR roads removed	0.71	0.76	0.73

B.4: Method performance in identifying static and moving vehicle detections from satellite imagery

The detection cleaning method (using distance, speed, road edge filters and removal of SLR roads) was effective at removing erroneous and parked vehicles from the data set, shown by the F1 score (compared to the SM validation data set) improving from 0.27 to 0.73. The validation shows that whilst the cleaning method is efficient at removing parked and erroneous detections, approximately 3% more SM vehicles were removed (recall reducing from 0.79 to 0.76) with 24% of static and moving vehicles missing from the vehicle speed data set. There are many other reasons why vehicles may be missing aside from the vehicle detection and speed method, such as the limitations of satellite imagery captured from a bird's eye view resulting in parts of road and vehicles obstructed from view by buildings, vegetation and shadow and therefore not easily discerned.

The different vehicle cleaning steps removed parked and erroneous detections from the data set based on assumptions of the environment and urban morphology of Barcelona. The road edge method to remove parked vehicles was challenging as firstly, not all roads in the city have cars parked on one or both sides of the road. Therefore, the road edge data set was prone to removing static or moving vehicles which were in the outer lanes of a road. Secondly, some OSM road segments cross junctions and therefore detections travelling through a junction may have been removed. Another aspect that proved difficult was cleaning vehicle detections and estimating speed for SLR roads. This was because the streets were narrow, where buildings, vegetation and shadow often obstructed the road from view and contributed to the noisy image scenes influencing the methods ability to select pixels relating to vehicle movement. Additionally, a very large proportion of the vehicles on SLR roads were parked and the vehicles that were moving were more likely to be below the limit of detection (approximately 12.5 km/h). Barcelona is a densely built-up city with high-rise buildings, in areas with lower building heights and density, the methods presented here should create less challenges.

Supplementary C: Vehicle speed results

Estimated speeds for detections over a junction across a primary and secondary road is shown in **Figure C1**. All the speeds are less than 100 km/h, with the majority of vehicles around the junction with correct 0 km/h speeds. A few moving vehicles have speeds of less than 50 km/h and one detection at the front of the queue has a speed of 69 km/h which is likely to be correct. There are a few static vehicles with incorrect fast speeds which are likely to have occurred due to other vehicles queuing in the same lane as the vehicle influencing the pixels extracted for the speed estimation.

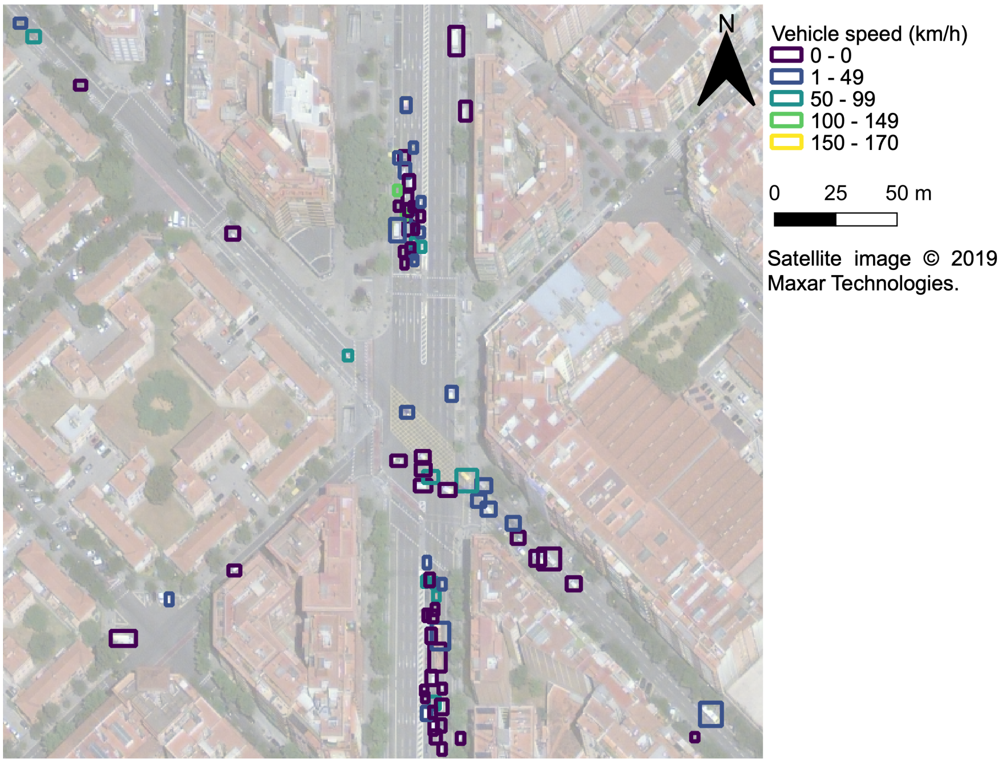

Figure C1. Estimated vehicle speeds for detections in the final data set after the four cleaning steps, an example of a junction for a primary and secondary road.

Figure C2 shows another example of detections over a junction on two primary roads with multiple lanes of traffic. The image shows that estimated speeds are between 0 km/h and 150 km/h with the majority of vehicles on primary roads away from junctions have speed estimates in the categories of 1 km/h to 49 km/h and 50 km/h to 99 km/h. Most static vehicles at the junction have speed estimates of 0 km/h however, there are a few high speeds relating to erroneous detections over road markings.

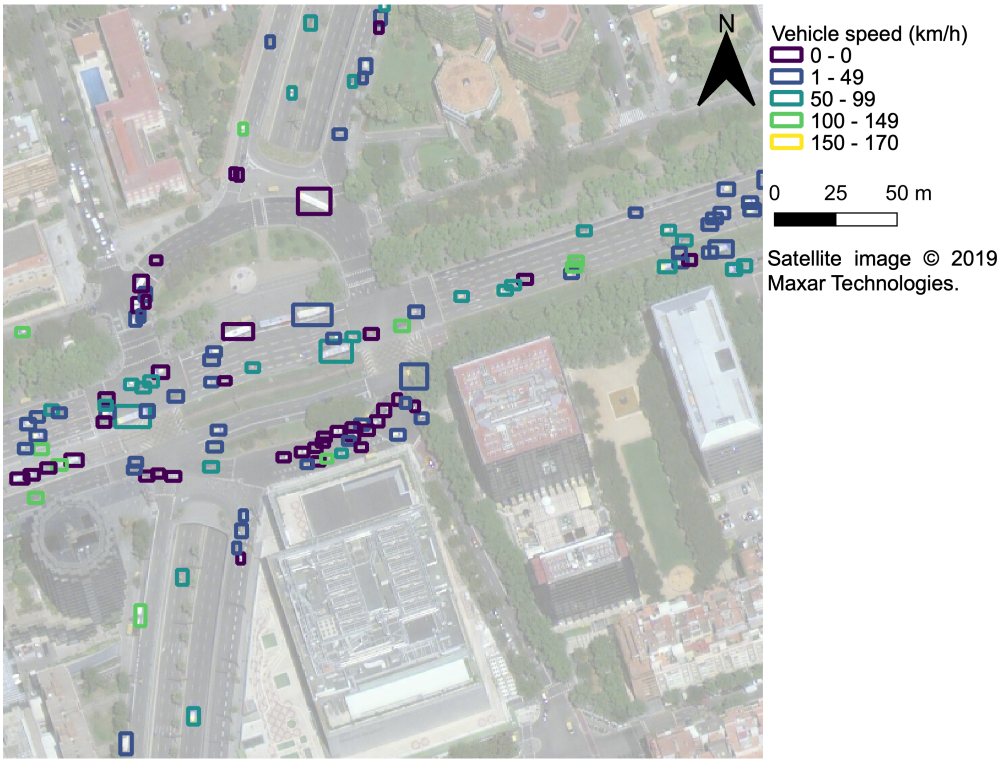

Figure C2. Estimated vehicle speeds for detections in the final data set after the four cleaning steps, an example of a junction on two primary roads.

Supplementary references:

Ester, M., Kriegel, H.P., Sander, J., Xu, X. 1996. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, AAAI Press, Portland, Oregon. 226–231. doi:10.1016/B978-044452701-1.00067-3.

Khan, K., Rehman, S.U., Aziz, K., Fong, S., Sarasvady, S. 2014. DBSCAN: Past, Present and Future. The Fifth International Conference on the Applications of Digital Information and Web Technologies (ICADIWT 2014). 232–238.

Schubert, E., Sander, J., Ester, M., Kriegel, H.P., Xu, X., 2017. DBSCAN revisited, revisited: Why and how you should (still) use DBSCAN. ACM Transactions on Database Systems. 42. doi:10.1145/3068335.

Sheehan, A., Beddows, A., Green, D.C., Beevers, S. 2023. City scale traffic monitoring using worldview satellite imagery and deep learning: A case study of Barcelona. Remote Sens. 15. doi:10.3390/rs15245709.