Australasian Transport Research Forum 2023 Proceedings
29 November – 1 December, Perth, Australia
Publication website: http://www.atrf.info

Changing priority: The before-after evaluation of an active transport project using Video Analytics

Sergio Banchero[1], Yuchao Sun[1*], Daniel Demiris[1], James Pearse[2], Craig Wooldridge[3]
[1]Planning and Transport Research Centre (PATREC), Perth, WA
[2]Hatch (formally WSP), Perth, WA
[3]Main Roads Western Australia, Perth, WA
[*]Email for correspondence (presenting author): chao.sun@uwa.edu.au

1. Introduction

Around Australia and many other developed nations, governments are upgrading road infrastructure to make it more friendly for pedestrians and cyclists, as well as other forms of active transport. However, there is a lack of quality data in assessing the effectiveness of different treatments due to technical difficulties in detecting road user movements. While pneumatic tubes can measure the speeds and volumes of cyclists relatively reliably, this technology poses a potential trip hazard to some path users and is not capable of accurately measuring pedestrian volumes. In this paper, we demonstrate how Video Analytics (VA) can be used to fill in this gap using a case study of priority change at three bike path crossings. The observed changes in vehicle speed and their give-way compliance were analysed, together with some challenges and lessons learnt during the project.

2. Background

The bike path on the South Perth Foreshore is a popular destination for cyclists. However, it was previously constructed as vehicle-priority, requiring cyclists to give way to vehicles coming in and out of two nearby car parks. To create a better riding experience, the City of South Perth and Main Roads Western Australia upgraded three of its crossings, Douglas Avenue, Coode Street, and Witcomb Place, providing cyclists with priority over motorists (Figure 1).

Main Roads have previously collected videos at these locations before the construction in 2021. The videos were collected by SurveyTech on Wednesday 17th March 2021 and Sunday 21st March 2021, both from 6 AM to 7 PM (and are referred to as 'before' videos). The 'after' videos were collected by SurveyTech on Wednesday 29th June and Sunday 26th June. The weather on all four days was sunny. The Planning and Transport Research Centre (PATREC) hosted by the University of Western Australia (UWA), has been subcontracted by WSP to analyse the changes using VA.

This is the first wave of 'after' surveys to establish immediate changes. Future surveys are planned to establish the long-term trend.

The project has been approved by UWA's Human Research Ethics Committee. It did NOT involve any face recognition. We blur people's faces when they are distinguishable in the video or screenshot to protect their privacy.

Figure 1: Crossing locations and the comparison of 'before' and 'after' the upgrade

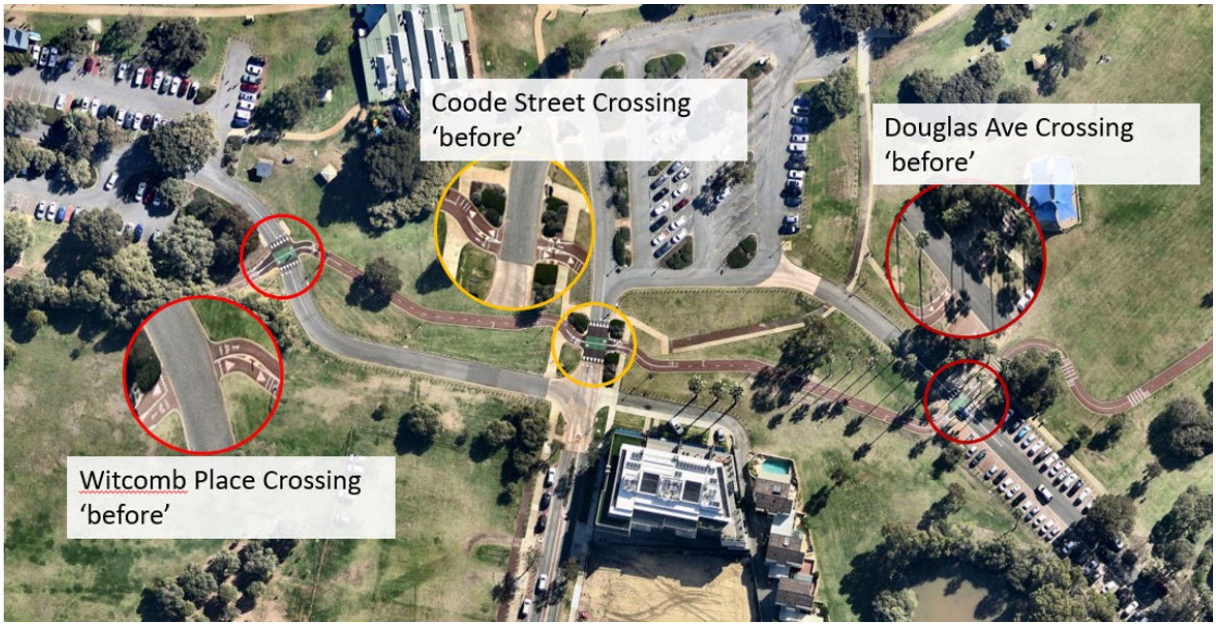

Figure 2 provides a detailed view of the Coode Street crossing before and after the treatment. The other two sites share similar features after the upgrade, such as a raised platform for the bike path and various markings including sharks teeth' to alert drivers of the cyclist priority.

Figure 2: Snapshots of the recorded video at Coode Street before and after the upgrade (vehicle speed detection zones are marked in blue and cyclist/pedestrian detection zones are marked in green)

(a) before (b) after

3. Method

Our video analytics pipeline consists of four main parts: preprocessing, detection, tracking, and postprocessing. Preprocessing includes marking areas of interest (e.g., speed detection zones shown in Figure 2), video stabilization, and projecting pixel distances into real distances. Detection uses a modified Yolo v5 deep learning model (Jocher et al., 2021). to identify and classify objects (e.g., pedestrians, cyclists, and vehicles) in video frames. Two models were used for the detection task: a model trained on the COCOs dataset (Lin et al., 2014) to detect vehicles and bicycles and a model trained on the Crowdhuman (Shao et al., 2018) dataset for people (pedestrians and cyclists). The tracking algorithm Bytetrack (Zhang et al., 2021) follows the detected object throughout the video sequence, assigning a unique ID and tracking its location, size, and trajectory over time. The postprocessing involves constructing an Origin-Destination (OD) matrix for each object class, calculating speeds, estimating give-way compliance, and conducting before-and-after comparisons (see Section 4).

Pixels in the camera view are affected by its perspective so they cannot be directly used for inferring distance-based measures. Therefore, vehicle positions were projected back to a 2D plane by matching a set of key reference points to an aerial photo. Given the fisheye distortion of the camera, relatively small areas were used to minimize the non-linearity of this distortion. Each object's speed is calculated based on its tracked distance within the detection zone and its corresponding time which results in an estimate of the average speed for traversing the zone. It is more accurate than using the distance between edges of the detection zone because objects do not necessarily travel in a straight line and their first and last detections do not always happen right at the edge.

To measure compliance, we identified pairs of objects (e.g., cyclist and vehicle) in both a 'give-way' and a 'priority' zone (Figure 3). If a vehicle in the 'give-way' zone and a cyclist in the 'priority' zone appeared simultaneously, and their trajectories intersected at a point in space, we labelled the pair as either 'complied' or 'failed to comply.' If the cyclist crossed the meeting point before the vehicle gave way, the pair was labelled 'complied'; otherwise, it was labelled 'failed to comply.' The set-up of the zones was based on the 'commonsense' and the understanding of traffic rules, which has subjectivity, but we tried to keep each site's before and after analysis zones consistent so at least the potential biases are consistent.

Figure 3: Give-way compliance zone at Coode Street and an example of non-compliance

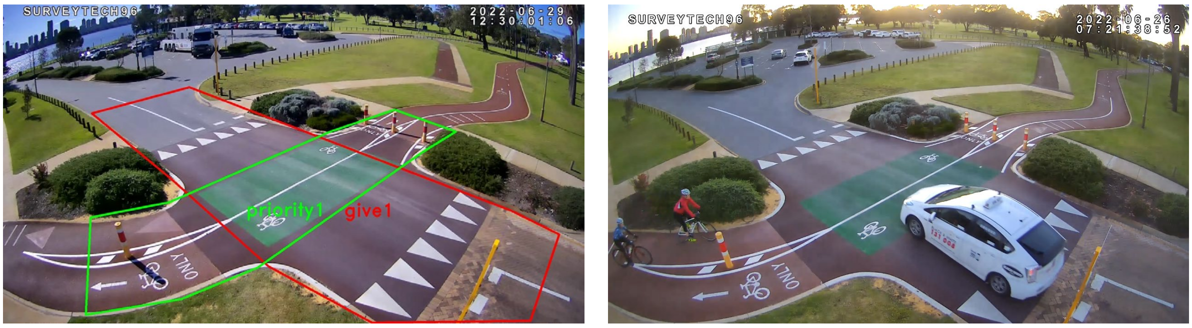

(a) compliance zones (b) a taxi failed to give way to cyclists

4. Results and validation

4.1 Changes in vehicle speeds

All sites had a consistent reduction in vehicle crossing speed after the upgrade. A one-tailed t-statistic was conducted to ensure statistical significance. One tail was used as it was expected that the mean for the 'after' upgrade would be lower, which is confirmed by the statistically significant p-values in Table 1. Meanwhile, there has also been a small increase (between 1 – 3 km/h) in the average cyclist speed across all sites in all survey days (details are omitted from this abridged paper).

Table 1: Before and after average vehicle speeds and p-values for one-tailed t-test

Location (day)	Before Average (km/h)	After Average (km/h)	Change in Speed (km/h)	P-Value
Coode (weekday)	22	17	-5	2.39e-106
Coode (weekend)	17	15	-2	7.24e-223
Witcomb (weekday)	20	15	-5	3.70e-83
Witcomb (weekend)	18	15	-3	8.82e-67
Douglas (weekday)	25	15	-10	1.09e-98
Douglas (weekend)	22	13	-9	1.64e-143

Figure 4 shows that the observed speeds approximate the normal distribution at Coode Street. In addition to the reduction in average speed, the percentage of vehicles above 20 km/h also dropped significantly on both the weekday and the weekend. Other sites have similar patterns.

Figure 4: Vehicle speed distributions at Coode St (the two figures are shown with different scales on the y-axis)

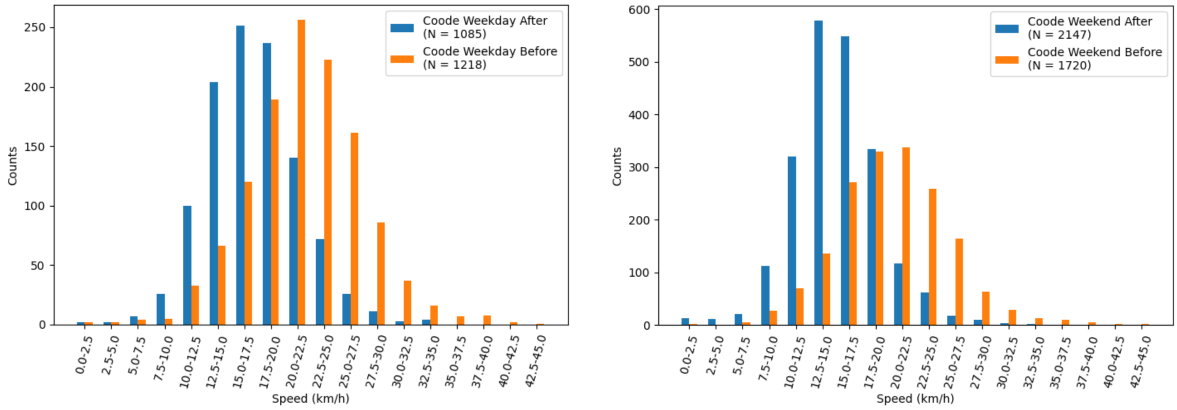

(a) Weekday (b) Weekend

4.2 Give-way compliance

Although the lower vehicle speed is a positive change, there were still a number of vehicles that failed to give way to cyclists, as shown in Table 2. This makes some cyclists hesitant to take priority. Although no before and after comparison can be made since it was vehicle priority before the change, the same zones will be used in future follow-up surveys to ensure consistency.

Table 2: Vehicle's give-way compliance rate after the priority change (Coode Street Crossing)

	Complied	Failed to comply	Total
Weekday	52	12 (19%)	64
Weekend	210	60 (22%)	270

The non-compliance videos were automatically extracted so that the project team can make further observations. Although the give-way compliance zones were designated with subjectivity, most captured cases were deemed reasonable to the team. In addition to the non-compliance cases similar to Figure 3b, many other observations were made from the extracted videos. Figure 5 contains two examples. Figure 5a captures a dangerous moment where both vehicles started moving forward after allowing a few path users to pass, despite more oncoming cyclists. Consequently, the silver vehicle nearly hit the cyclist in front of it. Figure 5b shows an example of a vehicle encroaching into the raised platform while giving way to cyclists, which was a commonly observed behaviour.

Figure 5: Captured moments of bad driver behaviour at Coode Street

(a) An impatient driver nearly hit the cyclist (b) A give-way vehicle encroached into the platform

4.3 Lane discipline of cyclists

We also used heat maps to represent movement patterns. A trajectory is produced for each tracked object, which has a set value for colour and brightness, and their values mix additively when they overlay. A trajectory's colour is determined by where it points, according to the colour ring on the top-left corner (Figure 6). For example, cyclists travelling towards the 2 o'clock direction have a magenta trajectory while the opposite direction has cyan. These two are dominant movements in Figure 6, suggesting most cyclists stayed on the left side of the path. However, there is also a significant amount of bright white stripes that are produced when the two opposite trajectories overlap. They are the results of many cyclists using the opposite lane for overtaking or straightening their paths at the crossing, so their trajectories blend with the other cyclists who were on the correct path. Although the heatmaps have no temporal dimension and those opposite movements may not have happened at the same time, it still highlights the issue of poor lane discipline. The tight curve that was designed to slow down cyclists and scooters might have exacerbated the problem because people prefer travelling in a straight line, especially fast cyclists. This is not an issue most of the time but we have observed relevant close encounters. By comparison, Douglas Avenue Crossing has the least amount of through movements using the wrong lane out of the three crossings. Therefore, it is worth considering revising the design guidelines so that the location and sharpness of the curves can be more carefully chosen considering the observed cyclist behaviour.

Figure 6: Heatmap of cyclist movements for Coode Street Crossing (Sunday 26th June 2022)

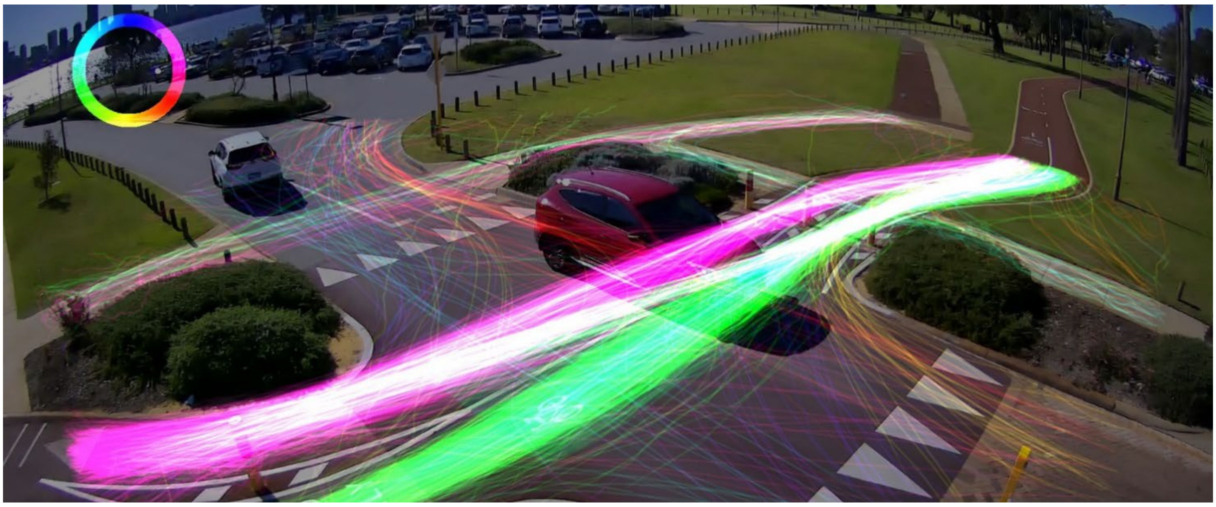

4.4 Validation

Manual validation by counting and classifying users of the path was conducted. Results show that the algorithm achieves reasonable accuracy. Three separate hour-long segments from two recordings (Coode Street Weekday Before and Weekend After) were chosen for validation.

Table 3 shows good overall accuracy and it is best for mid-day. The accuracy in the evening is lower because of sun glare.

Table 3: Accuracy of video analytics (cyclists only)

Time	Manual Counts	VA Count Accuracy
7 AM – 8 AM	372	93%
11 AM – 12 PM	307	97%
5 PM – 6 PM	131	88%
Weighted Average		94%

Manual validation for vehicle counts and their speed is still to be conducted but they are expected to be better than the cyclist results reported in Table 3. This is because vehicles are larger than people and usually visually different from the algorithm's point of view. It means that they are easier to reidentify using features such as colour, if occlusion occurs.

5. Discussion

The first wave of 'after' surveys reported here showed that the change to cyclist priority has reduced the speed of the vehicles crossing the shared path, which is positive. Because multiple modifications were made, it was not possible to differentiate between the effect of the give-way sign versus the effect of the raised platform. More surveys are planned in the future to confirm it as a long-term change. They will also help establish whether vehicle give-way compliance will be improved as drivers get used to the change. Better signage could also be used to remind drivers to look out for approaching cyclists and scooters, especially where their line of sight is blocked by nearby trees.

Although it is not a controlled study, which is often hard to do in traffic analysis, the team argues that post-project evaluations are needed to ensure the desired outcomes are achieved and enable evidence-based decision-making for future projects. It will also be useful to compare the effectiveness of different treatment types.

This project also highlights the value of applying VA in analysing traffic behaviour, especially when it involves active transport modes. Not only it can detect different transport modes with good accuracy, but also provides good insights into their spatial interactions. A project report (about 60 pages long) contains many additional observations that have not been included here.

Based on our experience, the current VA models are highly accurate for humans and vehicles but are less accurate for bicycles. As active transport modes are becoming more popular, future development should focus on increasing the detection accuracy with more and better training data. A more objective method for measuring give-way compliance should also be developed, which ideally can also assess the potential severity of non-compliant cases using speed analysis. Lastly, some field validation of the speed estimation's error range will add to the scientific rigour.

References

Jocher, G., et al. 2020. YOLOv5 by Ultralytics (Version 7.0) [Computer software]. https://doi.org/10.5281/zenodo.3908559

Lin, TY., et al. 2014. Microsoft COCO: Common Objects in Context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds) *Computer Vision – ECCV 2014. ECCV 2014. Lecture Notes in Computer Science, vol 8693*. Springer, Cham. https://doi.org/10.1007/978-3-319-10602-1_48

Shao, S., Zhao, Z., Li, B., Xiao, T., Yu, G., Zhang, X., Sun, J. 2018. CrowdHuman: A Benchmark for Detecting Human in a Crowd. arXiv:1805.00123.

Zhang, Y., Sun, P., Jiang, Y., Yu, D., Weng, F., Yuan, Z., Luo, P., Liu, W., Wang, X., 2021. ByteTrack: Multi-Object Tracking by Associating Every Detection Box. arXiv:2110.06864.