CENTRE FOR
AUTOMOTIVE
SAFETY RESEARCH

THE UNIVERSITY
of ADELAIDE

Video capture and analysis of cyclists using infrastructure in the ACT through machine learning

JRR Mackenzie, G Ponte

CASR REPORT SERIES
CASR202
September 2022

Report documentation

REPORT NO.	DATE	PAGES	ISBN	ISSN
CASR202	September 2022	39	978-1-925971-36-1	1449-2237

TITLE

Video capture and analysis of cyclists using infrastructure in the ACT through machine learning

AUTHORS

JRR Mackenzie, G Ponte

PERFORMING ORGANISATION

Centre for Automotive Safety Research
The University of Adelaide
South Australia 5005
AUSTRALIA

SPONSORED BY

ACT Road Safety Fund
GPO Box 158
Canberra City ACT 2601
AUSTRALIA

AVAILABLE FROM

Centre for Automotive Safety Research
http://casr.adelaide.edu.au/publications/researchreports

ABSTRACT

The implementation of infrastructure to reduce traffic conflicts and improve road safety for cyclists is critical. However, potential benefits resulting from strategic interventions can only be achieved if there is corresponding compliance or good utilisation of that infrastructure. Traditional methods for evaluating cycling behaviours and interactions with the traffic system either involve expensive roadside observations, which can influence cycling behaviours through the observer effect, or by using induction or pneumatic tube system which can only yield count information and cannot be used in mixed traffic situations. This study used long duration portable video cameras to, somewhat covertly, record mixed road traffic, with the captured video subsequently being processed with bespoke machine learning software. This innovative process required negligible set-up or manual processing time and a researcher was only required to do a desktop analysis on 6% of the total traffic video recorded. Agencies requiring evaluations of cyclist interactions with specific infrastructure or behaviours at specific locations could use this rapid and cost effective process to assist with data collection and analyses to inform or optimise their cycling safety strategies.

KEYWORDS

Traffic video analysis, machine learning, cyclist detection, cyclist safety

Summary

An important step in continuously improving road safety for cyclists (as vulnerable road users) is to ensure that the safety infrastructure that has been provided for their protection is performing well. This project sought to trial and showcase a video-based system, augmented with machine learning, to analyse cyclists' use of infrastructure in the ACT.

A suitable camera system, that was battery powered for deployment in locations of interest and able record at a quality required for use with a previously developed machine learning tool, was identified. Four sites in the ACT where cyclist interactions with road infrastructure were of interest were selected in consultation with cycling safety stakeholders.

After all relevant permissions and approvals were obtained, the camera systems were installed at the four sites to record during daylight hours for a period of seven days. The 378 hours of recorded footage was then processed by the machine learning tool and reduced in length by almost 95% to summary videos containing only the periods where cyclists were detected in the camera scene.

Image sets of the detected cyclists were also generated, and these were subsequently filtered and analysed by a researcher. This analysis consisted of categorising each detected cyclist's perceived gender and helmet use.

Analysis of the summary videos enabled the categorisation of various cyclist behaviours of interest, with respect to their use of safety infrastructure, particular to each of the sites. Some of these behaviours included cyclists' interaction with designated painted crossing point infrastructure, bicycle turn holding boxes, and bike lanes that terminate suddenly.

The overall system was considered to have worked well. Camera setup time was negligible and analysis of the generated image set was able to produce a rapid audit of the number and characteristics of cyclists in a particular area. The summary videos were useful in enabling an understanding of cyclist behaviours and interactions with road infrastructure at each site. These insights highlighted several instances where there may be the potential for improvements to cyclist safety.

It was noted that the formal analysis of the summary videos could take some time. This could be alleviated somewhat by recording for shorter periods at locations where there are a lot of cyclists. Additionally, it would be important to carefully consider the goals of any future analyses to determine whether the formal categorisation process was necessary or whether the summary videos would be sufficient.

Contents

1 Introduction 1

2 Identification of a suitable camera system 3

3 Selection of locations for surveillance 6

4 Installation of cameras for data collection 13

5 Video processing 15

6 Results 18

6.1 Site A 18

6.2 Site B 21

6.3 Site C 24

6.4 Site D 27

7 Discussion 29

Acknowledgements 32

References 33

Appendix A – Stakeholder engagement letter 34

Appendix B – Camera signage 35

1 Introduction

Cyclists are classified as vulnerable road users when riding on public roadways due to the speed and mass differential between themselves and vehicles, as well as their lack of physical protection in the event of a collision. As such, road authorities often install specific types of road infrastructure that are designed to assist cyclists in avoiding potential conflicts that could result in collisions. Some examples of this are bike lanes, separated paths, or dedicated turning zones and crossing points.

However, the benefits of these infrastructure treatments can only be maximised when cyclists use them – and use them correctly. On occasions, there can be a disconnect between infrastructure that has been installed with good intentions and the way in which cyclists prefer to utilise the road network. For example, a cyclist may decide to ride on the road rather than use a separated path because they prefer not to interact with pedestrians and dog walkers or prefer the quickest transport route. In addition, the proper use of some types of infrastructure, or even the presence of the infrastructure, is not always obvious to cyclists. Because of the complex nature of the road system, there are also many locations where cyclist infrastructure ends suddenly, or cannot be installed to the optimum standard, and this can result in unexpected situations for cyclists to negotiate.

With the aim of continuously improving the safety of the road network, it is important that road authorities collect information about how cyclist safety infrastructure is utilised. This enables an understanding of the performance of various types of infrastructure, whether cyclists are interacting with infrastructure correctly, and what residual safety issues may remain.

To address this need, the Centre for Automotive Safety Research (CASR) and the Australian Institute for Machine Learning (AIML) developed a system that can process recorded video of traffic, detect all cyclists in the video (using a faster RCNN object detection algorithm) and produce a summary video showing only video segments containing the detected cyclists (Ponte et al., 2014). In addition, the system can produce an image set of the cyclists that were detected.

There are many benefits to using video to capture cyclists in traffic. Traditional traffic counters often use pneumatic tubes to count cyclists, but this cannot be deployed in mixed traffic situations. Induction loops designed to specifically count cyclists are also used but suffer from the same issue. Pneumatic tubes and induction loops are installed at a fixed point so cannot be moved and cannot give any characteristics about the cyclists riding on those paths (such as helmet use or bicycle type). Video systems for counting cyclists exist, but the resolution of current systems do not facilitate automatic detection and counting of cyclists.

While it is possible to observe cyclists by having a researcher on site, cyclists may change their behaviour when they know they are being watched. This is known as the 'observer effect' and means that the observed behaviour may be modified rather than naturalistic. Furthermore, manual observation is error prone, limited to what is observed on the day, labour intensive, and costly.

The machine learning software tool that has been developed by CASR/AIML can be used to count cyclists in mixed traffic, track cyclists and their travel path, export images of the detected cyclists for characterisation (gender, helmet use, hi-vis use, etc.), and the videos can be further analysed to examine cyclist behaviours and interactions with motorised traffic or road infrastructure. The system also eliminates the observer effect and the need for a researcher to spend hours reviewing video footage while waiting for cyclists to come into view. Instead, the researcher can work more efficiently by concentrating on reviewing cyclists' behaviour and coding variables/events of interest from the condensed videos that always contain a cyclist.

This project sought to demonstrate the capabilities of using machine learning to analyse video that has been recorded with the intent of capturing the behaviour of cyclists interacting with infrastructure in the ACT.

Section 2 describes the process that was undertaken to identify a modular, battery powered, camera system that could be deployed to record footage of cyclists in traffic or on shared paths over several days. In Section 3, details regarding consultation with cyclist safety stakeholders, in order to select multiple locations where cyclist interactions with road infrastructure would be of interest, are provided. Then, in Section 4, the process of obtaining the necessary permissions for the deployment of cameras for recording in public areas in the ACT are provided, along with a summary of the subsequent deployment activities. Processing of the recorded video using the machine learning system and the analysis of the generated outputs, with the goal of documenting the characteristics of passing cyclists and interactions with road infrastructure, are described in Section 5. The results of these analyses are presented in Section 6 followed by a discussion of the project outcomes in Section 7.

As with any project that involves the use of cameras recording in public areas, there were ethical and privacy issues to consider. The activities that were undertaken during this project received approval from the University of Adelaide's Human Research Ethics Committee (H-2018-102) and this approval was valid for the entire duration of the project. Some specific requirements of the ethics approval were that no audio would be recorded and that any published images of a cyclist, driver, or other road user would be blanked out or blurred so that the individual could not be identified.

2 Identification of a suitable camera system

The first stage of this project was to identify a portable, robust, and reliable video camera system, capable of recording several days of video footage.

The requirements for a suitable video camera system were defined as follows:

- Capable of recording at least 5 days of daytime video footage, with specific regard to battery capacity and memory storage capacity;
- Robust in weather such as rain, heat, cold;
- Records video of sufficient quality that can be successfully processed by cyclist detection machine learning software;
- Able to be set to automatically record only in daylight hours;
- Can be safely and easily attached to roadside infrastructure typically found in the ACT; and
- Compliant with any requirements for privacy provisions in the ACT.

Initial development and testing was undertaken using GoPro video cameras. An attempt was made to incorporate a GoPro camera, with a suitable power supply, into a weatherproof enclosure. However, it was quickly discovered that the GoPro camera would over-heat (and shut off) after a few hours and was not reliable for long duration recording.

Efforts were then shifted to developing a bespoke camera system using a microcomputer (Raspberry Pi) fitted into an enclosure and connected to a suitable video camera. Progress on this task was good and a functioning prototype was developed that was able to record video for 72 hours at HD quality.

Concurrently, contact was made with a traffic survey company (Surveytech) who were in the process of developing their own modular traffic survey camera systems that were likely to meet the needs of this project.

Testing of the bespoke camera system developed by CASR found it to be suitable for the project, but the system set-up process was complex and prone to errors. While it would have been possible to reduce the complexity of setup with further development, it was acknowledged that the Surveytech camera system was just as capable (while being easier to use) and so attention was focussed on the testing of that system.

A test unit from Surveytech was acquired (on-loan) with a battery capacity sufficient to record 30 hours of continuous video. Units with larger battery capacities were available for purchase, but only the smaller unit was available for loan at the time.

The Surveytech unit consists of a camera, digital video recorder (DVR), and battery built into a weather-proof enclosure (rated at IP68) that features pole mounting attachment points. The unit can be connected to a computer or smartphone via WiFi and a live view can observed. The video camera system can be set to record at specific hours of the day and can be powered down automatically between scheduled recording times to conserve battery life. The quality of the video is 1920 x 1080 pixels and the bitrate can be adjusted from 600 to 8000 Mbps, while the frame rate can be adjusted from 5 frames per second (fps) up to 30 fps.

The video quality was tested with recordings at 1000, 2000 and 3000 Mbps as well as at 20 fps, 25 fps and 30 fps. As a result of this testing, it was found that a bit rate of 2000 Mbps at 25 fps produced a video file of sufficient quality for machine learning software to consistently detect cyclists. With these settings, a one-hour recording produced a video file consuming around 850 Megabytes of storage.

The ability of the unit to record for only daylight hours was tested by scheduling video to be recorded from 7am to 6pm. The unit was left for three days and then inspected. As expected, the unit collected 11 hours of footage on day one and two, then a little over 3 hours of footage on day three before the battery was depleted. This represents a total of 25 hours of recorded video and indicates that approximately 5 hours of battery time was depleted while the unit was in standby mode.

The camera was trialled for a day at one of the busier intersections in Adelaide, the intersection of Frome Road and North Terrace. An aerial view from Google maps is presented in Figure 2.1 which also shows the camera was mounted on a traffic pole approximately 3 metres above ground level. Figure 2.2 shows the 135-degree field of view captured by the camera with subsequent processed by machine learning software correctly identifying two cyclists within the traffic scene.

After testing the Surveytech camera system, it was deemed suitable for the goals of the project and sufficient funding was available to purchase four 130-hour capacity units with single cameras (see Figure 2.3). These units enabled four sites in the ACT to be monitored in daylight hours for a period of 5-10 days.

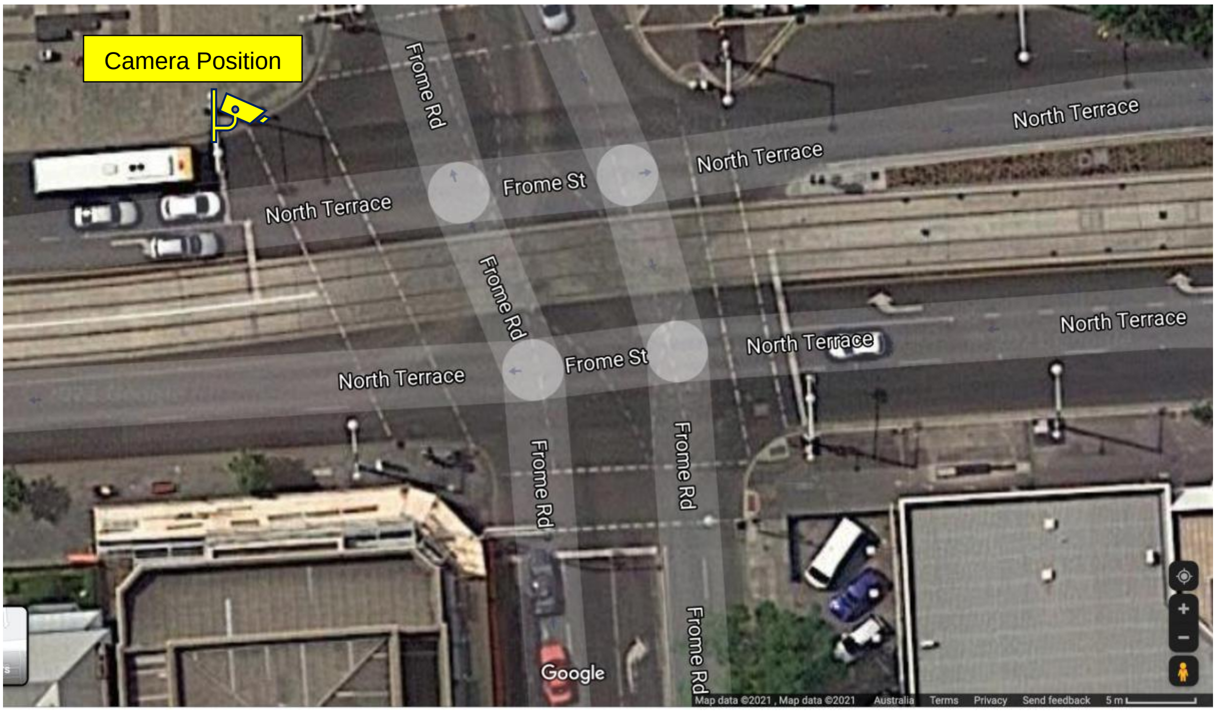

Figure 2.1
Aerial view of Frome Road and North Terrace, showing the intersection layout and camera position

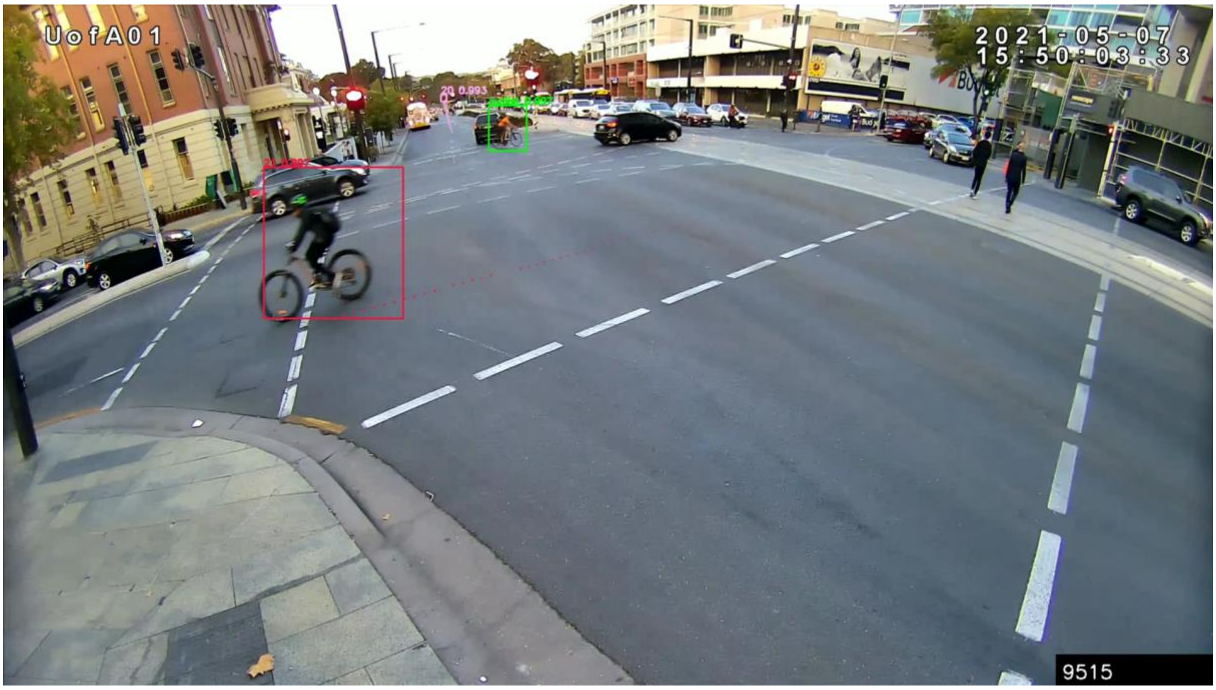

Figure 2.2
Captured camera field of view of traffic on Frome Road and North Terrace, with two cyclists being correctly identified by the machine learning software

Figure 2.3
Surveytech 130-hour modular traffic surveillance camera system

3 Selection of locations for surveillance

In order to determine suitable locations for the installation of the cyclist surveillance cameras, a group of ACT cyclist safety stakeholders were consulted. With the assistance of the Road Safety Policy Team at Transport Canberra and City Services (TCCS), the following stakeholder groups were identified:

- ACT Veterans' Cycling Club Inc
- Pedal Power ACT
- Amy Gillett Foundation
- Cycling and Walking Australia New Zealand
- TCCS Communications Team
- TCCS Active Travel Team
- TCCS Bicycle Advisory Group

A letter was sent to each stakeholder on 17 August 2021 outlining the objectives of the project and requesting their assistance to recommend locations within the ACT where the surveillance of cyclist interacting with infrastructure/traffic would be of interest. A copy of the letter is shown in Appendix A.

Responses were received from two stakeholders and together they recommended 30 locations, listed in Table 3.1 and shown in Figure 3.1 below.

As there were four cameras available for use with the project, the next task was to select four locations from the list. To assist with this process, each location was reviewed using Google Street View to identify the existing road/cyclist infrastructure, categorise the type of cyclist behaviours that were to be monitored, and explore the availability of street lighting poles suitable for camera mounting. Additionally, Strava cycling heat maps were consulted to identify the likely frequency of cyclists passing through each location as high, medium, or low.

Because this project is serving as a trial of the applicability of using machine learning to augment video surveillance and analyse cyclist behaviour across a range of scenarios/situations, it was deemed appropriate to select locations that would monitor a variety of cyclist behaviours. Locations with the greatest frequency of passing cyclists were prioritised during the selection process and locations where no suitable lighting poles were available for mounting cameras were excluded from consideration.

Based on the above criteria, locations 4, 5, 10, and 18 were chosen for surveillance with the video camera systems. Hereafter, for reasons of convenience, these locations are referred to as Site A, Site B, Site C, and Site D respectively. Images of the selected locations/sites are shown below in Figures 3.2 to 3.5.

Table 3.1
Stakeholder recommended locations for surveillance of cyclists' use of infrastructure (selected locations highlighted in grey)

Location number	Longitude	Latitude	Stakeholder description of location and interest in infrastructure usage	Cyclist travel
1	149.12527	-35.28667	The lake side of the path bridge over Parkes Way from the ferry terminal	N/A
2	149.12534	-35.28465	Marcus Clarke Street with a tricky junction	High
3	149.10802	-35.25864	Fairfax and Dryandra at the top end of O'Connor, travelling towards the city. The poor-quality cycle path peters out just as you reach the bus shelter and it's cross country from there to Dryandra St.	High
4	149.09464	-35.31439	Adelaide Ave and Cotter Rd (city bound). A lot of fearless riding goes on here during morning peak hour - made worse a year ago when the Cotter Rd was given its own lane onto Adelaide Ave.	High
5	149.12966	-35.27566	Crossing Northbourne at Cooyong/Barry Drive. Zig-zag tram line crossings, poorly-timed pedestrian lights.	High
6	149.13220	-35.18504	Gungahlin Town centre, Hibberson Street where it is narrowed into one lane and cyclists are forced onto the sidewalk. Cyclists are forced to either ride in among the vehicles or with pedestrians in narrow busy spaces.	High
7	149.12610	-35.18347	Gungahlin, Gundaroo Drive near petrol stations, where the bike path crosses and people head into the shops, gym etc. Very busy, lots of vehicles and cyclists, but no safe way for cyclists to navigate the area.	High
8	149.09230	-35.36450	Beasley St north/ Mawson Drive/Athllon Drive Intersection. Signalised intersection with main path (C4) crossing, cyclists often cut the corner at the slip lane islands, cyclists may not obey crossing 'green man' even though automatic activation, fast downhill approach on one side .	Medium
9	149.09275	-35.38432	Athllon Drive and Sulwood Roundabout. Major path (C4) crosses at large roundabout, difficult to judge speed of vehicles and direction of travel, fast downhill approach on one side.	Medium
10	149.08363	-35.42784	Barr-Smith Ave/Drakeford Dr Intersection. Drakeford Drive is busy arterial, Barr-Smith main collector road, on-road cycle lanes on Drakeford Dr, difficult for riders to see vehicles on Barr-Smith as they approach crossing, no pedestrian crossing, drivers have priority. On road cyclists must cross the left turn lane, can be confusion with who gives way, slip lanes with merge lane on major road means vehicles turn left at speed.	Medium

Location number	Longitude	Latitude	Stakeholder description of location and interest in infrastructure usage	Cyclist travel
11	149.16530	-35.22494	Hwy at the fed Hwy roundabout which is likely to increase traffic, in particular trucks along Antill st, making it even more of a black spot.	High
12	149.12570	-35.23257	Intersection of Barton Hwy and Ellenborough st West bound. Eastbound traffic turning into EB taking risky gaps, tuning in front on oncoming traffic and cyclists.	High
13	149.13484	-35.24415	Motorists think they are being courteous to each other by leaving a gap and allowing a vehicle or two to enter from a side street, but when your oncoming and approaching a side street motorists often “don’t see” you. Worst sports are, Nth bourne and Swinden St Southbound.	High
14	149.12728	-35.28958	Commonwealth Ave and Albert and Corkhill St Nthbound. When oncoming and approaching a side street motorists often “don’t see” you.	High
15	149.02694	-35.32542	Cotter Rd from Eucumbene Dr to where you start to descend down to the Cotter recreation area, a proper hard shoulder/bike lane on both sides.	Medium
16	149.05358	-35.32851	Dixon Drive Holder turning onto Streeton Drive Northbound	High
17	149.07803	-35.42328	Bike lane disappears along Drakeford drive	Medium
18	149.07446	-35.40273	Bike lane disappears along Drakeford drive	Medium
19	149.06915	-35.39198	Bike lane disappears along Drakeford drive	Low
20	149.12383	-35.27444	Barry Drive city bound where bike lane is pushed off road and merged with footpath in a section where cyclists are going 60kph	High
21	149.02694	-35.32542	Cotter road, riders coming from Stromlo going straight and vehicles from Eucumbene pulling out. Lots of shadows so drivers maybe looking past the cyclist.	Medium

Location number	Longitude	Latitude	Stakeholder description of location and interest in infrastructure usage	Cyclist travel
22	149.05102	-35.20394	Intersection of Spalding with Kingsford Smith drive - it's a stop sign and I've nearly been taken out twice because at tradie-o-clock the sign reads as "give way but don't look for cyclists".	Low
23	149.06136	-35.19864	Magrath Cres intersection with Kingsford Smith (southern end). There's a visibility/speed/timing issue with that one, and cyclists end up being behind the A-pillar as vehicles approach the stop sign.	Low
24	149.18726	-35.31213	Piallago ave and roundabouts outside Brindabella park offices. Bike lanes disappear, and vehicles change lanes into the left when it diverges into two lanes. Vehicles change lanes to the left lane without indicating and there is no overt way for cyclists to signal heading onto the road as the verge ends. Vehicles then try and overtake bicycles while going through the roundabout.	Medium
25	149.16557	-35.22547	Antill Rd between Federal Hwy and Knox St, Watson. No bike lane road.	High
26	149.09699	-35.32429	Change in traffic condition from a bike lane, to no bike lane along Kent street, Deakin.	Medium
27	149.10534	-35.31288	Merging onto Adelaide Ave from Hopetoun circuit travelling south. Road furniture before entering a bike lane.	High
28	149.08980	-35.17841	Kuringa Dr (top of Belco that runs to Barton Hwy). Heavy and aggressive traffic most of the time from early mornings on.	Low
29	149.09563	-35.22950	Baldwin dr and Maribyrnong ave (and maybe Ellenborough st) that runs through Kaleen area. Reasonable wide roads and two lanes in part but no bike lanes and heavy traffic.	Low
30	149.16644	-35.22751	Antill st Watson (from end of Hackett to the federal Hwy roundabout. This is ok when in an early morning bunch but with increasing housing in north watson and the use of Antill st as a quick alternative to the fed Hwy it has become busy. There is no bike lane or sealed side next to the white line.	Low

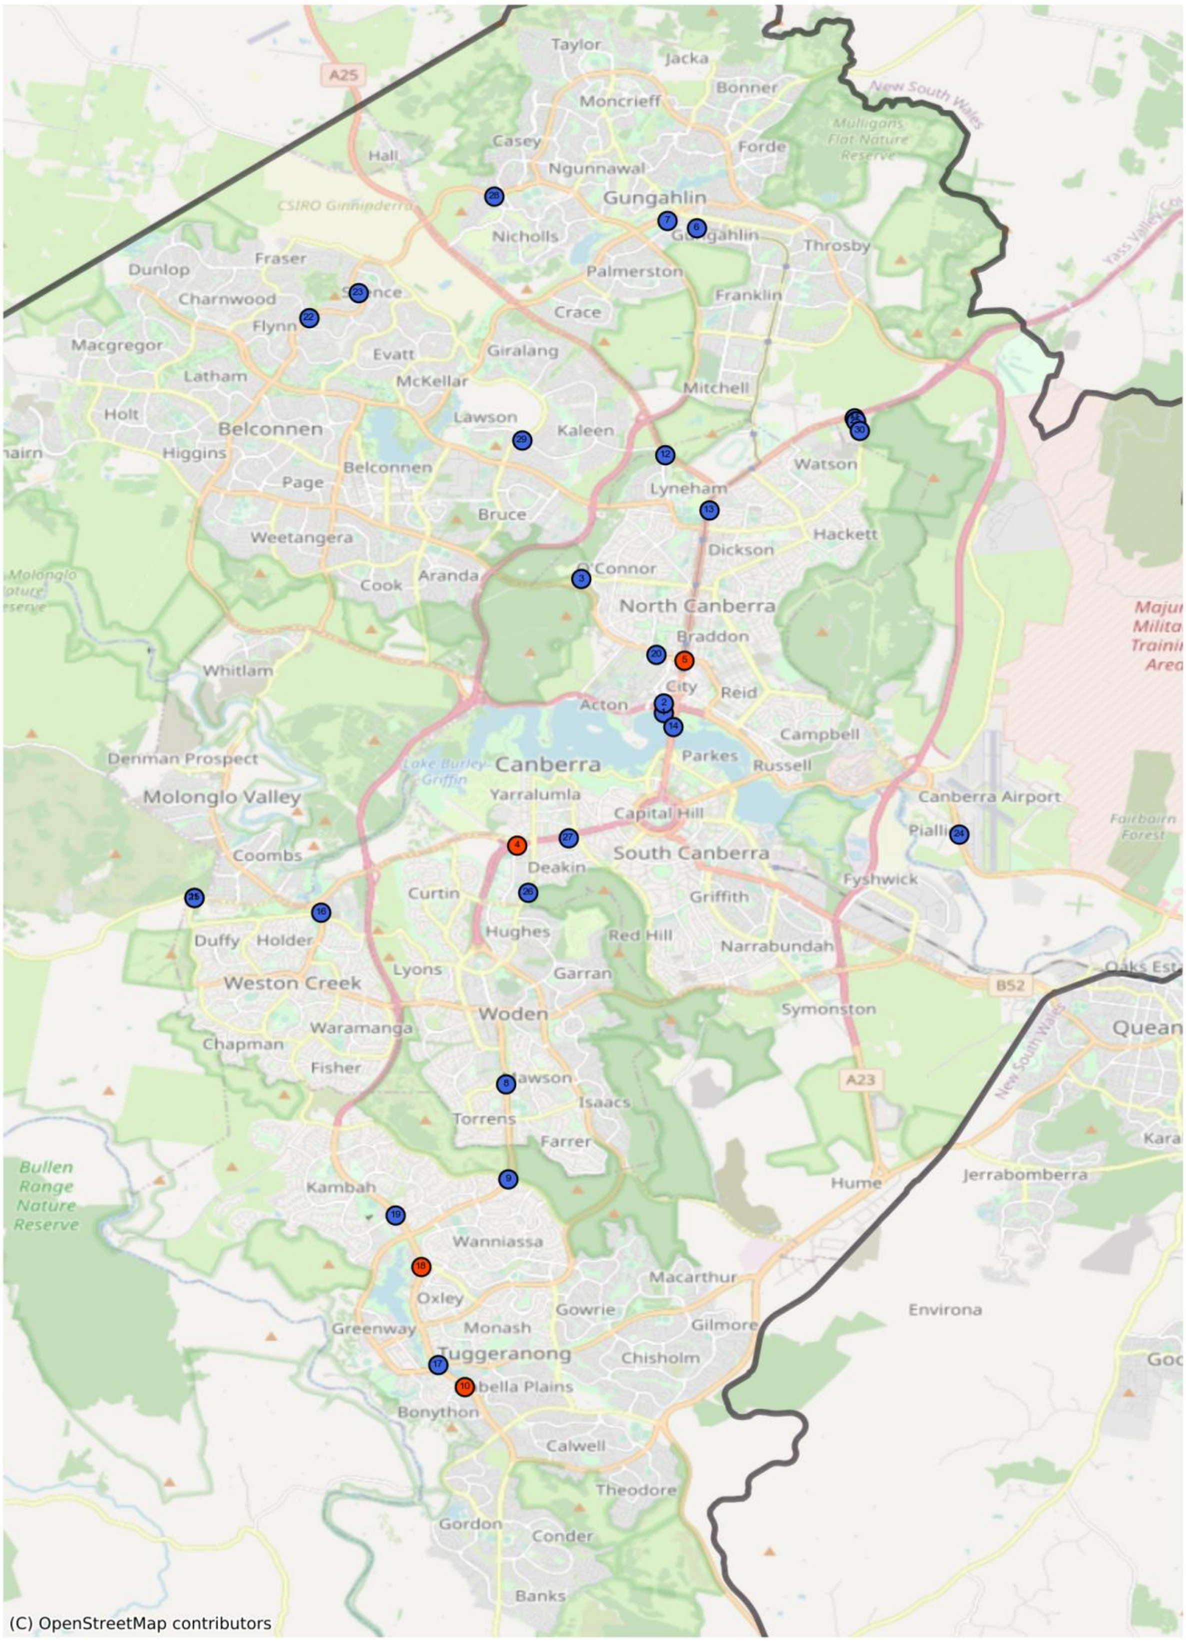

Figure 3.1
Stakeholder recommended locations for surveillance of cyclists' use of infrastructure
(selected locations shown in orange)

Figure 3.2
Birds-eye view and street view (Google Maps) of location 4, referred to as Site A

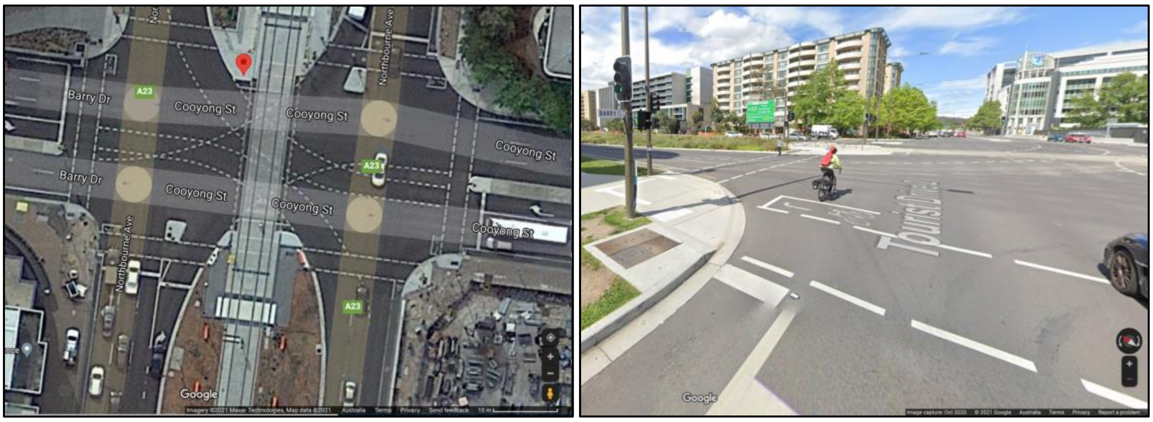

Figure 3.3
Birds-eye view and street view (Google Maps) of location 5, referred to as Site B

Figure 3.4
Birds-eye view and street view (Google Maps) of location 10, referred to as Site C

Figure 3.5
Birds-eye view and street view (Google Maps) of location 18, referred to as Site D

4 Installation of cameras for data collection

To ensure the correct protocols and procedures were followed prior to the installation of cameras, a meeting was held with representatives from several departments of the ACT Transport Canberra and City Services (TCCS). It was established that, to obtain the necessary permissions required by the project to install cameras for monitoring public spaces in the ACT, there were several steps that would need to be completed.

First, at each of the four selected sites, the specific streetlight poles that were being sought for use in mounting cameras were sent to TCCS. This allowed for the streetlight pole identification number to be determined and the availability of the pole (for use as a camera mount) to be reviewed. As a result of this process, one streetlight was identified as being with the Light Rail Corridor and an alternative streetlight at the same site was selected.

Second, a temporary traffic management (TTM) plan for the installation of the cameras at each site was developed and submitted by a local contractor RD Gossip Pty Ltd. The TTM plans (registered number 10092) were authorised for the period 9 December 2021, 7am to 31 March 2022, 5:00pm.

An application for the use of public unleased land was then submitted. This application noted the land would be used for the temporary installation of self-contained battery powered video recording devices, with accompanying signage (as per ACT Government Code of Practice for Closed Circuit Television Systems), on four designated light poles at 3m height for a period of one week. A permit for the designated land use was granted for the period 4 February 2022, 9am to 14 February 2022, 5pm. For contingency purposes (e.g. inclement weather), an extension to the permit duration was requested and granted until 1 March 2022, 5pm.

Finally, an appropriate sign, which complies with the ACT Code of Practice for CCTV use, was developed (Appendix B). Four of these signs, to accompany the four cameras, were manufactured at A4 size on 5mm PVC plastic.

After completing these steps, all the necessary approvals and requirements had been met for the installation of cameras in the ACT.

The installation of the cameras at the four locations in the ACT was contracted to RD Gossip Pty Ltd. The cameras were sent to RD Gossip and an information session was held in order to explain the goals of the project, the proper setup/use of the cameras, and the locations where the cameras were to be installed.

After checking that the weather forecast was favourable for the upcoming week, the cameras were installed at the four nominated sites on Friday 4 February 2022. At each site, it took two contractors less than 30 minutes to set up and install the video cameras. The cameras were set to record from 6:30am to 8:00pm each day for seven consecutive days, representing approximately 100 hours of video from each location. The mounting locations and viewpoints of the cameras are shown in Figures 4.1 – 4.4 below.

Over the course of the week, the cameras were checked to ensure they were working properly and had not been damaged.

The cameras were de-installed on Monday 14 February 2022. After checking to ensure they had successfully recorded the appropriate amount of footage, they were returned to Adelaide for analysis.

Figure 4.1
Position of camera (marked with arrow) and the actual camera view of the bike lane and road for Site A

Figure 4.2
Position of camera (marked with arrow) and the actual camera view of the bike lane and road for Site B

Figure 4.3
Position of camera (marked with arrow) and the actual camera view of the bike lane and road for Site C

Figure 4.4
Position of camera (marked with arrow) and the actual camera view of the bike lane and road for Site D

5 Video processing

After the four cameras systems were recovered from Canberra, all recorded video files were downloaded and backed up.

At each of the four monitored locations, the cameras recorded 13.5 hours of footage per day between the hours of 6:30am and 8:00pm. While nine days of video was recorded, the analysis was restricted to the seven-day period of Saturday 5th February to Friday 11th February 2022. Over this period, 94.5 hours of footage were recorded at each location, giving a total of 378 hours.

This captured video footage was then processed through the machine learning system to isolate the periods where a cyclist was present. The system achieves this by examining every 5th video frame and applying a trained computer vision algorithm to search for areas that it believes contain a cyclist. Once the system has examined all frames it compiles a summary video by extracting the original video frames where a cyclist was detected as well as a specified number of frames either side of each cyclist detection. All these frames are then complied into a single condensed video that represents a summary of all cyclist movements through the scene. Each summary video shows detected cyclists coming into view encapsulated with a bounding box and an indication of the nominal path the cyclist is travelling along. In addition to the summary video, the machine learning system also outputs images of each cyclist that it detects.

Once the footage from each site was processed the amount of video that required analysis was reduced as shown in Table 5.1 below.

Table 5.1
Change in length of recorded footage as a result of processing through the machine learning system

Location	Original video length	Summarised video length	Percentage reduction
Site A	94:30:00	5:23:25	94.3%
Site B	94:30:00	11:27:43	87.9%
Site C	94:30:00	3:29:30	96.3%
Site D	94:30:00	1:26:24	98.5%
TOTAL	378:00:00	21:47:02	94.2%

A researcher then analysed the images and summary videos to count the number of cyclists, gather details about their characteristics, and categorise their behaviours.

The first step was to view each of the images from each site to obtain a quick count for each day and collate characteristic aspects such as gender and helmet usage.

In instances where a recorded cyclist was riding abreast of another cyclist or was momentarily visually obscured (e.g. by a passing vehicle), the machine learning system would sometimes mistakenly believe it was observing a new cyclist rather than the same cyclist. This would result in multiple images of the same cyclist being generated. Additionally, the system could also occasionally mistakenly identify a cyclist were there was none. This was usually the result of mis-classifying a motorcyclist, but also sometimes a person in a ‘cyclist-like’ pose or a vehicle wheel combined with some other background shape. Some examples of this are shown in Figure 5.1.

As such, the analysis of the images involved some manual checking to filter out duplicates and any instances of non-cyclists. Even with this manual checking, there may still be some duplicates that are not removed, due to the small size of the images. However, the image review process is relatively quick, with a researcher able to process around 500 images in an hour, which allows for a rapid count and classification of cyclists passing through the recorded area.

Examining the summarised videos, and classifying various cyclist features or behaviours was much more time consuming, and could take up to one minute per cyclist, depending on the features being classified. On some occasions the length of the summary video segment was not sufficient to examine the entire transit of a cyclist through the scene and so the original video was examined. This often occurred when a cyclist was travelling slowly or when a cyclist had to stop and wait for an extended period of time.

At each site there were cyclist behaviours and/or interactions with infrastructure or vehicles that were highlighted as being of interest to the stakeholders that were consulted. These behaviours and interactions were categorised at each site by the researcher analysing the video. As this project was also serving as a demonstration of what data could be obtained using a video system, there were additional cyclist features or behaviours that were categorised at each site as well. The researcher examining the summary videos categorised the following cyclist behaviours and features at each site:

Site A

- Cyclist use of the specific crossing point
- Crossing behaviour of those that do not use the crossing point
- Instances of conflicts between crossing cyclists and vehicles
- Cyclist bicycle type
- Cyclist clothing type
- Cyclist conspicuity
- Cyclist use of rear light

Site B

- Cyclist use of bicycle turn holding boxes
- Instances of vehicles turning left and conflicting with cyclists in the bicycle lane
- Behaviour of cyclists in navigating light rail infrastructure

Site C

- Cyclist use of the specific crossing point
- Crossing behaviour of those that do not use the crossing point
- Instances of conflicts between cyclists and vehicles at the T-intersection
- Paths taken by cyclists

Site D

- Behaviours of cyclists in response to bike lane ending
- Instances of conflicts between cyclists and left turning vehicles
- Cyclist conspicuity
- Cyclists use of rear light
- Paths taken by cyclists

Figure 5.1
Examples of erroneous detections: pedestrian handrail and vehicle wheel combination, kick-scooter rider, and motorcyclist

6 Results

After reviewing and filtering the images generated by the machine learning system, there were over 4,000 cyclists detected passing through the four sites over the seven days of observation. The number and proportion of cyclists at each site that passed through on each day, were perceived as male or female, and were wearing helmets is shown in Table 6.1 below. Note that the image categorisation process was not perfect, due to factors such as human error or small image sizes, and the values shown in Table 6.1 were refined during the process of watching the summary videos. However, the change in the values as a result of this refinement was no larger than 10% for any of the sites or features of interest.

Encouragingly, helmet usage was generally 90% or above at each of the sites. Cyclists perceived as male by the researcher viewing the images made up about three quarters of the sample. The numbers of cyclists passing on each day of the week were unique for each site, with some showing greater numbers on weekdays while others had higher numbers on the weekend days.

Table 6.1
Summary of features based on cyclists detected from each of the four sites

Feature	Site A		Site B		Site C		Site D	
	Number	Percentage	Number	Percentage	Number	Percentage	Number	Percentage
Day (Date)								
Saturday (5/2/22)	102	14.7%	232	8.6%	165	15.7%	7	20.6%
Sunday (6/2/22)	75	10.8%	272	10.1%	214	20.4%	12	35.3%
Monday (7/2/22)	104	15.0%	394	14.6%	136	13.0%	3	8.8%
Tuesday (8/2/22)	113	16.3%	444	16.4%	118	11.2%	1	2.9%
Wednesday (9/2/22)	117	16.9%	448	16.6%	149	14.2%	5	14.7%
Thursday (10/2/22)	95	13.7%	437	16.2%	144	13.7%	3	8.8%
Friday (11/2/22)	87	12.6%	478	17.7%	123	11.7%	3	8.8%
Perceived gender								
Male	578	83.4%	1,904	70.4%	877	83.6%	26	76.5%
Female	89	12.8%	563	20.8%	156	14.9%	7	20.6%
Unknown	26	3.8%	238	8.8%	16	1.5%	1	2.9%
Helmet								
Worn	690	99.6%	2,433	89.9%	982	93.6%	34	100.0%
Not worn	3	0.4%	191	7.1%	40	3.8%	0	-
Uncertain	0	-	81	3.0%	27	2.6%	0	-
TOTAL	693	100.0%	2,705	100.0%	1,049	100.0%	34	100.0%

The summarised video from each site was then viewed by a researcher who categorised certain aspects of interest specific to each site. These additional aspects are described in the subsections below.

6.1 Site A

Site A was the road merge of Yarra Glen Road and Cotter Road. Yarra Glen Road is an 80 km/h arterial road with a dedicated bike lane, however, as Cotter Road merges with Yarra Glen Road, cyclists are expected to exit and merge across to Cotter Road to continue riding along Yarra Glen Road. The painted cyclist crossing point infrastructure encourages riders to make a sharp left turn which would require them to slow considerably. This location was described by stakeholders as having "a lot of fearless riding going on". A camera was placed at the site to observe how cyclists navigate the provided crossing point infrastructure.

The behaviour of each cyclist traversing the merge point at Site A was classified into three categories: (1) An appropriate cross, where the cyclist crossed at the designated crossing infrastructure; (2) an early merge, where the cyclist passed the designated crossing point then merged across the traffic lane into the bicycle lane prior to the second marked road chevron; and (3) a late merge, where the cyclist travelled beyond the second painted chevron marking then merged across the traffic lane into the bicycle lane at some later point. Examples of each of these merging behaviours are shown in Figures 6.1 – 6.3 below. It was noted that cyclists who crossed later had less protective buffer from vehicles on either side.

While categorising each cyclist's crossing behaviour, the presence and reaction of any nearby vehicles was also noted. Cyclists were categorised as crossing in the absence of any vehicles, giving way and then crossing behind a vehicle, or crossing in front of a vehicle. A cyclist was considered to have crossed in front of a vehicle when it was observed that a vehicle entered the video scene shortly after the cyclist had crossed. Instances where a vehicle stopped to allow a waiting cyclist to cross were also noted.

As shown in Table 6.2, only 10.8% of cyclists crossed Cotter Road using the designated crossing point (some first stopping and waiting for vehicle to pass), with most (66.4%) choosing to maintain their original trajectory (and travel speed) along Yarra Glen Road and merge across somewhere between the crossing point and the second painted chevron. The remainder (22.8%) chose to merge across even later, somewhere after the second painted chevron.

Most cyclists (69.0%) were able to cross in the absence of any vehicles approaching from Cotter Road, while 25.7% crossed after a vehicle had passed. However, 5.1% of cyclists crossed in front of a vehicle. Twenty-two of these events were deemed risky (in the opinion of the researcher analysing the video) and five of these resulted in a visually apparent defensive manoeuvre by a vehicle (e.g. braking). There were two instances where a cyclist was waiting and a vehicle stopped to allow them to cross in front.

Table 6.2
Characteristics of the cyclists and behaviours at Site A

Behaviour	Number	Percent
Crossing type		
Appropriate cross	75	10.8%
Early cross	460	66.4%
Late cross	158	22.8%
Crossing manoeuvre		
No vehicles	478	69.0%
Cyclist crossed after vehicle passed	178	25.7%
Cyclist crossed in front of vehicle	35	5.1%
Vehicle stopped for cyclist	2	0.3%
TOTAL	693	100.0%

Figure 6.1
Example of an appropriate merge behaviour (using the designated crossing infrastructure) at Site A

Figure 6.2
Example of an early merge behaviour at Site A

Figure 6.3
Example of a late merge behaviour at Site A

In addition to categorising behaviour, the bicycle type, clothing type, conspicuity, and rear light usage of each cyclist that passed were also categorised. These categorisations were consistent with those used in a previous cyclist observation by Raftery & Grigo (2013).

A summary of the characteristics of cyclists at Site A is shown in Table 6.3. Most cyclists rode a racing bike (69.3%) and wore full-cycling gear (58%). In terms of visibility of cyclists, 30.4% had a rear light illuminated (and this was generally in the morning period) and 37.2% and 29.1% respectively were classified as having high front or rear visibility. Note that fewer cyclists were considered to have had high rear conspicuity due to backpacks that degraded visibility (despite the clothing beneath being highly visible).

Table 6.3
Characteristics of the cyclists at Site A

Feature	Number	Percent
Bicycle type		
Racer	480	69.3%
Mountain bike	86	12.4%
Hybrid	82	11.8%
Mountain bike - electric	18	2.6%
Hybrid - electric	17	2.5%
Other	7	1.0%
Other - electric	3	0.4%
Clothing type		
Full-cycling	402	58.0%
Non-cycling	219	31.6%
Half-cycling	72	10.4%
Front conspicuity		
Low	435	62.8%
High	258	37.2%
Rear conspicuity		
Low	491	70.9%
High	202	29.1%
Rear light		
Yes and on	211	30.4%
No or off	482	69.6%
TOTAL	693	100.0%

6.2 Site B

Site B was the intersection of Northbourne Avenue with Cooyong Street and Barry Drive. The stakeholders noted that the installation of the light rail corridor resulted in cyclists having to zig-zag across from Cooyong Street to cross Northbourne Avenue. The camera installed at this site was therefore positioned at the North-East corner to record cyclists travelling along Cooyong Street across the light rail corridor. Additionally of interest was the frequency of cyclists using the right-turn bike box infrastructure and any instances of issues between cyclists and left turning vehicles.

Site B was the most complicated and challenging for analysis. Due to the complexity of the intersection, there were around 200 separate cyclist origin-designation routes within view of the camera. It should be noted that the camera was not able to view the entire intersection with enough detail to capture all cyclist movements. Therefore, the results below will refer to the 'observable cyclists' as those that were able to be seen from the position of the camera.

Figure 6.4 is labelled to show different origins, intermediate locations, and destinations that were observed by the camera and identified by the machine learning system. The majority of the observable cyclists used route A-to-J followed by route C-to-E.

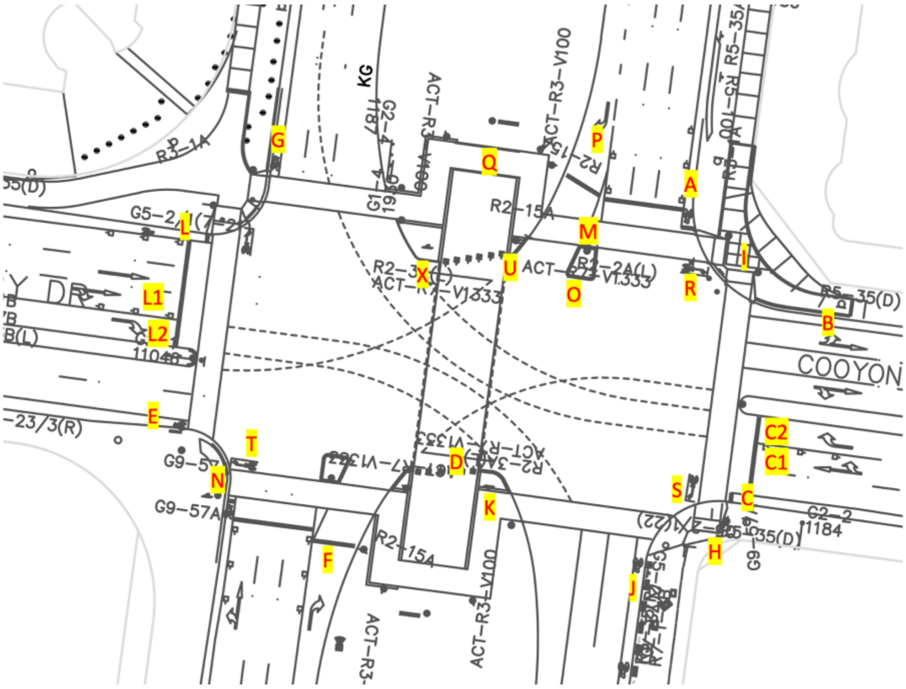

Figure 6.4
The different origins, intermediate passing locations and destinations that cyclists used, yielding around 200 route combinations

There were 284 cyclists originating at position C and travelling across the light rail corridor to position E. These cyclists either crossed along the roadway at position D or through the crosswalk at position K, as shown in Figure 6.5. As a result of the advanced stop-line, most of the 253 who travelled through position D reached the rail crossing point before any vehicles and without issue. However, in 16 cases cyclists arrived at point D simultaneously with a vehicle and, in the opinion of a researcher analysing the video, it was deemed that the vehicle passed the cyclist too closely. In another two cases cyclists were observed to be visibly cautious on the approach to the light rail crossing point due to the presence vehicles. The remaining 31 cyclists who originated at position C travelled via position K to traverse the rail corridor. Interestingly, over half (55%) of these cyclists disobeyed a red signal to get to position K.

Figure 6.5
Observable cyclist routes across light rail tracks

Figure 6.6 shows that there were two sets of observable routes where cyclists could turn right using either a right-turn lane or a bicycle box. A total of five cyclists used the bike holding box at S to turn right from position A. This is compared to 8 cyclists that used the turning lane at position P to turn right to position E.

Twenty-eight cyclists originally heading East on Barry Drive turned right to head South on Northbourne Avenue. Four cyclists, starting at position L, used the bike box at position R to travel towards position J. In comparison, eight cyclists used the middle lane (position L1) and five used the right lane (position L2) to turn right.

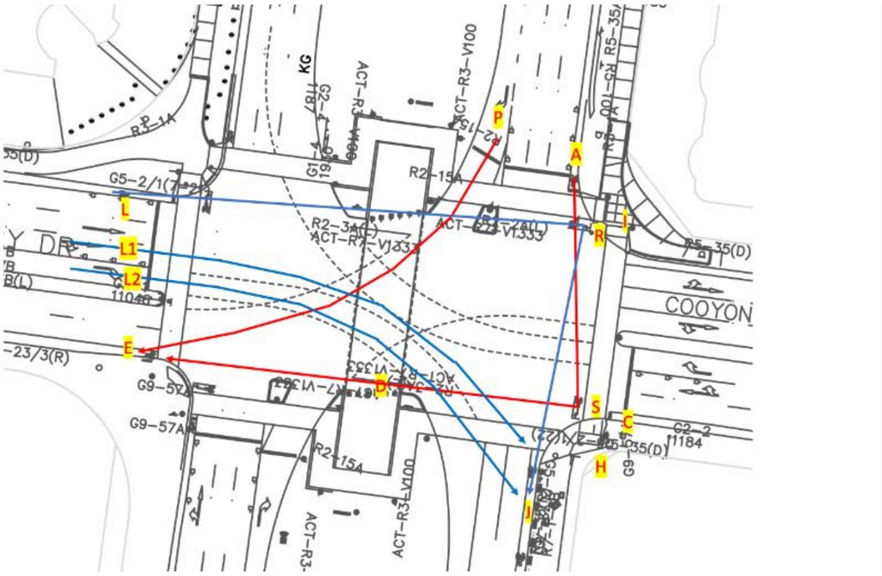

Figure 6.6
Origins and destinations for observable routes used by cyclists to turn right

The two observable left turn interactions between cyclists and vehicles are shown in Figure 6.7 below. There were 71 instances where a vehicle originally heading South on Northbourne Avenue turned left on Cooyong Street and had an interaction with a cyclist (LT1). In the majority of these interactions (54.9%), a left turning vehicle caused a cyclist to slow down in order to avoid a potential conflict. In 26.8% the turning vehicle blocked the bike lane such that a cyclist had to overtake on the right or wait for the vehicle to clear the lane.

While there were 284 cyclists originating at C heading to destination E, and there were very few issues with left turning vehicles (LT2). Nevertheless, there were 6 occasions where cyclists had to slow for a left turning vehicle, three occasions when a cyclist overtook a left turning vehicle that was blocking the bike lane, and one occasion when a cyclist had to stop and wait for a vehicle to clear the bike lane.

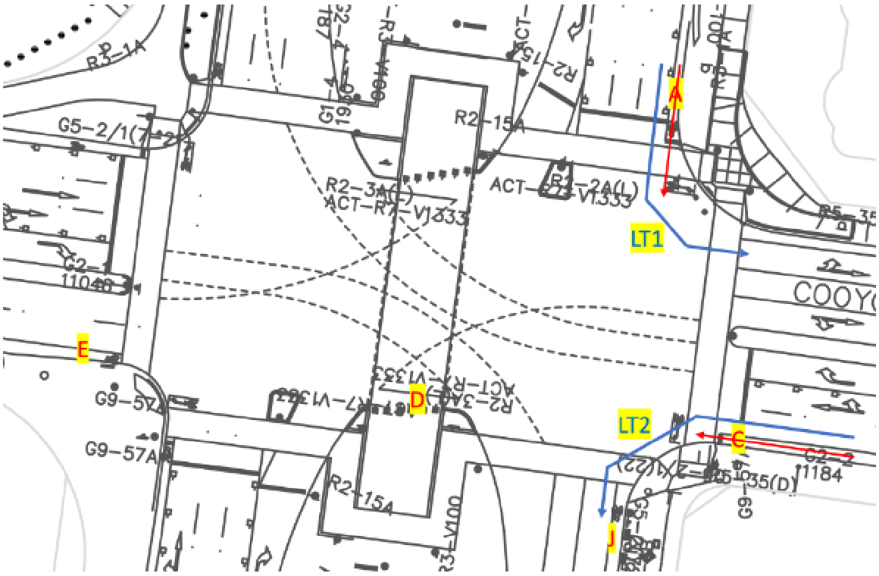

Figure 6.7
Observable left turn interactions between motor-vehicles and cyclists

6.3 Site C

Site C was the intersection of Drakeford Drive and Barr Smith Avenue. Drakeford Drive is an 80 km/h arterial road with a dedicated bike lane, but also has a separated shared path running parallel. Cyclists riding North-West along Drakeford Drive are expected to exit and merge across the Barr Smith Avenue left turn lane (which merges into Drakeford Drive) to continue riding along Drakeford Drive. A painted cyclist crossing point is provided for this purpose, which encourages riders to make a sharp left turn which would require them to slow considerably. Cyclists using the separated shared path have four locations where they must cross a road. This location was described by stakeholders as:

> *"...difficult for riders to see vehicles on Barr-Smith as they approach crossing, no pedestrian crossing, drivers have priority. On-road cyclists must cross the left turn lane, can be confusion with who gives way, slip lanes with merge lane on major road means vehicles turn left at speed".*

The camera at this location was installed to capture behaviours of cyclists and vehicles at the various intersections and to observe how cyclists crossed from Drakeford Drive over Barr Smith Avenue.

For the cyclists heading North-West and merging from Drakeford Drive across Barr Smith Avenue to continue heading north-west, four merge behaviours were identified. These were (A) utilising the designated painted crossing point, (B) travelling over to continue on the shared path, (C) merging across just past the designated crossing point, and (D) merging across a significant distance past the designated crossing point.

There were 121 cyclists observed who were riding North-West along the Drakeford Drive bicycle lane. The crossing behaviour of these cyclists were categorised as shown in Table 6.4. The majority of cyclists elected to perform either an early merge or a later merge and none were observed to have used the painted infrastructure crossing point in the intended manner.

Based on the movement of road vehicles and cyclists in the area, there were nine potential conflict areas identified. The vehicle travel paths and conflict areas are shown in Figure 6.8 below.

Table 6.4
Cyclist crossing behaviours at Site C

Crossing type	Number	Percent
Appropriate cross	0	0.0%
Shared path cross	1	0.8%
Early merge	82	67.8%
Late merge	38	31.4%
TOTAL	121	100.0%

As shown in Table 6.5, most cyclists did not have any interaction with a vehicle when travelling through Site C. There were 141 instances where a cyclist had to change their behaviour due to the presence, or the approach, of a vehicle. These behaviour changes included slowing, waiting for a vehicle to move past, or adjusting to travel around a stationary vehicle. Three interactions were considered, by the researcher analysing the videos, to be conflict events. On two occasions a cyclist had to brake abruptly at conflict area 4 to avoid a vehicle travelling on path X (with one cyclist losing balance). On another occasion, a child cyclist had to abruptly stop at conflict area 1 to avoid a left turning vehicle following path Z.

There were 26 instances where a vehicle changed their behaviour in response to the presence of a cyclist in the vicinity. In most cases this behaviour change was a vehicle slowing to allow a cyclist to pass. There were also a few instances where a vehicle that was blocking the shared path in conflict area 2 reversed to allow a cyclist to pass. In one case a vehicle braked abruptly as a result of a cyclist crossing over conflict area 2 without slowing.

In a small number of cases both the cyclist and the vehicle changed their behaviour. This was usually observed as both road users slowing and the vehicle then giving way for the cyclist to continue.

Table 6.5
Cyclist-vehicle interaction effects at Site C

Interaction effect	Number	Percent
No interaction	875	83.4%
Cyclist changed behaviour	141	13.4%
Vehicle changed behaviour	26	2.5%
Cyclist and vehicle changed behaviour	7	0.7%
TOTAL	1,049	100.0%

There were thirteen distinct routes taken by cyclists observed at the Site C and these are shown in Figure 6.9. The frequency of observations for each route is presented in Table 6.6 below. Over 85% of cyclists were observed to be using the off-road shared path, while 11.4% of cyclists were seen to be using the on-road bike lane.

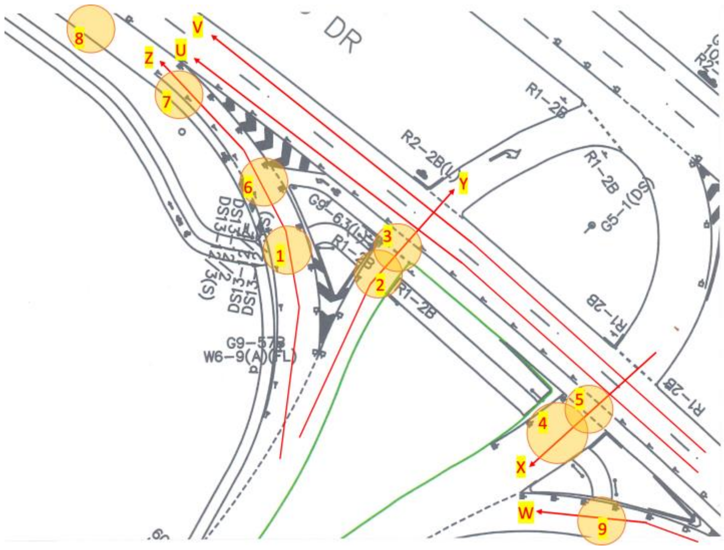

Figure 6.8
Road vehicle travel paths and potential cyclist conflict areas (circled 1 to 9) at Site C

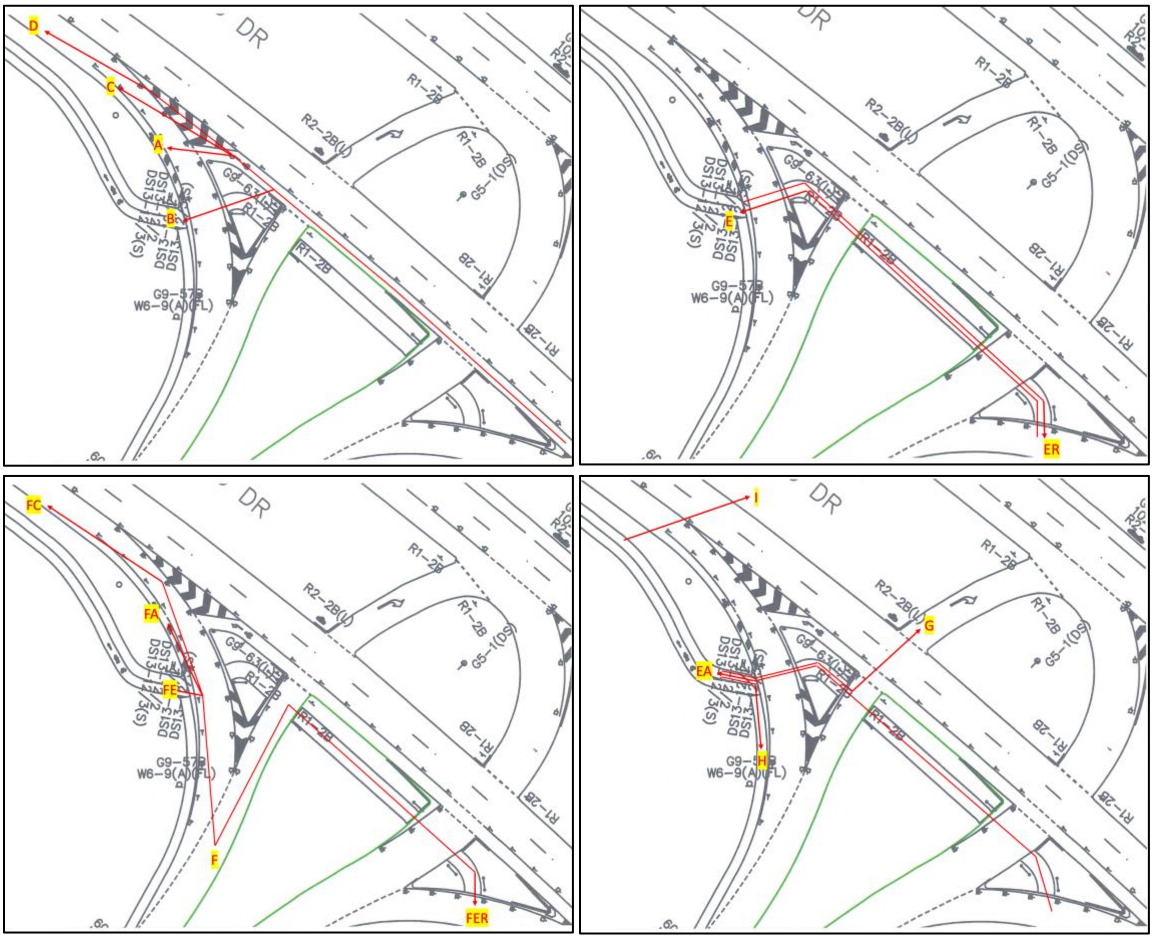

Figure 6.9
Cyclist routes observed at Site C

Table 6.6
Frequency of cyclist routes taken at Site C

Route	Number	Percent
ER	529	50.4%
E	371	35.4%
C	82	7.8%
D	38	3.6%
G	9	0.9%
FER	8	0.8%
FA	3	0.3%
FE	2	0.2%
H	2	0.2%
I	2	0.2%
EA	1	0.1%
FC	1	0.1%
B	1	0.1%
TOTAL	1,049	100.0%

6.4 Site D

Site D was the bike lane on Drakeford Drive as it terminated after the intersection with Mortimer Lewis Drive. The stakeholders stated that bike lanes which suddenly ended after an intersection were problematic for cyclists. A camera was installed so that cyclists using the bike lane running North on Drakeford Drive and terminating after Mortimer Lewis Drive could be observed and any issues noted.

Relatively few cyclists were observed at this site over the seven-day monitoring period. The routes taken by the observed cyclists are shown in Figure 6.10. and the frequency of cyclists using those routes is shown in Table 6.7. There were only 27 cyclists observed initially travelling along the bike lane on Drakeford Drive and proceeding through the intersection with Mortimer Lewis Drive (route A). Additionally, there were five cyclists who joined the terminating bike lane from routes B, C, and D and two cyclists who used the footpath parallel to the bike path (routes E and F).

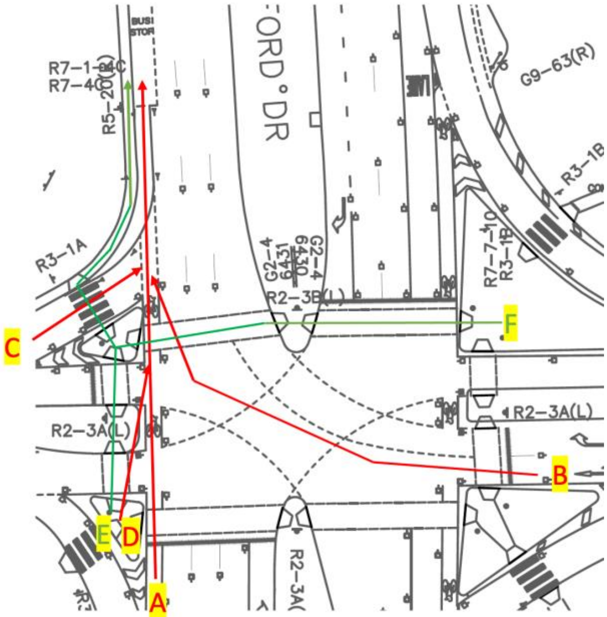

Figure 6.10
Cyclist routes observed at Site D

Table 6.7
Frequency of cyclist routes taken at Site D

Route	Number	Percent
A	27	79.4%
B	3	8.8%
C	1	2.9%
D	1	2.9%
E or F	2	5.9%
TOTAL	34	100.0%

The observed interactions between cyclists and vehicles at the termination of the bike lane for the 27 cyclists who originated from Drakeford Drive is shown in Table 6.8. While most cyclists navigated the route in the absence of any vehicles, there were three cyclists for which the terminating bike lane appeared to compromise cyclist safety with vehicles passing too closely (in the opinion of the researcher analysing the video). In terms of conflicts with left turning vehicles, there were no instances where a cyclist and a left turning vehicle were present concurrently so no conflicts were observed.

Table 6.8
Cyclist-vehicle interaction effects at Site D

Interaction effect	Number	Percent
No interaction	14	51.9%
Vehicle pass with sufficient space	10	29.4%
Vehicle pass with insufficient space	3	8.8%
TOTAL	27	100.0%

Table 6.9 shows the rear conspicuity and rear light status of the cyclists at Site D according to the classifications derived by Raftery & Grigo (2013). Around a third of cyclists had an illuminated rear light and 32.4% were classified as having high rear visibility.

Table 6.9
Characteristics of the cyclists at Site D

Feature	Number	Percent
Rear conspicuity		
Low	23	67.6%
High	11	32.4%
Rear light		
No or off	22	64.7%
Yes and on	12	35.3%
TOTAL	34	100.0%

7 Discussion

A camera-based system for the counting and classification of cyclists, which also provides the ability to categorise cyclist behaviour around, or interactions with, safety infrastructure has been developed. The aim of this project was to trial and showcase the capabilities of the camera system that was augmented with machine learning to capture and analyse cyclist interactions with infrastructure in the ACT.

A bespoke video recording system was developed for the project that consisted of a mini-computer, battery, data storage, and video camera mounted inside a weather-proof enclosure. However, this system was complex to set up and a newly released, commercially available system from SurveyTech was identified as an alternative. After conducting several tests with the SurveyTech system to ensure it would be suitable for the project, four units with long term recording capability were acquired.

Several cycling safety stakeholders were then consulted to obtain suggestions for where cyclist interactions with road infrastructure would be worth monitoring for potential problems. This consultation generated a list of 30 locations and, based on a set of criteria, four were chosen for monitoring with the four available camera systems.

A number of tasks were then completed to obtain the necessary permissions to be able to install the four cameras in the locations of interest for recording in public. These tasks included identifying specific street light poles upon which a camera was to be mounted, developing a suitable temporary traffic management plan, applying for use of public unleased land, and designing an appropriate sign to accompany the cameras. The camera installations were performed by two team members from local contractor (RD Gossips) and took less than 30 minutes at each site.

After recording during daylight hours for an entire week, the cameras were removed and the video footage was collated for processing (and backed up). There was a total of 378 hours' worth of video recorded across the four sites. This video data was then processed through the machine learning system in order to isolate only the sections where a cyclist was present within the camera scene. This processing produced summary videos that were almost 95% shorter than the original video data, as well as an image set of all cyclists detected at each site.

The image sets were filtered by a researcher to remove instances of duplicates and any non-cyclists. Concurrently, the images were classified to identify the characteristics of the detected cyclists. This process was rapid and provided a summary of the number of cyclists at each site and their characteristics.

Around 4,000 cyclists were detected, and their perceived gender ranged between 70% to 84% for males and 13% to 21% for females. Helmet compliance ranged between 90% to 100%. The proportion of cyclists detected per day of the week was unique for each site, indicating there was variety in the selected sites in terms of commuting riding (on weekdays) versus leisure riding (on weekend days).

The summary videos from each site were analysed by a researcher to classify various behaviours and other factors of interest, as well as to refine the output from the image analysis.

At Site A it was observed that few cyclists (10.8%) used the designated infrastructure to travel across the merging road. This is likely to be due to the angle on the painted infrastructure, which would require cyclists to slow considerably (or stop). This may be considered unfavourable to many cyclists who might want to maintain their speed when merging across a motor-vehicle traffic lane. Despite the low rate of usage for the infrastructure, most cyclists did choose to merge shortly after the designated cross point, when there was still a reasonable painted buffer area. Concerningly, over a fifth of cyclists merged late, exposing themselves to potential risk with less buffer area between themselves and vehicle traffic. Less than a third of cyclists at Site A were observed using a rear-light during the daylight hours or having high rear visibility, while 37% had high front visibility.

Site B was a complex intersection and cyclists were observed undertaking more than 200 different routes through the scene. When focussing on the routes of interest, it was found that motor-vehicle left turns were occasionally problematic for cyclists, particularly during a green cycle, and some risky left turns were observed. Some close passes were noted as a result of the light rail corridor which offset the bike lane towards the traffic lane slightly. Some cyclists appeared to avoid the offset bike lane and diverted to use the pedestrian crossing infrastructure instead. The bike boxes appeared to be under-utilised, with greater numbers of cyclists choosing to turn right from the traffic lanes.

At Site C it was observed that more cyclists used the bike paths than the bike lane. For the 121 cyclists that did use the bike lane, none were observed to have utilised the dedicated crossing point infrastructure provided for cyclists. As with Site A, this is assumed to have been due to the tight angle of the crossing point. Beyond the bike lane crossing point infrastructure, there were several other conflicts observed between cyclists and vehicles at various locations at Site C where the off-road path crossed a roadway.

Comparatively few cyclists were observed at Site D, but there were still three instances where there was a perceived issue in terms of vehicles passing too closely as a result of the bike lane terminating suddenly. Around a third of the observed cyclists had a rear light illuminated and were classified as having high rear visibility.

Overall, this project was successful, in that a video recording system, capable of recording 7-days of daylight footage was acquired, deployed and the recordings processed by the ML software and some useful basic analyses were undertaken. There were limitations in how the video cameras could be installed, using existing fixed light poles 3m above ground, with 135-degree field of view camera lenses, but this was expected given only one camera was installed at each location. For future investigations, more focussed examinations could be undertaken looking at specific routes, rather than all routes, or multiple cameras could be used.

This project has demonstrated that a camera system augmented with processing by a machine learning system has the capability to provide rapid insights into the behaviour of cyclists as specific locations of interest in the road network. The generated images of each cyclist detected by the system allows for a relatively inexpensive and rapid audit of the number and characteristics of all cyclists passing through a location. This project demonstrated examples of categorising by gender, helmet use, front and rear conspicuity, rear light usage, bicycle type, and clothing type, but any other visually identifiable characteristics of interest could also be categorised.

The generated summary videos were able to reduce the amount of recorded footage that would need to be viewed by around 95%. It was found that were was plenty of value in simply viewing the summary video from each site in terms of obtaining a general understanding of cyclist behaviour at that location. In some circumstances, this may be sufficient to provide confidence that a particular piece of cyclist infrastructure is functioning correctly or to provide some visual examples of sub-optimal use of infrastructure.

In this project the summary videos from each site were formally analysed by a researcher who categorised particular cyclist behaviours of interest. The time needed to view the summary videos and perform the categorisations varied for each type of behaviour and could take a significant amount of time in some cases. However, the result was a comprehensive (and analytical) understanding of the frequency of various cyclist behaviours at a particular location, which could be used for the purposes of deciding whether some road safety action was required. Additionally, it would provide clear evidence for what *type* of road safety action was required.

The camera system augmented with machine learning was developed as an inexpensive and rapid method of gathering data, and so the time required to analyse the summary videos does represent an issue. However, it should be noted that there is no requirement to record an entire week of footage at a location of interest. At locations where there are many cyclists present (such as at Site B) it would be suitable to record for just a few days, and this would provide sufficient footage to collate a representative sample of cyclist behaviours while viewing a reduced set of summary videos.

Each of the sites monitored in this project were unique, but that did not cause any major issues for the camera systems or with the identification of cyclists by the machine learning system.

One minor issue was found at Site B where the large intersection produced cyclist detections from many different directions of travel. The problem with this is that it resulted in many cyclist detections in areas that were not the focus of that site. Furthermore, consistent detections in some of the background areas of Site B were noted to have been unreliable as they consisted of only few pixels for analysis (i.e. very small cyclist images). A comprehensive audit of cyclists in those areas was thus unlikely to have been achieved.

Therefore, it is important for future deployments that the camera is directed towards the area of interest, and that multiple cameras should be utilised if a comprehensive audit across a large area is required.

Acknowledgements

This project was made possible with assistance from the ACT Road Safety Fund. We also acknowledge the valuable support and assistance of:

- Shelly Fraser (Roads ACT),
- Jane Carey (Roads ACT),
- Anand Joshi (Roads ACT),
- Josh Mahony (Roads ACT),
- Sean Sloan (Roads ACT),
- Darren Gerrard (Roads ACT),
- Phil Gritti (Roads ACT),
- Colin Evans (Roads ACT),
- Neil Pincombe (Roads ACT),
- Robert Dunn (RD Gossip Pty Ltd), and
- Budi Wijaya (RD Gossip Pty Ltd)

Additionally, we would like to thank the cyclist safety stakeholders who provided responses to our stakeholder engagement letter.

References

Ponte, G., Szpak, Z. L., Woolley, J. E., & Searson, D. J. (2014). *Using specialised cyclist detection software to count cyclists and determine cyclist travel speed from video*. Australasian Road Safety Research, Policing and Education Conference, Melbourne, Victoria, 12-14 November 2014.

Raftery, S. J., & Grigo, J. A. L. (2013). The conspicuity of South Australian cyclists: implications for safety (CASR113). Adelaide: Centre for Automotive Safety Research.

Appendix A – Stakeholder engagement letter

Dear <organisation contact person>,

Researchers at the Centre for Automotive Safety Research (the University of Adelaide) are undertaking a project entitled “video capture, observation and analysis of cyclists using various infrastructure in the ACT with machine learning”. The fundamentals of the project are to:

1. Consult with stakeholders to identify locations of interest in the ACT where there would be a safety benefit in observing cyclist behaviour;
2. Prepare secure mountings and temporarily install cameras at locations identified by the consultation for a period of one week;
3. Remove the cameras and process the recorded video using machine learning software developed at the University of Adelaide to isolate periods of footage where cyclists are present;
4. Analyse the footage to investigate cyclists’ interactions with infrastructure and traffic; and
5. Report on the outcomes of the video analysis and project goals.

As a cyclist safety stakeholder, we are seeking your input into possible locations where observation of cyclists would be of interest, along with details about what specifically should be observed at these particular locations.

As cyclist collisions are relatively rate, and the video observation will only occur for a period of one week, we are not seeking to record cyclist collisions. Rather, we would like to observe typical cyclist behaviour at particular locations to answer questions such as: what situations may result in increased risk, how is cyclist safety infrastructure being utilised, what types of cyclists are using particular routes (and how fast are they travelling), what proportion of cyclists disobey red light signals, or what proportion of cyclists are wearing safety equipment like helmets and hi-vis clothing.

We are open to any suggestions and would be particularly interested in any known infrastructure problem areas or where conflicts with other road users are common.

Our objective is to select four locations for monitoring but please feel free to recommend, in priority order, as many locations as you feel is appropriate (as well as the type of behaviour to be observed). The final list of locations will be selected based on the responses from all consulted stakeholders and will also need to account for the availability of suitable fixed infrastructure for mounting cameras.

It would be appreciated if we could have your response by Friday 3 September 2021.

Kind Regards,

Dr. Jamie Mackenzie (Research Fellow, Centre for Automotive Safety Research)

Appendix B – Camera signage

Figure B.1
Sign produced to accompany installed cameras (not to scale)