

Development of a Connected and Autonomous Vehicle Modelling Framework, with Implementation in Evaluating Transport Network Impacts and Safety

Author:

Virdi, Navreet

Publication Date:

2020

DOI:

<https://doi.org/10.26190/unsworks/22242>

License:

<https://creativecommons.org/licenses/by-nc-nd/3.0/au/>

Link to license to see what you are allowed to do with this resource.

Downloaded from <http://hdl.handle.net/1959.4/70492> in <https://unsworks.unsw.edu.au> on 2024-05-05



**Development of a Connected and
Autonomous Vehicle Modelling
Framework, with Implementation in
Evaluating Transport Network Impacts
and Safety**

By

Navreet Virdi

A thesis submitted in fulfilment
of the requirements for the degree of
Doctor of Philosophy

Presented to the
School of Civil and Environmental Engineering
At the University of New South Wales

Supervised by
Professor S. Travis Waller
Dr Hanna Grzybowska

Submitted on 19th November 2020

Dissertation Sheet

Surname : Virdi
Given Name : Navreet
Degree : PhD
Faculty : Engineering
School : Civil and Environmental Engineering
Thesis Title : Development of a Connected and Autonomous Vehicle Modelling Framework,
with Implementation in Evaluating Transport Network Impacts and Safety

Abstract

Transportation systems form critical links that connect developed cities with the broader world. They connect our residential, recreational, employment, and natural environments. Annual population and vehicle ownership growth place an increasing strain on transport systems, resulting in escalating levels of congestion and delay. Using new infrastructure and expansion is a problematic solution, as it often incentivises greater private vehicle use and worsens long-term congestion. Infrastructure expansion requires repeated and increasing levels of capital investment.

The Connected and Autonomous Vehicle (CAV) is an emerging technology that facilitates communication with infrastructure and other agents. CAVs address many inefficiencies of human driving by exhibiting instantaneous reaction times, smaller headways, and vehicle platooning. Their fundamentally different driving behaviour may render many infrastructure planning and modelling tools not applicable to future mixed fleets and CAVs. This thesis develops a comprehensive modelling framework for the emulation of CAV behaviour in microsimulation, with a focus on car-following, lane-changing, gap-acceptance, autonomous merging, and vehicle cooperation. The developed framework is implemented in a range of investigations aimed at better understanding the impact of mixed fleets and CAVs on vehicle kinematics, intersection performance, and safety. Uncertainties regarding CAV behaviour and motorway capacity, delay redistribution through signal optimisation, and the need for recalibrating macrosimulation modelling parameters are also investigated.

These investigations indicate that CAVs improve network performance, driver aggression (acceleration), and driver comfort (jerk). Low levels of penetration improved fleet operations, leading to increased throughput, increased capacity, reduced delay, and reduced likelihoods of accidents and conflicts. Average network delay is decreased significantly by redistributing the CAV travel time savings to all network agents, through signalling re-optimisation. Finally, this thesis demonstrates that macrosimulation modelling parameters used for human fleets show reduced predictive qualities when applied to mixed fleets and CAV environments. The performed recalibration provides significantly improved results in the predictive quality of volume-delay functions.

Declaration relating to disposition of project thesis/dissertation

I hereby grant to the University of New South Wales or its agents a non-exclusive licence to archive and to make available (including to members of the public) my thesis or dissertation in whole or in part in the University libraries in all forms of media, now or here after known. I acknowledge that I retain all intellectual property rights which subsist in my thesis or dissertation, such as copyright and patent rights, subject to applicable law. I also retain the right to use all or part of my thesis or dissertation in future works (such as articles or books).

.....
Signature

.....
Date

The University recognises that there may be exceptional circumstances requiring restrictions on copying or conditions on use. Requests for restriction for a period of up to 2 years can be made when submitting the final copies of your thesis to the UNSW Library. Requests for a longer period of restriction may be considered in exceptional circumstances and require the approval of the Dean of Graduate Research.

Originality, Copyright, and Authenticity Statement

Originality Statement

I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or substantial proportions of material which have been accepted for the award of any other degree or diploma at UNSW or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by others, with whom I have worked at UNSW or elsewhere, is explicitly acknowledged in the thesis. I also declare that the intellectual content of this thesis is the product of my own work, except to the extent that assistance from others in the project's design and conception or in style, presentation and linguistic expression is acknowledged.

Copyright Statement

I hereby grant the University of New South Wales or its agents a non-exclusive licence to archive and to make available (including to members of the public) my thesis or dissertation in whole or part in the University libraries in all forms of media, now or here after known. I acknowledge that I retain all intellectual property rights which subsist in my thesis or dissertation, such as copyright and patent rights, subject to applicable law. I also retain the right to use all or part of my thesis or dissertation in future works (such as articles or books). For any substantial portions of copyright material used in this thesis, written permission for use has been obtained, or the copyright material is removed from the final, public version of the thesis.

Authenticity Statement

I certify that the Library deposit digital copy is a direct equivalent of the final officially approved version of my thesis.

.....
Signature

.....
Date

Inclusion of Publications Statement

UNSW is supportive of candidates publishing their research results during their candidature as detailed in the UNSW Thesis Examination Procedure.

Publications can be used in their thesis in lieu of a Chapter if:

- The candidate contributed greater than 50% of the content in the publication and is the “primary author”, ie. the candidate was responsible primarily for the planning, execution and preparation of the work for publication.
- The candidate has approval to include the publication in their thesis in lieu of a Chapter from their supervisor and Postgraduate Coordinator.
- The publication is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in the thesis.

Please indicate whether this thesis contains published material or not:

- ☐ This thesis contains no publications, either published or submitted for publication.
- ☒ Some of the work described in this thesis has been published and it has been documented in the relevant Chapters with acknowledgement.
- ☐ This thesis has publications (either published or submitted for publication) incorporated into it in lieu of a chapter and the details are presented below.

Candidate’s Declaration

I declare that:

- I have complied with the UNSW Thesis Examination Procedure.
- Where I have used a publication in lieu of a Chapter, the listed publication(s) below meet(s) the requirements to be included in the thesis.

.....
Candidate's Name

.....
Signature

.....
Date

Abstract

Transportation systems form critical links that connect developed cities with the broader world. They connect our residential, recreational, employment, and natural environments. Annual population and vehicle ownership growth place an increasing strain on transport systems, resulting in escalating levels of congestion and delay. Using new infrastructure and expansion is a problematic solution, as it often incentivises greater private vehicle use and worsens long-term congestion. Infrastructure expansion requires repeated and increasing levels of capital investment.

The Connected and Autonomous Vehicle (CAV) is an emerging technology that facilitates communication with infrastructure and other agents. CAVs address many inefficiencies of human driving by exhibiting instantaneous reaction times, smaller headways, and vehicle platooning. Their fundamentally different driving behaviour may render many infrastructure planning and modelling tools not applicable to future mixed fleets and CAVs. This thesis develops a comprehensive modelling framework for the emulation of CAV behaviour in microsimulation, with a focus on car-following, lane-changing, gap-acceptance, autonomous merging, and vehicle cooperation. The developed framework is implemented in a range of investigations aimed at better understanding the impact of mixed fleets and CAVs on vehicle kinematics, intersection performance, and safety. Uncertainties regarding CAV behaviour and motorway capacity, delay redistribution through signal optimisation, and the need for recalibrating macrosimulation modelling parameters are also investigated.

These investigations indicate that CAVs improve network performance, driver aggression (acceleration), and driver comfort (jerk). Low levels of penetration improved fleet operations, leading to increased throughput, increased capacity, reduced delay, and reduced likelihoods of accidents and conflicts. Average network delay is decreased significantly by redistributing the CAV travel time savings to all network agents, through signalling re-optimisation. Finally, this thesis demonstrates that macrosimulation modelling parameters used for human fleets show reduced predictive qualities when applied to mixed fleets and CAV environments. The performed recalibration provides significantly improved results in the predictive quality of volume-delay functions.

Acknowledgements

Had it not been for the invaluable contribution of the following people, the arduous journey would have contained many more sleepless nights.

First and foremost, I would like to acknowledge and thank my supervisors, Professor S. Travis Waller and Dr Hanna Grzybowska. Your guidance has been invaluable and your provision of resources has made this work possible.

To my friend Dean Preston, your expertise in computer science and programming helped bring to reality a set of mathematical equations and algorithms. Without your assistance during the many times I had reached my 11th hour, I may still be playing with C++.

To my colleagues Esta Qiu and Mario Dizon Jr, undertaking our collaborative projects were an absolute pleasure. Your assistance in modelling experiments in this thesis provided valuable teaching and learning opportunities, and helped expand the capabilities of microsimulation.

To my parents Amrit and Harjit, words can not begin to express the gratitude I feel for your continued assistance, selflessness, and care in ensuring my comfort and wellbeing. You have been a constant source of inspiration, learning, guidance, safety, and love. For sharing many of my sleeplessness nights, stress, and burdening some of the pressure, I am grateful.

To my brother Karan, we may not see eye to eye on most things, but you nonetheless have been a source of inspiration, to continue learning and developing as an individual, and to see the world in your unique and creative perspective.

Finally, to my partner Sanveet, your understanding and support will be forever cherished, and your presence was a constant reminder for what is waiting on the other side.

List of Publications

Journal Publications

Virdi, N., Grzybowska, H., Waller, S. T., & Dixit, V. (2019). A safety assessment of mixed fleets with connected and autonomous vehicles using the Surrogate Safety Assessment Module. *Accident Analysis & Prevention*, 131, 95-111.

Sinha, A., Chand, S., Wijayaratna, K. P., Virdi, N., & Dixit, V. (2020). Comprehensive safety assessment in mixed fleets with connected and automated vehicles: A crash severity and rate evaluation of conventional vehicles. *Accident Analysis & Prevention*, 142, 105567.

Sinha, A., Chand, S., Wijayaratna, K. P., Virdi, N., & Dixit, V. (2020). Crash Severity and Rate Evaluation of Conventional Vehicles in Mixed Fleets with Connected and Automated Vehicles. *Procedia Computer Science*, 170, 688-695.

Conference Publications

Virdi, N., Grzybowska, H., & Waller, S. T., A New Era of City Planning: How the network benefits of Connected and Autonomous Vehicle (CAVs) can be leveraged to change the way our cities operate, *UNSW Engineering Postgraduate Research Symposium*, 2019, ISBN: 978-0-9953910-3-1.

Under Review Publications

Qiu, E., Virdi, N., Grzybowska, H., & Waller, S. T. (2020). Recalibration of the BPR Function for the Strategic Modelling of Connected and Autonomous Vehicles. *Transportmetrica B: Transport Dynamics*, Under Review.

Virdi, N., Dizon, M., Grzybowska, H., & Waller, S. T. (2019). An Investigation into the Feasibility of Adopting a Pedestrian-Priority Signalling Scheme in the Presence of Connected and Autonomous Vehicles. *Transportation Research Part A: Policy and Practice*, Under Review.

List of Figures

1	Remote sensing hardware of the CAV [Gates, 2017].	10
2	WAVE communication systems, comprising of the roadside unit, the on-board unit, and the service channels [Mercedes-Benz, 2014].	15
3	Thresholds for the regions within the Wiedemann psycho-physical car-following model [PTV Group, 2016].	34
4	The action regions of the psycho-physical Fritzsche car-following model [Fritzsche, 1994].	36
5	A diagrammatic representation of the dynamically assigned leaders and followers to an ego vehicle.	92
6	The structure of the regions that define CAV behaviour.	93
7	Graphical representation of P_{min} and P_{max}	96
8	Typical acceleration, velocity, and displacement profiles exhibited by CAVs as they approach an obstruction in the vehicles' trajectory.	99
9	Arrangement and notation used to describe the system of vehicles in a lane change.	104
10	Visual representation of the score as a function of $\Delta Speed$ and ΔGap , using $c_1 = 0.5$ and $c_2 = 0.5$	106
11	Visual representation of the caveats on discretionary lane changing.	110
12	Arrangement of the merge point and zone, and the network geometry in which this proposed control framework would provide benefit.	111
13	A hypothetical example of the changes in vehicle kinematic information for an agent subjected to an imprudent lane change.	119
14	Speed-Time graphs generated using the Gipps collision-avoidance car-following model [Gipps, 1981].	121
15	Effect on vehicle kinematics using the Gipps car-following model, assuming statistically distributed parameters attain their mean value.	122
16	An overtaking manoeuvre conducted in the Vissim microsimulator.	124
17	Example of conflict areas in a 4-way signalised intersection.	125

18	An example of a priority rule used at the approach arm of a roundabout. The three subfigures indicate how the gap between the priority (green line) and yielding (red line) vehicle can be altered.	125
19	Flowchart of the Gipps lane-changing model [Gipps, 1986]. . . .	130
20	Kinematic variable vs time plots for the scenario depicting the car-following for a platoon. Vehicles are generated at a spacing less than the minimum permitted distance and are approaching a red light 100m away. Each consecutive vehicle is represented by a darker shade.	135
21	Kinematic variable vs time plots for the lane-changing scenario, where an adjacent vehicle enters the platoon. Vehicles still react to a red light 100m away. Each consecutive vehicle is represented by a darker shade.	136
22	Network used for the kinematic assessment of the CAV framework.	139
23	Signal phasing structure of for the signalised intersections	140
24	Frequency histograms depicting the distribution of velocity, acceleration, and jerk of vehicles during network operation	144
25	The F-value distribution generated from 500 iterations of the ANOVA assessment of 100,000 data points, for each kinematic variable.	150
26	Velocity histograms for each sensitivity scenario, for all CPR cases.	154
27	Acceleration histograms for each sensitivity, for all CPR cases.	157
28	Jerk histograms for each sensitivity scenario, for all CPR cases.	159
29	The stochastic and memoryless Monte Carlo process used to determine the capacity of each intersection, as a result of increasing CPR.	175
30	Geometric layout and signal phasing for the signalised intersection with a filtered right turn.	176
31	Geometric layout and signal phasing for the signalised intersection with a dedicated right turn.	176
32	Number of conflicts points in different intersection arrangements.	177
33	The phasing diagram of different geometric arrangements, indicating that greater conflict points result in more phases.	177
34	Geometric layout of the roundabout intersection.	178

35	Geometric layout of the priority intersections.	179
36	Geometric layout for the environment used to assess green time and platooning on intersection performance.	179
37	Intersection performance for the signalised intersection with a filtered right turn, as CAV penetration increases.	181
38	Intersection performance for the signalised intersection with a dedicated right turn, as CAV penetration increases.	183
39	Intersection performance for the roundabout intersection, as CAV penetration increases.	184
40	Intersection performance for the priority intersection, as CAV penetration increases	186
41	Changes to platoon behaviour as CAV penetration and signal green-time increases.	187
42	Overview of the framework used for this study.	198
43	Context map, location photograph, and microsimulation model for each of the study environments.	205
44	Delay per agent at the Ultimo Road / Quay Street intersection for the three modelling scenarios, in 10% increments of CAV penetration.	208
45	Delay per agent at the Great Western Highway / Marsden Street Intersection for the three modelling scenarios, in 10% increments of CPR.	210
46	Comprehensive modelling results detailing average delay expe- rienced by vehicles, pedestrians, and system average for each CPR, signalling scheme, and intersection.	212
47	Modelling environments used for the different investigations. . .	226
48	Throughput vs Demand for the purposes of determining capac- ity for each CAV penetration and each weaving proportion. . .	229
49	Average network speed for each CPR (each line), DoS, and weaving proportion(each subplot).	231
50	Queueing delay for each CPR (each line), DoS, and weaving proportion (each subplot).	232
51	Average travel time for each CPR (each line), DoS, and weaving proportion (each subplot).	233
52	Environment 2 capacity for each CPR and weaving proportion.	234
53	Results for investigation 2, for each CAV penetration rate. . . .	236

54	Contextualisation of the BPR function within the wider strategic modelling framework.	246
55	Comparison of the two approaches proposed by Adebisi, against the Webster model and the HCM model [Adebisi, 2018].	258
56	A schematic of the network geometry used for this study.	260
57	The weaving length, speed limit, and weaving proportion used to differentiate the three testbeds.	260
58	The corridor throughput plotted against corridor demand, used to calculate the capacity of the modelling environment.	263
59	Plotting throughput vs demand, to determine the environment capacity for different CAV penetrations. The capacity is defined as the inflection point where a demand increase does not lead to an appreciable increase in throughput.	266
60	Optimised α and β values for each CPR and testbed.	271
61	The linear models fitted to the α and β parameters.	271
62	Framework for the flow of data in this study.	287
63	The definition of rear-end and crossing angle used by SSAM, and a scenario in the Vissim microsimulator where a conflict may be misidentified.	289
64	The calibrated microsimulation network of the Geelong area in Victoria, Australia.	292
65	GEH statistic for major network links.	294
66	Number of potential accidents within the environment, in 10% increments of CPR.	297
67	Distribution of highway vehicle spacing, with increasing CPR.	298
68	Number of potential accidents during different situations in the microsimulation environment, in 10% increments of CAV penetration.	299
69	Number of potential conflicts further segregated by conflict type.	302

List of Tables

1	Origin-destination matrix used for the peak hour (<i>veh/hr</i>). The zones correspond to those outlined in Figure 22.	140
2	Results of the internal consistency analysis conducted on this environment using 50 random seeds.	141
3	The population standard deviations for each dataset involved in the ANOVA assessment.	145
4	The skewness and kurtosis for each CPR and kinematic variable in the ANOVA assessment.	146
5	The skewness and kurtosis after being transformed.	148
6	Results of the ANOVA analysis and the critical F-values associated with $d_1 = 4$ and $d_2 = 99,995$	151
7	The mean value of each kinematic variable distribution.	160
8	Relative change in the mean value of the kinematic variable distribution for each sensitivity and CPR, compared to the base case.	161
9	The skewness of the kinematic variable distributions, for each sensitivity scenario and CPR.	162
10	Percentage change in the skewness, for each sensitivity scenario compared to the base case.	162
11	The F-Value for the ANOVA analysis comparing the sensitivity scenario with the base case. The critical F-value is 3.842 for $\alpha = 0.05$	163
12	Overall impacts of CAV behaviour on the signalised intersection with right turn filter.	182
13	Overall impacts of CAV behaviour on the signalised intersection with a dedicated right turn.	184
14	Overall impacts of CAV behaviour on the roundabout intersection.	185
15	Overall impacts of CAV behaviour on the priority intersection.	186
16	Observed traffic volumes from 11:30 AM to 12:30 PM.	205
17	Scaled volume (<i>veh/hr</i>), modelled volume (<i>veh/hr</i>), and GEH.	206
18	Vehicle kilometres travelled results for each seed.	207
19	Delay (<i>sec</i>) and LOS at the Ultimo Road / Quay Street intersection for all CPR and signalling scenarios.	209

20	Delay (<i>sec</i>) and LOS at the Great Western Highway / Marsden Street intersection for all CPR and signalling scenarios.	211
21	Base case OD matrices used for the modelling Environment 1.	227
22	Stability results for both modelling environments.	227
23	Environment 1 capacity for each CPR and weaving proportion.	229
24	Change in capacity relative to the 0% CAV base case.	230
25	Change in the Environment 2 capacity for each CPR case relative to the 0% CAV base case.	235
26	Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 5% weaving proportion.	238
27	Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 10% weaving proportion.	239
28	Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 20% weaving proportion.	239
29	Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 40% weaving proportion.	240
30	The base OD matrices used for each testbed in this study.	261
31	Network statistics used to establish model stability. Each value is the proportion of the standard deviation of all runs, to the median value. 50 iterations for each testbed was used.	262
32	The quantitative values used to generate Figure 58, also showing the rate of change for demand and throughput. This table is used to inform capacity.	264
33	The capacity for each testbed and each CAV penetration.	267
34	The change in capacity for each testbed and each CAV penetration rate, relative to the base case 0% CAV scenario.	267
35	Average travel time (<i>sec</i>) for each CPR, DoS and testbed, averaged for a randomly selected 30% of the synthetic data.	269
36	Difference in the average travel speed predicted by the BPR function, and the average travel speed found in using the synthetic data.	270
37	Difference in the average travel speed predicted by the recalibrated model, and the average travel time found using the synthetic data.	273

38	A direct comparison of the BPR function vs its recalibrated variant. A green value indicates the recalibrated model outperformed the original BPR, with a red value indicating the opposite.	274
39	Network scale calibration results using the GEH statistic method.	293
40	Tabular representation of the inter-platoon spacing distribution for the fleet along the highway.	298
41	Standard deviation for the number of conflicts for different intersections in the environment.	303
42	Number of conflicts for different intersections in the environment for the “M-M” interaction type, normalised for the volume of human vehicles remaining in the network.	304

List of Abbreviations

AADT	Average Annual Daily Traffic
ABS	Australian Bureau of Statistics
ACC	Adaptive Cruise Control
AICC	Autonomous Intelligent Cruise Control
AIM	Autonomous Intersection Management
ANOVA	Analysis of Variance
API	Application Programming Interface
B&B	Branch and Bound Algorithm
BITRE	Bureau of Infrastructure, Transport and Regional Economics
CAV	Connected and Autonomous Vehicle
CPI	Consumer Price Index
CPR	Connected and Autonomous Vehicle Penetration Rate
CSV	Comma Separated Value
CTM	Cell Transmission Model
DARPA	Defence Advanced Research Programs Agency
DDI	Diverging Diamond Interchange
DLC	Discretionary Lane Change
DLL	Dynamic Linking Library
DMV	Department of Motor Vehicles
DoS	Degree of Saturation
DSRC	Dedicated Short Range Communication
FHWA	Federal Highway Administration
FVD	Full Velocity Difference
GEH	Geoffrey E. Havers
GHR	Gazis-Herman-Rothery
GPS	Global Positioning System
HCM	Highway Capacity Manual
IEEE	Institute of Electrical and Electronics Engineers
IMU	Internal Measurement Unit
INS	Internal Navigation System
ITS	Information Technology Systems
JDF	Junction Delay Function
KPI	Key Performance Indicator
LiDAR	Light Detection and Ranging

LOS	Level of Service
MaaS	Mobility as a Service
MILP	Mixed-Integer Linear Programming
MLC	Mandatory Lane Change
MOBIL	Minimising Overall Braking Induced By Lane Change
NSW	New South Wales
OD	Origin-Destination
PET	Post-Encroachment-Time
PwC	PricewaterhouseCoopers
RADAR	Radio Detection and Ranging
RCRI	Rear-end Crash Risk Index
RMS	Roads and Maritime Services
RMSE	Root Mean Square Error
RNN	Recurrent Neural Network
SDK	Software Development Kit
SSAM	Surrogate Safety Assessment Module
STM	Strategic Transport Model
TfNSW	Transport for New South Wales
TTC	Time to Collision
UK	United Kingdom
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
VDF	Volume Delay Functions
VMS	Variable Message Sign
WAVE	Wireless Access in Vehicular Environments

List of Mathematical Notation

Best efforts have been made to keep mathematical notation in this thesis consistent. However, the repetition of symbols in varying context is unavoidable. For this reason, symbols are defined after each equation within the text. The most common use-case for each symbol is defined below. Unless explicitly stated, this thesis uses the International System of Units.

English

a	Vehicle acceleration.
B	Euler integral beta function.
c_x	Calibration parameter. The subscript x denotes that a single model may have multiple different calibration parameters.
C	Link capacity.
d	Vehicle deceleration. Also represents the mechanical drag coefficient in the AICC model.
g	Gravitational acceleration.
g_1	Distribution skewness.
g_2	Distribution kurtosis.
gap_x	Gap length surrounding the ego vehicle. Also represents the number of vacant cells in the Cellular Automata model.
H	Binary Heaviside function.
H_0	Null hypothesis for the ANOVA analysis.
H_a	Alternate hypothesis for the ANOVA analysis.
i	Identification number for an individual lane, link, or node.
j	Vehicle jerk.
J_d	Delay parameter used in the Davidson VDF.
k	Aerodynamic drag coefficient.
l	Length of the vehicle. Also represents average queue length in queueing theory.
L_w	Lane width.
m	Mass of the vehicle.
n	Identification number for an individual vehicle. Also represents population size in a statistical distribution.
O	Number of vehicles having departed from an intersection approach in a single platoon.

p	Politeness factor used in the MOBIL car-following model.
P	Probability of a specified outcome. Also represents vehicle power in the Fadhloun and Rakha car-following model.
P_{max}	Parameter in the proposed CAV emulation framework, provides an upper bound on vehicle acceleration and reaction distances.
P_{min}	Parameter in the proposed CAV emulation framework, provides a lower bound on vehicle acceleration and reaction distances.
r	Vector of parameters for the driver specific random term used in discrete-choice modelling. Also represents the perceived risk in the Pollatschek et al., gap-acceptance model.
S	Set of all action for a player in a Nash's equilibrium game.
t	Discretised time increment. Also represents the truncated power function, and travel time on a link in strategic modelling.
U	Utility of an option in discrete-choice modelling. Also represents the weighting factor in a range of models.
v	Velocity of the ego vehicle.
v_f	Free flow speed.
v^*	Desired velocity of the ego vehicle.
V	Agent volume.
W	Wait time.
x	Position of the ego vehicle. Also represents the horizontal cartesian coordinate in the Tang et al., model.
x^*	Nash's equilibrium solution.
X	Vector of explanatory variables, used in discrete-choice modelling. Also represents the input vector in neural network models.
X_o	Spring speed-dependant relaxation length.
\bar{x}	Population mean.
y	Vertical cartesian coordinate in the Tang et al., model.
$Z_1(g_1)$	Transformed dataset skewness.
$Z_2(g_2)$	Transformed dataset kurtosis.

Greek

α	Vector of driver-specific random term parameters, used in discrete-choice modelling. Also used as a calibration parameter in VDFs.
β	Vector of parameters, used in discrete-choice modelling. Also used for the alternate specific constant, and a calibration parameter in VDFs.
Δt	Time headway between the ego vehicle and its leader or follower.
Δv	Difference between the velocity of the ego vehicle and its leader or follower.
Δx	Difference between the position (spacing) of the ego vehicle and its leader or follower.
ϵ	Random error term.
η	Driveline efficiency of the vehicle.
Λ	Approach or lane arrival rate.
μ	Coefficient of friction. Also represents the logit model scale parameters in Farah et al., gap-acceptance model, and the departure rate in queueing theory.
ϕ	Membership function in fuzzy logic models. Also representing the cumulative standard-normal distribution.
ρ	Density of traffic flow. Also used for approach utilisation rate in queueing theory.
ρ_{max}	Maximum density of the traffic flow.
ρ_{ref}	Road curvature.
σ	Steering angle.
τ	Vehicle (or driver) reaction time. Also used to represent the spring dampening coefficient in the Li et al., car-following model.
θ	Binary term used for conditional components of models. Also represents vehicle heading in the Tang et al., model.

Contents

Dissertation Sheet	i
Originality, Copyright, and Authenticity Statement	ii
Inclusion of Publications Statement	iii
Abstract	iv
Acknowledgements	v
List of Publications	vi
List of Figures	x
List of Tables	xiii
List of Abbreviations	xiv
List of Mathematical Notation	xvi
1 Introduction	1
2 Connected and Autonomous Vehicles, A Review	8
2.1 Hardware and Capabilities of the CAV	10
2.2 Software Control of CAVs	18
2.3 Recent Trials using CAV Technology	20
2.4 Summary	23
3 Mathematical Models for the Microsimulation of Human Driven Vehicle Behaviour	24
3.1 Car-Following Algorithms	24
3.1.1 Gazis-Herman-Rothery Models	25
3.1.2 Collision Avoidance Models	28
3.1.3 Psycho-Physical Models	33
3.1.4 Fuzzy Logic Models	37
3.1.5 Recurrent Neural Networks	41
3.2 Lane-Changing Algorithms	41
3.2.1 Rule Based Models	42

3.2.2	Discrete Choice Based Models	46
3.2.3	Artificial Intelligence Models	51
3.3	Gap-Acceptance Algorithms	56
3.4	Summary and Model Comparison	65
4	Mathematical Models for the Microsimulation Emulation of CAV Behaviour	67
4.1	Car-Following Models	67
4.2	Lane-Changing Models	75
4.3	Gap-Keeping Models	81
4.4	Autonomous Intersection Models	84
4.5	Summary and Concluding Remarks	90
5	Development of a Mixed-Fleet Microsimulation Emulation Framework	91
5.1	CAV Behavioural Emulation Framework	91
5.1.1	Car-Following Component	92
5.1.2	Gap-Acceptance Component	103
5.1.3	Lane-Changing Component	107
5.1.4	Cooperative Merging Component	110
5.1.5	Interacting with Different Vehicle Types	114
5.2	Microsimulation Emulation of Human Vehicles	115
5.2.1	Car-Following Models	116
5.2.2	Gap-Acceptance Models	123
5.2.3	Lane-Changing Models	127
5.3	Framework Summary	131
6	Vehicle Kinematics and Parameter Sensitivity	133
6.1	Kinematic-Time Plot Assessment	134
6.2	Network Evaluation of Vehicle Kinematics	137
6.2.1	Network Geometry	138
6.2.2	Kinematic Results	141
6.3	ANOVA Assessment of Network Kinematics	145
6.3.1	Verification of ANOVA Assumptions	145
6.3.2	Methodology of the ANOVA Analysis Process	148
6.3.3	ANOVA Assessment Results	150
6.4	Sensitivity Assessment of Framework Parameters	152

6.5	Summary and Conclusion	164
7	An Investigation into the Implications of CAVs on Isolated Intersection Performance	165
7.1	Literature Review of CAV intersection Modelling	166
7.2	Experimentation Methodology	174
7.2.1	Experimentation Structure	174
7.2.2	Modelling Environment	175
7.2.3	Result Evaluation Metrics	180
7.3	Experimentation Results	181
7.4	Discussion	188
7.5	Conclusion	191
8	Intersection Phasing Optimisation for a Pedestrian-Priority Signalling Scheme	192
8.1	Literature Review of Vehicle and Pedestrian Interaction	194
8.2	Experimentation Framework	198
8.2.1	Pedestrian Emulation	200
8.2.2	Signal Optimisation Process	201
8.3	Case Studies	203
8.3.1	Study Area	204
8.3.2	Model Volumes	205
8.3.3	Model Calibration	206
8.3.4	Experimentation Structure	207
8.4	Results	208
8.4.1	Pedestrian-Dominant Environment	208
8.4.2	Vehicle-Dominant Environment	210
8.4.3	Results Summary	211
8.5	Discussion	213
8.6	Conclusion	216
9	Effect of CAV Behaviour and Selective Cooperation on Motorway Capacity	217
9.1	Literature Review of CAV Behaviour Impacts on Motorway Environments	218
9.2	Experimentation Methodology	223
9.2.1	Experimentation Structure	224

9.2.2	Modelling Environment	225
9.2.3	Result Evaluation Metrics	228
9.3	Experimentation Results	228
9.4	Discussion	241
9.5	Conclusion	242
10	Recalibration of the BPR Volume Delay Function for the Strategic Modelling of CAVs	243
10.1	Review of the Relevant Literature	245
10.1.1	Strategic Modelling in Transport Planning	245
10.1.2	Volume Delay Functions and Capacity Models used in Macroscopic Strategic Modelling	247
10.1.3	Recent Attempts at the Strategic Modelling of CAVs	251
10.2	Experimentation Development	259
10.2.1	Modelling Environment	260
10.2.2	Corridor Capacity	262
10.2.3	Experimentation Structure	265
10.3	Experimentation Results	265
10.3.1	Testbed Capacity	265
10.3.2	Verify BPR Model Performance	267
10.3.3	BPR Model Recalibration	271
10.3.4	Corrected Model validation	272
10.3.5	Commentary on the Conical and Davidson VDFs	275
10.4	Discussion	275
10.5	Conclusion	278
11	Safety Assessment of Mixed Fleets using the Surrogate Safety Assessment Module	279
11.1	Literature Review of Vehicle Conflicts and Safety	280
11.1.1	Safety Studies of Human Driven Vehicles	281
11.1.2	Safety Studies of Autonomously Driven Vehicles	283
11.2	Experimentation Framework	286
11.3	Case Study	290
11.3.1	Microsimulation Network	290
11.3.2	Model Calibration	292
11.3.3	Surrogate Safety Assessment Module Parameters	295
11.3.4	Experimentation Structure	295

11.4 Results	296
11.4.1 Conflicts on Highways	296
11.4.2 Conflicts at Intersections	298
11.5 Discussion	304
11.6 Conclusion	308
12 Summation and Concluding Thoughts	309
References	317

1 Introduction

The transport network forms the lifeblood of a city, it connects the largest of central commercial districts with the most remote residential towns, from the tallest of skyscrapers to our beloved outback, nature, and wilderness. However, this vital resource is struggling to maintain its flow. Population growth, private car adoption, and an increasing desire to travel are all contributing to excessive congestion. Transportation systems have facilitated explosive growth in advanced and sophisticated interconnected cities. Urban growth has brought about the modern technological age, an era of both automation and convenience. Transportation, mobility, and the network have all benefited from refinement; from humble beginnings as horse-drawn carriages on dirt roads to modern iterations of Connected and Autonomous Vehicles (CAVs) on highly engineered and optimised pavements. Congestion is not a new problem, nor is it a problem with a myriad of options and economical solutions. The age-old strategy of build, expand, grow, sprawl, has worsened the compounding traffic problem. CAV adoption may be a means of addressing congestion, without expensive investment in new infrastructure.

Each advancement in mobility has aided in alleviating the problems of its predecessor. During the late 19th century, London alone used 50,000 horses per day. Regulators of the time were concerned with the generation of waste by-product associated with this mode of transit. The quantity was unmanageable and posed significant health risks, also affecting the sanitation of water and attracting pests and disease-carrying insects [Johnson, 2015]. The solution was not one of policy or behavioural change, but a new technology that fundamentally changed the landscape of mobility to a mode that did not exhibit the same limitations. Mass production of an affordable automobile arose as the most viable solution. The privately-owned automobile is now the cause of unsustainable fuel consumption and pollution. Human operators facilitate inefficiencies in network operation, leading to congestion and delay. Finally, safety concerns plague modern city transportation systems, with regulators unable to balance system efficiency against system safety.

As the automobile addressed the concerns of the horse-drawn carriage, society verges on a technological precipice where advancements in the clean-burning

electric CAV may be the solution to problems posed by the human-driven combustion engine automobile. CAVs are envisaged as self-driving vehicles capable of network navigation with little or no human intervention. While the road to automation started many decades ago, with the invention of the electric starter to replace the hand-crank in 1910 [Boyd, 1968], more recent advancements in automation have given rise to the possibility of completely driverless systems in the near future.

The consumer and commuter behavioural changes that CAVs may incite are difficult to predict, most predominantly due to limited empirical information. Potential improvements in travel speeds and times delivered by coordinated vehicle operation, instantaneous reaction times, and platooning may improve network performance and inadvertently increase travel demand. If capacity and network throughput do not increase proportionately, this scenario could lead to worsening congestion. Alternatively, if the private vehicle ownership model becomes antiquated and “Mobility as a Service” (MaaS) prevails, the sharing economy may reduce network demand through increased vehicle occupancy. The benefits of this outcome could couple with the potential advantages of CAV operation and provide substantial improvements for network performance. Predicting the performance of a mixed fleet or CAV network is difficult due to a range of uncertainties. Information availability is limited regarding the tangible impact of CAVs on capacity, throughput, and traffic flow. User attitudes toward vehicle ownership are also evolving, potentially reducing reliance on privately owned vehicles. Additionally, government legislative frameworks, manufacturing capabilities, and penetration rates increase the uncertainty in predicting future network performance.

Further difficulty in predictive efforts arise when considering the unique and intricate development of different cities. Geography, climate, culture, and wealth have all contributed to the development of localised transportation paradigms. Consider the United States, which has a historical focus on road and highway infrastructure construction. This investment in roads is the result of an expansive landmass and a monetary policy that used investment in road infrastructure to reduce the effects of the great depression by creating jobs and valuable long-term assets. Singapore however, faces resources scarcity and limited available landmass, freshwater, food, agriculture, and are catering for

an ageing population. These issues have forced Singapore to create planning policies that better incentivise mass transit. Parts of Europe, such as Paris and London, have developed the same way, with narrow roads in high-density urban spaces placing restrictions on speed limits and infrastructure allocated to motorways. An integration strategy for CAVs suitable for one geographic location may be inappropriate for another.

Industry and academia are reacting appropriately to this technological change. 60% of long-term large urban area transportation plans include CAVs to some degree, an increase from just 6% in 2015 [Cottam, 2018]. Infrastructure planning efforts involve using strategic macrosimulation and operational microsimulation modelling, to assess the performance of transport networks and the effects of future infrastructure development. Macrosimulation modelling generally relies on the four-step modelling process, grounded in probabilistic mathematical models that provide aggregate performance metrics. Microsimulation modelling uses mathematical models to emulate agent interactions in discrete time steps over a modelling time horizon. Both forms of modelling rely on a set of calibrated formulations that appropriately depict the decision-making process of the modelled agents. The literature contains a vast array of formulations used for modelling agents in a microsimulation environment; there is an absence of an appropriate and holistic framework for CAVs, capable of testing the impact of different CAV behaviour and implementation policies on network performance.

Thesis Objectives

This thesis develops a versatile algorithm capable of emulating the behaviour of CAVs in a microsimulation environment. The algorithm emulates idealised CAV behaviour, taking full advantage of the offerings of Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communication protocols. The former allows CAVs to communicate with one another, allowing them to operate cooperatively, manage microscopic space negotiation, and transfer information between themselves. The later allows CAVs to communicate with the wider network, collect information on downstream problems, receive optimal routing information, or take part in large-scale cooperative operation determined by a centralised controller.

Human drivers rely on rudimentary heuristics to process information while driving, allowing them to make quick, and often sub-optimal, decisions. While CAVs could behave in the same way, greater operational efficiency can be derived from V2V and V2I. The cooperation and communication delivered through V2V would allow instantaneous vehicle reaction times to external stimuli, or for vehicles to maintain a significantly shorter spacing. V2I allows vehicles to surrender operational control to a centralised framework, potentially improving fleet efficiency in space-negotiation activities such as merging and weaving. These vehicle characteristics form critical assumptions in the proposed CAV emulation framework, and so the capabilities of V2V and V2I must be accordingly understood. V2V and V2I will result in vehicle operation that greatly differs from that of human vehicles, meaning that existing mathematical model used for macrosimulation and microsimulation may not be applicable for mixed or fully autonomous fleets. For this reason, it is necessary to design a comprehensive modelling framework that reconsiders not only the three staples of microsimulation; car-following, lane-changing, and gap acceptance, but also reevaluates vehicle cooperation, altruistic behaviour, and acceleration and braking characteristics.

Thesis Contributions

Further contribution of this thesis lies in its use of the developed algorithm to evaluate the impact of CAVs in a range of different studies. The literature presents numerous detailed studies assessing CAV behaviour on the high-level intersection and network impacts, which are used in this thesis to benchmark the developed algorithm. In addition to benchmarking and performance validation, the intersection and network assessments provide further knowledge into the effects of cooperative behaviour on performance such as average travel times, throughput, capacity, and vehicle kinematics. The additional investigations conducted as part of this thesis are explained as follows.

The first investigation uses the algorithm to understand further the effects of CAV cooperation on a small network, explicitly focusing on vehicle kinematics and occupant comfort. Much of the literature has made assumptions regarding CAV behaviour without validating or assessing their consequence. For this reason, the first investigation in this thesis focuses specifically on the impact

of CAV behaviour on the vehicle kinematics, validating whether the use of such behaviour is appropriate. The next study then investigates intersection performance as a result of incremental CAV penetration.

The kinematic and network assessment is followed by a study that assesses the feasibility of reallocating infrastructure resources from vehicles to other agents in the network. This investigation determines the implications of re-optimising traffic signal times to transition from a vehicle-priority model to a pedestrian-priority model. Such a re-optimisation may be feasible once the CAV Penetration Rate (CPR) surpasses a threshold. The results of this investigation determine whether distributing travel time savings through CAV operation to other intersection agents, leads to a global intersection improvement in either the busiest or pedestrian environments and the most sparse. If benefits in neither environment are found for total system delay, this concept can naturally be dismissed. However, the contrary would provide conceptual validity and warrant further investigation.

The subsequent investigation assesses changes to highway capacity and merge junction throughput. Effects of CAV behaviour are assessed under a range of scenarios including altruism, selective cooperation, and a trajectory forecasting algorithm used to adjust vehicle kinematics before vehicle arrival at the motorway merge point. Network demand in this investigation is incrementally loaded to determine the impacts of induced demand and providing mobility to previously precluded members. This investigation intends to provide significantly greater insights into CAV behaviour for motorway weaving and merge zones, which are often the bottlenecks in interconnected motorway systems.

The following investigation in this thesis aims to address the parameters involved in macrosimulation mathematical models. While a microsimulation modelling framework is useful in determining network impacts, it is limited in its ability to draw generalised implications of larger city areas. For this reason, a macrosimulation model is often used, especially when evaluating the future long-term impacts of “State Significant” projects. This investigation first evaluates the applicability of current mathematical models and parameters to mixed fleet and CAV networks. If the models are found to be deficient, a recalibration effort is conducted to adjust the parameters for mixed fleet

environments. Finally, the predictive qualities of the recalibrated models are determined to ensure generalised results. This investigation provides a simple means of delivering continued use of traditional strategic modelling parameters for future networks with mixed fleets. Traditionally, strategic modelling parameters are obtained from the Highway Capacity Manual (HCM).

This thesis concludes with a safety assessment, quantifying the change in conflict likelihood on networks as a result of CAV integration. Many administrations worldwide are striving for zero deaths on roads. The initiative in New South Wales (NSW), Australia is called “Towards Zero” [Transport for New South Wales, 2018]. NSW has been progressing towards this goal through legislative change, decreasing speed limits, and imposing greater penalties either as fiscal fines or driving restrictions for minor infractions. Aggressive approaches of this nature may not be necessary if a safety assessment of CAV operation shows a reduction in network conflicts. This investigation uses vehicle trajectory data generated from microsimulation network operation to identify the likelihood of conflicts.

In summation, the contribution of this thesis lies in its development and use of a custom CAV behavioural control algorithm. The algorithm is then used in a range of studies to assess the consequences of CAV adoption for society, the economy, and the environment. The studies cover a wide range of complexities, where some aim to quantify changes to road capacity, junction throughput, and safety, while others aim to assess the feasibility of untested concepts such as the implementation of pedestrian-priority models at intersections. Ultimately, this thesis and its component studies quantify many of the unknowns surrounding road network development and will be a useful resource for the evolution of planning policies and legislative frameworks. The findings of the experimentation conducted in this thesis play a small but critical role in understanding the effects of this disruptive technology, and how best to incorporate it into our multi-modal transport paradigm.

Thesis Outline

Chapter 1 provided an introduction to this thesis, the remainder is structured as follows. Chapter 2, Chapter 3, and Chapter 4 provide an overview of CAVs,

microsimulation human vehicle emulation models, and CAV behavioural models, respectively. Chapter 5 provides the first major contribution of this thesis through the development of a model and algorithm used for CAV behaviour emulation in microsimulation, in conjunction with the emulation of human vehicles. Chapter 6 investigates CAV kinematics in a microsimulation setting, acting as validation for the developed models and algorithms. Chapter 7 investigates the impact of CAV integration on isolated intersection performance. Chapter 8 assesses the potential and impact of redistributing intersection delay through signal optimisation by transitioning from a vehicle-priority signalling regime to a pedestrian-priority regime. Chapter 9 assesses the impact of CAV integration on motorway capacity and weaving and merge junction throughput. Chapter 10 investigates the applicability of strategic modelling parameters to mixed fleet networks, the scope of recalibration of these models, and the predictive quality of the recalibrated models. Chapter 11 investigates changes to network safety as a result of CAV integration, assessing conflict likelihoods. Finally, Chapter 12 concludes with a summation of the work and concluding remarks.

2 Connected and Autonomous Vehicles, A Review

The automotive industry leverages technological and methodological advancements to better refine their product. The automobile has slowly transitioned from a purely mechanical system to one with well-integrated automated processes. The rudimentary self-propelled steam-powered vehicle (1769) has been transformed by strides in power delivery through the internal combustion engine (1807), in ignition through the electric starter (1911), and in pressure delivery through hydraulic braking and steering systems (1922 and 1926), to its modern counterpart staple of mobility and movement.

Automation provides independence and self-sufficiency to a system, granting it decision-making capabilities with varying degrees of complexity [Shladover, 2018]. Major developments in vehicle automation began with cruise control in 1950. Cruise control allowed a vehicle to maintain a set speed indefinitely. Antilocking braking systems in 1985 meant that vehicles could also autonomously regulate pressure delivery to the wheels during emergency braking events. Global Positioning Systems (GPS) in 1992 automated route selection and navigation. Most modern vehicles are fitted with assistive features such as rear-view cameras (2009), lane identification and lane-keeping, adjacent vehicle detection, adaptive cruise control (ACC), and autonomous parallel parking capabilities. These advancements have made the private vehicle faster, safer, more efficient, and more reliable.

These autonomous features are passive, restricting their operation to only alerting a driver. Driving requires functions that extend beyond decision-making, as drivers must also acquire information from their surroundings in support of environmental perception. A vehicle that is capable of doing both is considered autonomous. Systems that aid in environmental perception are explained in detail later in this section. To extend the nuance of the nomenclature, a vehicle that communicates either with other vehicles, the network, or a cloud server is considered connected. A vehicle can be automated without being connected, and vice versa. However, the general public and industry expect that the natural maturity of this technology will result in vehicles that are proficient at both.

The overuse of terms such as “autonomous vehicle” and “self-driving car” has obscured their specific meaning. Without reservation, these terms have been applied to CAV technology in all areas of the complexity spectrum. When discussing CAV technology, the vehicles are generally divided into six levels ranging from L0 (completely human) to L5 (completely autonomous). The degree and extent of the permitted driver disengagement define the level of the CAV. Lower order CAVs allow short periods of disengagement in specific settings. Higher-order CAVs allow sustained disengagements in a broader array of settings and driving scenarios. Similarly, the degree of connectivity of a vehicle is defined by the quantity and type of information it exchanges. For example, low-level connectivity would constitute uni-directional communication, where a vehicle passively receives information from its surroundings. Higher-level connectivity would constitute bi-directional communication where the vehicle can return kinematic, location, or other information to the system

The applications of V2V communication include [Shladover, 2018];

- Cooperative collision warnings and hazard alerts, collision mitigation, and active braking.
- Cooperative adaptive cruise control, with tighter vehicle-following and enhanced traffic flow stability.
- Close-formation automated platooning, enabling aerodynamic drafting.
- Automated manoeuvre negotiation at merging locations or intersections.
- Transit bus connection protection.

These applications are time-critical and safety-oriented, meaning that the communication must be fast, low latency, and reliable. V2I communication can enable the following additional features and policies [Shladover, 2018];

- Providing traffic signal status information in real-time for in-vehicle display, signal violation warning, or green wave speed advisories to drivers, extended to traffic signal priority requests.
- Providing traffic and weather condition information and real-time routing advisories to drivers.
- Fleet management functions of vehicle routing and scheduling.
- Access control to closed facilities.

- End of queue warnings.
- Active support for lane guidance.
- Vehicle probe data applications to provide detailed traffic information (speed, volume, travel time, queue length, and stops) or road surface condition information (pavement roughness, or slippery conditions).
- Mayday and concierge services (such as OnStar).
- Electronic toll collection and parking payments.

The remainder of this chapter begins with a detailed explanation of the CAV, modern hardware that allows it to operate, and what capabilities this implies. This information is critical in justifying the CAV behavioural assumptions embedded in the microsimulation algorithm developed in Chapter 5.1. This chapter then outlines CAV control software and concludes with an outline of recent trials worldwide using CAVs, and aggregate insights obtained.

2.1 Hardware and Capabilities of the CAV

The hardware of this emerging technology is under development, though the essential elements remain consistent between manufacturers. The hardware is tasked with three primary duties; Exteroceptive Sensing, Proprioceptive Sensing, and Communicative Sensing [Campbell et al., 2018], and is simply shown in Figure 1;

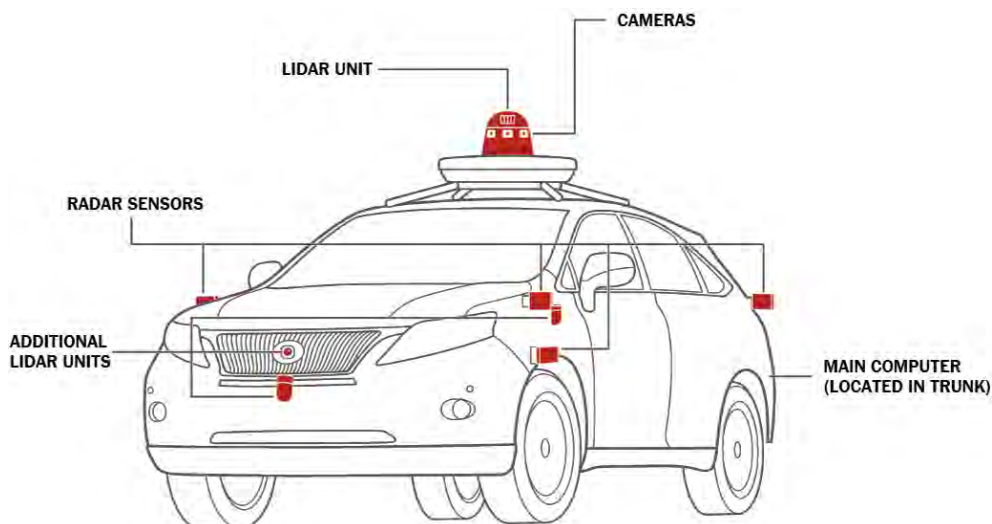


Figure 1: Remote sensing hardware of the CAV [Gates, 2017].

Exteroceptive Sensoring

Exteroceptive Sensoring is used for environmental perception. This family of sensors creates a digital image of the vehicle's locale, allowing it to calculate distances to surrounding objects. Sensors in this family include LiDAR (Light Detection and Ranging), RADAR (Radio Detection And Ranging), optical cameras, and ultrasonic devices.

LiDAR measures distance by emitting a pulse of light and measuring its return time. The vehicle's surroundings are visually mapped by emitting at a $150kHz$ frequency (150,000 pulses per second) up to a distance of $250m$. Precision and accuracy are altered by adjusting the pulse frequency and the fidelity of the return pulse measuring equipment. The cost for this detailed 360-degree visibility is high, with the development of solid-state LiDAR and infrared LiDAR potentially reducing costs. Solid-State LiDAR integrates many of the mechanical components of a traditional LiDAR into a single microchip, resulting in a smaller, cheaper, easier to calibrate, less likely to fail, and easier to mass-produce product [Higgins, 2015].

RADAR uses radio waves to measure the distance, angle, and velocity of objects. The RADAR emits radio waves at different frequencies, with higher frequencies providing higher fidelity information. RADARs are appropriate for short-to-medium range, serving approximately $50m$ - $100m$. They complement the information gathered from LiDARS as they are cheaper, easier to work with, and already fitted in modern vehicles for features such as adaptive cruise control and collision detection.

Cameras produce a digital image through passive light sensors. Cameras are advantaged in their ability to detect colour and texture, allowing CAVs to perceive their environments better. Where a LiDAR or RADAR can determine the presence of a traffic light or street sign, a camera can further determine its colour or the specificities of the instruction. They allow a CAV to analyse their position relative to road markings. A camera is cheap and readily available, but the computational power needed to infer relative distances and velocities compared to LiDAR is very high. Also, camera sensors are sensitive to low-intensity light and may not work appropriately in poor weather conditions.

Ultrasonic: measures distances to nearby objects by emitting a sound wave at a specific frequency and measuring the time delay in detecting the return wave. They are generally the cheapest of these sensors, work well in poor weather, and are regularly used as parking aids. They are regarded as the most accurate for close-proximity distance measurements but have multiple sources of error for longer distances. Constructive or destructive interference can alter the nature of the ultrasonic sound wave. Wave characteristics are also subject to the temperature and humidity of the medium through which they travel, although, most ultrasonic sensors algorithmically correct for errors.

Proprioceptive Sensing

Proprioceptive Sensing calculates information regarding the vehicle system. This family of sensors monitors the vehicle, calculating its motion, kinematics, and detailed location. The data informs future movement of the vehicle, including motor speed, wheel position, and joint angles.

GPS is a satellite-based radio navigation system that provides Ephemeris and Almanac data used for geolocation and time. GPS uses trilateration to determine location by concurrently requesting position information from at least four satellites. The location of the GPS receiver is determined by superpositioning the information from all connected satellites and calculating the overlapping region. Commercial-grade GPS can be accurate to within $1m$, but all GPS receivers need a direct line of sight and can be obstructed by tall buildings and trees in urban environments. Survey grade GPS receivers are accurate to within centimetres but cost significantly more.

Internal Measurements Units (IMU) measure forces using three accelerometers, angular rates using a gyroscope, and magnetic fields using magnetometers. By placing these sensors orthogonally to one another, the IMU makes measurements in all three directions of motion. The IMU can only provide information about the motion of a vehicle and not its location. The Internal Navigation System (INS) uses data gathered from the IMU to calculate velocity, altitude, and angular positions. The sensors suffer from accumulated error and lead to drift, which is corrected using the vehicles GPS.

Encoders are electro-mechanical devices that convert the linear or angular movement of a shaft or rotor, into a linear or angular transduction. Encoders are generally attached to vehicle wheels and odometers to measure distances and velocities in real-time. They are susceptible to errors due to wheel slippage in low traction environments. Encoders are combined with an INS and GPS to provide low-cost location information.

Communicative Sensoring

Communicative Sensoring exchanges information between the system and surrounding vehicles. This family of sensors is used to transmit and receive information about the vehicle or surrounding agents. They also provide digital, visual, and auditory cues to other agents. Techniques for Information Technology Systems (ITS) and communications include;

- 5.9 GHz Dedicated Short Range Communication (DSRC) is a technology that is similar to Wi-Fi but is designed specifically for transportation applications. DSRC transmits time-critical and safety-critical messages over a small licenced and protected portion of the electromagnetic spectrum with low latency.
- Wi-Fi can be used in a supportive role for minor ITS functionality. It has high connection latency and experiences message degradation through packet loss and delays when communication channels are congested. For this reason, it can not be used for critical information communication.
- Cellular refers to 4G or 5G telecommunications and their future iterations. Cellular has the advantage of a privately developed and maintained network of infrastructure, requiring little involvement from public entities. Private ownership and maintenance allow telecommunications to be used in most transport networks without developing new and dedicated infrastructure.
- Satellite is used in remote locations where cellular service is not available, and the provision of communications infrastructure is technically or fiscally infeasible. Latency issues with Satellite communications mean they are not appropriate for all applications.
- Bluetooth can provide only very short-range and low-bandwidth communication, so it should only be reserved for a supportive role.

The CAV Ecosystem and Perceived Impacts

Junior, the 2007 Defence Advanced Research Programs Agency (DARPA) Urban Challenge entry by Stanford [Levinson et al., 2011] used similar early-stage versions of this hardware. Objects surrounding the trajectory of the vehicle were detected using the Velodyne HDL-64E S2 rotating 64-beam LiDAR. This LiDAR maps 2.2 million points per second, ideal for generating digital contour maps of the vehicle’s surroundings. The LiDAR was positioned atop the vehicle, giving it a detection range of 120m with an accuracy of 2cm [Velodyne Corporation, 2019]. Numerous RADARs positioned on the front, rear, and sides of the vehicle attained information regarding nearby objects. A SICK LD-LRS LiDAR scanner on either side adequately covered blind-spots while an Applanix POS LV 420 internal GPS guided the vehicle. This system is used for geospatial data acquisition and can provide geolocation to within 20cm [Applanix Corporation, 2015]. The hardware was controlled and operated by an onboard Xeon computer, running a Linux operating system. A 12-core server executed the vision and laser algorithms, while a 6-core server handled planning, controls, and low-level communication. The vehicle used the powerful onboard computing capabilities to consider a wide range of possible manoeuvres and react quicker compared to its human driver counterpart.

CAVs have access to a significant amount of data that human drivers are otherwise ignorant to, or unable to process. CAV hardware communicates with devices embedded in the greater network through V2V and V2I (generalised as V2X). V2X communication is based on the communication standards defined for Wireless Access in Vehicular Environments (WAVE) under the Institute of Electrical and Electronics Engineers (IEEE) standards 1609 (dedicated short-range communication) and 802.11p (wireless local area network, medium access control, and physical layer communications) [Milanes et al., 2012]. WAVE systems comprise of three elements, the roadside unit, the onboard unit, and the service channels.

The roadside unit is installed in existing transportation infrastructure such as traffic lights and signals. Milanes et al., used the roadside unit as a control station and a means of data processing for their simulation of an unsignaled intersection [Milanes et al., 2012]. The roadside unit was expected to “classify

the vehicles within the driving area, identify stopped vehicles as obstacles, inform each vehicle about its vicinity, provide information in advance about the road layout, manage each specific area using priority levels according to the kind of vehicle to avoid potential collisions and to identify in advance traffic situations of risk”. Each vehicle provides the roadside unit with data, which it then processes, and returns safe paths and manoeuvres for the vehicle. The centralised controller prevents collisions and the formation of barricades through the intersection. Figure 2 shows a hypothetical rendering of the three elements of V2V and V2I systems in operation;




Figure 2: WAVE communication systems, comprising of the roadside unit, the on-board unit, and the service channels [Mercedes-Benz, 2014].

The control system maintains a data exchange between the vehicles, reducing communication packet loss, and ensuring vehicle safety. Milanese et al., showed that even with 100 stations competing for a single channel, the delay in communication for low priority vehicles was $515\mu s$, near negligible and insignificant for real-time traffic control. The onboard unit maintains connectivity between the CAV and the roadside unit. The service channels transmit information between the vehicles (V2V) or between the infrastructure (V2I). The small lag in information transfer allows the microsimulation modelling discrete time step to be reduced to its computational minimum.

Despite the myriad of complementary sensors compensating for each other’s weaknesses, further development and testing is necessary. Perception in poor weather and unfavourable lighting conditions is problematic as both vision-based systems and LiDAR-based systems are unable to see road markings. Dense rain and snow can cause “phantom obstacles” to be detected, making distance measurements to obstacles inaccurate. Lens-flares, large shadows, and poor lighting can disrupt data acquisition from visual sensors, with a priori information potentially outdated [Van Brummelen et al., 2018]. The limitations of sensors need to be well understood. Long-range automotive RADAR has a typical range of $4.5s - 7.5s$ and $150m - 250m$ at highway speeds of $120km/h$ [Al Henawy & Schneider, 2011]. The Insurance Corporation of British Columbia instructs human drivers to look at least twelve seconds ahead to ensure safe driving [ICBC, 2015]. Often, redundant sensors validate and correct one another, aiding in faulty sensor identification through an expensive and inefficient means [Ding, 2005]. Modern methods employ cross-checking, where a sensor’s current behaviour is benchmarked against its behaviour during fault-free operations [Heredia et al., 2008].

The network resulting from the amalgamation of civil transport infrastructure and technological innovation will be unlike any that has existed before, differing from its traditional and contemporary form, which has formed the status quo for the last century [Hancock et al., 2019]. How might a future city operate when 30% of its downtown traffic no longer circle the city in search of parking [Shoup, 2007], or the notion of private vehicle ownership is antiquated? Many theories have been posed as to what the implications are for a future centred on autonomous transport. Some are described in brief as follows.

Transport Logistics (food): Trucking and freight accounted for 28% of the food supply chain, and caused 71% of its emissions. Also, fresh produce lost 8% of its weight in the transportation process. Wastage is reduced by automating mechanical processes, optimising routing, less idling of engines, and less engine cold starts by long operating autonomous drivers [Heard et al., 2018]. The platooning of trucks is also estimated to reduce energy intensity by 10% - 15% [Wadud et al., 2016]. Even minor reductions in the cost of transporting goods could cater to a previously uneconomical unfulfilled demand.

Safety: The National Highway Traffic Safety Administration indicates that human-related errors such as distraction, fatigue, and emotional driving account for 94% of accidents [Van Brummelen et al., 2018]. Data obtained from the Australian Bureau of Statistics (ABS) and Transport for New South Wales (TfNSW) indicates that negligent behaviour such as driving under the influence, fatigue, and speeding comprises of 31.3% of accidents [Transport for New South Wales, 2017]. Automating the driving process should also mitigate human-caused error and accidents.

Increased Vehicle Demand: Subject to legislation, CAVs could provide mobility to a segment of the market previously precluded from vehicle ownership or operation (elderly, disabled, young, inebriated), which could increase road demand. Additionally, reallocating in-vehicle travel time to productive work could reduce the penalty that operating a vehicle in congestion poses on the decision to travel or selection of mode and route [Adnan et al., 2018].

Effects on Public and Active Transport: In a shared mobility economy, narrower car lanes, less on-street parking, and fewer roadways would increase the utility of travelling by public and active transport modes [Krechmer et al., 2016]. Alternatively, an improved vehicle travel experience may accelerate urban sprawl, with larger trip distances being infeasible by public and active transport [Litman, 2017]. Additionally, the safety of active transport users is under scrutiny, as studies have shown that CAV sensors are unable to detect pedestrians and cyclists, nor can they accurately predict the future intentions of these agents [Botello et al., 2019]. This problem is made worse by the “enhanced immunity fallacy”, where another agent believes that enhanced CAV technology will maintain safety [Harper et al., 2016b].

Ethics and Morality: Minor errors in sensors and hardware can lead to potentially fatal incidents. While the ideal outcome is for a CAV to react well in advance and eliminate the need for catastrophic accident mitigation, their programming still needs to accommodate ‘mayday’ circumstances. The source of failure may also be remote, through malicious interference, as these technologies are susceptible to manipulation [Hashim & Omar, 2017].

Over-Reliance on Machines: An over-reliance on machines may lead to a degra-

dation of critical skills and loss of competencies [Bazilinskyy et al., 2015]. It can also lead to an exacerbation of the obesity crisis. A distrust and over-trust in machinery can lead to disuse and abuse, respectively.

Servicing the First and Last Mile: The-first and last-mile problem relate to the time, effort, and delay involved in trips that originate (first mile) or destinate (last mile) in a populated city centre. Transportation, freight, and logistics trips expend significant manual labour in the first and last mile, with proportionally less progress than the remaining trip. By automating this portion of the trip, greater economic output is derived from activities constrained by traversing city centres.

This thesis intends to provide validity and numerical backing through microsimulation to these theories, and test the implications of a range of other policies and applications of CAVs. Real-world trials are sparse, expensive, and timely. By creating a microsimulation framework capable of emulating mixed fleet and CAV behaviour, policy and integration structures can be mass-trailed in theoretical feasibility studies. Experimentation with frameworks, policy, and regulation would outpace the development of the technology, which does not often happen with legal frameworks left to play catchup. A microsimulation framework is critical as full-scale trials would otherwise be too dangerous or expensive to attempt while CAV technology develops from its infancy.

2.2 Software Control of CAVs

The previous section extensively explored the innovations in hardware that allow a CAV to operate. In conjunction with hardware, the software algorithms that control dataflow and interpret information play a vital role in CAV operation. Understanding the algorithms used to control real CAV behaviour can aid in developing microsimulation control algorithms. The remainder of this section explores recent developments in CAV control software and algorithms.

While industrial development of CAV control software remains commercial and confidential, the literature is proposing a range of techniques to bridge the gap between data acquisition from hardware, and response action by the CAV. Zhou et al., propose a reinforcement-learning based approach to control driving behaviour and improve real-time traffic efficiency, vehicle fuel consumption,

and driver safety [Zhou et al., 2019]. Reinforcement learning places an agent within an environment from which it receives feedback. At any given state, the agent receives all necessary inputs, computes a probabilistic action, and then receives a reward. The nature of the reward reinforces positive decisions and punishes incorrect decisions, allowing the agent to learn before transitioning to the next state.

Zhou et al., developed the Deep Deterministic Policy Gradient algorithm to determine the appropriate reward system for positive behaviour. Using observed data during runtime, the algorithm updates the critic and actor, which are responsible for evaluating the vehicles current circumstances, and updating the policy (reward) function, respectively. Using this framework, the authors demonstrated that CAVs could be trained to intelligently decelerate in response to both preceding vehicles and the status of traffic lights. The multi-factor consideration to appropriate acceleration rates led to gentler and safer decelerations as CAV platoons approached signalised intersections. Additionally, when observing environments with varying congestion, their framework demonstrated that the effect on average speed was slight, meaning that travel efficiency remained high regardless of the level of congestion.

Ma et al., developed a framework centred on the Shooting Heuristic algorithm to construct feasible CAV trajectories that optimise traffic performance [Ma et al., 2017]. The Shooting Heuristic algorithm smooths traffic flow by conducting speed harmonisation and reducing the occurrence of sudden speed drops that lead to start-stop conditions. The algorithm simplifies vehicle trajectories through discrete quadratic representation and then reconstructs new trajectories subject to physical limitations, safety, and signal timing.

Determining system performance used in the Shooting Heuristic is an infinite dimension problem with non-linear objectives and constraints, making its exact solution computationally difficult to determine. For this reason, the authors develop a framework to optimise system performance. The algorithm begins by initialising a set of control parameters, passing them to the Shooting Heuristic algorithm, and calculating a preliminary trajectory. System performance based on the preliminary trajectory is optimised by the Shooting Heuristic first using simplified analytical approaches. The trajectory is

then passed to a dedicated optimisation algorithm, which determines if the preliminary trajectory was already optimised. If not, it then uses a rule-based process to update the control parameters and restarts the process in the next iteration. Using this framework, Ma et al., demonstrated that the optimised CAV trajectories improve manual trajectories by 36.7%, 39.7%, and 67.1%, in terms of travel time, fuel consumption, and safety, respectively.

Li and Li addressed the limitations of previous trajectory optimisation algorithms by providing a complete solution that does not require runtime heuristics and simplifications, guaranteeing the optimality of the solution [Li & Li, 2019]. Constraints on the algorithm include acceleration bound by mechanical limits, positive monotonic speed restriction, communication distances remaining below a reliable threshold, and boundary conditions that explain behaviour during an initial time and final position. A meaningful trajectory can be computed by using a finite number of piece-wise quadratic equations, allowing the error to remain arbitrarily small and the trajectory to be optimised in accordance with the mechanical and spatial constraints.

Using this approach, the authors simulate a jam occurring in a platoon of seven vehicles. The simulation demonstrates that stable platoon behaviour reoccurs after 20s. When travelling at an initial velocity of $12m/s$, the platoon is able to absorb the lead vehicles deceleration without having to exceed maximum deceleration limits. The developed framework is expandable by adding additional constraints such as limitations on jerk, considering vehicle kinematic characteristics, assessing the impact of heterogeneous acceleration rates, and impacts of communication or implementation delay.

2.3 Recent Trials using CAV Technology

This section outlines in brief recent trials of CAVs. Assessing CAV hardware provides extensive insight into their ultimate capabilities. Results from real-world trials offer valuable access to information regarding limitations imposed by transport systems, engagement with other agents, or technical barriers not otherwise considered. The trials outlined below were conducted by vehicle manufacturers, government entities, and private enterprise. Unfortunately, more often than not, the results and lessons learnt of such trials are not publicised. Trials include;

- Ford and Baidu: Ford and Baidu have partnered in China to trail 45 L4 CAVs on 31*mi* of open road. This extent will be increased to 84*mi* within a year [Szymkowski, 2019]. Baidu is also the developer of V2X technology and aims to use this trial as a means of testing both CAVs and their communication framework.
- Queensland and Renault: The Queensland government is using a AUD\$1.5 million custom-built Renault Zeo 2 to better understand how CAVs will interact with the built environment. Additionally, the trial aims to determine what infrastructure changes are required to make the State CAV-ready [Schmidt, 2019].
- United Kingdom (UK): The UK federal government is investing GBP£25 million into three CAV trials for 2021, with many smaller CAV trials already underway. The vehicles are intended to be options for those that struggle to access public transport, and a means of revolutionising goods movement and freight logistics [Williams, 2018].
- Minnesota: Minnesota’s challenging weather allows the city to gather a unique perspective on the operation of CAVs. Trials in the state focus on the performance of sensing equipment and vehicles in snowy conditions and iced roads, the number of trials in which are sparse. The trials begin with testing operations on roads by the State authority and gradually increase in complexity to testing V2I connectivity and impacts on transit priority and fuel efficiency [Hietpas, 2017].
- Intel: Intel is working with the Israeli government to test their ‘Mobileye’ platform in Jerusalem. The trials test safety in heavily congested and aggressive driving environments. Mobileye aims to provide MaaS with vehicle, software, platform, mapping, and safety mechanisms. By trialling vehicles in unfamiliar situations, Intel hopes to gain better data to train its artificial intelligence algorithms [Reichert, 2019].
- Perrone Robotics and Virginia: Perrone Robotics equipped a vehicle with L4 autonomy and trialled it on Virginian roads. This trial progressed the complexity of public trials by allowing the vehicle to autonomously navigate complex roundabouts, intersections, and areas with pedestrians and cyclists. During the trial, there were no safety incidences, the vehicle did not forfeit control to a human at any point, and the service received a positive reception from the public [Perrone, 2019].

- Navya: Navya is a French CAV manufacturer focusing on different vehicle solutions for segments of the transportation market. In July, Navya was permitted by Japan to operate an autonomous bus without a steering wheel, pushing CAV trials closer to L5 autonomy [Navya, 2019].
- Drive Me by Volvo: The Drive Me project aims to align the motivations of the public, private, and academic sector to accelerate the development of CAVs. Headed by Volvo, the project initiated in 2013 and has now progressed to on-road trials of prototype commercial vehicles in Gothenburg, Sweden [Volvo, 2017].

This list serves only as an introduction to the number and diversity of CAV trials currently being conducted worldwide. iMOVE maintains a database of currently 104 entries that describe worldwide CAV trials [iMOVE, 2020]. Each trial aims to provide insights into a previously untested environment for CAV operation or collect data to strengthen and train learning algorithms. While many trials have operated successfully, others have resulted in unfortunate accidents and highlighted the current limitations of this technology.

The use of Tesla’s “autopilot” feature, though impressive, is not flawless. The first death using Tesla’s self-driving features occurred in 2016 when a vehicle operating autonomously collided with a tractor-trailer at an intersection. Reportedly, the vehicle’s sensors were unable to distinguish between the colour of the trailer and the colour of the overshadow sky, resulting in a side-on collision and death of the driver [Corfield, 2017]. This incident highlighted the importance of positioning the sensors on vehicles in appropriate locations. The fatality could have been avoided by an alert driver prepared to resume control of the vehicle, and a sensor system that was not misguided by poor lighting and colour conditions. Though Tesla makes no claims that their vehicles are self-driving, that is the implication and eventuality of autopilot.

The first death of a pedestrian occurred by an Uber self-driving vehicle in 2018. The pedestrian was walking a bicycle across a four-lane road when the CAV operating in autonomous mode struck and later killed the pedestrian. This fatality could have also been avoided by an attentive operator, or by a pedestrian crossing within the confines of a legal crossing area [Levin, 2018]. This incident brought to light two key findings, the first being that Uber’s CAVs are unable to identify jaywalkers. The system did not contain a capability to

determine when pedestrians may decide to inattentively and illegally cross the road. The other lesson from this incident is that the behaviour of pedestrians and cyclists is far more stochastic than the rule-based and relatively deterministic behaviour of vehicles. CAVs still require substantial development in understanding the probabilistic behaviour of other agents in the network and be prepared to react to changing circumstances.

Real-world testing has been limited to simple network operations and vehicle interaction, with trials restricted by a lagging legislative framework. This is where work such as this thesis provides significant contributions. By creating a digital environment in which to emulate CAV behaviour and conduct scenario testing, a range of different manufacturing decisions, CAV behaviour, implementation strategies, and optimal use cases can be evaluated. Though limited in its capabilities, microsimulation testing is faster, safer, cheaper, and can be conducted without awaiting the legal framework accompanying CAVs to progress.

2.4 Summary

This chapter detailed the nature of a CAV, the hardware it is comprised of, and the technical capabilities that this hardware enables. Understanding CAVs is critical as their behaviour can not be emulated without understanding their technical limitations. A review of recent progress in commercial CAV development indicates that while the hardware currently has limitations in reliability and guaranteeing safety, they enable vehicle reaction times and safety many order of magnitudes greater than what humans are capable of. Additionally, the technical capabilities facilitated by V2V and V2I protocols allow degrees of cooperative and decentralised movement that are otherwise unachievable by a human fleet. A framework developed for emulating CAVs must be centred on many of these benefits, while also providing flexibility and expandability for when further technological advances increase CAV capabilities.

3 Mathematical Models for the Microsimulation of Human Driven Vehicle Behaviour

Chapter 2 outlined the capabilities of CAV hardware. Understanding the technical limitations places appropriate constraints on behaviour emulated in the microsimulation environment. A microsimulation setting is a valuable tool for assessing the performance of transport networks, or in evaluating the traffic impact of proposed developments and network alterations. This style of modelling generates individual agents and allows them to interact within the confines of a virtual environment.

Three major algorithms control the behaviour of vehicles in microsimulation; the car-following model, the lane-changing model, and the gap-acceptance model. These time-discrete algorithms dictate the behaviour of each agent in the network for each time increment. This chapter, in brief, outlines the development of these algorithms and mathematical models, and the data used to emulate human vehicles in a microsimulation setting. A Literature review of human-oriented models is essential as a well-developed model may already apply for CAV behaviour. In the absence of an appropriate model, a literature review would highlight model features to avoid when developing a modelling framework for CAVs. This literature review explains each family of models and then provides limited examples using impactful or unique contributions.

3.1 Car-Following Algorithms

Car-following models play a vital role in ensuring that microsimulation models reproduce real traffic flows and movement. The literature contains numerous car-following models, each offering a nuanced and novel means of accurately modelling the car-following behaviour of human motorists in a range of driving conditions [Brackstone & McDonald, 1999]. Car-following models are generally categorised into one of four types. Gazis-Herman-Rothery (GHR) models are built on the principle of stimulus-response, where a follower reacts to a leader. Collision Avoidance models are commonly developed using Newton's equations of motion and aim to maintain safe following distances. Psycho-Physical models state that drivers react to small changes to speed based on external cues, and discretise the behavioural domain into regions. Fuzzy Logic models use rule-based logic to discretise driver behaviour and conditionally

apply actions in microsimulation [Olstam & Tapani, 2004]. Additionally, an example of a recurrent neural network car-following model is also provided, which is a model style generally used by the literature in emulating CAVs, but can also be applied to human behaviour. This section explains each type of model in greater detail.

3.1.1 Gazis-Herman-Rothery Models

GHR models are built on the proposition that a driver’s acceleration is proportional to the following distance and velocity difference between the ego and lead vehicle. This type of model controls the behaviour of the ego vehicle by quantifying its response relative to its leader [Olstam & Tapani, 2004]. GHR models take the following form [Brackstone & McDonald, 1999];

$$a_n(t) = cv_n^{n+1}(t) \frac{\Delta v(t - \tau)}{\Delta x(t - \tau)} \quad (1)$$

Where, $a_n(t)$ is the acceleration of vehicle n at time t , Δv and Δx are the difference in speed and location between the ego and lead vehicle, respectively, τ is the time increment for the time-dependant model, and c is a calibration constant.

Though most GHR models maintain the assertion that car-following behaviour is proportional to velocity and headway, they vary in the complexity of this relationship. Δv and Δx may themselves be functions of other parameters that influence human behaviour. Outlined in brief below are examples of car-following models developed using the GHR approach.

Greenberg Model

The Greenberg model is one of the earlier derivations of the fundamental diagram and is based on the theory of the continuity of a compressible fluid [Greenberg, 1959]. Greenberg proposed that traffic behaves like a continuous fluid and so may be modelled using the theory of fluid dynamics. He did, however, state that this approach may not be appropriate for low traffic density scenarios. His fundamental diagram model is defined as follows;

$$v = c \cdot \ln(\rho) \quad (2)$$

Where, v is the velocity of the vehicle, c is a constant sensitivity calibration parameter, and ρ is the normalised traffic flow.

Though the Greenberg model is capable of replicating the high-level behaviour of vehicles in an uncongested network, it does not capture the large number of influential factors that impact driver car-following behaviour. By its nature, the fundamental diagram accurately represents aggregate fleet responses to density and can holistically replicate the behaviour of a stream. However, when circumstances deviate from the norm such as variations in driver characteristics or obstructions in vehicle trajectories, the model is unable to reproduce this naturally occurring stochasticity in busy environments.

Grazis et al., Model

Grazis et al., altered the Greenberg model by adjusting the sensitivity of lateral acceleration, making it proportional to the headway rather than a constant value [Gazis et al., 1959]. The authors theorised that headway played a psychological role in governing the acceleration of a driver. Their alteration is presented in the following expression;

$$v = f(\Delta x) \cdot \ln\left(\frac{1}{l \cdot \rho}\right) \quad (3)$$

Where, v is the velocity of the vehicle, $f(\Delta x)$ is a hyperbolic function of the headway that governs sensitivity, l is the vehicle length, and ρ is the density of the traffic flow.

This change captures a deterrence to acceleration when headways and vehicle spacing is small. By doing so, the model introduced localised fluctuations in the traffic stream as drivers now reacted to individual and varying space headways. Experimental verification showed that this alteration better replicated the test data than Greenberg’s original model.

Zang et al., Model

Zhang et al., evaluated the impact of the lead vehicles tail light on the driving behaviour of the ego vehicle [Zhang et al., 2018]. Their model is an extension of the Full Velocity Difference (FVD) model [Jiang et al., 2001], which is lim-

ited in its ability to describe non-standard phenomena in traffic streams, but otherwise has a compact form and performs well under standard conditions. Zhang et al., proposed that a vehicle spacing exists wherein the tail light of the preceding vehicle affects driving behaviour. Using this assumption, they expanded the FVD model to the following expression;

$$a_n(t) = f(v_n(t), \Delta x_n(t), \Delta v_n(t), \dots) + \theta_n(t)f(n, t) \quad (4)$$

Where $a_n(t)$ is the acceleration of vehicle n at time t , $v_n(t)$ is the velocity, $\Delta x_n(t)$ is the spacing between the ego and lead vehicle, $\Delta v_n(t)$ is the relative velocity, $\theta_n(t)$ is a binary term to indicate whether the tail light of vehicle $n-1$ is red, and $f(n, t)$ is a conditional function designed to determine whether the state of the lead vehicles tail light should impact the behaviour of the ego vehicle. $f(n, t)$ is defined as follows;

$$f(n, t) = f_0 \tanh\left(1 - \frac{\Delta x_n(t)}{\Delta x_n(0)}\right) H(-\Delta v_n(t)) \Delta v_n(t) \quad (5)$$

Where, H is the Heaviside function which equals 1 if $x \geq 0$ and 0 otherwise, and $\Delta x_n(t)$ is the spacing between the lead and ego vehicle.

Zhang et al., used a loop road of 1.5km length and 100 cars, then introduced a perturbation for 5s . Using the FVD model, vehicles experienced a speed change of 0.2m/s to 13.3m/s . When using the altered model with $f_0 = 0.5$, the exhibited speed change varied between 2.74m/s and 7.2m/s . This 34.4% reduction in the oscillation amplitude showed that the proposed model assists in reducing erratic behaviour and smoothening traffic flow.

Spring-Mass-Damper-Clutch Model

Recent developments in GHR models have drawn inspiration from other areas of physics and mechanics. Li et al., developed a car-following model using the mechanics of a “Spring-Mass-Damper-Clutch” system, where the ego vehicle driving behaviour also influences the lead vehicle [Li et al., 2019]. The non-linear wave propagation analysis technique used in their study means that a traffic stream can be modelled as multiple chained “Spring-Mass-Damper-Clutch” systems, allowing the proposed car-following model to be scaled and used for macroscopic assessment. Their car-following model is as follows;

$$m_n a_n(t) = k_n(\Delta x_n(t - \tau) - X_o(v_n(t - \tau))) + c_n \Delta v_n(t - \tau) \quad (6)$$

Where, m_n is the mass of vehicle n , $a_n(t)$ is the acceleration at time t , $v_n(t)$ is the vehicle velocity, τ is the spring dampening coefficient, X_o is the spring speed-dependant relaxation length, and c_n is the clutch time delay parameter.

Under standard driving conditions like on a highway, the velocity of the vehicle is bound between the maximum and minimum spring relaxation length, meaning that $X_o = x'_n \cdot v_n$, where x'_n is the gradient of the speed-dependent relaxation length curve. Equation 6 then becomes;

$$a_n(t) = \frac{k_n}{m_n(\Delta x_n(t - \tau) - x'_n v_n(t - \tau))} + \frac{c_n}{m_n \Delta v_n(t - \tau)} \quad (7)$$

This model, like the other GHR models, is built on the inherent assumption that drivers have an infinite perception and diminishing reaction to external stimuli. The Li et al., model fundamentally incorporates wave propagation as a means of modelling traffic streams. While both assumptions may be appropriate for human driving fleets, neither will suffice to model the instantaneous reaction times or platooning behaviour of CAVs.

3.1.2 Collision Avoidance Models

Collision avoidance models differ to GHR in that the car-following behaviour is not external stimuli dependant, but in response to creating and maintaining a safe distance from the lead vehicle. For this reason, most are manipulations of Newton's laws of motion [Brackstone & McDonald, 1999]. Examples of collision avoidance models include the following.

Kometani and Sasaki Model

Kometani and Sasaki were first to use collision avoidance as the driving motivation for car-following behaviour [Kometani & Sasaki, 1959]. Their model is not governed by a system of stimulus and response, but uses Newton's laws of motion to specify a safe, collision-free following distance. The original formation of the model is as follows;

$$\Delta x(t - \tau) = c_1 v_{n-1}^2(t - \tau) + c_2 v_n^2(t) + c_3 v_n(t) + c_4 \quad (8)$$

Where, $\Delta x(t - \tau)$ is the vehicle displacement during a time increment τ , c_1 to c_4 are calibration parameters, and $v_n(t)$ is the velocity of vehicle n at time t .

A floating car survey of two vehicles on a city street was used to obtain data for calibration of this model. 22 replications of the experiment yielded 310s for analysis. Best fit parameters produced an r^2 value of 0.75. When repeating this experiment with two vehicles on a faster test track, the results produced an r^2 of 0.25 and 0.95 for the two vehicles.

This initial attempt at developing a collision avoidance car-following model indicated that it was able to replicate human driving behaviour in select circumstances. Kometani and Sasaki were limited in their access to data and test scenarios. Perhaps with vehicle interaction data from different speed zones and traffic densities, the authors could have increased the scope of training data used to calibrate the model. The resulting car-following model would be more diverse, and experience reduced error in differing circumstances.

Gipps Model

The Gipps car-following model was derived by setting restrictions on the vehicle kinematics, based on driver performance [Gipps, 1981]. The model assumes that the following vehicle will adjust and maintain a headway to avoid collisions in a sudden emergency stop. The model is presented in Equation 9;

$$v(n, t + \tau) = \min(v_a(n, t + \tau), v_b(n, t + \tau)) \quad (9)$$

Where, $v(n, t + \tau)$ is the velocity of the vehicle at time t , τ is the reaction time of the driver, $v_a(n, t + \tau)$ is the velocity of the vehicle during acceleration and $v_b(n, t + \tau)$ is the velocity of the vehicle during deceleration.

The driver's acceleration towards a desired speed is provided in Equation 10;

$$v_a(n, t + \tau) = v(n, t) + 2.5a(n)\tau(1 - \frac{v(n, t)}{v^*(n)})\sqrt{0.025 + \frac{v(n, t)}{v^*(n)}} \quad (10)$$

Where, $v(n, t)$ is the velocity of vehicle n at time t , $v^*(n)$ is the desired speed of the road section and $a(n)$ is the maximum acceleration of the vehicle.

The change in vehicle velocity is bound by the section speed limit and the vehicle acceleration capabilities. Spatial constraints also bound vehicle motion. The deceleration equation is presented in Equation 11;

$$v_b(n, t + \tau) = d(n)\tau - \sqrt{d(n)^2 - d(n)(2x(n-1)\tau - l(n-1)) - x(n, t)\tau - \frac{v(n-1, t)^2}{d'(n-1)}} \quad (11)$$

Where, $d(n)$ is the maximum deceleration tolerated by the vehicle $x(n, t)$ is the position of vehicle n at time t , $l(n-1)$ is the effective length of the preceding vehicle and $d'(n-1)$ is an estimation of the preceding vehicle's deceleration.

During this models formulation, Gipps introduced three parameters to incorporate human decision making. The first is human reaction time. The second is a prediction of the preceding vehicles deceleration rate. The third was incorporated into his derivation and prevented underestimating the intentions of the preceding driver. These assumptions cause the model to exhibit sudden and erratic accelerations based on the movement of a leader. The model does not monitor adjacent lanes, and so the imprudent lane change of adjacent leaders causes a large deceleration.

General Model, MOBIL

The general minimising overall braking induced by lane change (MOBIL) model captures the implications of a lane change on surrounding vehicles [Kesting et al., 2007]. Kesting et al., state that despite other lane-changing models being adept at replicating the decision-making process of a driver intending to change lanes, they lack the ability to incorporate implications for the other vehicles in the utility functions that govern advantageous lane changes. Such a model and style of driving would help bridge the cognitive dissonance that is present between maximising ones own utility in a lane change and the impact this decision has on the wider network.

This model has three major advantages. The first advantage is that the lane-changing model is no longer responsible for calculating traffic conditions. Calculations of the gaps, relative velocities, and accelerations are transferred to the car-following model, allowing the lane-changing model to be simpler

and more concise. The second advantage is that when longitudinal and lane-changing models are consistent with each other, the lane-changing model is also collision-free if the car-following model is collision-free. This assumption again simplifies the formulations for this proposed model. Finally, the braking deceleration forced on another vehicle during the lane change manoeuvre can be used as a surrogate measure of safety. This allows the authors to unify safety and motivation into a single model.

This model differs from other car-following models in two ways. The first way is that its utility function for defining a good lane change is built on acceleration factors alone. The second is that this model incorporates a “politeness factor”. The purpose of the “politeness factor” is to vary the intentions of the lane change, from purely egotistic to more altruistic [Kesting et al., 2007]. The “politeness factor” changes the consideration that a vehicle has for its impacts on surrounding vehicles, to the point where this factor can mandate a lane change not occur unless there is a benefit for the system. The model is presented in the following expression;

$$\frac{\overline{a_n} - a_n}{(driver)} + p \left(\frac{\overline{a_{n+1}} - a_{n+1}}{(new\ follower)} + \frac{\overline{a_{n-1}} - a_{n-1}}{(old\ follower)} \right) > \Delta a_{th} \quad (12)$$

Where, the first two terms represent the utility gained from the possible lane change for vehicle n , and p is the “politeness factor”.

Models of this nature raise questions regarding the cooperativeness of vehicles in real-world traffic streams. Much of traffic, mobility, and transport is built on the economic principle of maximising one’s own utility. Drivers often do not have the capability nor the willingness to assess the impacts of their actions on the wider network in real-time.

Fadhloun and Rakha Model

Fadhloun and Rakha developed a model that emphasises human-in-the-loop [Fadhloun & Rakha, 2019]. The model captures driver perception and controls the inaccuracies and errors of other car-following models. The model consists of three major components, one for steady-state, one for collision avoidance, and one for vehicle dynamics. The collision avoidance component is responsible for

ensuring safe vehicle operation by intervening when a collision would otherwise be unavoidable. The collision avoidance model is an unaltered form of the kinematic equation of motion and is presented in Equation 13;

$$v_{n+1}^{CA} = \sqrt{v_n^2 + 2d_{max}(\Delta x - \Delta x_\rho)} \quad (13)$$

Where, v_{n+1}^{CA} is the velocity calculated using the collision avoidance model, v_n is the speed of the following vehicle, d_{max} is the maximum deceleration, Δx is the headway spacing between the ego and lead vehicle, and Δx_ρ is the jam density spacing.

The steady-state speed is calculated using the nonlinear functional form of the Van Aerde model [Van Aerde, 1995]. It is presented in Equation 14;

$$v_{n+1}^{VA} = \frac{-c_1 + c_3v_f + s_{n+1} - \sqrt{(c_1 - c_3u_f - \Delta x)^2 - 4c_3((\Delta x - c_1)u_f - c_2)}}{2c_3} \quad (14)$$

Where, v_{n+1}^{VA} is the velocity calculated using the steady-state model, c_1 to c_3 are modelling calibration parameters, and v_f is the free flow speed.

Finally, the vehicle dynamics model is used to ensure that its physical and mechanical capabilities bound the behaviour of the vehicle. Using Newton's second law of motion, the acceleration of the vehicle is the difference between the tractive forces and resistive forces (aerodynamics, rolling resistance, and grade resistance) acting on the vehicle, divided by its mass. The dynamics model is presented in Equation 15;

$$a_{n+1} = \frac{\min(3600, f(\frac{v_{n+1}}{v_{max}}) \cdot \frac{\eta c_4 P}{v_{n+1}}, mg\mu) - R_{n+1}}{m} \quad (15)$$

Where, a_{n+1} is the acceleration provided by the vehicle dynamics models, $f(\frac{v_{n+1}}{v_{max}})$ is the throttle function which uses initial velocity (v_{n+1}) and final velocity (v_{max}), η is the driveline efficiency, c_4 is a parameter accounting for gear-shift impacts at slow speeds, P is the vehicle power, m is the vehicle mass, g is gravitational acceleration, and μ is the coefficient of friction.

This model improved on previous iterations of collision avoidance models by modelling the driver throttle and brake pedal input, smoothening the acceleration profile of vehicles to be consistent with empirical data, explicitly capturing driver perception, and controlling inaccuracies and errors.

3.1.3 Psycho-Physical Models

Psycho-physical models are an extension of GHR models, where “action points” and thresholds are added to discretise car-following behaviour [Brackstone & McDonald, 1999]. In a GHR model, objects far from a vehicle will continue to have a small, but non-negligible, impact on its behaviour. Psycho-physical models aim to quantify regions between which the car-following behaviour changes. The reasoning is that driver behaviour changes as the nature of their interaction with other agents in the network changes. This interaction can be characterised by speed differences, headways, proximity to turns, or other external stimuli. This subsection outlines a range of novel and significant psycho-physical models.

Evans and Rothery Model

Evans and Rothery are credited with creating the first psycho-physical car-following model in 1973 [Evans & Rothery, 1973]. This style of modelling addresses the GHR models limitations, which did not consider a drivers ability to quantitatively assess headway, changes in headway, and relative speed, all of which are considered important in defining driver car-following behaviour.

Evans and Rothery conducted field experiments in which examiners asked participants to take the role of passengers in a test setting. The examiners then asked participants whether the lead vehicle was travelling faster than, at the same speed, or slower than the ego vehicle. This experiment was repeated at varying headways between $80ft$ and $550ft$. The experiment showed that when relative velocity divided by spacing exceeded $\pm 0.03s^{-1}$, it could be detected 75% of the time. Some experiments showed that subjects might be receptive to the change in pitch of the vehicle as it accelerates. Further investigations showed that this mechanism was effective in predicting acceleration for vehicles with a non-rigid suspension.

Wiedemann Model

The Wiedemann car-following model uses decision points and regions to define vehicle behaviour [Wiedemann, 1974]. The regions are presented diagrammatically in Figure 3 [PTV Group, 2016];




Figure 3: Thresholds for the regions within the Wiedemann psycho-physical car-following model [PTV Group, 2016].

The thresholds are defined as follows;

- **AX:** Desired spacing between stationary vehicles. This distance is the sum of the lead vehicles length and the bumper-to-bumper distance between vehicles.
- **BX:** Minimum following distance when the speed difference between the lead and ego vehicle is low.
- **SDX:** Maximum following distance until which the actions of the lead vehicle will effect those of the ego vehicle.
- **SDV:** Critical threshold where the ego vehicle notices it is approaching a slower travelling lead vehicle.
- **OPDV:** Point where a driver determines that their speed is lower than that of the lead vehicle.

The “No Reaction” state occurs when the headway to the lead vehicle is too large. The “Collision” region imposes the maximum deceleration when the headway falls below AX, to prevent an accident. The “Deceleration” region is designed to match the speed of the ego vehicle with the lead vehicle as the car-following headway approaches BX. The “Reaction” region defines behaviour when a vehicle approaches a noticeably slower region and is forced to decelerate.

This model is limited by its tendency to segment driving behaviour based on the prevalence of information to the driver. V2V and V2I will greatly extend not only the information available to the vehicle but also the quantity, quality, and accuracy of that information. For this reason, adjusting the parameters, mathematical models, and regions of psycho-physical models such as the Wiedemann model may not allow for the emulation of CAV behaviour in a simple and adjustable framework. A psycho-physical model reverts to a GHR model when the prevalence of information available to CAVs pushes the region thresholds to infinite.

Fritzsche Model

The Fritzsche car-following model is used in the Paramics microsimulator, and shares similarities with the Wiedemann model [Fritzsche, 1994]. Fritzsche claimed that modelling the propagation of waves in traffic streams using thermodynamics and fluid mechanics is only reasonable for moderately dense traffic (<30 cars/lane/km). Traditionally, thermodynamics and fluid mechanics models along with the fundamental diagram were used to relate velocity, density, and flow, and extrapolated to calculate capacity. Fritzsche claimed that this process introduced errors as the correlation between the models and data could not be adequately established.

Fritzsche proposed a psycho-physical car-following model that segregated the perception of speed into two regions, the perception of headway into four regions, and the perception of gaps into two regions. Perception of positive and negative speed difference is defined as follows;

$$\Delta v_{PTN/PTP} = \pm c_{PTN/PTP}(\Delta x)^2 \mp c_2 \quad (16)$$

Where, $\Delta v_{PTN/PTP}$ is the perception threshold distance for positive velocity changes (*PTP*) and negative velocity changes (*PTN*), c and c_2 are modelling parameters and Δx is the headway to the preceding vehicle.

The thresholds describing the ego vehicles reaction to a car-following headway is described below as one of four states, and diagrammatically presented in Figure 4;




Figure 4: The action regions of the psycho-physical Fritzsche car-following model [Fritzsche, 1994].

- Desired Distance (AD): Distance that the ego vehicle maintains.
- Risky Distance (AR): Limiting distance headway before the ego vehicle decelerates heavily to avoid a collision.
- Safe Distance (AS): Smallest required headway before the ego vehicle will positively accelerate, limiting speed permitted.
- Braking Distance (AB): Maximum distance required to stop during maximum deceleration.

The “Danger” region indicates that the headway is smaller than AR and so maximum deceleration is necessary. The “Closing in” region indicates that the lead vehicle is within perception range of a negative speed change (PTN), but the headway is greater than AR. This situation forces the ego vehicle to

decelerate and match the velocity of the lead vehicle when the headway reaches AR. “Following 1” and “Following 2” define regions of inaction or statistically induced variance in acceleration. Finally, “Free Driving” indicates a region where the lead vehicle’s velocity change is imperceivable or positive (PTP) and the headway is sufficiently large (AD), meaning that the ego vehicle can accelerate freely if travel speed limits permit.

Winsum Model

Winsum stated that car-following models lacked a human component, as they attempted to emulate human behaviour solely from an engineering perspective [Van Winsum, 1999]. He proposed that drivers apply a heuristic as they attempt to cope with limitations in reflexive and adaptive capabilities. This heuristic allows a driver to interpret vast amounts of data to make real-time decisions affecting vehicle operation. Using psychological assumptions and measurements, Winsum proposed the following expression [Van Winsum, 1999];

$$a_n = 1.04 \cdot c_1 \cdot \left(\frac{\frac{1.8}{\tan(1.1 \times a_{n-1} \times \tan(\frac{1.8}{\Delta t_{n-1} V_n}))}}{a_{n-1} \cdot \sqrt{2 \times \frac{\Delta t_{n-1} v_n - \frac{1.8}{\tan(1.1 \times a_{n-1} \times \tan(\frac{1.8}{\Delta t_{n-1} V_n}))}}{-a_{n-1}}}} \right)^{0.72} + c_2 + \epsilon \quad (17)$$

Where, a_n is the acceleration of the current vehicle, c_1 and c_2 are calibration constants, Δt_{n-1} is the time headway to the preceding vehicle, and ϵ is a random error term representative of uncaptured influences.

Though Winsum did not test this model in a microsimulation environment or verify its accuracy; his concluding remarks stated that this model was still too simple to capture all of the factors affecting a driver’s car-following behaviour.

3.1.4 Fuzzy Logic Models

Fuzzy logic models convert strict mathematical models and definitions into qualitative rule-based reasoning. For example, rather than defining a minimum headway, fuzzy logic applies categorical classifications such as “critical”, “close”, “comfortable”, and “far”. Each category has an associated probability distribution, with the distributions overlapping adjacent categories. Car-following would then become a matter of IF “close” AND “getting closer”

THEN “brake”. Fuzzy logic allows a system to better represent the disaggregate and stochastic behaviour of its constituents.

Kikuchi and Chakroborty Model

Kikuchi and Chakroborty were the first to propose the use of fuzzy logic to model car-following [Kikuchi & Chakroborty, 1992]. Where the other models used deterministic rules to dictate the reaction of the following vehicle in response to its lead vehicle, Kikuchi and Chakroborty claimed that car-following was an aggregation of vague rules that differed by circumstance. Fuzzy logic modelling uses language and descriptive rules with a probability distribution to model the imprecise human reasoning process.

The value of fuzzy logic inference is realised when the premise is compatible with the rule governing it. Only then is accuracy in a predicted consequence achieved. Consider the example provided by the authors;

$$\begin{array}{lll}
 \text{Input :} & x \text{ is somewhat } A & (x = A') \\
 \text{Rule :} & \text{if } x \text{ is } A \text{ then } y \text{ is } B & (R : x = A - > y = B) \\
 \text{Conclusion :} & y \text{ is somewhat } B & (y = B')
 \end{array} \quad (18)$$

The appropriateness of the conclusion is contingent on the accuracy of the input and the validity of the rule. The input for the Kikuchi and Chakroborty model is;

- Distance between the lead and following vehicle
- Relative speed between the lead and following vehicle
- Acceleration (or deceleration) rate of the lead vehicle

Each rule of the fuzzy logic model discretises the input into six language-based categories, and each category has an associated membership function. By conditionally following the rules, a vehicle determines its acceleration. For example, if the following distance is adequate, and the relative speed is near zero, and the lead vehicle deceleration is mild, then the following vehicle should decelerate mildly. This model applies to both the lead vehicle and the following vehicle, and is mathematically represented in Equation 19 as;

$$a_n = \left(\frac{\Delta v + a_{n+1} \cdot \tau}{\Delta t} \right) + \beta_{\Delta x} \cdot \phi \quad (19)$$

Where, a_n is the acceleration of the following vehicle, Δv is the perceived relative speed, a_{n+1} is the perceived acceleration of the lead vehicle, τ is the time interval over which the model is run, Δt is the time over which the following vehicle will match the lead vehicle (set to 2.5s in this study), $\beta_{\Delta x}$ is the number of categories from which the perceived distance (Δx) deviated from adequate, and ϕ is the lateral translation made to the membership function for this deviation (set to 1 in this study).

The level of compatibility between the input and rules is determined using fuzzy inference. Though this study did not verify the membership functions through collected field data, it was able to show that using a fuzzy logic approach yields comparable results to using a deterministic modelling approach. Fuzzy logic models better explain deviations and dispersions in data as probabilistic stochasticity, something that deterministic models are unable to do.

McDonald et al., Model

The fuzzy logic model proposed by McDonald et al., shares similarities with the Kikuchi and Chakroborty model [McDonald et al., 1997]. The model uses two inputs to the decision-making process, the relative speed between the lead and follow vehicles (Δv) and the ratio of the separation between the two vehicles to the desired separation ($\Delta x / \Delta x_{max}$). The Δv and $\Delta x / \Delta x_{max}$ both have five language-based categories, with Δv ranging from “opening fast” to “closing fast”, and $\Delta x / \Delta x_{max}$ ranging from “much too far” to “much too close”. This segregation results in five driver response language-based categories ranging from “strong acceleration” to “strong deceleration”.

Calibration data is needed to adjust the membership functions and fine-tune the operation of the model. Data was collected by asking six drivers to drive along a motorway and give there subjective opinion of following distances. The membership functions of the language-based categories were defined using the collected data and a triangular distribution. The car-following model is defined in the following equation;

$$\begin{aligned}
c_1 + c_2 &= 1 \\
AS_{(i)}c_1 + AS_{(j)}c_2 &= AS_{(ij)}
\end{aligned} \tag{20}$$

Where, $AS_{(i)}$ is the average deceleration rate corresponding to Δv and $\Delta x/\Delta x_{max}$ in set i , $AS_{(j)}$ is the average deceleration rate corresponding to Δv and $\Delta x/\Delta x_{max}$ in set j , and $AS_{(ij)}$ is the resulting acceleration rate when Δv and $\Delta x/\Delta x_{max}$ belong to different sets i and j . c_1 and c_2 are both parameters used to determine the percentage that the base rates will fire when Δv and $\Delta x/\Delta x_{max}$ are in different base sets.

Using this approach, the calibrated model appropriately replicated field car-following behaviour. Their work found that $\Delta x/\Delta x_{max}$ (deviation between the desired and actual following distance) was most important to car-following behaviour.

Khodayari et al., Model

Khodayari et al., developed a fuzzy logic car-following model by using a Gaussian membership function for each fuzzy set, and three membership functions per input [Khodayari et al., 2011]. The inputs were relative distance, relative speed, velocity of the following vehicle, and instantaneous reaction delay.

While the distribution of these membership functions can be obtained from observed data as other studies have done, instantaneous reaction delay is more difficult to quantify. Khodayari et al., achieved this by observing the relative velocity and acceleration of lead and following vehicle pairs. As the lead vehicle accelerates or decelerates, the relative velocity changes. As a result, the following vehicles adjust their acceleration. The instantaneous reaction time is the time difference between the relative velocity changes and when the following vehicle's acceleration changes accordingly.

By including instantaneous reaction time as part of the inputs in their fuzzy logic model, Khodayari et al., experienced a 33% reduction in the root mean square error (RMSE) when compared to not considering this variable input. The model demonstrated a mean absolute percentage error of 23.92%.

3.1.5 Recurrent Neural Networks

Neural network models consist of the input layer, a number of hidden layers, and an output layer. Reinforcement learning is used to create and strengthen probabilistic pathways between the input and output through the hidden layers. This style of modelling is gaining popularity due to its parameter-free and data-driven nature. Neural network models are generally deficient at modelling traffic oscillation, a problem with human vehicle fleets that is expected to be less pertinent with CAV fleets. Zhou et al., developed a recurrent neural network (RNN) that addressed the oscillation limitation of other neural network based car-following models [Zhou et al., 2017].

RNNs are used in speech recognition, language translation, stock prediction, image recognition, and various other applications. RNNs operate by not being a memoryless process, they retain data from previous iterations and time increments to help inform future decision. This characteristic makes them well-suited for the applications above, where the nature of any particular word, sentence fragment, data point, or pixel is contingent on its surrounding, and the strength of influence diminishes as the distance between points increases.

Zhou et al., implement an RNN by sequencing weighted input data through concurrent hidden layers. The length of the data is unrestricted and variable. The RNN takes an input of spacing, relative speed, and vehicle speed for the current time step, and outputs acceleration for the next time step. Based on outputted acceleration, the vehicle position and velocity for the next time step are updated based on the kinematic equations of motion. When comparing the RNN modelling framework with other models, the RNN model demonstrates a 41.7% reduction in RMSE.

3.2 Lane-Changing Algorithms

The lane-changing model is a crucial aspect of microsimulation modelling, as it reconciles a drivers' short term and long term goals [Gipps, 1986]. Lane-changing algorithms define a set of conditions under which a vehicle desires to change lanes. The lane change is classified as one of two types, a Mandatory Lane Change (MLC) or a Discretionary Lane Change (DLC). An MLC is incited by an unavoidable need to change lanes such as obstructive objects in

the drivers' current lane or an upcoming turning movement. A DLC is one that may not be strictly necessary but will provide short term advantages to the driver. This short term advantage may result from overtaking a slower vehicle or passing a heavy vehicle.

Many lane-changing models implement decisions that are generally based on the outcomes of a utility function, which evaluate travelling in adjacent lanes. This utilitarian means of evaluating lane-changing does not allow for cooperation between vehicles, as the implications of a vehicle's decision on surrounding vehicles and the traffic stream are not evaluated or considered [Kesting et al., 2007]. Nor do many models allow a DLC when the vehicle is within a zone strictly reserved for MLCs [Toledo et al., 2003].

Lane-changing models are primarily classified as one of four types, rule-based, discrete-choice, incentive-based, and artificial intelligence models. The remainder of this subsection provides a brief outline of each type and the novel models proposed throughout literature.

3.2.1 Rule Based Models

Rule-based models follow rigid frameworks that outline the lane change triggering thresholds. Rules are often static and comprehensive, dictating the specific conditions under which a lane change occurs. This subsection outlines in detail some of the unique contributions to rule-based lane-changing.

Gipps Lane-Changing Model

The 1986 Gipps lane-changing model evaluates the decision to change lanes by answering the following three questions; "Is it possible to change lanes? Is it necessary to change lanes? Is it desirable to change lanes?" [Gipps, 1986]. Gipps considered the following factors as being the most significant contributors to a drivers decision to change lanes;

- "Whether it is physically possible and safe to change lanes" – Drivers will not change lanes if the perceived risk of changing lanes is greater than what the driver is willing to accept. The perceived risk is a function of the driver's position relative to the gap in the adjacent lane and the

velocity of both the preceding and following vehicles in the adjacent lane. A driver's willingness to accept greater risks will increase as the need to change lanes increases, caused by an approaching obstacle or turning movement.

- “The location of permanent obstructions” – Drivers accustomed to the road network are familiar with upcoming obstacles such as regularly parked cars. This knowledge may incentivise a driver to change lanes and ignore the temporary advantage that is offered by the lane with the obstacle. As mentioned above, this desire to change lanes will increase as proximity to the obstacle decreases.
- “The presence of transit lanes” – Transit lanes are specially reserved for use by public transport vehicles and other high occupancy transit vehicles. Generally, a transit vehicle will only exit the transit lane if it is obstructed. Conversely, other vehicles will only enter the transit lane if forced to by an upcoming turning movement.
- “The driver's intended turning movement” – The distance of the driver from its turning movement directly impacts a driver's willingness to change lanes. Beyond a particular distance, the proximity to the turning movement will not impact the driver's decision. This indifference reduces as the turning movement approaches.
- “The presence of heavy vehicles” – drivers tend to avoid following heavy vehicles, primarily due to their slow acceleration rate. This behaviour is most prominent when the heavy vehicle is accelerating from rest, and almost non-existent when the heavy vehicles reach cruising velocity.
- “Speed” – The primary reason for changing lanes is to increase speed. If a driver perceives the adjacent lane as having greater speed and believes that travelling in this lane will result in a short term travel time or speed advantage, then changing lanes will appear as more appealing.

Though the Gipps lane-changing model is used in the Aimsun commercial microsimulator, it is still limited in its applicability to CAV operation. The model is rigid, meaning that despite the model having the potential for expansion through the definition of new rules, the new rules must also be rigidly designed. The rules currently contained within the model focus on the ego vehicle and provides no scope for cooperation with surrounding vehicles. The

final limitation is that the domain of the rules must be expansive, covering every situation a vehicle could experience in its movement in a microsimulation environment. This domain grows significantly when cooperation with more vehicles is considered in a well-connected environment.

Kita Model

Kita proposed modelling rule-based lane-change through the use of game theory [Kita, 1999]. Kita argued that the behaviour of the lane-changing vehicle and giving-way vehicle is not independent, and so should not be modelled as such. For this reason, Kita modelled the interaction between the two vehicles as a system following a two-person non-zero-sum non-cooperative game. Such a game particularly describes the situation where a vehicle on the mainline will actively change lanes in response to a merging vehicle.

The parameters of the game are as follows; the players consist of the two conflicting vehicles (even though lane changing actions have wider network impacts on vehicles in the near vicinity), a single game is played for each interaction, all games in the system are independent of one another, the drivers are aware of the pay-off matrix of other vehicles but do not communicate decisions with one another (i.e. a non-cooperative game with perfect information).

If the lane-changing vehicle is defined as P_1 and the through-vehicle in the mainline is defined as P_2 , then the strategy of P_1 (p) is given as $p = \{1 : \text{merge}, 2 : \text{pass}\}$ and similarly the strategy of P_2 (q) is given as $q = \{I : \text{giveway}, II : \text{do not giveway}\}$. The pay-off matrices are then;

$$[P_1] = \begin{matrix} & \begin{matrix} I & II \end{matrix} \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \end{matrix} \quad [P_2] = \begin{matrix} & \begin{matrix} I & II \end{matrix} \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \end{matrix} \quad (21)$$

Where;

- $a_{11} = c_1 t_3 + c_2$ where, t_3 is the time to collision (TTC) between the merging vehicle and the follower of the following vehicle in the target lane.

- $a_{12} = c_1 t_2 + c_2$ where, t_2 is the TTC between the merging vehicle and the following vehicle in the target lane.
- $a_{21} = a_{22} = c_3 t_1$ where, t_1 is the TTC between the merging vehicle and the end of the merging lane.
- $b_{11} = b_{21} = c_4 t_5 + c_5$ where, t_5 is the TTC between the following vehicle in the target lane and the merging vehicle's own following vehicle in its adjacent lane.
- $b_{12} = c_6 t_2$
- $b_{22} = c_7 t_4$ where, t_4 is the time headway between the leader and follower of the lane-changing vehicle in the target lane.
- $c_1, c_2, c_3, c_4, c_5, c_6, c_7$ are explanatory variables calibrated using a method of method likelihood estimation.

The model was tested on a case study in the Keiyo Expressway in Chiba Prefecture, Japan. Using 25 conflicts as training data for model calibration, the model had a correlation coefficient of 0.7 between the observed data and the model predictions. This level of correlation was deemed high enough by Kita to conclude the applicability of a game-theory approach to modelling lane-changing behaviour.

Cellular Automata Model

The Cellular Automata model, extended by Nagel et al., is a microscopic car-following model that also explicitly defines the conditions under which a lane change would occur [Nagel et al., 1998]. The base conditions are;

- $gap_n(t) < \min\{v_n(t) + 1, v_{max}\}$ where, $gap_n(t)$ is the number of empty cells ahead in the same lane, $v_n(t)$ is the speed of vehicle n at time t , and v_{max} is the maximum allowed velocity of the vehicle.
- $gap_{n,o}(t) > \min\{v_n(t) + 1, v_{max}\}$ where, $gap_{n,o}(t)$ is the number of empty cells in the target lane.
- $gap_{n,ob}(t) > v_{max}$ where, $gap_{n,ob}(t)$ is the number of empty cells behind the lane-changing vehicle in the target lane.

The first two conditions evaluate the adjacent lanes in search of favourable conditions. The third condition evaluates the target lane for sufficient space

for the lane change. The movement of the vehicle is defined as; *if* ($v < v_{max}$) *then* $v := v + 1$ (accelerate if possible); *if* ($v > gap$) *then* $v := gap$ (slow down if necessary); *if* ($v \geq 1$) *then with probability* p , *do* $v := v - 1$ (add variability and natural human stochasticism to the creation of a gap).

Being a macrosimulation model, the Cellular Automata model is not burdened with having to define the caveats of microsimulation lane-changing. A microsimulation rule-based lane-changing model explicitly defines how a vehicle reacts in every interaction type between agents and network objects. A macrosimulator, however, does not model specific agent interactions, as it models using aggregate link or node flows. For this reason, macrosimulation models lack the detail of their microsimulator counterparts. While a macrosimulation model for lane-changing can provide insights into the fundamental mathematical thresholds for evaluating a lane-change, the model is not directly transferable and nor is it appropriate for defining cooperative vehicle interactions between CAVs.

3.2.2 Discrete Choice Based Models

Discrete-choice models determine the utility of an action (or discrete choice) by aggregating quantitative variables and quantising qualitative variables. Parameters that affect a decision are aggregated to form a utility function that quantitatively provides the cost of that decision. The specificities of competing decisions are entered into the utility function in-turn, calculating an overall utility for each option. The utility of each competing decision is then placed into a utility model, such as the logit model or probit model. The utility model provides the probability of selecting each available option. This subsection outlines novel and unique discrete-choice lane-changing models.

Target Lane Model

The target lane model assesses the utility of all available lanes and moves the driver to the lane with the highest utility [Toledo et al., 2003]. This multi-lane assessment provides the driver with the greatest selection of target lanes, particularly on multi-lane highways. The utility of each lane is evaluated using the following equation;

$$U_{int}^{TL} = \beta_i^{TL} X_{int}^{TL} + r_n \alpha_{int}^{TL} + \epsilon_{int}^{TL} \quad \forall i \in (lane_1, \dots, lane_n) \quad (22)$$

Where, U_{int}^{TL} is the utility of lane i as a target lane TL for driver n at time t , β_i^{TL} is the vector of parameters that affect the utility of lane i , X_{int}^{TL} is the corresponding vector of explanatory variables, r_n is the vector of parameters for the driver-specific random term α_{int}^{TL} , and ϵ_{int}^{TL} is the random term associated with the target lane utilities.

Target lane model utilities are governed by specific lane attributes such as density, speed of traffic, presence of heavy vehicles, number of lanes required to move from the current lane to the target lane, and the appropriateness and distance of the target lane to the next turning movement. The error terms of the utility function are assumed to be independent, and identically Gumbel distributed. This assumption allows use of the multinomial logit model to estimate the probability of selecting each available lane [Toledo et al., 2003]. The multinomial logit model is as follows;

$$P(TL_{int} = i * v_n) = \frac{e^{(U_{int}^{TL} i * v_n)}}{\sum_{j \in TL} e^{(U_{int}^{TL} j * v_n)}} \quad (23)$$

Where, $P(TL_{int} = i * v_n)$ is the probability of selecting lane i as the target lane, $U_{int}^{TL} i * v_n$ is the utility derived by selecting lane i , and j is the set of all possible target lanes.

To estimate the coefficients of explanatory variables in the utility function, vehicle trajectory data was collected in a section of the I-395 southbound highway in Arlington, Virginia. Using the collected dataset and the method of maximum likelihood, the calculated utility function is as follows;

$$\begin{aligned} U_{int}^{TL} = & \beta_i - 0.011\rho_{int} + 0.119v_{int} + 0.022\Delta x_{int}^{front} \theta_{int}^{adj} \\ & + 0.115\Delta v_{int}^{front} \theta_{int}^{adj} - 2.783\theta_{int}^{tailgate} \theta_{int}^{CL} + 1.000\theta_{int}^{CL} \\ & - 2.633\Delta CL_{int} + \beta_i^{path} [x_{nt}^{exit}]^{0.371} - 0.980\theta_{nt}^{nextexit} \Delta Exit_i \\ & - \alpha_1 v_n + \epsilon_{int}^{TL} \end{aligned} \quad (24)$$

Where, β_i is the constant for lane i , ρ_{int} and v_{int} is the lane specific density and speed respectively, Δx_{int}^{front} and Δv_{int}^{front} is the spacing and relative speed of the preceding vehicle in lane i respectively, θ_{int}^{adj} equals 1 if i is the current or

directly adjacent lane and 0 otherwise, $\theta_{int}^{tailgate}$ equals 1 if the vehicle n is being tail-gated and 0 otherwise, ΔCL_{int} is the number of lane changes required to move from the current lane to lane i , β_i^{path} is the path plan impact coefficient for lane i , x_{nt}^{exit} is the distance to the exit for driver n , $\theta_{nt}^{nextexit}$ equals 1 if the driver intends to take the next exit and 0 otherwise, and $\Delta Exit_i$ is the number of lane changes required to move from lane i to the exit lane.

To validate this model, Toledo et al., implemented it in microscopic traffic simulation and tested it against data obtained for a specific section of the I-80 motorway in Berkeley, California. Comparison between the proposed target lane model and the default shift directional model showed a 29.79% improvement in the RMSE of the speed, and a 13.41% improvement in the ability to emulate real driver lane selection processes.

The target lane model shows improvement in environments where drivers can select from numerous potential target lanes. The model retains a heavy influence of human characteristic and does not conduct an assessment on the impacts of the vehicle decision on other vehicles in its immediate surroundings. While this additional assessment is not necessary to emulate selfish human behaviour, it may be critical in emulating cooperative CAV behaviour and deriving maximum utility from automated mechanical systems.

Ahmed Model

Ahmed proposed a lane-changing model designed to reflect the heterogeneity in motorist driving behaviours[Ahmed, 1999]. Ahmed considered the explanatory variables that affect driver behaviour, and segregated the lane-change manoeuvre into three categories; MLC, DLC and forced merging. An MLC is dictated by an upcoming turning movement, forcing a vehicle into the correct lane. A DLC is the result of a vehicle unsatisfied with its current situation. A forced merge occurs in heavy conditions when the target lane gap is not sufficient but is actively increased by the aggressive behaviour of the lane-changing vehicle. The probability that a driver performs a lane change is given by Equation 25;

$$P_t(LC|v_n) = \frac{1}{1 + e^{-X_n^{LC}(t)\beta^{LC} - c^{LC}\alpha_n}} \quad (25)$$

Where, LC is the lane change type (mandatory, discretionary, or forced lane change), $P_t(LC|v_n)$ is the probability of executing a lane change for vehicle n and time t , X_n^{LC} is the vector of explanatory variables that affect the decision to change lanes, β^{LC} is the corresponding vector of parameters, α_n is the driver-specific random term, and c^{LC} is the calibration parameter for α_n .

To trigger a lane change, minimum gap acceptance criteria must also be met. The critical minimum gap is defined as;

$$gap_n^{cr,n\pm 1} = e^{X_n^{cr,n\pm 1}(t)\beta^{n\pm 1} + c^{n\pm 1}\alpha_n + \epsilon_n^{n\pm 1}(t)} \quad (26)$$

Where, $gap_n^{cr,n\pm 1}$ is the critical gap to the leader or follower in the target lane for vehicle n at time t , $X_n^{cr,i}$ is the vector of explanatory variables that affect gap $n \pm 1$, $\beta^{n\pm 1}$ is the corresponding vector of parameters, α_n is the driver-specific random term, $c^{n\pm 1}$ is the calibration parameter for α_n , and $\epsilon_n^{n\pm 1}$ is a normally distributed random term ($N = \sigma_{n\pm 1}^2$).

The probability of executing the lane-change was given in Equation 25, the probability of accepting a lane-change is as follows;

$$\begin{aligned} & P_n(gap \text{ acceptance} \mid \alpha_n) \\ &= P_n(lead \text{ gap acceptable} \mid \alpha_n) P_n(lag \text{ gap acceptable} \mid \alpha_n) \\ &= P_n(gap_n^{lead}(t) > gap_n^{cr,lead}(t) \mid \alpha_n) P_n(gap_n^{lag}(t) > gap_n^{cr,lag}(t) \mid \alpha_n) \end{aligned} \quad (27)$$

Where, $gap_n^{lead}(t)$ is the probable lead gaps in the target lane, and $gap_n^{lag}(t)$ is the probable lag gaps in the target lane.

Ahmed verified the model by implementing it in the MITSIM macroscopic traffic simulator. The RMSE percentage error decreased from 5.81% to 1.56%, and the error in traffic counts decreased from 9.08% to 7.53%. When implemented in MITSIM (along with a proposed acceleration car-following model), the proposed lane-changing model better reflected the observed data in all categories, with the Theil's inequality (a measure of fitness, with lower coefficients implying a better fit) reducing from 0.050 to 0.039.

During free flow conditions ($\rho \leq 19veh/km/lane$), the model parameters showed a significant t-statistic. However, the model parameters were deter-

mined by calibrating to a case study that contained curvature in the roadway upstream of the data collection point, the presence of a weaving segment, and two exits downstream of the data collection point. These geometric features may have effected driving behaviour, and hence the calibration of the model.

Sun and Elefteriadou Model

Sun and Elefteriadou designed a stated preference survey to evaluate a driver's decision when faced with a range of different scenarios [Sun & Elefteriadou, 2011]. The survey investigated the impact of situations including upcoming turns, terminating lanes, stopped buses, another vehicle merging into the driver's lane, slower-moving vehicles ahead, shorter queue lengths in adjacent lanes, the influence of trucks and heavy vehicles, being tailgated by another vehicle, and pavement conditions, on a driver's decision to change lanes.

Their study found that the factors which affect lane-changing differed by the current situation of the driver. For example, a driver that is currently waiting behind a bus is most influenced by the level of traffic congestion and queueing ahead, location of the next bus stop, distance to the bus, and the number of people at the bus stop. A driver that observes a queue advantage in the adjacent lane is most influenced by the queue length difference, the distance to the next turn, and the congestion on the target lane. A vehicle experiencing tailgating is most influenced by the current speed limit, the level of congestion on all lanes, and the relative position of the driver to the entire corridor. While each scenario has its corresponding utility function, the utility function for the stopped bus scenario is provided in Equation 28 for illustrative purposes;

$$V(LC) = \beta_0 + \beta_1 Cgst + \beta_2 Queue + \beta_3 LocStop + \beta_4 Dist + \beta_5 NPson + \alpha_1 DriverTypeA + \alpha_2 DriverTypeB + \alpha_3 DriverTypeC \quad (28)$$

Where, *DriverType* are all dummy variables, *Cgst* is the level of congestion, *Queue* is the queue ahead, *LocStop* is the distance to the next stop, *Dist* is the distance to the bus, and *NPson* is the number of boarding passengers.

Sun and Elefteriadou also developed a lane-changing model through data collated from revealed preference surveys into the reaction of drivers under various conditions [Sun & Elefteriadou, 2012]. The authors state that previous

models which divide lane-changing into MLC or DLC may not be sufficient for modelling lane-changing in urban environments, congested situations, and arterial roads. The authors claim that drivers are exposed to situations in congested urban environments that greatly differ from those on a motorway network, and so the mitigating factors triggering a lane change also differ.

Their study used an instrumented vehicle to collect information on a diverse group of 40 drivers, varying in age, gender, driving experience, occupation, vehicle ownership, and perceived aggressiveness. Drivers navigated the route while the system recorded details regarding potential lane changes, attempted but not successful lane changes, and completed lane changes. An appropriate classification scheme was then applied to the data. The first classification scheme clustered data based on driver aggressiveness using the K-means algorithm, which attempts to naturally cluster data points. The second classification scheme was based on driver behaviour. An aggressiveness index for each driver was calculated by observing the difference between the number of attempted and completed lane changes. The results were again clustered using the K-means algorithm.

The results of the analysis found that segregating (and aggregating) the driving behaviour into four key groups based on aggressiveness was sufficient in better predicting the lane-changing intentions of drivers. Microsimulators such as CORSIM and Aimsun currently use up to 10 different categories, without providing a justification for this selection, nor drawing parallels to real-world behaviour. Sun and Elefteriadou suggest that conducting a stated preference survey on how aggressive a driver is and segregating results into four categories, would improve the modelling results as compared to current fleet segregations used in commercial microsimulators.

3.2.3 Artificial Intelligence Models

Artificial intelligence lane-changing models introduce randomness into the lane-changing process and provide weighting to the subjective considerations of different variables affecting motorists. They are often classified into five key categories. Hidden Markov models infer unknown information from known observable information by constructing probabilistic models that relate the two. Neural network models use training data and reinforced learning to create a

network of probabilistic pathways between the input and output. Regression models use simple mathematical models to fit relationships between inputs and outputs. Cognitive models help approximate the cognitive process of human drivers during decision-making events. Finally, fuzzy logic systems are built on probabilistic reasoning that rely on qualitative descriptors [Tang et al., 2018].

Ding et al., Model

Ding et al., developed a back-propagation neural network to predict the lane-changing behaviour of motorists [Ding et al., 2013]. Neural networks are so-called as they attempt to mimic the human neural system. A large interconnected system of parallel, non-linear, adaptive nodes connects through pathways. As data moves between the layers from node to node, certain pathways are strengthened, leading to a natural probabilistic set of outcomes for a specific input. This style of reinforced learning mimics the way humans learn and apply experience to new problems.

Neural networks have demonstrated great predictive and problem-solving capabilities through their pattern identification, signal processing, self-optimising, and self-organising tendencies. These characteristics also make neural networks ‘black boxes’, where self-correcting weighting factors for nodes and pathways transforms input data, making the process intractable and unpredictable.

The process of calculating the weighting factor for paths between nodes is referred to as the learning law (or learning algorithm) [Ding et al., 2013]. The learning law relies on the principle of reward and punishment, where weighting factors are strengthened when the neural network reacts appropriately to an input (reward). Conversely, the weighting factors are reduced when the output is an undesirable reflection of the input (punishment). The mathematical expression for a back-propagation neural network system is defined as;

$$Y_j = f\left(\sum_{i=1}^n U_{i,n} X_i - c_n\right) = f(N_n) \quad (29)$$

Where, X is the input vector of variables, U_n is the path connection cost for neuron n , c_j is the threshold to obtain an output, N is the input of neurons in the system, and f is the transfer function.

The training data in this study consisted of 40 sets of lane-changing samples, each with four variables (prior position, velocity, acceleration, and time headway). The output is a prediction of the vehicle state in the lane-changing process in the next 1s. The training process used was the Levenberg-Marquardt algorithm, where network weights are iteratively adjusted to minimise output errors. When using a prediction horizon of 1s, this modelling framework exhibited an RMSE of 0.0184. The RMSE increased to 0.2273 when the prediction horizon increased to 2s.

Ding et al., demonstrated the effectiveness of using back-propagation neural networks to predict and model lane changes. However, they also demonstrated the difficulty and limitations of such an approach. The training data required to develop the neural network needs to be extensive and complete, which may be difficult to obtain in many circumstances. Additionally, the neural network will exhibit the same bias as the training data, whether the bias is intentional and known or hidden to the modeller. Finally, the neural network will be adept at handling scenarios which it has seen repeatedly or scenarios that are similar to what it has previously experienced. But if this modelling approach is applied to an unfamiliar environment where motorist behaviour deviates significantly from the training data, the risk of incorrect outputs increases.

Tang et al., Model

Tang et al., developed an adaptive fuzzy neural network model designed to predict steering angles [Tang et al., 2018]. Fuzzy logic system based models transform qualitative descriptors into quantitative values through probability distributions. The limitations of fuzzy logic systems include non-comprehensive rules defining the fuzzy inference, a lack of learning mechanisms to correct errors in the interface, and limited incorporation of factors in the analysis. Tang et al., corrected these limitations by developing a fuzzy neural network with adaptive learning abilities. The movement of vehicles is modelled as follows;

$$\begin{aligned} x(\tau_s + 1) &= x(\tau_s) + v(\tau_s) \cdot \tau \cdot \cos(\theta(\tau_s)) \\ y(\tau_s + 1) &= y(\tau_s) + v(\tau_s) \cdot \tau \cdot \sin(\theta(\tau_s)) \\ \theta(\tau_s + 1) &= \theta(\tau_s) + v(\tau_s) \cdot \tau \cdot \tan(\theta_s(\tau_s)/l) \end{aligned} \tag{30}$$

Where, θ is the vehicle heading, x and y represent the cartesian coordinates of

the vehicle, x_0 and y_0 represent its centroid, θ_s is the steering angle, v is the instantaneous velocity, l is the length of the vehicle, τ is the computational time steps, and τ_s is the simulation time step.

The modelling framework consists of a fuzzy neural network receiving instantaneous kinematic and geolocation data. The framework then interacts with an adaptive learning algorithm which uses the accuracy of the outputs to update the fuzzy rules and parameters of the membership functions. The fuzzy rules take the form;

$$\begin{aligned}
R_i : & IF(\Delta x \text{ and is } A_i) \text{ and } (\Delta y \text{ is } B_i) \text{ and } (\Delta v \text{ and is } C_i) \\
& \text{and } (\theta \text{ is } D_i) \text{ and } (acc \text{ is } E_i) \\
& THEN y_i \text{ if } f_i(\Delta x, \Delta y, \Delta v, \theta, acc)
\end{aligned} \tag{31}$$

Where, A, B, C, D, E are the fuzzy sets relating to Δx , Δy , relative velocity, heading angle, and acceleration, respectively. The membership functions are defined by trapezoids displayed in Equation 32;

$$m(x) = \begin{cases} x - c_1/c_2 - c_1 & c_1 \leq x < c_2 \\ 1 & c_2 \leq x < c_3 \\ c_4 - x/c_4 - c & c_3 \leq x < c_4 \\ 0 & x < c_1 \text{ or } c_4 < x \end{cases} \tag{32}$$

Where, ϕ is the membership function for different variables, and c_1 to c_4 are parameters defining the shape of trapezoids bounding each fuzzy rule.

The learning algorithm uses the least-squares error method to minimise the RMSE. The algorithm identifies rules that make the greatest contribution to the output. The most impactful fuzzy sets divide into two, and each is replaced with a new set of rules, adding 2^5 rules to the database. The process repeats, continually updating the membership functions for the fuzzy system.

Ting et al., compared the performance of this adaptive fuzzy neural network approach with other approaches such as neural networks, support vector machine, hidden Markov models, and multivariable linear regression. For vehicles driving at 60 km/hr , the adaptive fuzzy neural network approach showed improvements up to 59% compared to the other methods. Similar improvements were seen when speed increased to 80 km/hr and 100 km/hr .

Hou, et al. Model

Hou et al., developed a lane-changing model that leverages Bayes classifiers and decisions-trees to model lane-changing decision-making [Hou et al., 2013]. Bayes classifiers are an extension of Bayes theory of conditional probability where the classifiers predict the likelihood that a specific outcome, record, or data point, belongs to a particular membership class. The equation for calculating the conditional probability is given as;

$$P(y_i|X) = \frac{p(X|y_i)P(y_i)}{p(X)} \quad (33)$$

Where, X is the input vector, $P(\cdot)$ is the conditional probability and $p(\cdot)$ is the probability density function. If y_1 and y_2 represent the merge and no merge classes, and if for example $P(y_1|x) > P(y_2|x)$, then x is classified to y_1 .

The risk of misclassification is reduced by using a penalty term (λ_{ki}). This penalty term changes the Bayes classification rule to;

$$x \text{ belongs to } y_i \text{ if } l_{12} = \frac{p(x|y_1)}{p(x|y_2)} > (<) \frac{p(y_2)\lambda_{21}}{p(y_1)\lambda_{12}} \quad (34)$$

Hou et al., used the k nearest neighbour method to estimate the class-conditional probability density functions. This method does not force an assumption regarding the distributional form, unlike the maximum likelihood method. The Bayes classifiers were then combined with a decision-tree model. At each node in the tree, the models determine if $x_i \geq a$, where a is a threshold value. The size of the tree plays a key role in the effectiveness of this method, where a small tree greatly increases the likelihood of misclassification, and a large tree overfits the model to the specific training dataset.

The Bayes classifiers used were Δv_{lead} and Δv_{lag} (the speed difference between the ego and lead vehicle, and ego and lag vehicle, respectively), Δx_{lead} and Δx_{lag} (the distance to the lead vehicle and lag vehicle, respectively), and Δx_{lane} (the distance to the end of the merging lane). This method demonstrated an accuracy of 92.3% for merging events and 79.5% for non-merging events. When using the method of decision trees, the accuracy reduced to 80.8% for the merging events and increased to 84.3% for the non-merging events. The two methods were combined, using a “majority voting” rule, resulting in a merging accuracy of 79.3% and a non-merging accuracy of 94.3%.

For safety reasons, it is more important to accurately identify a non-merge event than it is to identify a merge event. Misclassifying a merge event as a non-merge event results in a missed opportunity for the vehicle and the repercussions are minor. Hou et al., demonstrated that by adjusting the misclassification penalty parameters, the accuracy of non-merge events increases, but with a trade-off of reduced accuracy for merging events.

3.3 Gap-Acceptance Algorithms

Gap-acceptance models evaluate the gap in the target lane and determine its appropriateness for lane-changing. They may either exist as standalone models or in conjunction with their lane-changing model counterparts. While the lane-changing model identifies whether or not a vehicle should change lanes, the gap-acceptance model identifies whether the target gap is appropriate. For this reason, many gap acceptance models focus on stream stability and platooning. Violating the buffer space of lead and following vehicles in the target lane (imprudent lane change) is not only dangerous for the lane-changing vehicle, but can lead to erratic and sudden reactions by vehicles in the target lane.

TSS Gap-Acceptance Model

The Gipps gap-acceptance model, as modified by TSS for implementation in the Aimsun microsimulator, calculates the minimum critical gap between the preceding and following vehicle [Transport Simulation Systems, 2014]. If a gap falls below the critical minimum gap, then it is rejected, and the vehicle awaits an acceptable gap to surface in the target lane traffic stream. The minimum upstream gap is calculated using Equation 35;

$$\begin{aligned} gap_{up}(t) \geq & \max(0, \frac{v_{lc}(t)^2}{2d_{lc}} + 0.5v_{up}(t)\tau_{up} + \max(0, -\frac{v_{up}(t)^2}{2d_{up}} \\ & + a_{up}(1 - 0.5a_{up})d_{up}\tau_{up}^2 + (1 - a_{up})v_{up}(t)\tau_{up})) \end{aligned} \quad (35)$$

Where, $gap_{up}(t)$ is the minimum upstream gap, v_{lc} is the ego vehicle velocity, d_{lc} is the ego vehicle deceleration, $v_{up}(t)$ is the upstream vehicle velocity during time t , τ_{up} is the upstream vehicle reaction time, and a_{up} and d_{up} are the upstream vehicle acceleration and deceleration, respectively.

The minimum downstream gap is found using Equation 36;

$$gap_{dw}(t) \geq \max(0, \frac{v_{dw}(t)^2}{2d_{dw}} + 0.5v_{lc}(t)\tau_{lc} + \max(0, -\frac{v_{lc}(t)^2}{2d_{lc}} + a_{dw}(1 - 0.5a_{dw})d_{lc}\tau_{lc}^2 + (1 - a_{dw})v_{lc}(t)\tau_{lc})) \quad (36)$$

Where, $gap_{dw}(t)$ is the minimum downstream gap, v_{lc} and d_{lc} is the velocity and deceleration of the current vehicle respectively, and $v_{dw}(t)$, τ_{lc} , a_{dw} and d_{dw} is the velocity during time t , reaction time, acceleration, and deceleration of the downstream vehicle, respectively.

The Gipps gap-acceptance model provides practical ranges to characterise a safe gap. However, conceptually, the model is difficult to reconcile with the way that human drivers evaluate a gap. The model indefinitely considers the kinematics of target lane vehicles on gap acceptability, regardless of the distance between the vehicles. Additionally, the Gipps gap-acceptance model assumes that a rejected gap will not later become acceptable, despite increasing driver urgency with progressing time.

Lee Model

Vehicle reaction time plays a critical role in gap-acceptance models and has two effects on the traffic stream. The first is that it causes the wave-like reaction lag often observed when vehicles embark from a standstill. The reaction time forces each following vehicle to experience a slight delay before it accelerates, increasing the gap between itself and the preceding vehicle. The reaction time also has safety implications, as it forces larger gaps in high-velocity scenarios.

Lee models the minimum gap as a random variable using a log-normal distribution (Equation 37) [Lee, 2006]. The log-normal ensures that the calculated gaps are always non-negative, as follows;

$$\ln(gap^{i,cr}) = \beta^i X^i + r_n \alpha^i + \epsilon^i \quad i \in (lead, lag) \ \& \ (right, left) \quad (37)$$

Where, $gap^{i,cr}$ is the critical gap in the direction of lane change i , β^i and X^i is the vector of parameters and explanatory variables respectively, r_n is the parameter for the driver-specific random term α^i , and ϵ^i is a random term where $\epsilon^i \sim \mathcal{N}(0, \sigma_g^2)$.

For the gap to be appropriate, the driver must accept both the lead gap and the lag gap. The following equation gives the probability of changing lanes;

$$\begin{aligned}
P(\text{change in direction } i, r_n) &= P(l = i, r_n) \\
&= P(\text{accept lead gap} | i, r_n) \times P(\text{accept lag gap} | i, r_n) \\
&= P(\text{gap}^{\text{lead } i} > \text{gap}^{\text{lead } i, cr} | i, r_n) \times P(\text{gap}^{\text{lag } i} > \text{gap}^{\text{lag } i, cr} | i, r_n)
\end{aligned} \tag{38}$$

Where, i is the chosen target lane, $\text{gap}^{\text{lead } d}$ and $\text{gap}^{\text{lag } d}$ are the available lead and lag gaps in the direction of i respectively, l is the lane-changing action, and r_n is the parameter for the driver-specific random term.

Assuming that critical gaps follow the log-normal distribution, the probability that a gap will be acceptable is given by;

$$\begin{aligned}
P(\text{gap}^{\text{lead } i} > \text{gap}^{\text{lead } i, cr} | i, r_n) &= P(\ln(\text{gap}^{\text{lead } i}) > \ln(\text{gap}^{\text{lead } i, cr} | i, r_n)) \\
&= \phi\left(\frac{\ln(\text{gap}^i) - \beta^i X^i - r_n \alpha^i}{\sigma_i}\right)
\end{aligned} \tag{39}$$

Where, $\phi[\cdot]$ is the cumulative standard-normal distribution.

The computation of the critical minimum gap value isn't particularly pertinent to modelling the interactions between CAVs. An ideal model for calculating critical gap values would focus on the psychological implication of small gaps on the driver, or the efficiency of the traffic stream. Focusing on collision avoidance may play a reduced role as near-instantaneous access to vast amounts of data could render collisions unlikely.

Farah et al., Model

Farah et al., used synthetic data generated from a driving simulator and data from a stated preference survey to develop a detailed critical gap-acceptance model [Farah et al., 2009]. The questionnaire consisted of two parts, with the first part collecting socioeconomic data about the participants. This data included age, gender, marital status, education, income, and previous car accidents. The second part consisted of a 44-question assessment of questions grouped into four categories, where respondents were able to select a score

on a six-point scale. The categories included recklessness and carelessness, anxiousness while driving, anger and hostility, and patience and carefulness.

The participants drove four different uninterrupted rural highway scenarios, each with a length of approximately $7.5km$. The simulator collected vital information in $0.1s$ increments, information including longitudinal and lateral position, speed and acceleration of the subject vehicle and all other vehicles in the scenario, times and location of passing manoeuvres, distances between vehicles, and relative speeds. The data collected from the experiment was used to develop a gap-acceptance model. The model has a binary outcome where drivers either accept or reject a gap, by comparing it to an unobserved critical gap. This model is given in Equation 40;

$$X_{n,t} = \begin{cases} 1 & \text{if } gap_{n,i} \geq gap_{n,i}^{cr} \\ 0 & \text{if } gap_{n,i} < gap_{n,i}^{cr} \end{cases} \quad (40)$$

Where, i is a unique gap identifier, $X_{n,i}$ is the choice variable for driver n , $gap_{n,i}$ is the available gap, and $gap_{n,i}^{cr}$ is the critical gap value.

The critical gap is defined as the point where a driver is indifferent to accepting or rejecting a gap. Since this value is not directly observable, it was modelled as a random variable given by the following equation;

$$gap_{n,i}^{cr} = X_{n,i} \cdot \beta + \epsilon_{n,i} \quad (41)$$

Where, $X_{n,i}$ is the vector of explanatory variables, β is the vector of associated parameters, and $\epsilon_{n,i}$ is a random error term following the logit distribution. The probability of accepting a gap is then determined using a binary logit choice model, given in Equation 42;

$$P_{n,i}(accept\ gap) = \frac{1}{1 + e^{-\mu \cdot (G_{n,i} - X_{n,i} \cdot \beta)}} \quad (42)$$

Where, μ is a scale parameter for the model, which is inversely proportional to the standard deviation of the critical gap distribution.

Farah et al., tested the performance of five different models. Model 1 included traffic condition variables, model 2 further added geometric road design, model 3 and model 4 further added driving style scores and socioeconomic characteristics, respectively, and finally, model 5 contains all of the parameters in model

1 to 4. Likelihood tests showed that model 4 best fit the gap-acceptance data observed from the driving simulators. The fitted model 4 is as follows;

$$\begin{aligned} gap_{n,i}^{cr} = & 34.01 - 0.30SS + 5.15FG + 0.42FS - 0.14OS \\ & - 2.35RG - 7.00Age_1 - 4.95Age_2 - 2.84G \\ & + 0.23P + 1.05Km - 7.5E - 5CD \end{aligned} \quad (43)$$

Where, the parameters represent various behaviour, geometric, and socioeconomic characteristics. Model 4 has a level of significance of $p-value = 0.4645$.

Farah et al., successfully demonstrated the influence of socioeconomic variables and personal driving factors on accepting gaps. While these factors are strongly representative of human driving, they play a supplementary role in CAV operation. Either CAVs will operate more predictably, following rigid logic, in which case these factors are redundant in CAV gap-acceptance. Or, understanding that these parameters influence human gap-acceptance could form the basis of CAV gap-acceptance behaviour. The second behavioural design decision would allow a smoother transition from human drivers to CAVs and allow their decisions to be more random, distributed, and human-like.

Kim et al., Model

Kim et al., proposed dynamically updating the critical gap parameter in real-time, making it contingent on the location of the vehicle in the merging lane [Kim et al., 2008]. This approach differed from other gap-acceptance models that generally hold the critical gap value constant. The authors conducted a comprehensive study into the effects of leading and trailing gaps, target lane density, vehicle location on the target lane, and relative speed, to evaluate the factors that greatest contributed to motorist gap-acceptance behaviour. Data was collected on two different merge junctions on a motorway in Seoul, Korea.

For the leading and trailing gap, analysis showed that 15%, 50%, and 85% of drivers accepted a leading gap of 0.35s, 0.8s, and 1.9s, and a trailing gap of 0.6s, 1.4s, and 3.3s, respectively. The small range in variability showed that difference in gap-acceptance between drivers was 0.3s. The authors concluded that small variances did not warrant treating lead and lag gaps as variable.

Impact of traffic density in the target lane on motorist gap-acceptance was assessed by partitioning the data into four density scenarios, ranging from $20veh/km$ to $50veh/km$. In each category, the driver data showed how far into the acceleration lane a motorist travelled before changing lanes. The results demonstrated that the location distribution of the lane changes appreciably varied with lane density in the target lane, with drivers often accepting a gap at the start of the merging section in low traffic density.

The chi-square (χ^2) statistical test was performed to determine whether vehicle location on the acceleration lane affected motorist gap-acceptance. The test evaluated whether distributions of the accepted gap location were comparable, using a confidence level of 95%. The results showed that the total length of the acceleration lane, and distance of the motorist to the end of the acceleration lane, both influenced gap-acceptance behaviour. Finally, Kim et al., found that drivers accept a gap when relative speed is within $\pm 15km/hr$.

Using this information, Kim et al., developed a table which segmented the acceleration lane into eight sections, and the driving fleet into ten driver types. The collected data was utilised to calculate the critical minimum gap for each driver type and each acceleration lane segment. These critical gap values had the advantage of varying with distance to the end of the acceleration lane, providing this framework with the advantage of using a dynamic critical gap distance. The microsimulator checks the vehicle location, then evaluates if the lead and trailing gap and relative speed are within the thresholds defined for the current segment, with a positive outcome resulting in an accepted gap.

Evaluation of this gap-acceptance framework demonstrated strong statistical evidence that the modelled data represented observed data, both for determining gap-acceptance concerning density on the target lane and concerning relative speed between vehicles. While statistical significance was established, the authors did not evaluate the predictive qualities of the modelling framework in a new environment or with different data. Frameworks such as this risk overfitting calibrated parameters to the measured data and then reproducing the observed data during runtime. If such an occurrence did affect the modelling results, then the gap-acceptance modelling framework of Kim et al., may require substantial calibration to each location in which it is implemented.

Hwang and Park Model

Hwang and Park used the probit discrete choice model to calculate the likelihood of accepting a gap [Hwang & Park, 2005]. Hwang and Park state that the use of a utility function in conjunction with the probit discrete choice model can appropriately reflect human gap-acceptance decision-making. The mathematical expression of the model is as follows;

$$\begin{aligned} P_a &= \Phi\left(\frac{V_a - V_r}{\sigma}\right) \\ P_r &= 1 - P_a \end{aligned} \tag{44}$$

Where, P_a and P_r is the probability of accepting or rejecting a gap respectively, and Φ is the cumulative normal distribution function.

Field data was used to determine the explanatory variables that affect gap-acceptance and warrant inclusion in the model. The data consisted of vehicle location, length, velocity, and headway, in 0.5s increments. Initial explanatory variables consisted of total gap, lead and lag gap, front gap, remaining distance, subject heavy vehicle (dummy variable), object heavy vehicle (dummy variable), and velocity difference between the subject, lead, and lag vehicle. Parameter estimation was conducted on all variables, and those that showed low statistical significance were iteratively removed. The final model consisted of seven explanatory variables that affect gap-acceptance; lead and lag gap, front gap, remaining distance, subject heavy vehicle (dummy variable), and object heavy vehicle (dummy variable).

The style of modelling and data analytics in this study provides useful insights into the human behavioural component of gap-acceptance. While the model itself is simplistic, the statistical analysis on factors that influence gap-acceptance helps inform geometric design and driving policy. It also helps inform the factors that CAVs must be sensitive to if they are to mimic human drivers and integrate well in a mixed fleet, or avoid if they are idealised drivers. As with any other study that is contingent on data collection from a specific source and the use of curve fitting, this approach must be cautious of the inclusion of outliers, bias data, or data is not representative of general network operation.

Unsignalized Intersection Gap-Acceptance

Amina and Maurya assessed the applicability of gap-acceptance models to complex unsignalized and uncontrolled intersections [Amin & Maurya, 2015]. These intersections either lack a clear definition of priority movements, or drivers simply ignore priority and enter the intersection regardless of convention or oncoming traffic. The authors question the appropriateness of the assumption in many gap-acceptance models, primarily that drivers are rational and consistent. This assumption implies that gap-acceptance models generally do not permit a gap to be accepted that is smaller than one they have already rejected, and it implies that microsimulation agents will repeatedly perform the same actions when presented with the same situation.

The authors collected data from a four-legged uncontrolled intersection located in a semi-urban area of Ahmedabad (India). Five video cameras were used to gather critical information such as vehicle arrival rate/time, accepted and rejected gap/lag time, speed and type of conflict vehicles, waiting time of minor street vehicles at stop lines, driver's age and gender, and vehicle occupancy. The data was then analysed using several different critical gap calculation approaches. Their applied approaches that have not been discussed in this literature review include;

- The Raff method for calculating critical gap states that the critical gap time is when the frequency of gaps rejected below that time is equal to the frequency of gaps accepted above that time, when the drivers are indifferent to accepting a gap [Raff et al., 1950]. This method is similar to the Greenshield method, which plots the frequency of accepting or rejecting a gap against the gap size, and the critical gap is that which shows equal acceptance and rejection [Gattis & Low, 1999].
- The lag method and Harder's method [Brilon et al., 1999] both use lags or gap times to calculate the critical gap. These methods observe the size of the gaps that drivers on the minor approach arm accept, and assumes that this gap time is an appropriate reflection of the critical gap of the system. This method is effective in undersaturated traffic environments.
- Ashworth's method states that if the critical gap time (t_c) and follow-up time (t_f) for vehicles in a system were both normally distributed,

then the accepted gap time (t_a) could be modelled as $E(t_c) = E(t_a) - q_p Var(t_a)$, where q_p is the mainstream flow [Ashworth, 1968].

- Wu's model presents a probability equilibrium between accepting and rejecting gaps [Wu, 2006]. The advantage of this approach is that it does not require the assumption of a probability distribution for observed gaps, or the assumptions that the observed gaps are an appropriate reflection of the system.

Using these methods, the authors compared theoretical results with observed measurements. The theoretical models substantially overpredicted the critical gap time for an unregulated and unsignaled intersection, with observations showing that the minor approach forced the mainline approach to slow down and yield at times. This study demonstrated a clear limitation in many of the theoretical models presented for calculating gap-acceptance, their reliance on stringent rules and traditional intersection operation. These models may also be over-representative of the environments from which they were developed and calibrated, showing their lack of versatility in new modelling environments.

Pollatschek et al., Model

Pollatschek et al., presented a model that focused on gap-acceptance on secondary lanes when waiting at an unsignaled intersection, directly addressing the issue raised by Amina and Maurya [Pollatschek et al., 2002]. The model balances the risk of accepting small gaps against the disbenefit of not accepting them, achieved by accounting for perceived risk and waiting time. The authors developed a relationship between the increase in accepting a gap and the fall of perceived risk, when waiting time increases. The average waiting time for a gap is calculated as follows;

$$W(i) = i[e^{t_{nr}/i} - (1 + \frac{t_{nr}}{i})] \quad (45)$$

Where, $W(i)$ is the wait time for an average perceived gap on the mainline i , and t_{nr} is the gap time above which there is a negligible risk of an accident.

The perceived risk is then calculated as;

$$r(t) = \begin{cases} \infty & t \leq t_{sa} \\ r'_0(t - t_{sa})^{-\alpha} & t > t_{sa} \end{cases} \quad (46)$$

Where, $r(t)$ is the perceived risk, and t_{sa} is the gap time under which an accident is virtually guaranteed.

Finally, an estimation of the average gap on the mainline is given by;

$$gap_i = \frac{u \cdot gap_0 + \sum_{j=1}^i t_j}{u + 1} \quad (47)$$

Where, gap_0 is the priori estimate of the average gap, gap_i is the updated estimate, u is the shape parameter of the gamma distribution function that is used to estimate the priori state of the traffic, and t_j is the rejected gaps.

This approach to gap-acceptance modelling quantifies the risk a motorist perceives, without considering many of the stochastic factors that may affect a drivers tendency towards taking risky actions. Mood, trip purpose, random variability, time of day, vehicle occupancy, trip length, familiarity with the environment, and countless other factors may affect the motorist's willingness to accept risks. The authors acknowledged the limitations of taking capacity values from this model and applying them to other geographic regions, or those that exhibit a different driving behaviour profile. Additionally, CAV operations would not be affected by many of the factors listed above, and so risk calculations would be more deterministic.

3.4 Summary and Model Comparison

The literature review contains a myriad of diverse model types, each with their own advantages and disadvantages. GHR models calculate acceleration as a proportional response to following distance and velocity difference between the ego and lead vehicle. While mathematically simple in nature, the major limitation of these models is continual minute adjustments to acceleration in response to external stimuli, regardless of the distance between the ego and lead vehicles, unless a hard-cap is placed on the model. This tendency places an increasing computational strain on hardware during emulation when the number of vehicles in the simulation increases. Additionally, such an approach is not appropriate for CAVs due to the computational and mechanical strain of repeated minor adjustments to control actuators. Collision avoidance models generally suffer from the same limitations as GHR models, but they differ in that acceleration is in response to a safe gap as opposed to a leaders kinematics.

Psycho-physical models address the limitation of GHR and Collision Avoidance models by using “action points” and thresholds to trigger discrete events during car-following. The thresholds are calibrated based on observed site-specific data in varying circumstances such as urban and rural, congested and free flow, freeway and arterial, or multi-lane and single-lane. The psycho-physical car following modelling technique is an attractive approach to use with CAVs, especially when combined with the rule-based lane-changing modelling approach. The modular nature of psycho-physical and rule-based models allows infinite discretisation and incorporation of modules to allow the model to react to any number of circumstances. This characteristic makes them applicable to CAV behaviour, where CAV reactions may change depending on a vehicle’s future intentions or desire to cooperate. As the features for CAVs are developed, a psycho-physical and rule-based model would give way to expansion and easy modification.

Fuzzy-logic, machine learning, and artificial intelligence models are an emerging and promising field. These models use a data-driven process to calibrate and train learning algorithms autonomously, converting the modelling process from a calibration exercise to a data analysis exercise. The limitation of these models is in their requirement of extensive and diverse quantities of data. The neural networks that depend on reinforcement learning must be exposed to the same circumstances numerous times before the probabilistic pathways reliably emulate appropriate human decision making. A model calibrated for one environment may not be applicable for another due to differing driving behavioural characteristics, meaning that the difficult data acquisition and calibration process must be repeated. The nature and workings of the neural networks are also hidden from the modeller, meaning that mistakes can not be traced or manually corrected. Finally, this style of modelling is not appropriate for CAV emulation in the near-term, where the quantity of data required for calibration and validation is not available. However, they may form the most appropriate and reliable form of emulation models in the long-term, when training data is in abundance.

4 Mathematical Models for the Microsimulation Emulation of CAV Behaviour

Growing interest in CAV technology has resulted in increased academic response. The literature contains numerous different techniques, models, and algorithms for the microsimulation emulation of CAV behaviour. Many newly developed models follow a proactive approach, where understanding vehicle technology dictates behaviour rather than simply replicating it. The proactive approach is possible as CAVs follow algorithms in their operations. Although, the intractable nature of machine learning algorithms and artificial intelligence may result in CAV behaviour being as stochastic as human driving behaviour, undermining the proactive approach.

This chapter contains a review of the methods used to model CAV behaviour, with a more detailed description provided for specific unique and novel approaches. The chapter is subdivided into four sections, exploring car-following, lane-changing, gap-acceptance, and automated intersection operation. While autonomous intersections are only possible in a 100% CAV fleet, exploring how they are modelled can provide insights into space negotiating techniques for CAVs during regular network operations.

4.1 Car-Following Models

The car-following models for human vehicles presented in Chapter 3 differ from those for CAVs primarily in their treatment of kinematic variables. While human vehicle models assume the kinematics of surrounding vehicles and capture inherent stochastic behaviour, CAV models rely on the availability of accessible and correct kinematic information. CAV models use the expected capabilities of V2X as critical points in their development. V2X allows models to focus on elements such as information exchange, trust, platooning, and data verification, elements that were not necessary for the human variant of car-following models. The remainder of this section presents in brief detail novel and unique examples of CAV car-following models.

AICC Algorithm

Congestion in highway environments is a growing concern in major cities worldwide. Historically, the highway congestion problem has been addressed by increasing the capacity through expansion. However, industrialisation and accelerating urban sprawl renders this solution no longer viable. The AICC algorithm aims to reduce headways between vehicle and significantly increase the capacity of all roads, while still maintaining occupant safety and comfort [Ioannou & Chien, 1993]. Ioannou and Chien developed a car-following algorithm that maintains constant headways and following distances. Equation 48 shows the formulation used to define a safe distance for each vehicle;

$$\Delta x_{d_n} = \Delta x_{m_n} + \Delta x_{a_n} \quad (48)$$

Where, Δx_{m_n} is the minimum distance required to avoid collisions, and Δx_{a_n} is a parameter included to ensure greater safety and comfort.

Equation 49 provides the expression for Δx_{m_n} , which considers the maximum allowable jerk (J_{max}) during acceleration and braking;

$$\begin{aligned} \Delta x_{m_n} = & \frac{v_n^2 - v_{n-1}^2}{2a_{n-1}} + v_n \left(\tau + \frac{a_n + a_{n-1}}{J_{max}} + \frac{1}{a_{n-1}} \left(a_n \tau + \frac{a_n(a_n + a_{n-1})}{J_{max}} \right. \right. \\ & \left. \left. - \frac{1}{2} \frac{(a_n + a_{n-1})^2}{J_{max}} \right) \right) + \left(\frac{1}{2} a_n \tau^2 + \frac{a_n(a_n + a_{n-1})^2}{2J_{max}^2} \right. \\ & \left. - \frac{1}{6} \frac{(a_n + a_{n-1})^3}{J_{max}^2} + \frac{a_n(a_n + a_{n-1})}{J_{max}} \tau \right. \\ & \left. + \frac{1}{2a_{n-1}} \left(a_n \tau + \frac{a_n(a_n + a_{n-1})}{J_{max}} - \frac{1}{2} \frac{(a_n + a_{n-1})^2}{J_{max}} \right)^2 \right) \end{aligned} \quad (49)$$

Where, v_{n-1} and v_n are the velocities of the lead and following vehicles respectively, a_{n-1} and a_n are the maximum deceleration of the lead and following vehicles respectively, J_{max} is the maximum jerk, τ is the time required to detect the stopping manoeuvre (sensor communication delay for CAVs), and P_n is the vehicle engine input.

The following system of equations describes the vehicle kinematics;

$$\begin{aligned}
\frac{d}{dt}x_n(t) &= \dot{x}_n(t) = v_n(t) \\
\frac{d}{dt}\dot{x}_n(t) &= \ddot{x}_n(t) = a_n(t) \\
\frac{d}{dt}\ddot{x}_n(t) &= b(\dot{x}_n, \ddot{x}_n) + a(\dot{x}_n)P_n(t)
\end{aligned} \tag{50}$$

Where, $a(\dot{x}_n)$ and $b(\dot{x}_n, \ddot{x}_n)$ are given by the following equations;

$$\begin{aligned}
a(\dot{x}_n) &= \frac{1}{m_n \tau_n(\dot{x}_n)} \\
b(\dot{x}_n, \ddot{x}_n) &= -2 \frac{k_{d_n}}{m_n} \dot{x}_n \ddot{x}_n - \frac{1}{\tau_n(\dot{x}_n)} (\ddot{x}_n + \frac{k_{d_n}}{m_n} \dot{x}_n^2 + \frac{d_{m_n}(\dot{x}_n)}{m_n})
\end{aligned} \tag{51}$$

Where, x_n , v_n and a_n is the position, velocity and acceleration of the n^{th} vehicle respectively, m_n is the vehicle mass, τ_n is the interpretation and reaction time of human drivers and communication sensor delay for CAVs, k_{d_n} is the aerodynamic drag coefficient, and $d_{m_n}(\dot{x}_n)$ is the mechanical drag coefficient.

In simulating a highway network, Ioannou and Chien demonstrated that the “human look ahead model”, “human linear follow-leader model” and “human optimal control model” resulted in flow rates of $500veh/hr$, $1000veh/hr$ and $2200veh/hr$, respectively. Alternatively, the AICC model produced flow rates of $3500veh/hr$, resulting in a 59% capacity increase over the best simulated human driving models.

Car-following algorithms such as AICC indicate that increasing CPR also increases transportation network capacity without demanding significant additional infrastructure. Cities could economically cater to increasing demands without sacrificing land and resources through the creation of more highways or lanes. The introduction of CAVs in this manner would lead to immediate economical remediation of the world’s major congestion problems.

Although, Ioannou and Chien did not investigate the performance of the AICC algorithm in mixed fleets, nor did they investigate the impact of mixed fleets on highway capacity. While homogeneous behaviour from completely human and completely CAV fleets may result in stable performance, the challenge is in understanding how heterogeneous fleets and CAV-human interactions will affect network and highway performance.

Gu  riaux et al., Model

Gu  riaux et al., proposed a multi-anticipative car-following model that improved traffic stability by weighting the past influences from surrounding vehicles [Gu  riaux et al., 2016]. The model is divided into three layers to reliably consider information from the vehicle surroundings and limit disturbances caused by behavioural decisions. The physical layer governs vehicle dynamics and car-following behaviour. The communication layer controls information exchange based on proximity to other vehicles and reliability rules, expressed as probabilities. Finally, the trust layer models reliability in the information received and evaluates its trust in the surrounding vehicles.

Human-driven vehicles are emulated using the IDM car-following model, which gives acceleration as a function of the vehicle’s current velocity, headway, and relative velocity [Treiber & Kesting, 2013]. The model is as follows;

$$a_n = c_1(1 - (\frac{v_n}{v_0})^{c_2} - (\frac{x_n}{x_0})^2) \quad (52)$$

Where, the first term v_n/v_0 compares the current speed (v_n) to the desired speed (v_0), the second term x_n/x_0 compares the desired spacing (x_n) to the current spacing (x_0), and c_1 and c_2 are calibration terms. The desired spacing is given by;

$$x_n = x_0 + \max(0, v_n \Delta t + \frac{v \Delta v}{2\sqrt{a \cdot d}}) \quad (53)$$

Where, x_n is the current vehicle gap, x_0 is the minimum gap, v_n is the current velocity, Δt is the time headway, Δv is the relative velocity, and a and d are the comfortable acceleration and deceleration rate, respectively.

The bilateral multi-anticipative model provided in Equation 54 emulates CAV behaviour. Multi-anticipative refer to the nature of this model, as it considers the bi-directional impact of multiple lead and surrounding vehicles. The behaviour of the vehicle is then a response to the kinematic requirements imposed by each vehicle. Such an approach appreciably improves flow stability [Monteil et al., 2014]. The model is as follows;

$$a_n = f_{IDM}(v_n, \sum_j U_{nj} \Delta x_{n+j}, \sum_j U_{nj} \Delta v_{n+j}) \quad (54)$$

Where, f_{IDM} is the IDM function presented in Equation 52, and U_{nj} is the interaction coefficient that represents the weighting on the information between agent n and j (an indication of trust between the two vehicles).

While evaluating this framework, certain vehicles were fitted with “faulty” sensors that provided higher errors of up to 50% in speed and headway. Trust is computed based on a weighted average of the past 6,000 interaction results between the vehicle and other agents. Guériau et al., address cooperation uncertainties in average lane speeds and stability by changing CPR and maintaining consistent CAV behaviour.

Without CAVs in the network, their microsimulation modelling indicated that vehicles on the merge lane slowed to a complete halt as they waited for appropriate gaps in the mainline to arise. However, with 50% CAVs in the network, both the merging vehicles and the mainline maintain near free flow speeds. The centralised roadside unit implementing the control protocol homogenised the speed and headway among the onramp and the mainline lanes, leading to improved throughput, traffic flow, and traffic stability.

While it is valuable to understand how changing CPR effects networks, static behaviour between scenarios may not be an appropriate assumption. CAVs can alter behaviour in response to technological advancement, implementation policy changes, and increasing fleet integration. This ability to change behaviour means that it may be valuable to vary the degree to which a CAV operates as a CAV in the network, meaning that CAV behaviour would become more distinct as their penetration rate increased. Such an approach would smoothen the transition between driving styles in mixed fleets, without an abrupt introduction of a vehicle type with a stark difference in behaviour.

Jia and Ngoduy Model

Jia and Ngoduy developed a platooning-based cooperative driving model, with specific consideration of inter-vehicle communication [Jia & Ngoduy, 2016]. This work expanded traditional car-following models, which are generally a function of vehicle speed, headway, and speed of the lead vehicle. Jia and Ngoduy refined CAV behaviour using the capabilities and limitations of the

IEEE 802.11p networking standard, and the following assumptions;

- Platoons exhibit smaller intra-vehicle, and larger inter-vehicle, spacing.
- All vehicles are assumed to be identical.
- Vehicle speed and position is measured precisely.
- Each vehicle has a fixed transmission range and a frequency of $10Hz$.

The general equation for the car-following acceleration model is;

$$a_n(t) = f(v_n(t), \Gamma_1(\Delta x_n(t)), \Gamma_2(\Delta v_n(t)), \dots) \quad (55)$$

Where, $n \in N_1(t)$ for Γ_1 , $n \in N_2(t)$ for Γ_2 , ...)

Where, $a_n(t)$, $v_n(t)$, and $\Delta x_n(t)$ is the acceleration, velocity, and distance headway for vehicle n at time t . N_1, N_2 are time dependant topologies which consist of a varying number of vehicles depending on the surroundings of the ego vehicle. Γ_1, Γ_2 represent the corresponding control algorithm for the parameters over which they act. This model is expandable to include any number of variables that may affect CAV operation.

The state of the platoon leader takes the average of the vehicle states in the preceding platoon. The following vehicles then attain the state of the platoon leader. The consensus algorithm provides this state and is given by;

$$\begin{aligned} u_{i,k}(t) = & \sum_{j=1}^N a_{ij} \{ \\ & \gamma_1 [x_{j,k}(t - \tau_j) - x_{i,k}(t) + b_i v_0(t - \tau_0(t)) \tau_j \\ & \quad + (1 - b_i) v_{0,k}(t - \hat{\tau}_0) \tau_j - (i - j) \cdot s] \\ & + \gamma_2 [v_{j,k}(t - \tau_j) - v_{i,k}(t)] \} \\ & + \beta b_i \{ \gamma_1 [x_{0,k}(t - \tau_0(t)) - x_{i,k}(t) + v_{0,k}(t - \tau_0) \tau_0(t) - i \cdot s] \\ & + \gamma_2 [v_{0,k}(t - \tau_0(t)) - v_{i,k}(t)] \} \\ & + \beta (1 - b_i) \{ \gamma_1 [x_{0,k}(t - \hat{\tau}_0(t)) - v_{0,k}(t - \hat{\tau}_0) (\hat{\tau}_0 - \tau_0) - x_{i,k}(t) - i \cdot s] \\ & + \gamma_2 [v_{0,k}(t - \hat{\tau}_0(t)) - v_{i,k}(t)] \} \end{aligned} \quad (56)$$

Where, a_{ij} is the $(i, j)^{th}$ entry in the adjacency matrix, β , γ_1 and γ_2 are control parameters, $x_{0,k}(t - \hat{\tau}_0)$ and $v_{0,k}(t - \hat{\tau}_0)$ is the previous known location and

velocity of the platoon leader respectively (used in the case of lost or corrupted communication), τ_j is the delays to vehicles, and b_i is an indicator of whether the platoon data is reaching all members of the platoon.

From Equation 56, the first and second expressions represent the estimated position and velocity error between vehicle i and j , respectively. The third and fourth expressions represent the error in position and velocity received from the leader, respectively if the communication is received. Finally, the fifth and sixth expressions are the same as the third and fourth, but instead apply to the last successful communication received by the vehicle.

Jia and Ngoduy’s model provides a novel means of emulating platoon behaviour and the effect of intra- and inter-vehicle communication. This modelling framework may be beneficial in verifying the technology that will form the backbone of the future mobility landscape. However, using such a protocol for microsimulation purposes may not be appropriate as this work focuses on the communication process, as opposed to the behavioural outcomes and implications of vehicle communication.

Talebpour and Mahmassani Model

Talebpour and Mahmassani developed a set of mathematical models to address the limitations of existing models in a mixed fleet setting [Talebpour & Mahmassani, 2016]. Their study uses three major models to emulate vehicles with no communication capabilities, vehicle-to-vehicle communication, and CAV communication. The models are as follows;

- Vehicles with no communication capability: A model developed by Hamdar et al., [Hamdar et al., 2008] and extended by Talebpour et al., [Talebpour et al., 2011] is used for this situation. The model is based on collision avoidance due to its high weighting on driver decision making.
- Vehicle-to-vehicle communication: The IDM [Kesting et al., 2010] emulates communication between vehicles, notifying vehicles of the downstream driving environment, including weather and road conditions.
- CAVs: Models to describe CAV operation are based on the assumptions that a CAV may only travel fast enough to still come to a complete stop

within its sensing distance. The CAV maintains a headway and velocity sufficient to facilitate coming to a complete standstill should the leader chose to decelerate at its maximum rate.

To maintain safety, the maximum velocity of a CAV is given by;

$$v_{max} = \sqrt{-2a_n^{decc} \Delta x} \quad (57)$$

Where, v_{max} is the maximum velocity, and a_n^{decc} is the maximum deceleration for the vehicle. $\Delta x = \min[sensor\ detection\ range, \Delta x_n]$. Δx_n is the headway between the vehicle and its leader, and is given by;

$$\Delta x_n = (x_{n-1} - x_n - l_{n-1}) + v_n \tau + \frac{v_{n-1}^2}{2a_{n-1}^{decc}} \quad (58)$$

Where, x_{n-1} and x_n are the position of the lead and ego vehicle, l_{n-1} is the length of the lead vehicle, and τ is the reaction time of the ego vehicle.

This safety constraint and the acceleration of the vehicle is given by a model developed by Van Arem et al., as follows [Van Arem et al., 2006];

$$a_n^d(t) = c_a a_{n-1}(t-\tau) + c_v (v_{n-1}(t-\tau) - v_n(t-\tau)) + c_x (\Delta x_n(t-\tau) - \Delta x_{min}) \quad (59)$$

Where, $a_n^d(t)$ is the deceleration rate of vehicle n at time t , Δx_n is the headway between the vehicle and its leader, Δx_{min} is the minimum allowed headway, and c_a , c_v and c_x are calibration parameters that alter the influence of the lead vehicle on the following vehicle.

These modelling approaches provide a robust solution to emulating CAV behaviour in a microsimulation network setting; however, they do not make special consideration for lateral communication and inter-vehicle cooperation. These models are effective at emulating platooning behaviour, and their framework comprehensively covers mixed fleets and vehicles with limited communication capabilities. However, the rigidity of the framework makes it difficult to fine-tune and modify intricate details, which would aid its integration with a microsimulator. Additionally, lateral communication and active cooperation during network operations are not considered.

4.2 Lane-Changing Models

The human-oriented lane-changing models differ from models for CAVs in that CAV models often place a greater focus on system utility. Human lane-changing models reflect the utilitarian and selfish actions of human drivers, and do not aim to create situations of local optimum performance. Instead, they focus on modelling the needs of the individual. CAV lane-changing models place a greater influence on quantifying the impacts of the actions of one vehicle on the localised system. They tend to design actions in a way that either requires cooperation from surrounding vehicles, or changes behaviour based on the impact of actions on surrounding vehicles. Outlined in this section are a range of novel and unique lane-changing frameworks for CAVs.

Antoniotti et al., Model

Antoniotti et al., presented a merge control protocol for the emulation of CAVs at highway merge junctions [Antoniotti et al., 1997]. The controller dictating vehicle operations aims to meet the following hierarchical objectives;

- There must be no accidents.
- All vehicles must merge.
- Traffic must remain as smooth as possible.

The first condition monitors the gap between vehicles, ensuring that vehicles trajectories do not overlap. The second condition forces vehicles in the main-line to yield to merging vehicles, forcing “niceness” in the CAV operation. The control protocol is programmed in the SHIFT language, with a vehicle deceleration between $-0.5g$ and $0.1g$. The controller maintains a specific time headway between vehicles, so the acceleration is calculated as follows;

$$a_n = c_1 \times (\Delta t_n / \Delta t - 1) + c_2 \times (v_n - v_{n-1}) \quad (60)$$

Where, c_1 is the headway calibration parameter (set at 7), $\Delta t_n / \Delta t$ is the ratio of the actual to desired time headway, v_n is the velocity of vehicle n , and c_2 is the speed calibration parameter (set at 0.5).

The vehicle starts in the “Enter Merge Lane” state, where it attains an initial speed of $22m/s$. It then moves to the “Align to Gap” state after $240m$,

where the vehicle adjusts its velocity to the mainline. This state sets vehicle acceleration to the minimum value imposed by the headway to its current leader and the lead vehicle in the target lane. The vehicle then enters the “Merge State”, where the lag vehicle in the target lane reduces its relative velocity to less than $4m/s$, and the vehicle changes lanes once the appropriate gap has been created. Once the lane-changing manoeuvre is completed, the vehicle rests in the “Cruise” state.

This lane-changing protocol was evaluated by implementing it on a $10km$ highway environment with three entry ramps, tested under two demand scenarios (low demand and high demand). Under the low demand scenario, the control protocol safely and efficiently controlled vehicle operation. The speed of vehicles did not exceed the limit of $28m/s$, nor did they drop below $24.5m/s$. However, the high demand case highlighted the limitation of this modelling approach. The simplistic hard limits set on acceleration caused successive vehicles to reach their maximum deceleration rate without having appropriately adjusted their kinematics in retaliation to the preceding vehicle. This occurrence resulted in collisions under saturated conditions.

While this protocol is simple and effective in low demand scenarios, it is unscalable to higher demands. However, their results indicate that the use of CAVs could reduce the congestion on highways, to the point of complete mitigation. This algorithm would also mean that ramp metering techniques would no longer be required to maintain flows into merging junctions, which currently operate as congestion control mechanisms for many onramp arrangements.

Liu et al., Model

Liu et al., developed a means of predicting vehicle trajectory during lane-changing using a hidden Markov model [Liu et al., 2014]. Their work improved the early detection of accidents and warning systems for drivers. Though the framework was not explicitly developed for use in microsimulation, no part of the framework precludes its implementation in microsimulation. The framework segregates the trajectory of the vehicle into three segments; the beginning stage, the cutting in stage, and the ending stage.

The beginning stage observes the behaviour of the vehicle and classifies it as either “dangerous” or “normal”. A stochastic process using a hidden Markov model classifies the vehicle by considering the historical states of the vehicle. The hidden Markov model (λ) is expressed as;

$$\lambda = \{N, M, P_\pi, P, e\} \quad (61)$$

Where, N is the series of discrete hidden states, P is the probability of transitioning from one state to another, P_π is the initial probability of each state when the systems begin before reaching the observation interval, and e is the likelihood of a vehicle in a specific state emitting a particular observation.

The hidden Markov model uses prior observations and the forward algorithm to determine the probability of observing a particular outcome, given the behaviour it has already observed. The forward algorithm is initialised using the initial state distribution (π). The probability of subsequent observations is then calculated using Equation 62;

$$a_j(k) = e_j(x_k) \sum_{i=1}^N a_i(k-1)T_{ij} \quad (62)$$

Where, $a_j(k)$ is the probability of observing the subsequent observation k , $e_j(x_k)$ is the probability of observing the current observation x_k , and T_{ij} is the likelihood of transitioning between states at time k . The likelihood of observing the sequence of events x is given by,

$$P(x|\lambda) = \sum_{i=1}^N a_i(K) \quad (63)$$

Where, K is the set of observations.

Liu et al., trained the algorithm using the SHRP2 sample dataset and the 100 Car near-crash dataset. The authors found that the trained model correctly predicted the trajectory of “normal” and “aggressive” drivers 92.5% and 80.0% of the time, respectively. While the use of a hidden Markov model and the forward algorithm for trajectory prediction showed strong results, the innate data requirements of such a framework were not discussed.

The authors successfully demonstrated that driver behaviour affects lane-

changing trajectories, meaning that a CAV operating under such a prediction framework would require training under a vast array of conditions. The vehicles would need to witness driver behaviour pertaining to different driving conditions such as network geometry, congestion, weather effects, and time of day, and also consider different behavioural types in response to each of these stimuli. The framework is adept at forecasting trajectories, but the authors do not discuss the data required and the risk of over-training the algorithm.

Ho et al., Model

Ho et al., used the bicycle model for vehicle guidance during lane-changing [Ho et al., 2009]. The proposed algorithm generates a virtual curve that connects the current lane with the target lane. Their framework begins with passing the vehicle angle, data from the lane-keeping controller, and the virtual curve, to the bicycle model. The bicycle model then estimates the vehicles lateral position and returns the value to the controller. The controller uses the returned values to adjust the vehicle relative to its expected trajectory.

The bicycle model simplifies vehicle representation by combining the wheels on each axel into a single wheel. Displacement sensors placed in front and behind the vehicle centre of gravity maintain lateral displacement. The state-space model describing the lateral dynamics of the vehicle with respect to steering angle and road curvature is given in Equation 64 [Ackermann, 2012];

$$\frac{d}{dt} \begin{bmatrix} y_f \\ \dot{y}_f \\ y_r \\ \dot{y}_r \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ a_{21} & a_{22} & -a_{21} & a_{24} \\ 0 & 0 & 0 & 1 \\ a_{41} & a_{42} & -a_{41} & a_{44} \end{bmatrix} \begin{bmatrix} y_f \\ \dot{y}_f \\ y_r \\ \dot{y}_r \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ b_{21} & b_{22} \\ 0 & b_{32} \\ b_{41} & b_{42} \end{bmatrix} \begin{bmatrix} \sigma_f \\ \rho_{ref} \end{bmatrix} \quad (64)$$

Where, σ_f is the steering angle, ρ_{ref} is the reference road curvature, a_{ij} and b_{ij} are vehicle parameters, and y_f and y_r is the lateral displacement between the front and rear sensor to the reference lane, respectively.

The reference road curvature is given by;

$$\rho_{ref} = \frac{\frac{d^2 y}{dx^2}}{\sqrt{1 + (\frac{dy}{dx})^2}} \quad (65)$$

Where, y is the lateral position, and x is the longitudinal position.

The lateral position is given by;

$$y(x) = L_w \times (10(\frac{x}{\Delta x})^3 - 15(\frac{x}{\Delta x})^4 + 6(\frac{x}{\Delta x})^5) \quad (66)$$

Where, L_w is the lane width, Δx is the desired longitudinal distance.

The desired longitudinal distance is given by;

$$d = v \sqrt{\frac{L_w}{a_{max}} \times (60(\frac{x_m}{\Delta x}) - 180(\frac{x_m}{\Delta x})^2 + 120(\frac{x_m}{\Delta x})^3)} \quad (67)$$

Where, v is the velocity of the vehicle, a_{max} is the maximum vehicle acceleration, and x_m is the point of maximum curvature along the vehicle trajectory.

Ho et al., evaluated the control algorithm through simulation and experimentation by developing a miniature autonomous vehicle. The authors demonstrated that simulation provides reasonably close results to physical experimentation. This finding is significant as the authors developed a single control protocol that could be applied in both real-world CAV operation and within simulation.

A coherent control protocol between simulation and implementation renders the development of heuristics, simplifications, and approximations unnecessary for emulating CAV behaviour. While this study is computationally simple and tailored to an environment with limited markings and vehicle connectivity, it assumes vehicles operate individually and do not leverage CAV communications capabilities, especially in cooperative lane-changing environments.

Nie et al., Model

Nie et al., proposed a decentralised framework for lane-changing decision-making that used incentive-based modelling to predict the future decisions and states of vehicles [Nie et al., 2016]. The framework is subdivided into three modules, prediction, decision-making, and decision-coordination. The prediction module uses instantaneous vehicle kinematic information in a cooperative car-following model to predict the following state of all vehicles. The prediction is then forwarded to the decision-making module, which forms a decision that is optimally beneficial for the vehicle. Finally, the decision is

broadcast to the surrounding vehicles via the coordination module, leading to a final coordinated action.

The prediction model is as follows;

$$\begin{aligned}
a_n &= -\frac{1}{c_t}v_n + V\left(\sum_{i=0}^m c_r \Delta x_{n+i}\right) + W\left(\sum_{i_0}^m c_u \Delta v_{n+i}\right) \\
V\left(\sum_{i=0}^m c_r \Delta x_{n+i}\right) &= \frac{1}{c_t}v_{max}\left(\tanh\left(c_s \sum_{i=0}^m c_r \Delta x_{n+i}\right) \right. \\
&\quad \left. + \tanh\left(c_s\left(\sum_{i=0}^m c_r \Delta x_{n+i} - \Delta x_s\right)\right)\right) \\
W\left(\sum_{i_0}^m c_u \Delta v_{n+i}\right) &= c_v \sum_{i=0}^m c_u \Delta v_{n+i}
\end{aligned} \tag{68}$$

Where, a_n and v_n is the acceleration and velocity of vehicle n , Δv_{n+i} and Δx_{n+i} is the relative velocity and relative spacing of the vehicle to its leader, c_r and c_u are interaction coefficients between the vehicle and its leader, c_t is a relaxation term as velocity approaches the maximum limit, c_s is a smoothing coefficient, c_v is the sensitivity to relative velocity, and Δx_s is the safe headway between vehicles.

The decision-making module evaluates whether the utility gained from a lane change is greater than the utility of keeping the current lane. The equation underpinning the module is given by;

$$\begin{aligned}
U(sc, C, T) &= (\bar{a}_{sv} - a_{sv}) + p(\min(\bar{a}_t - a_t \mid t \in N_T)) + \\
&\quad q(\max(\bar{a}_c - a_c \mid n \in N_c)), \quad T \in \{L, R\} \\
N_s &= \{j \in V_s : 0 \leq \|x_{sv} - x_j\| \leq L\}, \quad S = C \text{ or } T \\
TS &= \arg \max_{T \in \{L, R\}} U(sv, C, T) \\
\text{Subject to } U(sv, C, T) &> \Delta a_{th}
\end{aligned} \tag{69}$$

Where, sv , C , and T is the subject vehicle, current lane, and target lane respectively, L and R represents left and right direction target lane, U is the overall advantage of changing lanes, N_C and N_T are the sets of following vehicles in the current and target lane respectively, \bar{a} is the acceleration of the

vehicle should a lane change occur, V_s is the set of all vehicles in the target lane, and a_{th} is the switching threshold.

Nie et al., conducted simulations using this framework on two separate motorway environments, the first containing an onramp, and the second without the onramp. The environments have a flow of $600veh/hr$ across two lanes, well below saturated conditions. Simulations with a duration of $300s$ showed that the proposed method improves platoon stability and prevents the formation of stop-and-go driving conditions from uncooperative and imprudent lane changes. The benefits of homogenising the driving fleet are also seen in improved traffic flow stability. The framework decreases the lane-changing time of the vehicle to half and improves the spatial distribution of lane changes, rather than having them concentrated at a single location.

4.3 Gap-Keeping Models

Gap-acceptance models are generally part of the cooperative lane-changing or car-following models for CAVs. As shown in Section 4.1 and Section 4.2, microsimulation models for CAVs use V2V and V2I as an integral part of behavioural emulation. For this reason, the notion of gap-acceptance becomes somewhat antiquated, as vehicle cooperation forces the creation of appropriate gaps. Instead, this section presents the gap-keeping models designed to maintain inter- and intra-vehicle platoon spacing. Studies of this nature are sparse, as they have been addressed in detail under the car-following and lane-changing sections of the literature review.

Naranjo et al., Model

Naranjo et al., developed an ACC model using fuzzy logic to maintain the gap between a vehicle and its leader [Naranjo et al., 2003]. Using the ACC and fuzzy logic algorithm, the authors generated vehicle reaction information and forwarded it to the vehicle hardware. The control and gap-keeping algorithm were implemented in a mass-produced vehicle to investigate its performance. While the algorithm was not explicitly designed for use in microsimulation, its algorithmic nature makes it readily adaptable.

ACC is restricted in its applicability, showing reliable performance in a narrow

bandwidth of speeds. To address this limitation, the authors used fuzzy logic to reinforce the ACC algorithm. The algorithm starts by using the Center of Mass method, which takes a weighted average of the inputs to generate an output. So;

$$x_{out} = \frac{\sum_i U_i x_i}{\sum_i U_i} \quad (70)$$

Where, x_i is the value of rule i , U_i is the weighting of rule i , and x_{out} is the output variable.

The vehicle speed error and acceleration estimate are passed to the fuzzy controller. Speed error is defined as the difference between the current speed and the desired speed, and acceleration is the change in speed for the lowest time interval possible by the computer clock. The algorithm produces the accelerator pedal pressure needed to reach the appropriate velocity in response to the vehicles current speed and headway. The membership functions define a “null” point and then use ORBIX, a language that allows fuzzy logic rules to be written in natural language. ORBIX allows the membership function to contain three categories, the null point, less than the null point, and more than the null point. This structure results in the following four rules controlling vehicle motion;

- IF *speed error* MORE THAN *null* THEN *accelerate up*
- IF *speed error* LESS THAN *null* THEN *accelerate down*
- IF *acceleration* MORE THAN *null* THEN *accelerate up*
- IF *acceleration* LESS THAN *null* THEN *accelerate down*

The algorithm was implemented in a commercial vehicle on a closed road system with a speed limit of $80km/hr$. The equipped vehicle followed a human driver. During experimentation, the CAV reached a minimum gap of $0.06s$ and a minimum distance of $1.86m$. The speed-time diagram of the CAV mimicked the human vehicle while remaining partially time-shifted. The probabilistic fuzzy-logic means of controlling gap raises questions regarding safety as ACC already shows safety concerns when operating under $30km/hr$. The authors showed acceptable performance using fuzzy-logic in low-velocity environments, but the probabilistic nature of the framework could potentially cause issues in congested stop-start conditions.

Bang and Ahn Model

Bang and Ahn proposed a rule-based platoon movement model based on swarm intelligence [Bang & Ahn, 2017]. Swarm intelligence describes the aggregated movement of individual agents such as the flocking of birds or the schooling of fish. Swarm intelligence models are defined by three key regions, the repulsion zone, the alignment zone, and the attraction zone. The repulsion zone immediately surrounds the agent and precludes other agents from residing within it, acting as a means of maintaining safety. The alignment zone is where vehicles aim to reconcile their kinematic behaviour and follow the unified movement. Within the attraction zone, agents attempt to minimise the spacing between themselves and move closer together, towards the swarm centroid.

Bang and Ahn used the mechanics of a spring and damper system to model the three zones of the swarm intelligence model. The natural attraction and repulsion caused by the extension and contraction of spring systems represent vehicle motion in the attraction and repulsion zone. The damping system models the alignment zone. The specificities of the spring damper car-following system have already been discussed in Section 4.1, and so only the CAV platooning concept using swarm intelligence is provided here.

The reaction of the swarm is akin to the reaction of gas particles placed under pressure. While the spring constant and traffic density are low, the vehicles are spatially spread and occupy the corridor. Increasing the vehicle density is mimicked by increasing the spring constant, resulting in the vehicles travelling closer to one another and providing stronger influence on each other's behaviour. As the flow and density increases, the vehicles continue to form larger and tighter platoons. Bang and Ahn set the spring constant directly proportional to flow, using a maximum, quadratic, or cubic function. This characteristic allows the platoon to adaptively change its behaviour in response to the number of members in the platoon.

To evaluate the framework, the authors implemented it in a motorway section with an onramp. Five different demand scenarios were assessed, ranging from 500veh/hr to 2500veh/hr . As the congestion in the environment increased, the time needed for the vehicles to cluster and form a platoon decreased. From

the 500veh/hr to 2500veh/hr case, the clustering time fell from 572s to 61s . The authors, however, do not report any network metrics to indicate what impact platoon clustering has on network performance.

Swarm intelligence forces vehicle to cluster by attracting them to the cluster centroid. While vehicles further back of the platoon are pulled forwards towards the centroid, the vehicles leading the platoon are pulled back. This behaviour raises two critical concerns. The first, is it necessary to draw back leading vehicles in an attempt to join a platoon? Aggregating vehicles in the network is a benefit as it homogenises vehicles behaviour and reduces stochastic movement. However, the benefit must be weighed against the disbenefit of pulling back lead vehicles. The second concern is that forcing the formation of platoons by proactively altering vehicle kinematics unnecessarily reduces the number and size of gaps in the network, which may inadvertently reduce network performance in high lane-changing scenarios. Most frameworks advocate for the formation of platoons. However, they do so in a natural way without forcing their formation. Without providing network performance metrics, it is difficult to comment on the benefit of using swarm intelligence.

4.4 Autonomous Intersection Models

Many intersections are designed with a cycle time of 120s to 140s , with phase times of as little as 35s . If a vehicle arrives during the commencement of the stop phase, it can experience a delay up to 100s . The use of CAVs, centralised control intersections, and V2I communication is hypothesised to significantly reduce delay time, frequency of stops, and vehicle emissions at intersections ([Kamal et al., 2013], [Qian et al., 2015] and [Makarem et al., 2012]). An unsignaled intersection is generally proposed as operating on a priority basis, with certain vehicles being given right of way through the intersection. The resulting vehicles adjust their trajectories to accommodate for the priority vehicles. The intersection remains as the centralised source of instructions and information, dictating the movement of each vehicle through it.

Many automated intersection designs depend on a 100% fleet penetration of CAVs. While approaches such as this still provide valuable insights into unique CAV control and cooperation protocols, their real-world applications are limited in the short-term. The autonomous intersection control protocols that

focus on mixed fleet operations provide more realistic real-world short-term intersection impacts of CAV behaviour. Especially those frameworks that increase the degree of complexity by including other agents such as cyclists and pedestrians. The literature review in this section provides in brief detail the unique and novel solutions to automated intersection operation.

Makarem et al., Model

Makarem et al., modelled the autonomous intersection using second-order dynamics and gradient-descent optimised navigation functions. The motion of each vehicle is modelled using Equation 71. Using second-order dynamics increases simulation realism taking into account inertial, acceleration, and deceleration forces. The model is as follow;

$$a_i = \frac{1}{m_i} k_i \quad (71)$$

Where, a_i is the vehicle acceleration, k_i is the control input, and m_i is the vehicle mass.

The motion function then feeds into the navigation function, which dictates the direction of the vehicle. The navigation function, shown in Equation 72, is designed to attain a high value as the vehicle veers off track;

$$\phi_i = \lambda_1 (x_{i,goal} - x_i)^2 + \lambda_2 \sum_{i \neq j} \frac{1}{\beta(q_i, q_j)} \quad (72)$$

Where, the first term is the distance of the vehicle from its goal and attains a higher value as the vehicle moves off its optimal path, and the second term is a penalty function that attains a higher value as the vehicle moves into the sensing (safety) zone of other vehicles.

The repulsive force between vehicles (β) is given in Equation 73;

$$\beta_0(q_i, q_j) = \begin{cases} 3(\frac{\|q_i - q_j\|}{\sigma})^2 - 2(\frac{\|q_i - q_j\|}{\sigma})^3 & \text{if } \|q_i - q_j\| < \sigma \\ 1 & \text{else} \end{cases} \quad (73)$$

During each time step, the vehicle traverses its optimum trajectory according to the gradient-descent method. Using the above described underlying framework, Makarem et al., simulated an unsignaled intersection in Aimsun

and Matlab. The key parameters used to assess intersection performance were vehicle average speeds, number of stops, vehicle throughput, fuel consumption, and vehicle emissions. This vehicle control algorithm demonstrated improvements in intersection performance over conventional adaptive signalling schemes, particularly during high flow scenarios. There was no change in average speed, a 69% reduction in the number of stops, a 20% increase in throughput, a 58% reduction in fuel consumption, a 25% reduction in carbon monoxide emissions, and a 50% reduction in nitrous oxide emissions.

Their proposal focused on decentralised control of CAVs through intersections, only feasible at 100% CPR. While testing was conducted on a relatively simple intersection (single lane approach arms), the feasibility and practicalities of a decentralised control approach are questioned when implementing this framework in larger intersections. Consider a common intersection with four lanes per approach arm. The location of vehicles on the x or y plane often overlaps at multi-lane intersection approaches. In this case, the self-contained sensing equipment in CAVs such as RADAR and LiDAR are not capable of appropriately tracking trajectories during vehicle motion. RADAR and LiDAR are currently capable of mapping the unit's immediate surroundings, with limited ability to penetrate surfaces and see behind obstacles. This limitation of the hardware highlights the safety implications of a decentralised approach, where a centralised control unit would not have the same limitations.

Kamal et al., Model

Kamal et al., proposed a similar solution to modelling autonomous intersections [Kamal et al., 2013]. However, they adopted a centralised approach with a control algorithm that maintained the trajectory of eight vehicles within the intersection during each time step. The control algorithm is the cost function shown in Equation 74;

$$\begin{aligned}
 j = & \sum_{t=0}^{T-1} (w_v \sum_{l \in L} \sum_{n=1}^M (v_n^l(t+1) - v_d)^2 + w_u \sum_{l \in L} \sum_{n=1}^M (u_n^l(t))^2 \\
 & + \sum_{p=1}^4 \sum_{l \in L} \sum_{n=1}^M \sum_{k \in L} \sum_{m=1}^M f_{ln,km}^p(t+1))
 \end{aligned} \tag{74}$$

Where, the first cost term reflects a vehicle deviating from their desired path,

the second cost term is concerned with the acceleration of the vehicle, and the third cost term is associated with the risk of collisions.

The risk of collisions is defined by the following risk function;

$$f_{ln,km}^p(t) = H\delta_{ln,km}^p e^{-(a_l^n(x_n^l(t)-z_p^l)^2 + a_k^m(x_m^k(t)-z_p^k)^2)} \quad (75)$$

Where, $f_{ln,km}^p(t)$ is the risk of vehicle ln and vehicle km colliding at point p at time t , H , a_l^n , and a_k^m are calibration constants, and $\delta_{ln,km}^p$ is defined as;

$$\delta_{ln,km}^p = \begin{cases} \theta_{ln}^p \theta_{km}^p & \text{if } k \neq l \\ 0 & \text{otherwise} \end{cases} \quad (76)$$

Where, θ_i^p is a binary value for vehicle i passing through collision point p .

By simultaneously controlling eight vehicles through the intersection, and with a heavy emphasis on comfort and safety, Kamal et al., demonstrated a 106% growth in traffic flow with 0% of the traffic turning left, and a 93% growth in traffic flow with 20% of the traffic turning left. These improvements indicate the potential advantages of implementing CAV technology in a centrally controlled intersection.

Makarem et al., and Kamal et al., both showed that intersection improvements are attained when the focus of the centralised control mechanism is on performance and optimal trajectory, or when the focus is on occupant comfort and environmental implications.

AIM Model

A comprehensive study conducted at the University of Texas [Kockelman et al., 2017] investigated the multifaceted implications of CAVs. Among their studies included a microsimulation-based assessment of CAV penetration using the Autonomous Intersection Management (AIM) model [Dresner & Stone, 2008].

AIM emulates the DSRC capabilities of CAVs through information exchange with an intersection manager. The intersection manager calculates real-time vehicle trajectory and conducts a comparison with a reservation table. The reservation table is responsible for storing past requests made by preceding

vehicles. If the trajectory does not conflict with currently reserved space-time, the intersection manager guarantees the vehicle safe-passage and appends its request to the reservation table. However, if a request is denied, then the vehicle is forced to decelerate and request a later time.

The limitation of this study is that CAV behaviour assessment in microsimulation is limited only to isolated intersections. A comprehensive study should also consider a network of intersections and the motorway environment. This is because isolated intersections risk creating an idealised environment, where true arrival and departure rates and vehicle platoon formation is not appropriately emulated. Additionally, the AIM model does not emulate CAV platooning, considered one of the most advantageous characteristics of CAV operation.

Fayazi and Vahidi Model

Fayazi and Vahidi developed a mixed-integer linear programming (MILP) solution for optimally scheduling CAVs through an unsignaled intersection [Fayazi & Vahidi, 2018]. MILP uses linear subject constraints to optimise the objective function, with some or all solutions to the objective function's inputs required to be integers. The intersection controller receives speed and location information and returns an intersection access time for each vehicle. The vehicles then use localised planning and mapping to adjust their trajectories in relation to their allotted access time.

The objective function is to minimise the access time of the last vehicle (and consequently maximise intersection throughput). The cost function of vehicle access to the intersection is given by;

$$J = w_1(t_{access,j} - t_0) + w_2\left(\sum_{i=1}^n |t_{access,i} - t_{access,des,i}|\right) \quad (77)$$

Where, J is the total cost function, the first term represents the penalty of not receiving a requested access time (delay), $t_{access,j}$ is the expected arrival time, t_0 is the current time, and the second term is the penalty for a vehicle travelling at an undesired speed to meet its access time.

The objective function is subject to three constraints. The speed and acceler-

ation for the vehicle must remain within limits; for this reason, a restriction is placed on the access time of the vehicle into the intersection. A minimum headway must be maintained between vehicles along the same movement. Finally, a safe gap among conflicting vehicles must ensure collision-free operation.

Linear programming requires continuity in both the objective function and constraints. Fayazi and Vahidi address discontinuity by using the “big-M” method. This method uses a dummy binary variable to conditionally switch from one function to another, using the switch to cover discontinuous portions of the domain. Other techniques include subdividing absolute value terms into two constraints (one for positive and negative), or providing additional constraints to limit the search domain prior to the discontinuity.

The MILP solution was compared against two other frameworks. The first framework involved CAVs using image recognition to identify the signal state manually and making a decision. The other framework used V2I to inform CAVs of the signal state from preceding vehicles. The results indicate what the MILP solution results in up to 98.9% fewer vehicle stops, 99.5% less stop delay, and an average travel time decrease of 29%.

It is apparent from the results that moderating vehicle trajectories and arrival times to reduce start-stop delay has beneficial impacts on intersection performance. However, Fayazi and Vahidi did not comment on the effect that this behavioural change had on queue lengths. It is trivial to understand that if the number of stops of a vehicle is reduced or eliminated, the corresponding delay that is derived explicitly from stops will also decrease. Reporting on stop delays, in this case, becomes a disingenuous metric. The authors did not report on the length of queues arising from keeping vehicles in constant motion. The approach arm of the intersection was 500m long, which is typically significantly longer than the distance between adjacent intersections in urban environments. If MILP intersection control results in larger queue lengths, then it can not be considered an appropriate solution, regardless of its delay advantages.

4.5 Summary and Concluding Remarks

Through a detailed review of the literature, a number of models have been proposed for the emulation of CAV behaviour in microsimulation. While many of these contributions have implemented their models in experimental testbeds and demonstrated substantial improvements in intersection or network performance, they currently have a range of limitations that must be addressed.

The data-driven nature of artificial intelligence, machine learning and neural network models makes them difficult and near-impossible to calibrate or validate. They require vast quantities of varied data to serve as training data. This data requirement, unfortunately, limits their use until after CAVs are well integrated into transport systems and the necessary data is widely available. Meaning that they are not appropriate for predictive and planning purposes while CAV technology is in its infancy and the data is unavailable. However, the purely data-reliant nature of these models may make them more appropriate for emulation and planning in the long term.

Validation is an important component of any model development, which requires real-world data to serve as a comparison point. Without real-world data available for validation, many of the proposed models in literature neglect this vital step. They instead use the justification of temporary assumptions, reasonableness, and appropriateness to bypass the validation process and use their proposed models for predictive efforts. This approach still provides substantial value to the literature, as it provides an insight into potential outcomes given that specific assumptions hold. While this style of model development is beneficial for forecasting significantly into the future, CAV technology is under rapid development and approaching network integration. For this reason, a more robust modelling framework is necessary that doesn't contain rigid assumptions, but instead as fluid, interchangeable, and modifiable modules that evolve as more information regarding CAVs becomes available.

Finally, many of the developed models do not undergo stress-testing in congested environments, experimentation in mixed-fleets, or facilitate integration into larger frameworks that address multiple components of vehicle behaviour emulation. For these reasons, benefit would be derived from developing a unique framework that addresses these limitations.

5 Development of a Mixed-Fleet Microsimulation Emulation Framework

This chapter contains two major sections. Section 5.1 develops in detail a new algorithm for emulating CAV behaviour in a microsimulation setting. Section 5.2 then intimately describes the way in which human vehicle behaviour is emulated. Both sections together form the entirety of the emulation framework used for mixed-fleets in microsimulation in the remainder of this thesis.

5.1 CAV Behavioural Emulation Framework

Chapters 3 and 4 provided an extensive review of the variety of models and modelling techniques used to emulate human vehicle and CAV behaviour in microsimulation, respectively. The question may then be asked why a framework for emulating CAV behaviour is not developed using the myriad of models already available? Chapter 3 discusses the inherent assumptions regarding human behaviour incorporated in human vehicle microsimulation models. These assumptions force the vehicle to exhibit behaviour such as perception-reaction lag, compression waves in traffic streams arising from erratic behaviour, uncooperativeness in lane-changing and space negotiation, and preventing the formation of platoons. While new models designed specifically for CAVs address many of these issues, they also have a range of limitations. Limitations including unreasonable amounts of calibration data and time, inflexible and rigid definitions of behaviour, or models that were created in isolation and do not have the capabilities to work as part of a larger framework.

This section outlines the development of a CAV emulation framework. The strength of this framework is in its simplicity, unique ability to scale, and small number of parameters that require calibration. Additionally, the components of this framework are modular; they can be removed, altered, or appended to, when greater information regarding CAV behaviour is available. Finally, the greatest strength of this framework is its provision of a unified and cohesive description of car-following, lane-changing, and gap-acceptance for CAVs. Additionally, this framework contains a trajectory optimisation component for collision-free merging at junctions. The novelty of this framework lies in its co-operative nature by simultaneously assigning the vehicle multiple leaders and followers. The behaviour of the vehicle then subtends to the set of constraints

that is most conservative in terms of safety and cooperation. This process is explained in greater detail throughout Section 5.1.1.

5.1.1 Car-Following Component

The car-following model defines how a vehicle reacts to other agents in its vicinity or an obstruction in its trajectory. This framework maintains safety through altruistic interactions with other vehicles by assigning the ego-vehicle multiple leaders and followers. The primary leader and follower are defined as the direct leader and follower of the ego vehicle in the current lane. Secondary (and tertiary) leaders and followers depend on the future intentions of the vehicle. For a vehicle desiring a lane change, the secondary leader and follower are the vehicles surrounding the target gap in the target lane. This approach simultaneously allows the framework to calculate the necessary acceleration imposed by safety constraints to all vehicles and subtends to the most conservative. Refer to Figure 6 for a diagrammatic representation;




Figure 5: A diagrammatic representation of the dynamically assigned leaders and followers to an ego vehicle.

The remainder of this section contains a detailed explanation of the framework’s car-following component. It initially defines behaviour concerning a trajectory obstacle, such as traffic lights and physical obstructions, and concludes by outlining CAV behaviour relative to other vehicles.

Reacting to an Obstructed Trajectory

This section outlines how a CAV reacts to an obstacle in its trajectory. The proximity of the vehicle to the obstacle is divided into four regions, each governing the behaviour of the vehicle differently based on safety constraints and urgency. The four regions are diagrammatically presented in Figure 6;




Figure 6: The structure of the regions that define CAV behaviour.

The four regions are;

- Region 1: Closest region to the obstruction. In this region, the vehicle decelerates at its maximum rate to come to a complete stop. Under most circumstances, the vehicle will have no acceleration and velocity by the time it reaches Region 1, by design of the framework.
- Region 2: Region in which safety governs the behaviour of the vehicle. The vehicle uses a safety-based deceleration rate calculated using the kinematic equations to ensure appropriate deceleration with respect to its leader. This approach avoids collision with the obstacle and brings the vehicle to a complete stop as it approaches Region 1.
- Region 3: Region in which the obstacle influences the behaviour of the vehicle. The vehicle begins to decelerate while still making arrangements to leave the obstructed lane if a lane change is necessary.
- Region 4: Furthest region from the obstruction in the vehicle trajectory. In this region, the obstruction has no implications on vehicle behaviour.

Region 1

The behaviour of the vehicle in Region 1 is trivial. Let x define the distance between the vehicle and the obstruction. Then, Region 1 is defined by $0 \leq x < x_{stop}$. x_{stop} is the distance by which the vehicle comes to a complete stop before the obstruction. The acceleration enforced by the other regions ensures that the vehicle velocity approaches zero as x approaches x_{stop} . But if this is not the case, then when $x \leq x_{stop}$;

$$a_{n,t+\tau} = d_{max} \quad (78)$$

Where, $a_{n,t+\tau}$ is the acceleration of vehicle n in the next time increment $t + \tau$, and d_{max} is the maximum permitted deceleration.

x_{stop} is a distance set by the modeller. For this work;

$$x_{stop} = 1 \quad (79)$$

Region 2

Region 2 is defined by $x_{stop} \leq x < x_{critical}$, and controls the motion of the vehicle based on safety. The vehicle must come to a complete stop as x approaches x_{stop} . The deceleration is governed by the kinematic equation of motion;

$$a_{n,t+\tau} = \frac{-v_{n,t}^2}{2 \times (x_{n,t} - x_{stop})} \quad (80)$$

Where, $a_{n,t+\tau}$ is the acceleration of vehicle n in the next time increment $t + \tau$, $v_{n,t}$ is the current vehicle velocity, and $x_{n,t}$ is the distance between the vehicle front bumper and the obstacle.

$x_{critical}$ is determined by placing a limit on the rate of change of acceleration ($\delta a / \delta t = (a_{i,t+\tau} - a_{i,t}) / \tau$), denoted as the jerk (J_{max}). The jerk often is a reflection of passenger comfort. Acceleration and velocity as a function of jerk are given by;

$$a = \int J_{max} dt = J_{max} t$$

$$v = \int J_{max} t dt = \frac{1}{2} J_{max} t^2$$

At the desired velocity v_{max} , and ensuring the vehicle reaches standstill as x approaches x_{stop} , the time over which this deceleration can occur is given by $t = (x_{critical} - x_{stop}) / v_{max}$. Using $v = v_{max}$ and $t = (x_{critical} - x_{stop}) / v_{max}$, we can derive that $v_{max} = 0.5 J_{max} ((x_{critical} - x_{stop}) / v_{max})^2$. Therefore;

$$x_{critical} = \left(\frac{2v_{max}^3}{J_{max}} \right)^{1/2} + x_{stop} \quad (81)$$

Region 3

In Region 3, the vehicle is aware of and influenced by the obstacle. A parabolic profile models the rate of change of this acceleration profile. With limited information regarding manufacturer specifications for CAVs, the selection of

the acceleration profile is somewhat arbitrary, so long as the chosen profile is reasonable. No one profile is more valid than another while available information regarding vehicle operations is sparse. The parabolic profile is chosen for this work for several reasons. Firstly, acceleration approaches $0m/s^2$ as x approaches $x_{critical}$ and $x_{approach}$, providing the profile with a smooth transition into adjacent regions. Additionally, the rate of change of a parabolic profile (in this case representing Jerk) can be mathematically restricted to a maximum value (J_{max}) due to the accessible nature of its mathematical representation and derivatives.

The symmetric nature of a parabola with negative concavity results in two occurrences congruent with traditional driving characteristics. The first is that acceleration approaches a minimum value as the vehicle spacing approaches a high enough value to enter Region 4 where proximity to the lead vehicle is irrelevant, or low enough to reach Region 2 where safety parameters and factors govern acceleration. Reaching a minimum acceleration value allows the vehicle to smoothly transition into the adjacent Region. Secondly, the parabola by definition has the greatest rate of change (jerk) at its extreme points and with the interface between adjacent regions. The interface points coincide with real vehicles embarking from constant velocity motion such as cruising speed and reacting to a lead vehicles motion, where jerk is often highest.

The acceleration profile of Region 3 is designed to limit a kinematic response. Acceleration response is confined to $P_{min} \times a_{max} \leq a_{n,t+\tau} \leq P_{max} \times a_{max}$. P_{max} and P_{min} are parameters designed to curtail the domain of acceptable acceleration values for regular vehicle operation. P_{min} prevents a CAV from continuously reacting, accelerating, and decelerating, making its behaviour more human-like. The decision to not make a CAVs reaction continuous, but rather stimuli and need dependant, is appropriate considering the computational and mechanical demand that continuous and minute changes in behaviour would have on the vehicle. This condition resembles many of the thresholds used in psycho-physical car-following models for emulating human behaviour. P_{max} reserves the maximum acceleration of the vehicle for emergencies rather than regular operation. Refer to Figure 7 for a diagrammatic representation of the P_{max} and P_{min} parameters;




Figure 7: Graphical representation of P_{min} and P_{max} .

The acceleration in Region 3 takes the form;

$$a_{n,t} = -\phi(x_{n,t} - x_{critical})(x_{n,t} - x_{approach}) + P_{min} \times a_{max}$$

Where, ϕ will be calculated momentarily. The maximum value of $a_{n,t} = P_{max} \times a_{max}$ occurs at $x_{n,t} = (x_{critical} + x_{approach})/2$. Therefore;

$$\begin{aligned} P_{max} \times a_{max} &= -\phi\left(\frac{1}{2}(x_{critical} + x_{approach}) - x_{critical}\right) \\ &\quad \times \left(\frac{1}{2}(x_{critical} + x_{approach}) - x_{approach}\right) + P_{min} \times a_{max} \end{aligned}$$

$$P_{max} a_{max} = -\frac{1}{4}\phi(x_{approach} - x_{critical})(-x_{approach} + x_{critical}) + P_{min} a_{max}$$

$$P_{max} a_{max} - P_{min} a_{max} = \frac{1}{4}\phi(x_{approach} - x_{critical})^2$$

$$\phi = 4a_{max}(P_{max} - P_{min})(x_{approach} - x_{critical})^{-2}$$

So;

$$\begin{aligned} a_{n,t+\tau} &= -\frac{4a_{max}(P_{max} - P_{min})}{(x_{approach} - x_{critical})^2}(x_{n,t} - x_{critical})(x_{n,t} - x_{approach}) \\ &\quad + P_{min} a_{max} \end{aligned} \quad (82)$$

To determine the value of $x_{approach}$, jerk is used to limit the gradient of the acceleration profile. So;

$$\frac{d}{dx}a_{n,t} \leq J_{max}$$

$$J_{max} \geq \frac{d}{dx} \left(-\frac{4a_{max}(P_{max} - P_{min})}{(x_{approach} - x_{critical})^2} (x^2 - x(x_{critical} + x_{approach}) + (x_{critical}x_{approach})) + P_{min}a_{max} \right)$$

$$J_{max} \geq -\frac{4a_{max}(P_{max}-P_{min})}{(x_{approach}-x_{critical})^2} (2x - (x_{critical} + x_{approach}))$$

In the most extreme scenario, $\frac{d}{dx}a_{i,t} = J_{max}$. Therefore to account for the most extreme scenario;

$$J_{max} = -\frac{4a_{max}(P_{max}-P_{min})}{(x_{approach}-x_{critical})^2} (2x - (x_{critical} + x_{approach}))$$

$$-4a_{max}(P_{max}-P_{min})(2x-x_{critical}-x_{approach}) = J_{max}(x_{approach}-x_{critical})^2$$

$$\begin{aligned} 4a_{max}(P_{max} - P_{min})x_{approach} - 4a_{max}(P_{max} - P_{min})(2x - x_{critical}) \\ = J_{max}x_{approach}^2 - 2J_{max}x_{approach}x_{critical} + J_{max}x_{critical}^2 \end{aligned}$$

$$\begin{aligned} 0 = (J_{max})x_{approach}^2 + (-2J_{max}x_{critical} - 4a_{max}(P_{max} - P_{min}))x_{approach} \\ + (J_{max}x_{critical}^2 + 4a_{max}(P_{max} - P_{min})(2x - x_{critical})) \end{aligned}$$

For convenience, let $C_1 = J_{max}x_{critical}$, $C_2 = 4a_{max}(P_{max} - P_{min})$, and $C_3 = 2C_1 + C_2$. Then;

$$0 = (J_{max})x_{approach}^2 - C_3x_{approach} + (C_1x_{critical} + C_2(2x - x_{critical}))$$

$$x_{approach} = (C_3 + \sqrt{C_3^2 - 4J_{max}(C_1x_{critical} + C_2(2x - x_{critical}))})/2J_{max}$$

$x_{approach}$ is inversely proportional to x , so the maximum value of $x_{approach}$ occurs at the minimum value of x ($x = x_{critical}$). Then;

$$\begin{aligned} x_{approach} &= \frac{C_3 + \sqrt{C_3^2 - 4J_{max}(C_1x_{critical} + C_2(2x_{critical} - x_{critical}))}}{2J_{max}} \\ &= \frac{2C_1 + C_2 + \sqrt{4C_1^2 + 4C_1C_2 + C_2^2 - 4J_{max}x_{critical}(C_1 + C_2)}}{2J_{max}} \end{aligned}$$

$$\begin{aligned}
&= \frac{2C_1 + C_2 + \sqrt{4C_1^2 + 4C_1C_2 + C_2^2 - 4C_1(C_1 + C_2)}}{2J_{max}} \\
&= \frac{2C_1 + C_2 + \sqrt{4C_1^2 + 4C_1C_2 + C_2^2 - 4C_1^2 - 4C_1C_2}}{2J_{max}} \\
&= \frac{2C_1 + C_2 + \sqrt{C_2^2}}{2J_{max}} \\
&= \frac{C_1 + C_2}{J_{max}}
\end{aligned}$$

So;


$$x_{approach} = \frac{J_{max}x_{critical} + 4a_{max}(P_{max} - P_{min})}{J_{max}} \quad (83)$$

Region 4


The obstacle is of no interest to the vehicle, and so a reaction is not warranted. The vehicle can continue to behave with its current kinematics or the kinematics imposed by another condition in this framework.

Summary


In summary, CAVs following this framework will exhibit an acceleration, velocity, and displacement profile as pictured in Figure 8. Figure 8 does not provide a probability distribution for vehicle kinematics, but provides a representative profile for how CAV kinematics change as the vehicles proximity to vehicles and obstacles changes. A precise profile is derived when the framework parameters are set.



(a) Typical CAV acceleration profile.



(b) Typical CAV velocity profile.



(c) Typical CAV displacement profile.

Figure 8: Typical acceleration, velocity, and displacement profiles exhibited by CAVs as they approach an obstruction in the vehicles' trajectory.

For $x_{n,t} < x_{stop}$,

$$a_{n,t+\tau} = a_{max}$$

Where, x_{stop} is manually defined.

For $x_{stop} \leq x_{n,t} < x_{critical}$,

$$a_{n,t+\tau} = \frac{-v_{n,t}^2}{2 \times (x_{n,t} - x_{stop})}$$

Where, $x_{critical} = (\frac{2v_{max}^3}{J_{max}})^{1/2} + x_{stop}$

For $x_{critical} \leq x_{n,t} < x_{approach}$,

$$a_{n,t+\tau} = -\frac{4a_{max}(P_{max}-P_{min})}{(x_{approach}-x_{critical})^2}(x_{n,t}-x_{critical})(x_{n,t}-x_{approach}) + P_{min}a_{max}$$

Where, $x_{approach} = \frac{J_{max}x_{critical} + 4a_{max}(P_{max} - P_{min})}{J_{max}}$

For $x_{n,t} \geq x_{approach}$,

$$a_{n,t+\tau} = 0$$

General Car-Following Model

The general car-following model, adapted from Van Arem et al., [Van Arem et al., 2006], describes vehicle behaviour in other circumstances. Other studies [Talebpour & Mahmassani, 2016] have also used this model, and is as follows;

$$a_{n,t+\tau} = c_a a_{n-1,t} + c_v(v_{n-1,t} - v_{n,t}) + c_x(\Delta x_{n,t} - \Delta x_{min}) \quad (84)$$

Where; c_a , c_v and c_x are unitless calibration parameters that govern the influence of the lead vehicle on the ego vehicle. When increasing these parameters, the ego vehicle will have a greater and more extreme reaction to the changing kinematics of the lead vehicle. c_a , c_v and c_x are set to values of 1.0, 0.58 and 0.1 respectively, refer to [Van Arem et al., 2006] for a detailed explanation of the parameters. $\Delta x_{n,t}$ and Δx_{min} is the current spacing between the lead and ego vehicle, and the minimum permitted spacing, respectively.

The remainder of this section explains how the car-following framework directs vehicle kinematics, based on the vehicle's current circumstances and position in the network.

Vehicle as Platoon leader

The leader of the platoon has no preceding vehicle, rendering $a_{n-1,t}$ and $\Delta x_{n,t}$ mute. Using Equation 84, acceleration for this vehicle is calculated as the factored difference between the desired and current velocity, given as;

$$a_{n,t+\tau} = c_v(v_{desired,t} - v_{n,t}) \quad (85)$$

Vehicle as Platoon Follower, without Lane Changing

A platoon follower faces a range of situations that may require an acceleration adjustment. These include, speed surpassing the maximum limit, the headway

falling below the minimum value, or actively reacting to the behaviour of the lead vehicle and maintaining safe operation. The CAV's response to these circumstances is outlined below.

Velocity Constraint: If the vehicle velocity is outside the desired range, the velocity component of the general car-following model is applied, similarly to the application of the model to the platoon leader, given as;

$$a_{n,t+\tau} = c_v(v_{desired,t} - v_{n,t}) \quad (86)$$

Headway constraint: If the headway between the vehicle and its leader falls below the minimum permitted headway, the maximum deceleration is used to establish safe following distances. Alternatively, if the headway increases to greater than the minimum headway, then the headway component of the general car-following model is applied as follows;

$$a_{n,t+\tau} = c_a a_{n-1,t} + c_x(\Delta x_{n,t} - \Delta x_{min}) \quad (87)$$

Safety Constraint: The velocity and headway constraints alone are limited in that unsafe situations can arise even when neither restriction has been triggered. Consider the scenario where the ego vehicle is travelling slightly below the speed limit, and the lead vehicle has stopped at a traffic light downstream, with a headway marginally higher than the minimum value. In this case, both constraints would recommend that the vehicle accelerate when, in reality, a strong deceleration is required. For this reason, a safety constraint is responsible for preventing collisions, given as;

$$a_{n,t+\tau} = c_a a_{n-1,t} + \frac{v_{n-1,t}^2 - v_{n,t}^2}{2(\Delta x_{n,t} - \Delta x_{min})} \quad (88)$$

This safety condition is triggered when $v_{n,t}$ surpasses specific thresholds. This prevents the constraint from being unnecessarily implemented such as when a lead vehicle is significantly far from the ego vehicle. This constraint is triggered when $v_{n,t} \geq v_{desired} \times P$, where P is a linear interpolation between P_{max} and 1 when $x_{critical} \leq x < x_{approach}$, and P is a linear interpolation between P_{min} and P_{max} when $x_{stop} \leq x < x_{critical}$. In summary;

$$P = \begin{cases} \frac{x - x_{crit}}{x_{app} - x_{crit}}(1 - P_{max}) + P_{max} & \text{for } x_{crit} \leq x < x_{app} \\ \frac{x - x_{stop}}{x_{crit} - x_{stop}}(P_{max} - P_{min}) + P_{min} & \text{for } x_{stop} \leq x < x_{crit} \end{cases} \quad (89)$$

Vehicle as Platoon Follower, with Lane Changing

If the platoon follower intends to change lanes or facilitate the lane change of another vehicle, the acceleration criteria are adjusted to accommodate the stringent safety conditions of the lane change. This section presents the mathematical models to describe car-following with lane-changing. A derivation and description of the mathematical model is provided in Section 5.1.2, which discusses the gap-acceptance portion of the framework. The adjusted velocity, headway and safety constraint are outlined below, with an additional constraint that addresses behaviour as the vehicle approaches the end of its link segment while being situated in the incorrect lane.

Velocity Constraint: The velocity condition continues to regulate vehicle behaviour based on current and desired speed, but has the additional role of satisfying the spatial safety constraint of a lane change. The additional component forces the creation of a gap deemed appropriate for a lane change (refer to Section 5.1.2 for the gap-acceptance criteria).

$$a_{n,t+\tau} = c_v \times P \times \frac{v_{n-1,t}}{1 + \frac{1}{2} \frac{n}{m}} - v_{n,t} \quad (90)$$

Where;

$$P = \frac{x - x_{critical}}{x_{approach} - x_{critical}} \text{ for } P \in [0, 1] \quad (91)$$

n and m are fixed parameters explained in Section 5.1.2. n is the number of time increments reserved to act as a spatial buffer between vehicles, and m is the duration over which the lane-changing process occurs. The factor P moderates the importance of the velocity constraint on lane-changing. It is contingent on the spacing between the ego vehicle and its leader and follower in the target lane. If the leader and follower are far ($x \geq x_{approach}$), then their velocity has no bearing on the lane-changing decision, and an acceleration adjustment is not necessary, meaning that $P = 0$. If the vehicles are close ($x \leq x_{critical}$), then the velocity of the vehicles must be considered at all times and $P = 1$. For $x_{critical} < x < x_{approach}$, P is calculated using Equation 91. In summary;

$$P = \begin{cases} 1 & \text{for } x \leq x_{critical} \\ \frac{x - x_{critical}}{x_{approach} - x_{critical}} & \text{for } x_{critical} < x < x_{approach} \\ 0 & \text{for } x \geq x_{approach} \end{cases} \quad (92)$$

Headway Constraint: The headway condition now also has an added component designed to meet the safety constraints of the lane change.

$$a_{n,t+\tau} = c_x(\Delta x_{n,t} - \Delta x_{min} - \frac{1}{2}n \times v_{n-1,t}) \quad (93)$$

Safety Constraint: The safety condition is calculated here the same way as in the platoon follower case without lane-changing.

End of Link Constraint: When a vehicle requires a lane change, the vehicle must be mindful of the approaching end of the link. This constraint only applies if the vehicle is not in its target lane. The calculation of acceleration in response to a trajectory obstacle can also be used to calculate the vehicle response to an ending link segment. This calculation was explained in Section 5.1.1. Conditionally triggering this constraint reduces unnecessary deceleration in retaliation to a lane drop. The conditional trigger is as follows, let x be the distance between the vehicle and the ending link segment;

For $x > x_{approach}$, the proximity to the link end is of no consequence to the vehicle's behaviour, which continues to search for an appropriate target gap.

For $x_{critical} < x \leq x_{approach}$, if $v_{n,t} > v_{max} \times \frac{x - x_{critical}}{x_{approach} - x_{critical}}(1 - P_{max}) + P_{max}$, then acceleration is dictated by the logic outlined in Section 5.1.1. Else, the proximity of the lane drop does not affect the behaviour of the vehicle.

For $x_{stop} < x \leq x_{critical}$, if $v_{n,t} > v_{max} \times \frac{x - x_{stop}}{x_{critical} - x_{stop}}(P_{max} - P_{min}) + P_{min}$, then acceleration is dictated by the logic outlined in Section 5.1.1. Else, the proximity of the lane drop does not affect the behaviour of the vehicle.

5.1.2 Gap-Acceptance Component

Gap-acceptance models evaluate the appropriateness and acceptability of target gaps during lane-changing. In this framework, an acceptable gap is determined by a spatial headway requirement and a speed requirement. Figure 9 shows the arrangement of the ego, lead, and following vehicles. When the ego

vehicle targets the gap between the lead and following vehicle in the adjacent lane, the following notation is used;




Figure 9: Arrangement and notation used to describe the system of vehicles in a lane change.

The remainder of this section outlines the mathematical formulations for evaluating the acceptability of a target gap.

Headway and Velocity Constraints

The headway condition states that the size of the target gap must be sufficiently large to accommodate the lane-changing vehicle, the minimum headway, and a speed dependant safety factor designed to act as a buffer between vehicles during a lane change. This condition is mathematically given by;

$$gap_{min,n,t} = L_n + 2x_{min} + n \times v_{n,t} \quad (94)$$

Where, $gap_{min,n,t}$ is the minimum distance required to accept a gap, L_n is the length of the lane-changing vehicle and $n \times v_{n,t}$ is the additional safety factor.

This safety factor is a distance calculated by multiplying the current travel velocity ($v_{n,t}$) by a time n , meaning that higher travel speeds demand a greater buffer distance. n is the time for which the safety buffer is provided and should be calibrated using a safety assessment where the likelihood of vehicle conflict during lane-changing is reduced to an acceptable level.

The velocity of the leader and follower in the target lane must be such that over the life of the lane change, their behaviour does not violate the minimum safety distance incorporated in the gap acceptance criteria ($n \times v_{n,t}$). This restriction implies that both the target lane leader and follower may close the gap to the ego vehicle by a distance of $\frac{1}{2}n \times v_{n,t}$.

A gap of $\frac{1}{2}n \times v_{n,t}$ metres is permitted to close at a rate of $v_{n\pm 1,t} - v_{n,t}$ metres per second, over a time period of m seconds, where m is the duration of the lane change manoeuvre. The velocity requirement of the leader and follower in the target lane can then be calculated relative to the ego vehicle velocity. $(v_{n\pm 1,t} - v_{n,t}) \times m > \frac{1}{2}n \times v_{n,t}$. Then;

$$v_{n+1,t} > v_{n,t} \left(1 + \frac{1}{2} \frac{n}{m}\right) \quad (95)$$

Similarly;

$$v_{n-1,t} < v_{n,t} \left(1 - \frac{1}{2} \frac{n}{m}\right) \quad (96)$$

The velocity constraint becomes less relevant as the distance between the ego vehicle and the lead or following vehicle in the target lane increases. The weighting factor presented in Equation 92, in Section 5.1.1, can again be used to reduce the impact of the velocity constraint;

$$P = \begin{cases} 1 & \text{for } x \leq x_{critical} \\ \frac{x - x_{critical}}{x_{approach} - x_{critical}} & \text{for } x_{critical} < x < x_{approach} \\ 0 & \text{for } x \geq x_{approach} \end{cases} \quad (97)$$

Where, x is the distance between the front bumper of the ego vehicle to the rear bumper of the lead vehicle, or the rear bumper of the ego vehicle to the front bumper of the following vehicle.

Evaluating Multiple Gaps

V2V and V2I infrastructure are further leveraged to allow vehicles to intelligently and actively select a gap in the target lane. Rather than the ego vehicle targetting the immediate adjacent gap, the ego vehicle actively selects a gap that reduces its impact on the target lane. Impact on the target lane is defined as the level of deceleration that arises in the target lane as a result of the ego vehicle changing lanes.

The ego vehicle evaluates and compares the different downstream gaps by assigning a score to each gap, using impact on acceleration in the target lane as a proxy for the score. The minimum score is 0, for gaps that fail the gap acceptance criteria outlined in Section 5.1.2. A gap becomes more favourable as the impact reduces, with a higher score indicating a more favourable gap.

If all gaps have a score of 0, then the adjacent gap is selected. A proxy for the score for each gap in the target lane is determined as follows;

$$\begin{aligned} score \approx a_{n-1,t} &= \frac{v_{n,t}^2 - v_{n-1,t}^2}{2 \times (gap_{n,t} - gap_{min})} \\ &\approx \frac{v_{n,t}^2 - v_{n-1,t}^2}{gap_{n,t} - gap_{min}} \\ &\approx \frac{\Delta Speed}{\Delta Gap} \end{aligned}$$

So;

$$score = (\Delta Speed)^{c_1} (\Delta Gap)^{c_2} \quad (98)$$

Where, $\Delta Speed$ is the difference in the travel speed between the ego vehicle and the following vehicle ($v_{n,t} - v_{n-1,t}$), ΔGap is the difference between the minimum required spacing and the current spacing between the ego and following vehicle ($x - x_{min}$), and c_1 and c_2 are weighting factors that alter the significance of the two components of the score with respect to one another.

The gap with the lowest score is selected as the target gap. Figure 10 displays the relationship between the score, $\Delta Speed$ and ΔGap ;

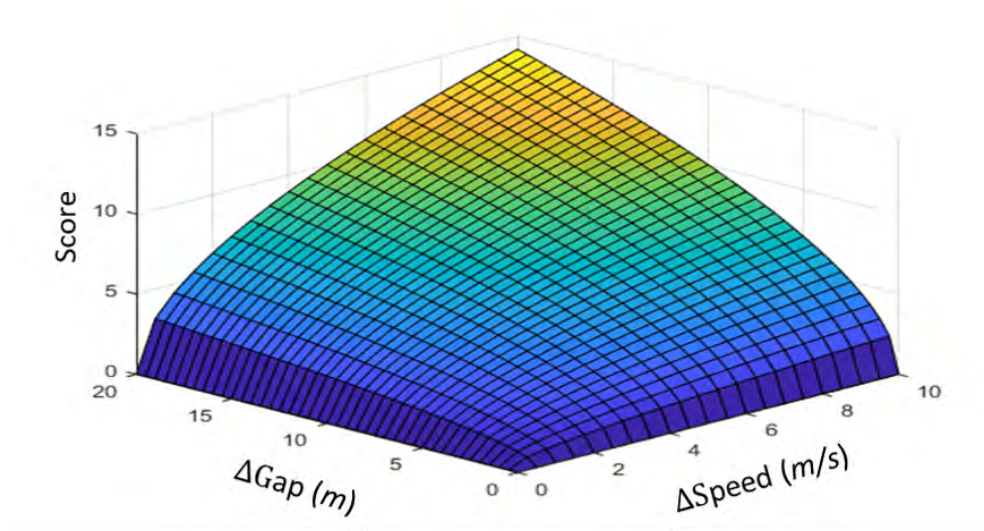


Figure 10: Visual representation of the score as a function of $\Delta Speed$ and ΔGap , using $c_1 = 0.5$ and $c_2 = 0.5$.

5.1.3 Lane-Changing Component

Lane-changing is the act altering wheel alignments to shift into the adjacent lane. This framework does not suggest or control the mechanics of the vehicle wheel alignment. This framework also does not indicate when an MLC is necessary, the microsimulator controls both actions. The microsimulator identifies the need for an MLC and notifies this CAV behavioural control framework, which then moderates vehicle kinematics until the vehicle is ready for the lane change. Control is then reverted to the microsimulator, which sets the appropriate wheel alignments. Once the lane change is completed, control of the vehicle operation is once again returned to this framework.

This framework does, however, trigger a DLC based on a novel criterion that evaluates the impact on the target lane. The DLC may be conducted when the travel speed of the adjacent lane is sufficiently better than the current vehicle travel speed. This section outlines in greater detail, the circumstances that trigger a DLC.

Discretionary Lane Change Evaluation

For a DLC to be deemed acceptable, it must increase the overall utility of the localised system. The localised system consists of the lane-changing vehicle and the following vehicle in the target lane. The utility in the final state of the system must be greater than the utility of the initial state.

An ego vehicle travelling with an initial velocity of $v_{n,t}$ and a desired velocity of $v_{desired}$, has an initial dis-utility of $d_{n,t} = -v_{desired} + v_{n,t}$. This dis-utility is the difference between the vehicles current and desired travel speed. Similarly, the following vehicle in the target lane has a dis-utility of $d_{n-1,t} = -v_{desired} + v_{n-1,t}$. Therefore, the total initial dis-utility of this localised system is;

$$d_{(n \& n-1),t} = v_{n,t} + v_{n-1,t} - 2v_{desired} \quad (99)$$

Where, $d_{(n \& n-1),t}$ is the total initial dis-utility of the system containing the ego vehicle n and the following vehicle $n - 1$ in the target lane.

The final utility for the localised system is again the summation of the final utility of the ego and following vehicle, both of which are derived below.

The expected velocity of the ego vehicle is the average of its maximum and minimum, both of which were calculated in Section 5.1.2. From that section;

$$v_{n+1,t} > v_{n,t}(1 + \frac{1}{2} \frac{n}{m}) \longrightarrow v_{n,t} < v_{n+1,t}/(1 + \frac{1}{2} \frac{n}{m})$$

$$v_{n-1,t} < v_{n,t}(1 - \frac{1}{2} \frac{n}{m}) \longrightarrow v_{n,t} > v_{n-1,t}/(1 - \frac{1}{2} \frac{n}{m})$$

Then;

$$\begin{aligned} E[v_{i,t+m}] &= \frac{1}{2} \left(\frac{v_{i+1,n}}{1 + \frac{1}{2} \frac{n}{m}} + \frac{v_{i-1,n}}{1 - \frac{1}{2} \frac{n}{m}} \right) \\ &= \frac{v_{i+1,t}}{2 + \frac{n}{m}} + \frac{v_{i-1,t}}{2 - \frac{n}{m}} \\ &= \frac{mv_{i+1,t}}{2m + n} + \frac{mv_{i-1,t}}{2m - n} \\ &= \frac{v_{i+1,t}m(2m - n) + v_{i-1,t}m(2m + n)}{4m^2 - n^2} \end{aligned}$$

The final dis-utility of the ego vehicle is calculated as;

$$d_{n,t+m} = -v_{desired} + \frac{v_{n+1,t}m(2m-n) + v_{n-1,t}m(2m+n)}{4m^2 - n^2}$$

For the following vehicle, the deceleration caused by the lane-changing vehicle was presented in Section 5.1.1 and is calculated as $a_{n-1,t} = c_x(\Delta x - \Delta x_{min})$. Using the kinematic equation;

$$a_{n-1,t} = \frac{v_{n,t}^2 - v_{n-1,t}^2}{2\Delta x}$$

$$c_x(\Delta x - \Delta x_{min}) = \frac{v_{n,t}^2 - v_{n-1,t}^2}{2(\Delta x - \Delta x_{min})}$$

$$v_{n,t}^2 - v_{n-1,t}^2 = 2c_x(\Delta x - \Delta x_{min})^2$$

$$v_{x,t} = \sqrt{v_{x-1,t}^2 + 2c_x(\Delta x - \Delta x_{min})^2}$$

The final dis-utility for the following vehicle is then;

$$d_{n-1,t+m} = -v_{desired} + \sqrt{v_{n-1,t}^2 + 2c_x(\Delta x - \Delta x_{min})^2}$$

The final dis-utility for the system of two vehicles becomes;

$$d_{(n \ \& \ n+1),t+m} = \frac{v_{n+1,t}m(2m-n) + v_{n-1,t}m(2m+n)}{4m^2 - n^2} + \sqrt{v_{n-1,t}^2 + 2c_x(\Delta x - \Delta x_{min})^2 - 2v_{desired}} \quad (100)$$

For the DLC to be acceptable, the final utility must be greater than the initial utility. So;

$$\begin{aligned} d_{(n \ \& \ n+1),t} &< d_{(n \ \& \ n+1),t+m} \\ v_{n,t} + v_{n-1,t} - 2v_{desired} &< \frac{v_{n+1,t}m(2m-n) + v_{n-1,t}m(2m+n)}{4m^2 - n^2} \\ &+ \sqrt{v_{n-1,t}^2 + 2c_x(\Delta x - \Delta x_{min})^2 - 2v_{desired}} \\ v_{n,t} &< -v_{n-1,t} + \frac{v_{n+1,t}m(2m-n) + v_{n-1,t}m(2m+n)}{4m^2 - n^2} + \sqrt{(v_{n-1,t}^2 + 2c_x(\Delta x - \Delta x_{min})^2)} \end{aligned}$$

So, for a DLC to be accepted;

$$v_{n,t} < f(v_{n-1,t}, v_{n+1,t}, \Delta x) \quad (101)$$

Where;

$$\begin{aligned} f(v_{n-1,t}, v_{n+1,t}, \Delta x) &= -v_{n-1,t} \\ &+ \frac{v_{n+1,t}m(2m-n) + v_{n-1,t}m(2m+n)}{4m^2 - n^2} \\ &+ \sqrt{v_{n-1,t}^2 + 2c_x(\Delta x - \Delta x_{min})^2} \end{aligned} \quad (102)$$

The velocity of the ego vehicle must fall below a threshold governed by the velocity of the lead and following vehicle in the target lane, and the gap between them, before a DLC is considered a betterment for the system.

Caveats on Triggering the Discretionary Lane Change

In addition to the criteria outlined above, other conditions must also be met before the DLC is conducted. The additional measures are introduced to implement this framework in microsimulation, where rules must be explicitly defined to avoid unintended vehicle behaviour. The additional caveats are demonstrated in Figure 11 and explained below;




Figure 11: Visual representation of the caveats on discretionary lane changing.

The additional caveats are as follows;

- The preferred lane must not be the current lane. The preferred lane is provided by Vissim when the vehicle has an upcoming MLC. The purpose of this caveat prevents the vehicle from travelling further away from the target lane when there is an upcoming MLC.
- The distance to an obstacle in the vehicle trajectory is greater than a certain distance away ($x_{dlc,min}$). This caveat prevents a DLC that may be deemed unnecessary or irrational. For example, if vehicles are waiting at a traffic light and the adjacent lane has negligible fewer vehicles, then a DLC to gain negligible space should not be seen as a reasonable decision. For this thesis, $x_{dlc,min} = 10m$.
- The final caveat is on the velocity of three vehicles involved in the DLC system. $v_{n,t} \leq P_{min} \times v_{desired}$, $v_{n+1,t} \geq P_{max} \times v_{desired}$ and $v_{n-1,t} \geq P_{max} \times v_{desired}$. These thresholds necessitate a reasonable requirement for a lane change. Minor differences in travel speed between two adjacent lanes, or a slightly lower travel speed than the desired speed, should not be justification enough to trigger a DLC.

5.1.4 Cooperative Merging Component

The final component of this framework is to coordinate vehicles through merge points cooperatively. This capability also leverages the offerings of V2V and V2I infrastructure. Traditionally at a merge point, one-of-two scenarios occur;

- Auxiliary Lane: The merging lane will exist as an auxiliary lane that may either be sustained or taper downstream. This arrangement gives vehicles adequate time to merge into the adjacent lane and also acts as a source of storage and capacity. The drawback of this approach is that additional resources must be provided for the provision of an auxiliary lane.

- **Priority Servicing:** Traffic on the merge lane will be forced to yield to the mainline, causing a queue in the merge lane to build if the arrival rate is greater than the service rate. While this approach does not require additional infrastructure, its primary drawback is that if the mainline is approaching saturation, then the merging vehicles may be forced to yield indefinitely.

By forecasting trajectories, vehicle behaviour upstream can be adjusted to avoid conflict at a merge point downstream. This approach has the benefits of not requiring additional infrastructure for an auxiliary lane, nor does it force vehicles to stop and yield completely.

The merge zone precedes the merge point, within which, vehicles have their trajectory adjusted. The modeller defines the length of the merge zone. The larger the merge zone, the greater the time that a vehicle will have to adjust its trajectory. If a conflict is identified at the last moment due to a short merge zone, the conflicting vehicles will experience high deceleration rates or will collide with one another at the merge point. Refer to Figure 12 for a diagrammatic representation of the merge zone, merge point and the zipper merge network geometry in which this method for merging would be appropriate.




Figure 12: Arrangement of the merge point and zone, and the network geometry in which this proposed control framework would provide benefit.

Vehicle requesting passage through the merge point are placed on a reservation table. The objective function to optimise the reservation table is to minimise the maximum delay (refer to Equation 103). This objective function is chosen as it provides the highest equity to the system, ensuring that neither the mainline nor the merging lane has sufficiently more priority. By minimising the maximum delay, the system does not operate optimally from a total system

delay perspective. However, it does provide the most politically acceptable solution, where all vehicles are treated equally, and optimisation is performed at the local conflict level. The objective function is given as;

$$objective = \sum_{t=1}^T \sum_{j=1}^n \min_{i=1}^2 \left(\frac{d_{i|j,t}}{v_{adjusted}} - \frac{d_{i|j,t}}{v_{i|j,t}} \right) \quad (103)$$

Where, t is the simulation time step for all time steps $t \in [1, T]$, j is the specific conflict pair for all conflict pairs $j \in [1, n]$, i is the ID for the vehicle in conflict j for $i \in [1, 2]$, d is the merge point distance, and v is the vehicle velocity.

Once the vehicles enter the merge zone, their kinematic information is recorded to the reservation table. The reservation table is optimised each time increment. Optimisation involves the identification of conflicts, which are resolved by minimising the maximum delay. The optimised reservation table returns velocities to the vehicles, which are implemented in the following time step. This process is divided into two components, where the first component creates and maintains the reservation table using C++. The pseudocode for this component is provided in Algorithm 1;

Algorithm 1: Pseudocode cooperative merge algorithm; reservation table component

```

import:  $S = \{Reservation\ Table\}$ 
initialise:  $S' = 0, \hat{n} = \{n\}$ 
while  $\hat{n} \neq 0$  do
    select a vehicle  $n$  from  $\hat{n}$  to update
    if  $i \notin S$  then
        | Create new entry  $n$  and add to  $S'$ 
    end
    else if  $n \in S$  then
        | Update entry  $n$  from  $S$  and add to  $S'$ 
    end
    Remove  $n$  from  $\hat{n}$ 
    Remove  $n$  from  $S$ 
end
if  $S \neq 0$  then
    | Vehicles remaining in  $S$  have exited the merge zone
end
return  $S'$ 

```

Where, S is the set of all entries on the reservation table, S' is the incrementally built updated reservation table, and \hat{n} is the set of all vehicles n currently in the merge zone.

The second component conducts the optimisation process externally using Python. The pseudocode for this component is provided in Algorithm 2;

Algorithm 2: Pseudocode cooperative merge algorithm; optimisation component

```

import:  $S = \{Reservation\ Table\}$ 
initialise:  $S' = 0, \hat{j} = \{j\}$ 
while  $\hat{j} \neq 0$  do
    select  $j$  from  $\hat{j}$  to update
    for  $x_{1,j} > x_{2,j}$  do
        | calculate total pair delay  $d_{(1,2),j}$ 
    end
    for  $x_2 > x_1$  do
        | calculate total pair delay  $d_{(2,1),j}$ 
    end
    if  $d_{(1,2),j} > d_{(2,1),j}$  then
        | adjust kinematics of vehicle 1
        | append  $x_{1,j}$  and  $j_{2,j}$  to  $S'$ 
    end
    else if  $d_{(2,1),j} > d_{(1,2),j}$  then
        | adjust kinematics of vehicle 2
        | append  $x_{1,j}$  and  $j_{2,j}$  to  $S'$ 
    end
    Remove  $j$  from  $\hat{j}$ 
end
return  $S'$ 

```

Where, \hat{j} is the set of all conflict pairs in the reservation table, j is the conflict pair currently under investigation, $x_{n,j}$ is the kinematic information of vehicle n in conflict pair j , and $d_{(p,q),j}$ is the conflict pair delay of vehicle p preceding vehicle q in conflict pair j .

Delay is calculated as the difference in travel time from the vehicles current location to the end of the merge zone, travelling at both its current speed

$(v_{n|j,t})$ and its adjusted speed in response to the conflict ($v_{adjusted}$). Where, t is the current time increment in a modelling period of T , j is the vehicle conflict pair currently under consideration from all vehicle conflict pairs n and i is the vehicle currently under consideration from the conflict pair.

5.1.5 Interacting with Different Vehicle Types

The algorithm contains several provisions to facilitate the interaction between different vehicle types. The primary interaction of concern is between a human vehicle and a CAV. In the case where a human vehicle is the following vehicle, the human vehicle car-following model explained in Section 5.2.1 maintains appropriate spacing and headway to the lead vehicle. The model remains consistent regardless of whether the leader is another human or a CAV. This decision is predicated on the assumption that a human vehicle is incapable of discerning whether a leader is a CAV or another human vehicle.

However, when a CAV is following a human vehicle, the algorithm institutes a greater following distance. While CAVs are technologically capable of driving at short following distances, human drivers are not psychologically capable of the same. This characteristic is demonstrated by the high penalty placed on short following distances in many human car-following models presented in Section 3.1. For this reason, the algorithm has a provision that allows a greater minimum car-following distance when a CAV follows a human vehicle.

The cooperative merge component of the algorithm presented in Section 5.1.4 is contingent on V2I communication protocol, as all vehicles in the system forfeit control to a centralised controller while navigating through a merge point. Human vehicles can be given direction through VMS systems, however, their compliance can not be guaranteed. For this reason, the algorithm assumes that the trajectory of the human vehicle is unalterable. When a CAV interacts with a human vehicle at a merge point, the CAV will always adjust its trajectory to prevent the conflict.

Finally, the last provision made by the algorithm is to facilitate the implementation of a hierarchical priority structure. The algorithm allows the CAV (or human) fleet to be segregated into any number of subclasses, with the subclasses engaging in selective cooperation. This provision allows experimen-

tation involving non-complete cooperation between all vehicle manufacturers. Human vehicles are granted the highest subclass rank, meaning CAVs are forced to cooperate with and facilitate the actions of human vehicles.

The same provision in the algorithm that facilitates subdividing the CAV fleet also allows each subclass to have different behavioural characteristics. This can be implemented by changing the algorithm parameters for each subclass. Such a decision may be desirable when investigating the effect of fleet heterogeneity and the consequence of manufacturers designing slightly different CAVs. However, the experimentation in this thesis uses a homogenous CAV fleet unless explicitly stated. The reason for this arises from the highly regulated nature of the vehicle manufacturing sector. Communication protocols are standardised under IEEE standards, and vehicles contain numerous standardised features such as airbags, seatbelts, ABS, and crumple zones. Additionally, operating motor vehicles contains a myriad of standardised network conventions and legislations. It is unreasonable to anticipate that CAVs will enter such a highly regulated industry and be given the freedom to diversify their operation. In fact, CAVs may be subject to significantly higher degrees of standardisation as many of the human factors involved in human driving are mechanised.

5.2 Microsimulation Emulation of Human Vehicles

For experimentation in microsimulation, the CAV control framework presented in Section 5.1 must be accompanied by a human vehicle emulation framework. This work does not redesign the behaviour of human vehicles. The literature, as presented in Chapter 3 has extensively explored feasible and innovative means for emulating human behaviour, many of which are used in commercial software and for evaluating state-significant infrastructure projects.

This thesis uses both the Vissim and Aimsun commercial microsimulators to assess the intersection, corridor, and network implications of CAVs and mixed fleets under a range of settings. The Vissim and Aimsun microsimulators have been in commercial development for many years, continually experiencing refinement and enhancement. The mathematical models that underpin their vehicle behaviour are both well documented and available in the literature. Additionally, their programming interfaces allow for a high level of control over vehicle behaviour. The remainder of this section describes the nature of

the human vehicle behavioural models used for microsimulation. All models presented in this chapter are used in this thesis without modification.

5.2.1 Car-Following Models

The Vissim platform uses the psycho-physical Wiedemann car-following model, while the Aimsun platform uses the collision-avoidance Gipps car-following model. Both are explained in greater detail in the following subsections.

Wiedemann Car-Following Model

The Wiedemann car-following model was explored as part of the literature review in Section 3.1.3. To briefly reiterate, the Wiedemann car-following model is a psycho-physical model, meaning that it subdivides the domain of human reaction into subregions. The subregions calculate acceleration differently, depending on the urgency and safety requirements of the vehicle. Refer to Figure 3 for a diagrammatic representation of the different regions. The regions operate as follows;

- Free Driving: In this state, the driver aims to maximise its travel speed. The preceding vehicle does not affect the behaviour of the current vehicle.
- Approaching: This state arises when a vehicle approaches a driver with a lower speed. Deceleration is set to result in no difference in speed between the current and lead vehicle when the driver has reached the safe following distance.
- Following: In this state, the vehicle follows the lead vehicle without changing acceleration. The safe following distance is kept constant. This behaviour is akin to the platooning behaviour of CAVs. However, this state is not stable and is only maintained when the lead vehicle has a lower desired velocity than the follower, but the follower is still travelling with speed within its tolerance and does not conduct a DLC.
- Braking: This state arises when the distance to the preceding vehicle falls below safe. The vehicle applies a medium to heavy deceleration to recreate a safe following distance. This situation arises if the lead vehicle abruptly changes behaviour or if an adjacent vehicle changes lanes in front.

For implementation in microsimulation, the regions follow a stringent mathematical structure and thresholds define different actions in various regions. The parameters for the Wiedemann car-following model and the values used for this thesis are as follows;

- CC0 Standstill Distance ($1.5m$) – Average standstill distance between vehicles. It has no variance.
- CC1 Headway Time ($0.9s$) – Time headway used to calculate the average following distance.
- CC2 Following Variation ($4.0m$) – Headway by which the safe following distance is allowed to surpass before the following vehicle accelerates within maximum link speeds.
- CC3 Threshold for Entering Following ($-8.00s$) – Time taken to reach the safe following distance when a slower lead vehicle is registered.
- CC4 Negative Following Threshold ($-0.35m/s$) – Sensitivity of the vehicle to the lead vehicles negative changes in velocity.
- CC5 Positive Following Threshold ($0.35m/s$) – Sensitivity of the vehicle to the lead vehicles positive changes in velocity.
- CC6 Speed Dependency of Oscillation ($11.44m \cdot s$) – Influence of distance on speed oscillations, with 0 indicating that speed oscillations are independent of distance.
- CC7 Oscillation Acceleration ($0.25m/s^2$).
- CC8 Standstill Acceleration ($3.50m/s^2$) – Desired acceleration when starting from a standstill, limited by the maximum acceleration.
- CC9 Acceleration with ($1.50m/s^2$) – Desired acceleration at $80km/hr$, limited by the maximum acceleration.

These parameters refine behaviour when localised deviations from standard behaviour are observed. However, the default parameters recommended by Vissim are performing well, and so adjustments to these parameters are not warranted. The mathematical models that define the regions of this model are calculated as follows [Aghabayk et al., 2013];

$$A_x = L + CC0 \quad (104)$$

Where, A_x is the collision threshold and L is the length of the lead vehicle.

$$B_x = A_x + CC1 \times v \quad (105)$$

Where, B_x is the deceleration threshold and v is the velocity of the vehicle if it is slower than the lead vehicle. Otherwise, it is the velocity of the lead vehicle with an error term randomly generated between -0.5 and 0.5.

$$CLDV = \frac{CC6}{17000} \times (\Delta x - L)^2 - CC4 \quad (106)$$

Where, $CLDV$ is the reaction threshold and Δx is the distance headway between the current and lead vehicle.

$$SDV = -\frac{\Delta x - B_x - CC2}{CC3} - CC4 \quad (107)$$

Where, SDV is the perception threshold.

$$SDX = B_x + CC2 \quad (108)$$

Where, SDX is the unconscious reaction threshold.

$$OPDV = -\frac{CC6}{17000} \times (\Delta x - L)^2 - \delta \times CC5 \quad (109)$$

Where, $OPDV$ dictates the upper bound for when unconscious reaction still applies, and δ is a dummy variable that is 1 when the subject speed is greater than $CC5$ and 0 otherwise.

The Wiedemann model for human vehicles in urban settings and the default parameters recommended by the commercial microsimulator are as follows;

- Average Standstill Distance (2.00m) – Average distance between two vehicles with a deviation of 1.0m normally distributed about 0m.
- Additive Part of Safety Distance (2.00) – Used in the calculation of the safety distance (explained in Equation 110).
- Multiplic Part of Safety Distance (3.00) – Used in the calculation of the safety distance (explained in Equation 110).

The safety distance (b) mentioned above is a function of the additive part (b_{add}) and the multiplic part (b_{mult}). The safety distance is calculated as;

$$b = (b_{add} + b_{mult} \times z) \times \sqrt{v} \quad (110)$$

Where z has a range of 0 to 1, normally distributed about 0.5 with a standard deviation of 0.15, and v is the vehicle velocity (m/s).

Consider the following hypothetical example of vehicle kinematics when a vehicle is subject to an imprudent lane-change;




Figure 13: A hypothetical example of the changes in vehicle kinematic information for an agent subjected to an imprudent lane change.

The vehicle is initially in a stable condition prior to an imprudent lane change occurring. The lane change causes the headway between the ego vehicle and its leader to immediately fall below an acceptable level. This occurrence is demonstrated by the horizontal black line in Figure 13. The vehicle then responds by decelerating and reducing its velocity, recreating the appropriate headway between itself and the new leader. The overcompensation by deceleration is corrected through acceleration, and again a reduction in the headway between the vehicles occurs. This oscillatory pattern continues until new homeostasis is established.

The degree of oscillation in Figure 13 is exaggerated for demonstrative purposes. While a collision-avoidance model such as the Gipps car-following model makes minute continuous alterations to vehicle kinematics, a psycho-physical model such as the Wiedemann car-following model makes larger corrections when a vehicle's situation crosses specific thresholds. This characteristic means it is common during microsimulation to observe a vehicle exhibit large decelerations at merge junctions and during imprudent lane-changing.

Gipps Car-Following Model

The Gipps car-following model was explained in brief in Section 3.1.2. To reiterate, the Gipps collision-avoidance car-following model was developed by restricting vehicle kinematics, and ensuring that headway is maintained that can facilitate vehicles braking at maximum deceleration. The velocity of the vehicle using this model is calculated by;

$$v(n, t + \tau) = \min(v_a(n, t + \tau), v_b(n, t + \tau)) \quad (111)$$

Where, $v(n, t + \tau)$ is the velocity of the vehicle at time t , τ is the reaction time of the driver, $v_a(n, t + \tau)$ is the velocity of the vehicle during acceleration and $v_b(n, t + \tau)$ is the velocity of the vehicle during deceleration.

The driver's acceleration towards the desired speed is calculated by;

$$v_a(n, t + \tau) = v(n, t) + 2.5a(n)\tau \left(1 - \frac{v(n, t)}{v^*(n)}\right) \sqrt{0.025 + \frac{v(n, t)}{v^*(n)}} \quad (112)$$

Where, $v(n, t)$ is the velocity of vehicle n at time t , $v^*(n)$ is the desired speed of the road section and $a(n)$ is the maximum acceleration of the vehicle.

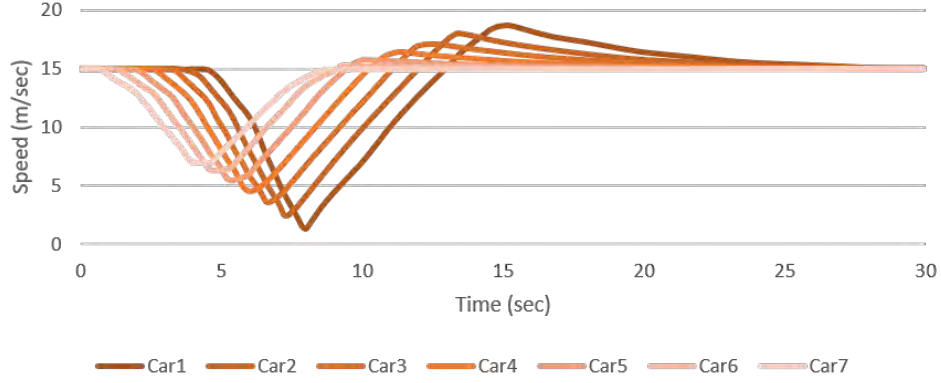
The deceleration equation is presented in the following Equation;

$$v_b(n, t + \tau) = d(n)\tau - \sqrt{\begin{matrix} d(n)^2 - d(n)(2x(n-1)\tau - l(n-1)) \\ -x(n, t)\tau - \frac{v(n-1, t)^2}{d'(n-1)} \end{matrix}} \quad (113)$$


Where, $d(n)$ is the maximum deceleration tolerated by the vehicle $x(n, t)$ is the position of vehicle n at time t , $l(n-1)$ is the effective length of the preceding vehicle and $d'(n-1)$ is an estimation of the preceding vehicle's deceleration.

Gipps evaluated the model by using parameters deemed appropriate for human drivers. Parameters were set to a maximum acceptable acceleration ($a(n)$) of 2.0 m/s^2 , a maximum acceptable deceleration ($d(n)$) of -3.0 m/s^2 , a desired velocity (v_n) of 20.0 m/s and a reaction time (τ) of $2/3 \text{ s}$. The deceleration parameter acts as a mechanism for dampening disturbances that perturbed through the flow. Figure 14a display the speed-time graphs for seven successive vehicles when the deceleration of the preceding vehicle is under-predicted

($d'(n-1) < d(n)$). Similarly, Figure 14b displays the graph for when the deceleration of the lead vehicle is over predicted ($d'(n-1) > d(n)$).



(a) Speed-Time graph under predicted lead vehicle deceleration ($d'(n-1) < d(n)$)



(b) Speed-Time graph for over predicted lead vehicle deceleration ($d'(n-1) > d(n)$)


Figure 14: Speed-Time graphs generated using the Gipps collision-avoidance car-following model [Gipps, 1981].

When $d'(n-1) < d(n)$ (Figure 14a), the deceleration intention of the preceding vehicle was underestimated. This prediction forces the following vehicles to brake harder, resulting in the headway to increase. To compensate, the following vehicles then increase their velocities to reduce the headway. This leads to the formation of a wave-like reaction lag that can be observed propagating through a traffic stream. When $d'(n-1) > d(n)$ (Figure 14b), the deceleration intentions of the preceding vehicle was overestimated. This prediction forces the following vehicles to brake at a lower rate, allowing stability to re-establish within the traffic stream sooner.


For microsimulation implementation, the following parameter values are used;

- τ - Reaction time, set to $2/3s$.
- a_n - Maximum acceleration, determined by a normal distribution for each vehicle to be $a_n \sim N(1.7, 0.3^2) \text{ m/s}^2$.
- v_{max} - Maximum velocity for a vehicle, determined by a normal distribution to be $v_{max} \sim N(20, 3.2^2) \text{ m/s}^2$.
- b_n - Maximum deceleration for a vehicle, set at $-2.0a_n \text{ m/s}^2$.
- \hat{b} - Deceleration of the vehicle, calculated by $\min\{-3.0, (b_b - 3.0)/2\} \text{ m/s}^2$.

By fixing the probabilistic parameters to their mean values, the Gipps car-following model can be described as being a function of the vehicles current velocity ($v_n(t)$), the proceeding vehicles current velocity ($v_{n-1}(t)$), and the headway between the two vehicles ($x_{n-1}(t) - s_{n-1} - x_n(t)$). This then allows the model to be graphically represented. Figure 15 below shows the acceleration profile for the Gipps car-following model as a function of the vehicles current speed when the velocity of the preceding vehicle is 18m/s ;



(a) Acceleration in the next time increment, as a function of velocity and headway.



(b) Change in acceleration as a function of velocity and headway. (c) Change in velocity as a function of velocity and headway.

Figure 15: Effect on vehicle kinematics using the Gipps car-following model, assuming statistically distributed parameters attain their mean value.

Using the Gipps car-following model, vehicles travel at their maximum velocity until a threshold is reached. The threshold is a combination of speed and headway, where vehicle operation is assumed to be safe. Once this threshold is reached, the model implements a high deceleration rate to rectify the safety issue, a behaviour that is seen in Figure 15 as a sudden drop. This figure indicates that the change in acceleration and change in velocity is consistently low, but suddenly and drastically varies the moment that the critical point in velocity and headway is reached.

Gipps identified the sudden and extreme reaction of this model as a limitation. During normal car-following behaviour, the model will not have the opportunity to reach these threshold values, as vehicle acceleration will be moderated well in advance. However, these situations may occur spontaneously at merge junctions and lane changes, where vehicles from the adjacent lane imprudently enter into small headways. Strong deceleration rates resulting from such actions perturbs the traffic and creates a ripple effect downstream of the braking location. These perturbations can significantly reduce the efficiency of a traffic stream and result in unnecessary delays, similarly to human driving conditions. These characteristics make the Gipps car-following model appropriate for human behaviour.

5.2.2 Gap-Acceptance Models

The gap-acceptance models in both the Vissim and Aimsun platform act as a boolean check for the gap size, and determine acceptability. This section explains in greater detail the gap-acceptance models derived from the Wiedemann models for Vissim, and the Gipps models for Aimsun.

Vissim Gap-Acceptance Process

The default gap-acceptance process used by the commercial microsimulator is retained in its native state for manual vehicle operation [PTV Group, 2016]. Gap-acceptance in Vissim is contingent on the speed of the lane-changing vehicle and the target lane following vehicle. Gap-acceptance is used for three major actions in Vissim, lane-changing, negotiating space through conflict areas, and reacting to priority rules. The remainder of this subsection explains in brief, gap-acceptance in each of these circumstances.

Gap-Acceptance in Lane-Changing

Lane-changing can be divided into two types, the first is a regular lane change, and the second is an overtaking manoeuvre. For a regular lane change, the modeller dictates the minimum acceptable gap, set by default to $0.5m$. Minimum gap is also subject to the safety distance determined by Equation 110.

An overtaking manoeuvre involves changing to the adjacent lane, accelerating beyond a vehicle or a group of vehicles, and changing back to the original lane. Refer to Figure 16 for an example;




Figure 16: An overtaking manoeuvre conducted in the Vissim microsimulator.

The overtaking manoeuvre requires a vehicle to change to the adjacent lane and then return downstream, making it inherently more dangerous than a traditional lane change. For this reason, Vissim requires a gap downstream of the last vehicle being overtaken that is at least the length of the vehicle plus two times the safety distance defined by Equation 110.

Gap-Acceptance in Conflict Areas

Conflict points in Vissim identify locations in the microsimulation network where the trajectory of two opposing vehicles may cross. The rules at a conflict point dictate which of the two vehicles has priority. Refer to Figure 17 for an example of a priority rule. The priority rules can be configured in one of three ways. The first is to provide neither movement with a priority (Figure 17a). The second is to provide a specific movement with priority over another (the East-West movement in Figure 17a), and the last is to force a movement to yield to another (the East-West movement in Figure 17c). This figure does not highlight all conflicts, as turning movements would also conflict with through movements;

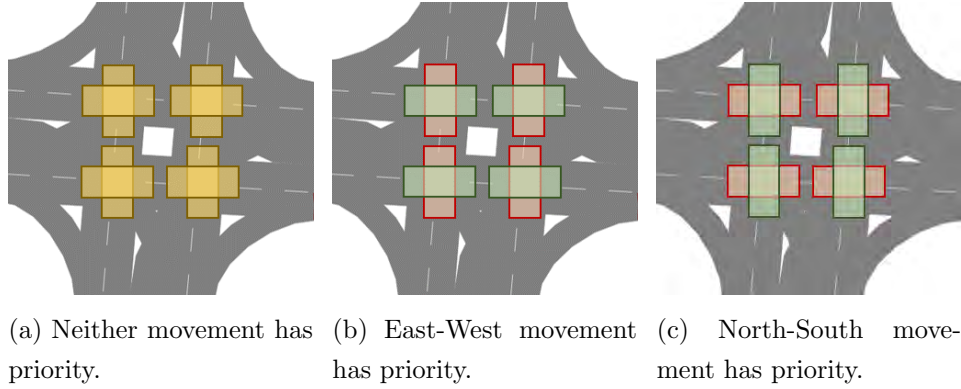


Figure 17: Example of conflict areas in a 4-way signalised intersection.

As vehicles approach a conflict area, they become aware of other vehicles also approaching the conflict point. Depending on which of the three alternatives the conflict point is configured in, the vehicle will adjust its trajectory and kinematics to ensure the gap between vehicles increases to greater than the minimum required gap. The minimum gap is calculated using Equation 110.

Gap-Acceptance in Priority Rules

Priority rules differ from conflict areas in that they allow the modeller to specify a distance and time headway, defining in greater detail the nature of an acceptable gap. Figure 18 provides an example of how priority rules are structured;

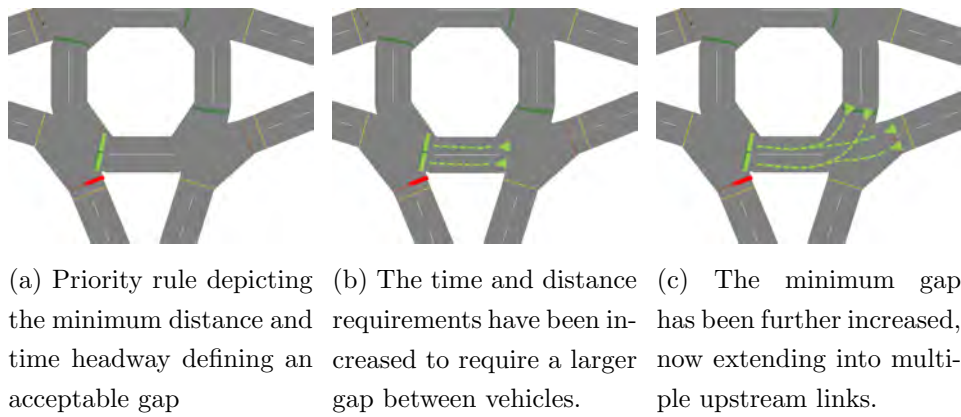


Figure 18: An example of a priority rule used at the approach arm of a roundabout. The three subfigures indicate how the gap between the priority (green line) and yielding (red line) vehicle can be altered.

When configuring the priority areas, the modeller has control over three parameters; distance headway, time headway, and maximum speed. Referring to Figure 18, the distance headway defines how far behind the solid green line a conflicting vehicle must be for a vehicle waiting at the red stop line to accept a gap. The green triangle indicates distance headway. Time headway works similarly, except calculates the threshold headway as a time. The third parameter (maximum speed) defines a threshold beyond which the priority rule no longer considers distance headway and only uses time headway.

Modellers use the priority rule approach over the conflict areas approach to manually define the gap-acceptance criteria as opposed to vehicles adhering to a model. This approach may be necessary for forcing the model and vehicle behaviour to better adhere to observed behaviour, for calibration and validation purposes. If no such unique behaviour is observed, the recommendation is to use conflict areas over priority rules to define gap-acceptance.

Gipps Gap-Acceptance Model

The Gipps gap-acceptance model was already presented as part of the literature review in Section 3.3. To reiterate, the Gipps gap-acceptance model determines the acceptability of a gap based on the following three requirements; the gap is positive, the computed speed is positive, and the imposed deceleration is smaller than the maximum desired deceleration. The minimum gap is calculated between the proceeding and following vehicle. The minimum upstream gap is evaluated using the following equation;

$$Gap_{up}(t) \geq \max(0, \frac{v_{lc}(t)^2}{2d_{lc}} + 0.5v_{up}(t)\tau_{up} + \max(0, -\frac{v_{up}(t)^2}{2d_{up}} + a_{up}(1 - 0.5a_{up})d_{up}\tau_{up}^2 + (1 - a_{up})v_{up}(t)\tau_{up})) \quad (114)$$

Where, $Gap_{up}(t)$ is the minimum upstream gap, v_{lc} is the ego vehicle velocity, d_{lc} is the ego vehicle deceleration, $v_{up}(t)$ is the upstream vehicle velocity, τ_{up} is the upstream vehicle reaction time, τ_{lc} is the ego vehicle reaction time, d_{up} is the upstream vehicle deceleration, and a_{up} is the upstream vehicle acceleration.

The minimum downstream gap is evaluated using the following equation.

$$Gap_{dw}(t) \geq \max(0, \frac{v_{dw}(t)^2}{2d_{dw}} + 0.5v_{lc}(t)\tau_{lc} + \max(0, -\frac{v_{lc}(t)^2}{2d_{lc}} + a_{dw}(1 - 0.5a_{dw})d_{lc}\tau_{lc}^2 + (1 - a_{dw})v_{lc}(t)\tau_{lc})) \quad (115)$$

Where, $Gap_{dw}(t)$ is the minimum downstream gap, $v_{dw}(t)$ is the downstream vehicle velocity, d_{dw} is the downstream vehicle deceleration, and a_{dw} is the downstream vehicle acceleration.

5.2.3 Lane-Changing Models

Vissim provides a distinction in vehicle behaviour for mandatory and discretionary lane-changing, and provides the user with a range of parameters that control the nature of the lane change. Aimsun, on the other hand, uses the rule-based Gipps lane-changing model to evaluate discretionary and mandatory lane changes. This section discusses both models.

Vissim Lane-Changing Model

An MLC is conducted when a vehicle must be situated in the appropriate lane to travel to the next node or link. Vissim allows the vehicle to use the maximum acceptable deceleration rate for itself and imposes this rate on the following vehicle in the target lane to create a sufficiently large gap. As the vehicle approaches the end of the link or node and is still situated in the incorrect lane, the maximum permitted braking rate increases. This rate ensures the vehicle comes to a standstill as it reaches the emergency distance before the end of the connector.

A DLC is conducted if the vehicle finds greater speed or space in an adjacent lane. While a modeller can not affect vehicle “aggressiveness”, the modeller can control the safety distance that is required for a DLC to be accepted. A suitable gap depends on the speed of the lane-changing vehicle and the speed of the following vehicle in the target lane. Greater detail regarding the safety constraints of an acceptable gap has already been discussed in Section 5.2.2.

The modeller has control over the following lane-changing model elements;

- Free lane selection (True): Whether vehicles are free to choose any lane as a target lane, or if changing into a specific lane is restricted.

- **Slow Lane Rule (False):** The slow lane rule, used in countries such as Australia and Germany, requires that vehicles return to the slow lane after conducting an overtaking manoeuvre. Cooperation can then be increased between vehicles if this rule is in effect, forcing vehicles to create acceptable gaps actively.
- **Maximum Deceleration ($-4m/s^2$):** The upper bound of the deceleration rate used by vehicles during a lane-changing event.
- **Accepted Acceleration ($-1m/s^2$):** The lower bound of the deceleration rate imposed on the vehicle and the following vehicle of the target lane during a lane-changing event.
- **Deceleration Slope ($-1m/s^2$ per $100m$ travelled):** The deceleration slope changes how quickly the accepted acceleration rate approaches the maximum deceleration rate, as the vehicle itself approaches the emergency stop point before the connector.

Gipps Lane-Changing Model

The Gipps lane-changing model was introduced in brief during the literature review in Section 3.2.1, and is explained in greater detail here. To reiterate, the Gipps lane-changing model evaluates the process of changing lanes by assessing “Is it possible to change lanes? Is it necessary to change lanes? Is it desirable to change lanes?” [Gipps, 1986]. An explanation of the factors in Gipps lane-changing model is provided in the remainder of this subsection.

The first question posed by the model is whether a lane change is physically possible. A lane change is deemed physically impossible if it meets either of the following three conditions. The first is that the target lane physically does not exist, such as moving left into the third lane of a two-lane road. The second is that the target lane contains an obstruction upstream. The third is that changing into the target lane causes a deceleration that violates the limits imposed by the Gipps car-following model.

If it is possible to change lanes, the vehicle will evaluate whether it has an upcoming MLC. If the MLC is “close”, the vehicle will make attempts to position itself in the correct lane. Close is defined as $10s$ of travel.

As the vehicle approaches an MLC and is not situated in the correct lane, its urgency to change lanes increases. Urgency is reflected in changes to the vehicle's willingness to brake harder and accept smaller gaps, given by;

$$d_n = (2 - \frac{x_{turn} - x_n(t)}{10v_n^*}) \times b_{n,max} \quad (116)$$

Where, d_n is the drivers braking rate, x_{turn} is the location of the turn, x_n is the position of the vehicle, v_n^* is the free-flow speed of the driver, and $b_{n,max}$ is the maximum braking rate of the driver.

Vehicles are also precluded from driving in a transit lane, except where an obstructed vehicle trajectory forces the action. In this case, a vehicle is allowed to change lanes into the transit lane 10s before the obstruction. While a non-transit vehicle is driving in the transit lane, the lane-changing model continues to check if the adjacent lane is clear and if the vehicle is safe to return to its original lane. The requirement to return, however, is ignored if the vehicle is close to an MLC required from the transit lane.

When the vehicle is not close to an intended lane change, it is considered to be positioned in the middle region. The middle region extends to 50s before the MLC. In the middle region, the vehicle can still change lanes, but subject to the condition provided in Equation 117. This condition ensures that the vehicle does not stray too far from a lane necessary for a downstream MLC;

$$(l_p - l_1)(l_i - l)(l_i - I) = 0 \quad (117)$$

Where, l_p is the preferred lane determined through the lane-changing model, l_i is the current lane, and I is the total number of lanes.

If previous factors have not forced a lane change, the vehicle is able to conduct a DLC. A DLC provides the vehicle with an advantage over its current situation. The vehicle first assesses competing lanes for a downstream obstruction that may affect its travel conditions less than the obstruction in its current lane. If none are found, it inspects for heavy vehicles. If none are found, the vehicle then checks whether competing lanes provide a travel speed advantage, calculated using the Gipps car-following model. Better travel performance in the adjacent lane and safe gap conditions allows the vehicle to change lanes.

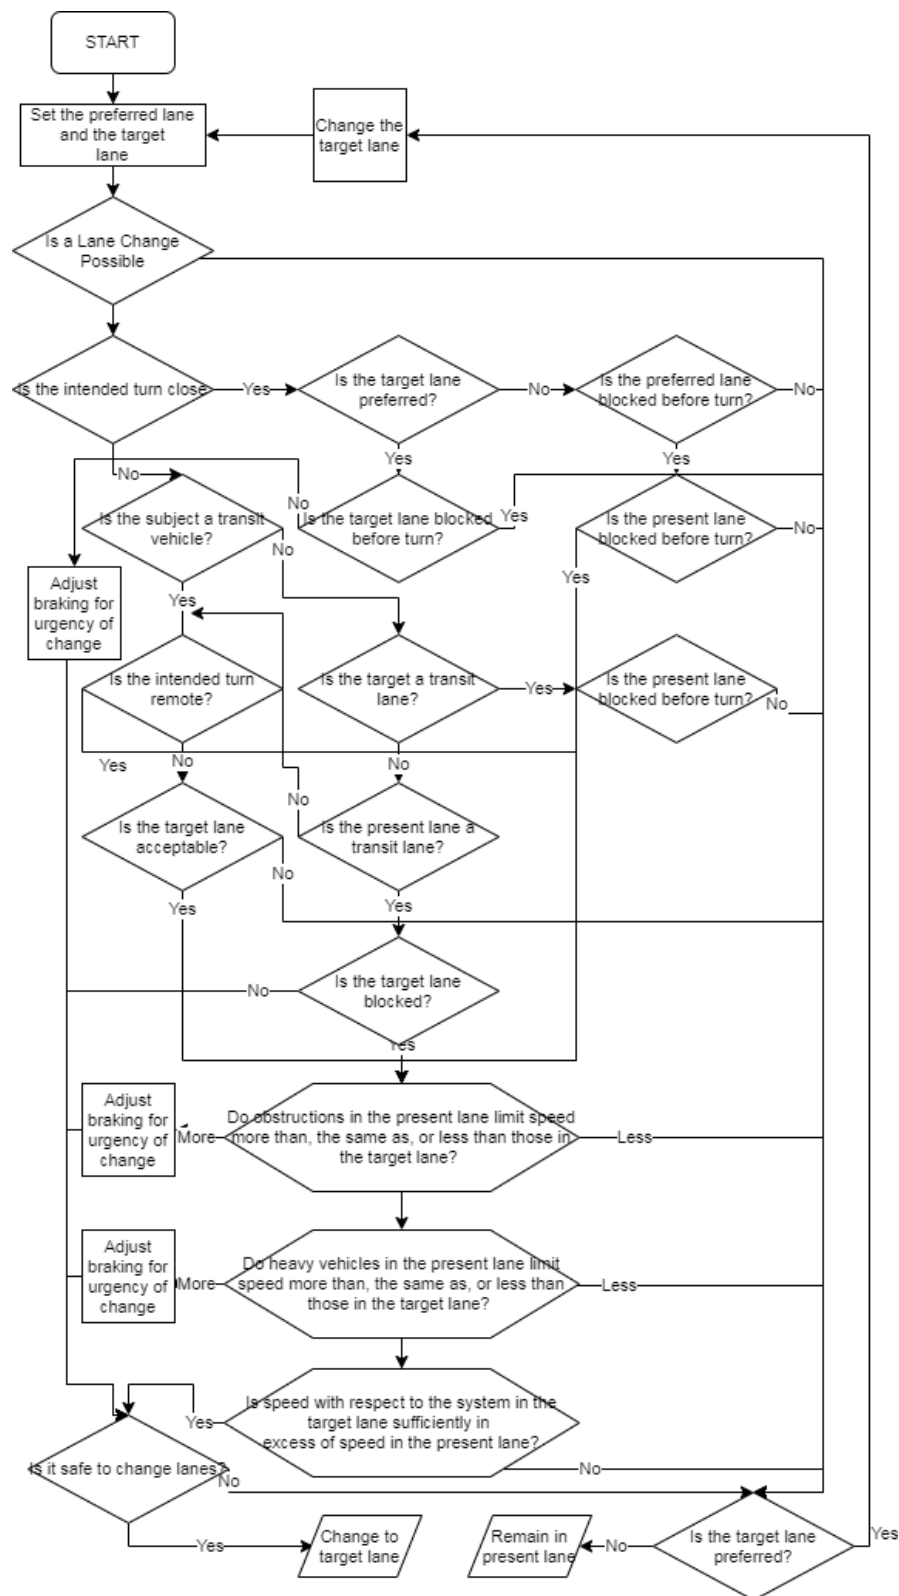


Figure 19: Flowchart of the Gipps lane-changing model [Gipps, 1986].

Using the aforementioned conditions, the vehicle ensures that safety constraints posed by the gap-acceptance model (Section 5.2.2) are not violated. If safety is satisfied, the vehicle executes the lane-changing manoeuvre. Figure 19 provides the rule-based structure of the Gipps lane-changing model.

5.3 Framework Summary

This chapter consisted of two sections, the first developed a detailed modelling framework for CAVs, and the second provided the accompanying human emulation framework. The CAV framework, the major contribution of this thesis, transcribed many of the V2V and V2I benefits explored in Chapter 2 into a set of mathematical models and rule-based modules. The benefit of this framework is in its modular design. The mathematical models are interchangeable when more information regarding CAV behaviour is available. Additionally, new calibration parameters proposed as part of the model development are limited and tractably related to real-world physics.

This framework leverages shorter vehicle spacing, instantaneous reaction times, cooperation during lane-changing, decentralised vehicle control through merge junctions, interplay with human vehicles, all while conservatively maintaining safety and passenger comfort. Model validation is a critical component of any model development. However, real-world operational data is not available to conduct validation. For this reason, the following chapter provides a detailed kinematic assessment of the framework and comparison with human vehicles, providing validity to the performance of the developed framework.

The second section of this chapter detailed the emulation of human vehicles in microsimulation. Since the investigations in this thesis focus on the impact of mixed-fleets, it is critically necessary to also emulate their behaviour. Emulating human behaviour is not a unique contribution of this thesis, and so this component of the framework uses models proposed by the literature and already implemented in commercial microsimulators.

Acknowledgement of Publication

This work successfully facilitated the publication of two journal article;

Sinha, A., Chand, S., Wijayaratna, K. P., Viridi, N., & Dixit, V. (2020). Comprehensive safety assessment in mixed fleets with connected and automated vehicles: A crash severity and rate evaluation of conventional vehicles. *Accident Analysis & Prevention*, 142, 105567.

Sinha, A., Chand, S., Wijayaratna, K. P., Viridi, N., & Dixit, V. (2020). Crash Severity and Rate Evaluation of Conventional Vehicles in Mixed Fleets with Connected and Automated Vehicles. *Procedia Computer Science*, 170, 688-695.

6 Vehicle Kinematics and Parameter Sensitivity

Abstract: Models require validation before they are fit for forecasting and predictive works. The process of validation compares data generated by the model with observed data collected from real-world experimentation. Data pertaining to CAV operation is not available, especially for CAVs in a network setting. For this reason, validation in the traditional sense is not possible to perform. This chapter uses a range of other methods to validate the proposed CAV emulation model in the absence of real-world data. The first method generates kinematic-time plots to qualitatively demonstrate the stability of the emulation model and the consistency of CAV behaviour in platoons. The second method conducts a thorough statistical analysis of vehicle kinematics during network operation using trajectory data. The analysis demonstrates that as network CPR increases, vehicles accelerations and vehicle jerks do not statistically experience a change, which vehicle velocity experiences an increase. These results strongly demonstrate that the CAV control algorithm does not vary vehicle behaviour by artificially manipulating existing parameters such as acceleration rates and speed limits. The underlying kinematics of the CAV has been kept consistent with the human vehicle, meaning that velocity improvements and network efficiency increases are the result by design of the algorithm. Finally, this chapter concludes with a detailed sensitivity assessment of the parameters proposed in the framework development.

This chapter evaluates the kinematic performance of the CAV emulation framework defined in Chapter 5.1 and the human vehicle emulation framework defined in Chapter 5.2. This evaluation provides an insight into CAV vehicle operation, ensuring that vehicle safety and comfort has not been sacrificed in the pursuit of network efficiency and performance. The assessments in this chapter also serve as a form of pseudo-validation. Validation is the process of benchmarking modelled data against real-world observed data, where a fit-for-purpose model would limit deviations between the two data sets. In the absence of real-world data, traditional forms of validation can not be conducted. This chapter conducts a comprehensive statistical study into the kinematic performance of human, mixed, and CAV fleets. By demonstrating that statistically there is no change in the underlying acceleration and jerk kinematic variables, improvements in network performance are solely the re-

sult of the developed algorithms performance, and not through manipulation of superficial vehicle and driving parameters.

A kinematic assessment of the CAV emulation framework developed in this thesis is carried out using three approaches. The first is to create a space-time plot of a vehicle platoon under normal operation and lane-changing. These plots reveal whether platoon stability is maintained and helps highlight erratic behaviour arising due to the framework. The second is to perform a kinematic assessment on a small network. Vehicle kinematics consist of the velocity (m/s), acceleration (m/s^2), and jerk (m/s^3). Kinematics provides valuable insight into vehicle comfort as they are a proxy for forces experienced by passengers. Finally, the third assessment is to perform a sensitivity analysis of the parameters introduced in this framework (P_{min} , P_{max} and n).

Vehicle passengers are more sensitive to vehicle forces, as drivers can brace against the steering wheel. Also, drivers generally are more comfortable in situations over which they have direct control. For both these reasons, an adverse change in vehicle kinematics would indicate inappropriate parameters in the behavioural control framework. Also, a kinematic assessment will highlight whether vehicles are behaving more erratically and closer to their extremes. The sensitivity analysis intends to determine the degree to which the parameters influence vehicle behaviour.

6.1 Kinematic-Time Plot Assessment

The kinematic-time plot assessment provides a qualitative indication of stability and reasonable behaviour. The quantitative assessment can be found under Section 6.3 and Section 6.4, where the framework is implemented in a microsimulation environment. Before the kinematic-time plots can be created, the base case parameters must first be defined. For the remainder of this thesis, except where explicitly stated, $P_{max} = 0.75$, $P_{min} = 0.25$, $n = 2\tau$, $\tau = 0.1s$, $a_{max} = 4m/s^2$, $J_{max,approach} = 1m/s^3$, and $J_{max,critical} = 3m/s^3$.

Ten CAVs are successively introduced to a road with a speed limit of $60km/hr$. The vehicles are generated at a headway of $1.5m$ and instructed to maintain a minimum headway of $2.0m$. The minimum headway for CAVs is set to $0.5m$ (refer to Section 6.2 for a justification of this headway value). The value is

increased to $2.0m$ here to create the kinematic-time plots, as a short headway makes the individual vehicles difficult to discern. A red-light traffic light is placed $100m$ downstream from the generation point of the vehicles. Figure 20 and Figure 21 provide the kinematic-time diagrams for a car-following scenario and a lane-changing scenario respectively;

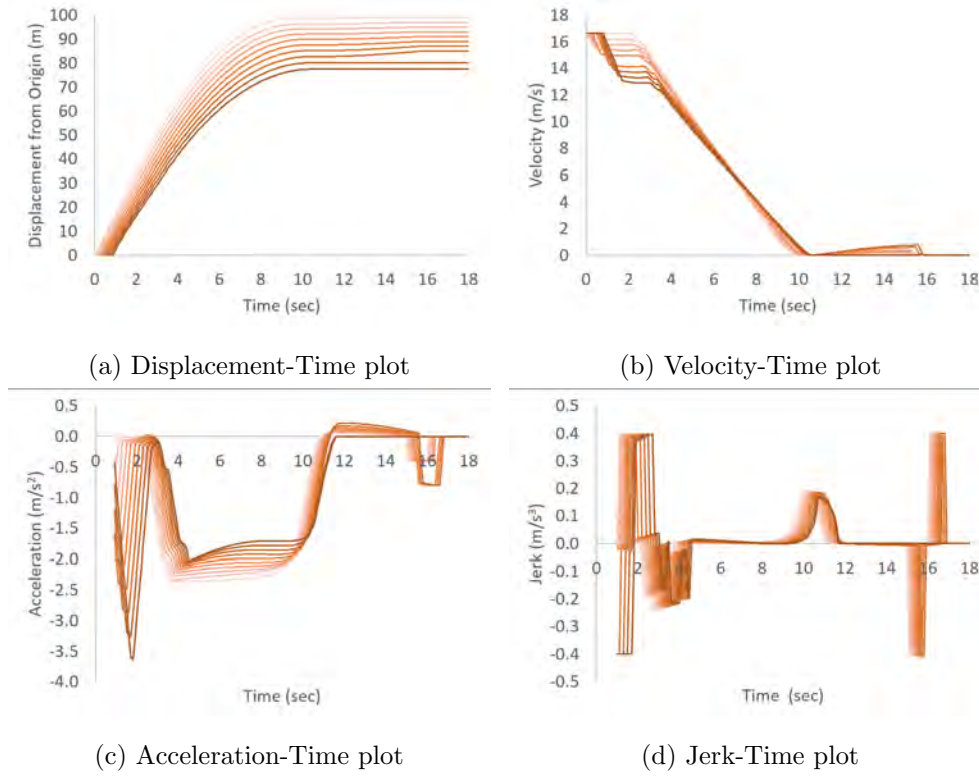


Figure 20: Kinematic variable vs time plots for the scenario depicting the car-following for a platoon. Vehicles are generated at a spacing less than the minimum permitted distance and are approaching a red light $100m$ away. Each consecutive vehicle is represented by a darker shade.

The displacement-time plot from Figure 20 indicates that spacing between vehicles increases as time increases. Vehicles are generated into the environment at a headway ($1.5m$) lower than their minimum permitted value ($2.0m$), causing them to decelerate immediately. The acceleration-time plot in the same figure effectively displays the transition in vehicle behaviour, from reacting to a lead vehicle to responding to a trajectory obstruction. The initial deceleration is in response to the headway violation. Each consecutive vehicle has a cumulatively smaller headway to its lead vehicle, leading to a consequently

higher deceleration. Vehicle velocity stabilises before vehicles transition to the $x_{approach}$ zone and must start reacting to the encroaching traffic light. Figure 20 indicates that vehicle behaviour does not surpass the maximum defined thresholds, nor does erratic behaviour arise. Each consecutive vehicle exhibits behaviour resembling its lead vehicle, attesting to the stability of the framework. The kinematic-time plots for the lane-changing scenario are as follows;

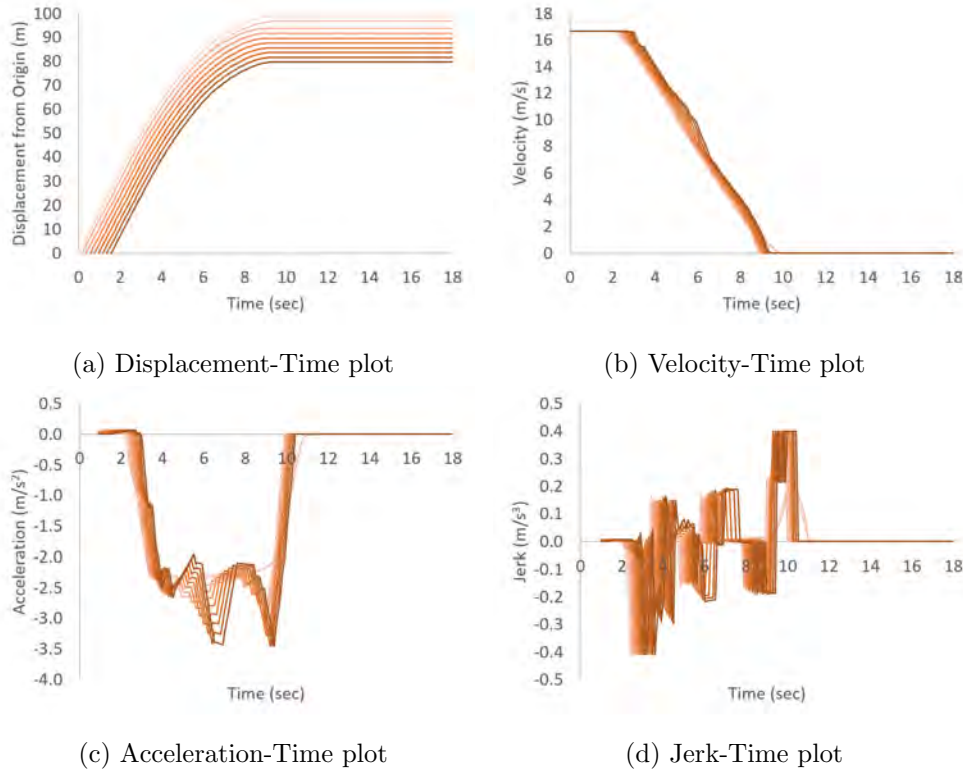


Figure 21: Kinematic variable vs time plots for the lane-changing scenario, where an adjacent vehicle enters the platoon. Vehicles still react to a red light 100m away. Each consecutive vehicle is represented by a darker shade.

For the lane-changing scenario depicted in Figure 21, vehicles are introduced to the environment at a headway of $3.0m$, and vehicle three travels in parallel to vehicle two until $2.0s$. At $2.0s$, vehicle three communicates its intention to change lanes, at which point, deceleration of all vehicles can be seen in the acceleration-time plot. Once an adequate gap has been created, by approximately $6.0s$, the displacement-time plot shows the entrance of vehicle three through the branching of the line representing vehicle two.

Figure 21 represents the unique and critical element of this modelling framework. When the lane-change is triggered at $t = 3s$, the platoon reacts as a singular unit and adjusts acceleration collectively. This behaviour is in direct contrast to the velocity-time plots generated for the Gipps car-following model presented in Figure 14. Since this is a discrete-time model, with vehicles reacting to direct leaders, the platoon still exhibits a delay of one time-increment ($\tau = 0.1s$). The perception-reaction lag could be further reduced by having a vehicle react to the platoon leader, as opposed to their direct lead vehicles. However, this approach would only provide further benefits for long platoons. At a simulation resolution of $\tau = 0.1s$, a platoon of ten vehicles will experience an average delay of $0.5s$ due to the discrete-time nature of this framework. The following section uses vehicle trajectory data generated from microsimulation for a more detailed quantitative kinematic assessment.

Figure 20 and Figure 21 display the changes in kinematic variables relative to longitudinal displacement from the origin. They do not, however, provide the plots for lateral displacement as the vehicles change lanes. The reason why the later have been omitted is that lateral displacement is controlled by the microsimulator and not by the proposed algorithm. When the model developed in this thesis approves a lane-changing request, control of the vehicle reverts to the microsimulator which then controls wheel alignments and lateral displacement. Once the lane change manoeuvre is completed, vehicle control reverts to the developed algorithm. The kinematic variable plots relative to longitudinal displacement demonstrated the contribution of the developed framework. Lateral displacement is not a contribution of this thesis, and so similar plots would repeat the work of commercial microsimulators. Additionally, the way in which the microsimulator control lateral displacement is unique to each commercial modelling package.

6.2 Network Evaluation of Vehicle Kinematics

A microsimulation modelling approach is used to determine the kinematic performance of a network that transitions from human-driven vehicles to CAVs. The microsimulator used is Vissim 9.09, commercially developed by PTV [PTV Group, 2016]. The framework for CAV emulation presented in Chapter 5.1 is implemented using the Vissim Application Programming Interface (API). The API is written in C++ and allows the modeller to control every

aspect of vehicle behaviour between entering and exiting the network.

The framework parameters are consistent with those presented in Section 6.1; however, the minimum headway between vehicles is $0.5m$. CAVs are expected to communicate at a rate of 1 communication per $515\mu s$ [Milanes et al., 2012], which is approximately 2000 cycles per second. In an extreme scenario, the CAVs are operating at $22.2m/s$ at a headway of $0.5m$. For a collision to occur in the worst-case scenario, the lead vehicle would need to decelerate at its maximum rate ($-8m/s^2$), and communication would need to experience complete disruption for $1.12s$ consecutively, 2,240 consecutive communication cycles.

The intention of this work is not to assess or validate the claim that CAV communication will not be disrupted for $1.12s$. This seems unreasonably high for the mass deployment of automated vehicles. Other studies [Latrech et al., 2018] have used a minimum headway comprised of the standstill distance and a factor of the current travel velocity. A faster-moving vehicle does not justify a decrease in safety, only an increase in consequences should an accident occur. When assuming that CAV sensing equipment is faultless, the velocity component of minimum headway becomes redundant, and the minimum following distance can be reduced to the standstill distance.

The remainder of this subsection provides a detailed description of the modelling environment, an assessment on model stability, and an analysis of the network kinematic performance.

6.2.1 Network Geometry

The network environment for this kinematic assessment (provided in Figure 22) consists of two signalised intersections, one roundabout, one priority intersection, and a merge junction. The approach arm of each intersection is separated by $100m$, with a network speed limit of $60km/hr$. Driving convention follows Australian laws, with vehicles adhering to the road's left side.






Figure 22: Network used for the kinematic assessment of the CAV framework.

The top-left signalised intersection contains a filtered right turn on all approaches, which yields to the conflicting through movements. The bottom-left signalised intersection differs in that it contains a dedicated right turn lane with a corresponding green-time in the signal phasing. The roundabout intersection follows Australian standards where vehicles navigate through it clockwise. The priority junction gives priority to the East-West movement, with the North-South movement yielding and waiting for a sufficient gap. Finally, the merge junction is defined as the dual-lane mainline situated on the right and the onramp merging lane situated on the left. The merge lane yields to the mainline.

The signalling structure is presented in Figure 23. For the top-left signalised intersection phasing (Figure 23a), each “b” phase is separated by a 4s inter-phase time with no movements. The total cycle time is then 60s. For the bottom-left signalised intersection phasing (Figure 23b), each phase is separated by a 4s inter-green time that contains non-conflicting movements from the preceding and following phases. This structure results in a total cycle time of 100s. The phasing diagrams are as follows;



(a) Signalling structure used for the top left signalised intersection. The green colour indicates movements with priority, while the yellow movements are those that yield to opposing conflicting movements.



(b) Signalling structure used for the bottom left signalised intersection.

Figure 23: Signal phasing structure of for the signalised intersections

The model contains a $30min$ warmup period with a $60min$ assessment period following the warmup. The warmup is loaded with 80% of the demand compared to the peak modelling period. Seeing as this is a hypothetical network, the demand was set visually to ensure that the network reached capacity but did not cause latent demand and off-network queueing. The origin-destination (OD) matrix used in this study is provided in Table 1;

O\D	1	2	3	4	5	6	7
1	0	93.75	93.75	93.75	93.75	93.75	93.75
2	75	0	75	75	75	75	75
3	93.75	93.75	0	93.75	93.75	93.75	93.75
4	93.75	93.75	93.75	1875	93.75	93.75	93.75
5	93.75	93.75	93.75	93.75	0	93.75	93.75
6	93.75	93.75	93.75	93.75	93.75	0	93.75
7	93.75	93.75	93.75	93.75	93.75	93.75	0

Table 1: Origin-destination matrix used for the peak hour (veh/hr). The zones correspond to those outlined in Figure 22.

Microsimulation modelling techniques use stochastic probability distributions in parameters to model day-to-day variations in traffic networks. If the variability in parameters such as network loading rates results in drastically different results, then the model can not be used to reliably forecast future scenarios [RMS, 2013]. For this reason, model stability must be established for the environment. Model stability was established using 50 random seeds to verify the model’s internal consistency. Table 2 provides aggregated network results for all seeds;

	Avg Speed (km/hr)	Total Dist (km)	Total Travel Time (sec)	Total Delay (sec)	Total Stops	Serviced Demand
Median	28	3,270	420,583	232,809	10,710	5,647
STD	1	11	21,186	20,963	1,455	21
STD/Med	4.5%	0.3%	5.0%	9.0%	13.6%	0.4%

Table 2: Results of the internal consistency analysis conducted on this environment using 50 random seeds.

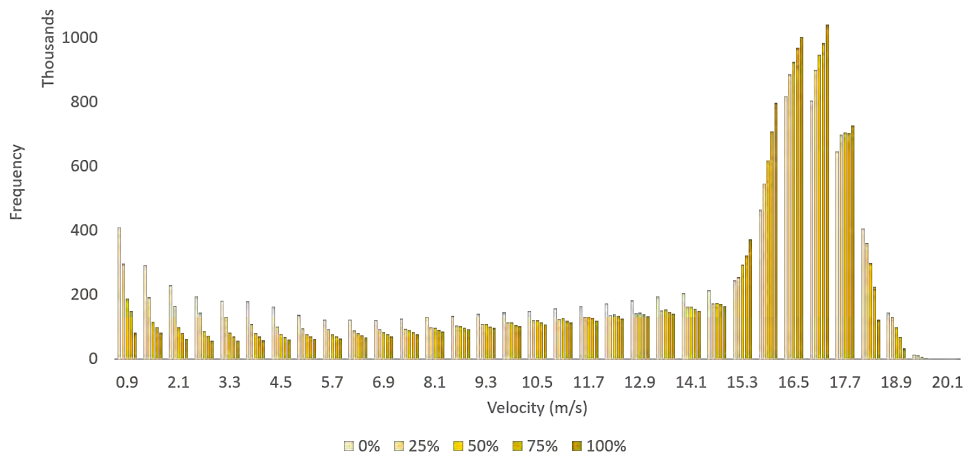
The standard deviation as a percentage of the median result is below 15% for all performance metrics, with the majority falling under 5%. Serviced demand has a standard deviation to median value ratio of 0.4%, indicating that all demand enters the network. The total distance and total travel time have a ratio of 0.3% and 5.0% respectively, indicating that irregularities in network operation such as grid-locking do not occur. The model is considered stable using these results. The median seed is determined by calculating the distance of each seed from the median value. Seed number 2461 exhibited the most median behaviour and is used for the kinematic and sensitivity analysis.

6.2.2 Kinematic Results


Five iterations of the model were run that transitioned the fleet from human-driven to CAV in 25% increments. The microsimulator during runtime produces a trajectory file that contains the vehicle ID of every vehicle in the network, along with instantaneous speed and instantaneous acceleration for each 0.5s time increment. The Surrogate Safety Assessment Module (SSAM) developed by the US Federal Highway Administration (FHWA) was used to convert the raw trajectory data into a Comma Separated Value (CSV) file. The raw data is then converted to a 2D array using Python, with vehicle ID

on the x-axis, timestamp on the y-axis, and the array containing either instantaneous velocity, acceleration, or jerk. Python is then used again to create a frequency histogram for each of these kinematic variables.

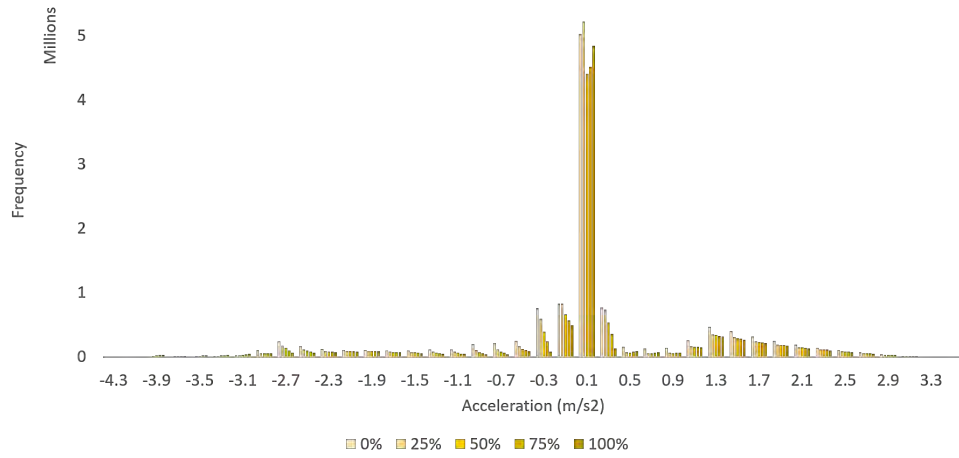
Using the modelling environment outlined in Section 6.2.1 and the data generation methodology outlined above, Figure 24 provides the frequency distribution for velocity, acceleration, and jerk obtained from the model runs. The darkening colour of the columns indicates an increasing CPR;




(a) Velocity frequency distribution histogram for 25% CPR increments.



(b) Frequency histogram providing greater detail in the low velocity results. (c) Frequency histogram providing greater detail in the high velocity results.




(d) Acceleration frequency distribution histogram for 25% CPR increments.




(e) Frequency histogram providing greater detail in the low acceleration results.

(f) Frequency histogram providing greater detail in the high acceleration results.



(g) Jerk frequency distribution histogram for 25% CPR increments.



(h) Frequency histogram providing greater detail in the low jerk results. (i) Frequency histogram providing greater detail in the high jerk results.

Figure 24: Frequency histograms depicting the distribution of velocity, acceleration, and jerk of vehicles during network operation

The frequency histograms in Figure 24 provide several interesting insights into the kinematic performance of mixed fleets and CAVs. Based on visual inspection, the velocity distribution tends to skew to the right as CPR increases. This outcome implies that vehicles on average tend to travel faster as the mixture of CAVs in the fleet increases. The mean value of the fleet velocity between different CAV scenarios increases from $11.70m/s$ in the 0% CAV base case, to $12.67m/s$, $13.41m/s$, $13.68m/s$, and finally to $14.02m/s$ in the 100% CAV case. Though velocity provides a valuable insight into network performance, where a higher average velocity indicates a better performing network, it provides little indication into passenger comfort. A vehicle under constant velocity, regardless of the velocity, experiences no force and hence no discomfort to the passengers. For this reason, it is necessary to assess the impact of CAV integration on acceleration and jerk.

Visually observing the acceleration and jerk distributions in Figure 24, a trend is more difficult to identify. It appears as though increases in CAV penetration result in fleet behaviour showing a higher frequency of acceleration and jerk around $0.1 m/s^2$ and $0 m/s^3$ respectively. The figure also shows a reduction in the frequency of higher and lower acceleration and jerk results, aggregating in the centre of the distribution. The mean value of the different CAV penetrations provides no further clarity, with acceleration having a mean value of

0.9 m/s^2 for all CAV penetrations, and jerk having a mean value of 0 m/s^3 for all CAV penetrations. To gain greater clarity in the results, the following subsection conducts a thorough quantitative assessment.

6.3 ANOVA Assessment of Network Kinematics

Section 6.2.2 used primarily visual inspection to conclude the kinematic impacts of CAVs. Due to the size of the dataset, this approach does not provide a reliable or meaningful means of assessment. For a more in-depth analysis of the impact of changing CPR on network kinematics, an analysis of variance (ANOVA) is used. An ANOVA assessment verifies the presence of a difference between two or more different datasets. The ANOVA assessment is relatively simple, and can only indicate a binary true or false answer to whether there is a similarity between the datasets, or not [Lane, 2007].

6.3.1 Verification of ANOVA Assumptions

The use of an ANOVA assessment requires the verification of three assumptions, which are first explained and verified before the ANOVA assessment is explained and conducted.

Assumption 1: Homogeneity of Variance

The first assumption states that the standard deviation of the populations involved in the assessment is equal. To demonstrate validation of this assumption, refer to Table 3;

CAV Penetration	Standard Deviation		
	Velocity	Acceleration	Jerk
0%	5.96	1.13	0.19
25%	5.68	1.09	0.45
50%	5.10	1.19	0.42
75%	4.85	1.23	0.45
100%	4.47	1.25	0.45

Table 3: The population standard deviations for each dataset involved in the ANOVA assessment.

The standard deviations, while not the same, are considered similar enough to

carry out the ANOVA assessment. The ratio between the standard deviation to the median value of the columns in Table 3 yields a result of 10.7%, 5.0%, and 22.9% for the velocity, acceleration, and jerk, respectively.

Assumption 2: Normality of Population Distribution

The second assumption states that the population distribution of each dataset is normal. While there are numerous ways of establishing normality of a dataset [Ghasemi & Zahediasl, 2012], this thesis uses the D’Agostino’s K^2 Test, which uses skewness and kurtosis as goodness-of-fit metrics. Skewness and kurtosis are calculated using Equation 118 and Equation 119, respectively;

$$g_1 = \frac{m_3}{m_2^{3/2}} = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^3}{(\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2)^{3/2}} \quad (118)$$

$$g_2 = \frac{m_4}{m_2^2} - 3 = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^4}{(\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2)^2} - 3 \quad (119)$$

Where, g_1 is the skewness, g_2 is the kurtosis, m_z is the z^{th} moment of probability about the mean, n is the population size, x_i is the i^{th} sample in the population, and \bar{x} is the population mean.

The skewness and kurtosis values for each CAV penetration are provided in Table 4;

CAV Penetration	Velocity		Acceleration		Jerk	
	Skewness	Kurtosis	Skewness	Kurtosis	Skewness	Kurtosis
0%	-0.58	-1.17	-0.77	3.12	1.52	212.14
25%	-0.89	-0.67	-1.13	5.78	0.59	44.93
50%	-1.17	0.09	-1.32	5.72	0.80	55.64
75%	-1.30	0.49	-1.48	6.07	0.76	54.08
100%	-1.47	1.10	-1.63	6.60	0.81	55.02

Table 4: The skewness and kurtosis for each CPR and kinematic variable in the ANOVA assessment.

Skewness stays well below the critical value of ± 1.96 for a significance level of $P < 0.05$, however, the kurtosis value exceeds this critical value. This method is not always reliable for measuring kurtosis for a dataset greater than

200 samples [Ghasemi & Zahediasl, 2012]. Large datasets generally require a transformation to provide a better estimate of skewness and kurtosis.

The skewness is transformed using;

$$Z_1(g_1) = \delta \operatorname{asinh}\left(\frac{g_1}{\alpha\sqrt{\mu_2(g_1)}}\right) \quad (120)$$

Where, $Z_1(g_1)$ is the transformed skewness, $\mu_2(g_1)$ is presented below in Equation 122, and δ and α are calculated as follows;

$$\begin{aligned} W^2 &= \sqrt{2\gamma_2(g_1) + 4} - 1 \\ \delta &= 1/\sqrt{\ln W} \\ \alpha^2 &= 2/(W^2 - 1) \end{aligned} \quad (121)$$

Where, $\gamma_2(g_1)$ is presented below in Equation 123. $\mu_2(g_1)$ is calculated as;

$$\mu_2(g_1) = \frac{6(n-2)}{(n+1)(n+3)} \quad (122)$$

$\gamma_2(g_1)$ is calculated as;

$$\gamma_2(g_1) = \frac{36(n-7)(n^2+2n-5)}{(n-2)(n+5)(n+7)(n+9)} \quad (123)$$

The kurtosis is transformed using;

$$Z_2(g_2) = \sqrt{\frac{9A}{2}} \left(1 - \frac{2}{9A} - \left(\frac{1 - 2/A}{1 + \frac{g_2 - \mu_1(g_2)}{\sqrt{\mu_2(g_2)}} \sqrt{2/(A-4)}}\right)^{1/3}\right) \quad (124)$$

Where, $\mu_1(g_2)$, $\mu_2(g_2)$, and A are given in Equation 125, Equation 126, and Equation 127, respectively.

$\mu_1(g_2)$ is given by;

$$\mu_1(g_2) = -\frac{6}{n+1} \quad (125)$$

$\mu_2(g_2)$ is given by;

$$\mu_2(g_2) = \frac{24n(n-2)(n-3)}{(n+1)^2(n+3)(n+5)} \quad (126)$$

A is given by;

$$A = 6 + \frac{8}{\gamma_1(g_2)} \left(\frac{2}{\gamma_1(g_2)} + \sqrt{1 + \frac{4}{\gamma_1(g_2)^2}} \right) \quad (127)$$

Where, $\gamma_1(g_2)$ is given by;

$$\gamma_1(g_2) = \frac{6(n^2 - 5n + 2)}{(n + 7)(n + 9)} \sqrt{\frac{6(n + 3)(n + 5)}{n(n - 2)(n - 3)}} \quad (128)$$

Table 5 below provides the transformed skewness and kurtosis values. The values are now significantly closer to the ± 1.96 threshold, with the majority falling within this threshold. The kurtosis values for the jerk datasets still show a high value, indicating that the jerk distributions contain a narrower peak than the traditional long-drawn tails of the normal distribution. Refer to the following table for the transformed values;

CAV Penetration	Velocity		Acceleration		Jerk	
	Skewness	Kurtosis	Skewness	Kurtosis	Skewness	Kurtosis
0%	-1.35	-1.87	-1.75	2.70	2.57	7.69
25%	-1.92	-0.44	-2.34	3.46	1.01	6.96
50%	-2.29	0.74	-2.54	3.44	1.34	7.78
75%	-2.47	1.15	-2.74	3.52	1.25	7.91
100%	-2.70	1.64	-2.94	3.64	1.34	8.19

Table 5: The skewness and kurtosis after being transformed.

Assumption 3: Independence of Observations

The final assumption states that the process generating the data is stochastic and the observations are independent of one another. Since each vehicle in the network is generated through a memoryless process, there is no dependency of data generated from one vehicle to another. This property of microsimulation holds true as the vehicle fleet transitions from human to CAVs. Additionally, the random seed is altered between consecutive runs. For these reasons, the final assumption of the ANOVA assessment process is validated.

6.3.2 Methodology of the ANOVA Analysis Process

The ANOVA analysis uses the right-skewed F-Distribution and begins by defining the null and alternative hypothesis. If the F-value is greater than a critical F-value for a specific level of significance, the null hypothesis should

be rejected. For this assessment, the null hypothesis is that the mean of the distribution for a given kinematic variable for each CPR is statistically the same. The alternate hypothesis is that they are statistically different.

$$\begin{aligned} H_0 : \mu_{x,0\% \text{ CAV}} &= \mu_{x,25\% \text{ CAV}} = \mu_{x,50\% \text{ CAV}} = \mu_{x,75\% \text{ CAV}} = \mu_{x,100\% \text{ CAV}} \\ H_a : \text{means are not equal} \end{aligned}$$

Where, H_0 is the null hypothesis, H_a is the alternate hypothesis, μ is the mean of the dataset, and x is the kinematic variable currently under investigation.

The F-value is calculated using Equation 129;

$$F = \frac{MST}{MSE} \quad (129)$$

Where, MST is the mean sum of squares between datasets, and MSE is the mean sum of squares within datasets.

The mean sum of squares between datasets (MST) is calculated using;

$$MST = \frac{SST}{p-1} \quad (130)$$

Where, SST is the sum of squares between datasets, and p is the number of datasets involved in the analysis.

The sum of squares between datasets (SST) is calculated using Equation 131;

$$SST = \sum_{i=1}^p n_i (x_i - \bar{x})^2 \quad (131)$$

Where, n_i is the number of data points in dataset i , x_i is the mean of dataset i , and \bar{x} is the mean of all data involved in the analysis.

The mean sum of squares within datasets (MSE) is calculated using;

$$MSE = \frac{SSE}{n-p} \quad (132)$$

Where, SSE is the sum of squares within datasets.

The sum of squares within datasets (SSE) is calculated using Equation 133;

$$SST = \sum_{i=1}^p (n_i - 1) S_i^2 \quad (133)$$

Where, S_i is the standard deviation of dataset i .

6.3.3 ANOVA Assessment Results

100,000 data points were extracted from each dataset to form a sample, which was used to conduct the ANOVA analysis. This process was repeated for 500 iterations for each kinematic variable. The F-value for each iteration and each kinematic variable is provided in Figure 25 as a box plot;




Figure 25: The F-value distribution generated from 500 iterations of the ANOVA assessment of 100,000 data points, for each kinematic variable.

Without context and interpretation, the values in Figure 25 provide little information. To determine whether the null hypothesis should be accepted or rejected, the F-value must be compared to a critical F-value. The critical F-value is drawn from the F-distribution, the probability density function of which is provided in Equation 134 [Abramowitz & Stegun, 1948];

$$f(x; d_1, d_2) = \frac{1}{B(\frac{d_1}{2}, \frac{d_2}{2})} \left(\frac{d_1}{d_2}\right)^{\frac{d_1}{2}} x^{\frac{d_1}{2}-1} \left(1 + \frac{d_1}{d_2}x\right)^{-\frac{d_1+d_2}{2}} \quad (134)$$

Where, x is the probability, d_1 is the first degree of freedom ($d_1 = p - 1$), d_2 is the second degree of freedom ($d_2 = n - p$), and B is the Euler integral Beta Function. B is calculated as follows [Abramowitz & Stegun, 1948];

$$B(d_1, d_2) = \int_0^1 t^{d_1-1} (1-t)^{d_2-1} dt \quad (135)$$

Where, $t \mapsto t_+^n$ is the truncated power function, calculated by Equation 136.

$$t_+^n = \begin{cases} t^n & \text{if } t > 0 \\ 0 & \text{if } t \leq 0 \end{cases} \quad (136)$$

Using the F-distribution tables for a one-way ANOVA analysis is a valid and simpler means of calculating the critical F-value. The ANOVA tables require df_1 (4), df_2 (99,995) and the level of confidence (α) to find the corresponding critical F-value. The level of confidence dictates how extreme the results of the ANOVA analysis must be to reject the null hypothesis. Higher α values increase the likelihood of committing “Type I” error, rejecting the null hypothesis when it is true. Table 6 provides the results of the ANOVA analysis for the kinematic parameters, a range of α values, and the associated critical F-value for $d_1 = 4$ and $d_2 = 99,995$ for reference. This table also reports the P-value, which indicates the likelihood that the null hypothesis is correct;

Kinematic Variable	Velocity	Acceleration	Jerk
F-value	3205.812	2.470	0.783
P-Value	0	0.0425	0.536

α	0.001	0.005	0.01	0.05	0.1	0.2
Critical F-value	4.621	3.717	3.321	2.373	1.945	1.497

Table 6: Results of the ANOVA analysis and the critical F-values associated with $d_1 = 4$ and $d_2 = 99,995$.

The consensus on an appropriate P-value and α value is generally 0.05, although the threshold used is generally circumstance dependant. The results in Table 6 indicate that as CPR increases, there is a 0% probability that the velocity distribution of the network stays the same, and a 4.25% and 53.6% probability for the acceleration distribution and jerk distribution, respectively.

The results of the kinematic assessment indicate that as CAV penetration increases, the velocity distribution of the network appreciably changes. However, there is far more evidence to show that acceleration and jerk distribution do not change. This outcome demonstrates that the forces experienced by passengers do not change between human-driven vehicles and CAVs. Consistency in forces, and by extension comfort, demonstrates that the design of the algorithm for the emulation of CAVs is conservative and does not introduce unreasonable or unrealistic driving conditions.


6.4 Sensitivity Assessment of Framework Parameters

The three parameters introduced in this framework are $P_{max} = 0.75$, $P_{min} = 0.25$, and $n = 0.2$. Changing these parameters is not anticipated to have an appreciable impact on traffic performance. This expectation is because the parameters alter acceleration restrictions, and acceleration was already shown in Figure 24d to sit well within the predefined thresholds. The sensitivity analysis in this section is conducted by independently varying each parameter by $\pm 10\%$ and $\pm 25\%$. This results in a total of 13 scenarios when also including the base case. The scenarios are labelled as follows;


- Scenario 1: $P_{max} = 0\%$, $P_{min} = 0\%$ and $n = 0\%$
- Scenario 2: $P_{max} = -25\%$, Scenario 3: $P_{max} = -10\%$
- Scenario 4: $P_{max} = +10\%$, Scenario 5: $P_{max} = +25\%$
- Scenario 6: $P_{min} = -25\%$, Scenario 7: $P_{min} = -10\%$
- Scenario 8: $P_{min} = +10\%$, Scenario 9: $P_{min} = +25\%$
- Scenario 10: $n = -25\%$, Scenario 11: $n = -10\%$
- Scenario 12: $n = +10\%$, Scenario 13: $n = +25\%$

The overall impacts of the sensitivity analysis are not difficult to predict. Increasing the parameters related to acceleration will result in higher network velocities, acceleration rates, and jerks. Without collecting primary data from CAVs, these parameters are not possible to calibrate. But understanding the influence of minor changes in their values on vehicle comfort is valuable.


The analysis is conducted on the same network explained in Section 6.2.1, using the same signalling structure and random seed. The outputted trajectory data is once again converted to CSV using SSAM and analysed using Python. The following histograms show the frequency distribution for the kinematic performance of the network (velocity, acceleration, and jerk), for CPR in 25% increments, resulting in a total of 12 distinct histograms;




(a) Velocity frequency distribution histogram for 25% penetration increment of CAVs.




(b) Greater detail for low velocity results. (c) Greater detail for high velocity results.




(d) Velocity frequency distribution histogram for 50% penetration increment of CAVs.




(e) Greater detail for low velocity results. (f) Greater detail for high velocity results.




(g) Velocity frequency distribution histogram for 75% penetration increment of CAVs.



(h) Greater detail for low velocity results. (i) Greater detail for high velocity results.



(j) Velocity frequency distribution histogram for 100% penetration increment of CAVs.




(k) Greater detail for low velocity results. (l) Greater detail for high velocity results.


Figure 26: Velocity histograms for each sensitivity scenario, for all CPR cases.

Visual inspection of the velocity frequency histograms indicates that, as anticipated, the average velocity in the network shows marginal changes in response to the different sensitivity scenarios. For a 25% CPR, there is a noticeable leftward shift in the histogram for Scenario 4, 8, and 13 (+10% P_{max} , +10% P_{min} , and +25% n). By forcing the vehicles to decelerate at a higher rate, it is apparent that vehicles were more likely to experience a lower velocity. As CPR increased to 50%, 75%, and 100%, a similar observation was found. Although, the difference between these scenarios and the base case was significantly less than in the 25% CPR scenario.


The following figure provides the same histograms for the acceleration results;



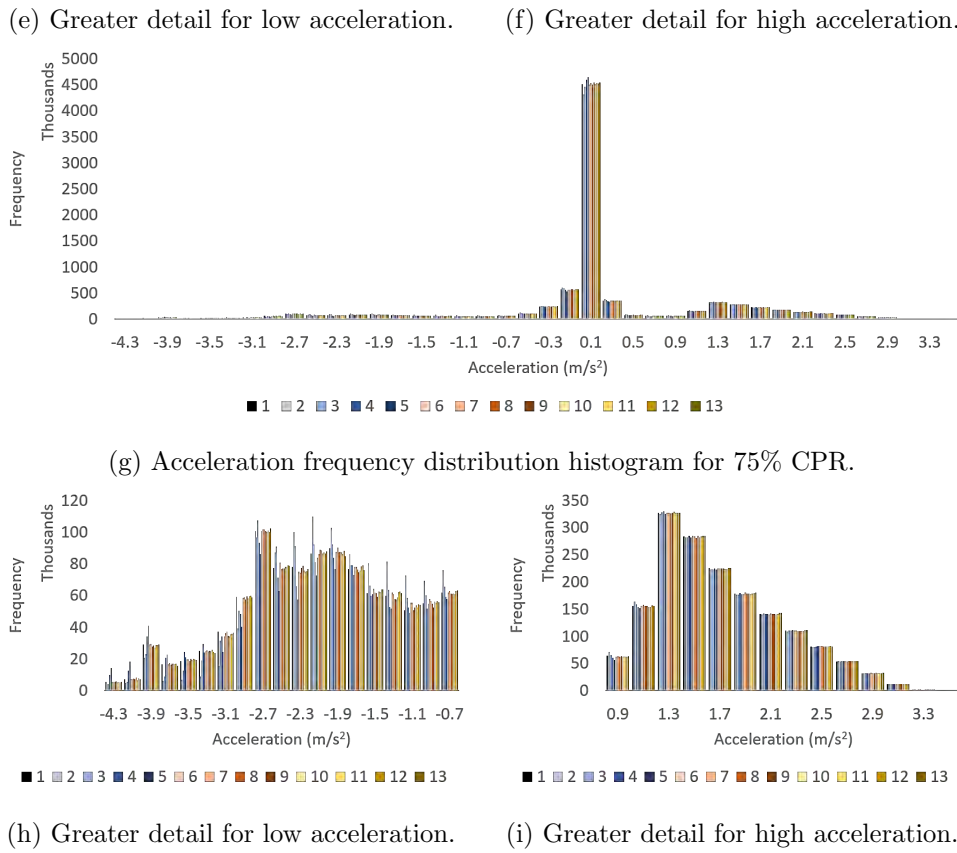


(a) Acceleration frequency distribution histogram for 25% CPR.




(b) Greater detail for low acceleration.




(c) Greater detail for high acceleration.





(j) Acceleration frequency distribution histogram for 100% CPR.




(k) Greater detail for low acceleration.

(l) Greater detail for high acceleration.



Figure 27: Acceleration histograms for each sensitivity, for all CPR cases.

The effects of the parameters sensitivities are more apparent on acceleration than they are on velocity. Figure 27 shows a minor rightward drift for Scenario 4, 8, and 13 (+10% P_{max} , +10% P_{min} , and +25% n). This behaviour is as expected, increasing the parameters directly related to acceleration result in higher accelerations in the network. As CPR increases, the difference in the acceleration profiles between sensitivity scenarios reduces, with the exception of Scenario 2 and Scenario 3. These scenarios reduce the maximum acceleration during regular operation and show that lower acceleration rates are far more frequent than in the base case Scenario 1.

The following figure provides the sensitivity results for the jerk experienced by the vehicles;




(a) Jerk frequency distribution histogram for 25% penetration increment of CAVs.





(b) Greater detail for low jerk results.

(c) Greater detail for high jerk results.




(d) Jerk frequency distribution histogram for 50% penetration increment of CAVs.




(e) Greater detail for low jerk results.

(f) Greater detail for high jerk results.




(g) Jerk frequency distribution histogram for 75% penetration increment of CAVs.




(h) Greater detail for low jerk results.

(i) Greater detail for high jerk results.



(j) Jerk frequency distribution histogram for 100% penetration increment of CAVs.



(k) Greater detail for low jerk results.

(l) Greater detail for high jerk results.

Figure 28: Jerk histograms for each sensitivity scenario, for all CPR cases.

The comments made regarding the velocity and acceleration histograms are also applicable to jerk histogram. The results of the sensitivity analysis were anticipated, and are summarised as follows;

- Increasing n (safety distance incorporated into lane-changing) results in a greater frequency of low velocity, acceleration, and jerk results. This result is anticipated as increasing n forces vehicles to provide larger gaps for cooperative lane-changing, forcing decelerations at a greater rate.
- Increasing P_{max} also results in more lower-range velocity, acceleration, and jerk results. This outcome is in response to the framework imposing greater decelerations, resulting in vehicles attaining lower velocities quicker and increasing the frequency of low-velocity outcomes.
- Increasing P_{min} results in the same outcome as increasing n and P_{max} .

The outcomes mentioned above are based on visual inspection of the frequency histograms. This method does not provide a quantitative analysis, nor does it indicate statistical significance. Despite its limitations, visual inspections provide indications of trend. Table 7 provides the mean of each histogram. The scale of the numbers and the small absolute change between them makes trends difficult to discern. Table 8 provides the percentage change between the mean value of the kinematic variable and the base case (Scenario 1) result;

Scen	Velocity (m/s)				Acceleration (m/s ²)				Jerk (m/s ³)			
	25%	50%	75%	100%	25%	50%	75%	100%	25%	50%	75%	100%
1	12.67	13.41	13.68	14.02	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
2	12.81	13.29	13.59	13.80	0.09	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
3	12.78	13.30	13.65	13.94	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
4	12.57	13.44	13.70	14.09	0.08	0.09	0.09	0.09	0.00	-0.01	-0.01	0.00
5	13.12	13.46	13.76	14.13	0.09	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
6	12.85	13.41	13.68	14.02	0.09	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
7	12.65	13.47	13.65	14.03	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
8	12.82	13.40	13.69	14.03	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
9	12.96	13.38	13.68	14.04	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
10	13.08	13.46	13.70	14.05	0.09	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
11	12.95	13.42	13.70	14.04	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
12	12.87	13.37	13.65	14.03	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00
13	12.47	13.40	13.66	14.02	0.08	0.09	0.09	0.09	-0.01	-0.01	-0.01	0.00

Table 7: The mean value of each kinematic variable distribution.

Scen	Velocity (m/s)				Acceleration (m/s ²)				Jerk (m/s ³)			
	25%	50%	75%	100%	25%	50%	75%	100%	25%	50%	75%	100%
2	1.0%	-0.9%	-0.7%	-1.6%	2.9%	-1.7%	-1.3%	-2.2%	3.8%	-0.9%	2.9%	-2.7%
3	0.8%	-0.8%	-0.2%	-0.6%	0.4%	-0.6%	-1.0%	-0.3%	2.2%	-2.8%	2.0%	1.3%
4	-0.8%	0.2%	0.2%	0.5%	-1.1%	0.6%	0.2%	0.1%	-6.8%	2.5%	4.9%	-2.5%
5	3.5%	0.3%	0.6%	0.8%	1.4%	-1.2%	0.2%	0.5%	7.3%	-3.1%	2.2%	4.2%
6	1.3%	0.0%	0.0%	0.0%	1.4%	0.1%	0.1%	0.0%	4.1%	-1.7%	-0.3%	4.6%
7	-0.2%	0.4%	-0.2%	0.1%	-1.4%	0.8%	-0.5%	0.0%	-4.2%	-0.7%	0.4%	-1.7%
8	1.2%	-0.1%	0.1%	0.0%	1.2%	0.1%	0.3%	-0.2%	0.1%	-3.6%	0.8%	0.8%
9	2.3%	-0.2%	0.0%	0.1%	1.1%	-0.9%	0.2%	-0.2%	7.1%	0.0%	1.3%	-0.4%
10	3.2%	0.3%	0.2%	0.2%	2.1%	0.4%	0.7%	-0.1%	10.9%	1.5%	0.6%	4.3%
11	2.2%	0.1%	0.1%	0.1%	0.6%	0.2%	-0.3%	-0.3%	5.0%	1.3%	0.7%	8.9%
12	1.5%	-0.3%	-0.2%	0.1%	0.2%	-0.2%	-0.4%	0.2%	5.4%	0.2%	0.3%	8.8%
13	-1.6%	-0.1%	-0.1%	0.0%	-0.9%	-0.3%	1.3%	-0.3%	-5.3%	-2.9%	-3.5%	2.1%

Table 8: Relative change in the mean value of the kinematic variable distribution for each sensitivity and CPR, compared to the base case.

Table 8 provides a greater insight into which sensitivity scenarios have a measurable impact on network vehicle kinematics. This table indicates that the network velocity doesn't see a considerable difference when CPR is greater than 25%. At 25% CPR, Scenario 5, 8, 9, 10, 11, and 12 show an apparent increase in the mean velocity. These scenarios correspond to an increase in the three parameters, which results in higher accelerations and higher travel velocities. The results in Table 8 demonstrate that at 50% CPR and greater, the sensitivity cases don't show a major difference against the base case. This result may be attributed to a higher average velocity in the network resulting in vehicles spending less time in the network, which would, in turn, reduce the congestion, lane-changing, weaving, and degree of deceleration required by the remaining vehicles. The investigation in Chapter 7 verifies this outcome.

The skewness of each distribution has been calculated using Equation 118 and presented in Table 9. The skewness demonstrates the lean of the histogram, with a negative skewness value indicating a left lean relative to the mean value. A greater negative value indicates that a distribution is more left-leaning;

Scen	Velocity (m/s)				Acceleration (m/s ²)				Jerk (m/s ³)			
	25%	50%	75%	100%	25%	50%	75%	100%	25%	50%	75%	100%
1	-0.89	-1.17	-1.30	-1.47	-1.13	-1.32	-1.48	-1.63	0.59	0.80	0.76	0.81
2	-0.92	-1.11	-1.25	-1.36	-1.07	-1.31	-1.48	-1.66	0.66	0.74	0.70	0.71
3	-0.92	-1.12	-1.29	-1.43	-1.10	-1.30	-1.48	-1.63	0.69	0.76	0.71	0.76
4	-0.87	-1.18	-1.31	-1.50	-1.21	-1.37	-1.54	-1.68	0.56	0.80	0.69	0.79
5	-1.04	-1.19	-1.34	-1.53	-1.19	-1.44	-1.62	-1.76	0.83	0.83	0.70	0.76
6	-0.94	-1.16	-1.30	-1.47	-1.10	-1.31	-1.48	-1.64	0.64	0.84	0.72	0.76
7	-0.89	-1.18	-1.29	-1.47	-1.15	-1.32	-1.49	-1.65	0.55	0.85	0.73	0.82
8	-0.94	-1.16	-1.31	-1.47	-1.12	-1.32	-1.48	-1.64	0.64	0.75	0.72	0.79
9	-0.98	-1.16	-1.30	-1.48	-1.10	-1.33	-1.50	-1.65	0.70	0.74	0.71	0.78
10	-1.03	-1.18	-1.31	-1.48	-1.10	-1.30	-1.48	-1.63	0.76	0.80	0.71	0.80
11	-0.98	-1.17	-1.30	-1.48	-1.11	-1.31	-1.48	-1.63	0.69	0.83	0.73	0.77
12	-0.95	-1.15	-1.29	-1.47	-1.12	-1.31	-1.48	-1.63	0.62	0.77	0.72	0.75
13	-0.83	-1.16	-1.29	-1.47	-1.17	-1.32	-1.49	-1.64	0.49	0.81	0.75	0.77

Table 9: The skewness of the kinematic variable distributions, for each sensitivity scenario and CPR.

The percentage change in skewness compared to the base case scenario is presented in Table 10;

Scen	Velocity (m/s)				Acceleration (m/s ²)				Jerk (m/s ³)			
	25%	50%	75%	100%	25%	50%	75%	100%	25%	50%	75%	100%
1	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
2	3.5%	-4.7%	-4.0%	-7.5%	-4.8%	-0.7%	0.1%	1.7%	11.7%	-7.4%	-7.0%	-12%
3	3.0%	-3.6%	-0.9%	-2.5%	-2.4%	-1.3%	-0.2%	-0.2%	17.0%	-4.8%	-5.8%	-6.3%
4	-3.0%	1.0%	1.2%	2.4%	6.9%	3.8%	3.8%	3.1%	-4.7%	-0.2%	-8.1%	-2.6%
5	15.9%	2.3%	3.4%	4.0%	5.8%	9.3%	9.7%	7.9%	41.2%	2.9%	-7.6%	-6.8%
6	5.4%	-0.1%	-0.1%	0.0%	-2.3%	-0.8%	0.2%	0.4%	8.8%	4.8%	-4.9%	-6.6%
7	-0.9%	1.5%	-0.6%	0.2%	1.7%	-0.2%	0.4%	0.7%	-6.7%	5.9%	-3.9%	1.0%
8	4.9%	-0.4%	0.7%	0.1%	-1.0%	0.0%	0.1%	0.1%	9.0%	-6.0%	-4.9%	-3.6%
9	9.8%	-0.6%	0.4%	0.5%	-2.0%	1.1%	1.7%	0.8%	19.9%	-7.8%	-6.2%	-4.3%
10	15.0%	1.4%	0.6%	1.0%	-2.5%	-1.5%	0.3%	-0.5%	29.4%	-0.7%	-5.6%	-2.1%
11	9.8%	0.7%	0.5%	0.7%	-1.4%	-0.5%	-0.2%	-0.4%	18.3%	3.8%	-3.6%	-5.1%
12	6.3%	-1.2%	-0.6%	0.2%	-0.8%	-0.9%	0.1%	0.0%	5.9%	-3.8%	-4.7%	-7.6%
13	-7.4%	-0.5%	-0.4%	-0.1%	3.7%	0.0%	0.6%	0.6%	-16%	0.3%	-1.0%	-5.1%

Table 10: Percentage change in the skewness, for each sensitivity scenario compared to the base case.

The colour distribution in Table 10 closely resembles the colour distribution

of Table 8. The increase in network velocity in Scenario 5, 8, 9, 10, 11, and 12, have already been mentioned numerous times. Table 8 indicates that for 25% CPR, there is a general rightward shift to the lean of the frequency. However, this lean shifts leftward as CPR increases further. The expected cause of this lean shift is that as CAV behaviour increases the average velocity of the network, this results in an average increase in acceleration and corresponding increases in jerk. However, at CPR greater than 25%, the saturation level of the network falls significantly, enough to mitigate congestion. The reduced levels of congestion expectedly reduce weaving, lane-changing, interactions between vehicles, and expectantly reduce the required acceleration and jerk. Although, this explanation is explored in greater detail in Chapter 7.

Finally, an ANOVA analysis of 500 iterations is conducted on each scenario, comparing the base case scenario to the sensitivity scenarios in-turn. Each iteration randomly samples 100,000 data points from the dataset. This methodology means that $df_1 = 1$ and $df_2 = 99,998$. The null hypothesis is that the mean value of the sensitivity frequency distribution is the same as that of the base case scenario, with the alternative hypothesis being that they are different. At a level of significance of $\alpha = 0.05$, the critical F-value is 3.842. The mean F-Value of each scenario is provided in Table 11;

Scen	Velocity (m/s)				Acceleration (m/s ²)				Jerk (m/s ³)			
	25%	50%	75%	100%	25%	50%	75%	100%	25%	50%	75%	100%
2	86.91	20.32	1.44	5.48	1.82	0.97	1.24	1.08	0.94	0.99	0.99	0.91
3	2.34	13.39	0.95	1.73	1.01	0.82	0.99	0.97	0.92	0.88	0.92	1.07
4	295.6	3.15	1.08	1.12	1.73	1.05	1.04	1.04	0.98	0.97	1.02	0.88
5	333.5	4.50	2.43	2.72	1.18	1.08	1.17	1.24	0.94	0.97	0.98	0.96
6	63.43	6.35	3.64	1.27	1.15	1.03	1.12	0.98	1.08	0.97	0.86	1.06
7	43.24	34.07	2.02	1.06	1.23	0.87	0.99	0.92	1.08	0.88	0.94	1.04
8	21.84	1.08	2.87	1.09	1.10	0.99	0.99	0.99	0.95	1.11	0.87	0.93
9	165.0	5.19	1.43	0.95	1.11	0.94	0.92	1.06	1.04	0.89	0.95	0.99
10	396.2	17.29	3.40	1.11	1.49	0.88	1.01	0.91	0.97	0.98	1.03	0.96
11	124.5	2.86	0.98	1.01	1.13	0.97	0.97	0.93	0.94	1.12	0.94	1.07
12	110.2	3.12	1.32	1.80	1.09	0.89	0.99	1.13	0.97	0.89	0.99	0.99
13	420.6	1.07	3.07	0.99	1.53	1.01	0.94	0.94	0.91	0.92	0.90	1.03

Table 11: The F-Value for the ANOVA analysis comparing the sensitivity scenario with the base case. The critical F-value is 3.842 for $\alpha = 0.05$

In Table 11, red indicates scenarios where the null hypothesis should be rejected and yellow represents scenarios where it should be accepted. The ANOVA analysis confirms that changing the parameters P_{max} , P_{min} and n by $\pm 10\%$ or $\pm 25\%$ has minimal impacts on vehicle behaviour. The network velocity tends to increase as these parameters are increased. Still, the ANOVA analysis, and visual inspection of the frequency histograms, indicates that the underlying kinematics of acceleration and jerk in vehicle operations experiences negligible (statistically insignificant) change. This finding indicates that despite the emulation framework replicating expected CAV behaviour and improving network performance, it does without sacrificing passenger comfort.

6.5 Summary and Conclusion

In this chapter, a detailed kinematic and sensitivity analysis of the developed CAV modelling algorithm was conducted. The chapter started with a qualitative confirmation of platoon stability through the generation of kinematic-time plots for a platoon of ten vehicles conducting a lane change (Section 6.1). It then conducted a qualitative assessment of network kinematics while incrementally transitioning the fleet to CAV (Section 6.2 and Section 6.3). Finally, this chapter concluded with a sensitivity analysis of the parameters introduced in the developed algorithm (Section 6.4). The kinematic-time plots attest to the stability of the emulation algorithm, during both regular car-following and lane-changing. Acceleration and jerk rates do not surpass critical predefined values, nor does erratic behaviour and deviations in vehicle motion occur during sudden deceleration events. These results qualitatively demonstrate the algorithms stability. The quantitative results of the network kinematic assessment indicate that as CAV penetration increases, the velocity distribution of the network appreciably changes while acceleration and jerk distributions stay consistent. This outcome demonstrates that the forces experienced by passengers do not change between human-driven vehicles and CAVs. Consistency in forces, and by extension comfort, demonstrates that the design of the algorithm for the emulation of CAVs is conservative and does not introduce unreasonable or unrealistic driving conditions, providing it further validation. Finally, the ANOVA analysis confirms that changing the parameters P_{max} , P_{min} and n by $\pm 10\%$ or $\pm 25\%$ has minimal impacts on vehicle behaviour. This finding indicates that by design of the algorithm, its replication of passenger safety and comfort is not contingent on coincidentally calibrated parameters.

7 An Investigation into the Implications of CAVs on Isolated Intersection Performance

Abstract: Intersections form a major component of road transport systems. Priority-based navigation and traffic signals render the capacity of an intersection significantly lower than the link section, causing them to act as arterial network bottlenecks. This chapter assesses the impact of CAV penetration on traditional isolated intersections. The investigated intersections include signalised intersection with dedicated turn, signalised intersection with filtered turn, roundabout intersection, and priority give-way intersection. The intersections are investigated in the Vissim microsimulator using the developed CAV control algorithm. Modelling results demonstrate that throughput increases by up to 110%, delay reduces by up to -63%, queue length decreases by up to -28%, and average speed increases by up to 130%. However, each intersection type did not improve to the same degree, raising the concern that disproportionate enhancements to parts of the network may redistribute the location of bottlenecks and pinch points. Additionally, this chapter investigates the formation and movement of platoons contingent on intersection green time. The results indicate that higher CAV benefits would better from lower green and cycle times, providing each signalised movement with a higher turnover rate.

CAVs verge on becoming household objects, as hardware and computational capabilities meet the safety and performance expectations held for these vehicles. CAV technology could revolutionise transportation as it promotes car-sharing and the notion of MaaS [Hietanen, 2014]. On the other hand, CAVs may increase the utility of a private vehicle, whereby in-vehicle travel time is better utilised, and private vehicle mode share and congestion increases [Davidson & Spinoulas, 2015]. Government regulation, policy and intervention will strongly direct the development and implementation of CAVs [Fajardo et al., 2011]. However, an implementation framework requires development before the transport network is ready to cater to these vehicles. Well informed policy-related decisions to develop a sustainable future are difficult to make due to the lack of understanding regarding future CAV behaviour and lack of structure surrounding CAV implementation.

This chapter addresses uncertainties surrounding the intersection impacts of

CAVs, with a specific focus on cooperation and altruistic behaviour. The literature presents a range of algorithms used to emulate CAV behaviour and assess intersection performance. Although, these attempts either emulate CAV behaviour through a rudimentary change in commercial software parameters, or a non-complete behavioural emulation algorithm focusing on part of vehicle behaviour in microsimulation.

The remainder of this chapter is structured as follows; Section 7.1 contains a brief literature review of recent attempts at emulating CAV behaviour and assessing intersection performance. Section 7.2 introduces the microsimulation environment used for experimentation and the structure of the tests. Section 7.3 and Section 7.4 provide the results and discussion, respectively. Finally, Section 7.5 provides a brief summation of the work and the conclusions drawn.

7.1 Literature Review of CAV intersection Modelling

The literature proposes numerous means of evaluating the intersection performance of mixed fleets and CAV networks. Attempts vary in complexity with some opting to change commercial microsimulator variables, and others developing complex intersection control mechanisms. This section provides a brief explanation of the unique and novel attempts at quantifying the impact of CAVs on intersection performance.

Yang Study

Yang investigated the network and intersection impacts of CAVs by integrating them into the four-step modelling process [Yang, 2017]. The four-step modelling process starts by segregating the network into zones and determining how many trips each zone produces (trip production). Next, the number of trips each zone attracts is determined (trip distribution), resulting in the formation of an OD matrix. The OD matrix is subdivided into a matrix for each available transit mode (mode choice). Finally, the OD matrix is applied to the network in a traffic assignment process to determine link flows.

Yang incorporated CAVs in trip production and trip generation by multiplying existing generation and attraction models by a factor 0.05. The existing models produce (or attract) trips contingent on household size and household

income, meaning that zones with greater household sizes or incomes are more likely to produce more CAV trips. CAVs were introduced into mode choice by providing them with a separate utility function used in the nested logit structure. The utility function for CAVs is as follows;

$$U_{DA} = \frac{c_{ivt}t_{sov}\alpha + c_{term}T_{ij} + c_{opcost}D_{sov}C_{op} + c_{pcost}(Toll_{sov} + c_{wkp}(i))}{c_{na}} \quad (137)$$

Where, U_{DA} is the driver utility, c_{ivt} , t_{sov} , and α are the coefficient, actual, and reduction rate, for in-vehicle travel time, respectively. c_{term} and T_{ij} is the terminal time coefficient, and terminal time from zone i to zone j , respectively. c_{opcost} , D_{sov} , and C_{op} are the auto operation coefficient, travel distance, and auto operation cost, respectively. c_{pcost} , $Toll_{sov}$ and $c_{wkp}(i)$ are the parking cost coefficient, toll cost, and parking cost in zone i , respectively.

During trip assignment, the CAVs differed from human vehicles through a 50% reduction in in-vehicle travel time, parking fees, and vehicle reaction time, also leading to a proportional decrease in headway and increase in capacity.

Using this modelling approach, Yang demonstrated that a 10% CPR leads to a -4.22% reduction in travel time and a +4.86% increase in travel speed. As CPR increases to 90%, these improvements become -22.15% in travel time and +32.12% in travel speed. Yang's modelling also demonstrated that CAVs induce more trips and cannibalise trips from competing mass transit modes. A 10% CPR results in a -2.53% change in transit trips and a +0.27% increase in total trips. These changes increase to -10.16% reduction in transit trips and +4.51% increase in total trips as CPR approaches 90%.

Yang's study indicated the potential for CAVs to induce more vehicles into already congested environments. For this reason, it is critical to understand what CAV behaviour implies for the arterial and local road environment. Whether traditional intersections will be able to cater for the increased demand must be determined to plan for a mode of transit that will quickly dominate the transit fleet due to its inherent attractive qualities. This chapter verifies the improvements found by Yang in a macrosimulation environment by emulating agent interactions.

Patel et al., Study

Patel et al., used a multi-class cell transmission model (CTM) to investigate how reduced reaction times affect travel time and congestion on highly congested arterial roads and freeway networks [Patel et al., 2016]. The speed of vehicles in the network was a function of free-flow speed, capacity, and backward wave propagation, given by;

$$u(k_1, \dots, k_M) = \min\{u^f, \frac{q^{max}(\frac{k_1}{k}, \dots, \frac{k_M}{k})}{k}, w(\frac{k_1}{k}, \dots, \frac{k_M}{k})(k^{jam} - k)\} \quad (138)$$

Where, $u(k_1, \dots, k_M)$ is the speed of different vehicle classes M , u^f is the free flow speed, $q^{max}(k_1/k, \dots, k_M/k)$ is the capacity function given the class proportion k_M/k , $w(k_1/k, \dots, k_M/k)$ is the backward wave speed function of class proportion k_M/k , and k^{jam} is the jam density.

Flow in the cell transmission model is modelled as;

$$y_i^m(t) = \min\{n_{i-1}^m(t), \frac{n_{i-1}^m(t)}{n_{i-1}(t)} Q_i(t), \frac{n_{i-1}^m(t)}{n_{i-1}(t)} \frac{w_i(t)}{u^f} (N - \sum_{m \in M} n_i^m(t))\} \quad (139)$$

Where, $y_i^m(t)$ is the transition flow of class m from cell i at time t , $n_{i-1}^m(t)$ is the number of vehicle, and $Q_i(t)$ is the cell capacity.

Link capacity is given by;

$$q^{max} = \frac{u^f}{u^f \sum_{m \in M} \frac{k_m}{k} \Delta t_m + l} \quad (140)$$

Where, Δt_m is the reaction time of class m , and l is the vehicle length.

Finally, the backward wave speed is calculated using;

$$w = \frac{l}{\sum_{m \in M} \frac{k_m}{k} \Delta t_m} \quad (141)$$

The model was implemented on two arterial road networks, a freeway environment and a model replicating downtown Austin. On the arterial network, CAVs used a priority reservation-based system to traverse intersections and conflict points. The modelling indicates that on the first arterial network

where intersections are further apart, the use of a reservation-based protocol for intersection control reduced travel times. However, on another arterial network, the reservation-based protocol was only beneficial at 50% demand. Once demand was increased to 75%, the traditional priority signalling scheme was found to lower travel times. The authors found that shorter CAV reaction times (50% less) were responsible for the improvement in the first environment. They also found that suboptimal reservations made by the controller in the second environment due to the close proximity of adjacent intersections resulted in worsening performance.

Prior to saturation, the arterial roads models showed that at 25% CPR, total travel time reduced by as much as 50%. The improvements found in the freeway environment were more modest, though, with better performance in the higher CPR scenarios. Under saturated conditions (100%) demand, a 72% reduction in travel times was seen. The less drastic improvements in freeway travel times were attributed to the fact that freeways do not contain signal delays, and so can not take full advantage of the reservation controller on onramps. Results on the downtown Austin model showed that intersection reservations reduced travel times by 55%, and reduced reaction times resulted in a travel time reduction of 78%.

Patel et al., indicate that CAV network integration has significant potential travel time implications. However, the inclusion of a reservation-based centralised controller for intersections may be inappropriate for evaluating short term intersection effects of CAVs. While such an advancement may be the eventual conclusion to the evolution of intersections, its implementation in the short term while this technology is untested and in its infancy is extremely unlikely, primarily due to safety and logistical concerns. For this reason, a robust modelling exercise of traditional intersection control schemes with CAV operation should be conducted, as is presented in this chapter.

Guler et al., Study

Guler et al., used an enumeration based approach to minimise intersection delay [Guler et al., 2014]. The approach modelled two one-way streets and equipped certain vehicles with CAV hardware. A signal controller enumerated

through different release patterns of vehicles into the intersection in accordance with an objective function. The study considers the impacts of platooning and adaptive signalling regimes, while also varying the demand and the CPR.

The algorithm works as follows. All vehicles currently in the intersection zone are placed in a set. The intersection controller then enumerates through all possible release combinations of vehicles through the intersection. The departure pattern is that which results in the minimum delay. Equation 142 is used to determine the penalty (time taken) for a vehicle to traverse the intersection. This expression appropriately incorporates the effects of vehicle platooning on vehicle departures and delay;

$$P_{n,k} = \max\left(\frac{s}{v}, \frac{-a\frac{1}{\rho_{max,m}}(O_{n,k} - 1) + \sqrt{(a\frac{1}{\rho_{max,m}}(O_{n,k} - 1))^2 + 2as}}{a}\right) \quad (142)$$

Where, $P_{n,k}$ is the penalty for vehicle n in departure pattern k , s is the length of the conflict zone, v is the assumed intersection crossing speed, a is the vehicle acceleration, $\rho_{max,m}$ is the saturation flow rate of approach m , and $O_{n,k}$ is the number of successive vehicles that have departed from the current approach in the current platoon.

Experimentation scenarios had a demand of $1000veh/hr$, $1500veh/hr$, and $2000veh/hr$. Demand was distributed over the two approaches, where a perfectly balanced demand would have a demand ratio of one, and the ratio would approach zero as the demand became weighted to one approach. The authors tested a range of intersection control mechanisms, including first-in-first-out, a traditional signalling scheme, and their developed minimum delay algorithm. The results of the modelling found that in both the low demand and high demand scenarios, the minimum delay algorithm far outperformed the first-in-first-out approach, leading to delay reduction of over 50% in both the balanced and unbalanced demand scenario. However, when compared to a fixed-time signalling scheme, the minimising delay algorithm showed little improvement, usually performing on par with a traditional signalling scheme.

While the minimum delay algorithm showed improvement over the first-in-first-out approach, this algorithm is not an appropriate means of intersection control. The first-in-first-out approach is inherently flawed in mixed-fleet en-

vironments because the signal control must provide a minimum green time to account for human reaction time and error. When the reservation table for a first-in-first-out approach repeatedly alters between approaches, platoons are unable to form and so the minimum green time must be repeated. However, the minimise delay algorithm prioritises a platoon depending on wait time, in much the same way as actuated signalling already does, accounting for why the authors observed little to no benefit from adopting this algorithm as compared to traditional intersection signals.

Le Vine et al., Study

Le Vine et al., investigated the relationship between rider comfort and intersection capacity [Le Vine et al., 2015]. The investigation developed a microsimulation environment in which the authors altered vehicle driving parameters to measure the impact on capacity. The study area consisted of a 4-arm approach signalised intersection with a free flow speed of 50km/hr . Signals used 90s cycle times with two phases. Traffic demand on all approaches was divided in the ratio 1:3:1 for the left turn, through movement, and right turn respectively. A traffic density of 20s/veh ($250\text{veh/hr/approach}$) used in the study found an intersection capacity of 1793veh/hr .

The base case intersection had an average vehicle delay of 20s . The first modelling scenario has a CPR of 25%, and CAV kinematics were constrained to that similar to a Light Rail vehicle (1.34m/s^2 longitudinal acceleration and 1.47m/s^2 lateral acceleration). Under this scenario, the average vehicle delay increased by 5%. When CAV kinematics were altered to match that of Heavy Rail (0.58m/s^2 longitudinal acceleration and 0.49m/s^2 lateral acceleration), delay increased by 36%. The delay increases corresponded to capacity decreases of 4% and 18%, respectively for the two scenarios.

This study emulated CAV behaviour in Vissim by changing vehicle acceleration rates and did not attempt to account for many of the other anticipated benefits of CAVs such as platooning and quicker reaction times. By decreasing acceleration, it is not difficult to see why intersection delay increased and capacity decreased. Although, this study does raise critical concerns regarding CAV operation. The majority of studies conclude that CAV penetration will

lead to more efficient driving and higher road capacity. Le Vine et al., emphasise that if CAVs are programmed more conservatively, and treat occupants as passive passengers (like on public transport), then acceleration rates on par with traditional public transport vehicles may not lead to intersection improvements. A conservative driving fleet could reduce road capacity, contrary to the belief of much of the literature.

Elhenawy et al., Study

Elhenawy et al., proposed a completely autonomous and unsignaled intersection that negotiated space between vehicles based on game theory and Nash's equilibrium. [Elhenawy et al., 2015]. A non-zero non-cooperative game was played where the drivers have the option to either yield or not yield. It is played with two players, where a player anticipates the behaviour of vehicles conflicting with its trajectory. The payoff matrix is standard, where the non-yielding vehicle is considered the winner, both vehicles yielding is considered a tie, and both vehicles not yielding is considered a loss. Each vehicle in the game can accelerate, decelerate, or continue at a constant velocity.

The game is structured as follows. Once each vehicle is within 200m of the centralised intersection controller, it provides its kinematic details to the controller. The controller then uses the values of the closest vehicle to the intersection to determine the set of feasible actions that all other vehicles can perform. The controller creates a game matrix for the four vehicles and simulates each possible set of actions. A payoff table provides the value for each outcome. Nash's equilibrium is a solution to the game matrix which satisfies the following equation;

$$\forall n, x_n \in S_n : f_n(x_n^*, x_{-n}^*) \geq f_n(x_n, x_{-n}) \quad (143)$$

Where, n is the specific player with action x_n , S_i is the set of all actions of player n , $f_n(x_n, x_{-n})$ is the score of player n playing action x_n and player $-n$ playing action x_{-n} , and x^* is the solution to Nash's equilibrium.

Elhenawy et al., implemented the algorithm in a single-lane four-way-approach intersection with a speed limit of 40km/hr. After conducting 30,000 simulations, the authors found that on average, the algorithm reduced travel time

by 49% and delay of 89%. This study indicates the potential upper bound for travel time savings with a 100% CPR. This approach has significant potential to coordinate adjacent intersections and calculate a globally optimal solution to a game played at the network level. However, it is limited in the short-term as it can not be adopted for a mixed fleet environment.

Sharon and Stone Study

Sharon and Stone extended the AIM protocol to hybrid AIM (H-AIM) [Sharon & Stone, 2017], as AIM was only effective at CPR greater than 90%. Alternatively, H-AIM demonstrated intersection performance benefits from as little as 1% CPR. The main difference between the AIM and H-AIM protocol is that when trajectories of vehicles overlap, the AIM protocol immediately denies the request of the second vehicle. H-AIM will consider the request and only immediately deny it if the second vehicle conflicts with a human. To identify a human vehicle, the intersection controller manually counts the number of vehicles on a given lane and then compares that count to the number of reservation requests received. If the count is greater than the reservation table, then the controller assumes the difference is comprised of human vehicles.

Vehicles are assigned a turning policy based on whether they are human or CAV. A turning policy identifies all safe movements for a particular vehicle class. For example, it would be unwise to permit a human vehicle in the left lane to turn right, whereas a CAV may be able to safely and reliably do so. Additionally, the H-AIM protocol prevents the allocation of a safe and unsafe policy to a pair of human and CAVs, as this would again place the human vehicle in dangerous situations.

Sharon and Stone applied the H-AIM protocol to a four-way intersection with a travel speed of $25m/s$. Results from their modelling indicate that during low traffic levels ($\leq 300veh/hr$) and high traffic levels $\geq 700veh/hr$, using safe turning policies results in worsening performance for human vehicles. The only benefit was brought to human vehicles when the most conservative and safest lane policy was used, which restricted each lane to a single unshared turning movement. CAVs, on the other hand, benefited from adopting an unsafe turning policy, regardless of the CPR.

The H-AIM protocol showed benefits over the AIM protocol when the most conservative lane policy was used, where each lane has a dedicated and unshared turning movement. The practicality of such a solution should be questioned. Dedicated lanes with unshared movements are only viable with balanced traffic flow and balanced phase green-time. Once the balance is disrupted, excessive queueing will begin to arise in the oversaturated movement, which would have otherwise been distributed over many lanes. For this reason, the H-AIM protocol may not be appropriate for use in most intersections.

7.2 Experimentation Methodology

A review of the literature indicates that prior work has focused on sophisticated implementations of CAV infrastructure to evaluate the impact of CAV behaviour on intersection performance. The decentralised control of an intersection is only feasible at a 100% penetration, which is many decades away. Decentralised controllers and optimisation algorithms fail to take into account the economic impacts of developing, manufacturing, and retrofitting the infrastructure with the required hardware. This chapter addresses the implications of placing CAVs into the currently used intersection infrastructure and transport systems.

This chapter addresses a gap in the literature by retaining existing signalling and priority structures, but transitioning the fleet gradually to autonomous. The chapter critically investigates the impact of CAV behaviour on intersection performance, providing a detailed assessment of metrics such as delay, throughput, queue length, and average speed. An investigation is conducted on a range of intersections such as a signalised intersection with a filtered right turn, signalised intersection with a dedicated right turn, roundabout, and priority junction. This chapter then assesses the impact of vehicle platooning and signal green-time on intersection throughput. The remainder of this section provides in greater detail the structure of the experimentation.

7.2.1 Experimentation Structure

The investigation is conducted on four individual intersections, refer to Section 7.2.2 for more detail about each intersection. For each model environment, the capacity is determined by incrementally loading the network and observing when throughput ceases to increase with increased demand. Once capacity is

found, 250 iterations of microsimulation are conducted in each environment to develop a Demand vs Throughput relationship. This relationship is required to find intersection capacity.

Bias is avoided by using a stochastic and memoryless Monte Carlo method. This method randomly generates a random seed RS where $RS \in [1, 9999]$, and demand where Degree of Saturation (DoS) $DoS \in [0, 1.1]$ for each iteration. This process is repeated for each 20% increment of CPR, providing a comprehensive image of the intersection capacity, and is outlined in Figure 29;

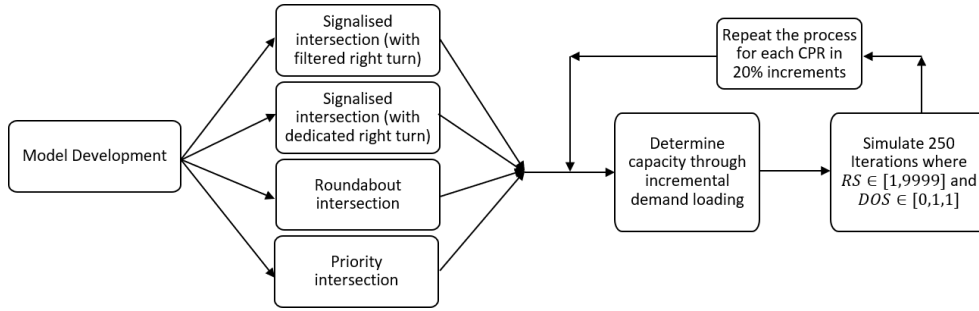



Figure 29: The stochastic and memoryless Monte Carlo process used to determine the capacity of each intersection, as a result of increasing CPR.

7.2.2 Modelling Environment

The modelling environments consist of a signalised intersection with a filtered right turn, a signalised intersection with a dedicated right turn, a roundabout section, and a priority junction. The remainder of this subsection provides greater detail about each of these environments.


The signalised intersection with a right turn filter (Figure 30a) contains two lanes per approach. The left lane shares the through and left-turn movement, and the right lane shares the through and right-turn movement. Each approach arm is $200m$ in length, but the intersection node for data collection purposes extends $100m$ beyond the intersection stop line. The signalling structure uses a split structure approach (Figure 30b) with the right turn forced to filter through the opposing movements. The demand at each approach is in the ratio 2:3:1 between the left turn, through movement, and right turn, respectively. Demand is generated at $50km/hr$, the speed limit of the environment. Figure 30 provides the environment geometry and phasing structure;



(a) Geometric layout. (b) Phasing arrangement. Green and yellow arrows indicate priority and giving way movements, respectively.

Figure 30: Geometric layout and signal phasing for the signalised intersection with a filtered right turn.

The signalised intersection with a dedicated right turn (Figure 31a) contains three lanes per approach, with each turning movement (left, through, and right) having a dedicated lane. The signalling structure differs from the previous intersection in that it uses a double-diamond phasing scheme (Figure 31b), where the right turn at each approach is given dedicated priority green-time in the signal cycle. Other elements of this environment including the demand ratio, speed limit, data collection node, and approach arm length, are the same as the previous environment. The environment geometry and phasing structure are provided in the following figure;




(a) Geometric layout. (b) Phasing arrangement.

Figure 31: Geometric layout and signal phasing for the signalised intersection with a dedicated right turn.

To retain generality in the results, critical consideration has been given to designing the experiments for this investigation. The approach arms of a signalised intersection operate independently, meaning that an intersection


with more approach arms will have an increased number of turning movements and corresponding conflict points. Consider the following conflict diagrams for a two-way, three-way, and four-way intersection;



(a) Two-way intersection contain 0 conflict points. (b) Three-way intersection contain 6 conflict points. (c) Four-way intersection contain 24 conflict points.


Figure 32: Number of conflicts points in different intersection arrangements.

The impact of increasing the number of conflicts is that the phasing arrangements must be redesigned to minimise conflicts in any given phase. Increasing the number of phases while maintaining a set cycle time means that each phase now has a lower green time. Therefore the impact of intersection geometry at signalised intersections is the resulting change to phase green times. Consider the following phasing structures appropriate for each intersection.



(a) Two-way phasing diagram.

(b) Three-way phasing diagram.



(c) Four-way phasing diagram.

Figure 33: The phasing diagram of different geometric arrangements, indicating that greater conflict points result in more phases.

The impact of changing phase times is a critical component of this investigation and is discussed in greater detail towards the end of this subsection. Other lane allocations also exist in intersection configurations. For example, Melbourne, Australia has provisions for “hook-turns” and “p-turns”. The hook-turn lane allocation places a vehicle into the left-most lane, forces them to wait for conflicting traffic to clear, then permits them to conduct a right turn across all lanes in the intersection. The p-turn occurs at an intersection where a right turn is banned. A vehicle is forced forward towards the adjacent intersection, where they conduct a u-turn, travel back to the previous intersection, and conduct a left turn from the opposing direction. These lane allocations are rare and not permitted in NSW due to their inherently more dangerous configuration. For this reason, they have been excluded from the investigation and the remaining signalised lane allocation configurations reflect the most commonly used in NSW.

The roundabout intersection (Figure 34), contains two lanes per approach, with the left lane sharing the left and through movement, and the right lane sharing the through and right movement. Each approach arm is 200m in length, and the intersection node for data collection extends 100m beyond the intersection stop line. The demand at each approach is in the ratio 2:3:1 between the left turn, through movement, and right turn, respectively. No U-Turning vehicles are included in this environment. Demand is generated into the network at 50km/hr, but the speed is reduced to 30km/hr as vehicles circulate the roundabout, increasing back to 50km/hr as they exit. The geometry is as follows;




Figure 34: Geometric layout of the roundabout intersection.

The priority junction (Figure 35), contains a major East-West movement, and minor North-South movement. The Northern and Southern approach arms are forced to yield to the major movement. The North and South approach allows both left and right turn, with a through movement possible by making a left turn followed by a right turn when travelling southbound, vice versa when travelling northbound. Other elements of this environment, including the demand ratio, speed limit, data collection node, and approach arm length, match the previous environments. The environment geometry is as follows;




Figure 35: Geometric layout of the priority intersections.

Finally, the environment used to assess the impact of green-time and CAV platooning on throughput is provided in Figure 36. This environment is a $1km$ approach arm with a signal used to form a queue. Demand is effectively infinite, as vehicles are produced into the network at the maximum loading rate permitted by Vissim. The cycle time is $180s$ and alternates between a green phase ranging from $5s$ to $80s$, and an off phase for the remainder of the cycle time. The geometry is as follows;



Figure 36: Geometric layout for the environment used to assess green time and platooning on intersection performance.

For this experiment, the phase time is set from $5s$ to $80s$ in $1s$ increments, and CAV penetration was increased from 0% to 100% in 20% increments.

The results for the signal timing assessment can not be generalised further for the following two reasons. Australian signal design guidelines mandate a minimum of 4s for the amber phase in a traffic sequence. If the green time is set to the minimum of 1s, the resulting phase time is 5s, acting as the lower bound. It is uncommon to see cycle times greater than 180s, as this results in extensively long delay times for waiting motorists. When using the most basic split phasing arrangement, the 180s cycle time is divided by two, resulting in 90s for each phase. If the results demonstrate strong trends between a phase time of 5s to 80s, then additional assessment to 90s will not be necessary.

7.2.3 Result Evaluation Metrics

This study uses a range of performance metrics to understand CAV impacts. The performance metrics and justification for their use are itemised as follows.

Demand vs CPR vs Throughput curve: This metric provides a means of understanding the benefits of platooning and shorter reaction times. The maximum throughput is intersection capacity. Understanding how it changes with CPR will provide an insight into how much induced demand a CAV network can absorb, or how much benefit it can provide to existing saturated environments.

Demand vs CPR vs Delay curve: This curve indicates whether increases in intersection DoS lead to the same proportionate increase in delay, in both low and high CPR scenarios. Junction delay functions (JDF) model the delay experienced in response to DoS. Developing a demand vs CPR vs delay curve will indicate if traditional JDF relationships still hold as CPR increases.

Demand vs CPR vs Average Speed curve: Understanding average travel speed provides an accessible understanding of delay by contextualising it with the real-world implication for a driver.

Demand vs CPR vs Average Queue Length curve: Understanding the impact of CAVs on queue length is critical in understanding whether intersection design standards need to be updated, depending on the discharge rate of CAVs and their propensity to clear queues.

Signal Discharge Rate vs CPR curve: Human-driven vehicles have a perception-

reaction lag when moving from a standstill at an intersection. As the traffic signal green-time is reduced, proportionately more of the green-time is wasted in this lag as opposed to moving vehicles through the network. CAV platooning eliminates this lag, meaning that signal green-times could be reduced. Such a reduction would ideally reduce the overall phase time and reduce the delay experienced by the average vehicle at the intersection.

7.3 Experimentation Results

Figure 37 provides results for the signalised intersection with a right turn filter;

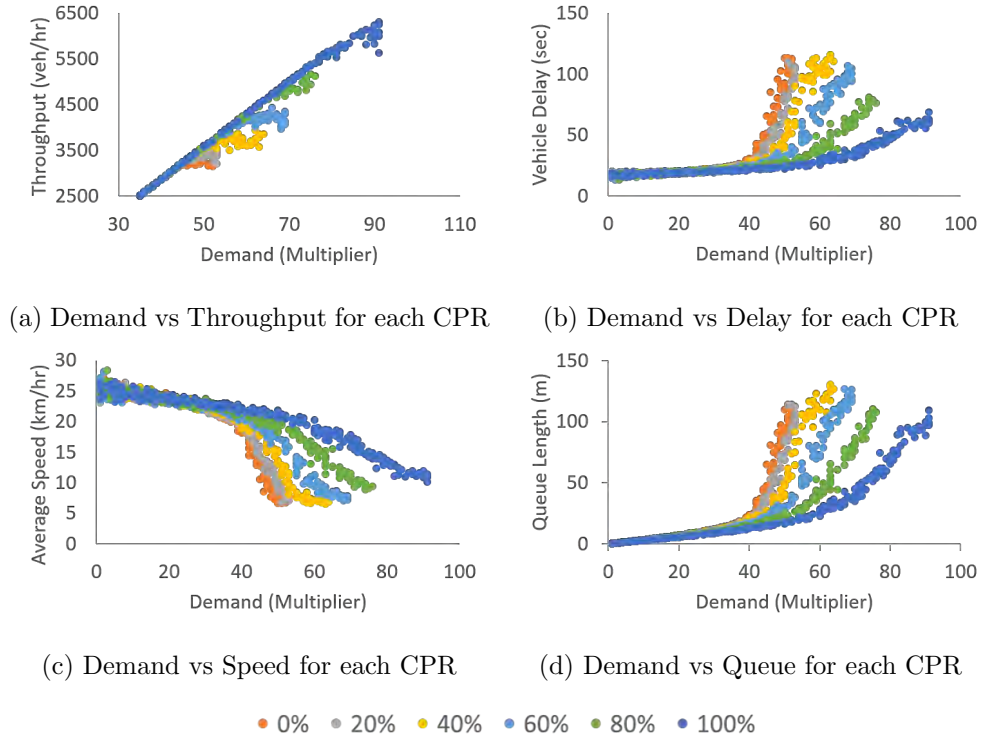


Figure 37: Intersection performance for the signalised intersection with a filtered right turn, as CAV penetration increases.

Figure 37a displays a benefit to vehicle capacity when CPR increases. During unsaturated conditions, the increase in throughput is linearly proportional to demand increase, as is expected. Once capacity is reached, indicated by the linear increase transitioning to a flat line, increases in demand do not increase throughput. Higher levels of CPR improves the resilience of the intersection, shown by its ability to accommodate more demand before reaching capacity.

Refer to Table 12 for the quantitative increase in intersection capacity.

The vehicle delay graph (Figure 37b) shows a similar degree of increased resiliency. Unsaturated conditions provide negligible increases to vehicle delay. Once the intersection approaches capacity, demonstrated by the sharp inflection point, each additional vehicle has a greater impact on intersection delay, resulting in the well-understood compounding increase in delay effect. It is noteworthy that CAVs have three clear impacts on intersection delay; they postpone the onset of the inflection point and allow the intersection to carry more demand before delay increases, they slow the compounding effect of additional vehicles demonstrated by the lower gradient of the curve after the inflection point, and they reduce the ultimate delay experienced by the vehicles demonstrated by the lower curve tip. Similar observations are made when assessing the average queue at the intersection (Figure 37d). Refer to Table 12 for the quantitative decrease in intersection delay and average queue length.

The increase in average travel speed for the different CPR scenarios follows the same curve used to model junction delay. This occurrence is important to note as this means that JDFs for strategic modelling may only need a recalibration rather than a complete re-derivation. Otherwise, average travel speed shows the same benefits as the other metric. CPR increases intersection resiliency, allowing vehicle travel at higher speeds with greater network demand loading. Additionally, when speed does decrease, it remains higher than it would otherwise fall to, at lower a CPR. Refer to Table 12 for the quantitative increase in average travel speeds, and changes to the other performance metrics;

CPR	0%	20%	40%	60%	80%	100%
Throughput (veh/hr)	3323	3436	3789	4239	5047	6098
Change from Base		3.40%	14.02%	27.57%	51.88%	83.51%
Delay (sec)	113.23	110.06	115.9	106.44	80.73	68.47
Change from Base		-2.80%	2.36%	-6.00%	-28.70%	-39.53%
Queue Length (m)	113.25	114.23	130.52	126.39	110.19	109.35
Change from Base		0.87%	15.25%	11.60%	-2.70%	-3.44%
Speed (km/hr)	6.68	6.85	6.55	7.05	8.89	10.16
Change from Base		2.49%	-2.00%	5.48%	33.10%	52.09%

Table 12: Overall impacts of CAV behaviour on the signalised intersection with right turn filter.

Figure 38 provides results for signalised intersection with a dedicated right turn;

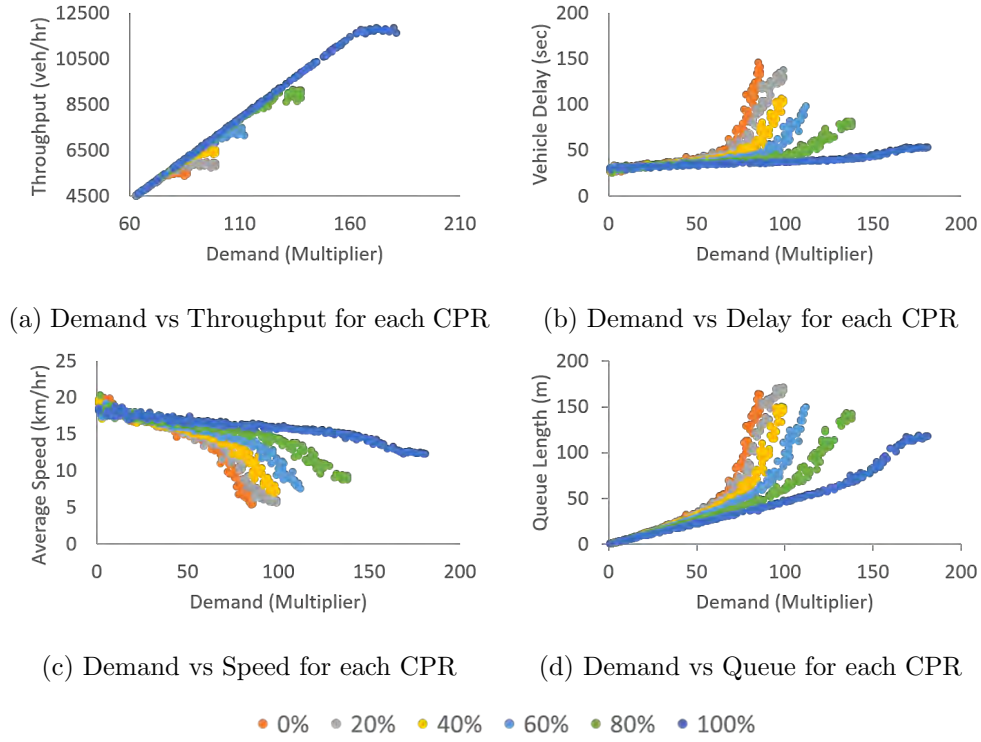


Figure 38: Intersection performance for the signalised intersection with a dedicated right turn, as CAV penetration increases.

The results for the signalised intersection with the dedicated right-turn are similar to that of the intersection with the filtered right turn. CAV operation increases the resiliency of the network and allows it to cater to significantly greater demand. The intention of investigating both the filtered right turn and the dedicated right turn is that a reasonable expectation would be for right-turning vehicles waiting for gaps in the opposing stream to stifle the right lane in a shared-lane arrangement.

Intuition may suggest that finding appropriate gaps becomes difficult when CPR is higher, as inter-vehicle spacing is small. When combined with a high demand, it would lead to both lower throughputs and higher delays. The results in Figure 38 and the summary in Table 13 demonstrate that this intuition is correct. While the filtered right turn intersection saw substantial benefits from CAV operation, greater benefits are achieved from a geometric realign-

ment and providing the right turn with a dedicated lane and signal time. The following table provides the quantitative change in performance metrics;

CPR	0%	20%	40%	60%	80%	100%
Throughput (veh/hr)	5574	5984	6543	7356	8989	11718
Change from Base		7.36%	17.38%	31.97%	61.27%	110.23%
Delay (sec)	145.85	137.44	106.44	98.42	81.39	53.56
Change from Base		-5.77%	-27.02%	-32.52%	-44.20%	-63.28%
Queue Length (m)	164.2	171.8	150.62	149.72	143.5	118.56
Change from Base		4.63%	-8.27%	-8.82%	-12.61%	-27.80%
Speed (km/hr)	5.35	5.64	7.05	7.53	8.83	12.29
Change from Base		5.43%	31.81%	40.93%	65.21%	129.95%

Table 13: Overall impacts of CAV behaviour on the signalised intersection with a dedicated right turn.

Figure 39 provides the results for the roundabout intersection;

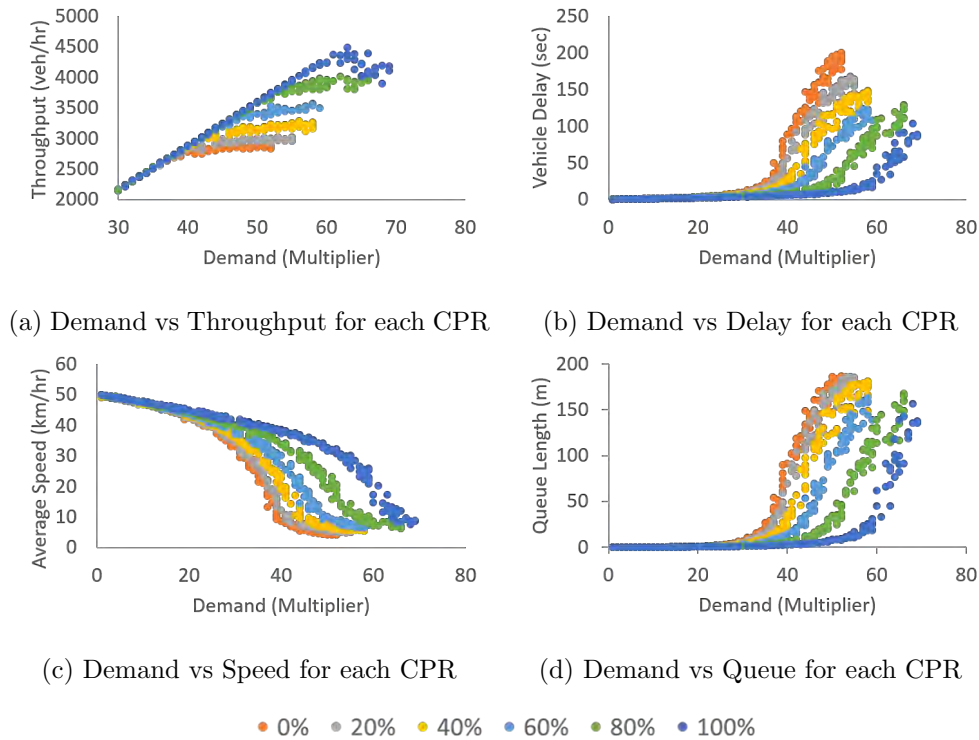


Figure 39: Intersection performance for the roundabout intersection, as CAV penetration increases.

The performance of the roundabout intersection lies between that of the sig-

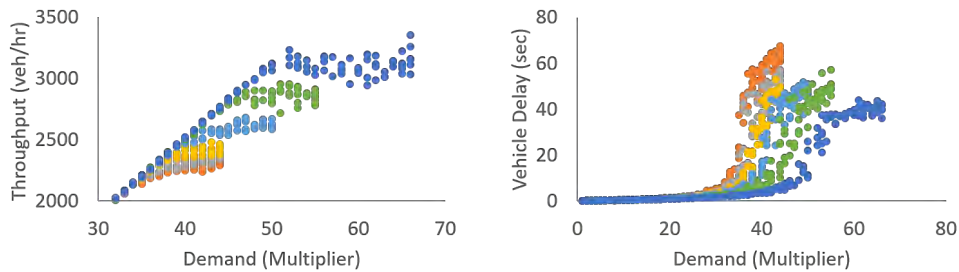
nalised intersection with a filtered right turn, and signalised intersection with a dedicated right turn, for all metrics. During runtime, vehicles were observed entering the roundabout as platoons of multiple vehicles, helping increase the throughput. Although, since the demand from all approaches was balanced, a significantly long platoon was not able to form which could have otherwise warped the results of the roundabout intersection.

250 iterations of each CAV penetration scenario using different random seeds indicates that roundabouts are more sensitive to random variability and slight changes in vehicle arrival patterns. This finding is demonstrated by the lack of distinction in the results between CPR scenarios (shown in Figure 39b and Figure 39d specifically). While the previous two intersection types demonstrated a clear distinction between CPR scenarios, the roundabout results are less clear for CPR under 40%, indicating the contingency of roundabout performance on beneficial vehicle arrival patterns, particularly for queue length and delay. Refer to Table 14 for the quantitative change in network metrics;

CPR	0%	20%	40%	60%	80%	100%
Throughput (veh/hr)	2851	2998	3205	3465	3931	4291
Change from Base		5.16%	12.42%	21.54%	37.88%	50.51%
Delay (sec)	200.34	167.43	147.85	123.55	128.76	103.29
Change from Base		-16.43%	-26.20%	-38.33%	-35.73%	-48.44%
Queue Length (m)	187.07	186.39	181.55	163.77	167.95	156.81
Change from Base		-0.36%	-2.95%	-12.46%	-10.22%	-16.18%
Speed (km/hr)	4.15	4.89	5.46	6.40	6.17	7.47
Change from Base		17.74%	31.62%	54.20%	48.73%	79.93%

Table 14: Overall impacts of CAV behaviour on the roundabout intersection.

Figure 40 provides the results for the priority intersection;



(a) Demand vs Throughput for each CPR (b) Demand vs Delay for each CPR

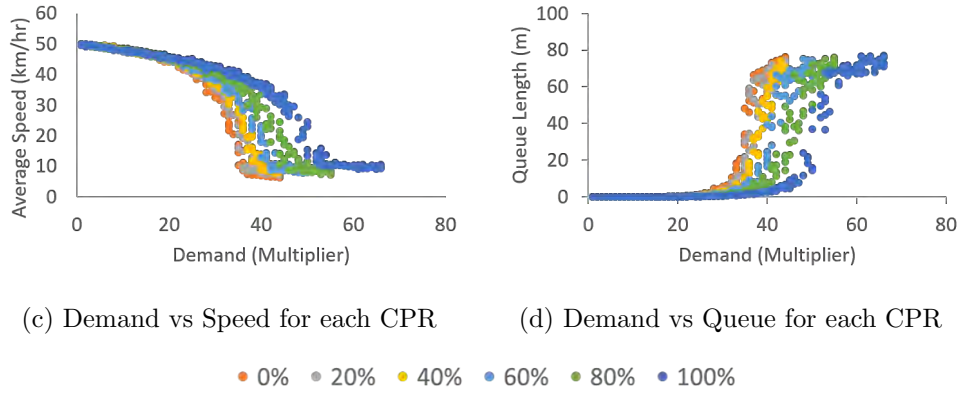


Figure 40: Intersection performance for the priority intersection, as CAV penetration increases

Referring to Figure 40d, while other intersections demonstrated that the incremental increase in queue length after reaching capacity for the 100% CAVs scenario was generally less than that for the mixed fleets scenario, the priority section did not demonstrate this property. Although, despite the queues being the same length, vehicles traversing the priority junction in platoons caused the queues to dissipate quicker. Figure 39b demonstrated this outcome through the lower delay time for higher CAV penetrations despite the same queue length.

Additionally, without a clear turn-based priority, the priority junction suffers the same effect as the roundabout intersection, where random variability in arrival patterns results in a lack of distinction in intersection performance while CPR is below 40%. Refer to Table 15 for the quantitative change in network metrics;

CPR	0%	20%	40%	60%	80%	100%
Throughput (veh/hr)	2303	2354	2428	2619	2878	3168
Change from Base		2.21%	5.43%	13.72%	24.97%	37.56%
Delay (sec)	67.42	57.17	53.04	51.73	57.2	43.8
Change from Base		-15.20%	-21.33%	-23.27%	-15.16%	-35.03%
Queue Length (m)	76.41	74.72	74.2	75.33	76.16	77.22
Change from Base		-2.21%	-2.89%	-1.41%	-0.33%	1.06%
Speed (km/hr)	6.29	7.25	7.73	7.89	7.25	9.06
Change from Base		15.33%	22.92%	25.54%	15.28%	44.15%

Table 15: Overall impacts of CAV behaviour on the priority intersection.

Finally, Figure 41 provides the results for how vehicle throughput and other metrics change in relation to the green-time offered to vehicles;

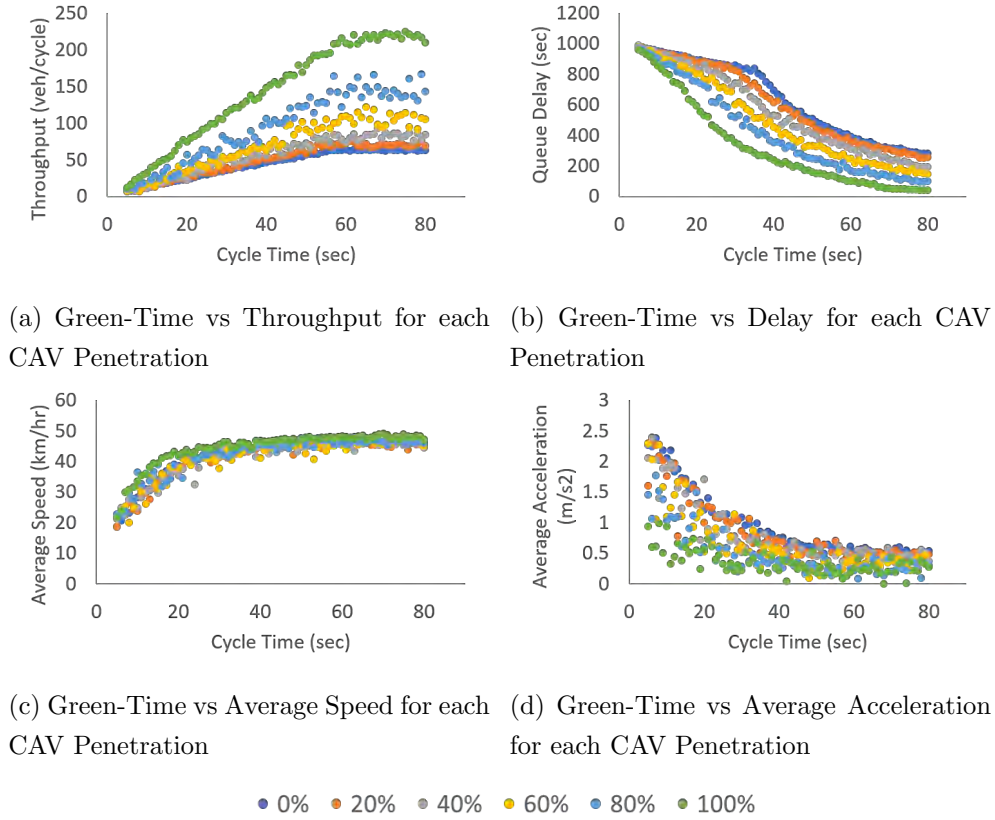


Figure 41: Changes to platoon behaviour as CAV penetration and signal green-time increases.

The results in Figure 41 differ from the other intersections in that this is not an assessment of a specific intersection type, but rather the performance of vehicles and platoons in response to different signal green-times. These results provide numerous noteworthy insights. Figure 41a indicates that increasing the green-time results in higher throughput, an expected finding in itself. As CPR increases, the rate of increase in throughput also increases. This occurrence is attributed to the effect of instantaneous reaction times and shorter headways on vehicle platooning.

Perception and reaction lag is a pertinent source of inefficiency when vehicles debark from a traffic light, which is why providing more green-time provides higher throughput per cycle (indicated by the linear increase in Figure 41a).

Once vehicles join the end of the queue at free flow speed as the perception reaction lag dissipates through the traffic stream, throughput per cycle ceases to increase, indicated by the flattening of the throughput per cycle curve in Figure 41a.

Referring to Figure 41b, the importance of eliminating the perception-reaction lag inefficiency is demonstrated more clearly. Delay decreases linearly with increasing green-time until a critical point is reached. This critical point is when vehicles join the end of the queue and traverse the intersection in the same cycle, and join the queue at speeds greater than standstill. As green-time further increases, the speed at which additional vehicles join increases, resulting in the hyperbolic style curve observed. Throughput would be at its maximum if all vehicles joined the queue at free flow speeds, which would result in the linear portion of the curve in Figure 41b being eliminated. The same figure shows that as CPR increases, the queue delay approaches this optimal outcome with the shortening of the linear perception-reaction lag section of the delay.

The explanation provided above regarding the observed behaviour in throughput and delay is further demonstrated in Figure 41c. This figure demonstrates that a 100% CPR, the intersection operates with higher average speeds resulting from eliminating the perception-reaction lag and more vehicles joining the end of queues at higher speeds. The average acceleration, shown in Figure 41d, is lower as CAV penetration increases, which is also an expected outcome. Firstly, the CAVs in this thesis are designed to have lower acceleration limits than the human vehicles, and also adhere to these limits more stringently. Additionally, CAVs joining the back of queues at higher speeds means that acceleration for shorter periods is sufficient to reach free flow speeds.

7.4 Discussion

The comprehensive microsimulation modelling at each intersection indicates that CAVs provide a substantial benefit to intersection performance. The signalised intersection with a filtered right turn experienced a throughput increase of 84% as CPR increased to 100%. Similarly, the other intersections, signalised intersection with a dedicated right turn, the roundabout intersection, and the priority intersection, experienced a throughput increase of 110%,

51%, and 38%, respectively. Herein potentially lies the first logistical challenge of CAV integration into transport networks. Though at first glance a network-wide throughput increase appears to be a benefit, this increase in throughput may result in the movement or formation of bottlenecks, as each intersection type does not experience the same proportional increase. Consider the interface between a residential precinct and the arterial road network. While the arterial road network experiences up to a 110% increase in throughput at signalised intersections, the local residential road network primarily consisting of roundabouts and priority intersections will only increase in throughput by up to 51%. Without proper planning to cater for this additional demand reaching low capacity intersections, bottlenecks may arise in new areas.

The queue length is an important metric as it greatly informs design decisions. Queue length demonstrated an appreciable reduction in the signalised intersection with a dedicated right turn and in the roundabout intersection. Naturally, an assumption may be that greater throughput results in shorter queues. However, this chapter indicates that this is only true at two intersections, meaning that many of the white papers and thought pieces that advocate for reclaiming road space from vehicles for the purposes of other agents, may not be possible. Assuming that with higher CPR, vehicle demand also increases to maintain intersection operation at capacity, design changes can not be made to reduce the storage space available to vehicles. However, if the demand were to not increase with higher CPR, then queue lengths would fall substantially, and road space could be reclaimed.

Delay also experiences a drastic reduction, with the signalised intersection with a filtered right turn, signalised intersection with a dedicated right turn, roundabout intersection, and priority intersection experiencing a reduction in delay of -40%, -63%, -48%, and -35%, respectively, as CPR approaches 100%. A reduction in delay may not have unforeseen implications such as those mentioned for throughput and queue length. However, shorter travel times often incentivise changes to travel behaviour, with people electing to make more trips by private vehicle or changing their existing travel mode to a private vehicle. Such a behavioural shift brought on by better road network performance would result in increased congestion. While this chapter demonstrates that an undersaturated CAV network is more resilient and able

to cater to a higher level of demand, even an oversaturated intersection is subject to collapse and degradation in performance eventually.

Average travel speed also increased in all environments, with the intersections experiencing an increase in average travel speed of 52%, 130%, 80%, and 44%, respectively. While the percentage improvement in average speed is substantial at 100% CPR, lower CAV penetrations show significantly less and even negligible levels of average travel speed improvement. At priority-controlled junctions such as the roundabout intersection and priority intersection, the arrival of vehicles plays a large role in the performance of the environment. As random seed values change, the arrival patterns of vehicles and the emergence of gaps changes, occasionally leading to opportune scenarios and other times increasing vehicle delay. For this reason, low CPR scenarios showing benefits in average travel speeds were primarily the result of the opportune alignment of vehicle arrivals and gaps in opposing traffic streams.

The effect of CAV platooning on intersection throughput indicates that as green-time increases, throughput also increases. This finding is expected and holds true for human vehicles too. However, as CPR increases, throughput improvement occurs through the mitigation of the perception-reaction lag of human drivers. Increasing CPR results in throughput increasing, and delay decreasing, proportionately faster with each additional second of green-time.

Referring to Figure 41b, the greatest reduction in delay for each consecutive second of green-time occurs when green-time is below 40s. After 40s, improvement to delay still occurs, although the rate of improvement diminishes. When observing 0% CPR at 40s green-time, the base case scenario still operates well below optimal conditions. This result implies that signal times could actually be reduced as CPR increases, allowing phases to occur more frequently and increasing the turnover between movements. Short signal times can not be used now for human vehicles due to the high proportion of green-time wasted in perception-reaction by human drivers. But they should be implemented with CAVs, and result in shorter delays for vehicles. A higher turnover of phases would also reduce the required infrastructure storage space at intersections.

7.5 Conclusion

In this chapter, the impact of CAV penetration on intersection performance was explored. Intersection performance increased substantially, with throughput increasing by up to 110%, delay reducing by up to -63%, queue length decreasing by up to -28%, and average speed increasing by up to 130%. However, this chapter demonstrated that proportionally, all intersections do not experience the same degree of improvement, raising the concern that disproportionate enhancements to parts of the network may redistribute the location of bottlenecks and pinch points. Additionally, improved travel conditions may result in the induction of demand and mitigate any improvements possible through CAV operation.

The other significant contribution of this chapter was in its assessment of the platoon formation and movement in response to changing phase green-times. This chapter demonstrated that large portions of the green-time were wasted by human drivers in perception and reaction, while CAVs can utilise the time better. This outcome means that shorter phase times are possible as CPR increases, increasing the turnover of phases, reducing the delay per approach arm, and ideally reducing the infrastructure required as storage space at intersections.

8 Intersection Phasing Optimisation for a Pedestrian-Priority Signalling Scheme

Abstract: Rapid progress has been occurring in the development of CAVs as a means of solving existing transportation problems. This chapter assesses the feasibility of transitioning existing vehicle-priority signalling schemes to a pedestrian-priority scheme. Such a change could lead to the redevelopment of signalling schemes in densely populated city centres, resulting in behavioural changes in transit and mobility. This study is undertaken using microsimulation to assess the travel times and delays of agents in the system. It uses two intersections as testbeds. The first represents a pedestrian-dominated environment and the second represents a vehicle-dominated environment. The custom behavioural model for CAVs is integrated into the microsimulator Aimsun to emulate CAV behaviour. Testing consists of developing three signalling scenarios. The first is the base case that reflects current signal operation. The second aims to minimise vehicle delay, and the third aims to minimise pedestrian delay. The results indicate that transitioning from the current vehicle-priority signalling scheme to a pedestrian-priority scheme results in a substantial increase in the vehicle and total system delay in a vehicle-dominated environment, regardless of the CAV penetration. However, the pedestrian-dominated environment showed reductions in total system delay from such a transition when CAV penetration exceeded 40%. This study indicates that through the efficiency of CAV operation, signal green-time can be reallocated from vehicles to pedestrians in pedestrian-dominated environments without increasing total system delay. This finding has substantial implications for how pedestrian-focused infrastructure, such as recreational and commercial precincts, central business districts, and public transit hubs, are designed.

As cities continue to proliferate, transport issues such as congestion, road accidents, emissions, parking requirements, and mobility constraints consequently worsen. These issues may be addressed through the network integration of technological advancements in mobility, such as CAVs. Several studies in the literature have attempted to understand the impact of this technology on the economy, environment, and society. However, the literature indicates that attempts to homogenise the disjointed intentions of designing traffic signals for vehicles and pedestrians have not been investigated in detail, especially

with consideration to CAVs and mixed fleets. Previous chapters in this thesis demonstrated that CAVs improve intersection and network performance. This chapter intends to investigate whether the travel time savings gained from CAVs can be redistributed to other agents in the network (mainly pedestrians) through a re-optimisation of intersection traffic signalling.

This chapter intends to determine the effect on intersection performance by converting current vehicle-priority signalling schemes to a pedestrian-priority signalling scheme. The change would theoretically create a safer and more comfortable crossing experience for pedestrians. Whether this leads to a reduction in conflicts is not within the scope of this chapter, as Chapter 11 assesses the impact of CAV operation on vehicle conflict rates. Changing intersections from vehicle-priority to pedestrian-priority would result in worsened vehicle performance at most intersections. For this reason, such a transition may only be possible by leveraging the operational advantages offered by CAVs, such as instantaneous reaction times and platooning.

The investigation is conducted through microsimulation and use of a commercial traffic simulator (Aimsun). The study focuses on investigating the mixed fleet transition between a fully human-driven (0% CPR) and a fully automated (100% CPR) traffic fleet, in 10% increments of CAVs. The study uses two calibrated and validated intersections located in Sydney, Australia. The first is Ultimo Road / Quay Street intersection, and the second is Great Western Highway / Marsden Street intersection. The former testbed represents a pedestrian-dominated environment, while the latter is a vehicle-dominated environment. The intention of testing both extremes of the spectrum is to identify whether such a signalling regime change is beneficial in any case. If neither case shows an improvement in total system performance, then a pedestrian-priority signalling scheme warrants no further investigation. However, if either case shows a benefit, with the most likely being the pedestrian-dominated environment, then further investigations should also be undertaken to determine economic, environmental, and safety benefits arising from such a signalling structure transition.

This chapter contributes to the literature through an assessment of signal configuration and CAV operation on user experience, measured as delay. The

chapter considers the impacts of signals on the three major agents expected to share intersections in the near-term; human vehicles, CAVs, and pedestrians. The novelty of this chapter is in its intensive emulation of CAV behaviour to determine the travel implications on the agents as signal phasing is re-optimised. In addition, this study contributes to the literature by recommending the CAV penetration at which specific signalling schemes are best used.

The immediate benefactors of the results and conclusions produced by this study would be government agencies and decision-makers for transport-related infrastructure, as this study aids in understanding the impacts of CAVs in pedestrian-dominated and vehicle-dominated environments. The results would ideally highlight the need for informed intersection design and infrastructure layout that considers the implications on all agents in the system.

The remainder of this chapter is structured as follows. Section 8.1 contains a summation of the literature and the work completed to date. Section 8.2 outlines the framework for the experimentation conducted in this study and provides in greater detail the process used for pedestrian emulation. Section 8.3 explains the specificities of the case studies. Section 8.4 and Section 8.5 provide the results and discussion, respectively. Finally, Section 8.6 concludes the chapter with a summation of the work.

8.1 Literature Review of Vehicle and Pedestrian Interaction

Several studies have quantified the extent of pedestrian conflicts and accidents in road networks. The Governors Highway Safety Association in the United States [Retting, 2017] showed that on average, 16 pedestrians are fatally injured per day. The European Transport Safety Council [Adminaite et al., 2015] state that pedestrians account for 21% of road fatalities in the European Union. East Asia and India also show unreasonably high rates of pedestrian conflicts, with India showing 60% [Mohan et al., 2009] and the region of East Asia showing 20% of all accidents involving a pedestrian, respectively. These reports show that pedestrian conflicts are not confined to a specific geographic location or social behaviour, but are a universal problem. The World Health Organisation [Organization, 2015] showed that pedestrian conflicts worldwide range from 12% to 38%, with countries averaging 25%. Additionally, 54% of conflicts occur during legal activities [Kumar & Parida, 2011].

To reiterate, this study does not intend to quantify changes in pedestrian conflicts using a pedestrian-priority signalling scheme. It aims to determine whether doing so is feasible in the context of delay, travel time, and general intersection performance. The remainder of this section outlines recent attempts at understanding the interaction between pedestrians and vehicles.

Daganzo and Knoop investigated the effect of giving pedestrians total right of way, with permission to cross at any point and with complete priority [Daganzo & Knoop, 2016]. Vehicles were modelled using kinematic wave theory and the Newell’s car-following model on an infinitely long road with pedestrian crossings. Their experimentation found that for low pedestrian volumes, the road capacity is inversely proportional to the square of the pedestrian volume. Their formulations indicate that more pedestrians crossing results in less volume at each crossing point, and hence greater road capacity between them, but also results in increased vehicular delay. A similar study by Anderson and Knoop found that low pedestrian volumes resulted in decreased vehicle travel speeds and increased delays when pedestrians are given right of way [Anderson & Knoop, 2018]. The contribution of this chapter is to extend the work of Daganzo and Knoop through the investigation of whether any reduction in the pedestrian delay is sufficient to offset the effect on vehicle delay, an outcome that would lead to the better design of intersections in pedestrian-dominated city centres. While a system allowing unfettered pedestrian crossing is appealing, especially in dense city centres, its practicality and safety are questioned in environments containing human drivers.

Many studies have investigated the impact of shared spaces, rather than providing either transit mode priority at crossings. Shared space schemes are designed to decrease disorder on roads, augment safety awareness, diminish the separation between vehicles and pedestrians, and attract people to walking or cycling. Anvari et al., presented a microscopic mathematical model for shared spaces by integrating the car-following model and the shortest path algorithm based on the flood fill algorithm, into the Social Force Model [Anvari et al., 2016]. This model evades conflicts and is successful in incorporating interaction factors between pedestrians and vehicles. Their experiment observed higher vehicle flow, higher vehicle speed, and greater segregation between pedestrians and vehicles in the study area. Their study indicated that

the shared space concept is highly dependent on context, with factors such as flow, ratio, speed of both pedestrians and vehicles, and infrastructure design all influencing shared space performance. This chapter aims to determine whether signal optimisation is a viable means of reaching global improvement, opposed to the context-specific benefits found with shared space schemes.

Frosch conducted a similar study using data from Grumbein's Island on West Virginia University [Frosch, 2017]. The study implemented a shared space in microsimulation, where interaction between vehicles and pedestrians happens frequently. Simulations in Vissim were performed to calculate network performance measures such as travel time, delay time, queue length, and the impact of the proposed shared space on the upstream and downstream intersection. The results of his study found that travel time decreased by 13%. Additionally, the average delay reduced from 60s to 20s, and queues were observed to dissipate quicker compared to the base scenario. The shared space concept reduced the network impact of the congestion zone. Shared spaces can only be implemented in low vehicle flow scenarios and at intersections without long crossing distances for pedestrians. For this reason, the study in this chapter employs a signal optimisation approach to determine if pedestrians can benefit at intersections where shared space is not feasible.

Other studies have focused on the interaction between human vehicles and pedestrians. Arkell quantified the impact of pedestrians on network performance using an Auckland Light Rail access site as the case study [Arkell, 2017]. Light rail vehicle headways of $3min$ and $5min$ were both investigated during the AM and PM peak periods. This study found that pedestrians had a significant impact on the average delay experienced by the Light Rail vehicle. The presence of pedestrians resulted in vehicle delays, where vehicles had to wait for pedestrians to clear the intersection in unprotected phases. The travel time of the Light Rail significantly increases in both the AM and PM peak due to the right of way of pedestrians at intersections. This chapter aims to extend the works of Arkell by investigating the network impact of optimising for each agent in turn. Doing so will lead to the development of a spectrum bound by pedestrian-dominated locations and vehicle-dominated location. Such a spectrum would demonstrate system optimality and the consequences of extreme solutions that discriminately benefit a single agent.

Jayaraman et al., investigated the interaction between CAVs and pedestrians using Uncertainty Reduction Theory [Jayaraman et al., 2018]. This theory asserts that more information about a system leads to less uncertainty. Their study investigated signalised and unsignalised environments. They found that at unsignalised crosswalks, CAV driving behaviour had minimal impact on pedestrians' trust. The authors determined that CAVs could integrate into a shared-space environment with ease, as pedestrians have little to no reaction to the proximity of a CAV. Nevertheless, this chapter focuses on the signalised crossing environment where the implementation of a shared zone is not feasible.

Hagenzieker et al., conducted a study in which cyclists assessed different interactions between bicycles and vehicles (both human and automated vehicles) [Hagenzieker et al., 2020]. The subjective insights provided by the participants indicated that cyclists do not expect to feel safer around a CAV as compared to a human vehicle, nor do they have the certainty that CAVs would stop for them. Several other studies ([Visser et al., 2017], [Blau, 2015], [Lundgren et al., 2017], [Lagström & Malmsten Lundgren, 2016], [Habibovic et al., 2014]) have also used stated-preference surveys to quantify the hesitation of cyclists and pedestrians in interacting with CAVs. Their results show that pedestrians are less likely to cross the road when the driver of an approaching vehicle behaves uncommonly or is inattentive. These studies generally indicate that pedestrians are cautious in interacting with CAVs. Visser et al., determined that outward displays of communication from the CAV, such as lighting and signage, could be used to improve the comfort experienced by pedestrians [Visser et al., 2017]. Ideally, by prolonging the green-time provided to pedestrians for crossings, this discomfort can be alleviated. While this study does not determine changes to comfort and safety, it investigates whether changing the signalling scheme is possible without intersection performance impacts.

Several studies have conducted microsimulation of CAVs, with most considering only the 100% CPR case. Few studies have considered gradually loading, which better reflects a realistic transition period. Minelli et al., modelled the implications of altering mode share on the travel time of the system using the Paramics microsimulator [Minelli et al., 2015]. Their result indicates that a fully autonomous network experiences improvements in travel times, but a network with increased demand results in significant delays. The authors at-

tributed this to too many CAVs concurrently choosing the shortest path. The current study aims to extend the work of Minelli et al., by determining if the CAV operation performance improvements can be reinvested into the system through signal optimisation and achieve a better pedestrian outcome.

8.2 Experimentation Framework

The framework used in this study consists of a collection of interconnected modules (Figure 42), each responsible for generating, using, analysing or transferring data. The traffic data and network data modules provide constraints on the simulation. The traffic data is used to calibrate the model and generate the agents during runtime; it consists of both vehicle and pedestrian agent volumes over the modelling period. The network data provides the geospatial constraints on the movement of the agents. These constraints consist of geometric movement, speed limits, and signalling controls. The microsimulator allows the agents to interact with one another. The microsimulator forwards vehicle information to the external algorithm, which then uses this information to calculate and return appropriate actions for the CAVs. Finally, the simulation results are then passed from the microsimulator to the analysis platform for processing. Refer to the following figure for a diagrammatic representation of the interconnection between modules;

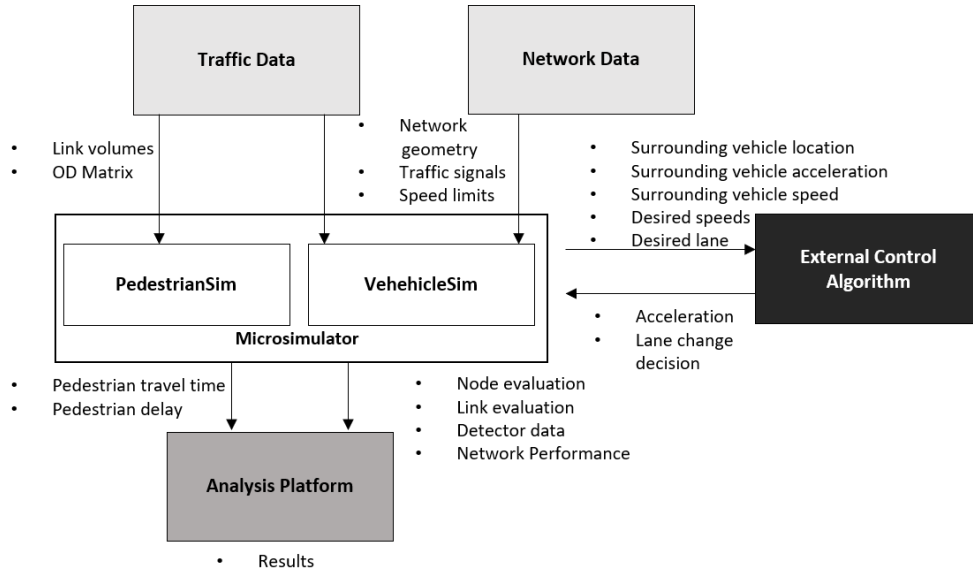


Figure 42: Overview of the framework used for this study.

Traffic Data: Vehicle and pedestrian volumes were both obtained from site visits. This traffic data is used to adjust the OD matrix iteratively. Vehicle turning counts were calibrated in 15-minute intervals, and pedestrian counts were aggregated into 1-minute intervals. Greater detail about the two study areas is provided in Section 8.3.

Network Data: Open Street Map and Google Maps geospatial data imported to the microsimulator were used to create the geometric network. Also, the Street View features of these geospatial imagery platforms provide additional network information such as speed limits, signalling schemes, and phasing configurations, which were calibrated and adjusted using the video recordings from the site visits.

Microsimulator: The microsimulator is a combination of Aimsun 8.1 for simulating vehicle travel and delay (here forth referred to as VehicleSim) and a discrete-time based microscopic simulator for calculating pedestrian delay (here forth referred to as PedestrianSim). The microsimulation process is responsible for facilitating the interaction between the agents in the environment and the movement of data between the different modules. VehicleSim provides the external algorithm with the geolocation, speed, acceleration, and lane-changing intention of each vehicle for the current time increment. In return, the external algorithm provides VehicleSim with the acceleration and lane-changing status for the next time increment. Additionally, the microsimulator records information regarding network performance during runtime and provides data on individual agents, links, intersections, and the network.

External Algorithm: The external algorithm, developed in Chapter 5.1, is used to dictate the CAV behaviour. The algorithm receives information from VehicleSim related to the current vehicle position and behaviour. It then uses this information to determine the appropriate acceleration and lane-changing behaviour for the CAV for the next time increment. The algorithm interacts with VehicleSim through a dynamic linking library (DLL), developed using the Aimsun MicroSDK (software development kit).

8.2.1 Pedestrian Emulation

Pedestrians are emulated using a discrete-time based simulation tool (PedestrianSim) developed in Microsoft Excel. This tool achieves two objectives; the first is to estimate pedestrian delay for each CPR scenario, and the second is to optimise phasing structures that result in minimum pedestrian delay for the pedestrian-priority scenarios.

PedestrianSim emulates pedestrian behaviour in $1s$ increments over a $1hr$ modelling period and calculates the delay experienced by the pedestrians in response to each signal phasing scheme. Pedestrian delay is equal to the duration of the red light that each pedestrian agent experiences. For this reason, stochastic distributions to describe the characteristics of the pedestrian fleet (such as age, walking speed, walking gate, and trajectory) are not necessary. A 100% compliance rate for the traffic signals is assumed.

PedestrianSim works as follows. Input into PedestrianSim is the arrival volumes of the pedestrians in $1min$ bins, and the traffic light phasing times (green-time, red-time, and phase-time). The arrival of the pedestrians between $1min$ bins is deterministic and already known based on real site visits and primary data collection. Within the $1min$ bins, the arrival of the pedestrians at the second-to-second discretisation is random. Random arrival is created by assigning each $1s$ time step with a random number. The number of pedestrians to arrive during that $1s$ time step is equal to the proportion of the random number assigned to that time step, against all time steps in that $1min$ bin. This process is outlined in Equation 144;

$$V_{t,b} = V_b \times \frac{RND_{t,b}}{\sum_{i=1,b}^{60,b} RND_{i,b}} \quad (144)$$

Where, $V_{t,b}$ is the pedestrian volume for the i^{th} second in the b^{th} bin, V_b is the total pedestrian arrival for the b^{th} bin, and $RND_{t,b}$ is the random number associated with the i^{th} second in the b^{th} bin.

Random arrival over the $1min$ bins avoids bias for any particular signalling schemes. If the pedestrian arrives while the traffic signal is green, then a delay of 0 is registered for that pedestrian. If the pedestrian arrives during a red light, then the delay is counted in $1s$ increments until the light changes

to green. The total pedestrian delay is then the summation of the delay experienced by all pedestrians during the simulation period.

8.2.2 Signal Optimisation Process

The optimisation objective function is to minimise average delay for the agent under investigation. For the vehicle-priority signalling scenario, the objective is to minimise vehicle delay. Similarly, for the pedestrian-priority signalling scenario, the objective is to minimise pedestrian delay.

The branch and bound (B&B) algorithm is used to optimise the signalling for each scenario tested. The B&B algorithm works by incrementally moving through a decision tree where the set of all solutions is partitioned into subsets of solutions. The decision tree contains all feasible cycle times and green-time arrangements for the signalling structure, and the algorithm operates by exploring each branch in succession. If a better solution is found in the next iteration along a branch, then the solution branch is retained and further investigated. Otherwise, the branch is terminated, and further solutions are not investigated along it. This process continues until all branches have been explored, with the optimal solution providing the lowest total system delay. The pseudocode for the B&B algorithm is provided in Algorithm 3;

Algorithm 3: Pseudocode for the B&B Algorithm

```

set:  $L = \{X\}$ 
initialise:  $\hat{x}$ 
while  $L \neq \emptyset$  do
    select a subproblem  $S$  from  $L$  to explore
    if a solution  $\hat{x}' \in \{x \in S \mid f(x) < f(\hat{x})\}$  can be found then
        | Set  $\hat{x} = \hat{x}'$ 
    end
    if  $S$  cannot be pruned then
        | Partition  $S$  into  $S_1, S_2, \dots, S_r$ 
        | Invert  $S_1, S_2, \dots, S_r$  into  $L$ 
    end
    Remove  $S$  from  $L$ 
end
return  $\hat{x}$ 

```

Where, L is the unexplored search space, \hat{x}' is the solution under investigation, and S is the branch under investigation. Selection of the first iteration affects the performance of the B&B algorithm.

Three components affect the algorithm; the search strategy, branching strategy, and pruning rules. The search strategy, highlighted in the third line of the pseudocode, influences the order of node selection. The search strategy used is called best-first search. This strategy selects the next iteration as the phasing structure that generates the minimum delay time. The branching strategy refers to how each subproblem is divided, with this study using the wide branching strategy. This strategy divides the current branch under investigation into all possible sub-branches by incrementing the green-time by $\pm 1s$. The pruning rule indicates whether to proceed further along a specific subproblem. The bounding function serves as the pruning rule, with this study electing to discard (prune) subsequent nodes along branches that increase previously calculated delay time. A subproblem may appear more than once from a specific parent node in an iteration. If the subproblem has already been explored for that CPR, the branch is pruned.

If the first iteration is chosen poorly, many modelling iterations are needed to reach the optimal solution. For this reason, queueing theory with a Markovian arrival and departure rate is used to determine the first iteration.

The wait time of a system is a function of the arrival rate and the average queue length, given by Equation 145;

$$w_q = \lambda_q^{-1} l_q \quad (145)$$

Where, w_q is the wait time of the approach arm, λ_q is the arrival rate of the approach arm, and l_q is the average queue length of the approach. The average queue length (l_q) is calculated by Equation 146;

$$l_q = \frac{\rho^{c+1}}{(c-1)!(c-\rho)^2} \cdot P_0 \quad (146)$$

Where, ρ is the utilisation rate of the approach, c is the number of lanes of the approach, and P_0 is the probability of the approach having no queue. The utilisation of an approach is determined using Equation 147;

$$\rho = \frac{\lambda}{c \times \mu} \quad (147)$$

Where, μ is the departure rate of the approach per lane. The probability of the approach lane having no queue is calculated using Equation 148;

$$P_0 = 1 + \rho + \frac{\rho^2}{2(1 - \frac{\rho}{c})} \quad (148)$$

Using Equation 145 to 148, the average queue length of a movement at the intersection becomes a function of the arrival rate (λ), the number of lanes (c) and the departure rate (μ). Therefore, the total intersection delay is the summation of the delay experienced by each vehicle in each turning movement of the intersection, and is calculated using Equation 149;

$$Delay = \sum_{x=1}^n A_x \times \lambda^{-1} \times \frac{(\frac{\lambda}{c \times \mu})^{c+1}}{(c-1)! (c - \frac{\lambda}{c \times \mu})^2} \times (1 + \frac{\lambda}{c \times \mu} + \frac{(\frac{\lambda}{c \times \mu})^2}{2(1 - \frac{\lambda}{c^2 \times \mu})}) \quad (149)$$

Where, x is the ID of the turning movement, n is the total number of turning movements in the intersection, and A_x is the arrival rate of turning movement x . The approach arrival rate (λ) and departure rate (μ) are calculated using Equation 150 and Equation 151, respectively;

$$\lambda_x = \frac{A_x \times CycleTime_x}{3600} \quad (150)$$

$$\mu_x = \frac{D_x \times CycleTime_x}{3600} \times \frac{GreenTime_x}{CycleTime_x} \quad (151)$$

Where, D_x is the discharge rate of turning movement x . Equation 149 to 151 are used to develop a single equation to calculate the total delay experienced by the intersection. The delay equation becomes a function of the green-time given to each turning movement. Using a system of partial differential equations, the minimum delay and corresponding phase green-time and cycle-time are calculated, also forming the first iteration to the B&B algorithm.


8.3 Case Studies

This investigation uses two case studies, both located in Sydney, Australia. The first is the pedestrian-dominant Ultimo Road/Quay Street Intersection located near the Sydney CBD. The second is the vehicle-dominant Great Western Highway/Marsden Street Intersection located in the western suburbs of


Sydney. The purpose of selecting two vastly distinct locations is to investigate pedestrian-priority at both ends of the spectrum. Finding no benefit in a vehicle-dominant environment would not dismiss the possibility of finding benefit in a system dominated by pedestrians. Alternatively, finding benefits exclusively in a pedestrian-dominated environment would not provide insight into the range of intersections where a pedestrian-priority signalling scheme could be beneficial.

8.3.1 Study Area

Primary data at both intersections were collected between 11:30 AM and 12:30 PM. Historical traffic data provided by the RMS was used to convert the observed data into average weekday AM peak hour volumes [RMS, 2018]. Figure 43 provides both a context map and the point-of-view of the cameras used to record vehicle and pedestrian activity.



(a) Pedestrian-dominant environment context map.




(b) Vehicle-dominant environment context map.



(c) Pedestrian-dominant environment site visit first person view from the recording equipment.



(d) Vehicle-dominant environment site visit first person view from the recording equipment.



(e) Pedestrian-dominant microsimulation model environment. (f) Vehicle-dominant microsimulation model environment.

Figure 43: Context map, location photograph, and microsimulation model for each of the study environments.

8.3.2 Model Volumes

Traffic and pedestrian volumes are provided in Table 16 (*veh/hr*);

Ultimo Road / Quay Street Intersection												
Time	C4			C1			C3			C2		
	L	T	R	L	T	R	L	T	R	L	T	R
11:30-11:45	98	32	15	20	16	18	26	69	29	13	12	37
11:45-12:00	77	66	18	25	16	20	47	77	28	10	25	45
12:00-12:15	88	42	34	37	34	28	55	73	26	15	20	39
12:15-12:30	85	48	31	29	35	28	58	69	35	22	38	37
Total (turn)	348	189	97	111	101	94	187	288	119	60	95	158
Total	634			306			594			313		
Great Western Highway / Marsden Street Intersection												
Time	C4			C1			C3			C2		
	L	T	R	L	T	R	L	T	R	L	T	R
11:30-11:45	126	142	26	40	40	103	5	47	7	45	137	193
11:45-12:00	144	167	19	38	34	82	8	47	8	40	151	177
12:00-12:15	129	169	25	53	42	126	4	41	12	33	160	158
12:15-12:30	134	197	25	60	64	132	8	53	11	36	163	169
Total (turn)	533	676	95	192	181	443	27	188	39	154	612	696
Total	1304			816			254			1462		

Table 16: Observed traffic volumes from 11:30 AM to 12:30 PM.

Traffic data was collected in 15min intervals and 1min intervals for vehicles and pedestrians, respectively. In Table 16, L, T, and R is the left turn, through movement, and right turn volumes respectively. Using the RMS historical traffic volumes database, a peak hour factor of 1.46 is used at Ultimo Road / Quay Street and 1.37 is used at Great Western Highway / Marsden Street to convert observed volumes to average AM peak-hour volumes [RMS, 2018].

8.3.3 Model Calibration

The RMS Traffic Modelling Guidelines suggests using the GEH goodness-of-fit calibration criteria [RMS, 2013]. This GEH statistic is a Chi-squared statistic designed to be tolerant of large errors in low flows. It is often used to compare two sets of traffic counts, the observed counts and the modelled counts. Refer to Section 11.3.2 and Equation 171 for a detailed empirical description of the calibration process, where it is applied to a network as opposed to an intersection. The GEH statistic is required to be less than 10 for 100% of turning movements and less than 5 for 85% of turning movements. The scaled observed volume (*veh/hr*), modelled volume (*veh/hr*), and GEH for each turning movement are provided in Table 17;

Turn		Ultimo Rd / Quay St			Great Western Hwy / Marsden St		
		Observed	Modelled	GEH	Observed	Modelled	GEH
C1	Left	162	172	0.79	730	693	1.39
	Through	147	152	0.36	926	955	0.94
	Right	137	126	1.02	130	114	1.49
C2	Left	88	92	0.46	263	271	0.50
	Through	139	130	0.76	248	259	0.69
	Right	231	263	2.04	607	589	0.73
C3	Left	273	269	0.27	211	206	0.38
	Through	420	413	0.36	838	817	0.76
	Right	174	165	0.67	954	932	0.71
C4	Left	508	482	1.18	37	38	0.22
	Through	276	280	0.26	258	255	0.17
	Right	142	139	0.25	53	67	1.76

Table 17: Scaled volume (*veh/hr*), modelled volume (*veh/hr*), and GEH.

Regression analysis is also conducted to avoid prioritisation of larger or smaller turning movements. The regression analysis involves fitting a linear model originating at 0 to the observed volumes and modelled volumes data. The r^2

value for Ultimo Road / Quay Street and Great Western Highway / Marsden Street is 0.9978, and 0.9995, respectively, well above the 0.95 requirements stated in the RMS modelling guidelines.

Finally, model stability is established by demonstrating that the total vehicle kilometres travelled between random seeds does not vary significantly. The RMS modelling guidelines recommend the seed numbers when creating calibrated models. The vehicle kilometres travelled results are provided in Table 18. This table indicates that the models are sufficiently stable, meaning that they are fit for purpose in this study;

Seed	Ultimo Rd / Quay St		Great Western Hwy / Marsden St	
	VKT	Diff from Median	VKT	Diff from Median
28	369	0.7%	1033	0.0%
560	367	0.0%	1034	0.0%
2849	355	-3.3%	1017	-1.6%
7771	381	3.8%	1048	1.4%
86524	362	-1.2%	999	-3.4%
5321	372	1.5%	1026	-0.7%
137	370	0.9%	1031	-0.2%
98812	370	0.8%	1044	1.1%
601027	363	-1.1%	1033	0.0%
559	364	-0.8%	1053	1.9%

Table 18: Vehicle kilometres travelled results for each seed.

8.3.4 Experimentation Structure

The experimentation structure is as follows. Ten random seeds are modelled for each scenario, and the median seed results are used for reporting. CAV penetration is incrementally increased at both intersections in 10% increments, starting from 0% completely human-driven to 100% completely autonomous. Each CAV penetration scenario tests three signalling regimes, consisting of the base case signalling structure, the structure that minimises vehicle delay (vehicle-priority), and the structure that minimises pedestrian delay (pedestrian-priority). These optimal signalling regimes are found using the B&B optimisation algorithm outlined in Section 8.2.2. Fairness is maintained by imposing restrictions on the structure of phase lengths, as follows;

- The length of the pedestrian crossing phase must not be less than the

time taken to cross the intersection at a walking speed of $1.2m/s$.

- The inter-green phase within a cycle must be 4s long.
- All crossings must be given green-time in each cycle.

8.4 Results

This section presents the microsimulation modelling results as average delay and level of service (LOS). LOS contextualises the delay and provides a scale of significance to changes in delay between scenarios. RMS's LOS criteria is an amendment to the original HCM criteria [Roads & Authority, 2002].

8.4.1 Pedestrian-Dominant Environment

Figure 44 and Table 19, provide the average delay and LOS results for the three signalling regimes in the pedestrian-dominant environment;




Figure 44: Delay per agent at the Ultimo Road / Quay Street intersection for the three modelling scenarios, in 10% increments of CAV penetration.

When implementing the base case or vehicle-priority signalling scheme, the system shows little performance improvement with increasing CPR. However, when a pedestrian-priority signalling scheme is used, the delay per agent reduces with increasing CPR. This experimentation shows that system benefits are attained when transitioning to a pedestrian-priority signalling model at a CPR greater than 40%.

The base case signalling scheme with 0% CAVs has an average delay per agent

of 45.4s. As CPR increases, the average delay per pedestrian remains constant, and the average delay per vehicle declines at a compounding rate of -3.1% with each 10% increment in CPR. The total system delay per agent, considering both vehicles and pedestrians declines at a compounding rate of -1.15%.

The vehicle-priority signalling regime shows an oscillating pattern about the results of the base case signalling regime as CPR increases from 0% to 100%. In this signalling regime, the pedestrians experience a compounding delay increase of 0.57% per agent for each 10% increment of CPR. This increase is attributed to the re-optimisation of cycle times for vehicle performance, which experience a compounding delay decrease of -3.14%. When considering the system as a whole and combining pedestrian and vehicle delays, each 10% increment of CPR results in a compounding delay decrease of -0.84%.

The pedestrian-priority signalling scheme showed significant delay reductions with increasing CPR, for all agents. The vehicles, pedestrians, and total system experience a compounding delay decrease of -5.68%, -3.04%, and -4.92%, respectively. Despite the compounding delay decrease rate being greater than that of the vehicle-priority scheme for vehicles, the absolute delay in the pedestrian-priority scheme is substantially higher. Vehicle delay under the pedestrian-priority scheme is 41s and 23s for the 0% and 100% CPR cases, respectively. However, under the vehicle-priority scheme, these delays are 19s and 13s, respectively. While CAV operation aids in reducing vehicle delay, the increase in absolute delay caused by the signal scheme change does not return to base case levels. Table 19 provides the delay and LOS for each scenario;

CAV %	0	10	20	30	40	50	60	70	80	90	100
Base Case	45.4	44.9	44.4	43.1	42.8	42.3	41.9	41.7	40.9	40.7	40.4
Vehicle Priority	45.4	44.7	43.2	41.7	44.2	41.8	42.7	39.9	41.7	38.7	41.8
Pedestrian Priority	55.8	56.9	49.8	45.2	42.2	39.3	39.4	39.6	36.8	33.8	33.7
LOS F	LOS E		LOS D			LOS C			LOS B		LOS A
D > 80s	55s < D < 80s		35s < D < 55s			20s < D < 35s			10s < D < 20s		D < 10s

Table 19: Delay (*sec*) and LOS at the Ultimo Road / Quay Street intersection for all CPR and signalling scenarios.

The LOS results indicate that while the base case and vehicle-priority signalling regime both experience a reduction in total system delay with increasing CAV penetration, they remain within the same LOS bracket, a rating of D. However, the pedestrian-priority scheme causes a reduction in delay significant enough to increase the intersection by two brackets from an E to a C. This table contextualises the results and helps determine the significance of the delay reduction.

8.4.2 Vehicle-Dominant Environment

The vehicle-dominant Great Western Highway / Marsden Street Intersection shows different results to that of the pedestrian dominated environment. The modelling results (refer to Figure 45 and Table 20) indicate that in vehicle-dominated environments, transitioning to a pedestrian-priority model is not beneficial for intersection performance regardless of the CAV penetration. Refer to the following figure for the change in delay for each signalling scheme in the vehicle-dominant environment;




Figure 45: Delay per agent at the Great Western Highway / Marsden Street Intersection for the three modelling scenarios, in 10% increments of CPR.

When retaining the base case signalling regime, the pedestrians experience no change in delay, and the vehicles experience a compounding delay decrease of -4.56%. This results in a total system compounding delay decrease of -4.12%.

Implementation of the vehicle-priority signalling regime results in small changes

to total system delay with changing CAV penetration. Compounding vehicle delay decreases by -4.79%. The difference in delay between the base case and vehicle-priority signalling regime is -4.05% and -6.39% for the 0% and 100% CPR cases, respectively. The pedestrians experience a compounding delay decrease of -1.28%, resulting in a total system delay decrease of -4.45%.


Under the pedestrian-priority signalling regime, the vehicles, pedestrians, and total system experienced a compounding delay decrease of -4.32%, -11.91% and -4.48%, respectively. While increasing CPR aided in reducing vehicle delay, the delay between the pedestrian-priority case and base case experienced by vehicles increased by 90.0% and 94.9% for 0% and 100% CPR, respectively. The LOS table also indicates the substantial increase in delay experienced by the system under the pedestrian-priority regime regardless of CPR;

CAV %	0	10	20	30	40	50	60	70	80	90	100
Base Case	70.1	61.3	57.7	54.3	52.5	51.5	50.5	49.8	48.6	46.5	46.0
Vehicle Priority	67.5	61.4	57.4	55.1	52.6	48.2	46.9	45.6	44.6	43.7	42.8
Pedestrian Priority	126	112	104	101	104	97.8	103	108	82.7	80.5	79.9
LOS F	LOS E		LOS D		LOS C			LOS B		LOS A	
D > 80s	55s < D < 80s		35s < D < 55s		20s < D < 35s			10s < D < 20s		D < 10s	


Table 20: Delay (*sec*) and LOS at the Great Western Highway / Marsden Street intersection for all CPR and signalling scenarios.

8.4.3 Results Summary


Section 8.4.2 and Section 8.4.1 made mention of the results for pedestrians and vehicles separately. The figures and tables above only show the aggregated average delay per agent for the network. Figure 46 provides a detailed summation of the delay experienced per vehicle and per pedestrian for each signalling and CPR scenario;



(a) Pedestrian-dominant environment Base case signalling scheme. (b) Vehicle-dominant environment Base case signalling scheme.



(c) Pedestrian-dominant environment Vehicle-priority signalling scheme. (d) Vehicle-dominant environment Vehicle-priority signalling scheme.



(e) Pedestrian-dominant environment Pedestrian-priority signalling scheme. (f) Vehicle-dominant environment Pedestrian-priority signalling scheme.

Figure 46: Comprehensive modelling results detailing average delay experienced by vehicles, pedestrians, and system average for each CPR, signalling scheme, and intersection.

This data provides significant insights into CAVs, pedestrians, and signalling schemes. Increasing CAV penetration from 0% to 100% reduces the delay for all agents, regardless of the intersection type and the signalling scheme used. The vehicle-dominated environment only showed total system performance improvements under the vehicle-priority scheme. The improvements to vehicle operation offered through CAV technology were not sufficient to offset the increase in delay caused by transitioning to a pedestrian-priority model. Such a transition is possible at a pedestrian-dominated environment where pedestrian volumes account for a larger proportion of the system. At approximately 40% CAVs, total system performance under a pedestrian-priority scheme fell below that of a vehicle-priority and base case signalling scheme. However, total system performance improvements in the pedestrian-dominant environment using a pedestrian-priority signalling scheme were the result of significant sacrifice to vehicle performance.

8.5 Discussion

The results show a global decrease in the delay experienced by each agent in the system between a CPR of 0% and 100%, indicating that regardless of the signalling scheme adopted, CAV operation will have benefits for both pedestrians and vehicles. In the pedestrian-dominated environment, this benefit amounts to a compounding delay decrease of -3.10% with each 10% increment in CAV penetration, and -4.56% in the vehicle-dominated environment.

The vehicle-priority signalling scheme offers vehicles benefit at both locations. However, the pedestrians experience a reduction in delay only at the vehicle-dominated location. When comparing the base case signalling regime to the vehicle-priority scheme, there are no additional benefits to vehicle delay at either location. The reason for this is because intersections in Australia are designed to minimise delay and maximise LOS for vehicles, with system performance being measured through vehicle delay exclusively, effectively meaning that the base case already operates similarly to a vehicle-priority scheme.

Using the vehicle-priority signalling scheme, pedestrians experience a compounding decrease in delay of -1.28% in the vehicle-dominated environment, but an increase of 0.57% in the pedestrian-dominated environment. Comparing delay in absolute terms, this increase in delay amounts to 1.9s per pedes-

trian. The disparity in the volume of pedestrians at the two intersections is the cause of this increase. At Ultimo Street / Quay Street, the pedestrian volumes are higher, meaning that when using a random arrival pattern, the volume of pedestrians arriving during a red light is greater. At the vehicle-dominated environment, the likelihood of arriving at a red light is the same. However, lower proportions of pedestrians reduce the impact on total pedestrian delay.

Using a pedestrian-priority signalling scheme has total system benefit only at the pedestrian-dominated environment, although still increases the delay for vehicles in both environments. Pedestrians experience greater benefit, with compounding delay reducing at -3.04% and -11.91% at the pedestrian-dominated and vehicle-dominated locations, respectively. When comparing against the base case signalling scheme, the pedestrian-priority scheme results in a -59.9% and -80.7% decrease in pedestrian delay at the pedestrian-dominated and vehicle-dominated locations, respectively, with 100% CAVs. The vehicles experience a substantial increase in delay. For the 0% CPR case, vehicle delay increases by 44.3s (124%) and 54.5s (90%) at the pedestrian-dominated and vehicle-dominated environments, respectively. These figures change to 18.5s (71%) and 36s (95%) when CPR increases to 100%.

The use of a pedestrian-priority signalling scheme in a vehicle-dominated environment should be disregarded. The increase in vehicle delay unreasonably outweighs the benefit delivered to pedestrians, especially considering the relatively smaller proportion of pedestrians at such a location. However, the use of this scheme in a pedestrian-dominated environment indicates that benefits to the system delivered through CAV operation can be reallocated to green-time for pedestrian crossings. Doing so does raise equity concerns. Minimising the total system delay disproportionately affects the minority (the vehicles), worsening their situation for the betterment of the majority (the pedestrians).

A CAV proportion of 40% is the first case where the total system delay of the pedestrian-priority signalling scheme falls below that of the other two schemes. At this penetration, the delay per agent decreased by -1.35% (-0.4s). This reduction is an amalgamation of a pedestrian delay decrease of -49% (-16s) and a vehicle delay increase of 80% (24.5s). Criteria must be developed to determine whether such a course of action will better the system and is

equitably appropriate to do so. Criteria for determining system betterment could include;

- Trip purpose and value of time: Determining approximate trip purposes and then assigning a value of time is a means of converting delay values to monetary and economic impacts. The appropriate course of action would then be that which minimises economic impacts.
- Safety: This study has neither considered the vehicle safety implications of shortening vehicle green times, nor the pedestrian safety implications of lengthening crossing times. Such an investigation would determine the likely change in accidents for vehicles and pedestrians. Many studies determine the economic impact of an accident based on severity, so such studies can be applied to convert safety to a monetary economic impact.
- Commercial traffic: Changes to commercial traffic has historically justified infrastructure spend. Consider the AUD\$290million cost of the Newcastle Light Rail project, designed to increase urban amenities and pedestrian activity in the central business district [Paris, 2017]. This chapter demonstrated that while changing from a vehicle-priority signalling scheme to a pedestrian-priority scheme does not have infrastructure costs, it has economic costs in the form of vehicle delay. A cost-benefit analysis into the economic benefit of increased pedestrian traffic and spending, against the financial burden of increased vehicle delays, could again provide further clarity into the feasibility of the signalling change.
- Other Key Performance Indicators (KPI): Improving pedestrian comfort has other societal benefits. An improved pedestrian network incentivises walking and naturally reduces dependence on private vehicles. This, in turn, leads to reduced vehicle congestion and improved health. Additionally, an improved pedestrian network would lead to higher pedestrian traffic around commercial businesses and increase economic output. KPIs such as these quantify indirect and intangible impacts of decisions.

This study demonstrated that the transition from a vehicle-priority signalling scheme to a pedestrian-priority scheme has benefits for the system. However, it has raised the equability concerns around redistributing delay for the betterment of the majority at the detriment of the minority. Further investigations

into the economic impacts of delay, societal benefits of changes in conflict rates, or potential benefits in increased physical activity and foot-traffic around commercial businesses, may provide a strong argument into the implementation of such a scheme.

8.6 Conclusion

This study quantified the impacts of transitioning from a vehicle-priority phasing structure to a pedestrian-priority phasing structure. The feasibility of doing so was evaluated by increasing the CAV penetration in 10% increments from 0% to 100%. Two different intersections, a pedestrian-dominated location and a vehicle-dominated location, were used to assess the impact of three different signalling regimes, the base case scheme, a vehicle-priority scheme, and a pedestrian-priority scheme.

The study found that regardless of the signalling scheme adopted, increasing CAV penetration results in a decrease in delay for all agents in the network. In a pedestrian-dominated environment, a pedestrian-priority signalling scheme can improve the pedestrian LOS without significant detriment to the vehicle LOS. At 40% penetration of CAVs, the average delay per agent is -1.35% lower when comparing a pedestrian-priority scheme to the base case regime. This is the results of a -48.5% reduction in pedestrian delay and an 80.1% increase in vehicle delay. The trend continues to 100% CAV penetration, where the total system delay is reduced by -16.6%, comprised of a -60% reduction in pedestrian delay and a 71.1% increase in vehicle delay. No such benefits are found in transitioning the vehicle-dominated environment to a pedestrian-priority signalling scheme, regardless of CAV penetration.

Acknowledgement of Publication

This work is currently under review for publication - Viridi, N., Dizon, M., Grzybowska, H., & Waller, S. T. (2019). An Investigation into the Feasibility of Adopting a Pedestrian-Priority Signalling Scheme in the Presence of Connected and Autonomous Vehicles. Transportation Research Part A: Policy and Practice, Under Review.

9 Effect of CAV Behaviour and Selective Cooperation on Motorway Capacity

Abstract: This chapter conducts two investigations to explore the impact of CAVs on motorway weaving and merge junctions. The first investigation involves incrementally loading a weaving section and merge section with traffic, to determine the resiliency of these bottlenecks to congestion with increasing CPR. This investigation demonstrates that when weaving proportions are low (5%), a 100% CAV results in an 83% throughput increase. When the weaving proportion is increased to unreasonably high levels (40%) inefficiencies resulting from the CAV emulation framework result in decreasing throughput as CPR increases. This outcome is the result of stringent cooperation requirements between vehicles globally reducing the average travel speed of vehicles to cater to the large proportion of lane-changing vehicles. In the high weaving proportion scenario, human behaviour concentrates the delay with vehicles waiting on the secondary lanes of the weaving section, whereas CAVs distribute the delay among all agents and result in overall worse system performance. In contrast, the merge junction showed a 131% throughput increase with a 100% CAV fleet. The second investigation segregates the CAV fleet in subclasses and forces hierarchical priority through the weaving section, forcing vehicles to selectively cooperate. The purpose of this investigation is to assess the impact of potential policies that provide vehicles with priority through the infrastructure. Segregating vehicles by class and then instituting a class hierarchy showed to provide little to no disbenefit to network performance, assuming that weaving proportions stayed below 20% and demand remained below saturation ($DoS \leq 0.8$).

The literature, governments, and the private sector all regard CAV technology as highly disruptive to current transportation systems. Many patents such as that by Bloomquist et al., [Bloomquist et al., 2005] outline potential V2V and V2I components retrofitted to infrastructure in preparation for CAV operation. These advancements would accelerate CAV integration into transport systems.

As a direct consequence of recent advancements, significant research is being conducted into the potential impacts of CAVs on highway environments. However, these studies often confine the study area to a disjointed subsection of the highway. While experimentation of this nature provides valuable insights into

the impact of CAVs, doing so often creates an idealised environment which obscures additional critical information by the extents of the microsimulation study area. For example, queue formation and dissipation become difficult to access when the approach length of intersections or the ramp length of highways is inadequate to cater for demand. This becomes a more significant concern in studies that either assess the impact of increased traffic demand or discover that altruistic or cooperative merging operation leads to a reduction in capacity or travel speeds. Another limitation in many studies is the use of standard commercial software parameters in emulating CAV behaviour.

This chapter uses the developed CAV emulation algorithm to investigate the impact of CAV penetration and demand increase on the performance of highway weaving sections and merge junctions. Additionally, this chapter examines the impact of segregating the CAV vehicle fleet into separate classes and providing certain classes with priority over others. CAV technological advancement is being driven by private sector investment, so an investigation into the impacts of monopolistic behaviour such as premium services is necessary.

The remainder of this chapter is structured as follows. Section 9.1 provides an in-depth review of the literature and relevant studies. Section 9.2 explains the framework for experimentation in the microsimulation environment. Section 9.3 details the results and Section 9.4 provides a discussion. Finally, the chapter concludes with Section 9.5, providing a summation.

9.1 Literature Review of CAV Behaviour Impacts on Motorway Environments

The literature contains studies conducted on the potential capacity and throughput impacts of CAVs, with a heavy emphasis on the highway environment. This section contains a summation of relevant work done to date. The literature presents two opposing views on the impact of CAVs on throughput. While all sources acknowledge that CAV operation will increase throughput in the short-term on highways, contention arises on the longevity of this improvement. The literature also presents opposing views on the source of bottlenecks and point of failure in future mixed fleet networks.

Lioris et al., developed a fluid network model that made two predictions [Li-

oris et al., 2016]. The first is that if the saturation flow of every movement is increased by a factor γ , then the network can support a throughput increase of γ . The second prediction is that despite this increase in throughput, delay and travel time remain unchanged. Their case study included a Los Angeles network with 16 intersections and 73 links. The authors demonstrated that when signal configurations are held constant, a 100% increase in network demand leads to an average queue length increase of 76%. However, when the signal configuration is altered to adopt a control regime that permits 6 phase changes per cycle, the queue increase is 2% compared to the base case fixed signal configurations. This change still represents an increase of 99% if the alternate signalling regime is applied to the base case demand. The study concluded that employing an adaptive signal control with vehicle platooning may cause the storage capacity of links to be the limiting factor for network throughput. The limitation of this study is in its use of a mesoscopic simulator. The mesoscopic simulator accurately models arrival and departure rates at intersections, but queues on links are stored as point queues that discharge into the intersection at the saturation flow rate. By using a fully microscopic model in this chapter, an additional degree of realism in representing the interactive effects of adjacent intersections can be incorporated. The interactive effects alter arrival and departure patterns, especially in congested environments.

Hu et al., proposed a “Polite Lane-Changing” protocol to achieve optimal performance when considering both network efficiency and network safety [Hu et al., 2012]. They investigated different lane-changing behaviour on efficiency and safety in the highway environment. The safest scenario was for CAVs to never change lanes, and the most efficient scenario was for CAVs to always change lanes. Neither option is practical, so the most optimal scenario was the hybrid scenario that used the Politeness Factor to evaluate the upstream effects of changing lanes. Hu et al., differentiated the behaviour of CAVs and regular vehicles by setting CAV reaction time to 0. Their study found that a CAV penetration of 50% with a politeness factor of 0.6 to 0.8 leads to the most optimal performance. Increasing CAV penetration further in highly saturated environments was found to have an adverse effect on the network, similar to the findings of Lioris et al. The limitation of this study was in its use of a strategic simulator as opposed to an agent-based microsimulator. Discrete-time microsimulators provide an additional degree of realism by emulating

agent interaction and network response such as queuing and platooning.

Maarafi had similar findings by emulating the behaviour of CAVS in the Vis-sim microsimulator [Maarafi, 2015]. CAV behaviour was emulated by altering the parameters of the Wiedemann car-following models, and experimentation was conducted on a calibrated model of the I-79 motorway. Maarafi found that a 60% CPR led to a 17% improvement in motorway throughput. Their study also found that increasing demand by 40% from the base case leads to a throughput decrease of 18%. Maarafi attributes this behaviour to the blocking of onramps and offramps caused by human vehicles. This chapter differs through its implementation of a custom framework for CAV behaviour as opposed to changing default software parameters. Also, the scope is extended by also considering the effect of segment weaving proportions and merge geometric alignment.

Van Arem et al., found similar results when conducting a study using a similar methodology to the studies mentioned above [Van Arem et al., 2003]. The authors demonstrated that when the time headways generally maintained by human drivers are also adopted for CAVs (1.5s), a CPR greater than 40% caused performance to deteriorate in situations that also had high demand. Network elements such as lane drops and merges had the highest performance when vehicles using ACC were segregated from the fleet and placed in a separate lane. Adopting a significantly shorter headway (0.5s) showed improved traffic flow. This arrangement led to an improvement in traffic efficiency, with maximal lane flows ranging from 2100veh/hr to 2900veh/hr. This fleet consisted of 20% human vehicles, 20% of vehicles using ACC and 60% using AICC. The authors used a discrete-time based methodology to update various modules of their simulation model periodically. This chapter aims to extend current understanding by conducting an analysis based on real-time agent interaction in a microsimulation setting. This chapter also alters the proportion of vehicles at merge junctions and uses two distinct types of merge junctions for experimentation.

Liu et al., concluded that the throughput capacity of the merge junction would limit motorway throughput as opposed to the storage capacity of the onramp [Liu et al., 2018]. Their study involved developing a custom behavioural con-

trol algorithm for CAVs and emulating the behaviour in a microsimulation environment using Aimsun. The test environment consisted of one 4-lane motorway with demand incrementally increasing until motorway saturation. CAVs were introduced to the environment in 20% increments. Their study found that motorway throughput increases by 46% to 70% over current capacities at a 100% CPR [RMS, 2017]. However, the study found a reduction of 13% caused by bottlenecking of the onramp merge junctions. This chapter follows a similar methodology and extends the scope of the simulation environment to include both a traditional weaving section and a zipper merge section. A conventional weaving section that provides a temporary auxiliary lane has a higher capacity but also has higher embedded construction and maintenance costs. The zipper merge has the advantage of requiring a smaller infrastructure footprint, but sacrifices merge capacity. By investigating both arrangements, a better understanding of failure points and potential future bottlenecks can be derived, as opposed to only investigating a subset of geometric arrangements.

Alternatively, studies such as that conducted by Letter and Elefteriadou reported no evidence of a decrease in network performance [Letter & Elefteriadou, 2017]. They presented a longitudinal freeway merge control algorithm to maximise the travel speed for CAVs. The algorithm facilitated communication between vehicles and roadside units to optimise the trajectories of vehicles. They modelled four different demand scenarios, three different CAV penetration rates and three different safe time gap settings depending on driver aggressiveness. In the experimentation, the throughput was shown to be equal to the demand during undersaturated conditions and equal to the theoretical maximum achievable value during oversaturated conditions. The average travel speed, total travel time, and average travel time per distance improved between 3% and 7% for the free flow conditions. As the demand increased, the improvements in travel speed improved by 61.4%. The contribution of this chapter is to use agent-based discrete-time modelling to represent interactions between agents better. The limitation of strategic-based modelling approaches is the disengagement of the performance of adjacent intersections. Queueing at one intersection has no impact on the approach arms of adjacent intersections. This interaction downstream on the highway may potentially limit the benefits gained from merge points and weaving sections upstream.

Baskar et al., propose a model-based predictive control approach for coordinating CAVs [Baskar et al., 2008]. In this study, all CAVs are capable of intelligent speed adaptation, ACC, and dynamic route planning and guidance. The control architecture consists of high-level controls that provide information about the region and network, roadside units that control localised behaviour like on a highway section, platoon controllers that execute platoon manoeuvres, and vehicle controls that translate commands from the rest of the architecture into control signals for vehicle actuators. Their approach determines the state of the system and aims to predict the behaviour over a time horizon. An open-loop control optimisation problem is then solved to determine the actions required to minimise the adverse effects of vehicle actions, constrained by their optimality criteria. Using this approach, the authors found that on a 10km highway with two lanes and no onramps or offramps, total system travel time improved by 10%. The limitation of this study is that without conflicting traffic from onramps and offramps, an essential degree of realism when evaluating the impacts of cooperation is missing from the system. On-ramps form a critical source of vehicles that disrupt well-formed platoons, and offramps force cooperation through weaving and lane-changing. This chapter aims to address this by contextualising the motorway weaving and merging section with the surrounding supporting ramp infrastructure.

Several studies have empirically demonstrated the effect of increasing human vehicle demand leading to reductions in throughput on motorways. Cassidy and Bertini used observed data for the Gardiner Expressway in Ontario, Canada, to demonstrate that excessive demand resulted in a throughput decrease of between 4% and 20% on three separate occasions [Cassidy & Bertini, 1999]. A similar effect was again demonstrated by Cassidy and Rudjanakanoknad through observations of the Northbound 805 Freeway in San Diego, United States [Cassidy & Rudjanakanoknad, 2005]. This effect was also demonstrated for human vehicles by other studies ([Banks, 1991], [Hall & Agyemang-Duah, 1991], [Oh & Yeo, 2012], [Persaud et al., 1998], and [Srivastava & Geroliminis, 2013]). This chapter will quantify whether this effect persists for CAVs.

The literature indicates that some studies find CAV penetration and demand increase will lead to bottlenecking in the network and reductions in throughput, other studies show continued improvement. The studies that mention a

reduction in capacity are then further divided based on the cause of this reduction. While some studies attribute the bottleneck to link capacity, others attribute it to the intersection and merge junction throughput. This chapter adds to the current state-of-the-art by addressing the limitations of the studies mentioned above. The chapter contextualises the operation of the motorway and identifies the critical failure point of weaving sections and zipper merge junctions that are subject to changing demand, changing CAV penetration, and changing weaving proportion.

9.2 Experimentation Methodology

The literature review has identified two distinct findings regarding CAV operations and motorway sections. The first offers an optimistic outcome, where CAV operation increases motorway capacity. The other provides a cautionary warning, where collapses in merge junctions and ramps result in bottlenecks and reductions in capacity. These investigations lack a detailed assessment of the role played by the weaving and merging section of the motorway. The weave and merge section cause bottlenecking and delays and therefore requires a thorough investigation before commenting on motorway capacity.

This chapter investigates many of the gaps still present in literature. Firstly, this chapter assesses the impact of CAV fleet integration on the capacity of weaving sections on motorways. Weaving sections are often the lowest throughput component of motorways, so if their capacity does not also increase, then any increase in motorway section capacity is essentially meaningless. This chapter then investigates the impact of increasing congestion on weaving section performance, assessing key metrics such as travel time and speed, for changing levels of CPR, demand, and weaving proportions.

Next, this chapter assesses the performance of a motorway zipper merge, where the number of upstream lanes is greater than the number of downstream lanes. Finally, this chapter concludes with an investigation into the effect on motorway performance of segregating the vehicles into different classes and instituting a priority hierarchy. The purpose of this investigation is to gain insights into how network performance changes if different “brands” of CAVs are precluded from cooperating, or if CAV and infrastructure operators implement a “premium” priority service.

The remainder of this section explains the nature of the experimentation conducted in this chapter. It provides a detailed description of the modelling environments and the metrics used for evaluating network performance.

9.2.1 Experimentation Structure

As mentioned previously, this chapter conducts three distinct investigations into the effect of CAV behaviour on motorway weaving sections. This subsection explains the three investigations in greater detail.

Investigation 1: Effect of CAV Behaviour on Weaving Sections

The first investigation involves evaluating the impact of CAVs on motorway weaving section performance. This investigation incrementally loads the microsimulation area with demand to find capacity. Capacity is defined as the point where additional demand does not lead to an equal increase in throughput through the merge section. This process is repeated for CAV penetration increments of 25%, and a weaving proportion of 5%, 10%, 20%, and 40%.

The first investigation quantifies the effect of vehicle cooperation on weaving section performance. This thesis has designed CAVs to operate with altruism, giving way, and cooperating. How such a framework will perform in a congested environment is difficult to predict. The key difference between this methodology and those used in other investigations is that here the proportion of weaving vehicles is altered. Holding the weaving proportion constant does not provide an insight into the breadth of situations to which CAVs will be exposed. Additionally, if the weaving proportion is low, then vehicles essentially do not interact with one another, meaning that the cooperative capabilities of a framework are not stressed.

Investigation 2: Effect of CAV Behaviour at Zipper Merge Sections

The second investigation assesses the performance of a zipper merge under varying CAV penetration. A zipper merge consists of two lanes upstream seamlessly merging into a single lane downstream. Vehicles in both lanes offset their positions from one another to ensure that trajectories do not cross at the downstream merge point. Vehicles from both lanes alternate as they

move through the merge point, hence the name zipper merge. This style of merge is often used when two major highway or arterial segments join.

The purpose of this investigation is to assess the benefit derived by the trajectory forecasting cooperative merge algorithm developed in Section 5.1.4. The algorithm forecasts vehicle trajectories upstream and adjusts vehicle kinematics, to prevent conflicts at the downstream merge point and create a smoother and conflict-free merging experience. The algorithm should improve merge point throughput and average speed. Performance improvement may be derived from well-informed conflicting vehicles performing minute kinematic adjustments upstream, rather than large braking actions at the merge point.

Investigation 3: Effect of Conditional CAV Cooperation on Weaving Sections

The final investigation assesses the impact of selective cooperation at a merge junction. A 100% CAV fleet is incrementally segregated into 2, 3, 4, and 5 different vehicle classes. The different vehicle classes are selective in their choice to cooperate with another vehicle class. For simplicity, the vehicle classes are ordered chronologically, and a vehicle will not cooperate with another vehicle that has a lower class number.

This investigation assesses the impact of “premium infrastructure”. Uber already segregates their market by offering more premium vehicles for a higher price. Infrastructure is also segregated, where access to certain motorways is granted only for toll-paying customers. For these reasons, it is not unreasonable to expect segregation in the CAV market, where vehicle or infrastructure operators provide customers with “premium” paths and priority through the network. Additionally, this investigation assesses the impact of vehicle manufacturers denying cooperation requests from other vehicle brands. Stratification of the transport fleet is emulated by dividing the microsimulation vehicle fleet into different vehicle classes and allowing vehicle types to reject the cooperation requests of lower-ranking vehicles.

9.2.2 Modelling Environment

Figure 47 shows the two distinct modelling environments used in this study;



(a) Model 1 is used for investigation 1 and 3.



(b) Model 2 is used for investigation 2.

Figure 47: Modelling environments used for the different investigations.

Environment 1 is a weaving section where two major arterial roads or motorways interact briefly, before segregating downstream. During the $300m$ long interaction area, the vehicles change lanes to position themselves in the appropriate lane before the downstream segregation. Environment 2 differs in that the number of downstream lanes is less than the number of upstream lanes, meaning that one lane from each approach arm merges and vehicles collide at a conflict point. Both environments have an approach arm of $230m$ and a speed limit of $80km/hr$. Detectors for collecting network performance metrics are placed on all lanes, downstream of the merge point.

To prevent collisions in Environment 2, Figure 47b demonstrates the conflict area used to assess gap-acceptance for human vehicles. Section 5.2.2 provides a detailed explanation of the operation of conflict areas. CAVs operating under the framework developed in this thesis do not obey the conflict area, but instead, operate under the algorithm proposed in Section 5.1.4.

The base demand matrix for Environment 1 is provided in Table 21, and is symmetrical between both approach arms in all environments. The base case matrix is multiplied by a scaling factor when used in this study. Figures in this chapter that refer to a demand multiplier are referring to scaling of these base case matrices. Environment 2 only has 1 exit, so a matrix is not necessary. The base case volume for Environment 2 is $50veh/hr$, $25veh/hr$ from both approach arms. Both environments have a $30min$ warm-up period

which contains 80% of the demand of the 1hr peak modelling period used for assessment. The base case matrix for Environment 1 is as follows;

Environment 1					
5% Weaving Proportion			10% Weaving Proportion		
O\D	North Exit	South Exit	O\D	North Exit	South Exit
North Entry	50	2.6	North Entry	50	5.6
South Entry	2.6	50	South Entry	5.6	50
20% Weaving Proportion			40% Weaving Proportion		
O\D	North Exit	South Exit	O\D	North Exit	South Exit
North Entry	50	12.5	North Entry	50	33.3
South Entry	12.5	50	South Entry	33.3	50

Table 21: Base case OD matrices used for the modelling Environment 1.

50 iterations of each environment were run using a 0% CPR (base case), and a demand DoS of 0.9, to establish the model stability of both environments. The random seed number was randomly generated from a uniform distribution from 1 to 99,999. The standard deviation as a proportion of the median of 50 seeds for the critical network metrics (average delay, average travel speed, average travel distance, and average travel time) is provided in Table 22;

	Average Delay	Average Speed	Average Travel Distance	Average Travel Time
5% Weaving	21%	8%	0%	9%
10% Weaving	21%	7%	0%	8%
20% Weaving	23%	8%	0%	9%
40% Weaving	43%	9%	0%	12%
Environment 2	17%	6%	0%	6%

Table 22: Stability results for both modelling environments.

The stability analysis demonstrates that the modelling environments show low variability between random seeds, except for the average delay metric. Variability in average delay is expected when modelling a congested environment operating at a high DoS. The distribution of gaps can either provide lucrative opportunities for fortunate vehicles or cause vehicles to be severely hindered by the absence of gaps. This behaviour is further demonstrated by average

delay variability increasing as the weaving proportion in the network increases.

9.2.3 Result Evaluation Metrics

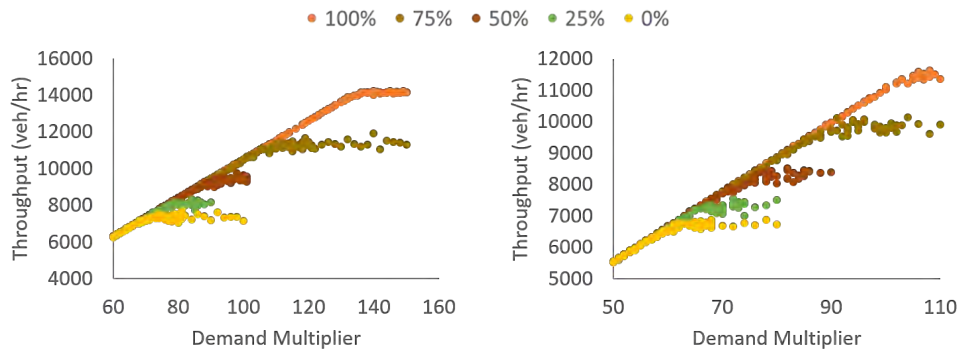
Data is collected from all environments using data collection points and travel time counters. The data collectors are placed 1m downstream of the merge points. They record several aggregated fleet metrics such as average vehicle speeds, average acceleration, average queuing delay, and the number of vehicles traversing the counter. The travel time counters range the span of the corridor and provide aggregate average travel times from each origin to each destination. The results in Section 9.3 report average travel speed, average travel time, and queueing delay resulting from each investigation and case.

9.3 Experimentation Results


This subsection presents the modelling results, segregated by investigation.

Investigation 1: Effect of CAV Behaviour on Weaving Sections

Investigation 1 assess the impact of CAV behaviour, as developed in this thesis, on the performance of motorway weaving sections. The investigation begins by incrementally loading the network to establish capacity. Figure 48 demonstrates the change in throughput in response to the change in demand for each CAV penetration, segregated by weaving proportion;



(a) Throughput vs Demand for each CPR and 5% weaving proportion. (b) Throughput vs Demand for each CPR and 10% weaving proportion.



(c) Throughput vs Demand for each CPR and 20% weaving proportion. (d) Throughput vs Demand for each CPR and 40% weaving proportion.

Figure 48: Throughput vs Demand for the purposes of determining capacity for each CAV penetration and each weaving proportion.

Capacity is defined as the point along the curve marking a substantial change in gradient, from constant to 0. This point indicates that additional demand is not traversing the weaving section and capacity has been reached. Table 23 below provides the demand multiplier corresponding to the capacity in the modelling environment. The demand multiplier is used to scale the base case matrices provided in Table 21 for the different DoS scenarios;

CPR	Weaving Proportion			
	5%	10%	20%	40%
0%	75	64	49	29
25%	82	68	54	33
50%	97	78	60	35
75%	112	91	66	36
100%	137	107	72	33

Table 23: Environment 1 capacity for each CPR and weaving proportion.

The first trend identified by the capacity assessment is that as CAV penetration increases, so does the environment capacity. The environment contains no traffic signals and vehicles are generated into the environment at maximum speed and with a large headway controlled by the microsimulator. This means that vehicles do not have the opportunity to form tight platoons with minimum headway, as followers do not have the opportunity to catch up to leaders. Therefore, any benefit gained in throughput and capacity is solely the result of communication and efficient lane-changing. CAVs are still able to change

lanes at lower headways and cooperate to create space.

Table 24 provides the change in capacity relative to the base case;

CPR	Weaving Proportion			
	5%	10%	20%	40%
25%	9%	6%	10%	14%
50%	29%	22%	22%	21%
75%	49%	42%	35%	24%
100%	83%	67%	47%	14%

Table 24: Change in capacity relative to the 0% CAV base case.

The other trend identified by this assessment is that when weaving proportion is high, increasing the CAV proportion results in reduced weaving section throughput. This outcome contradicts many of the results previously found in this thesis, where increasing CAV proportions has resulted in greater throughputs. This specific outcome is seen when the weaving proportion increases to 40%, a number that is unreasonably high and included for the purpose of stress-testing both the CAV emulation algorithm and road infrastructure. The reason why a high weaving proportion adversely affected vehicle throughput is due to the cooperative and altruistic design of the CAV control algorithm. Consider that when 4-lanes worth of traffic travel through 4 lanes with no interaction and weaving, each lane has an effective utilisation rate of 100%. Consider now that in a poorly designed network, the same 4-lanes of traffic are forced to all change lanes and weave through the middle two lanes. The weaving lanes now have an effective utilisation rate of 200% each, with the remaining lane having a utilisation rate of 0% each. These scenarios represent extreme cases. The modelled case has a 40% weaving proportion, meaning that the two weaving lanes have a utilisation rate of 140% each.

The additional demand on the lane abides by the cooperative and altruistic nature of the algorithm, meaning that all CAVs are decelerating to give way to other CAVs also navigating through the weaving lanes. Normal human behaviour would force merging vehicles to wait in congested environments, resulting in low travel time for high demand links and higher travel times for low demand links, therefore resulting in a low weighted average travel time. The CAVs, however, are designed to minimise the maximum travel

time, resulting in an even and equitable distribution in delay, leading to a worsening of performance in environments with significant vehicle interaction and weaving.

Having determined environment capacity, this study now assesses a range of scenarios. $CPR \in [0\%, 100\%]$ in 25% increments, $DoS \in [0.1, 1.2]$ in 0.1 increments, weaving proportions of 5%, 10%, 20%, and 40%, with 20 random seeds for each of these cases. This experiment setup results in a total of 4,800 iterations. The results for the average network speed for each weaving proportion is provided in Figure 49;

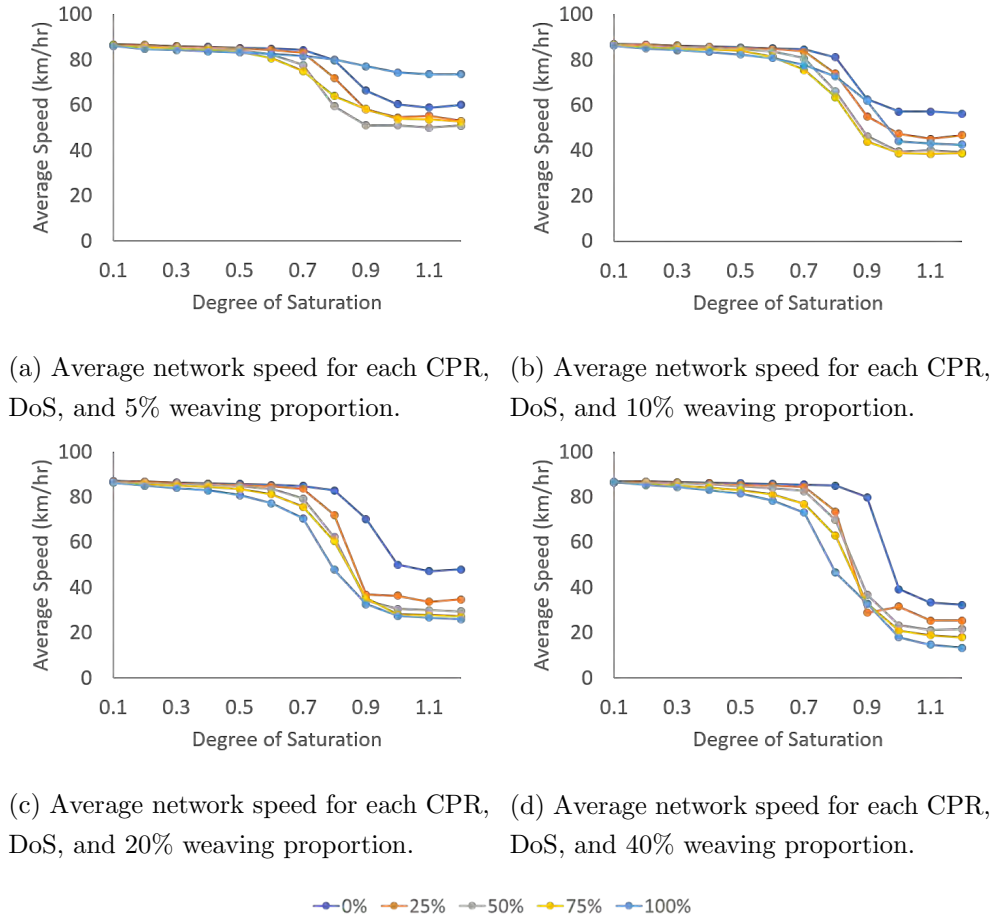


Figure 49: Average network speed for each CPR (each line), DoS, and weaving proportion(each subplot).

Figure 49 demonstrates that when the weaving proportion is low (Figure 49a), a 100% CAV network outperforms a mixed fleet and human network, exhibit-

ing a speed that is 22% higher than the 0% CAV network and 46% higher than the mixed fleet networks. As the proportion of weaving vehicles increases, the 100% CAV scenario quickly falls to the worst-performing environment, the reason for which has already been discussed above. While the performance of all mixed fleet scenarios drops when weaving proportion increases, the higher CAV scenarios are more severely impacted in the high demand situations.

The queuing delay only considers the time that vehicles spend queueing. The queue state is initiated when a vehicles speed drops below $5km/hr$. Figure 50 provides the average queuing delay for each weaving scenario;

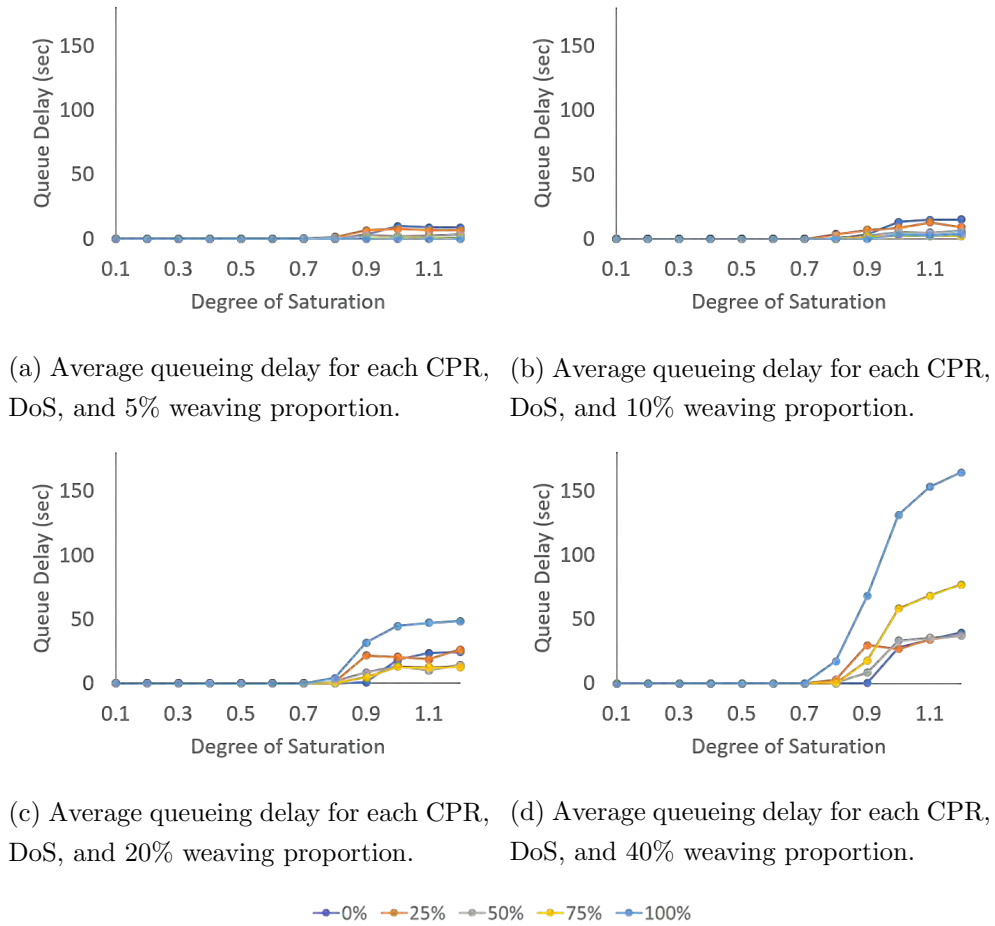


Figure 50: Queueing delay for each CPR (each line), DoS, and weaving proportion (each subplot).

The queueing delay in Figure 50 indicates a similar outcome to the average

network speed metric. Lower CPRs show a queueing delay at low levels of weaving (5% and 10%), but the 100% CAV fleet avoids a queueing delay. Better performance of the 100% CPR fleet results from its design to minimise maximum delay in cooperative environments. While the human vehicles in mixed fleet environments offload delay to other network agents and maintain a higher average speed, the 100% CAV scenario distributes the delay among all vehicles. This behaviour, although results in a lower average speed, keeps the speed of all agents greater than the threshold defining queueing. However, the high weaving scenarios oversaturate the balance of delay distribution and queueing increases to beyond that in the self-serving human vehicle approach. Finally, Figure 51 provides the average travel time results for each scenario;

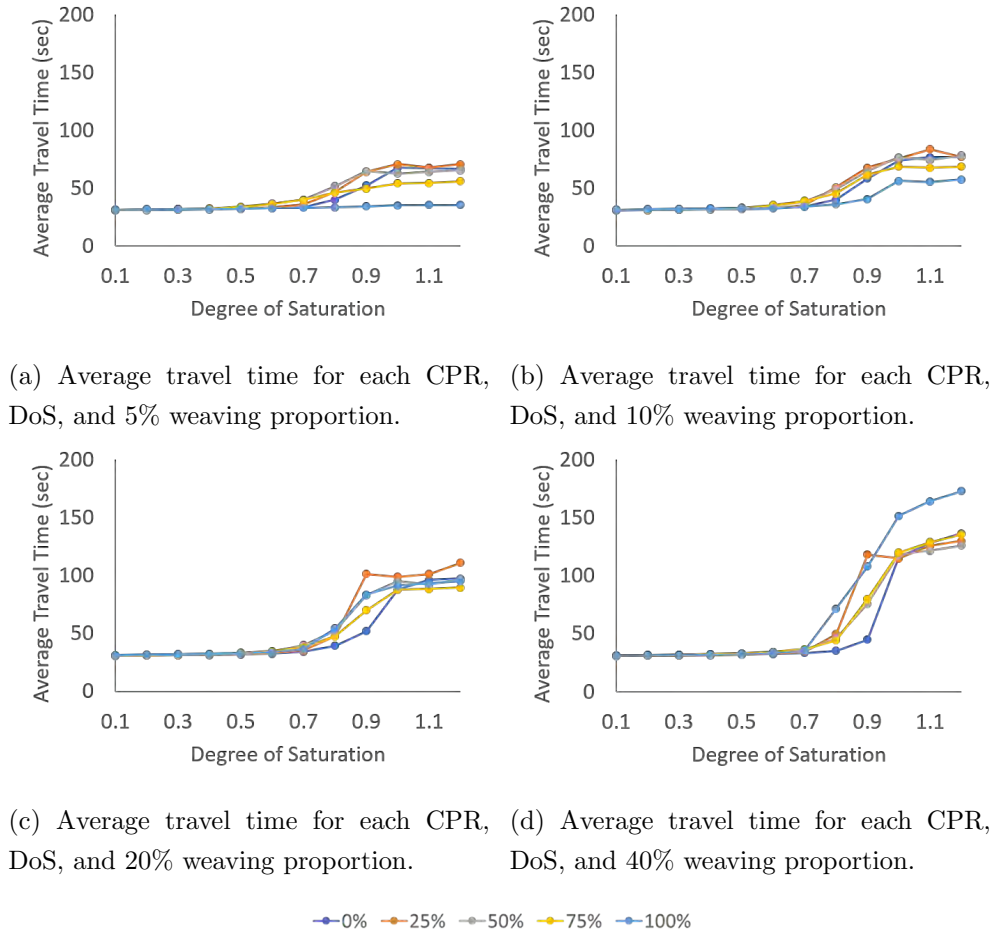


Figure 51: Average travel time for each CPR (each line), DoS, and weaving proportion (each subplot).

Figure 51 confirms the outcome demonstrated by the other two metrics. A selfish behavioural attitude is appropriate in high weaving scenarios to reduce average travel times. The 100% fleet shows better or on-par performance to that of the mixed fleet until weaving proportions reach 40%. The average travel times presented in Figure 51 provide unique insight. While the other figures demonstrate that a 0% CAV network outperforms the mixed fleet environments in terms of average travel speed, the average travel time is significantly lower between DoS of 0.7 and 1.0. This outcome confirms that the tendency for human vehicles to avoid cooperation in weaving segments results in better average speeds (as the minimum value of speed is bound) but leads to worse overall network travel times (as delay can indefinitely increase).

Investigation 2: Effect of CAV Behaviour at Zipper Merge Sections

Investigation 2 assess the impact of CAV penetration on zipper merges, where the first-in-first-out style of scheduling is generally adhered to by human drivers. The capacity of the merge for each CAV penetration is found in the same way as it was for Investigation 1. Figure 52 provides the throughput vs demand plot for each CAV penetration scenario;




Figure 52: Environment 2 capacity for each CPR and weaving proportion.

Once again, capacity is defined as the point where the gradient of the line changes. The reason why Figure 52 does not exhibit a horizontal-like curve after reaching capacity is due to the geometric nature of the zipper merge en-

vironment. This environment has four lanes upstream and three lanes downstream. While two of the upstream lanes are forced to merge into a single lane, the other two remain unobstructed, meaning that additional demand continues to flow unhindered in the outer two lanes while the inner two lanes are funnelled through a merge. For this reason, Figure 52 displays a reduction in the curve gradient as opposed to a horizontal line demonstrated by the previous environment. Table 25 provides the quantitative capacity values for this environment, as well as the capacity change relative to the base case;

CPR	0%	25%	50%	75%	100%
Capacity	108	122	144	185	250
Change from Base		13%	33%	71%	131%

Table 25: Change in the Environment 2 capacity for each CPR case relative to the 0% CAV base case.

The second environment shows a clear and predictable trend, as CPR increases, as does the capacity of the zipper merge. This outcome is the direct result of the trajectory forecasting cooperative merge algorithm proposed in Section 5.1.4. The algorithm forecasts the trajectory of vehicles and adjusts their kinematics upstream, avoiding conflicts downstream. The human vehicles, however, operate using conflict areas and occasionally must come to a complete stop in search of an appropriate gap, significantly limiting throughput. The cooperative algorithm shows an increase of 131% in throughput for a 100% CAV fleet compared to a human fleet.

Having determined the environment capacity, this study now assesses a range of scenarios where $CPR \in [0\%, 100\%]$ in 25% increments and $DoS \in [0.1, 1.3]$ in 0.1 increments. The weaving proportion is redundant here as demand is equal on both approach arms, which results in the highest possible number of vehicle conflicts at the merge point. Additionally, multiple seeds are not used in this investigation. Computational time for trajectory forecasting is high due to the excessive data transfer between different modules, programming languages, and storage media. A modelling iteration in Investigation 2 is approximately 240 times longer than that in Investigation 1, running at slower than real-time. For this reason, a single seed is used in this investigation. The experiment setup results in a total of 60 iterations.

The experimentation results are provided below in Figure 53 for average travel speed, average queue delay, and average travel time;

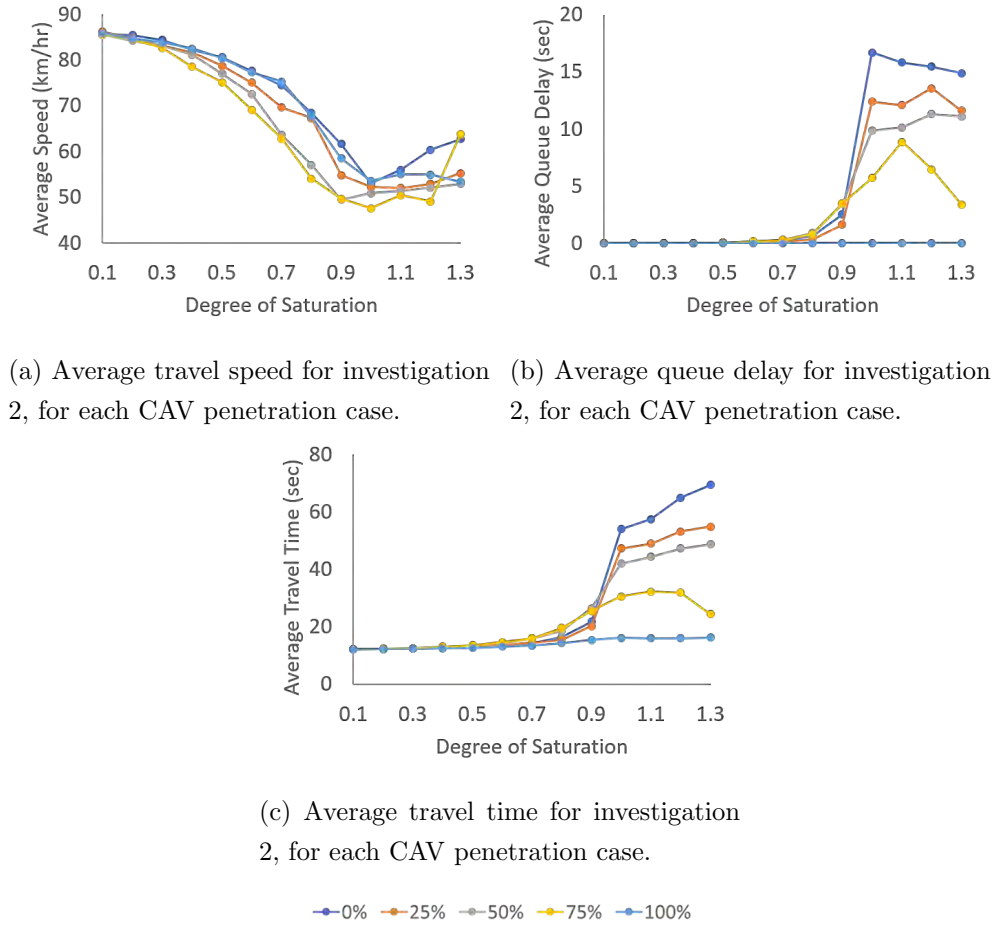


Figure 53: Results for investigation 2, for each CAV penetration rate.

Figure 53 demonstrates significant improvements in vehicle performance. Figure 53a shows that the average speed of vehicles between the 0% CAV and 100% CAV cases marginally varies. However, both cases far outperform the mixed fleet case, having an 18% and 5% higher average speed in the DoS 0.7 and DoS 1.0 cases, respectively. The 0% CAV and 100% CAV penetration cases have one source of inefficiency each, whereas all mixed fleet environments have two. The inefficiency in the 0% CAV scenario is that all vehicles on the secondary approach must give way to all vehicles on the primary approach, leading to excessive queueing and delay on the secondary approach. The source of inefficiency in the 100% CAV case is that vehicles must negotiate

space to reduce overall maximum delay, leading to inefficiencies for vehicles with altered trajectories. The mixed fleet scenario, however, is subject to both sources of inefficiency. The CAVs must not only give way to other conflicting CAVs, but also to human vehicles, meaning that both the primary and secondary approach is restricted. To remediate this problem, the CAVs need to be designed to be more aggressive and not always play the role of the passive agents in a human-CAV interaction.

Figure 53b indicates that despite the 0% CAV case having near highest average travel speeds, all CAV cases far outperform the 0% base case in average queueing delay. The trajectory forecasting algorithm is designed to adjust vehicle kinematics before the merge point, which leads to reduced braking and hence reduced start-stop conditions. The continuous flow of vehicles not only increases throughput but also reduces the time vehicles spend in queues. This outcome is reflected in Figure 53c, which demonstrates that despite having lower travel speeds, CAVs still spend less time queueing and traversing the merge junction. At a DoS of 1.0, the 100% CAV case had a 70% lower travel time than the 0% CAV case, despite showing near-identical average travel speeds at the data collection point directly after the merge zone.

Investigation 3: Effect of Conditional CAV Cooperation on Weaving Sections

Investigation 3 assesses the impact of selective cooperation on weaving section performance. Infrastructure funding is increasingly from the private sector, where companies own private toll roads, private bus services, private train lines, and soon to be privatised CAV fleets. Private transport infrastructure often undermines the purpose of transport, which is to provide accessible and equitable mobility to all, by segregating the market and creating premium services. The question should then be asked how the infrastructure would perform if CAVs were also segregated, with premium services being offered priority treatment in the network.

The environment used in Investigation 3 is the same as that in Investigation 1, so capacity does not need recalculation. This investigation segregates the CAV subset of the fleet into vehicle classes. The classes are ordered, and a vehicle in

a higher class will refuse to cooperate with a vehicle in a lower class. Human vehicles retain the highest class and receive cooperation from all CAVs. The CAV portion of the fleet is equally divided among all CAV classes.

The investigation assesses $CPR \in [0\%, 100\%]$ in 25% increments and $DoS \in [0.1, 1.2]$ in 0.1 increments, number of CAV vehicle classes varying from 2 to 5 in increments of 1 class, and weaving proportion of 5%, 10%, 20% and 40%. Each case is run using 3 random seeds. This experimentation structure results in a total of 2,880 iterations. Table 26 to Table 29 below provide the travel time (sec) results for each of the cases, with each table proving the results for a specific weaving proportion;

Weave 5%		DoS											
Classes	CPR	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	1.1	1.2
2	0%	30.9	31.0	31.3	31.5	32.0	32.5	34.0	39.5	51.3	64.5	65.5	68.8
	25%	30.7	31.1	31.6	31.9	32.2	33.0	36.0	43.5	54.7	62.4	66.7	66.5
	50%	30.9	31.2	31.7	32.0	33.1	35.4	41.9	49.8	59.4	63.7	62.8	65.3
	75%	31.2	31.5	32.0	32.5	33.0	38.6	42.1	46.4	55.1	56.6	57.1	56.6
	100%	31.3	31.8	32.1	32.3	32.5	32.9	33.5	34.1	35.3	37.0	36.5	36.7
3	0%	30.7	31.2	31.3	31.4	32.1	32.6	34.2	39.4	55.9	66.9	66.1	68.2
	25%	30.9	31.1	31.5	31.6	32.4	33.3	37.4	41.3	53.7	67.4	64.5	68.0
	50%	31.0	31.3	31.8	32.0	32.7	34.8	39.7	51.4	61.3	63.4	64.9	63.0
	75%	31.1	31.5	32.0	32.5	33.2	37.7	43.6	50.1	58.0	59.1	58.6	59.6
	100%	31.3	31.8	32.1	32.3	32.4	32.8	33.2	34.0	36.0	37.0	37.8	36.9
4	0%	30.8	31.1	31.3	31.5	32.0	32.3	33.8	39.1	53.5	65.7	64.7	67.9
	25%	30.8	31.2	31.4	31.8	32.1	32.9	34.3	45.1	54.8	65.1	67.2	66.6
	50%	31.0	31.3	31.6	31.9	33.0	35.2	42.3	52.9	61.6	63.4	64.1	62.2
	75%	31.1	31.4	31.9	32.5	33.5	36.5	44.6	46.8	54.7	59.5	58.6	59.0
	100%	31.3	31.8	32.1	32.2	32.6	32.8	33.4	33.9	35.2	36.9	36.8	37.2
5	0%	30.7	30.9	31.2	31.5	31.8	32.5	34.4	41.3	57.7	68.2	69.1	67.1
	25%	30.9	31.1	31.5	31.8	32.2	33.3	37.0	41.5	52.4	65.8	69.0	65.4
	50%	31.1	31.3	31.6	32.3	33.1	35.7	43.4	49.5	60.6	64.8	64.7	63.6
	75%	31.0	31.6	31.8	32.2	33.7	37.6	44.4	50.9	57.3	60.1	58.8	59.7
	100%	31.3	31.8	32.1	32.3	32.5	32.8	33.3	34.2	36.0	37.3	36.9	37.9

Table 26: Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 5% weaving proportion.

Weave 10%		DoS											
Classes	CPR	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	1.1	1.2
2	0%	30.8	31.0	31.1	31.5	31.8	33.0	33.9	40.2	50.5	76.8	75.2	77.7
	25%	30.8	31.0	31.3	31.5	32.2	33.0	35.1	42.8	59.3	71.9	78.4	76.4
	50%	30.9	31.3	31.7	32.0	32.5	34.3	39.3	45.1	73.7	75.8	73.7	77.4
	75%	31.0	31.5	31.9	32.3	33.5	35.0	41.0	51.0	70.5	72.5	75.3	71.7
	100%	31.2	31.8	32.1	32.5	32.8	33.6	35.1	39.3	59.6	62.6	67.3	67.6
3	0%	30.7	30.9	31.3	31.5	31.9	32.8	34.9	39.7	50.2	74.8	77.4	78.9
	25%	30.9	31.1	31.3	31.8	32.2	33.1	35.0	39.3	56.4	72.8	76.3	76.4
	50%	31.0	31.3	31.5	31.9	32.8	33.7	40.0	49.7	64.4	76.0	78.3	75.5
	75%	31.2	31.6	31.8	32.3	33.2	36.1	39.6	52.9	70.4	73.6	73.4	73.1
	100%	31.4	31.7	32.1	32.4	33.0	33.5	34.7	38.6	65.9	71.1	72.2	72.8
4	0%	30.8	31.1	31.1	31.5	31.8	32.7	33.7	39.5	50.1	72.3	77.7	77.7
	25%	30.8	31.0	31.3	31.7	32.1	33.3	34.9	38.8	57.1	72.9	75.5	75.9
	50%	31.0	31.3	31.5	32.1	33.1	34.6	38.1	45.7	64.4	77.2	81.2	79.2
	75%	31.1	31.4	31.8	32.2	33.4	35.0	40.6	55.5	67.7	74.7	79.4	78.3
	100%	31.4	31.7	32.2	32.5	33.0	33.7	34.8	37.5	61.9	72.4	78.0	76.2
5	0%	30.8	30.9	31.2	31.4	31.8	32.7	34.2	40.6	58.9	75.1	78.0	79.7
	25%	30.8	31.1	31.3	31.6	32.2	33.1	35.0	40.5	55.1	74.2	77.2	77.9
	50%	31.0	31.2	31.5	31.9	32.9	35.1	38.7	51.1	64.5	74.6	76.8	76.6
	75%	30.9	31.5	31.8	32.3	33.3	35.7	42.8	51.9	69.7	78.8	76.0	73.9
	100%	31.4	31.7	32.1	32.5	33.0	33.5	34.5	45.2	68.8	71.1	77.7	75.0

Table 27: Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 10% weaving proportion.

Weave 20%		DoS											
Classes	CPR	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	1.1	1.2
2	0%	30.8	30.9	31.1	31.5	31.9	32.4	34.0	37.7	55.7	90.1	94.6	101
	25%	30.8	31.0	31.4	31.6	32.3	33.4	36.4	45.0	76.9	88.6	96.0	98.8
	50%	31.0	31.2	31.5	31.9	32.7	34.0	38.6	51.3	91.3	97.0	94.4	97.9
	75%	31.1	31.4	31.9	32.3	33.0	35.3	39.2	56.0	91.8	96.6	99.1	99.3
	100%	31.2	31.7	32.1	32.6	33.3	34.6	44.7	92.4	108	111	121	114
3	0%	30.7	30.9	31.2	31.5	32.0	32.4	34.1	36.9	49.0	87.3	92.7	97.8
	25%	30.7	31.1	31.4	31.7	32.3	33.4	35.9	44.3	62.9	93.4	95.4	96.9
	50%	30.9	31.2	31.5	31.8	32.7	35.1	38.3	57.7	84.8	95.4	95.5	100
	75%	31.1	31.4	31.8	32.3	33.1	35.7	39.0	55.5	96.1	99.6	107	109
	100%	31.3	31.8	32.2	32.7	33.5	35.0	51.5	101	96.0	128	116	124
4	0%	30.8	30.9	31.1	31.5	31.9	32.8	34.1	39.2	49.1	84.8	95.1	100
	25%	30.8	31.0	31.3	31.6	32.2	33.0	35.6	45.4	80.9	90.9	94.0	98.1
	50%	30.9	31.2	31.5	32.0	32.4	34.5	38.7	52.0	94.6	94.7	97.6	99.6
	75%	31.1	31.3	31.8	32.3	33.3	35.0	41.1	82.5	100	103	108	106
	100%	31.2	31.7	32.1	32.6	33.5	34.8	40.0	94.0	99.1	130	140	124
5	0%	30.7	30.9	31.2	31.4	32.0	32.3	34.3	40.2	54.0	87.1	94.5	98.1
	25%	30.9	31.0	31.4	31.6	32.4	33.7	36.5	45.4	62.6	92.9	94.9	94.7
	50%	31.0	31.2	31.4	31.9	32.5	34.2	41.5	55.6	87.0	98.1	101	99.3
	75%	31.0	31.4	31.8	32.2	33.2	35.0	41.9	99.5	99.7	108	107	106
	100%	31.3	31.6	32.1	32.7	33.3	34.7	60.8	99.0	121	125	136	120

Table 28: Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 20% weaving proportion.

Weave 40%		DoS											
Classes	CPR	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	1.1	1.2
2	0%	30.8	30.9	31	31.4	31.7	32.5	33.4	35.5	43.3	115	131	139
	25%	30.9	31	31.3	31.6	32.1	33.4	35	43.3	95.6	105	123	122
	50%	31	31.1	31.4	31.7	32.4	33.6	36.4	43.6	103	122	130	161
	75%	31	31.4	31.6	32.2	32.8	34.3	41	69.8	132	146	154	159
	100%	31.2	31.7	32	32.4	33.2	34.6	67.5	113	202	172	222	232
3	0%	30.7	30.9	31	31.4	31.8	32.5	33.5	35.6	50.4	86.8	117	139
	25%	30.7	31.1	31.3	31.6	32.2	32.8	35.1	39.2	78.1	108	103	125
	50%	30.9	31.2	31.4	31.8	32.3	33.4	35.9	43.7	111	121	134	127
	75%	31.2	31.4	31.7	32.3	32.8	34.2	37.6	93.7	139	159	168	162
	100%	31.2	31.6	31.9	32.5	33.2	34.9	77.1	141	223	232	252	290
4	0%	30.7	30.9	31.1	31.4	31.8	32.4	33.5	34.9	44.5	119	129	142
	25%	30.8	31	31.3	31.5	32.3	33.2	34.4	39.3	68.2	114	124	124
	50%	30.9	31.1	31.4	31.8	32.4	33.6	35.8	55.3	99.2	103	129	129
	75%	31.1	31.3	31.6	32.1	33	33.8	45.2	88.1	139	158	156	182
	100%	31.2	31.7	32	32.4	33	34.6	36.9	160	192	209	247	245
5	0%	30.8	30.9	31.1	31.4	31.8	32.4	33.7	36.3	46.7	111	133	137
	25%	30.8	30.9	31.3	31.5	32.3	33.2	35	41.4	81.9	109	118	120
	50%	30.9	31.1	31.4	31.9	32.6	33.4	36.4	47.6	107	108	138	136
	75%	31	31.3	31.7	32.1	32.9	34	37.8	125	151	152	166	167
	100%	31.4	31.5	32	32.5	33.3	34.3	63.8	124	195	222	286	288

Table 29: Heatmap of average travel time for each CPR, DoS, CAV vehicle classes, and 40% weaving proportion.

The previous four heatmaps indicate that segregating the CAV fleet into vehicle classes and providing specific classes with priority does not impact network travel times in most cases. When the weaving proportion is below 20%, no difference is observed between the different vehicle class cases, regardless of the CPR or the DoS. As weaving proportion increases to 20%, increasing the number of vehicles classes adversely affects network performance when the network is operating near capacity (DoS greater than 0.8). At a DoS of 1.0, additional vehicle classes reduce travel speeds by 17%. The adverse impacts on system travel time increase as weaving proportion is increased further. At a 40% weaving proportion and DoS of 1.0, five vehicle classes increase travel time by 29% compared to the two-vehicle-class case.

This investigation demonstrates that while the infrastructure is under-utilised, segregating the vehicle class into premium services will not adversely affect weave section performance. Additionally, such a policy can also be implemented in busy areas where excessive weaving and lane-changing does not occur.

9.4 Discussion

Each investigation conducted in this chapter identified an interesting reaction by the network in response to a stimulus. The first investigation assessed the impact of CAVs on weaving sections as the proportion of weaving vehicles increased. This study found that cooperation and altruism in motorway junctions with high proportions of weaving vehicles lead to substantially worse performance than traditional selfish human behaviour. This finding indicates that a logical and structured hierarchy for negotiating space through weaving junctions would benefit all agents involved. While this thesis has demonstrated that the altruistic design of CAVs has great benefits in the arterial road network, it is apparent that the motorway junctions require greater complexity in defining cooperation conditions.

The second investigation demonstrated that CAV integration at zipper merges substantially increase the throughput of the junction, while also reducing the queueing delay experienced by agents. The merge junctions provide benefits as the trajectory is moderated before vehicles arrive at the merge zone. However, human vehicle involvement in mixed fleet settings substantially reduces the benefit of the system. This does, however, raise an interesting question regarding the capabilities of variable message signs (VMS), ITS, and other warnings signs offered to human drivers. Could the efficiency of this merge protocol exhibited with a 100% CAV fleet also be demonstrated by a mixed fleet where humans are given extensive information regarding imminent collisions at merge points and instructions on how to react? If so, then perhaps high levels of CAV penetration will not be necessary to derive the benefits observed in this study for zipper merges.

Finally, the third investigation assessed the impact of segregating the CAV fleet into vehicle classes and providing certain classes with priority over others. The investigation found that even when CAV fleets are subdivided into as much as 5 classes, networks experience no travel time impacts if they are operating under capacity, regardless of the weaving proportion. This indicates that if transport and mobility are privatised in an autonomous future, the segregation of services into “premium” and “basic” will not hinder network performance. Nor will a decision by manufacturers to prevent communication with one another. However, cooperation was shown to be highly beneficial

for vehicle kinematics (Chapter 6), intersection performance (Chapter 7) and road safety (Chapter 11).

High proportions of weaving vehicles were shown to be problematic in all investigations, substantially impacting the network performance. CAV behaviour, as designed in this thesis, relies on localised game-theory-type assessments where vehicles change lanes when a mutual benefit is found for the lane-changing vehicle and its follower in the target lane. This structure provides the vehicle with a microscopic assessment of its actions. A problem arises when CAVs enter the merge section and immediately broadcast their desire to change lanes. Surrounding vehicles respond by decelerating and creating adequate space, rendering the remaining 300m merge section under-utilised. If the game-theory style approach was extended over the entirety of the merge section using a centralised controller, the location of vehicles and conflicting weaving pairs could be distributed over a longer distance, severely reducing their impact on upstream vehicles and better utilising the available space.

9.5 Conclusion

This chapter investigated the impact of increasing CPR on weaving sections and zipper merges. Additionally, this chapter assessed the weaving section impact of CAV fleet segregation and preferential treatment for different vehicle classes. The assessment indicates that CAV behaviour can increase weaving section throughput by 83% with a 100% CAV penetration. However, this is assuming a relatively low weaving proportion of 5%. As the weaving proportion increases to 40%, altruistic and selfless CAV behaviour causes inefficiencies to arise in the network. Human behaviour concentrates the delay with the vehicles waiting on the secondary movement. In contrast, CAVs distribute the delay among all agents and result in an overall worse situation. This outcome, however, does not occur in zipper merges. The trajectory forecasting merge algorithm demonstrates that capacity can be increased by up to 131%. The average speed of vehicles was relatively the same between CPR scenarios. However, the time spent in queueing was substantially lower. Finally, segregating vehicles by class and then instituting a class hierarchy showed to provide little to no disbenefit to network performance, assuming that weaving proportions stayed below 20% and demand remained below saturation ($DoS \leq 0.8$).

10 Recalibration of the BPR Volume Delay Function for the Strategic Modelling of CAVs

Abstract: This chapter assesses the adequacy of the BPR volume delay function for the strategic modelling of Connected and Autonomous Vehicles (CAVs). Three testbed environments are simulated at 10% increments of CAV penetration rates (CPR) to observe network performance in mixed fleet environments. The microsimulation dataset is compared with the BPR travel time predictions to evaluate the need for recalibration. Where appropriate, the BPR modelling parameters are redefined as a function of the CPR. The predictive quality of the recalibrated model is then validated by comparing it against the BPR function on synthetic data. The numerical results indicate an overall improvement in travel time prediction using the recalibrated model, with a significant reduction in root mean square error from 15.16 to 8.86. The recalibrated model also outperformed the traditional BPR model in 67% of the 4620 cases used for validation, and better-predicted travel time by 5.43 times.

The implications of CAVs on transport infrastructure has been discussed in extensive detail throughout the previous chapters. This thesis investigated their effects on intersection and network performance, operational vehicle safety, motorway capacity, the contingency of road performance on vehicle cooperation, and optimisation of signalling for other network agents. However, it has not commented on how CAVs may affect the network planning process, which often relies on macroscopic citywide strategic modelling.

Transport modelling is critical in infrastructure planning and traffic management. The current transport planning framework does not consider the impact of CAVs on network performance. The Sydney Strategic Travel Model (STM) considers modes of travel by car, rail, ferry, bus, cycle and walk, and a combination thereof [Bureau of Transport statistics, 2012], failing to include future and emerging technology. CAV behavioural models have been proposed in the literature that rely on changing behavioural parameters such as sight distance, minimum acceptable gap lengths, and lane change preferences ([Tetamanti et al., 2016], [Qian et al., 2015], [Kamal et al., 2013], [Antoniotti et al., 1997], [Ioannou & Bose, 1999]). Such attempts do not appropriately capture the intricate and complex effects of CAV behaviour on network performance.

Strategic models often contain predefined parameters calibrated against local traffic data to capture field characteristics. The current suggested values for these parameters may be inappropriate for use with CAVs, as they were calibrated against localised historical data generated by human drivers. Simulation results from other studies suggest that 50% of the total travel time is reduced when at least 80% of the vehicles are equipped with V2X communication technology [Katsaros et al., 2011]. This thesis has demonstrated that CAVs increase intersection throughput by 110%, reduce intersection delays by 63%, and reduce queue lengths by 28%. With the same traffic volume, network delay induced by CAVs is lower than that of human vehicles and may result in the over-prediction of travel time in current strategic models. Travel time is a major contributing factor to vehicle routing in strategic models, the gross miscalculation of which, would render strategic models unusable.

This chapter investigates the validity of the continued use of current strategic modelling parameters for CAVs. Although many studies explored the implications of CAVs on transport planning ([Epting, 2019], [Beza & Zefreh, 2019], [Clements & Kockelman, 2017], [Harper et al., 2016a]), the impacts have been discussed more conceptually than empirically. This chapter aims to examine a common strategic model used in the current transport planning practice and propose necessary modifications to more accurately reflect CAV behaviour. This chapter focuses on the BPR volume delay function (VDF) in its form outlined in the HCM. It also provides a brief assessment and commentary of the Conical and Davidson functions. VDFs are models designed to mimic the deterioration in link travel speed as link volume increases. Their position in the strategic modelling framework is presented in Section 10.1.1, and their significance is outlined in Section 10.1.2.

The remainder of this chapter is structured as follows. Section 10.1 briefly describes the transport planning framework, reviews currently used volume delay functions, and new strategic models proposed for CAVs. Section 10.2 presents the methodology and framework to recalibrate the selected models and parameters. This section also presents the microsimulation environments and scenarios used for evaluation. Section 10.3 presents the results of this study, which includes the microsimulation results, mathematical predictions using the original strategic models, and the recalibrated models. Section 10.4

provides a discussion of the results and the limitations of the study, with possibilities for future research. Finally, Section 10.5 provides a summation and conclusion.

10.1 Review of the Relevant Literature

This section commences with a brief description of strategic modelling, to contextualise the use of VDFs and draw attention to their importance. A literature review of popular VDFs then follows. Finally, this section concludes with a review of recent attempts at modifying or proposing strategic modelling techniques for use with CAVs.

10.1.1 Strategic Modelling in Transport Planning

Transport modelling is an invaluable resource that helps inform governments of travel behaviour changes on a citywide scale, as a consequence of policy change or investment in civil transport infrastructure. The four-step model is a widely used demand forecast model which estimates future travel behaviour, network performance, and patronage on links. The four steps of the model are detailed as follows;

- **Trip Generation:** The first step uses macroeconomic data (population and employment) and land use data (zone boundaries and points of interest) to determine the number of trips produced by each zone.
- **Trip Distribution:** The second step distributes the trips by using the same macroeconomic and land use data to determine the relative attractiveness of each zone.
- **Mode Choice:** The third step uses generalised path costs to determine the capture rate of each mode between each origin-destination pair.
- **Route Assignment:** The final fourth step assigns the road traffic origin-destination matrix to the network through a process of equilibrating general path costs.

The four-step modelling process relies on developing a generalised cost function for each link and node in the network. By summing discretised link and node costs, a total path cost between origin-destination pairs is determined. The generalised path costs of competing routes are then equilibrated by adjusting vehicle routing volumes [Saw et al., 2015]. Link delays and node delays are

contingent on a range of factors including capacity, volume, heavy vehicle proportions, behavioural parameters, and localised parameters. The specific VDF for links addressed in this chapter is the BPR function, explained in greater detail in Section 10.1.2.

To contextualise the use and significance of the BPR function, refer to Figure 54. From this figure, it is apparent that determining the costs of paths in the network underpins the strategic modelling process. If the path costs are underestimated, this affects the attractiveness of the path and then affects route assignment and mode choice. If the performance of a major network motorway is then overestimated as a result of the initial error, it may have implications on the relative attractiveness of zones and affect the future distribution of citywide population and employment. The ramifications of calculation errors in the path costs can have significant implications on the outcome of business cases and the cost-benefit analysis of new transport infrastructure. For this reason, the appropriateness of current VDFs must be evaluated for applicability to CAV behaviour. Strategic modelling is structured as follows;

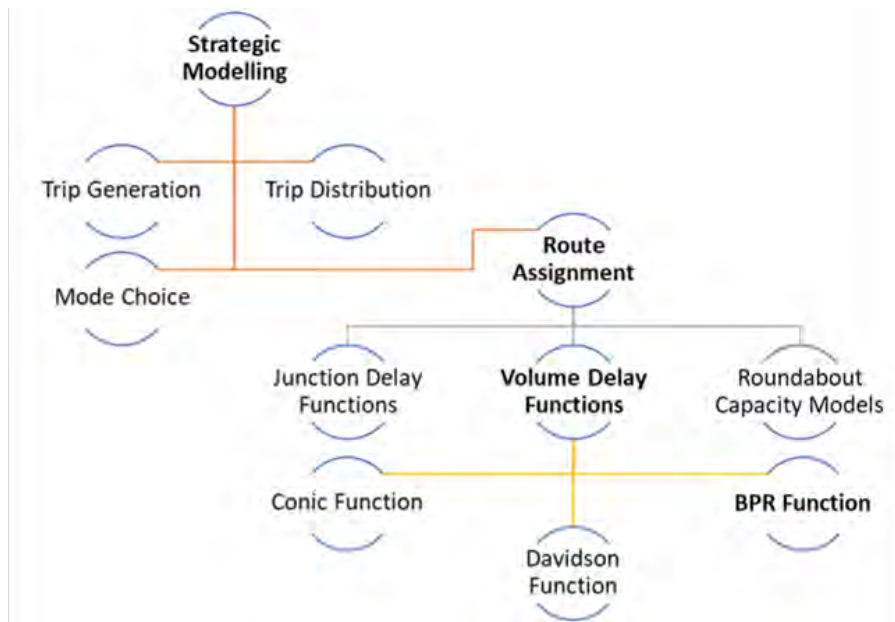


Figure 54: Contextualisation of the BPR function within the wider strategic modelling framework.

V2X technology grants CAVs access to information that was previously not available to human drivers, resulting in a behavioural shift in the vehicle fleet

[Wedel et al., 2009]. Additionally, V2X communication also enables cooperative manoeuvring, where vehicles travel coordinately under a centralised or decentralised decision-making strategy [Hobert et al., 2015]. This effectively mitigates factors of misjudgement, misinterpretation, and response time that are common in driver interactions. For these reasons, existing functions and parameters used for calculating link delay functions may not be appropriate and require investigation.

10.1.2 Volume Delay Functions and Capacity Models used in Macroscopic Strategic Modelling

This investigation into the applicability of current strategic modelling techniques to CAV behaviour focuses exclusively on VDFs. However, capacity is often a variable required by most, if not all, VDFs. For this reason, an appropriate way of determining link capacity must also be investigated. This section outlines currently used VDFs and techniques by which link capacity is determined.

Volume Delay Functions

VDFs describe the relationship between travel time (or cost) of a road link and the traffic volume. These functions are critical in most traffic assignment models by specifying the impact of road utilisation on the travel time. With an expected improvement in road capacity from CAVs, it is important to assess whether the currently proposed fundamental relationship between capacity and travel time still hold as the fleet transitions to CAVs.

The standard BPR function developed by the U.S. Bureau of Public Roads in the 1960s is defined as;

$$t = t_0(1 + \alpha(\frac{V}{C})^\beta) \quad (152)$$

Where, t is the link travel time, t_0 is the link travel time under free flow conditions, V/C is the volume to capacity ratio (similar to DoS), and α and β are calibration parameters, unique to each road geometry and environment.

The BPR function was developed by fitting a polynomial equation to a speed-flow curve observed on a single United States motorway [Manual, 2000]. It is

the most widely used function in transport modelling [Mtoi & Moses, 2014] due to its simplicity and minimal input requirements. Parameters α and β determine the shape of the function, and their values are often predefined based on assumptions and network characteristics. α regulates the magnitude of the travel time increase relative to volume increase. β controls the rate at which the link cost increases with the flow, to reach the magnitude defined by α . Higher values of α indicate that conditions on a particular link become much worse with increasing traffic volume, while higher values of β indicate that the road is unable to absorb the effects of increasing traffic and congestion effects become prominent sooner.

The Conical Volume Delay Function [Spiess, 1990] was introduced as an alternative to the BPR function and is defined as;

$$t = t_0(2 + \sqrt{\alpha^2(1 - \frac{V}{C})^2 + \beta^2} - \alpha(1 - \frac{V}{C}) - \beta) \quad (153)$$

Where, t is the travel time (or cost) of the studied road link, t_0 is the travel time under free flow condition, V/C is the volume to capacity ratio, and α and β are calibration parameters where $\alpha > 1$ and $\beta = (2\alpha - 1)/(2\alpha - 2)$.

The Conical VDF addresses inherent drawbacks of the BPR function, where the high exponent β value could lead to overflow conditions and loss of precision [Spiess, 1990]. High values of β in the BPR functions assign undue weight to overloaded links during the first few iterations of an equilibrium assignment, which can cause numerical problems. With high values of β , the estimated travel time for a link far under capacity becomes independent of the actual traffic volume as $(V/C)^\beta \rightarrow 0$. The equilibrium model consequentially locally reforms into an all-or-nothing assignment model and provides solutions that may not be unique to links.

Spiess suggests that the difference between the Conical function and the BPR function is small for $V/C < 1$, so the parameters from BPR could be directly transferred to the Conical function. Parameters α and β may require recalibration against traffic data generated from CAVs, to better assist the model's applicability to CAVs networks.

Davidson developed the Davidson Function in 1966 based on principles of

queuing theory [Akçelik, 1991]. The function is defined as;

$$t = t_0(1 + \frac{J_d \cdot x}{1 - x}) \quad \text{for } V < C \quad (154)$$

Where t is the travel time (or cost) of the studied road link, t_0 is the travel time under free flow condition, $x = V/C$ is the volume to capacity ratio, and J_d is the delay parameter.

The Davidson Function gained popularity over its ability to accommodate different traffic conditions and environments [Mtoi & Moses, 2014]. However, the function is unable to express travel time for traffic conditions where volume exceeds link capacity. For $V > C$, $1 - x < 0$, the travel time (t) for the link then decreases as traffic volume increases. This contradicts the fundamental speed-flow relationship, where an increase in traffic flow leads to a decrease in travel speed and an increase in travel time. Furthermore, Davidson's function has been a subject of discussion as the function implies the equivalence between free flow travel time and the reciprocal of the link capacity ($t_0 = 1/C$). The inherent inconsistency is unjustified from the viewpoint of delay formulation, and the ambiguity of parameter definitions also complicates the parameter calibration process [Akçelik, 1991].

Capacity Models

The HCM defines capacity as the maximum flow that can reasonably traverse the cross-section of a road segment. Capacity is an essential parameter in traffic state analysis, evident by its frequent occurrence in strategic modelling.

The Selected Maxima method states that capacity is the maximum observed flow over a period. This approach assumes that the actual road capacity is rarely observed due to the presence of external factors such as driver behaviour, weather conditions, and other factors that prevent idealised traffic performance. A deterministic maximum value gives the value of capacity with an aggregation of negative factors that reduce road capacity in accordance with the effect of localised conditions [Dervisoglu et al., 2009]. The calibration process estimates the deterministic maximum value and allows room for capacity reduction when investigating the impact of incidents.

The Fundamental Diagram method uses speed, flow, and density data to construct diagrams that display capacity. Mathematical models such as Greenshields and Van Aerdes can be used to fit the plotted data, with capacity being the maximum turning point of the fitted curve. Greenshields model is as follows [Rakha & Crowther, 2002];

$$\begin{aligned} u &= u_f - \frac{u_f}{k_j} k \\ q &= ku = u_f k - \frac{u_f}{k_j} k^2 \end{aligned} \tag{155}$$

Where, q is flow, k is density, u is the average space-mean speed, u_f is free flow speed, and k_j is the jam density.

Greenshields model assumes a linear relationship between speed and density. When plotting this model, the y-intercept is free flow speed, and the gradient is given by the ratio between free flow speed and jam density. By making use of the fundamental relationship $q = k\bar{v}_s$, a parabolic relationship is derived for the flow-density relationship. The highest point of the flow-density curve is the capacity(q_c), corresponding to the critical density (k_c).

Equation 156 outlines Van Aerde's method [Rakha & Crowther, 2002];

$$\begin{aligned} k &= \frac{1}{h} = \frac{1}{c_1 + \frac{c_2}{u_f - u} + c_3 u} \\ q &= ku = \frac{u}{c_1 + \frac{c_2}{u_f - u} + c_3 u} \\ c_1 &= \frac{u_f}{k_j u_c^2} (2u_c - u_f) \\ c_2 &= \frac{u_f}{k_j u_c^2} (u_f - u_c)^2 \\ c_3 &= \frac{1}{q_c} - \frac{u_f}{k_j u_c^2} \end{aligned} \tag{156}$$

Where, k is the density, h is headway between vehicles, c_1 is fixed spacing constant, c_2 is the first variable spacing constant, c_3 is the second variable spacing constant, u_f is free flow speed, and u_c is speed at capacity.

Van Aerde's method is based on a simple car-following model which uses the minimum headway distance between consecutive cars. Using the relationship $k = 1/h$, the model describes the headway as a combination of three terms;

a constant term, a term that depends on the difference between speed and free flow speed at any given time, and a term that depends on the speed at any given time. The Van Aerde model is used to fit the observed data, and capacity is determined by the maximum flow on the fitted speed-flow curve.

The Product Limit method, initially proposed by Kaplan and Meier in 1958, is used for lifetime data analysis such as estimating the fraction of patients living for a certain amount of time after treatment [Kaplan & Meier, 1958]. By considering a traffic breakdown as a failure event, when a sudden drop in traffic speed occurs due to demand exceeding capacity, the Product Limit method estimates capacity based on flow observations made over the observation period. The distribution function of capacity is given as;

$$F_c(q) = 1 - \prod_{i:q_i \leq q} \frac{k_i - d_i}{k_i} \quad \text{for } i \in \{B\} \quad (157)$$

Where, $F_c(q)$ is the distribution function of capacity c , q_i is traffic volume in interval i , k_i is the number of intervals with a traffic volume of $q_i \leq q$, d_i is the number of breakdowns at a volume of q_i , and B is the set of breakdown intervals where traffic is fluent in interval i but the average speed drops below the threshold speed in the next time interval $i + 1$.

The limitation of the Product Limit method is that it does not specify the type of distribution function required. Also, the value of the capacity distribution function will only be one if the product in the equation is zero. This outcome occurs when the maximum observed value of q belongs to B , when the flow is observed directly after a breakdown.

10.1.3 Recent Attempts at the Strategic Modelling of CAVs

Rather than assessing the validity of existing strategic modelling techniques, the literature has contributed new ways of modelling CAVs. These contributions range in complexity, from simple mathematical models to complex neural networks that require extensive data and calibration. This section provides a brief outline of the range of these contributions.

Multiclass Cell Transmission Model

Levin and Boyles developed a multiclass CTM to model dynamic traffic assignment for shared road scenarios with mixed fleets [Levin & Boyles, 2016]. A car-following model was then developed based on driver reaction times to estimate capacity and backwards wave speed. Consider a scenario where vehicle 2 follows vehicle 1 at a speed of u m/s with vehicle length l m. Vehicle 1 begins to decelerate at $t = 0$, and vehicle 2 decelerates at a rate of a , after a reaction time of Δt . The model describes the distance between vehicle 1 and vehicle 2 at time t using kinematics;

$$x_1(t) - x_2(t) = \begin{cases} u - \frac{1}{2}at^2 + L & \text{for } t \leq \Delta t \\ -at\Delta t + \frac{1}{2}a(\Delta t)^2 + L & \text{for } \Delta t < t \leq u/a \\ \frac{u^2}{2a} - ut + \frac{1}{2}a(t - \Delta t)^2 + L & \text{for } t > \frac{u}{a} \end{cases} \quad (158)$$

Where, u is the velocity, a is the acceleration, u/a is the time required to reach a complete stop (the position of vehicle 1 is constant for $t > u/a$), and L is the following distance.

Minimum distance between vehicle 1 and vehicle 2 occurs when both vehicles come to a stop, at $t = u/a + \Delta t$, when $t > u/a$. The minimum distance is then given by;

$$x_1(t) - x_2(t) = \frac{u^2}{2a} - u\left(\frac{u}{a} + \Delta t\right) + \frac{1}{2}a\left(\frac{u}{a}\right)^2 + L \quad (159)$$

To avoid collisions, the safe following distance needs to be greater than the sum of the minimum following distance given above and the vehicle length l ;

$$L \geq -\frac{u^2}{2a} + u\left(\frac{u}{a} + \Delta t\right) - \frac{1}{2}a\left(\frac{u}{a}\right)^2 + l = u\Delta t + l \quad (160)$$

Using the relationship between headway and density ($k = 1/L$) and the fundamental flow-density relationship ($q = ku$), capacity for free flow speed with maximum density is calculated using;

$$q_{max} = \frac{u_f}{u_f\Delta t + l} \quad (161)$$

To account for heterogeneous fleets, consider a different density, headway, and reaction time for each vehicle class m . Assuming that the vehicle proportion of each class (k_m/k) remains constant, the capacity for mixed fleets is;

$$q_{max} = \frac{u_f}{u_f \sum_{m \in M} \frac{k_m}{k} \Delta t_m + l} \quad (162)$$

Levin and Boyles also explored intersection capacity using larger conflict regions [Levin & Boyles, 2016]. The Conflict Region algorithm is a simplification of a reservation-based policy developed by Dresner and Stone designed for arbitrary vehicle prioritisation [Dresner & Stone, 2004]. Levin and Boyles modify the Conflict Region algorithm to accommodate LEMITM [Bento et al., 2013], which is a policy that allows reservation of space-time for all possible turning movements, and increases safety margins for human vehicles that lack the ability to communicate with the infrastructure. The developed algorithm follows the key assumption of LEMITM, that the intersection controller determines whether a vehicle waiting at an intersection is a CAV based on its ability to communicate with the controller digitally. The proposed model does not account for potential human errors. The Conflict Region algorithm omits the ability for the intersection controller to cancel granted reservations for CAVs if a human vehicle fails to enter the intersection in time.

Levin and Boyles conducted experiments using two classes of vehicles; human vehicles with a reaction time of 1s and CAVs with a reaction time of 0.5s. Vehicle lengths are set to 20ft. The total demand remains unchanged throughout the experiments, while CPR varies. Based on Equation 162, CAVs require 0.593 of the capacity that human vehicles require. In a single intersection study, the experiments assume a demand of 1200veh/hr at each intersection approach, with an assumed capacity of 1800veh/hr and a free flow speed of 60mi/hr. CPR increased in 10% increments, and each experiment had a 1hr duration with random vehicle departure times.

The results showed that average travel time decreased linearly with increasing CPR between 0% and 60%. The travel time is almost unchanged for CPR greater than 70%. Total travel time was selected as a measure of effectiveness, with the apparent trend showing that intersection capacity increases with increasing CPR. The minimal change in travel time above 70% penetration is explained by a sufficient increase in intersection capacity at that point to accommodate the network demand. Levin and Boyles's study provides insight into the potential network improvement with increased CAV penetration. However, the study's use of deterministic time headway that

may not adequately capture the stochasticity of traffic flow. Furthermore, the model does not consider vehicular interaction, as the model assumes vehicle behaviour to be independent of the type of vehicle preceding and following it. Finally, by not independently determining capacity, or increasing the degree of saturation, their experimentation reached undersaturated conditions by 70% CAV penetration and prevented showing any further improvement.

Markov Chain Model

Ghiasi et al., addressed the drawbacks of Levin and Boyles approach [Ghiasi et al., 2017]. Ghiasi et al., modelled mixed fleets on a one-lane highway segment, where vehicle types are distributed stochastically, and vehicle headways are distributed randomly depending on the corresponding vehicle type. This analytical stochastic model characterises traffic patterns by three critical factors; CPR, CAV platooning intensity, and mixed traffic headway settings. A Markov Chain model describes the vehicle type distribution in the mixed fleet scenarios while incorporating the two other parameters mentioned above.

Consider a stream of N vehicles along the highway segment, vehicle class is denoted by $A_n \in \{0, 1\}$, where $A_n = 1$ if the n^{th} vehicle is CAV and $A_n = 0$ if the vehicle is human. A_n is representative of the state variable at step n , with state-space $S := \{1, 0\}$. The CAV penetration rate P_1 describes the expected percentage of CAV in the traffic stream, where $P_1 := E(\sum_{n \in N} A_n / N)$. The percentage of human vehicles in the stream is then defined as $P_0 = 1 - P_1$, where $P_0 := E(\sum_{n \in N} (1 - A_n) / N)$. CAV platooning intensity is denoted by O , where $O = 1$ represents maximum platooning intensity where all CAVs are perfectly platooned in one group, $O = 0$ represents independent platooning where CAVs and human vehicles are randomly distributed, and $O = -1$ represents minimum platooning intensity where the number of CAV platoons are at a maximum.

Let t_{sr} denote the probability for a type-s vehicle followed by a type-r vehicle, so t_{10} would represent the probability for a CAV (1) followed by a human vehicle (0). t_{sr} for each interaction type is a function of the CAV penetration rate (P_1) and the degree of platooning (O), and is provided in Equation 163;

$$\begin{aligned}
t_{10}(P_1, O) &:= \begin{cases} P_0(1 - O) & \text{for } O \geq 0 \\ P_0 + O(P_0 - \min(1, \frac{P_0}{P_1})) & \text{for } O < 0 \end{cases} \\
t_{11}(P_1, O) &:= 1 - t_{10}(P_1, O) \\
t_{01}(P_1, O) &:= \begin{cases} P_1(1 - O) & \text{for } O \geq 0 \\ P_1 + O(P_1 - \min(1, \frac{P_1}{P_0})) & \text{for } O < 0 \end{cases} \\
t_{00}(P_1, O) &:= 1 - t_{01}(P_1, O)
\end{aligned} \tag{163}$$

The expected capacity of the Markov Chain model is given by;

$$\bar{q} = E(\frac{1}{h}) = E(\frac{N-1}{\sum_{n=1}^{N-1} h_n}) \tag{164}$$

Where, h_n is the time headway between vehicle n and vehicle $n+1$, and N is the total number of vehicles.

Due to the random headway distributions and the exponential number of vehicle type distributions, it is difficult to evaluate link capacity. For this reason, a closed-form analytical solution was proposed that approximates capacity. The closed-form solution is provided in Equation 165;

$$\bar{q} = \frac{N-1}{\sum_{n=1}^{N-1} E(h_n)} = \frac{N-1}{\sum_{n=1}^{N-1} h_{A_n A_{n+1}}} \tag{165}$$

The probability of the n^{th} vehicle of type-s, followed by a type-r vehicle is;

$$Pr(A_n A_{n+1} = sr) = P_s t_{sr}(P_1, O) \tag{166}$$

Then, the approximated expected headway can be calculated as $\frac{\sum_{n=1}^{N-1} h_{A_n A_{n+1}}}{N-1} = \sum_{s,r \in S} P_s t_{sr}(P_1, O) \bar{h}_{sr}$. Therefore, the closed-form of the capacity approximation is given by;

$$\tilde{q}(P_1, O, h) = \frac{1}{\sum_{s,r \in S} P_s t_{sr}(P_1, O) \bar{h}_{sr}} \tag{167}$$

Where h is the vector of expected headways; $h = [\bar{h}_{11}, \bar{h}_{10}, \bar{h}_{01}, \bar{h}_{00}]$.

The study shows that the approximate capacity is sufficient to represent actual capacity by investigating the theoretical relationship between the two capacity models and comparing numerical results from experiments.

Linear Regression Delay and Artificial Neural Network Model

Adebisi simulated the delay experienced by mixed fleet traffic at two signalised intersections by modelling different CPRs using the Vissim microsimulation [Adebisi, 2018]. The study developed a V2I communication algorithm that optimised travel dynamics by adjusting the speed and acceleration of the arriving CAVs based on their position and the state of the traffic signal, at each 0.1s time step. During each time increment, the algorithm;

- Collects vehicle data such as vehicle class, speed, acceleration, and position within a predefined communication range.
- Terminates communication with the human vehicles and proceeds to optimise the travel time dynamics for the CAVs. The optimisation involves adjusting trajectories to avoid collisions in vehicle movement.
- Ceases communication if the traffic light is red. Else, determines how much green-time is left before the red phase.
- Determines whether the vehicle can travel through the intersection before the red phase at its current speed. If not, calculates and communicates the new speed and acceleration information to the vehicle.
- Vehicles that cannot travel through the intersection before the red phase due to signalling and operational constraints will come to a complete stop at the approach.

Adebisi then proposed two delay models as a function of four variables; cycle length, effective green time, degree of saturation of the lane group, and CAV penetration rate. Cycle lengths for the intersection are fixed, but the effective green-time is expected to change with respect to CAV penetration. The saturation flow rate is obtained through the relationship $S = 3600/\bar{h}$, where \bar{h} is the saturation headway obtained by plotting the discharge headway at each signal head. Their study increased CPR from 0% to 90% in 10% increments, with 100% not being assessed as a completely autonomous fleet would render signals redundant. The two delay models, Linear Regression Delay model and Artificial Neural Net model, are outlined below.

The linear regression model takes the form;

$$y_i = \beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2} + \dots + \beta_{(p-1)} x_{i,p-1} + \epsilon_i \quad (168)$$

Where y_i is the control delay obtained through microsimulation, and x_i is the selected variable. Adebisi used MATLAB to perform multiple regression analysis and obtained $p \leq 0.05$, which showed the significance of the variables in the model at 95% confidence interval. An R-squared value of 0.76 and an analysis of variance showed an overall $p \leq 2.02 \times 10^{-6}$, suggesting a good fit.

Artificial Neural Networks are a collective of mathematical systems efficient in processing and interpreting large datasets to produce an estimation or approximation. The Artificial Neural Networks model used by Adebisi is a three-category layer model consisting of the input layer, numerous hidden layers, and the output layer. The network's nodes, dubbed neurons, compute input values (x_j) and their assigned weightings (w_j) within a non-linear activation function. The function is given by Equation 169;

$$y = F\left(\sum_{j=0}^N x_j w_j\right), x_0 \equiv 1 \quad (169)$$

All neurons within a layer are linked to all neurons within the consecutive layer via weight lines. The weight lines denote the degree of weighting needed to minimise errors. In the estimation of delay, the input layer requires variables from the four categories of cycle length, effective green-time, DoS, and CAV penetration rate. Variables are processed into the first hidden layer where the nodes and the collection of weight factors sort inputs via a training method which employs the backpropagation algorithm. The sort intends to produce the optimal combination of all layers exhibiting minimal overall errors.

Adebisi compared the proposed linear regression delay model and Artificial Neural Networks model with the Webster and HCM model. The study performed another simulation to obtain a set of data that is different from the data used to calibrate the proposed models. Delay data was collected for each lane group for the two intersections that were studied. The simulated and calculated delay output for a specific lane group is plotted in Figure 55;




Figure 55: Comparison of the two approaches proposed by Adebisi, against the Webster model and the HCM model [Adebisi, 2018].

Results from the comparison show that delay estimated by the proposed models are more accurate than the conventional delay models, with the estimate from the Artificial Neural Networks model closely aligning with the observed data (estimated error of approximately 3.2%). The linear regression model has an error of 8.1%, whereas the HCM and Webster’s model had a higher error of 32.2% and 33.5%, respectively. Adebisi’s work showed significant improvement in delay estimation with the proposed models, which further reiterates the inadequacy of existing strategic models in measuring mixed fleet delay.

However, there are limitations to the study which must be addressed. The linear regression model assumes a linear relationship between the delay and the selected variables, which may be an oversimplification of the traffic flow characteristics. Despite the accuracy of the Artificial Neural Networks model estimates, a major limitation of this model is that it does not explain the behaviour of the network. The system iteratively searches for a set of values which best fit the given data but fails to provide a physical or mathematically basis which supports the estimated traffic performance. This model is also computationally and data intensive. The study also fails to examine the predictive quality of the proposed model at a low or high degree of saturation where network behaviour is relatively unstable.

This chapter aims to extend the foundational work established by many of the authors and work explored throughout this literature review, by evaluating the applicability of the BPR path cost model to CAV behaviour. This chapter then proposes adjustments to the strategic model, if warranted, by recalibrating relevant modelling parameters and applying correction terms. This process results in a simple model, providing practitioners with an interim solution to modelling CAVs, without resorting to computationally intensive, data-hungry, and mathematically complex solutions presented in the literature. Additionally, this chapter verifies the predictive qualities of the adjusted models by benchmarking them against the current BPR model in new environments.

10.2 Experimentation Development

Based on the literature review, many have proposed models and techniques for the macrosimulation modelling of CAVs and mixed fleets. However, none have attempted to verify the applicability of existing methods, nor has an investigation been conducted into the effectiveness of adding a correction term to existing VDFs that more accurately reflects CAVs and mixed fleets. For this reason, this chapter will investigate how well the BPR VDF performs as the driving fleet is incrementally transitioned to 100% CAV. If the BPR function is found to be inadequate, then recalibrating the model parameters will be attempted, with a verification of the recalibrated model performance. The remainder of this section contains a thorough description of the methodology and experimentation setup.

The investigation begins with the development of a testbed. The environment is a simple “Type A Weaving” motorway corridor with three onramps and offramps. The intricate details of the environment and proof of its stability is provided in Section 10.2.1. Once stability is established, capacity must be determined for each CPR. The literature review in Section 10.1 provides a range of methods through which to calculate capacity, however, none were considered appropriate. The methodology for calculating corridor capacity is provided in Section 10.2.2. Finally, this section concludes with an explanation of the experimentation structure.

10.2.1 Modelling Environment

Three testbeds are used in this study to generate data. Each testbed is a hypothetical motorway section consisting of three sets of onramps and offramps. Figure 56 displays a single section; the full modelling environment consists of three of these sections joined consecutively;



Figure 56: A schematic of the network geometry used for this study.

The three testbeds differ in weaving area length, speed limit, and weaving proportion. Refer to Table 57 for the differences between the three testbeds;

		Segment 1	Segment 2	Segment 3
	Speed Limit	90 km/hr	80 km/hr	100 km/hr
Weaving length	Testbed 1	200 m	250 m	300 m
	Testbed 2	250 m	250 m	250 m
	Testbed 3	300 m	250 m	200 m
Weaving Proportion	Testbed 1	15.00 %	22.77 %	33.01 %
	Testbed 2	19.39 %	22.34 %	30.77 %
	Testbed 3	23.47 %	22.68 %	19.10 %

Figure 57: The weaving length, speed limit, and weaving proportion used to differentiate the three testbeds.

Each onramp uses a form of ramp metering to ensure that weaving vehicles do not cause a choke point to arise and artificially restrict flow. The ramp metering operates by holding vehicles on the ramp until the adjacent mainline average travel speed returns to a minimum of $50km/hr$. CAV behaviour is controlled by the framework developed in Chapter 5.1, and human behaviour is controlled by the microsimulator using the models in Chapter 5.2.

The naming convention for the zones in the base OD matrix is as follows. The mainline is the main motorway movement, with vehicles travelling from East (“Mainline On”) to West (“Mainline Off”). The ramps are numbered sequentially in order of appearance, starting with 1 in the East and 3 in the

West. “On” and “Off” ramps have the appropriate suffix appended to their name. So a vehicle travelling from the mainline to the third offramp would be labelled as “Mainline On to Ramp 3 Off”. Using this naming convention, Table 30 provides the base OD matrix for each of the testbeds. To moderate demand in the network, the base case matrices are multiplied by an integer scaling factor, similar to the study conducted in Chapter 9;

Testbed 1					
Origin\ Destination	Ramp 1 Off	Ramp 2 Off	Ramp 3 Off	Mainline Off	Total
Mainline On	6	12	13	57	88
Ramp 1 On	3	2	3	4	12
Ramp 2 On		1	3	6	10
Ramp 3 On			2	15	17
Total	9	15	21	82	127
Testbed 2					
Origin\ Destination	Ramp 1 Off	Ramp 2 Off	Ramp 3 Off	Mainline Off	Total
Mainline On	10	11	13	52	86
Ramp 1 On	3	2	3	4	12
Ramp 2 On		1	3	5	9
Ramp 3 On			2	9	11
Total	13	14	21	70	118
Testbed 3					
Origin\ Destination	Ramp 1 Off	Ramp 2 Off	Ramp 3 Off	Mainline Off	Total
Mainline On	10	11	6	55	82
Ramp 1 On	3	2	3	8	16
Ramp 2 On		3	3	6	12
Ramp 3 On			3	5	8
Total	13	16	15	74	118

Table 30: The base OD matrices used for each testbed in this study.

The number of weaving vehicles in the first weaving section is the summation of vehicles travelling from “Mainline On” to “Ramp 1 off” and vehicles travelling from “Ramp 1 On” to every destination bar “Ramp 1 Off”. This method of calculating weaving proportion amounts to a weaving proportion of 15%, 23%, and 33% in the first, second and final section, respectively, in Testbed 1.

Model stability is established by conducting 50 iterations of each testbed using 0% CAVs and a network demand at capacity. The standard deviations of the key network metrics (average delay, average travel time, average travel distance, and average speed) were calculated as a proportion of the median value. The network results of the stability analysis are provided in Table 31. The average delay has a higher deviation compared to the other metrics, an expected outcome for environments operating near capacity. The other metrics show little variance compared to their median values;

	Average Delay	Average Speed	Average Travel Distance	Average Travel Time
Testbed 1	20.12 %	5.18 %	0.25 %	5.43 %
Testbed 2	14.87 %	2.80 %	0.30 %	2.95 %
Testbed 3	13.93%	7.37%	0.88%	6.91%

Table 31: Network statistics used to establish model stability. Each value is the proportion of the standard deviation of all runs, to the median value. 50 iterations for each testbed was used.

Each iteration uses a warmup period of *30min*, where the network is loaded with 80% of the peak hour demand. The warmup period is then followed by a peak modelling period of *1hr*, which is used for network evaluation.

10.2.2 Corridor Capacity

Capacity plays a critical role in the BPR function. If the capacity of an environment is not known accurately, then using an approximation in the BPR function would compound error upon error, potentially leading to a disingenuously inaccurate result. For this reason, extensive care has been taken in determining the corridor capacity. While many of the approaches presented in Section 10.1 were explored to determine capacity, they were not deemed

appropriate for this study. Those approaches either required an unreasonably intimate knowledge of the corridor and specific driving behaviour, or they provided approximations, the accuracy of which would be difficult to verify.

This study used a modified version of the Selected Maxima method. Since the localised factors that affect capacity in this corridor were unknown, the network was incrementally loaded, and throughput was observed. Data were collected on each of the three offramps and the two mainline lanes leading to the end of the model. The environment was then incrementally loaded with traffic by multiplying the base OD matrix by a factor starting at 1 and incrementally increasing by 1. Total network throughput is calculated by summing the volume measured at each detector. An example plot of total throughput vs demand multiplier is provided in Figure 58;




Figure 58: The corridor throughput plotted against corridor demand, used to calculate the capacity of the modelling environment.

Incrementally loading the corridor reduces the complexity in identifying its capacity. Referring to Figure 58, capacity is defined as the point where an appreciable increase in corridor demand does not lead to an appreciable increase in throughput, when the gradient of the plot approaches 0. This methodology was sufficient in determining capacity in Chapter 9, however, it still contains a degree of ambiguity that is not acceptable considering the nature of this study. The additional information provided in Table 32 provides the percentage change in demand, the percentage change in throughput, a ratio of the two metrics, and a 3-value running average.

Demand Multiplier	Total Through	Δ Demand	Δ Through	$\frac{\Delta Demand}{\Delta Through}$	Running Average
37	4621	3%	2%	67%	92%
38	4776	3%	3%	124%	105%
39	4870	3%	2%	75%	89%
40	4974	3%	2%	83%	94%
41	5114	2%	3%	113%	90%
42	5153	2%	1%	31%	76%
43	5225	2%	1%	59%	68%
44	5170	2%	-1%	-45%	15%
45	5174	2%	0%	3%	6%
46	5109	2%	-1%	-57%	-33%
47	5135	2%	1%	23%	-10%
48	5335	2%	4%	183%	50%

Table 32: The quantitative values used to generate Figure 58, also showing the rate of change for demand and throughput. This table is used to inform capacity.

In Table 32, the pale yellow colour depicts the demand multipliers that result in an undersaturated corridor, the darker orange colour depicts the oversaturated environment, and the unique orange colour results in the corridor capacity. The criteria used for this evaluation is as follows. When $\Delta Demand / \Delta Through$ is near 100%, it indicates that the incremental demand loaded to the network was able also to reach the detectors at the exit points, and hence the network is still performing in an unhindered state below capacity. As the network is further loaded, this ratio drops, indicating proportionally less throughput than in the previous case.

Setting a threshold for the $\Delta Demand / \Delta Through$ ratio is not adequate for defining capacity. Due to the inherent stochasticity of microsimulation modelling, a particular run may show a decrease in throughput while the next may show an increase that more than sufficiently compensates for the previous decrease. For this specific reason, an average of the previous 3 results is used, dubbed as the “running average” in Table 32. While placing a threshold on the running average provided substantially more representative calculations for capacity, it also was limited in its usefulness when consecutive runs tend to oscillate around a value. In rare cases where the oscillation occurred, manual selection of the capacity multiplier was made.

10.2.3 Experimentation Structure

The experimentation structure is as follows;

- Establish Model Stability: The stability of the modelling environments must first be proven, to provide validity to any data synthesised from it. Model stability has already been demonstrated in Section 10.2.1.
- Determine the Capacity: The capacity for each modelling environment in 20% increments of CAVs is calculated using the methodology outlined in Section 10.2.2.
- Verify BPR Performance: Synthetic data is generated for the modelling environment in 20% CPR and 10% DoS increments. By comparing the synthetic data with the predictions made by the BPR function, the appropriateness of the BPR's use with CAVs and mixed fleets is evaluated.
- BPR Model Correction: If the BPR function is found to be deficient in appropriately calculating link delay for CAVs and mixed fleets, then a range of correction techniques will be attempted using the synthetic data as a means of calibration.
- Corrected Model Validation: The corrected model is then validated by comparing its predictive qualities against the BPR function in a new set of modelling environments.


10.3 Experimentation Results

This section provides the experimentation results. The results are segregated into the subsections mentioned in Section 10.2.3. Demand values in this section are a multiplication factor for the base case matrices provided in Table 30.


10.3.1 Testbed Capacity

Capacity for each testbed and each CPR was found by incrementally loading the network with demand and using the process identified in Section 10.2.2. Capacity found using a single random seed is not reliable, as random seeds and natural microsimulation stochasticity render any one modelling result insignificant unless it is benchmarked against other results. For this reason, 300 iterations of each testbed for each CPR were conducted, where the demand was randomly generated for $DoS \in [0, 1.1]$ and random seed (RS) was


randomly generated for $RS \in [1, 99999]$. Figure 59 provides the throughput recorded for each iteration;



(a) Demand vs Throughput for the first testbed, discretised by CAV penetration.



(b) Demand vs Throughput for the second testbed, discretised by CAV penetration.



(c) Demand vs Throughput for the third testbed, discretised by CAV penetration.

Figure 59: Plotting throughput vs demand, to determine the environment capacity for different CAV penetrations. The capacity is defined as the inflection point where a demand increase does not lead to an appreciable increase in throughput.

Figure 59 demonstrates that while each testbed is undersaturated, an increase in demand leads to an identical rise in throughput. For each CAV penetration, the environment reaches a point where the Demand vs Throughput plot becomes horizontal and additional demand does not increase throughput further. This point is defined as capacity. Table 33 provides the quantitative value for capacity, for each CAV penetration and each testbed.

	CAV Penetration Rate (%)										
	0	10	20	30	40	50	60	70	80	90	100
Testbed 1	44	46	49	53	55	58	63	68	73	80	82
Testbed 2	49	51	52	55	59	65	68	72	75	80	84
Testbed 3	47	49	52	55	59	64	69	74	77	83	85

Table 33: The capacity for each testbed and each CAV penetration.

The values in Table 33 are used to define the different DoS scenarios. The increase in capacity relative to the 0% CAV case for each testbed is provided in Table 34;

	CAV Penetration Rate (%)										
	0	10	20	30	40	50	60	70	80	90	100
Testbed 1		5%	11%	20%	25%	32%	43%	55%	66%	82%	86%
Testbed 2		4%	6%	12%	20%	33%	39%	47%	53%	63%	71%
Testbed 3		4%	11%	17%	26%	36%	47%	57%	64%	77%	81%

Table 34: The change in capacity for each testbed and each CAV penetration rate, relative to the base case 0% CAV scenario.

10.3.2 Verify BPR Model Performance

Verification of the BPR model requires synthetic data. Synthetic data was generated for all three testbeds, $CAVs \in [0\%, 100\%]$ in 10% increments,

$DoS \in [10\%, 100\%]$ in 10% increments, and 20 randomly generated iterations for each case. This structure amounts to a total of 6,600 iterations of modelling. The synthetic data is then aggregated into two groups, 30% reserved for the VDF model evaluation and recalibration, and the remaining 70% reserved for VDF model validation if it is warranted.

Table 35 below provides the average microsimulation travel time for each CPR and DoS scenario. This table shows that travel time increases as congestion in the network build, and travel time decreases with increasing CAV penetration. The first result is well understood, but it does demonstrate that the system reacts rationally to growing levels of congestion. The second finding has been already demonstrated in this thesis, specifically in Chapter 7 and Chapter 9. The travel time results for the environments in this study are as follows;

Testbed 1 - Average Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	78	79	79	80	80	82	86	95	124	214
10%	78	79	79	80	81	83	87	98	134	212
20%	78	79	79	80	81	83	90	99	137	201
30%	79	79	80	80	82	85	93	110	156	211
40%	78	79	80	80	82	84	93	101	141	174
50%	79	79	80	81	83	86	93	109	142	177
60%	79	80	81	81	84	85	98	101	144	149
70%	79	80	81	82	84	89	101	118	139	171
80%	79	80	81	82	86	88	99	106	141	141
90%	80	81	82	84	87	95	107	123	139	143
100%	80	81	82	84	89	93	111	114	128	131
Testbed 2 - Average Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	88	88	89	90	91	92	95	101	144	162
10%	88	88	89	90	91	93	96	107	144	169
20%	88	88	89	90	91	93	96	102	134	157
30%	88	89	90	91	92	94	97	109	137	166
40%	88	89	90	91	92	93	101	105	145	158
50%	89	89	90	91	93	96	103	119	148	166
60%	89	90	91	91	93	96	105	111	142	156
70%	89	90	91	92	94	97	104	115	139	156
80%	89	90	92	92	95	97	105	110	127	139
90%	90	91	92	93	95	98	105	112	126	137
100%	90	91	92	93	95	96	103	106	123	127

Testbed 3 - Average Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	70	71	71	72	74	75	80	88	116	162
10%	70	71	72	73	74	77	81	94	135	168
20%	71	71	72	73	74	77	83	93	130	150
30%	71	71	72	73	75	78	86	100	134	158
40%	71	72	72	73	76	78	89	95	139	144
50%	71	72	73	74	77	81	92	108	135	158
60%	71	72	73	74	78	79	94	99	134	134
70%	72	73	74	76	79	85	95	110	128	137
80%	72	73	74	76	79	82	93	99	120	125
90%	72	74	75	77	81	86	94	105	116	122
100%	73	74	76	78	82	85	93	98	112	116

Table 35: Average travel time (sec) for each CPR, DoS and testbed, averaged for a randomly selected 30% of the synthetic data.

Table 36 demonstrates the difference in travel time calculated through the synthetic data, and travel times predicted by the BPR function. For the BPR function to be used appropriately, its parameters must be calibrated for each use case. The BPR α and β values that minimised error between the synthetic data and the prediction made by the BPR function were used ($\alpha = 1.0122$ and $\beta = 4.1856$). Deviations in travel time predictions are as follows;

Testbed 1 – Difference between Synthetic and BPR Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	-1	-1	-1	0	1	5	9	14	4	-58
10%	-1	-1	-1	-1	1	4	8	11	-6	-56
20%	-1	-1	-1	-1	0	3	5	9	-9	-45
30%	-1	-1	-2	-1	0	2	2	-1	-28	-55
40%	-1	-2	-2	-1	-1	2	2	7	-13	-18
50%	-1	-2	-2	-2	-1	1	2	-1	-14	-21
60%	-1	-2	-3	-2	-2	1	-3	7	-16	6
70%	-2	-3	-3	-3	-3	-2	-5	-10	-11	-15
80%	-2	-3	-3	-3	-4	-1	-4	2	-13	15
90%	-2	-3	-4	-5	-6	-8	-12	-15	-12	12
100%	-2	-4	-4	-4	-7	-6	-16	-6	0	25

Testbed 2 – Difference between Synthetic and BPR Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	-1	-1	-2	-1	1	5	11	20	0	13
10%	-1	-1	-2	-1	1	4	11	15	-1	6
20%	-1	-1	-2	-1	1	5	11	19	9	18
30%	-1	-2	-2	-2	0	3	9	13	7	9
40%	-1	-2	-2	-2	0	4	6	16	-1	17
50%	-2	-2	-3	-2	-1	1	4	3	-4	9
60%	-2	-3	-3	-3	-2	2	1	10	1	18
70%	-2	-3	-4	-3	-2	0	3	6	4	19
80%	-2	-3	-4	-4	-3	1	2	12	16	36
90%	-3	-4	-5	-4	-3	-1	2	9	18	38
100%	-3	-4	-5	-4	-3	1	4	15	21	48

Testbed 3 – Difference between Synthetic and BPR Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	-1	-1	-1	-1	0	4	6	9	0	-21
10%	-1	-1	-1	-1	0	1	5	4	-19	-27
20%	-1	-1	-1	-1	0	2	3	5	-15	-10
30%	-1	-1	-2	-1	-1	0	0	-2	-18	-17
40%	-1	-2	-2	-2	-2	0	-3	3	-24	-3
50%	-1	-2	-2	-3	-3	-3	-6	-10	-19	-17
60%	-1	-2	-3	-3	-4	-1	-8	-1	-18	7
70%	-2	-3	-3	-4	-5	-6	-9	-13	-12	3
80%	-2	-3	-4	-4	-5	-4	-7	-1	-5	16
90%	-2	-3	-5	-6	-7	-8	-8	-7	-1	19
100%	-3	-4	-5	-6	-8	-7	-7	0	3	25

Table 36: Difference in the average travel speed predicted by the BPR function, and the average travel speed found in using the synthetic data.

Table 36 shows that despite the BPR parameters being calibrated to minimise error, the BPR tends to incorrectly predict travel time when either DoS or CPR increases. The RMSE for the BPR function is 11.56, 9.68 and 7.95, for the three testbeds respectively. Decreasing the weaving proportion from Testbed 1 to Testbed 3 reduced the error of the BPR corridor travel time predictions. To contextualise this error, consider the scenario that generated the highest error value of $48s/veh$. This scenario had a demand of $9,912veh$, implying a total system delay of $476,000s$ in the $1hr$ modelling period for this 3-link corridor. Extrapolating to a citywide network would result in substantially poor results. For this reason, recalibrating the BPR function parameters to better predict delay in mixed fleet environments will be attempted.

10.3.3 BPR Model Recalibration

The BPR function has no capabilities to cater to CPR. This section starts by proposing a unique α and β for each CPR and for each testbed, optimised to minimise the delay prediction error of the BPR function for 30% of the synthetic data. Figure 60 plots the resulting optimised parameters;




Figure 60: Optimised α and β values for each CPR and testbed.

A linear model is fitted to the average parameter value for each CPR. The linear best fit model for the α and β parameter is provided in Figure 61.




Figure 61: The linear models fitted to the α and β parameters.

The r^2 value describing the goodness of fit of the linear model is 0.9013 and 0.9821 for the α and β parameters respectively. The calibrated linear models for both parameters is provided in Equation 170;

$$\begin{aligned}\alpha &= -0.7302 \times CPR + 1.4193 \\ \beta &= -4.8811 \times CPR + 6.7691\end{aligned}\quad (170)$$

Where, $CPR \in [0, 1]$

10.3.4 Corrected Model validation

The recalibrated BPR parameters were used to develop an α and β model with respect to CPR. The predictive qualities of the recalibrated model are evaluated using the remaining 70% of the synthetic data. Provided in Table 37 is the difference in synthetically generated travel time and that predicted by the recalibrated model, similar to Table 36 developed for the BPR model;

Testbed 1 - Average Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	-1	-1	-1	-2	-2	-1	1	7	7	-26
10%	-1	-1	-2	-2	-2	-1	1	5	-2	-30
20%	-1	-1	-2	-2	-2	-1	0	5	-6	-24
30%	-1	-1	-2	-2	-2	-1	-2	-4	-26	-40
40%	-1	-2	-2	-2	-2	1	0	6	-11	-9
50%	-1	-2	-2	-2	-1	1	2	-1	-13	-18
60%	-1	-2	-2	-1	-1	3	-1	8	-16	4
70%	-2	-2	-2	-1	0	1	-2	-8	-12	-23
80%	-2	-2	-2	0	0	4	1	5	-15	1
90%	-2	-2	-1	0	1	0	-4	-11	-16	-7
100%	-2	-1	1	3	3	5	-6	-2	-7	0
Testbed 2 - Average Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	-1	-1	-2	-3	-3	-1	2	13	3	49
10%	-1	-2	-2	-3	-3	-1	3	9	3	35
20%	-1	-1	-2	-3	-2	0	5	15	13	40
30%	-1	-2	-3	-3	-2	0	5	10	10	26
40%	-1	-2	-3	-2	-2	2	4	15	1	27
50%	-2	-2	-3	-3	-2	1	3	3	-3	13
60%	-2	-3	-3	-2	-1	3	3	12	2	16
70%	-2	-3	-3	-2	0	4	7	9	3	10
80%	-2	-3	-2	0	2	7	8	15	13	21
90%	-3	-3	-1	2	5	8	11	13	13	16
100%	-2	-2	1	4	8	13	15	20	13	20

Testbed 3 - Average Travel Time (sec)										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	-1	-1	-1	-2	-3	-2	-1	3	3	8
10%	-1	-1	-2	-2	-3	-3	-1	0	-16	-4
20%	-1	-1	-2	-2	-2	-2	-1	2	-12	9
30%	-1	-1	-2	-2	-3	-2	-3	-4	-16	-4
40%	-1	-2	-2	-2	-3	-1	-5	2	-22	5
50%	-1	-2	-3	-3	-3	-3	-6	-10	-18	-14
60%	-1	-2	-3	-2	-3	0	-6	0	-18	5
70%	-2	-2	-3	-3	-2	-3	-6	-10	-13	-4
80%	-2	-3	-3	-2	-1	1	-2	2	-7	3
90%	-2	-2	-2	-1	0	-1	-1	-4	-5	2
100%	-2	-2	-1	1	1	3	1	4	-3	2

Table 37: Difference in the average travel speed predicted by the recalibrated model, and the average travel time found using the synthetic data.

To directly compare the performance of the original BRP model to the recalibrated variant, RMSE is used. Table 38 contains the ratio of RMSE for the recalibrated BPR function, to RMSE for the original BPR function. A value between 0 and 1 in green indicates that the recalibrated model outperformed the original BPR function. A value greater than 1 in red indicates the opposite, and the BPR outperformed the recalibrated function. The predictive qualities of the recalibrated model are as follows;

Testbed 1 – Original BPR RMSE / Recalibrated BPR RMSE										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	1.01	1.11	1.49	3.83	1.70	0.43	0.17	0.53	2.72	0.29
10%	1.01	1.09	1.42	2.84	2.16	0.32	0.24	0.61	2.41	0.50
20%	1.01	1.08	1.30	2.42	5.32	0.25	0.03	0.61	0.80	0.47
30%	1.01	1.05	1.22	1.71	7.87	0.68	0.48	2.90	0.91	0.69
40%	1.00	1.04	1.13	1.49	3.52	0.31	0.06	0.89	0.86	0.37
50%	1.00	1.01	1.03	1.08	1.32	0.49	0.84	0.52	0.89	0.89
60%	1.00	0.97	0.91	0.71	0.50	1.88	0.12	1.20	0.98	0.21
70%	0.98	0.90	0.75	0.49	0.17	0.00	0.05	0.72	1.07	1.83
80%	0.96	0.79	0.54	0.03	0.21	3.87	0.01	2.79	1.26	0.06
90%	0.89	0.65	0.28	0.09	0.25	0.02	0.34	0.72	1.36	0.87
100%	0.72	0.32	0.13	0.72	0.46	0.69	0.33	0.30	9.05	0.01

Testbed 2 - Original BPR RMSE / Recalibrated BPR RMSE										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	1.01	1.09	1.37	2.99	3.06	0.26	0.19	0.61	1.39	4.49
10%	1.01	1.06	1.35	2.26	4.42	0.32	0.26	0.63	2.61	31.33
20%	1.01	1.06	1.25	2.15	4.45	0.01	0.44	0.79	1.32	2.45
30%	1.00	1.05	1.19	1.68	116.15	0.04	0.51	0.79	1.34	4.42
40%	1.00	1.03	1.11	1.39	5.53	0.45	0.68	0.92	1.85	1.64
50%	1.00	1.01	1.02	1.06	1.21	0.53	0.93	1.18	0.73	1.56
60%	1.00	0.97	0.92	0.77	0.35	1.86	1.57	1.17	6.16	0.83
70%	0.99	0.92	0.76	0.49	0.57	39.99	2.87	1.79	0.65	0.56
80%	0.96	0.82	0.56	0.06	1.04	9.44	3.03	1.34	0.79	0.54
90%	0.90	0.66	0.30	0.35	1.53	7.77	3.93	1.52	0.66	0.42
100%	0.75	0.35	0.13	1.02	2.68	11.09	3.86	1.32	0.65	0.43
Testbed 3 - Original BPR RMSE / Recalibrated BPR RMSE										
CPR\DoS	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
0%	1.01	1.11	1.45	3.29	4.08	0.56	0.39	0.30	0.24	0.51
10%	1.01	1.08	1.38	2.41	7.51	0.82	0.26	0.14	0.71	0.08
20%	1.01	1.07	1.27	1.96	5.71	1.06	1.55	0.47	0.83	0.05
30%	1.01	1.05	1.19	1.52	3.21	6.09	2.57	1.92	0.88	0.45
40%	1.00	1.03	1.11	1.32	1.62	0.97	1.56	0.76	0.91	0.36
50%	1.00	1.01	1.02	1.05	1.08	1.08	1.03	0.98	0.95	0.79
60%	1.00	0.97	0.93	0.81	0.76	0.26	0.80	2.75	0.99	0.62
70%	0.99	0.92	0.81	0.65	0.49	0.50	0.61	0.84	1.07	11.83
80%	0.96	0.84	0.65	0.32	0.27	0.23	0.34	0.39	1.67	0.17
90%	0.91	0.70	0.44	0.20	0.15	0.10	0.21	0.49	2.42	0.04
100%	0.75	0.46	0.16	0.18	0.16	0.47	0.19	75.27	1.71	0.02

Table 38: A direct comparison of the BPR function vs its recalibrated variant. A green value indicates the recalibrated model outperformed the original BPR, with a red value indicating the opposite.

Through visual inspection of Table 38, it is clear that there are specific points where the recalibrated model performs significantly worse than the original BPR model. However, the overarching distribution of green does seem to indicate that the model as a whole operates better. Visual inspection alone is insufficient to conclude, a qualitative assessment follows.

When considering the aggregated synthetic data and applying the two BPR variants to each data point, the original model yields an RMSE of 15.16, whereas the recalibrated BPR yields an RMSE of 8.86. This change amounts to a model improvement of 42%. The improvement is comprised of 67% of

the data points showing an improvement in travel time prediction, with the remaining 33% showing a worse prediction. However, simply calculating the number of instances of improvement is also insufficient, as the 67% could have shown negligible improvement, while the remaining 33% could have shown significantly worse performance. By assessing the absolute improvement in travel time prediction, the significance of the improvement can be investigated. The average improvement of all data points is 543%. This figure indicates that the improvements in travel time prediction are 5.43 times better than any errors arising from the new BPR model. The predictive qualities of the new approach, while not flawless, far outweigh the predictive qualities of the original BPR model in a mixed fleet environment.

10.3.5 Commentary on the Conical and Davidson VDFs

A similar methodology was used to attempt a recalibration of both the Conical and Davidson VDFs, with neither showing positive results. This methodology was successfully applied to the BPR function is due to the simple nature of the model. The α and β parameters occur in the BPR function once, making their role both tractable and relatable to real-world phenomena. Therefore, altering the model has a predictable consequence that also reflects the change observed in real-world fleet operation or in synthetic data. The Conical and Davidson function, however, have a complex relationship between the travel time prediction and their calibration parameters (refer to Equation 153 and Equation 154). This complex relationship results in modifications to the model, as conducted for the BPR function, having unintended consequences. The methodology used in this study for recalibrating the BPR function is not appropriate for either the Conical function or the Davidson function.

10.4 Discussion

The results reported in Section 10.3 demonstrate that the traditional BPR function performs well for undersaturated conditions, but performs poorly as congestion or CPR increases. When the α and β parameters are recalibrated as functions of CPR, the predictive qualities of the BPR function improve. Before and after recalibration, the BRP function RMSE improved from 15.16 to 8.86. In 67% of cases, the recalibrated variant better-predicted travel time than the traditional BPR function. Finally, of the 4,620 simulation runs forming

the 70% of the dataset used for validation, the recalibrated BPR function better-predicted travel time by 5.43 times.

The implications of the results found in this chapter are two-fold. Firstly, this study demonstrates that the BPR function can benefit from recalibration and make its form more appropriate for use with mixed fleets and CAVs. Secondly, this study demonstrates that for legacy purposes, the BPR function does not need to be explicitly replaced with more complicated and sophisticated models developed specifically for CAVs. The literature review in Section 10.1 demonstrated that alternative approaches for mixed fleet and CAV VDFs relying on artificial intelligence and machine learning requires substantial quantities of data for training and calibration purposes. Other approaches proposed in the literature are not as simple and accessible as the BPR function. For this reason, validating its continued use is critical.

There are, however, questions raised in regards to the methodology used in this study that warrant further investigation. The original BPR function was developed by curve-fitting a model to data observed from a single motorway section in the United States. The small dataset used for calibration significantly narrows the use-cases in which the parameters recommended in the HCM are appropriate. For this reason, the HCM recommends recalibrating the BPR model parameter for each environment in which it is used. While this study used three different testbeds to synthesise the data, it still followed a similar methodology to that used in the development of the original BPR function. The detailed investigation conducted in Chapter 9, where motorway performance as a function of CAV penetration was assessed, demonstrated that motorway section performance is significantly susceptible to the proportion of weaving vehicles. This implies that the synthetic data used for calibration and validation in this study may also be subject to the same over-fitting as the original BPR function.

Additionally, this study demonstrated that the weaving proportion affected the optimal α and β parameters. Figure 60 showed that the variance in the optimum parameters is higher for low CAV penetrations (less than 20%). Therefore it is less reliable to use a curve-fitted model to calculate the parameter for lower CAV penetration rates. Table 37 reconfirms this expectation,

as travel time predictions in the low CAV penetration and high DoS region tend to show a greater error against the observed travel time. Altering the parameters to be functions of CPR may not be enough. If the parameters also show sensitivity to vehicle weaving proportions, speed limits, weaving segment length, and other factors, then also including these factors in the parameter function may yield improved results. However, each additional factor included in the assessment exponentially increases the data required and the modelling time, especially if a covariate assessment approach is used.

The CAV emulation framework developed as part of this thesis and presented in Chapter 5.1 is underpinned by the critical assumption of cooperativeness. The investigation in Chapter 9 demonstrated that cooperativeness at merge junctions significantly worsens performance when the proportion of vehicle weaving is high. This study used weaving proportions of up to 33%, which is relatively high compared to what real networks experience. Had the weaving proportion been lower, the recalibration and prediction efforts may have yielded further improvements than they already did.

During the recalibration process of the BPR parameters, a linear best-fit model was used as it provided a high r^2 value for both parameters and is simple. This decision raises two key questions. The first, would an alternate model have yielded better results? The second, what are the real-world implications or justifications for using a specific model, and the value of the parameters? This study demonstrated that α and β are both inversely linearly proportional to CPR. In the BPR function, lowering the α parameter reduces the gradient of the travel time increase between a DoS of 0 and approximately 0.85, resulting in a sharper increase in travel time between DoS 0.85 and 1. Raising the β parameter influences the magnitude of the travel time increase at a DoS of 1. To say that CAVs have an inversely linear relationship means that increasing the CPR reduces the travel time impact of proportionally additional demand. This relationship was verified in Chapter 7, indicating the derivation of an inverse relationship is consistent with the other findings in this thesis.

10.5 Conclusion

In this study, the applicability of the BPR VDF to mixed fleets and CAVs was investigated. Synthetic data was generated using 6,600 modelling iterations. 30% of the synthetic data was used to assess the parameters of the BPR function. When it was found that the BPR function showed errors as CPR or DoS increases, the same 30% of the data was used to generate a linear relationship between CPR, α , and β . The recalibrated BPR function now contained α and β parameters that were a function of CPR. Using the remaining 70% of synthetic data, the predictive qualities of the recalibrated BPR function parameters were assessed. This assessment demonstrated that the RMSE improved from 15.16 to 8.86. In 67% of cases, the recalibrated variant better-predicted travel time than the traditional BPR function. Finally, of the 4,620 simulation runs, forming the 70% of the dataset used for validation, the recalibrated BPR function better-predicted travel time by 5.43 times.

Acknowledgement of Publication

This work is currently under review for publication - Qiu, E., Viridi, N., Grzybowska, H., & Waller, S. T. (2020). Recalibration of the BPR Function for the Strategic Modelling of Connected and Autonomous Vehicles. *Transportmetrica B: Transport Dynamics*, Under Review.

11 Safety Assessment of Mixed Fleets using the Surrogate Safety Assessment Module

Abstract: The transportation network can provide additional utility by addressing the safety concerns on roads. On-road fatalities are an unfortunate loss of life and lead to significant costs for society and the economy. CAVs, envisaged as operating with idealised safety and cooperation, could be a means of mitigating these costs. This chapter provides insights into the safety improvements to be attained by incrementally transitioning the fleet to CAVs. This investigation is done by constructing a calibrated microsimulation environment in Vissim and deploying the custom-developed CAV control algorithm. CAVs are introduced to the environment in 10% increments, and safety performance is assessed using SSAM. The results of this study show that CAVs at low penetrations result in an increase in conflicts at signalised intersections but a decrease at priority-controlled intersections. The initial 20% penetration of CAVs is accompanied by a +22%, -87%, -62% and +33% change in conflicts at the signalised, priority, roundabout and DDI intersection respectively. CAVs at high penetrations indicate a global reduction in conflicts. A 90% CAV penetration is accompanied by a -48%, -100%, -98% and -81% change in conflicts at the signalised, priority, roundabout and DDI intersection respectively.

Safety in transport is at the forefront of technological design, innovation, and regulation. The NSW state government of Australia is one of many regulatory bodies focused on delivering a safer and more efficient road network. Through its “Towards Zero” initiative, the NSW state government aims to implement strategic and emerging technology to reduce the number of deaths and severe injuries on the road network [Transport for New South Wales, 2018].

Accidents on roads are highly correlated with driving behaviours such as jerk [Pande et al., 2017], perception of risk and reward [Dixit, 2013], and the real-time state of the network such as density [Alsalmi et al., 2018]. The ABS and TfNSW indicate that accidents caused by negligent behaviours such as driving under the influence, fatigue and speeding contribute around 30% to the total number of accidents [Transport for New South Wales, 2017]. These accidents are a small subset of the negligent behaviour conducted by motorists, saying that CAV technology will mitigate 30% of accidents is an initial and conserva-

tive estimate. An independent study conducted by PricewaterhouseCoopers (PwC) indicates this reduction could be as high as 90% [PwC, 2004].

Significant economic benefits accompany the safety benefits of CAV technology. Road accidents result in a substantial financial cost to society, as shown by the Australian Bureau of Infrastructure, Transport and Regional Economics (BITRE) in their 2006 inquiry into the financial implications of vehicular accidents [BITRE, 2009]. The inquiry concluded that road accidents and deaths cost the community AUD\$2.8 Billion per year. That is AUD\$3,180,598 from each fatality, AUD\$346,869 from each serious injury, and AUD\$17,511 from each minor injury. The breakdown of costs indicates that post-accident vehicle repair and output losses (the loss to society due to sudden death) contribute to over half the costs incurred through an accident. CAV technology has the potential to reduce accidents and associated costs significantly.

When assessing the situation from either a financial and economic perspective, or a societal impact perspective, potential safety improvements through CAV technology provides a strong business case. This chapter investigates further the effects on road safety resulting from CAV uptake using a microsimulation modelling approach and SSAM. The remainder of this chapter is structured as follows; Section 11.1 contains a summation of the literature and work completed to date. Section 11.2 introduces the framework used for this experimentation. Section 11.3 discusses the development of the microsimulation network. Section 11.4 and Section 11.5 contain a summation of the results and the discussion, respectively. Finally, Section 11.6 provides a conclusion.

11.1 Literature Review of Vehicle Conflicts and Safety

Knowledge regarding the safety implications of CAVs derived from microsimulation modelling is limited. Physical CAV components are rigorously tested during development to minimise component and vehicle failure. Testing CAVs in network settings is difficult due to the requirement of expensive resources and the ethical uncertainties of human involvement in the trial of emerging technology. This ethical uncertainty applies to both the field and simulative testing of cooperative operations. This section outlines the safety studies conducted regarding current practices that inform the design of infrastructure, as well as investigations into the safety impacts of CAVs. Also, this section

in brief outlines crash-prediction models that have been calibrated using observed data. These models provide insights into the geometric factors that affect human driving safety, provoking thought into whether the same factors will influence a CAV.

11.1.1 Safety Studies of Human Driven Vehicles

The review of human-focused safety studies is segregated by intersection type. This section covers priority intersections such as the roundabout, signalised intersections such as the 4-way lights, and diverging diamond interchanges.

Roundabout Intersection Studies

The roundabout has been a staple of network design since the 1950s in many developed countries. Its relatively cheap construction costs, low maintenance, and priority-controlled operation make it a great fit in low-flow and low-velocity environments. Vehicle trajectories are forced through a narrow path in a circular motion, with conflicting vehicles required to yield. A study in Victoria, Australia [Austroads, 1993] measured the casualty rates at 73 intersections before and after the installation of roundabouts. Before installation, these intersections were either give way, stop, or police controlled. Post-installation, the study found a 74% reduction in casualty rates, a 32% reduction in property damage, and a 68% reduction in pedestrian casualty.

Kim and Choi found that the likelihood of accidents could be predicted at roundabouts using a model dependant on the number of approaches, number of entry lanes, entry width, flare width, number of circulating lanes, and circulating lane width [Kim & Choi, 2013]. Their study indicated that speed was not a significant contributor to the rate of accidents because roundabouts generally act as pinch points, wherein they force deceleration to 50% of free flow speeds as vehicles travel through them.

Qin et al., created a negative binomial crash prediction model using data for roundabouts obtained in Wisconsin [Qin et al., 2011]. Their approach developed a range of models with dependent variables that included Annual Average Daily Traffic (AADT), number of legs, number of lanes, geographic location in Wisconsin, the configuration of the yield signage, inscribed circle diameter,

and outer circle diameter. The model that considered the geographic location of the roundabout best fit the observed data. This study demonstrated that the familiarity of the driver to the concept and operation of a roundabout was more influential to its safety performance than the geometric considerations given to its inner and outer diameters or other design elements.

Turner and Roozenburg used data from 104 roundabouts in New Zealand to develop a range of accident prediction models empirically describing the likelihood of accidents in differing circumstances [Turner & Roozenburg, 2006]. The models estimate the likelihood of accidents for situations such as entering and circulating vehicles, rear-end collisions, loss-of-control accidents, pedestrians, cyclists, and high-velocity roundabouts. Contrary to other accident prediction models, their study found that speed and flow play a more significant role in the occurrences of accidents in motor vehicles than the geometric design does. Their model suggests that if the mean circulating velocity of 20km/h were increased by 20%, the resulting number of accidents would increase by 38%.

Signalised Intersection Studies

Contrary to priority-controlled intersections, the signalised intersection operates by dictating right of way and minimising concurrent conflicting movements. However, this style of operation introduces the motorist to the “dilemma zone” [Papaioannou, 2007]. The dilemma zone is the amber light period where a driver must decide whether to break in preparation of the red light or proceed through the intersection. Papaioannou found that the factors affecting this decision included pavement condition, intersection layout, cycle length, position in a platoon, and vehicle speed. The study concluded that in a 60km/h zone, 26.3% of drivers exceed the speed limit to cross the intersection in the dilemma zone. Males on average are 14 times more likely to exceed the speed limit and pass through the intersection on an amber light, with the 85% percentile exceeding the speed limit by 26%. The compliant behaviour of CAVs may eliminate dangerous behaviour such as this.

The diverging diamond intersection (DDI) was first implemented in the United States in 2009 and has gained popularity in recent times. The operation of a DDI involves the traffic stream momentarily crossing to the opposing side

before returning to the correct side at another crossing point downstream. This configuration reduces the number of conflicting movements present in conventional diamond intersections. DDIs yields higher throughput, contain fewer conflict points, and are designed to be safer. Claros et al., measured the change in the frequency of accidents in six locations in Missouri before and after the implementation of a DDI [Claros et al., 2015]. This study found substantial reductions in the occurrences of accidents for all severity types. Fatal crashes showed a 59.3% to 63.2% reduction, property damage crashes showed a 33.9% to 44.8% reduction, and total crashes showed a 40.8% to 47.9% reduction. A comparable study conducted by Hummer et al., used a similar methodology and found supporting results [Hummer et al., 2016]. Their study demonstrated that replacing conventional intersections with DDIs resulted in an overall crash reduction of 33%.

Variance in the explanatory variables used in these crash prediction models makes it difficult to predict the effect of CAVs on intersection safety. Models such as that presented by Qin et al., place heavy emphasis on the importance of geometric design in roundabout safety, other models such as that by Turner and Roozenburg focus on the flow and speed of vehicles. CAVs are expected to traverse the network in platoons with significantly smaller headways and higher average speeds. The Turner and Roozenburg models imply that the likelihood of accidents has the potential to increase in the interim scenarios, as portions of the network begin to transition to autonomous. Similarly, the notion of the dilemma zone, though pertinent to motorists occupying the road today, is antiquated for a vehicle whose behaviour is governed by a set of deterministic algorithms. For this reason, a framework developed in this thesis for evaluating CAV behaviour is necessary to gain insights into safety implications. This framework is implemented to evaluate CAV safety by determining the likelihood of conflicts qualitatively. The economic impact of CAVs through safety improvements can then be determined.

11.1.2 Safety Studies of Autonomously Driven Vehicles

Assessing CAV safety is necessary for evaluating their societal and economic impact. Studies of this nature may accelerate the development of political and legal policies required to facilitate their adoption. Outlined in this section are the sparse and recent investigations regarding the safety implications of CAVs.

Carbaugh et al., assessed the severity and frequency of rear-end collisions for automated and human highway systems [Carbaugh et al., 1998]. Their modelling found that in their testbed, a typical human driver has a probability of collision of 0.87, whereas an alert human driver has a probability of collision of 0.11. The CAVs, on the other hand, were found to have a probability of collision of 0.028. They also collided at less than 30% the velocity of typical human drivers. This study indicated that CAVs are four times safer than human drivers. The testbed was confined to the highway environment, assessing the safety implications of transverse vehicle interactions. Also, their study used probabilistic models to identify conflicts between vehicles. The novelty of this chapter is in its use of microsimulation modelling to emulate CAV behaviour for each agent. The evaluation of safety is conducted via analysis of vehicle trajectories in forecasting position and identifying potential conflicts. Additionally, this chapter assesses CAV safety implications for intersections.

Deluka Tibljas et al., quantified CAV safety using a microsimulation environment [Deluka Tibljaš et al., 2018]. CAV behaviour was emulated by changing the parameters of the Wiedemann 99 model [PTV Group, 2016], and safety performance was assessed using SSAM. Their study concluded that the number of conflicts would increase with the introduction of CAVs. While both this chapter and their study use microsimulation with Vissim and the SSAM module for safety evaluation, the novelty of this chapter lies in its use of a custom and external control protocol for dictating CAV behaviour. Deluka Tibljas et al., adjusted the parameters of the default car-following model used by the microsimulation software. Refined results can be obtained by using a framework that has been developed specifically for CAVs, giving special attention to CAV behaviour, communication, and cooperation.

Rahman and Abdel-Aty used the Vissim commercial microsimulator and SSAM in their analysis of the effects of CAV behaviour on safety [Rahman & Abdel-Aty, 2018]. Their study used a calibrated model of the Holland East-West Expressway (SR408) in Orlando, Florida as the testbed, and the IDM car-following model to emulate the behaviour of CAVs. Their study used five surrogate metrics to assess safety; the standard deviation of speed, time-exposed TTC, time-integrated TTC, time exposed rear-end crash risk index (RCRI), and sideswipe crash risk. Their study showed a reduction in time-exposed

TTC and time-integrated TTC of 19% to 21% when CAVs were allowed to form platoons in all lanes. When CAV operation was confined to a managed lane, the study found a reduction of 26% to 28% for both metrics, respectively. This chapter uses a similar methodology; however, it employs a different CAV car-following model and extends the findings by also investigating safety at a range of intersections in an arterial roadway setting.

Kidando et al., used the 2017 crash data obtained from the Signal Four Analytics website maintained by the University of Florida, to qualitatively assess the likelihood of accidents [Kidando et al., 2018]. The crash data indicated that for their segment, 85.1% of accidents were at intersections, with the remaining 14.9% occurring on the freeway. Rear-end collisions accounted for 55% of all collisions. The authors qualitatively assessed each accident type and determining whether CAV operation and technology would be adequate in eliminating it. Kidando et al., concluded that the potential reduction in conflicts resulting from CAV operation is between 17% and 70%. This chapter uses microsimulation modelling and a custom CAV control algorithm to assess safety as opposed to a qualitative analysis of historical crash data. Additionally, the large range identified by Kidando et al., provides scope for further refinement and more specified results.

Rahman et al., attempted to investigate the impact of CAV operation on safety during reduced visibility conditions [Rahman et al., 2018]. They used the Interstate I-4 in Florida as their base model, the IDM car-following model for CAV behaviour, and the Vissim commercial microsimulator for modelling vehicle interaction. Their study used the standard deviation of speed, the standard deviation of headway, and RCRI as surrogate measures of safety. The “look ahead” parameter of the commercial software was considered most critical for calibrating to fog conditions. Also, the ten parameters of the car-following model were iteratively adjusted to find the best match between observed and modelled car-following behaviour during fog conditions. The modelling found the most significant improvement at a CPR of 100%, with significant decreases seen after 30% penetration. The limitation of this study was in its constraint of the assessment to the highway environment. The modelling environment used in this chapter extends the understanding of CAV safety by conducting simulations in a mixed urban and freeway setting.

As data becomes available regarding CAV operation, these models can be better calibrated. The California Department of Motor Vehicles (DMV) requires that disengagement and accident data be made public by CAV developers, in exchange for permits to operate the vehicles [California, State of, 2018]. A study into the 2015 data released by manufacturers [Dixit et al., 2016] assessed several metrics, including the trust placed in CAVs by the occupants. This study evaluated the degree of correlation between autonomous disengagements and manual disengagements. An autonomous disengagement event is where the occupant is required to take control of the vehicle due to a shortcoming or error on the CAVs behalf. A manual disengagement event is where the occupant willingly seizes control of the CAV. The study found that there was a high correlation (0.73) between the frequency of autonomous disengagements and manual disengagements, indicating that trust in this emerging technology is currently fickle. Further real-world and microsimulation testing is necessary to develop a complete picture regarding CAV safety. Studies such as this chapter that evaluate the safety impacts of emerging technology and help build confidence in its potential.

11.2 Experimentation Framework

The framework for this study uses traffic and network data to emulate CAVs in a microsimulation environment. During a simulation, the microsimulator records trajectory data which the SSAM module uses to highlight potential conflicts that arose during runtime. Outlined in Figure 62 is the data flow structure of this study, with a more detailed explanation of each component provided in this section.

The framework consists of four key components. The data (yellow) is used to provide an element of realism to the testing, by calibrating and validating the study area to a real road network. This data acts as spatial and behavioural constraints for vehicle operation. The microsimulator (grey) is used to emulate vehicle interactions based on the restrictions imposed by the network and traffic data. The microsimulator is also responsible for generating, recording, and forwarding data to both the external control algorithm and analysis platform. The CAV emulation algorithm (orange) is responsible for dictating CAV movement. The microsimulator passes vehicle geospatial and behavioural information to the external control algorithm, where appropriate acceleration

and lane-change decisions are determined and returned to the microsimulator. The analysis platform (red) processes microsimulator data.

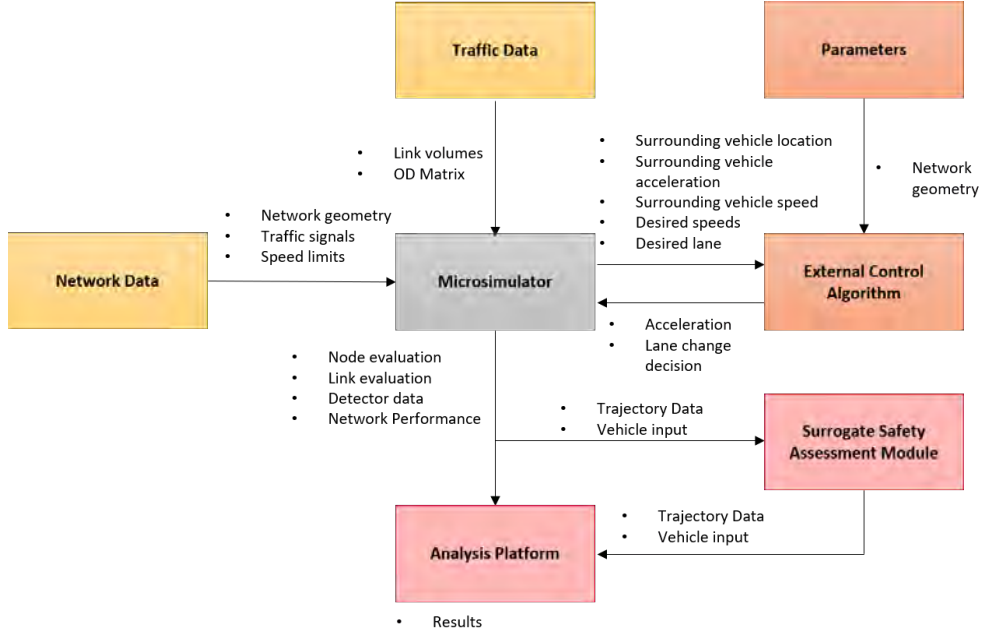


Figure 62: Framework for the flow of data in this study.

The Data

The traffic data consists of volume along links and a characteristic 24-hour traffic profile. Volume along links is available as AADT from public sources and converted to typical peak hour volumes using the characteristic profiles. This data is then used to calibrate an OD matrix manually. This process ensures a well-calibrated and fit for purpose base traffic model.

Google Maps, Bing Maps, and OpenStreetMap are used to obtain network data. Using their Satellite capabilities, these platforms form an underlay for creating the network geometry, ensuring correct road alignment and number of lanes. The Street-View capabilities of Google Maps aid in determining the signal phasing configuration. Cycle durations are set using the Roads and Maritime Services (RMS) signalling guidelines [RMS, 2016].

The Microsimulator

The microsimulator used for this study is Vissim 9-09 [PTV Group, 2016]. The microsimulator provides an environment in which the agents obtained from the traffic data can interact with each other. The interactions are confined in scope by restrictions imposed by the network data. The urban environment uses the Wiedemann 74 model, and the freeway environment uses the Wiedemann 99. The microsimulator also generates, stores, and forwarding real-time data including vehicle positions, current behaviour, and future intentions to the CAV algorithm. The algorithm uses this information to determine the appropriate course of action and returns the acceleration and lane-changing intentions of the vehicle for the next time increment.

The Surrogate Safety Assessment Module


SSAM is a tool created by Siemens Energy and Automation, Inc. with the FHWA [Gettman et al., 2008]. This tool uses trajectory data generated by microsimulators to identify potential conflicts, based on the definition of a conflict provided by the modeller. A trajectory file is created by the microsimulator during model runtime and contains information about the position and movement of each vehicle. Data in the trajectory file forms a subset of either the “Dimension”, “Timestep” or “Vehicle” class, explained as follows;

- The Dimension class contains information regarding the spatial characteristics of the observation area. “MinX”, “MinY”, “MaxX” and “MaxY” are used to define the rectangular bounding box of the microsimulation environment.
- The Timestep class contains a recording of the current time step since the commencement of the simulation. This variable allows SSAM to position the vehicles temporally.
- The Vehicle class contains information about the spatial characteristics of the vehicle. “VehicleID”, “Link ID”, “Lane ID”, “Front X”, “Front Y”, “Rear X”, and “Rear Y” are used to position the vehicle spatially. “Speed” and “Acceleration” forecast the movement of the vehicle.

Using the temporal and spatial information, SSAM determines whether the trajectory of the vehicle will interact with that of another, and reports infor-

mation regarding this interaction. This information includes time to conflict, speed during conflict, and speed after conflict. SSAM provides a range of criteria by which to define a collision. Maximum TTC is estimated based on the current location, speed, and trajectory of the two vehicles involved in the interaction. Maximum post-encroachment-time (PET) is the time between when the preceding vehicle and following vehicle last occupied the same space.

The rear end angle and the crossing angle are also used to identify potential conflicts. The rear end angle defines a potential collision during car-following and lane-changing. The crossing angle defines potential collisions in head-on scenarios, such as during manoeuvres through an intersection. Figure 63a provides a diagrammatic representation of these angles;



(a) Rear end angle and the crossing angle. (Source: [Gettman et al., 2008]) (b) A scenario where SSAM may inaccurately identify a conflict.

Figure 63: The definition of rear-end and crossing angle used by SSAM, and a scenario in the Vissim microsimulator where a conflict may be misidentified.

Due to the small headway kept between CAVs, SSAM tends to flag safe interactions for CAVs as a potential conflict. Consider the situation depicted in Figure 63b, where spatial restrictions in a congested environment force a lane-changing vehicle to remain in the centreline between both lanes. Once the front bumper of the vehicle has entered the adjacent lane, the following vehicles will no longer consider this vehicle as a leader and progress to seemingly drive through its rear. However, SSAM uses both the front and rear coordinates and flags these interactions as conflicts.

The Analysis Platform

The analysis platform is responsible for processing raw data obtained from the microsimulator and SSAM. The data from the microsimulator and SSAM include;

- Node evaluation results - Contain information regarding the performance of intersections by tracking the movement of vehicles through the node. This information is used to infer volume, density, speed, and delay.
- Link evaluation results - Provide similar information, but along links and for midblock locations.
- Detector data - Can be collected at any point in the network and is used to attain velocity, acceleration, and delay information.
- Network performance data - Provides a high-level aggregation of network statistics including total system travel time, total system delay, throughput, and volume.
- Trajectory data – Contains the geospatial and movement information for each vehicle in the network for each time increment. This information is used to infer potential conflicts.
- Vehicle input data – Creates a record of all vehicles entering the network and distinguishes between vehicle types (human-driven or CAVs). This information, with the trajectory data, classifies an interaction as either between humans, between CAVs, or containing both.

11.3 Case Study

The microsimulation environment is based on the Geelong area of Victoria, Australia. This location is chosen due to its hybridisation of both a highway environment and a residential urban environment. The area also has extensive publicly available data to use readily for calibration of the base model. Outlined in this section are the specificities of the case study and the structure of the experimentation.

11.3.1 Microsimulation Network

Calibration has been conducted to retain realism, with a DDI artificially incorporated into the environment. The reason for adding a DDI is because

this element of the transport network is increasing in popularity, resulting in it changing from being a rare occurrence in transport networks to becoming accessible by regular motorists. The DDI is safer than conventional intersections because it has fewer conflict points between interacting movements. So, if safety improvements can also be seen for an intersection arrangement that is already safer than conventional intersections, then this would further attest to the benefits of CAVs.

The elements of the network are as follows, with the contextualised modelling environment provided thereafter in Figure 64;

- **Signalised Intersection:** Four signalised intersections are present in the environment. One exists as a conventional four-way intersection, one is in the form of a DDI, and the other two are present at the motorway onramp and offramp.
- **Priority Junctions:** Eleven priority junctions are present in the form of four roundabouts and seven give-way junctions.
- **DDI Intersection:** A DDI has been artificially added to the environment in the top left of the study area. The DDI is a network element that is increasing in popularity and warrants investigation. Its geometric and signal configuration makes it a safer intersection arrangement than conventional intersections.
- **Highway Environment:** The study area contains the M1 Geelong Ring Road (Princess Freeway). This motorway contains two lanes in each direction with onramps accessing the motorway through a third tapered lane.




Figure 64: The calibrated microsimulation network of the Geelong area in Victoria, Australia.

11.3.2 Model Calibration

Calibration is the process of ensuring that modelled network behaviour aligns with observed network behaviour. The objective function for calibration is to minimise the Geoffrey E. Havers (GEH) statistic, which is a function of the observed and modelled volumes. The objective function is expressed as;

$$objective = \min\{GEH = f(M, C)\} \quad (171)$$

Where, M and C are the modelled and counted traffic volumes respectively. The GEH statistic is calculated using;

$$GEH = \frac{|M - C|}{0.5 \times (M + C)^{0.5}} \quad (172)$$

The GEH statistic measures the deviation between observed and modelled

traffic volumes. The absolute size of the measurement weights the deviations. This method ensures that smaller modelled volumes need to be closer to observed volumes, while larger volumes have a degree of tolerance.

Calibration followed the RMS modelling guidelines [RMS, 2013], where the OD matrix is iteratively adjusted in response to the deviation calculated between the observed and modelled flows. Similar methods are used in a range of other studies ([Oketch & Carrick, 2005], [Hollander & Liu, 2008], [Rahman et al., 2019], [Chu et al., 2003]). The GEH statistic is calculated for each turning movement or link flow independently. The portion of turning movements in the model containing a GEH statistic less than the threshold indicates the calibration quality. The guidelines recommend that 85% and 100% of volumes have a GEH statistic less than 5 and 10, respectively. Figure 39 displays the GEH Statistic for all major network links, calculated using a range of methods and all showing that the RMS criteria has been met.

Method	GEH
Sum of total volumes on links	1.05
Weighted average of the GEH on individual links against the observed volume	2.65
Weighted average of the GEH on individual links against the measured volume	2.69

Table 39: Network scale calibration results using the GEH statistic method.

AADT data used to calibrate the model is available publicly through the VicRoads Open Data platform [VicRoads, 2020]. This database provides extensive coverage for the majority of Victorian arterial and motorway roads. AADT was converted to peak-hour AM flow using the “Typical Hourly Traffic Volume” provided by the Victorian Government [Victoria Government, 2019]. The typical hourly volumes indicated an AADT to AM peak-hour volume conversion factor of 0.091.

Figure 65 provides the observed modelled volume for all major network links. This figure demonstrates that modelled volumes are sufficiently close to observed volumes, and the model is calibrated for the scope and purpose of this study. The calibrated demand was doubled during model runtime to ensure that latent demand is present in the model, not to skew results with an undersaturated network. Calibration results are as follows;



Figure 65: GEH statistic for major network links.

Signal cycle and phase times follow the RMS traffic signal design guidelines [RMS, 2016]. These guidelines incorporate a degree of optimality and safety in signal and phasing design. They reduce the number of conflict points arising in a cycle and ban risky arrangements such as dual turns filtering through opposing traffic movements. This ensures that baseline safety results are not artificially accentuated through the implementation of unoptimised phasing.

Travel time information along key links or queue length information at key intersections is a standard means of validating a microsimulation model. This study area is well under saturated and performs at near free flow speeds in the weekday peak period. The key routes in the microsimulation environment have an average travel speed that ranges between 94% and 107% of observed travel speeds, with observed travel speeds obtained from Google Travel Time data. A deviation of between -6% and +7% is considered acceptable for this microsimulation environment and is well within the $\pm 15\%$ threshold identified by the RMS modelling guidelines.

The default SSAM parameters are recommended by the United States FHWA and several other studies. The SSAM parameters are discussed in greater detail in Section 11.3.3. Calibration through traffic volumes and validation through travel times indicates that the OD matrix, network software parameters, and the vehicle behavioural parameters are adequately calibrated.

11.3.3 Surrogate Safety Assessment Module Parameters

The TTC and PET values were set to 1.5s and 5s respectively for the human vehicles, defaults recommended by the software to reflect human capabilities. The rear-end angle and crossing angle were set to 30° and 80° respectively, also defaults recommended by the software. These values have been calibrated and recommended by the United States FHWA ([Gettman et al., 2008], [Sabra et al., 2010]) and have been used in a number of other studies ([Stevanovic et al., 2013], [Wu et al., 2018], [Huang et al., 2013], [Ni et al., 2013], [Stamatiadis et al., 2013]). The standstill and following distance for CAVs were reduced to one-third. For this reason, the TTC and PET values defining CAV conflicts have also been reduced to one-third. The results are presented as a percentage difference between the base case and each scenario, meaning that biases inherent to the base case are also contained within each scenario tested.

11.3.4 Experimentation Structure

Between the different scenarios, CPR is increased in 10% increments from 0% being the base case with a fully human fleet to 100% being the fully autonomous case. Three random seeds are used in each of these scenarios. The environment location of the results reported in Section 11.4 are as follows;

- Signalised Intersection – Located on Hamilton Highway, at the North-East end of the study area.
- DDI North – Located on Shannon Avenue, the Northern end of the DDI.
- DDI South – Located on Shannon Avenue, the Southern end of the DDI.
- Roundabout – Located at the intersection of Barrabool Road and Shannon Avenue, at the South-East end of the study area.
- Highway – Large continuous roadway located in the West of the model environment.

11.4 Results

The results attained through microsimulation for the entire network are provided in Section 11.4.1 and Figure 66. The results are then disaggregated by intersection type in Section 11.4.2 and Figure 68. These results are the median of three random seeds. Using the median of multiple seeds accounts for the stochasticism of microsimulation modelling and the variability of day-to-day operation. The conflicts are further separated by the type of vehicle involved in the interaction. “M-M” represents a human vehicle following and interacting with another human vehicle, “A-M” represents a human vehicle following a CAV, and “M-A” represents a CAV following a human vehicle. Interactions involving a CAV following a CAV are excluded, as SSAM incorrectly identifies their behaviour as overly aggressive and potential conflicts when compared to human driving.

11.4.1 Conflicts on Highways

Figure 66 shows that a 10% CPR accompanies a 56% reduction in potential conflicts for the entire network. However, 84% of the 4,341 conflicts observed in this modelling environment occurred in segment midblocks during lane-changing and weaving actions. The remaining 16% of conflicts that occurred at intersections show results substantially less drastic than that observed for the entire network. The dramatic reduction in midblock conflicts is attributed to two factors. The first is that by increasing the proportion of CAVs in the network but holding the total demand constant, platooning operation increased the number of inter-platoon gaps in road segments. These gaps are then leveraged by human vehicles to conduct lane changes. The second cause is that by introducing CAVs into the fleet, the frequency of human vehicle interactions is reduced, naturally adding to the reduction in human vehicle conflicts. Refer to the following figure for the change in total network conflicts by interaction type;




Figure 66: Number of potential accidents within the environment, in 10% increments of CPR.

The inter-platoon headway distribution of vehicles in 10% increments of CPR is provided in Table 40, calculated using detectors throughout the highway. This table indicates that the average headway increases by 17.65% (0.94s). On an average Australian highway at 90km/hr or average suburban road at 60km/hr, this increase in headway amounts to a distance headway increase of 23.5m and 15.7m respectively, justifying the long-term reduction in conflicts with increasing CAV penetration.

The first 20% penetration of CAVs increases the average headway on the highway by 2.4% (0.13s). This amounts to a distance headway increase of 3.18m and 2.12m on highways and suburban roads, respectively. The headway increase is equivalent to approximately 106% the size of a small car or 71% the size of a medium car, which is substantial additional buffer room in lane-changing, especially when the CAVs are designed to be wholly altruistic and cooperative. Perhaps if a demand increase accompanied CAV penetration, average headways in the network would not increase as significantly, and conflict rates would stay consistent until higher penetrations of CAVs. Refer to the following table for the change in average headway as CPR increases;

CAV Penetration	Average Headway (sec)	Change in Average Headway	Headway Standard Deviation (sec)	Change in Headway Standard Deviation
0%	5.326		2.888	
10%	5.394	1.3%	2.980	3.2%
20%	5.453	1.1%	2.978	-0.1%
30%	5.420	-0.6%	2.991	0.4%
40%	5.326	-1.7%	3.087	3.2%
50%	5.453	2.4%	3.110	0.7%
60%	5.562	2.0%	3.087	-0.7%
70%	5.865	5.5%	3.180	3.0%
80%	5.967	1.7%	3.246	2.1%
90%	5.988	0.4%	3.213	-1.0%
100%	6.266	4.6%	3.299	2.7%

Table 40: Tabular representation of the inter-platoon spacing distribution for the fleet along the highway.

Figure 67 provides a visual representation of the change in headway distribution. The headways are segregated in 2s bins. This figure, in complement to Table 40, indicates that the mean headway drift towards larger values and but the occurrence of small (less than 2s) and large (greater than 6s) headways increases. The significance of this outcome is that when demand is held constant and CPR increases, the formation of more platoons and increasing inter-platoon gaps makes lane-changing easier for human vehicles. Refer to the following figure for the change in headway distribution with increasing CPR;




Figure 67: Distribution of highway vehicle spacing, with increasing CPR.

11.4.2 Conflicts at Intersections

When observing rates of conflicts at a granular intersection level (Figure 68), the increase in safety is significantly less. On the contrary, the microsimulation modelling indicates that while conflicts between human vehicles decline,

the conflicts between CAVs and human vehicles increases disproportionately. This outcome results in an initial increase in intersection conflicts at low CPRs, before a decline in conflicts for higher CPRs. For the first 20% of CAV penetration, the signalised, priority, roundabout, and DDI intersection show a change in conflicts of +22%, -87%, -62% and +33% respectively;

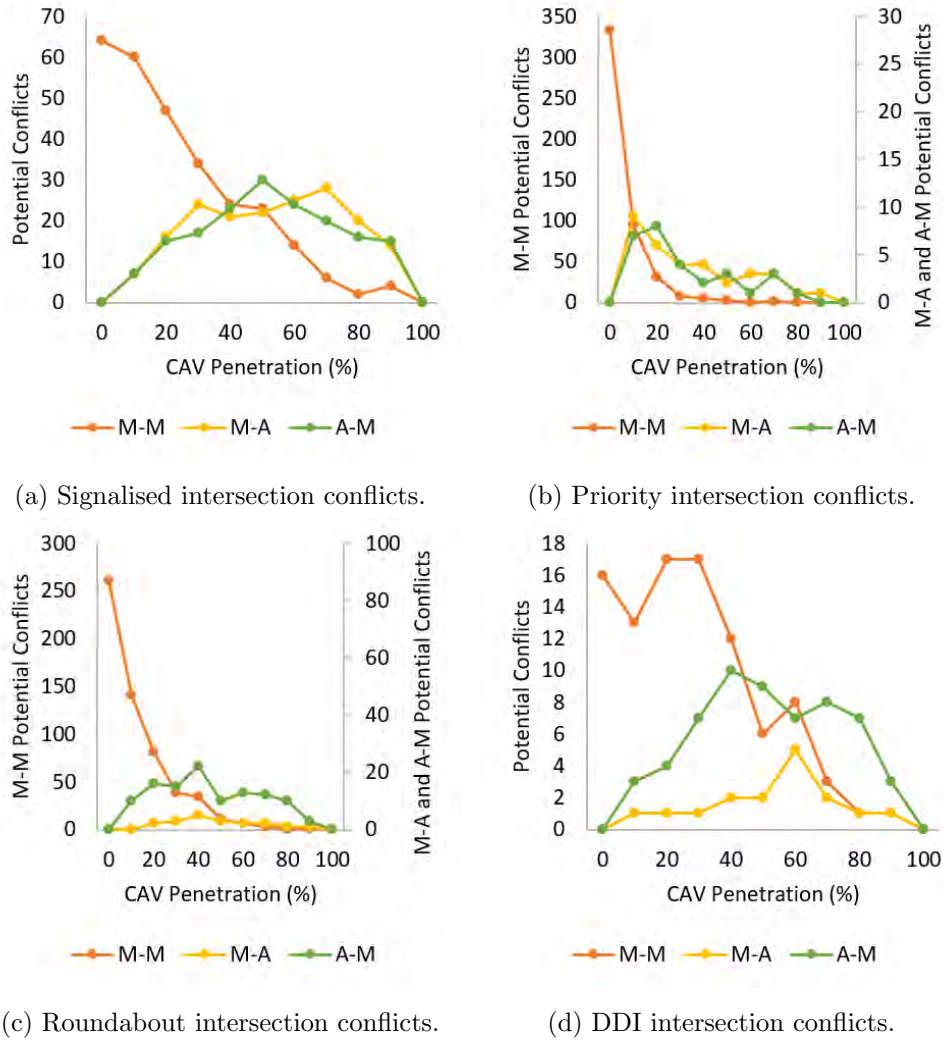


Figure 68: Number of potential accidents during different situations in the microsimulation environment, in 10% increments of CAV penetration.

Signal controlled intersections such as the signalised intersection and DDI intersection both show an increase in total conflicts for low CPRs, whereas the priority-controlled intersections such as the priority intersection and roundabout both show an immediate and significant reduction in conflicts. This

observation provides further confirmation of aggressive vehicle behaviour in the dilemma zone causing safety concerns at signalised intersections. The aggressive and dangerous behaviour during amber signalling phases continues to drive conflict rates in signalised environments.

Also, Figure 68a and Figure 68d indicate that while human following vehicles are responsible for the increase in conflicts at the DDI intersection; this is not the case with the signalised intersection, where CAV following vehicles drive the increase in conflicts. The differentiating factor between the two intersections is the geometry, where turning radii are significantly higher in the signalised intersection than the DDI intersection. This difference raises the question of whether geometry and turning radii continue to affect vehicle safety as the literature's crash prediction models suggest it does for human vehicles, or if this outcome occurred as a result of software limitations.


The commercial microsimulator returns headway as the front-bumper-to-front-bumper distance between vehicles. During a turn, however, the vehicles are not laterally in the same plane, causing the headway passed to the algorithm being marginally higher than actual. The algorithm compensates by accelerating and reducing the gap. SSAM, however, uses both the front and rear bumpers position in the identification of conflicts, with the marginal decrease in the gap between vehicles considered a potential conflict. This behaviour may result in SSAM identifying an artificial increase in conflicts for cases where CAVs are the following vehicle.

Between the 0% CAV base scenario and the 90% CPR scenario, signalised, priority, roundabout, and DDI intersections experience a reduction in conflicts of 48%, 100%, 98%, and 81%, respectively. The greater reduction in conflicts observed at the two priority intersections is attributed to the same factors mentioned above to explain the reduction in segment midblocks. Also, the stringent gap-acceptance and altruistic nature of vehicle cooperation reduces conflicts during lane-changes, such as those occurring at priority intersections.


Results for the DDI intersection presented in Figure 68d shows an increase in conflicts when transitioning from 10% to 20% CAVs, and 50% to 60% CAVs. The increase is attributed to two factors. The first is that the number of

conflicts at this intersection is relatively small, at 16 in the base scenario. Minor fluctuations caused by the variance between random seeds can potentially cause large relative changes to the small base case result. The second reason is that this behaviour occurs in the lower CPR scenarios where the likelihood of human drivers interacting with one-another is still high as CAVs have not dominated the fleet.


Figure 69 further segregates conflicts by type as either a rear-end conflict, lane change conflict, or crossing conflict. When the conflicts are segregated by type, the two signal-controlled intersections once again show distinct results from the two priority-controlled intersections. The priority intersection (Figure 69b) and roundabout intersection (Figure 69c) are dominated by rear-end collisions, with lane-changing and crossing conflicts being substantially lower. This outcome is consistent with the intentions of the CAV control algorithm. The CAV cooperative gap-acceptance protocol forces the vehicle to communicate and facilitate the formation of appropriate gaps, while also having a strict acceptance criterion for gaps. This leads to a low level of lane-changing conflicts shown in these figures;



(a) Signalised intersection conflicts.



(b) Priority intersection conflicts.



(c) Roundabout intersection conflicts.

(d) DDI intersection conflicts.




Figure 69: Number of potential conflicts further segregated by conflict type.

Network geometry explains the low level of crossing conflicts at the two priority junctions. A crossing conflict is defined as occurring between 80° and 180° relative to the vehicle (refer to Figure 63a). However, the concentric direction of motion around a roundabout makes such an angle rare and difficult to achieve. This also applies to the priority intersection, where vehicles travel near parallel to their target lane.

The signalised intersection and DDI show a significantly higher proportion of crossing and lane-changing conflicts. In low penetration scenarios, CAV behaviour introduces more conflicts than the human vehicle conflicts that it mitigates (Figure 69a). This figure also indicates that the CAV as the following vehicle contributes a larger share of the newly occurring conflicts. The CAV gap-acceptance criteria are based on a headway and speed thresholds, meaning that the CAVs do not behave differently depending on the following vehicle type. Allowing them to accept small gaps even with human vehicles, results in one of two outcomes. Either the aggressive $0.5m$ headway is not appropriate for a heterogeneous mixed fleet environment, or the two vehicles being in different planes of motion is resulting in a mismatch between the headway calculated by Vissim and the headway calculated by SSAM. Both have been

discussed prior when commenting on the results of Figure 68. In the interest of conservative safety, the headway of a CAV following a human vehicle is increased to $3.0m$ for the remainder of this thesis.

The lane-changing conflicts increase by up to 5 conflicts at the DDI intersection (Figure 69d). It is difficult to comment on the statistical significance of this outcome, considering that it may be the result of deviations between random seeds and the inherent stochasticity of microsimulation modelling. This outcome requires further investigation.

Table 41 provides the standard deviation for the results attained from the different random seeds. The standard deviation between seeds provides insight into the uncertainty and variability introduced to the network through CAVs. Relative to the observed values, the low standard deviations indicate that CAV presence does not cause uncertain and erratic or dangerous behaviour in human vehicles during mixed fleet scenarios.

Intersection Type	Interaction Type	Penetration											
		0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%	
Complete Network	M-M	225	50	46	11	4	8	16	8	1	1	0	
	A-M	0	10	4	11	12	12	9	7	3	3	0	
	M-A	0	0	7	8	5	3	6	4	7	5	0	
Signalised Intersection	M-M	3	6	7	3	5	0	1	3	2	1	0	
	A-M	0	2	3	2	5	6	6	2	1	2	0	
	M-A	0	2	4	2	4	5	5	5	4	3	0	
Priority Intersection	M-M	41	34	20	9	6	0	2	1	0	0	0	
	A-M	0	5	8	1	2	2	1	1	2	1	0	
	M-A	0	2	0	2	1	0	2	2	1	0	0	
Roundabout Intersection	M-M	5	30	6	9	6	1	2	0	1	1	0	
	M-M	0	0	1	0	1	0	1	1	1	0	0	
	A-M	0	0	1	2	3	1	2	2	2	1	0	
DDI Intersection	M-M	4	1	2	4	1	2	2	1	1	0	0	
	A-M	0	0	1	2	3	1	2	2	2	1	0	
	M-A	0	1	1	0	2	1	2	2	1	0	0	

Table 41: Standard deviation for the number of conflicts for different intersections in the environment.

Table 42 provides the results for only the “M-M” interaction type, normalised to represent the number of conflicts between human vehicles when factoring in the decreasing number of human vehicles in the microsimulation environment;

Intersection	CAV Penetration										
	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
Complete Network	4341	1966	976	580	388	270	180	133	75	70	0
Signalised Intersection	64	67	59	49	40	46	35	20	10	40	0
Priority Intersection	333	106	39	10	8	6	0	3	0	0	0
Roundabout Intersection	261	157	101	56	57	22	15	10	0	10	0
DDI Intersection	16	14	21	24	20	12	20	10	5	10	0

Table 42: Number of conflicts for different intersections in the environment for the “M-M” interaction type, normalised for the volume of human vehicles remaining in the network.

The decreasing trend in this table is at comparable rates to that presented in Figure 66 and Figure 68, indicating that the results are not skewed by a decreasing presence of human vehicles in the network. The modelling indicates that even if human drivers continue to operate in a mixed fleet environment, the likelihood of a human driver being involved in an accident with another human or autonomous vehicle decreases.

11.5 Discussion

A clear pattern emerges when observing the number of potential conflicts for the entire network. As CPR increases, the number of accidents involving human vehicles decreases. This result is attributed to several factors. Firstly, CAVs treat amber lights as red lights, eliminating the dilemma zone that has been established as the cause of many conflicts at signalised intersections [Papaioannou, 2007]. Additionally, CAV behaviour is designed as cooperative and altruistic. The gap-acceptance criteria ensures that vehicles do not merge or change lanes into a gap that results in unsafe headways, nor do the vehicles merge if it results in excessive braking for the merging or following vehicle.

Managing imprudent lane changes with CAVs through cooperation, limits variations in speed and reduces conflicts, which were a major source of conflicts in this environment. The effect would be greater in cases where the following vehicle is distracted, as this leads to an eight-fold increase in the likelihood of collisions [Carbaugh et al., 1998]. Finally, CAVs have access to complete and correct information regarding their surroundings, with vehicle reaction time reduced to the minimum simulation time step of the microsimulator. This information means that CAVs do not make assumptions regarding the spatial and behavioural characteristics of surrounding vehicles. Having access to this

precise information means that minimum safe headway requirements are never violated, further reducing the likelihood of potential conflicts.

The results indicate that CAVs in the short term have the potential to increase conflicts at intersections, findings consistent with the Deluka Tibljas et al., [Deluka Tibljaš et al., 2018] and Turner and Roozenburg [Turner & Roozenburg, 2006]. The 0.5m headway is inappropriate in a mixed fleet setting with substantial heterogeneous behaviour. Once homogeneity returns to the vehicle fleet in higher CAV penetrations, the use of this headway results in a decrease in conflicts. However, it should be noted that demand was kept constant throughout the study, meaning that an increase in CAVs leads to an increase in gaps in the network. The increase in gaps may be the cause of the reduction in conflicts, warranting a further investigation with a higher DoS.

In the highway environment, the results of this study show benefit to safety from a 10% CAV penetration. Other studies do not show significant improvements until a CAV penetration of 20% to 30% ([Rahman et al., 2019], [Papadoulis et al., 2019]). The limitations of the other studies are in their emulation of CAV behaviour. These studies have opted to use rudimentary autonomous features such as ACC and lane guidance to emulate what is a highly promising technology that is currently in its infancy. To “implement lower level automation features under a connected vehicle environment which is available in many vehicles in the market” [Rahman et al., 2019] has significant caveats. CAV technology will not be ready, implemented, or see significant market penetration for decades to come. Computational capabilities and data processing techniques improve at an exponential rate. The approach used in this study in applying highly refined CAV operations such as precise vehicle coordination, minimal headways, and complete cooperation may be optimistic. CAV operation may not match the assumptions made in this study either, but it provides new insights into defining and refining the assumptions of CAV behaviour for future studies.

The results also show that when CAV penetration increases and homogeneous operation returns to the network, potential conflicts are likely to decrease, consistent with much of the literature. The initial 10% penetration of CAVs results in an approximate 56% network-wide reduction in conflicts, with 8%

of accidents at this penetration involving a CAV. An accurate cost of a CAV is difficult to identify but is estimated to start at \$250,000 (USD) [LeVine, Steve, 2017]. A 56% reduction in network-wide conflicts (with 8% involving a CAV) still leads to an estimated economic saving of 22% of total accident costs when considering only vehicle replacement costs. As CAVs become cheaper, this economic benefit will increase substantially. However, using the higher purchase price, this amounts to a saving of approximately AUD\$786 million annually for Australia alone. This figure is calculated based on the AUD\$2.8 billion cost to the economy found in the BITRE study [BITRE, 2009], adjusted for inflation at 1.9% per annum using the consumer price index (CPI) [Australian Bureau of Statistics, 2020]. This figure is conservative and does not include the potential savings arising from prevented damage to property and infrastructure. This figure also does not consider the effects of reducing CAV prices, which would further increase economic savings.

The economic savings figure also does not consider the savings in infrastructure development. Most transport infrastructure contains redundancies designed to facilitate the movement of emergency vehicles or operate if an accident occurs. This emergency infrastructure consists of emergency-vehicle-only lanes on motorways or shoulders on arterials. These savings in construction will be significant for a country such as Australia, where development in metropolitan areas has occurred near existing transport corridors, without residual space to grow. All future projects in the Sydney and Melbourne long-term master plans contain vastly more expensive tunnelling components where cost increases exponentially with the number of lanes and tunnel diameter.

The DDI (Figure 68d), designed to be safer than conventional intersections, experiences a more gradual improvement in safety. Reductions in conflicts do not begin until CPR reaches 30%, indicating that the inherent safety of DDIs and lack of conflict points requires a higher penetration of CAVs to extract similar benefits to that extracted at lower penetration rates from traditional intersections. The signalised environment (Figure 68a) showed the highest increase in the number of conflicts of the four scenarios for the initial 10% penetration of CAVs. This has been attributed to more dangerous driving by human drivers in the dilemma zone, decreasing CAV safety at low headways.

When observing the standard deviations between seeds, a notable observation is that the standard deviation is higher for lower CPRs. This indicates that the presence of CAVs in a mixed fleet environment disrupts the homogeneity in both a completely human and heavily autonomous environment. While human behaviour and personal preferences regarding speed, headway, and other driving elements cause variances in a traffic stream, agents do not generally deviate substantially from accepted bounds. Speed is regulated by design limits, and headway is governed by risk aversion. CAVs also exhibit homogeneity, as they follow speed and headway requirements dictated by algorithms. However, whilst the CAVs in this study adhere to a $0.5m$ headway, the human vehicles standstill at a distance four times higher than this, and drive with a headway significantly higher. The difference in the fundamental behaviour of these vehicle types disrupts the order and uniformity of the fleet. Consequently, the seeds exhibit a high degree of variability regarding the prevalence and structure of gaps in conflicting traffic streams, which in turn increases the variability in the number of potential network conflicts. This behaviour, however, decreases as CAV penetration increases and uniformity in behaviour once again returns to the network.

A limitation of this chapter is in its decision to hold network demand constant as CPR increases. Section 11.3.2 mentions that the environment is undersaturated, and demand was doubled to ensure the presence of latent demand in the base network. However, Chapter 7 demonstrated that intersection throughput increases by up to 110% with increasing CPR. Despite the system starting in a congested state, it did not remain this way as CPR increased. Decreasing levels of congestion gave rise to reductions in vehicle interactions in the motorway section and at the priority intersections, which may have been a contributing factor to the decrease in conflicts at these intersections. Further investigation is warranted to provide additional insights into whether a safer network is the byproduct of an autonomous network or a less congested network.

The implementation of CAVs leads to the idealised vision of a zero-accident environment. In the interim, smart infrastructure and design decisions are still required to maintain safety. Crash prediction models are useful in assessing the safety of designs and will need recalibration to incorporate the presence and conflict characteristics of CAVs. A reduction in the likelihood of accidents

as observed here means that the way vehicular insurance is structured may also require reform. As society embraces the sharing economy with CAVs servicing travel requirements on-demand, and their tendency to be safer as shown in this study, insurance agencies may benefit from a restructure in the way that insurance is sold. Policies involving insuring the driver as opposed to insuring the vehicle may warrant investigation.

11.6 Conclusion

This chapter investigated the effect of CAV penetration on road network safety. This investigation was conducted using microsimulation modelling, with CAV behaviour emulated using the custom developed external control algorithm, and the likelihood of potential conflicts identified using the surrogate safety assessment module. The contribution of this chapter lies in its implementation of this custom control algorithm for CAV emulation and assessment of CAV safety based on microsimulation testing. The results indicate that while CAV operation shows a significant overall improvement in safety, this improvement is concentrated at segment midblocks. The signalised intersections show an increase in potential collisions for low penetration rates, while the priority intersections show an immediate and significant decrease. As CAV penetration increases, the potential conflicts in all settings declines. Reductions in potential conflicts are greater in priority-controlled intersections such as roundabouts and give-way environments, as compared to signalised intersections.

Acknowledgement of Publication

This work was successfully published as a journal article - Viridi, N., Grzybowska, H., Waller, S. T., & Dixit, V. (2019). A safety assessment of mixed fleets with connected and autonomous vehicles using the Surrogate Safety Assessment Module. *Accident Analysis & Prevention*, 131, 95-111.

12 Summation and Concluding Thoughts

The final chapter of this thesis provides a brief overview and summation of the work contained within, as well as parting thoughts.

The introductory chapter to this thesis established the necessity for this work. CAVs are an unquestionably disruptive force whose development is driven by the private sector. Often, development by the private sector, and without consultation with the government, outstrips the policy and legal frameworks set in place for emerging technology. The result is that governments often exhibit a lack of preparedness when faced with rapid change. Consider recent disruptions in the ride-sharing sector with Uber, the recreational housing sector with AirBNB, the rise of non-fiat currency with Bitcoin, or changes to the retail landscape with e-commerce. In each of these instances, private markets have progressed while policy and legislation have been retroactively amended. A robust modelling framework is critically necessary to practically test scenarios and shape the integration of CAV technology with society, to prevent a repeat pattern also occurring with CAVs.

After the introductory chapter, Chapter Two provided a review of CAV design and hardware, assessing their potential capabilities and features. Technology changes at a rapid and accelerating pace, so it is inappropriate to assume that the CAV technology of today is an initial reflection of its mature capabilities. Adaptations and implementations of autonomous and CAV technology are still being envisioned. However, baseline assumptions are necessary to define to initiate planning and modelling efforts, allowing policymakers and researchers to restrict the infinite domain of possibilities. Understanding CAV hardware allows for assumptions regarding the three key CAV features that affect behaviour, reaction time, platooning headway, and cooperation and communication.

Having understood the expected behavioural characteristics of CAVs in Chapter Two, Chapter Three conducted an in-depth review of the literature pertaining to human vehicle behaviour in microsimulation. The purpose of this literature review was three-fold; to understand the vast styles of models necessary to emulate vehicle behaviour in a microsimulation setting, to evaluate

techniques used in the development of previously influential models, and to assess the applicability of existing human vehicle models to the behaviour of CAVs. Microsimulation behavioural models fall into one of three major categories, car-following, lane-changing, and gap-acceptance. Models from as early as 1959 have used physical phenomena to describe human behaviour, drawing inspiration from fluid dynamics to spring and damper systems. Generally, human behavioural models aim to capture the stochastic behaviour of human drivers. This motivation leads to the development of models that contain the selfish, flawed, and inefficient tendencies exhibited by human drivers. For this reason, many of the human behavioural models proposed in the literature are not applicable for use with CAVs in microsimulation.

The literature has identified that human behavioural models do not appropriately model anticipated CAV behaviour, and so has proposed a range of new microsimulation models. Chapter Four provided a literature review detailing the novel and innovative models used to emulate CAV behaviour. While the research is relatively young as compared to that of human behaviour emulation, the literature already contains a vast array of modelling techniques for CAVs. Many of the proposed methods leverage modern computational capabilities and use complex techniques such as machine learning and artificial intelligence. Approaches such as these result in models that risk overfitting to training data, are difficult to generalise to a range of driving situations, impossible to trace through for errors and inconsistencies, and generally inaccessible to gain a sound understanding of the vehicle's decision making processes. For these reasons, while many of the models in the literature demonstrated accurately replicating CAV behaviour under limited circumstances, their tendency to focus on a specific aspect of CAV behaviour makes them difficult to aggregate as part of a wholistic modelling framework. The nature of models in the literature also increases the difficulty of applying them to a range of innovative scenarios and examining network performance under new policy and control mechanisms.

Having completed a detailed literature review and identifying the need for a simple, complete, and adaptable framework for microsimulation CAV emulation, Chapter Five developed a custom CAV control algorithm. The algorithm addresses the three key components of microsimulation behaviour, car-

following, lane-changing, and gap-acceptance. Additionally, the framework proposes a trajectory forecasting cooperative merge algorithm to guide vehicles autonomously through merge junctions downstream. The car-following component of the algorithm heavily relies on the instantaneous reaction time and communication capabilities of CAVs to chain vehicles together as platoons. The car-following algorithm proposes a novel way for platoon leaders to traverse the network, while platoon followers act as convoy units. The discretionary lane-changing component of the algorithm uses a cooperative game-theory-style approach that collectively maximises the utility of the system of vehicles directly effected by the lane change. The gap-acceptance component outlines a stringent set of rules for critiquing and accepting a gap that prioritises safety, comfort, and impacts on immediate vehicle externalities, once again highlighting the core focus on vehicle communication.

Chapter Five also explained the models used to emulate human vehicles. Emulation of human vehicles is necessary to complement the emulation of CAVs in mixed fleet settings, as mixed fleets will persist for decades before an autonomous fleet. Each study uses either the Aimsun or Vissim commercial microsimulators. For this reason, the native car-following, lane-changing, and gap-acceptance models used in both of these commercial packages are retained. The Aimsun package exclusively uses the Gipps models. The Gipps car-following model is a collision-avoidance model that restricts vehicle kinematics based on safety constraints. The Gipps lane-changing model is a rule-based model that qualitatively explains the scenarios under which a driver would initiate a lane change. Finally, the Gipps gap-acceptance model is derived from the car-following model, where the safe car-following thresholds are used to define minimum upstream and downstream gaps for lane-changing vehicles. Vissim uses the Wiedemann car-following model, which is a psychophysical model that discretises the domain of human car-following behaviour and describes each region with a unique mathematical equation. Vissim interconnects lane-changing and gap-acceptance. The modeller defines a range of parameters that control the gap-acceptance criteria and aversion to risk for network agents.

The framework developed in Chapter Five appropriately emulates CAV behaviour of instantaneous reaction times, coordinated and cooperative manoeu-

ving, and vehicle platooning. However, in pursuit of network efficiency, passenger safety and comfort must not be compromised. Chapter Six conducted a detailed kinematic assessment on the modelling framework and a sensitivity study on the framework parameters, a study that is generally not completed for proposed CAV models in the literature. The kinematic assessment begins with the development of kinematic-time plots for a platoon of ten vehicles for two scenarios. The first scenario is regular car-following as vehicles approached a stoplight, and the second scenario modelled a vehicle in the parallel lane conducting an imprudent lane change into the vehicle platoon. The kinematic-time plots for both scenarios demonstrated homogeneous coordinated platoonal action and kinematic reactions well within the confines of safe vehicle behaviour. Next, Chapter Six conducted a kinematic assessment using a small network of intersections and the trajectory data generated by the microsimulator. A detailed ANOVA demonstrated that while CAV penetration changed the high-level vehicle velocity, the underlying vehicle acceleration and jerk remain unchanged. This outcome implied that CAV behaviour results in more efficient driving without affecting driver aggression (acceleration) and driver comfort (jerk). Finally, this chapter concluded with a sensitivity analysis of the three parameters proposed in this framework. The analysis demonstrated that the parameters had a tractable and predictable change to network behaviour, which remained unaffected by minor changes in the parameters.

Chapter Seven conducted an investigation into the performance of intersections subjected to an incremental loading of CAVs. Infrastructure development is a timely and expensive process, and so requires forethought and planning. Without understanding the network impacts of CAVs, long-term infrastructure planning is unwise. This investigation consisted of four intersections, a signalised intersection with a filtered right turn, signalised intersection with a dedicated right turn, roundabout intersection, and priority intersection. CAVs were loaded onto the network in 20% increments, and key performance metrics such as throughput, delay, queue length, and speed were assessed. The study indicated that intersection performance increased substantially, with throughput increasing by up to 110%, delay reducing by up to -63%, queue length decreasing by up to -28%, and average speed increasing by up to 130%. This study also identified that not all intersections improve at the same rate or to the same magnitude. While any improvement may seem beneficial, dispro-

portionate improvements may lead to increased congestion and delay as high-efficiency signalised intersections in the arterial network deliver greater quantities of vehicles to low-efficiency intersections such as roundabouts and priority junctions in the local road network. The miss-match of adjacent intersections could lead to the development of new bottlenecks and network failure points. Additionally, this chapter demonstrated that the CAV behavioural algorithm developed in this thesis successfully eliminated the perception-reaction lag from vehicle standstill movement, leading to more efficient use of green-time and the potential to shorten intersection phasing while increasing the turnover of movements.

The previous studies have focused on CAV impacts on vehicle experience, and have found drastic improvements in capacity, speed, travel time, queuing, and safety. Rather than allowing private vehicles to retain these benefits, improving travel conditions, inducing demand, and again worsening travel conditions, consider reallocating CAV travel time savings to other network agents. Chapter Eight assessed the potential to redistribute CAV travel time savings by implementing pedestrian-priority signalling schemes through traffic signal re-optimisation. The investigation developed a microsimulation environment of two distinct areas, one that is dominated by vehicles and another that is dominated by pedestrians. At both environments, CAVs incrementally penetrated the fleet, and three distinct signalling regimes were examined. The regimes are the base case signalling regime used at each intersection today, the vehicle-priority signalling regime that minimises delay for vehicles, and the pedestrian-priority signalling regime that minimises delay for pedestrians. Additionally, a pedestrian simulator was developed that calculated the delay of randomly arriving pedestrians. This study found that in the vehicle-dominated environment, transitioning from a vehicle-priority model to a pedestrian-priority model was not feasible, regardless of the CAV penetration. The increase in delay for vehicles could not be offset by the reduction in delay for the pedestrians. The pedestrian-dominated environment, however, demonstrated that such a transition was possible when CAV penetration exceeds 40%. At this penetration, the average delay per agent is -1.35% lower when comparing a pedestrian-priority scheme to the base case regime, resulting from a -48.5% reduction in pedestrian delay and an 80.1% increase in vehicle delay. This study demonstrated that the efficiencies introduced by CAVs could be redis-

tributed to other agents in the network to create a more equitable mobility environment in dense city centres.

Motorways form major high capacity links in transport networks. Their proposal, development, widening, and maintenance is often met with mixed reception. As space in major cities worldwide diminishes, extracting greater utility from existing motorway infrastructure may cater to growing travel and population demand. Chapter Nine assessed the impact of CAV behaviour at weaving and merging junctions on motorways. Much of the literature claims that CAV behaviour will substantially improve capacity, while other parts of the literature claim that it will give rise to new congestion points and bottlenecks. This study, which consisted of three investigations, found both to be true. The first investigation determined the capacity of a weaving section in 25% increments of CAVs and for weaving proportions ranging from 5% to 40%. The results indicate that for low weaving environments, the section capacity can increase by up to 83%. However, CAVs are hindered in high weaving proportion environments, with capacity improvements being only 14% for a 40% weaving proportion. The forced cooperation between vehicles and the mandated altruism by design reduced traffic flow efficiency when lane-changing and weaving became excessive. The second investigation assessed the impact of CAVs on the performance of a zipper merge. Using the trajectory forecasting cooperative merge algorithm, CAVs increased throughput by up to 131%. The vehicles did not increase average speed, where the gain in throughput was the result of eliminating start-stop conditions at busy merge sites and maintaining a smooth and perpetual flow by actively adjusting trajectories from a centralised controller upstream of the merge junction. Finally, this study investigated the impact of segregating the CAV fleet into multiple vehicle types and offering certain vehicle classes preferential treatment in the network. This study indicates that providing vehicles with preferential treatment does not worsen network performance unless the environment is operating near capacity, or the environment has excessive weaving proportions (20% or above).

Microsimulation is a useful tool in evaluating network impacts and is used in conjunction with macrosimulation modelling. While microsimulation provides insights into localised decision impacts and agent interaction, macrosimulation provides traffic flows for a city-wide scope, making it critical in large-scale in-

frastructure planning. Chapter Ten investigated the BPR volume delay function for strategic modelling, and its applicability to mixed fleets and CAVs. CAVs will be well integrated into society's mobility options in the near future, so the integrity of macrosimulation modelling must be investigated for emerging technology. This study developed three microsimulation models, which synthesised a large quantity of vehicle demand and travel time data. Using 30% of the synthetic data, the predictive quality of the BPR function was tested for varying network degree of saturations and CAV penetration rates. The BPR function was shown to exhibit an RMSE of 11.56, 9.68 and 7.95 for the three modelling environments. The BPR function tended to incorrectly predict travel time for highly congested environments and higher CAV penetration rates. The same 30% of the data was used to recalibrate the function, deriving a linear relationship between both the α and β parameters and CAV penetration rate. The updated BPR function was validated using the remaining 70% of the synthetic data and showed an RMSE of 7.57, 9.98 and 5.36 for the three different modelling environments. The updated model better predicted travel time in 67% of cases and showed an improvement in travel time predictions of 542.6% for the entire validation dataset.

Following passenger comfort and network performance, safety is the next great barrier impeding technological adoption. Chapter Eleven used the surrogate safety assessment module to conduct a study into the potential conflict rates of CAVs. During a simulation, the microsimulator records the trajectory information for each vehicle for each time increment. The surrogate safety assessment module uses the trajectory data to forecast vehicle positions and estimate the likelihood of a collision. This study found that while CAV penetration is low, CAV behaviour could increase the likelihood of collisions between vehicles. Heterogeneous behaviour leads to the development of dangerous driving situations where the smaller gap-acceptance policies of CAVs result in unsafe conditions arising for human drivers. Additionally, stringent compliance with traffic lights at intersections and quicker reaction times increased the likelihood of crossing angle collisions at signalised intersections. However, once the network reaches a tipping point (approximately 40% - 50% CAV) and homogeneous behaviour returns to the network, the likelihood of conflicts also dissipates. Priority controlled intersections and roundabouts showed immediate and drastic reductions in potential conflicts as a result of CAV penetration.

The most likely accident at priority controlled junctions and roundabouts is a rear-end collision. The CAVs are considered to be faultless and maintain a generous headway to human vehicles, accounting for the significant and immediate reduction in conflicts in these environments.

This thesis added to the body of understanding by proposing a novel CAV behavioural control framework and unique trajectory forecasting algorithm for cooperative merging. In developing this framework and algorithm, assumptions were made regarding imprecise and unavailable vehicle behaviour parameters. The advantage of this framework is in its ability to change and alter independent modules when further refinement to behaviour needs to be made. The framework and algorithm were then applied in a range of studies that provided further understanding regarding passenger comfort, network performance, vehicle safety, motorway capacity, signalling optimisation, and strategic modelling parameters. While the overarching results indicate strong benefits in all regards from CAV integration, many chapters warn of potentially problematic effects if CAV integration is not planned and regulated.

However, the autonomous future is exciting and ripe with possibilities. Innovations to wasted in-vehicle time by vehicle manufacturers promises to change the way society interacts with vehicles. The correct economic circumstances may render private vehicle ownership antiquated in cities with strong public and active transport systems. The private vehicle once formed a major centroid for modern life, may revert to the importance given to any other household appliance or tool. Or major strides in CAV technology and affordability may rekindle the dwindling attachment of people to their private vehicles, inciting a new era of private vehicle trips, increased travel demand, and greater levels of congestion. A strong modelling framework and novel scenario testing for policy and integration, as that presented in this thesis, will both aid in ensuring efficient, equitable, environmental, and sustainable outcomes from connected and autonomous vehicles.

References

- [Abramowitz & Stegun, 1948] Abramowitz, M. & Stegun, I. A. (1948). *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office.
- [Ackermann, 2012] Ackermann, J. (2012). *Robust control: Systems with uncertain physical parameters*. Springer Science & Business Media.
- [Adebisi, 2018] Adebisi, A. E. (2018). *Delay at Signalized Intersections in Mixed Conventional and Connected Autonomous Vehicle (CAV) Traffic Conditions*. PhD thesis, Florida Agricultural and Mechanical University.
- [Adminaite et al., 2015] Adminaite, D., Allsop, R., & Jost, G. (2015). Making walking and cycling on europe’s roads safer.
- [Adnan et al., 2018] Adnan, N., Nordin, S. M., bin Bahrudin, M. A., & Ali, M. (2018). How trust can drive forward the user acceptance to the technology? in-vehicle technology for autonomous vehicle. *Transportation research part A: policy and practice*, 118, 819–836.
- [Aghabayk et al., 2013] Aghabayk, K., Sarvi, M., Young, W., & Kautzsch, L. (2013). A novel methodology for evolutionary calibration of vissim by multi-threading. In *Australasian Transport Research Forum* (pp. 1–15).
- [Ahmed, 1999] Ahmed, K. I. (1999). Modeling drivers’ acceleration and lane changing behavior.
- [Akçelik, 1991] Akçelik, R. (1991). Travel time functions for transport planning purposes: Davidson’s function, its time dependent form and alternative travel time function. *Australian Road Research*, 21(3).
- [Al Henawy & Schneider, 2011] Al Henawy, M. & Schneider, M. (2011). Integrated antennas in ewlb packages for 77 ghz and 79 ghz automotive radar sensors. In *2011 8th European Radar Conference* (pp. 424–427).: IEEE.
- [Alsalmi et al., 2018] Alsalmi, R., Dixit, V. V., & Gayah, V. V. (2018). On the existence of network macroscopic safety diagrams: Theory, simulation and empirical evidence. *PloS one*, 13(8).

- [Amin & Maurya, 2015] Amin, H. J. & Maurya, A. K. (2015). A review of critical gap estimation approaches at uncontrolled intersection in case of heterogeneous traffic conditions. *Journal of transport literature*, 9(3), 5–9.
- [Anderson & Knoop, 2018] Anderson, P. & Knoop, V. L. (2018). *Pedestrians and cars in urban networks: effect of various interaction strategies*. Technical report.
- [Antoniotti et al., 1997] Antoniotti, M., Deshpande, A., & Girault, A. (1997). Microsimulation analysis of automated vehicles on multiple merge junction highways. In *Systems, Man, and Cybernetics, 1997. Computational Cybernetics and Simulation., 1997 IEEE International Conference on*, volume 1 (pp. 839–844).: IEEE.
- [Anvari et al., 2016] Anvari, B., Bell, M. G., Angeloudis, P., & Ochieng, W. Y. (2016). Calibration and validation of a shared space model: case study. *Transportation Research Record*, 2588(1), 43–52.
- [Applanix Corporation, 2015] Applanix Corporation (2015). Poslv specifications. applanix.com/pdf/specs/POSLV'Specifications'dec'2015.pdf.
- [Arkell, 2017] Arkell, C. (2017). Legion for aimsun: modelling the interactions between pedestrians, light rail and general traffic. In *Australian Institute of Traffic Planning and Management (AITPM) National Conference, 2017, Melbourne, Victoria, Australia*.
- [Ashworth, 1968] Ashworth, R. (1968). A note on the selection of gap acceptance criteria for traffic simulation studies. *Transportation Research/UK*.
- [Australian Bureau of Statistics, 2020] Australian Bureau of Statistics (2020). 6401.0 - consumer price index, australia, dec 2019. abs.gov.au/ausstats/abs@.nsf/mf/6401.0.
- [Austroads, 1993] Austroads (1993). Roundabouts—guide to traffic engineering practice-part 6.
- [Bang & Ahn, 2017] Bang, S. & Ahn, S. (2017). Platooning strategy for connected and autonomous vehicles: transition from light traffic. *Transportation Research Record*, 2623(1), 73–81.

- [Banks, 1991] Banks, J. H. (1991). Two-capacity phenomenon at freeway bottlenecks: a basis for ramp metering? *Transportation Research Record*, (1320).
- [Baskar et al., 2008] Baskar, L. D., De Schutter, B., & Hellendoorn, H. (2008). Model-based predictive traffic control for intelligent vehicles: Dynamic speed limits and dynamic lane allocation. In *2008 IEEE intelligent vehicles symposium* (pp. 174–179).: IEEE.
- [Bazilinsky et al., 2015] Bazilinsky, P., Kyriakidis, M., & de Winter, J. (2015). An international crowdsourcing study into people’s statements on fully automated driving. *Procedia Manufacturing*, 3, 2534–2542.
- [Bento et al., 2013] Bento, L. C., Parafta, R., Santos, S., & Nunes, U. (2013). Intelligent traffic management at intersections: Legacy mode for vehicles not equipped with v2v and v2i communications. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)* (pp. 726–731).: IEEE.
- [Beza & Zefreh, 2019] Beza, A. D. & Zefreh, M. M. (2019). Potential effects of automated vehicles on road transportation: A literature review. *Transport and Telecommunication Journal*, 20(3), 269–278.
- [BITRE, 2009] BITRE (2009). *Cost of road crashes in Australia 2006*. Bureau of Infrastructure, Transport and Regional Economics.
- [Blau, 2015] Blau, M. A. (2015). *Driverless vehicles’ potential influence on cyclist and pedestrian facility preferences*. PhD thesis, The Ohio State University.
- [Bloomquist et al., 2005] Bloomquist, D., McVay, M., Larsson, E. G., & Dumas, C. (2005). Autonomous highway traffic modules. US Patent 6,900,740.
- [Botello et al., 2019] Botello, B., Buehler, R., Hankey, S., Mondschein, A., & Jiang, Z. (2019). Planning for walking and cycling in an autonomous-vehicle future. *Transportation research interdisciplinary perspectives*, 1, 100012.
- [Boyd, 1968] Boyd, T. (1968). The self-starter. *Technology and Culture*, 9(4), 585–591.

- [Brackstone & McDonald, 1999] Brackstone, M. & McDonald, M. (1999). Car-following: a historical review. *Transportation Research Part F: Traffic Psychology and Behaviour*, 2(4), 181–196.
- [Brilon et al., 1999] Brilon, W., Koenig, R., & Troutbeck, R. J. (1999). Useful estimation procedures for critical gaps. *Transportation Research Part A: Policy and Practice*, 33(3-4), 161–186.
- [Bureau of Transport statistics, 2012] Bureau of Transport statistics (2012). *Strategic Travel Model Standard Outputs*.
- [California, State of, 2018] California, State of (2018). Autonomous vehicle disengagement reports 2017. [dmv.ca.gov/portal/dmv/detail/vr/autonomous/disengagement report 2017](https://dmv.ca.gov/portal/dmv/detail/vr/autonomous/disengagement%20report%202017).
- [Campbell et al., 2018] Campbell, S., O’Mahony, N., Krpalcova, L., Riordan, D., Walsh, J., Murphy, A., & Ryan, C. (2018). Sensor technology in autonomous vehicles: a review. In *2018 29th Irish Signals and Systems Conference (ISSC)* (pp. 1–4).: IEEE.
- [Carbaugh et al., 1998] Carbaugh, J., Godbole, D. N., & Sengupta, R. (1998). Safety and capacity analysis of automated and manual highway systems. *Transportation Research Part C: Emerging Technologies*, 6(1-2), 69–99.
- [Cassidy & Bertini, 1999] Cassidy, M. J. & Bertini, R. L. (1999). Observations at a freeway bottleneck. *Transportation and Traffic Theory*, (pp. 107–146).
- [Cassidy & Rudjanakanoknad, 2005] Cassidy, M. J. & Rudjanakanoknad, J. (2005). Increasing the capacity of an isolated merge by metering its on-ramp. *Transportation Research Part B: Methodological*, 39(10), 896–913.
- [Chu et al., 2003] Chu, L., Liu, H. X., Oh, J.-S., & Recker, W. (2003). A calibration procedure for microscopic traffic simulation. In *Proceedings of the 2003 IEEE International Conference on Intelligent Transportation Systems*, volume 2 (pp. 1574–1579).: IEEE.
- [Claros et al., 2015] Claros, B. R., Edara, P., Sun, C., & Brown, H. (2015). Safety evaluation of diverging diamond interchanges in missouri. *Transportation Research Record*, 2486(1), 1–10.

- [Clements & Kockelman, 2017] Clements, L. M. & Kockelman, K. M. (2017). Economic effects of automated vehicles. *Transportation Research Record*, 2606(1), 106–114.
- [Corfield, 2017] Corfield, G. (2017). Tesla death smash probe: Neither driver nor autopilot saw the truck. theregister.co.uk/2017/06/20/tesla-death-crash-accident-report-ntsb.
- [Cottam, 2018] Cottam, B. J. (2018). Transportation planning for connected autonomous vehicles: How it all fits together. *Transportation Research Record*, 2672(51), 12–19.
- [Daganzo & Knoop, 2016] Daganzo, C. F. & Knoop, V. L. (2016). Traffic flow on pedestrianized streets. *Transportation Research Part B: Methodological*, 86, 211–222.
- [Davidson & Spinoulas, 2015] Davidson, P. & Spinoulas, A. (2015). Autonomous vehicles: what could this mean for the future of transport. In *Australian Institute of Traffic Planning and Management (AITPM) National Conference, Brisbane, Queensland*.
- [Deluka Tibljaš et al., 2018] Deluka Tibljaš, A., Giuffrè, T., Surdonja, S., & Trubia, S. (2018). Introduction of autonomous vehicles: Roundabouts design and safety performance evaluation. *Sustainability*, 10(4), 1060.
- [Dervisoglu et al., 2009] Dervisoglu, G., Gomes, G., Kwon, J., Horowitz, R., & Varaiya, P. (2009). Automatic calibration of the fundamental diagram and empirical observations on capacity. In *Transportation Research Board 88th Annual Meeting*, volume 15.
- [Ding et al., 2013] Ding, C., Wang, W., Wang, X., & Baumann, M. (2013). A neural network model for driver’s lane-changing trajectory prediction in urban traffic flow. *Mathematical Problems in Engineering*, 2013.
- [Ding, 2005] Ding, S. (2005). Model-based fault diagnosis in dynamic systems using identification techniques, silvio simani, cesare fantuzzi and ron j. patton, springer: London, 2003, 282pp. isbn 1-85233-685-4. *International Journal of Robust and Nonlinear Control*, 15(11), 509–512.

- [Dixit, 2013] Dixit, V. V. (2013). Behavioural foundations of two-fluid model for urban traffic. *Transportation Research Part C: Emerging Technologies*, 35, 115–126.
- [Dixit et al., 2016] Dixit, V. V., Chand, S., & Nair, D. J. (2016). Autonomous vehicles: disengagements, accidents and reaction times. *PLoS one*, 11(12), e0168054.
- [Dresner & Stone, 2004] Dresner, K. & Stone, P. (2004). A reservation-based multiagent system for intersection control. *IFAC Proceedings Volumes*, 37(8), 136–141.
- [Dresner & Stone, 2008] Dresner, K. & Stone, P. (2008). A multiagent approach to autonomous intersection management. *Journal of artificial intelligence research*, 31, 591–656.
- [Elhenawy et al., 2015] Elhenawy, M., Elbery, A. A., Hassan, A. A., & Rakha, H. A. (2015). An intersection game-theory-based traffic control algorithm in a connected vehicle environment. *2015 IEEE 18th international conference on intelligent transportation systems*, (pp. 343–347).
- [Epting, 2019] Epting, S. (2019). Automated vehicles and transportation justice. *Philosophy & Technology*, 32(3), 389–403.
- [Evans & Rothery, 1973] Evans, L. & Rothery, R. (1973). *Experimental measurements of perceptual thresholds in car-following*. Number 464.
- [Fadhloun & Rakha, 2019] Fadhloun, K. & Rakha, H. (2019). A novel vehicle dynamics and human behavior car-following model: Model development and preliminary testing. *International Journal of Transportation Science and Technology*.
- [Fajardo et al., 2011] Fajardo, D., Au, T.-C., Waller, S., Stone, P., & Yang, D. (2011). Automated intersection control: Performance of future innovation versus current traffic signal control. *Transportation Research Record: Journal of the Transportation Research Board*, (2259), 223–232.
- [Farah et al., 2009] Farah, H., Bekhor, S., Polus, A., & Toledo, T. (2009). A passing gap acceptance model for two-lane rural highways. *Transportmetrica*, 5(3), 159–172.

- [Fayazi & Vahidi, 2018] Fayazi, S. A. & Vahidi, A. (2018). Mixed-integer linear programming for optimal scheduling of autonomous vehicle intersection crossing. *IEEE Transactions on Intelligent Vehicles*, 3(3), 287–299.
- [Fritzsche, 1994] Fritzsche, H.-T. (1994). A model for traffic simulation. *Traffic Engineering+ Control*, 35(5), 317–21.
- [Frosch, 2017] Frosch, C. (2017). Evaluation of shared space to reduce traffic congestion: A case study on west virginia university’s downtown campus.
- [Gates, 2017] Gates, B. (2017). The race for self-driving cars. [nytimes.com/interactive/2016/12/14/technology/how-self-driving-cars-work.html](https://www.nytimes.com/interactive/2016/12/14/technology/how-self-driving-cars-work.html).
- [Gattis & Low, 1999] Gattis, J. & Low, S. T. (1999). Gap acceptance at atypical stop-controlled intersections. *Journal of Transportation Engineering*, 125(3), 201–207.
- [Gazis et al., 1959] Gazis, D. C., Herman, R., & Potts, R. B. (1959). Car-following theory of steady-state traffic flow. *Operations research*, 7(4), 499–505.
- [Gettman et al., 2008] Gettman, D., Pu, L., Sayed, T., Shelby, S., & Siemens, I. (2008). *Surrogate safety assessment model and validation*. Technical report, United States. Federal Highway Administration. Office of Safety Research and
- [Ghasemi & Zahediasl, 2012] Ghasemi, A. & Zahediasl, S. (2012). Normality tests for statistical analysis: a guide for non-statisticians. *International journal of endocrinology and metabolism*, 10(2), 486.
- [Ghiasi et al., 2017] Ghiasi, A., Hussain, O., Qian, Z. S., & Li, X. (2017). A mixed traffic capacity analysis and lane management model for connected automated vehicles: A markov chain method. *Transportation Research Part B: Methodological*, 106, 266–292.
- [Gipps, 1981] Gipps, P. G. (1981). A behavioural car-following model for computer simulation. *Transportation Research Part B: Methodological*, 15(2), 105–111.
- [Gipps, 1986] Gipps, P. G. (1986). A model for the structure of lane-changing decisions. *Transportation Research Part B: Methodological*, 20(5), 403–414.

- [Greenberg, 1959] Greenberg, H. (1959). An analysis of traffic flow. *Operations research*, 7(1), 79–85.
- [Guériau et al., 2016] Guériau, M., Billot, R., El Faouzi, N.-E., Monteil, J., Armetta, F., & Hassas, S. (2016). How to assess the benefits of connected vehicles? a simulation framework for the design of cooperative traffic management strategies. *Transportation research part C: emerging technologies*, 67, 266–279.
- [Guler et al., 2014] Guler, S. I., Menendez, M., & Meier, L. (2014). Using connected vehicle technology to improve the efficiency of intersections. *Transportation Research Part C: Emerging Technologies*, 46, 121–131.
- [Habibovic et al., 2014] Habibovic, A., Englund, C., & Wedlin, J. (2014). Current gaps, challenges and opportunities in the field of road vehicle automation.
- [Hagenzieker et al., 2020] Hagenzieker, M. P., van der Kint, S., Vissers, L., van Schagen, I. N. G., de Bruin, J., van Gent, P., & Commandeur, J. J. (2020). Interactions between cyclists and automated vehicles: Results of a photo experiment. *Journal of Transportation Safety & Security*, 12(1), 94–115.
- [Hall & Agyemang-Duah, 1991] Hall, F. L. & Agyemang-Duah, K. (1991). Freeway capacity drop and the definition of capacity. *Transportation research record*, (1320).
- [Hamdar et al., 2008] Hamdar, S. H., Treiber, M., Mahmassani, H. S., & Kesting, A. (2008). Modeling driver behavior as sequential risk-taking task. *Transportation research record*, 2088(1), 208–217.
- [Hancock et al., 2019] Hancock, P., Nourbakhsh, I., & Stewart, J. (2019). On the future of transportation in an era of automated and autonomous vehicles. *Proceedings of the National Academy of Sciences*, 116(16), 7684–7691.
- [Harper et al., 2016a] Harper, C. D., Hendrickson, C. T., Mangones, S., & Samaras, C. (2016a). Estimating potential increases in travel with autonomous vehicles for the non-driving, elderly and people with travel-restrictive medical conditions. *Transportation research part C: emerging technologies*, 72, 1–9.

- [Harper et al., 2016b] Harper, C. D., Hendrickson, C. T., & Samaras, C. (2016b). Cost and benefit estimates of partially-automated vehicle collision avoidance technologies. *Accident Analysis & Prevention*, 95, 104–115.
- [Hashim & Omar, 2017] Hashim, H. & Omar, M. (2017). Towards autonomous vehicle implementation: Issues and opportunities. *Journal of the Society of Automotive Engineers Malaysia*, 1(2), 111–123.
- [Heard et al., 2018] Heard, B. R., Taiebat, M., Xu, M., & Miller, S. A. (2018). Sustainability implications of connected and autonomous vehicles for the food supply chain. *Resources, conservation and recycling*, 128, 22–24.
- [Heredia et al., 2008] Heredia, G., Ollero, A., Bejar, M., & Mahtani, R. (2008). Sensor and actuator fault detection in small autonomous helicopters. *Mechatronics*, 18(2), 90–99.
- [Hietanen, 2014] Hietanen, S. (2014). Mobility as a service. *the new transport model*, (pp. 2–4).
- [Hietpas, 2017] Hietpas, J. (2017). Minnesota autonomous and connected vehicle perspectives. mtap.transportation.org/wp-content/uploads/sites/42/2017/09/Jay-AustinTXMinnesota.pdf.
- [Higgins, 2015] Higgins, S. (2015). Solid-state lidar: A new era of 3d scanning. spar3d.com/blogs/the-other-dimension/vol13no50-solid-state-lidar-a-new-era-of-3d-scanning.
- [Ho et al., 2009] Ho, M., Chan, P., & Rad, A. (2009). Lane change algorithm for autonomous vehicles via virtual curvature method. *Journal of Advanced Transportation*, 43(1), 47–70.
- [Hobert et al., 2015] Hobert, L., Festag, A., Llatser, I., Altomare, L., Visintainer, F., & Kovacs, A. (2015). Enhancements of v2x communication in support of cooperative autonomous driving. *IEEE communications magazine*, 53(12), 64–70.
- [Hollander & Liu, 2008] Hollander, Y. & Liu, R. (2008). The principles of calibrating traffic microsimulation models. *Transportation*, 35(3), 347–362.
- [Hou et al., 2013] Hou, Y., Edara, P., & Sun, C. (2013). Modeling mandatory lane changing using bayes classifier and decision trees. *IEEE Transactions on Intelligent Transportation Systems*, 15(2), 647–655.

- [Hu et al., 2012] Hu, J., Kong, L., Shu, W., & Wu, M.-Y. (2012). Scheduling of connected autonomous vehicles on highway lanes. *2012 IEEE Global Communications Conference (GLOBECOM)*, (pp. 5556–5561).
- [Huang et al., 2013] Huang, F., Liu, P., Yu, H., & Wang, W. (2013). Identifying if vissim simulation model and ssam provide reasonable estimates for field measured traffic conflicts at signalized intersections. *Accident Analysis & Prevention*, 50, 1014–1024.
- [Hummer et al., 2016] Hummer, J. E., Cunningham, C. M., Srinivasan, R., Warchol, S., Claros, B., Edara, P., & Sun, C. (2016). Safety evaluation of seven of the earliest diverging diamond interchanges installed in the united states. *Transportation research record*, 2583(1), 25–33.
- [Hwang & Park, 2005] Hwang, S. Y. & Park, C. H. (2005). Modeling of the gap acceptance behavior at a merging section of urban freeway. In *Proceedings of the Eastern Asia Society for Transportation Studies*, volume 5 (pp. e1656).: Tokyo: Eastern Asia Society for Transportation (EASTS).
- [ICBC, 2015] ICBC (2015). *Learn to Drive Smart: Your Guide to Driving Safely*.
- [iMOVE, 2020] iMOVE (2020). Smart mobility projects and trials across the world. imoveaustralia.com/smart-mobility-projects-trials-list.
- [Ioannou & Bose, 1999] Ioannou, P. & Bose, A. (1999). Automated vehicle control. In *Handbook of Transportation Science* (pp. 187–232). Springer.
- [Ioannou & Chien, 1993] Ioannou, P. A. & Chien, C.-C. (1993). Autonomous intelligent cruise control. *IEEE Transactions on Vehicular technology*, 42(4), 657–672.
- [Jayaraman et al., 2018] Jayaraman, S. K., Creech, C., Robert Jr, L. P., Tilbury, D. M., Yang, X. J., Pradhan, A. K., & Tsui, K. M. (2018). Trust in av: An uncertainty reduction model of av-pedestrian interactions. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 133–134).
- [Jia & Ngoduy, 2016] Jia, D. & Ngoduy, D. (2016). Platoon based cooperative driving model with consideration of realistic inter-vehicle communication. *Transportation Research Part C: Emerging Technologies*, 68, 245–264.

- [Jiang et al., 2001] Jiang, R., Wu, Q., & Zhu, Z. (2001). Full velocity difference model for a car-following theory. *Physical Review E*, 64(1), 017101.
- [Johnson, 2015] Johnson, B. (2015). The great horse manure crisis of 1894. historic-uk.com/HistoryUK/HistoryofBritain/Great-Horse-Manure-Crisis-of-1894.
- [Kamal et al., 2013] Kamal, M. A. S., Imura, J.-i., Ohata, A., Hayakawa, T., & Aihara, K. (2013). Coordination of automated vehicles at a traffic-lightless intersection. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)* (pp. 922–927).: IEEE.
- [Kaplan & Meier, 1958] Kaplan, E. L. & Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American statistical association*, 53(282), 457–481.
- [Katsaros et al., 2011] Katsaros, K., Kernchen, R., Dianati, M., Rieck, D., & Zinoviou, C. (2011). Application of vehicular communications for improving the efficiency of traffic in urban areas. *Wireless Communications and Mobile Computing*, 11(12), 1657–1667.
- [Kesting et al., 2007] Kesting, A., Treiber, M., & Helbing, D. (2007). General lane-changing model mobil for car-following models. *Transportation Research Record: Journal of the Transportation Research Board*, (1999), 86–94.
- [Kesting et al., 2010] Kesting, A., Treiber, M., & Helbing, D. (2010). Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1928), 4585–4605.
- [Khodayari et al., 2011] Khodayari, A., Kazemi, R., Ghaffari, A., & Braunschweig, R. (2011). Design of an improved fuzzy logic based model for prediction of car following behavior. In *2011 IEEE International Conference on Mechatronics* (pp. 200–205).: IEEE.
- [Kidando et al., 2018] Kidando, E., Moses, R., Ghorbanzadeh, M., & Ozguven, E. E. (2018). Traffic operation and safety analysis on an arterial highway: Implications for connected vehicle applications. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)* (pp. 2753–2758).: IEEE.

- [Kikuchi & Chakroborty, 1992] Kikuchi, S. & Chakroborty, P. (1992). Car-following model based on fuzzy inference system. *Transportation Research Record*, (pp. 82–82).
- [Kim et al., 2008] Kim, J.-T., Kim, J., & Chang, M. (2008). Lane-changing gap acceptance model for freeway merging in simulation. *Canadian Journal of Civil Engineering*, 35(3), 301–311.
- [Kim & Choi, 2013] Kim, S. & Choi, J. (2013). Safety analysis of roundabout designs based on geometric and speed characteristics. *KSCE Journal of Civil Engineering*, 17(6), 1446–1454.
- [Kita, 1999] Kita, H. (1999). A merging-giveway interaction model of cars in a merging section: a game theoretic analysis. *Transportation Research Part A: Policy and Practice*, 33(3-4), 305–312.
- [Kockelman et al., 2017] Kockelman, K., Boyles, S., Stone, P., Fagnant, D., Patel, R., Levin, M. W., Sharon, G., Simoni, M., Albert, M., Fritz, H., et al. (2017). *An assessment of autonomous vehicles: traffic impacts and infrastructure needs*. Technical report, University of Texas at Austin. Center for Transportation Research.
- [Kometani & Sasaki, 1959] Kometani, E. & Sasaki, T. (1959). Dynamic behavior of traffic with a nonlinear spacing-speed relationship. *Theory of Traffic Flow (Proc. of Sym. on TTF (GM))*, (pp. 105–119).
- [Krechmer et al., 2016] Krechmer, D., Blizzard, K., Cheung, M. G., Campbell, R., Bitner, J., Row, S., Alexiadis, V., Osborne, J., Jensen, M., Tudela, A., et al. (2016). *Connected Vehicle Impacts on Transportation Planning-Primer and Final Report*. Technical report, United States. Department of Transportation. Intelligent Transportation
- [Kumar & Parida, 2011] Kumar, P. & Parida, M. (2011). Vulnerable road users in multi modal transport system for delhi. *LTA*, 6, 38–47.
- [Lagström & Malmsten Lundgren, 2016] Lagström, T. & Malmsten Lundgren, V. (2016). Avip-autonomous vehicles’ interaction with pedestrians-an investigation of pedestrian-driver communication and development of a vehicle external interface. Master’s thesis.

- [Lane, 2007] Lane, D. M. (2007). Introduction to analysis of variance. *Introduction to Analysis of Variance*.
- [Latrech et al., 2018] Latrech, C., Chaibet, A., Boukhniher, M., & Glaser, S. (2018). Integrated longitudinal and lateral networked control system design for vehicle platooning. *Sensors*, 18(9), 3085.
- [Le Vine et al., 2015] Le Vine, S., Zolfaghari, A., & Polak, J. (2015). Autonomous cars: The tension between occupant experience and intersection capacity. *Transportation Research Part C: Emerging Technologies*, 52, 1–14.
- [Lee, 2006] Lee, G. (2006). Modeling gap acceptance at freeway merges.
- [Letter & Elefteriadou, 2017] Letter, C. & Elefteriadou, L. (2017). Efficient control of fully automated connected vehicles at freeway merge segments. *Transportation Research Part C: Emerging Technologies*, 80, 190–205.
- [Levin & Boyles, 2016] Levin, M. W. & Boyles, S. D. (2016). A multiclass cell transmission model for shared human and autonomous vehicle roads. *Transportation Research Part C: Emerging Technologies*, 62, 103–116.
- [Levin, 2018] Levin, S. (2018). 'uber should be shut down': friends of self-driving car crash victim seek justice. theguardian.com/technology/2018/mar/20/uber-self-driving-car-crash-death-arizona-elaine-herzberg.
- [LeVine, Steve, 2017] LeVine, Steve (2017). What it really costs to turn a car into a self-driving vehicle. qz.com/924212/what-it-really-costs-to-turn-a-car-into-a-self-driving-vehicle.
- [Levinson et al., 2011] Levinson, J., Askeland, J., Becker, J., Dolson, J., Held, D., Kammel, S., Kolter, J. Z., Langer, D., Pink, O., Pratt, V., et al. (2011). Towards fully autonomous driving: Systems and algorithms. In *Intelligent Vehicles Symposium (IV), 2011 IEEE* (pp. 163–168).: IEEE.
- [Li & Li, 2019] Li, L. & Li, X. (2019). Parsimonious trajectory design of connected automated traffic. *Transportation Research Part B: Methodological*, 119, 1–21.
- [Li et al., 2019] Li, Z., Khasawneh, F., Yin, X., Li, A., & Song, Z. (2019). A new microscopic traffic model using a spring-mass-damper-clutch system. *arXiv preprint arXiv:1903.04469*.

- [Lioris et al., 2016] Lioris, J., Pedarsani, R., Tascikaraoglu, F. Y., & Varaiya, P. (2016). Doubling throughput in urban roads by platooning. *IFAC-PapersOnLine*, 49(3), 49–54.
- [Litman, 2017] Litman, T. (2017). *Autonomous vehicle implementation predictions*. Victoria Transport Policy Institute Victoria, Canada.
- [Liu et al., 2018] Liu, H., Kan, X., Shladover, S. E., Lu, X.-Y., & Ferlis, R. E. (2018). Impact of cooperative adaptive cruise control on multilane freeway merge capacity. *Journal of Intelligent Transportation Systems*, 22(3), 263–275.
- [Liu et al., 2014] Liu, P., Kurt, A., et al. (2014). Trajectory prediction of a lane changing vehicle based on driver behavior estimation and classification. (pp. 942–947).
- [Lundgren et al., 2017] Lundgren, V. M., Habibovic, A., Andersson, J., Lagström, T., Nilsson, M., Sirkka, A., Fagerlön, J., Fredriksson, R., Edgren, C., Krupenia, S., et al. (2017). Will there be new communication needs when introducing automated vehicles to the urban context? In *Advances in human aspects of transportation* (pp. 485–497). Springer.
- [Ma et al., 2017] Ma, J., Li, X., Zhou, F., Hu, J., & Park, B. B. (2017). Parsimonious shooting heuristic for trajectory design of connected automated traffic part ii: computational issues and optimization. *Transportation Research Part B: Methodological*, 95, 421–441.
- [Maarafi, 2015] Maarafi, A. (2015). The impact of autonomous vehicles on freeway throughput.
- [Makarem et al., 2012] Makarem, L., Pham, M.-H., Dumont, A.-G., & Gillet, D. (2012). Microsimulation modeling of coordination of automated guided vehicles at intersections. *Transportation Research Record: Journal of the Transportation Research Board*, (2324), 119–124.
- [Manual, 2000] Manual, H. C. (2000). Highway capacity manual. *Washington, DC*, 2.
- [McDonald et al., 1997] McDonald, M., Wu, J., & Brackstone, M. (1997). Development of a fuzzy logic based microscopic motorway simulation model.

- In *Proceedings of Conference on Intelligent Transportation Systems* (pp. 82–87).: IEEE.
- [Mercedes-Benz, 2014] Mercedes-Benz (2014). Car-to-x communication. [mercedes-benz.com/en/innovation/connected/car-to-x-communication](https://www.mercedes-benz.com/en/innovation/connected/car-to-x-communication).
- [Milanes et al., 2012] Milanes, V., Villagra, J., Godoy, J., Simo, J., Pérez, J., & Onieva, E. (2012). An intelligent v2i-based traffic management system. *IEEE Transactions on Intelligent Transportation Systems*, 13(1), 49–58.
- [Minelli et al., 2015] Minelli, S., Izadpanah, P., & Razavi, S. (2015). Evaluation of connected vehicle impact on mobility and mode choice. *Journal of traffic and transportation engineering (English edition)*, 2(5), 301–312.
- [Mohan et al., 2009] Mohan, D., Tsimhoni, O., Sivak, M., & Flannagan, M. J. (2009). *Road safety in India: challenges and opportunities*. Technical report, University of Michigan, Ann Arbor, Transportation Research Institute.
- [Monteil et al., 2014] Monteil, J., Billot, R., Sau, J., & El Faouzi, N.-E. (2014). Linear and weakly nonlinear stability analyses of cooperative car-following models. *IEEE Transactions on Intelligent Transportation Systems*, 15(5), 2001–2013.
- [Mtoi & Moses, 2014] Mtoi, E. T. & Moses, R. (2014). Calibration and evaluation of link congestion functions: applying intrinsic sensitivity of link speed as a practical consideration to heterogeneous facility types within urban network. *Journal of Transportation Technologies*.
- [Nagel et al., 1998] Nagel, K., Wolf, D. E., Wagner, P., & Simon, P. (1998). Two-lane traffic rules for cellular automata: A systematic approach. *Physical Review E*, 58(2), 1425.
- [Naranjo et al., 2003] Naranjo, J. E., González, C., Reviejo, J., García, R., & De Pedro, T. (2003). Adaptive fuzzy control for inter-vehicle gap keeping. *IEEE Transactions on Intelligent Transportation Systems*, 4(3), 132–142.
- [Navya, 2019] Navya (2019). Navya’s autonom shuttles, first automated vehicles without steering wheel to operate on open road in japan. navya.tech/en/press/navyas-autonom-shuttles-first-automated-vehicles-without-steering-wheel-to-operate-on-open-road-in-japan.

- [Ni et al., 2013] Ni, Y., Ling, Z., & Li, K. (2013). A new evaluation method combining efficiency and safety: multimodal comprehensive level of service of signalized intersections. In *ICTIS 2013: Improving Multimodal Transportation Systems-Information, Safety, and Integration* (pp. 1447–1455).
- [Nie et al., 2016] Nie, J., Zhang, J., Ding, W., Wan, X., Chen, X., & Ran, B. (2016). Decentralized cooperative lane-changing decision-making for connected autonomous vehicles. *IEEE Access*, 4, 9413–9420.
- [Oh & Yeo, 2012] Oh, S. & Yeo, H. (2012). Microscopic analysis on the causal factors of capacity drop in highway merging sections.
- [Oketch & Carrick, 2005] Oketch, T. & Carrick, M. (2005). Calibration and validation of a micro-simulation model in network analysis. In *Proceedings of the 84th TRB Annual Meeting, Washington, DC*.
- [Olstam & Tapani, 2004] Olstam, J. J. & Tapani, A. (2004). *Comparison of Car-following models*, volume 960. Swedish National Road and Transport Research Institute Linköping, Sweden.
- [Organization, 2015] Organization, W. H. (2015). *Global status report on road safety 2015*. World Health Organization.
- [Pande et al., 2017] Pande, A., Chand, S., Saxena, N., Dixit, V., Loy, J., Wolshon, B., & Kent, J. D. (2017). A preliminary investigation of the relationships between historical crash and naturalistic driving. *Accident Analysis & Prevention*, 101, 107–116.
- [Papadoulis et al., 2019] Papadoulis, A., Quddus, M., & Imprialou, M. (2019). Evaluating the safety impact of connected and autonomous vehicles on motorways. *Accident Analysis & Prevention*, 124, 12–22.
- [Papaioannou, 2007] Papaioannou, P. (2007). Driver behaviour, dilemma zone and safety effects at urban signalised intersections in greece. *Accident Analysis & Prevention*, 39(1), 147–158.
- [Paris, 2017] Paris, M. (2017). Newcastle light rail nearly twice the price of canberra’s new tram system. newcastleherald.com.au/story/5128157/newcastle-light-rail-nearly-twice-the-price-of-canberras.

- [Patel et al., 2016] Patel, R., Levin, M. W., & Boyles, S. D. (2016). Effects of autonomous vehicle behavior on arterial and freeway networks. *Transportation Research Record*, 2561(1), 9–17.
- [Perrone, 2019] Perrone (2019). *Autonomous Shuttle Pilot*. Technical report, Perrone Robotics.
- [Persaud et al., 1998] Persaud, B., Yagar, S., & Brownlee, R. (1998). Exploration of the breakdown phenomenon in freeway traffic. *Transportation Research Record*, 1634(1), 64–69.
- [Pollatschek et al., 2002] Pollatschek, M. A., Polus, A., & Livneh, M. (2002). A decision model for gap acceptance and capacity at intersections. *Transportation Research Part B: Methodological*, 36(7), 649–663.
- [PTV Group, 2016] PTV Group (2016). *PTV Vissim 9 User Manual*, volume 9. PTV Group.
- [PwC, 2004] PwC (2004). Autofacts. *Executive perspectives. Managing risk in the automotive planning process*. London.
- [Qian et al., 2015] Qian, X., Gregoire, J., De La Fortelle, A., & Moutarde, F. (2015). Decentralized model predictive control for smooth coordination of automated vehicles at intersection. In *2015 European Control Conference (ECC)* (pp. 3452–3458).: IEEE.
- [Qin et al., 2011] Qin, X., Khan, G., Bill, A., & Noyce, D. A. (2011). Comprehensive safety evaluation of roundabouts in wisconsin. *Journal of Transportation Safety & Security*, 3(4), 289–303.
- [Raff et al., 1950] Raff, M. S. et al. (1950). A volume warrant for urban stop signs.
- [Rahman & Abdel-Aty, 2018] Rahman, M. S. & Abdel-Aty, M. (2018). Longitudinal safety evaluation of connected vehicles’ platooning on expressways. *Accident Analysis & Prevention*, 117, 381–391.
- [Rahman et al., 2019] Rahman, M. S., Abdel-Aty, M., Lee, J., & Rahman, M. H. (2019). Safety benefits of arterials’ crash risk under connected and automated vehicles. *Transportation Research Part C: Emerging Technologies*, 100, 354–371.

- [Rahman et al., 2018] Rahman, M. S., Abdel-Aty, M., Wang, L., & Lee, J. (2018). Understanding the highway safety benefits of different approaches of connected vehicles in reduced visibility conditions. *Transportation research record*, 2672(19), 91–101.
- [Rakha & Crowther, 2002] Rakha, H. & Crowther, B. (2002). Comparison of greenshields, pipes, and van aerde car-following and traffic stream models. *Transportation Research Record*, 1802(1), 248–262.
- [Reichert, 2019] Reichert, C. (2019). Ces 2019: Intel details autonomous vehicle trial in israel. zdnet.com/article/ces-2019-intel-details-autonomous-vehicle-trial-in-israel.
- [Retting, 2017] Retting, R. (2017). Pedestrian traffic fatalities by state. *Governors Highway Safety Association: Washington, DC, USA*.
- [RMS, 2013] RMS (2013). *Traffic Modelling Guidelines*. RMS.
- [RMS, 2016] RMS (2016). *Traffic Signal Design: Section 7 Phasing and Signal Group Display Sequence*, volume 1.3. Transport for New South Wales.
- [RMS, 2017] RMS (2017). *Motorway Design Guide*. RMS.
- [RMS, 2018] RMS (2018). Traffic volume viewer. rms.nsw.gov.au/about/corporate-publications/statistics/traffic-volumes/aadt-map.
- [Roads & Authority, 2002] Roads & Authority, T. (2002). Guide to traffic generating developments.
- [Sabra et al., 2010] Sabra, Z. A., Gettman, D., Henry, R. D., Nallamotheu, V., Sabra, W., et al. (2010). *Balancing safety and capacity in an adaptive signal control system—Phase 1*. Technical report, United States. Federal Highway Administration.
- [Saw et al., 2015] Saw, K., Katti, B., & Joshi, G. (2015). Literature review of traffic assignment: static and dynamic. *International Journal of Transportation Engineering*, 2(4), 339–347.
- [Schmidt, 2019] Schmidt, B. (2019). The driven. thedriven.io/2019/08/14/queensland-launches-most-advanced-autonomous-vehicle-trial-using-renault-zoe.

- [Sharon & Stone, 2017] Sharon, G. & Stone, P. (2017). A protocol for mixed autonomous and human-operated vehicles at intersections. *International Conference on Autonomous Agents and Multiagent Systems*, (pp. 151–167).
- [Shladover, 2018] Shladover, S. E. (2018). Connected and automated vehicle systems: Introduction and overview. *Journal of Intelligent Transportation Systems*, 22(3), 190–200.
- [Shoup, 2007] Shoup, D. (2007). Cruising for parking.
- [Spiess, 1990] Spiess, H. (1990). Conical volume-delay functions. *Transportation Science*, 24(2), 153–158.
- [Srivastava & Geroliminis, 2013] Srivastava, A. & Geroliminis, N. (2013). Empirical observations of capacity drop in freeway merges with ramp control and integration in a first-order model. *Transportation Research Part C: Emerging Technologies*, 30, 161–177.
- [Stamatiadis et al., 2013] Stamatiadis, N., Kirk, A., & Agarwal, N. (2013). An optimization tool for intersection design. In *2013 5th International Conference on Modeling, Simulation and Applied Optimization (ICMSAO)* (pp. 1–5).: IEEE.
- [Stevanovic et al., 2013] Stevanovic, A., Stevanovic, J., & Kergaye, C. (2013). Optimization of traffic signal timings based on surrogate measures of safety. *Transportation research part C: emerging technologies*, 32, 159–178.
- [Sun & Elefteriadou, 2012] Sun, D. & Elefteriadou, L. (2012). Lane-changing behavior on urban streets: An “in-vehicle” field experiment-based study. *Computer-Aided Civil and Infrastructure Engineering*, 27(7), 525–542.
- [Sun & Elefteriadou, 2011] Sun, D. J. & Elefteriadou, L. (2011). Lane-changing behavior on urban streets: A focus group-based study. *Applied ergonomics*, 42(5), 682–691.
- [Szymkowski, 2019] Szymkowski, S. (2019). China’s baidu rolls out self-driving cars for public trial. cnet.com/roadshow/news/china-baidu-self-driving-cars-public-trial.
- [Talebpour & Mahmassani, 2016] Talebpour, A. & Mahmassani, H. S. (2016). Influence of connected and autonomous vehicles on traffic flow stability and

- throughput. *Transportation Research Part C: Emerging Technologies*, 71, 143–163.
- [Talebpour et al., 2011] Talebpour, A., Mahmassani, H. S., & Hamdar, S. H. (2011). Multiregime sequential risk-taking model of car-following behavior: specification, calibration, and sensitivity analysis. *Transportation research record*, 2260(1), 60–66.
- [Tang et al., 2018] Tang, J., Liu, F., Zhang, W., Ke, R., & Zou, Y. (2018). Lane-changes prediction based on adaptive fuzzy neural network. *Expert Systems with Applications*, 91, 452–463.
- [Tettamanti et al., 2016] Tettamanti, T., Varga, I., & Szalay, Z. (2016). Impacts of autonomous cars from a traffic engineering perspective. *Periodica Polytechnica Transportation Engineering*, 44(4), 244–250.
- [Toledo et al., 2003] Toledo, T., Koutsopoulos, H., & Ben-Akiva, M. (2003). Modeling integrated lane-changing behavior. *Transportation Research Record: Journal of the Transportation Research Board*, (1857), 30–38.
- [Transport for New South Wales, 2017] Transport for New South Wales (2017). Crash and casualty statistics - nsw general view. road-safety.transport.nsw.gov.au/statistics/interactivecrashstats/nsw.html.
- [Transport for New South Wales, 2018] Transport for New South Wales (2018). Towards zero. towardszero.nsw.gov.au.
- [Transport Simulation Systems, 2014] Transport Simulation Systems (2014). *Aimsun 8 Dynamic Simulators Users' Manual*, volume 8.1. Transport Simulation Systems.
- [Treiber & Kesting, 2013] Treiber, M. & Kesting, A. (2013). Traffic flow dynamics. *Traffic Flow Dynamics: Data, Models and Simulation*, Springer-Verlag Berlin Heidelberg.
- [Turner & Roozenburg, 2006] Turner, S. & Roozenburg, A. (2006). Roundabout safety—influence of speed, visibility and design.
- [Van Aerde, 1995] Van Aerde, M. (1995). Single regime speed-flow-density relationship for congested and uncongested highways. In *74th Annual Meeting of the Transportation Research Board, Washington, DC*, volume 6.

- [Van Arem et al., 2003] Van Arem, B., Tampere, C. M., & Malone, K. (2003). Modelling traffic flows with intelligent cars and intelligent roads. *IEEE IV2003 Intelligent Vehicles Symposium. Proceedings (Cat. No. 03TH8683)*, (pp. 456–461).
- [Van Arem et al., 2006] Van Arem, B., Van Driel, C. J., & Visser, R. (2006). The impact of cooperative adaptive cruise control on traffic-flow characteristics. *IEEE Transactions on intelligent transportation systems*, 7(4), 429–436.
- [Van Brummelen et al., 2018] Van Brummelen, J., O’Brien, M., Gruyer, D., & Najjaran, H. (2018). Autonomous vehicle perception: The technology of today and tomorrow. *Transportation research part C: emerging technologies*, 89, 384–406.
- [Van Winsum, 1999] Van Winsum, W. (1999). The human element in car following models. *Transportation research part F: traffic psychology and behaviour*, 2(4), 207–211.
- [Velodyne Corporation, 2019] Velodyne Corporation (2019). Hdl-64e. velodynelidar.com/hdl-64e.html.
- [VicRoads, 2020] VicRoads (2020). Vicroads open data. vicroadsopendata-vicroadsmaps.opendata.arcgis.com/datasets/traffic-volume.
- [Victoria Government, 2019] Victoria Government (2019). Typical hourly traffic volume. data.vic.gov.au/data/dataset/typical-hourly-traffic-volume.
- [Vissers et al., 2017] Vissers, L. K., van der Schagen, S., & INLG van & Haggenzieker, M. (2017). Safe interaction between cyclists, pedestrians and automated vehicles: what do we know and what do we need to know?
- [Volvo, 2017] Volvo (2017). Autonomous driving. volvocars.com/en-kw/own/own-and-enjoy/autonomous-driving.
- [Wadud et al., 2016] Wadud, Z., MacKenzie, D., & Leiby, P. (2016). Help or hindrance? the travel, energy and carbon impacts of highly automated vehicles. *Transportation Research Part A: Policy and Practice*, 86, 1–18.
- [Wedel et al., 2009] Wedel, J. W., Schünemann, B., & Radusch, I. (2009). V2x-based traffic congestion recognition and avoidance. In *2009 10th Inter-*

- national Symposium on Pervasive Systems, Algorithms, and Networks* (pp. 637–641).: IEEE.
- [Wiedemann, 1974] Wiedemann, R. (1974). Simulation des strassenverkehrsflusses.
- [Williams, 2018] Williams, T.-A. (2018). Driverless cars in the uk 'in just three years': Government ploughs £25million into three autonomous vehicle trials for 2021. www.dailymail.co.uk/sciencetech/article-6416527/UK-autonomous-vehicle-trials-receive-25-million-grant.html.
- [Wu et al., 2018] Wu, J., Radwan, E., & Abou-Senna, H. (2018). Determination if vissim and ssam could estimate pedestrian-vehicle conflicts at signalized intersections. *Journal of Transportation Safety & Security*, 10(6), 572–585.
- [Wu, 2006] Wu, N. (2006). A new model for estimating critical gap and its distribution at unsignalized intersections based on the equilibrium of probabilities. In *Proceeding of the 5th international symposium on highway capacity and quality of service. Yokohama, Japan*.
- [Yang, 2017] Yang, H. (2017). Scenarios analysis of autonomous vehicles deployment with different market penetration rate.
- [Zhang et al., 2018] Zhang, J., Tang, T.-Q., & Yu, S.-W. (2018). An improved car-following model accounting for the preceding car’s taillight. *Physica A: Statistical Mechanics and its Applications*, 492, 1831–1837.
- [Zhou et al., 2017] Zhou, M., Qu, X., & Li, X. (2017). A recurrent neural network based microscopic car following model to predict traffic oscillation. *Transportation research part C: emerging technologies*, 84, 245–264.
- [Zhou et al., 2019] Zhou, M., Yu, Y., & Qu, X. (2019). Development of an efficient driving strategy for connected and automated vehicles at signalized intersections: A reinforcement learning approach. *IEEE Transactions on Intelligent Transportation Systems*, 21(1), 433–443.