

Programming Traffic: Modelling and Control of Microscopic Vehicle Manoeuvres

ZHAOHAN WANG



THE UNIVERSITY OF
SYDNEY

Supervisor: David Levinson
Associate Supervisor: Mohsen Ramezani

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

*Faculty of Engineering
The University of Sydney
Australia*

25 November 2025

Statement of Originality

I hereby affirm that the content presented in this thesis is entirely my own creation, to the best of my knowledge and belief. The thesis has not been submitted for any academic degree or evaluated for any other purpose previously. I confirm that all intellectual contributions in this work are solely mine, except where explicitly acknowledged. Furthermore, I certify that I have acknowledged all assistance received and sources consulted in preparing this thesis. I have read and understood the University of Sydney Student Plagiarism: Coursework Policy and Procedure, and I am aware that failure to comply with these guidelines can lead to proceedings against me for potential student misconduct under chapter 8 of the University of Sydney By-Law 1999 (as amended).

Author: Zhaohan Wang

Signature:

Date: Tuesday 25th November, 2025

Author Attribution Statement

This thesis encompasses content from both published and under-review journal articles, co-authored as indicated, and integrated with the permission of all co-authors. The papers included are listed as follows, in the order they appear within the thesis:

- (1) Wang, Z.* Ramezani, M., and Levinson, D. (2024) How mandatory are ‘Mandatory’ lane changes? An analytical and experimental study on the costs of missing freeway exits. *Transportation Research Part B: Methodological*, 186, 102994. In Chapter 3.
- (2) Wang, Z.* Ramezani, M., and Levinson, D. (2025). A game-theoretic two-dimensional driving model. *under review*. In Chapter 4.
- (3) Wang, Z.* Ramezani, M., and Levinson, D. (2025). Communication-free decentralised two-dimensional trajectory controller for autonomous vehicles on a lane-free road: A level-k game approach. *under review*. In Chapter 5.

My contributions to these papers encompassed methodology development, software implementation, data analysis, investigation, and the preparation of the original drafts.

Author: Zhaohan Wang

Signature:

Date: Tuesday 25th November, 2025

As the supervisor for the research conducted in this thesis, I affirm the accuracy of the above statements concerning the authorship.

Supervisor: David Levinson

Signature:

Date: Tuesday 25th November, 2025

As the associate supervisor for the research conducted in this thesis, I affirm the accuracy of the above statements concerning the authorship.

Associate supervisor: Mohsen Ramezani

Signature:

Date: Tuesday 25th November, 2025

Acknowledgements

My journey towards completing this PhD has not been easy, but it has certainly been a rewarding one, marked by transformation and growth. None of this would have been possible without the unwavering support of many individuals.

First and foremost, I offer my deepest gratitude to my supervisors, Professor David Levinson and Associate Professor Mohsen Ramezani, for their guidance, inspiration, and encouragement. One of the most valuable lessons I learned from David is to question assumptions rather than simply accept them. He helped me realise the importance of thinking critically instead of following others blindly. Mohsen, who once resembled James Dean, gave me all the support he promised at the beginning of my journey. His wisdom, gentleness, and decency have set a lifelong example for me. Despite the inevitable obstacles along the way, I feel fortunate to have shared those moments with David and Mohsen. I will always remember my nervousness when meeting them for the first time, the joy when our paper was accepted, and the heated discussions that brought our ideas to life. And after years of David's playful reminders about whether I had finished my thesis, I can finally, and proudly, say that I have.

I am also sincerely thankful to Dr. Emily Moylan and Dr. Andrés Fielbaum for their invaluable guidance and advice throughout these three and a half years. Their expertise and feedback have greatly enriched both the quality of my research and my understanding of transportation. I am grateful to my colleagues and friends: Yue Yang, Yang Gao, Changle Song, Haotian Wang, Tim Xian, Alireza Soltani, Zhuopeng Xie, Ruihao Zeng, Sylvia Zhao, Peace, and many others. It has been a privilege to know them and to share both laughter and tears along the way.

Last, and certainly not least, I extend my deepest gratitude to my family. My parents have been a constant source of strength and an unwavering presence in my life. Simply knowing they are behind me is empowering. And to my fiancée, Christine Sun, you are the light and warmth that carried me through every peak and trough.

Abstract

This dissertation explores microscopic decision by human drivers and builds control frameworks that optimise microscopic actions of autonomous vehicles. Existing studies often treat mandatory lane-changes as absolute, neglecting the possible option of not performing them. As such, we explore the incentive behind the manoeuvre by analysing and modelling the cost for not performing lane-changes for exiting freeways. The developed mathematical model is deterministic and can estimate the exit-missing cost with simple network-level variables. Both simulation and empirical results indicate that exit missing costs are not substantial.

Driving is a two-dimensional task, yet studies have often neglected the interdependence between the two dimensions. This study integrates longitudinal car-following with lateral lane-keeping and discretionary lane-changing behaviours to create a realistic model that replicates human-driven trajectories. Recognising the hierarchical nature of human decision-making, we incorporate a game-theoretic reasoning structure and validate the model using a naturalistic trajectory dataset. Model calibration is conducted heterogeneously through a Bayesian framework. The model is able to replicate macroscopic traffic patterns with high accuracies and can generate smooth human-like trajectories.

Building on the two-dimensional model, we design an optimal controller for autonomous vehicles operating in a lane-free freeway environment. Vehicles are assumed to be communication-free, with each making decisions independently. A game-theoretic formulation ensures that vehicles account for the potential actions of others. The full framework is evaluated at both the microscopic and macroscopic scales. Experiments demonstrate how reasoning depth and vehicle-size heterogeneity influence overall traffic efficiency. The proposed lane-free controller is also compared against a similarly defined lane-based controller.

Contents

Statement of Originality	ii
Author Attribution Statement	iii
Acknowledgements	v
Abstract	vi
List of Figures	xi
List of Tables	xiv
Nomenclature	xvi
Chapter 1 Introduction	1
1.1 Background	1
1.2 Research Objectives	3
1.3 Thesis Overview	4
Chapter 2 Literature review	8
2.1 Introduction	8
2.2 Microscopic Modelling of Human Driving	8
2.2.1 Car-following Models	9
2.2.2 Lane-changing Models	11
2.2.3 Coupled Car-following and Lane-changing Models	13
2.2.4 Research Gaps	15
2.3 Autonomous Vehicle Optimal Control	18
2.3.1 Lane-based AV Control	18
2.3.2 Lane-free AV Control	19
2.3.3 Research Gaps	21

2.4	Conclusions	23
Chapter 3 Costs of Missing Freeway Exits		24
3.1	Introduction	24
3.2	Analytical Costs for Missing Freeway Exits	26
3.2.1	Mean of Exit-Missing Costs	26
3.2.2	Standard Deviation of Exit-Missing Costs	33
3.3	Simulation Experiments of Exit-Missing Costs in the Top 50 Metropolitan Areas	35
3.3.1	Extraction of Exit-Missing Costs	36
3.3.2	Goodness of Fit with Analytical Costs	37
3.4	Empirical Evidence for Exit Missing Costs	43
3.4.1	Method for Empirical Analysis	44
3.4.1.1	Datasets and Pre-processing	44
3.4.1.2	Estimation of Exit-Missing Costs	47
3.4.2	Empirical Results	47
3.4.2.1	Costs of Missing Exit	48
3.4.2.2	Analysis on Relative Proportions	51
3.4.2.3	Analysis on Commuting Effect	54
3.4.2.4	Analysis on Peak Hour Effect	56
3.5	Conclusions and Limitations	61
Chapter 4 A Game-theoretic Two-dimensional Driving Model		67
4.1	Introduction	67
4.2	Methodology	68
4.2.1	Receding Horizon Cost Optimisation	69
	Speed cost	71
	Spacing cost	72
	Safety cost	73
	Steering and acceleration costs	76
	Lane-centring cost	77
4.2.2	Level-k Game Theory	77

4.2.3	Numerical Solutions	80
4.3	Calibration	83
4.3.1	Data	83
4.3.2	Calibration Procedure	83
4.3.3	Calibration Results	87
4.4	Validation	90
4.4.1	Validation on the Calibrated Track (Track 50)	92
4.4.2	Validation on Non-calibrated Tracks	96
	Track 01	96
	Track 60	97
4.4.3	Discussion	99
4.5	Conclusion	101
Chapter 5 Decentralised and Communication-free Control of Autonomous Vehicles on Lane-free Freeways		105
5.1	Introduction	105
5.2	Methodology	108
5.2.1	Optimal control problem	108
5.2.2	Continuation/GMRES Method	111
5.2.3	Level-k Game Theory	117
5.2.4	Parameter tuning	118
5.3	Simulation Results	120
5.3.1	Microscopic Behaviour	121
5.3.2	Macroscopic Behaviour	125
5.3.3	Comparing $K = 1$, $K = 2$, and Mixture	131
5.3.4	Comparing Levels of Heterogeneity in Vehicle Size	132
5.3.5	Comparing Lane-free with Lane-based	133
5.4	Conclusion	136
Chapter 6 Conclusion		138
6.1	Summary of Main Findings	138

6.2	Future Outlook	140
Appendix A	Appendix for Chapter 3	142
A1	Exits Analysed	142
A2	Route Deviation and Circuitry	143
A3	Categorical Analysis of Final Exits	146
Appendix B	Appendix for Chapter 4	152
B1	Fitting of Truncated Normal Distribution	152
B2	Correlation of Calibrated Parameters	153
B3	Calibrated Distributions of IDM+MOBIL Model	153
B4	Sensitivity of Parameters	154
Appendix C	Appendix for Chapter 5	156
C1	Visualisation and Animations of Simulated Scenarios	156
C1.1	Scenario 1	156
C1.2	Scenario 2	157
C1.3	Scenario 3	157
C1.4	Scenario 4	158
Bibliography		159

List of Figures

1.1	Thesis outline	5
3.1	Composition of exit-missing costs	27
3.2	Contributions of exits to the surrounding destinations	28
3.3	Scatter plots comparing the actual and predicted mean of costs	39
3.4	Scatter plots for the actual and predicted mean of costs for to-work and to-home trips	40
3.5	Scatter plots for the actual and predicted standard deviation of costs	41
3.6	Distributions of the time cost of missing exits	49
3.7	Distributions of the distance cost of missing exits	52
3.8	Distributions of the percentage time cost of missing exits	53
3.9	Distributions of the percentage distance cost of missing exits	55
3.10	Distributions for the time cost of missing exits in commuting and non-commuting trips	58
3.11	Distributions for the distance cost of missing exits in commuting and non-commuting trips	58
3.12	Distributions for the time cost of missing exits in peak hour and off-peak hour trips	60
3.13	Distributions for the distance cost of missing exits in peak hour and off-peak hour trips	60
4.1	Bicycle model for ego vehicle on lane-based road	70
4.2	Definition of safe velocity and sight distance	72
4.3	Two-dimensional TTC components	76
4.4	Lane-centring cost	77
4.5	Game opponent identification by an ego vehicle	78
4.6	Fréchet distance over calibration iterations	88
4.7	Calibrated distributions of parameters for 2DDM	89

4.8	Example trajectories from Track 50	91
4.9	Example space-time diagram from Track 50	92
4.10	Speed distributions from Track 50	95
4.11	Space-time velocity heatmaps from Track 50	95
4.12	Speed distributions from Track 01	98
4.13	Space-time velocity heatmaps from Track 01	98
4.14	Speed distributions from Track 60	100
4.15	Space-time velocity heatmaps from Track 60	100
5.1	Bicycle model for ego vehicle on lane-free road	110
5.2	A pair of ego-preceding vehicles	112
5.3	Initial vehicle configurations for scenario 1	122
5.4	States and control inputs for scenario 1	122
5.5	Initial vehicle configurations for scenario 2	124
5.6	States and control inputs for scenario 2	124
5.7	Initial vehicle configurations for scenario 3	126
5.8	States and control inputs for scenario 3	126
5.9	States and control inputs for scenario 4	127
5.10	Fundamental diagrams for AVs controlled at $K = 2$ and $K = 1$	129
5.11	Distributions of the travel times for AVs controlled at $K = 2$ and $K = 1$	131
5.12	Distributions of the travel times for a mixture of AVs controlled at $K = 2$ and $K = 1$	132
5.13	Distributions of the travel times for AVs at different size heterogeneity	133
5.14	Distributions of the travel times for AVs in lane-free and lane-based environments	134
5.15	Distributions of the travel times for AVs in lane-free and lane-based environments with 30% size heterogeneity	135
5.16	Relationship between vehicle size heterogeneity and traffic efficiency	135
A.1	Spatial location of all exits in the datasets	142
A.2	Original and alternative routes of trip 497 from sample GPS1019	143
A.3	Route deviation distributions between original and alternative routes	144
A.4	Circuitry distributions for original and alternative routes	145

A.5	Final exits of to-work trips for sample GPS2038	148
A.6	Decision tree of exit structures for commuting trips	151
B.1	Correlation matrices for calibrated parameters	153
B.2	Calibrated distributions of parameters for IDM+MOBIL	154
C.1	Animation of lane-free scenario 1	156
C.2	Animation of lane-free scenario 2	157
C.3	Animation of lane-free scenario 3	157
C.4	Animation of lane-free scenario 4	158

List of Tables

2.1	Calibration and validation of coupled car-following and lane-changing models of human driving	15
2.2	Summary of car-following, lane-changing, and coupled models	16
2.3	Summary of lane-based and lane-free AV controllers	22
3.1	Partial derivatives for all random variables in the cost functions.	34
3.2	Summary of the data pre-processing steps. VMT and Otrek together contribute to the I-35W dataset. For each dataset, both the numbers of trips and points are showcased, as well as the original and final data size. The Otrek data do not contain speed attributes; thus, the speed filter is omitted.	46
3.3	Number of exits extracted from datasets. VMT and Otrek together contribute to the I-35W dataset.	48
3.4	Summary of the time cost of missing exits. (** $p \leq 0.01$, * $p \leq 0.05$)	49
3.5	Summary of the distance cost of missing exits. (** $p \leq 0.01$, * $p \leq 0.05$)	52
3.6	Summary of the time cost of missing exits in terms of the proportion of the total trip. (** $p \leq 0.01$, * $p \leq 0.05$)	53
3.7	Summary of the distance cost of missing exits in terms of the proportion of the total trip. (** $p \leq 0.01$, * $p \leq 0.05$)	55
3.8	Summary table for the time cost of missing exits for commuting and non-commuting trips. (** $p \leq 0.01$, * $p \leq 0.05$)	57
3.9	Summary table for the distance cost of missing exits for commuting and non-commuting trips. (** $p \leq 0.01$, * $p \leq 0.05$)	57
3.10	Summary table for the time cost of missing exits for peak hour and off-peak hour trips. (** $p \leq 0.01$, * $p \leq 0.05$)	59

3.11	Summary table for the distance cost of missing exits for peak hour and off-peak hour trips. (** $p \leq 0.01$, * $p \leq 0.05$)	59
4.1	Parameters of the two-dimensional driving model	84
4.2	Statistics of the truncated normal fits to the calibrated parameters	89
4.3	Validation results for Track 50. LC denotes lane-change.	93
4.4	Validation results for Track 01. LC denotes lane-change.	96
4.5	Validation results for Track 60. LC denotes lane-change.	97
5.1	Fixed parameters for the controlled vehicles	120
A.1	The correlation between measurements of missing exits. The factors are: the extra travel time cost, the extra travel distance cost, the circuitry difference, and the routes deviation.	146
B.1	Statistics of the truncated normal fits to the calibrated parameters	154
B.2	Sensitivity analysis of calibrated parameters using percentage error of lane-change frequency. The analysis is conducted using the $K = 1$ model and Track 50 from the highD dataset.	155

Nomenclature

$\mathbb{1}$	The indicator function.
\mathbf{a}_0	Lower limits of the prior distributions in MB-SMC.
$a(t)$	Acceleration of the vehicle at time t .
\mathbf{b}_0	Upper limits of the prior distributions in MB-SMC.
B	A Borel set of the underlying point process Ψ .
C	The network circuitry.
ΔD	Distance between centroids of two vehicles.
Δd	Distance between projections of two vehicles.
$\Delta D_{e,j}$	Distance between the ego vehicle and an opponent vehicle, projected along the relative velocity vector.
$d(p, q)$	Euclidean distance between points p and q .
$D_{c,F}$	The distance cost of missing exits on freeway network.
$D_{c,S}$	The distance cost of missing exits on surface street network.
D_c	The distance cost of missing exits.
d_f	The intermediate gap length between freeway exits.
$D_F(P, Q)$	Discrete Fréchet distance between trajectory sequences P and Q .
D_s	Driver sight distance.
D_W	Wasserstein distance between two distributions.
$d_{e,j}$	Distance-to-collision between the ego vehicle and an opponent, projected along the relative velocity vector.
\mathbb{E}	The expectation of a random variable.
$\mathbf{f}(\cdot)$	The coupled differential equations describing vehicle dynamics.
f/F	Probability/Cumulative distribution function of a random variable.
G	A multivariate function representing either the time or the distance cost function.
$g(\cdot)$	Running cost function.

$g_i(\cdot)$	Cost component. Components include: speed, spacing, acceleration, steering, and centring.
$\mathbf{h}(\cdot)$	Collision-avoiding inequality constraint.
$\mathcal{H}(\cdot)$	Hamiltonian.
\mathbb{J}	Set of all vehicles perceived by the ego vehicle.
\mathbb{J}^c	Set of all opponent vehicles identified by the ego vehicle.
K	Level of predefined cognitive hierarchy of an ego vehicle.
k	Level of current cognitive hierarchy of a vehicle.
\bar{l}_{sample}	Average length of vehicles that passed a loop detector within a sampling period.
ℓ	The Lebesgue measure of subset B .
$\mathcal{L}(\cdot)$	Cost function.
l	Length of the vehicle.
l_B	Vehicle wheelbase.
l_f	Horizontal distance between the front wheel centre and the centre of mass of the vehicle.
l_r	Horizontal distance between the rear wheel centre and the centre of mass of the vehicle.
N	Number of intervals over the prediction horizon.
n_c	Number of intervals over the optimisation horizon.
$N_{\text{iteration}}$	Number of iterations till termination in MB-SMC.
N_{particle}	Number of particles to be generated in MB-SMC.
n_{sample}	Number of vehicles passing over a loop detector during a sampling period.
\mathbb{P}	The probability of an event occurring.
Q	Traffic flow calculated from a loop detector.
\mathbb{R}	Set of real numbers.
R_n	The random variable representing the distance between a randomly chosen point and its n^{th} nearest neighbour in the point process.
r_n	The distance between a particular point and its n^{th} nearest neighbour in a given point process. It can be thought of as a particular realisation of R_n .
$s(t)$	Current spacing between the ego car and the preceding car.
$s^*(t)$	Spacing threshold between the ego car and the preceding car.

s_0	Minimum standstill distance.
ΔT	Length of intervals over the time axis.
Δt	Length of intervals over the prediction horizon.
T	Length of prediction horizon.
t_0	Time index at the initial time step.
T_c	The time cost of missing exits.
T_{occupied}	Time that a loop detector is occupied within a sampling period.
t_r	Reaction time of the driver to an event.
t_{safe}	Safe time headway of a vehicle.
T_{sample}	Sampling period for computing the traffic states using a loop detector.
TTC	Time-to-collision.
$\tilde{\mathbf{u}}(t)$	Vector of dummy variables in the admissible control constraint.
$\mathbf{U}(t)$	Vector of optimal inputs and multipliers.
$\mathbf{u}(t)$	Vehicle control actions at time t .
\mathcal{U}	The admissible action space defined by a_{\min} , a_{\max} , δ_{\min} , and δ_{\max} .
$V(t)$	Vector of necessary optimality conditions.
$v(t)$	Velocity of the vehicle at time t .
v_f	The travel velocity on freeway network.
v_{safe}	The maximum safe speed at current road condition.
v_s	The travel velocity on surface street network.
$v_{e,j}$	Relative velocity between the ego vehicle and an opponent.
w	Width of the vehicle.
w_{proj}	Projected length of an ellipse.
$x(t)$	Position of the vehicle along x-axis at time t .
$y(t)$	Position of the vehicle along y-axis at time t .
$\mathbf{Z}(t)$	The extended state vector at time t .
$\mathbf{z}(t)$	The state vector at time t .
z_a	z-score of the lower truncation abscissas x_a .
z_b	z-score of the upper truncation abscissas x_b .
α_i	Weighting of a cost component.
$\beta(t)$	Slip angle of the vehicle at time t .
Γ	The gamma function defined as $\Gamma(z) = \int_0^{\infty} t^{z-1} e^{-t} dt$.

$\delta(t)$	Front wheel steering angle of the vehicle at time t .
ϵ	Semi-smooth Fischer-Burmeister constant.
ϵ_t	Finite difference step size.
$\zeta(t)$	Lane index of the vehicle at time t .
$\eta(t)$	Lagrangian multiplier for inequality constraints.
θ	Angle of slope that the line of projection makes with the horizontal axis.
ι	The simple counting measure of B , defined as $\iota(B) = \sum_{x_i \in \Psi} \mathbb{1}(x_i \in B)$.
Λ	The intensity measure of points in B defined as $\Lambda(B) = \mathbb{E}[\iota(B)]$.
$\lambda(t)$	Costate variable.
$\bar{\mu}$	Mean of a normal distribution.
$\mu(t)$	Lagrangian multiplier for equality constraints.
ξ	The intensity of the point process Ψ .
ρ	Traffic density calculated from a loop detector.
ρ_P	Sequence of indices traversing P .
$\bar{\sigma}$	Standard deviation of a normal distribution.
σ_{pert}	Standard deviations of the perturbation distributions in MB-SMC.
$\Delta\tau$	Sampling interval.
τ_Q	Sequence of indices traversing Q .
v	Stabilising constant.
$\Phi(\cdot)$	Cumulative distribution function of a general normal distribution.
$\phi(\cdot)$	Probability density function of a general normal distribution.
$\Xi(\cdot)$	The semi-smooth Fischer-Burmeister function.
Ψ	A point process in \mathbb{R}^2 .
$\psi(t)$	Heading direction of the vehicle at time t .
Ω	A set of y-positions of all lane markings.
ω	Maximum number of lanes.
Ω^+	The upper boundary of the road.
Ω^-	The lower boundary of the road.

Introduction

1.1 Background

Understanding how individual vehicles behave is important for improving traffic flow, safety, and the overall efficiency of road networks. Small adjustments by drivers, such as changing lanes or altering speed, can produce wider effects in traffic, contributing to congestion patterns or instability ([Ahn and Cassidy, 2007](#), [Gao and Levinson, 2023](#)). Human driving decisions are shaped by experience, perception, and individual judgement, but are also affected by the presence and movement of other vehicles. These choices are often inconsistent and difficult to predict.

Driver decisions are rarely made in isolation. A driver's choice to change lanes often depends on their current car-following conditions. Similarly, the decision to slow down or speed up might depend on whether a lane change is planned. Because of this, longitudinal and lateral movements are closely linked. This makes it important to model these behaviours together, rather than separately. At the same time, different drivers may respond differently to similar situations depending on their preferences. This adds further variability to traffic dynamics. Driving is also interactive by nature. Drivers constantly observe their surroundings and form expectations about how others will behave. Based on these anticipations, they make decisions that can reflect both cooperation and competition.

To manage both the longitudinal and lateral aspects of driving, drivers must weigh multiple objectives at once, including time efficiency, safety, comfort, and compliance with traffic rules. These trade-offs are especially apparent in decisions that involve lane changes, where the

need to act must be balanced against the risks and difficulty of the manoeuvre. For example, when preparing to exit a freeway, a lane change may seem necessary. But even then, drivers might delay or avoid the change if traffic conditions are not suitable. These situations suggest that lane changes are not always absolute. Drivers make judgement calls based on what they think is possible and worthwhile.

With the rise of autonomous vehicles (AVs), there is an opportunity to improve traffic conditions using systems that respond more consistently and safely than human drivers. Research has explored both connected and unconnected AV systems, as well as centralised and decentralised control strategies. Among these, communication-free decentralised approaches show particular promise, especially in the near to medium term when full AV penetration is unlikely. In these systems, each AV makes decisions independently, using onboard sensors to perceive its surroundings and predict the actions of nearby vehicles. Although decentralised methods may not achieve the full system-level optimality of centralised approaches, they offer greater robustness, scalability, and practicality. Importantly, even without communication, AVs can still anticipate and adapt to others' behaviour through game-theoretic reasoning.

Current studies have explored AV decision-making in both traditional lane-based environments and in newer lane-free settings that relax lane-keeping constraints. Lanes have long improved the safety and efficiency of human driving by reducing the driving task to mainly longitudinal control. In contrast, lateral manoeuvres such as lane-changes are more complex and risky, requiring simultaneous control in both directions. AVs, however, benefit from advanced sensors and fast response times, allowing them to operate effectively in more complex two-dimensional spaces. This has led to growing interest in lane-free traffic systems, where AVs are no longer confined to predefined lanes but instead move flexibly and adaptively in response to their surroundings ([Papageorgiou et al., 2021](#), [Troullinos et al., 2021](#)). In such settings, two-dimensional planning is essential for generating smooth and physically realistic trajectories that account for both directions of motion. Like human drivers, AVs must balance multiple objectives such as efficiency, safety, and comfort when making decisions. Furthermore, in communication-free settings, AVs should anticipate the behaviour of surrounding vehicles in order to interact safely and cooperatively.

These challenges can be addressed using methods such as control theory and game theory. Control theory provides a structured way to plan actions over time while satisfying physical and safety constraints. In autonomous driving, it is commonly used to compute trajectories by optimising a cost function that reflects driving goals. These methods have been applied in both lane-based and lane-free environments (Karafyllis et al., 2022b, Levy and Haddad, 2021, Wang et al., 2015, Zhang et al., 2013). Beyond AVs, the same framework can also be used to describe human driving in a physics-informed way. Like autonomous systems, human drivers often make decisions by weighing competing factors, even if those judgements are less precise. By adjusting the weights assigned to different cost components, variation in driver preferences and behaviours can be represented within this framework.

While control theory is typically used to optimise the behaviour of a single agent, game theory focuses on situations involving multiple interacting agents. In driving, this is important because drivers often anticipate how others will act and adjust their own behaviour in response. Game theory offers tools to formalise these kinds of interactions. Originally developed in economics and psychology to understand strategic decision-making, game theory assumes that each agent chooses actions that maximise their own payoff, given the possible strategies of others (von Neumann et al., 1944). In transportation research, game-theoretic models have been developed to study lane-changing behaviours (Liu et al., 2007, Talebpour et al., 2015). Both classical approaches, such as Nash (Nash, 1951) and Stackelberg (Stackelberg, 2010) equilibria, as well as behavioural extensions like level-k reasoning (Camerer et al., 2004) have shown promise in capturing interactive behaviour in traffic settings.

1.2 Research Objectives

This thesis aims to improve the modelling of microscopic driving behaviour and to develop decision-making frameworks for AV control. The thesis examines mandatory lane changes in tactical routing scenarios such as freeway exits. A mathematical model is developed to estimate the costs of missing an intended exit, effectively quantifying the consequences of not performing a lane change that is typically considered mandatory. The model is formulated

using network-level variables and is validated through both simulation and empirical GPS data.

The two-dimensional driving model is proposed to capture human decision-making at the trajectory level. The model assumes that drivers optimise a cost function over a short prediction horizon. Vehicle interactions are modelled using level-k game theory, which accounts for anticipatory behaviour in multi-agent settings. To reflect the diversity in driving styles, driver heterogeneity is captured through a Bayesian calibration process using a naturalistic trajectory dataset.

Building on the two-dimensional model, the thesis proposes an optimal control framework for AVs operating in a lane-free freeway environment. AVs are assumed to function without communication, relying only on onboard sensing and decentralised decision-making. The control approach integrates receding horizon optimisation with level-k game theory, allowing each vehicle to anticipate others' actions and compute safe, real-time trajectories. Beyond individual vehicle performance, the framework is tested at the macroscopic level by simulating large fleets under various traffic conditions and levels of vehicle heterogeneity. Furthermore, the average vehicle efficiency is evaluated and compared against a similarly defined lane-based control framework.

1.3 Thesis Overview

The remaining chapters of this thesis cover existing studies on microscopic driving behaviour and control, followed by three core contributions: a game-theoretic two-dimensional model of human driving, an analysis of the cost of missing intended freeway exits, and an optimal control framework for AVs operating in lane-free environments. The overall structure of the thesis, which is organised into six chapters, is illustrated in Figure 1.1.

Chapter 2 reviews existing studies on microscopic modelling of human driving and optimal control of autonomous vehicles. It covers both decoupled models, which separately address car-following and lane-changing behaviours, and coupled models that integrate the two. The

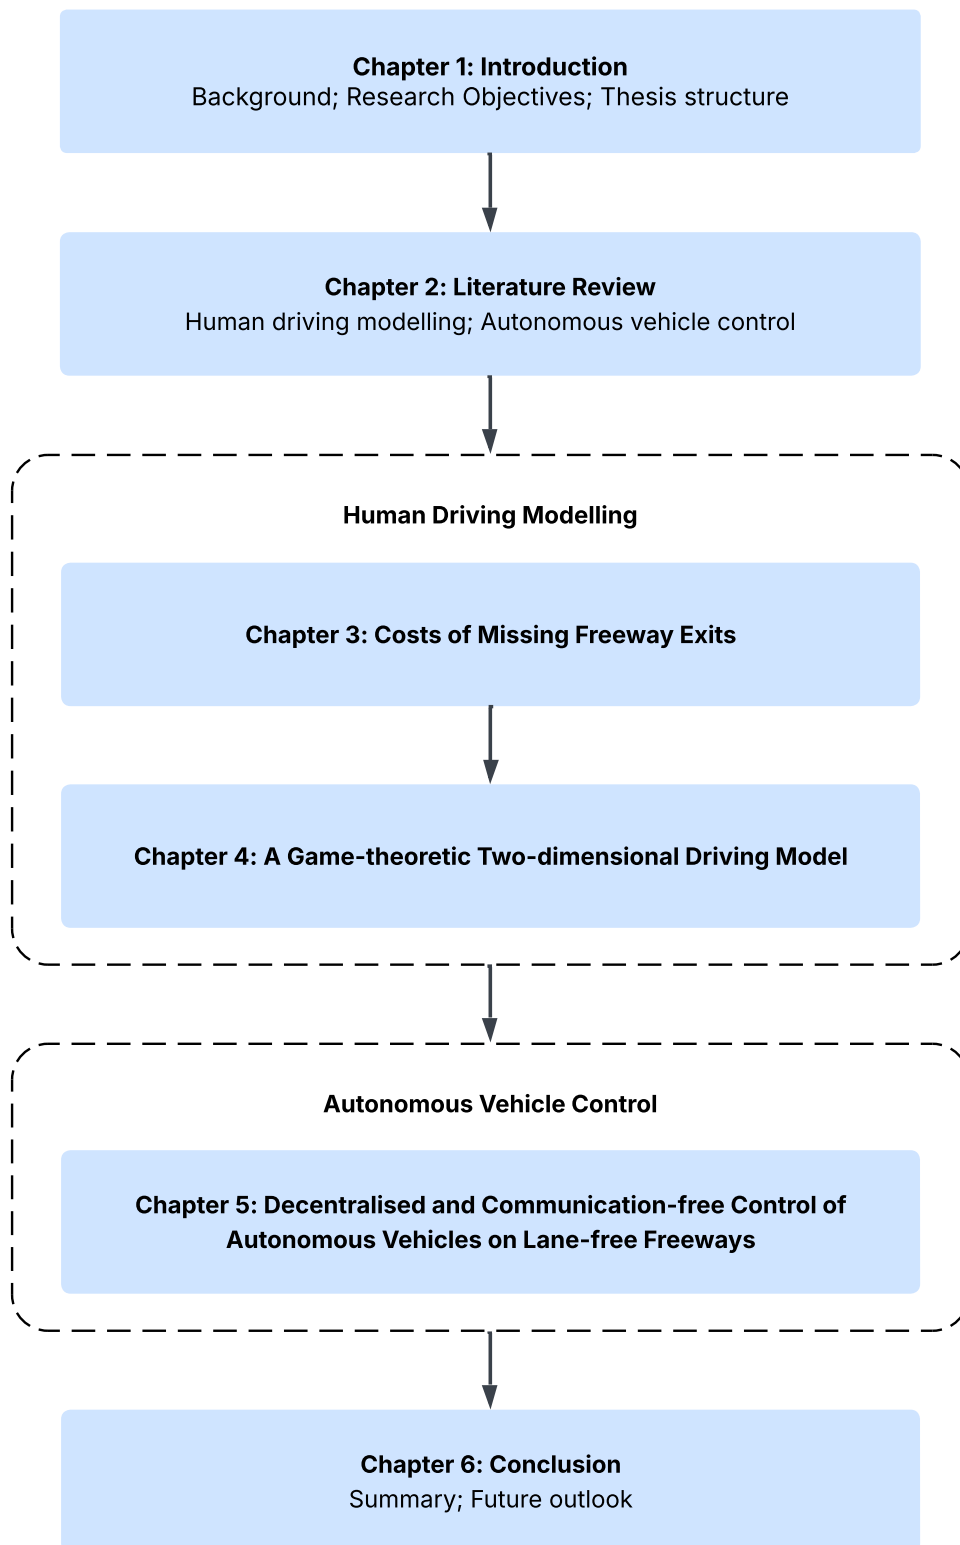


FIGURE 1.1. Thesis outline

chapter also discusses trajectory planning approaches for AVs in lane-based and lane-free environments. Based on these reviews, key research gaps are identified and potential directions for addressing them are outlined.

Chapter 3 examines mandatory lane changes in the context of tactical routing, specifically freeway exits. Rather than treating these lane changes as absolute, the chapter investigates the costs associated with failing to perform them, thereby shedding light on the trade-offs involved in mandatory lane-changing decisions. An analytical model is developed to estimate exit-missing costs based on network-level variables, and its outputs are validated using simulation results from the top 50 metropolitan areas in the US. An empirical analysis is also conducted using GPS data from the Minneapolis–St. Paul region. The results suggest that while there are measurable time and distance costs to missing exits, their magnitudes may be lower than often assumed.

Chapter 4 introduces a microscopic vehicle model that combines longitudinal and lateral decisions within a unified objective function. This two-dimensional driving model applies receding horizon optimisation to iteratively predict future movements. Strategic interactions between drivers are captured using level-k game theory, which reflects the hierarchical reasoning structure observed in human behaviour. Driver heterogeneity is incorporated through Bayesian calibration using the highD dataset, allowing the model to estimate distributions over individual cost preferences. The calibrated model successfully reproduces macroscopic traffic patterns with low error and demonstrates strong transferability when validated on data from different times and locations.

Chapter 5 proposes a trajectory control framework for AVs operating in a lane-free freeway setting. In this communication-free and decentralised approach, each vehicle independently determines its acceleration and steering actions using onboard sensors. The control problem is formulated as a nonlinear optimal control problem solved via the continuation/GMRES method. Strategic interactions are modelled using level-k game theory. The controller is evaluated through both microscopic and macroscopic analyses. At the microscopic level, it produces smooth, safe trajectories that respect all constraints. At the macroscopic level, large-scale simulations involving traffic bottlenecks show that the system achieves higher flow

efficiency than conventional human-driven traffic. Additional analyses examine the influence of reasoning depth and vehicle heterogeneity, and the controller's performance is compared with a similarly designed lane-based alternative.

Finally, Chapter 6 concludes the thesis by summarising the main findings from each chapter. It highlights the key contributions of the research and discusses potential future directions for applying and extending the proposed models and control frameworks.

Literature review

2.1 Introduction

Microscopic traffic models describe the behaviour of individual vehicles and their interactions within the traffic stream. These models form the foundation for traffic simulation, safety analysis, and the development of vehicle control systems. Two key behavioural components in this context are car-following and lane-changing, which capture longitudinal and lateral vehicle dynamics, respectively.

In addition to modelling human driving behaviour, the rise of autonomous vehicles has led to the development of trajectory planning algorithms that compute safe and efficient control actions in real time. These control approaches are often formulated as optimisation problems and have been applied in both lane-based and lane-free traffic environments.

This chapter reviews key developments in microscopic models of human driving and optimal controllers for autonomous vehicle trajectories. First, car-following, lane-changing, and integrated driving models are reviewed. Then, optimal control approaches for autonomous vehicles are discussed, with a focus on both lane-based and lane-free strategies. The relevant gaps are identified for both microscopic models and autonomous control frameworks.

2.2 Microscopic Modelling of Human Driving

This section reviews microscopic models of driving behaviour, focusing on car-following and lane-changing models. Most studies have modelled these two behaviours separately to keep

the modelling process simple (Hou et al., 2014, Toledo et al., 2003, Yang and Koutsopoulos, 1996). Only a few studies have considered both behaviours together, and those that do often use data-driven methods to estimate human driving patterns.

2.2.1 Car-following Models

Car-following models, primarily dealing with longitudinal control, are considered conceptually simpler due to a more constrained set of dynamics and have thus been more extensively studied. Gipps (1981) proposed a physics-based car-following model that considers limitations on a driver's braking rate to maintain a safe speed relative to the preceding car. This model simultaneously accounts for the driver's desired speed and the safe speed to mitigate the impact of slow-moving vehicles or obstructions that are far from the vehicle. This well-known model was extensively studied, modified, and validated by numerous researchers ((Ciuffo et al., 2012, Wilson, 2001, Yang et al., 2019) among others).

The Optimal Velocity Model (OVM) and its variants form another major family. In the original OVM (Bando et al., 1995), each driver aims for an "optimal" speed that depends on the current gap to the leader. The driver's acceleration is proportional to the difference between this optimal velocity and the vehicle's actual velocity. The OVM captures realistic phenomena such as free-flow traffic and stop-and-go waves, but it can predict unrealistically large accelerations in certain situations. To address this, Helbing and Tilch (1998) proposed a Generalised Force model, which adds a corrective term for negative speed differences (i.e. when the follower is faster than desired). Further refinement came with the Full Velocity Difference (FVD) model (Jiang et al., 2001), which explicitly includes the velocity difference as a second stimulus in the acceleration law. In the FVD model each vehicle's acceleration has two parts: one proportional to the gap-based optimal velocity error and another proportional to the actual speed difference. This makes the model more stable and realistic, combining the safe-braking logic with an optimal-velocity rule.

The Intelligent Driving Model (IDM), proposed by Treiber et al. (2000), is another popular deterministic car-following model. This time-continuous model describes the reaction of

a subject vehicle based on the stimulus from its preceding vehicle. The subject vehicle is assumed to consider both the safe separation from the vehicle in front and the desire to achieve the free-flow speed. The model was later calibrated based on three vehicle trajectories from a one-lane road in Stuttgart, Germany ([Kesting and Treiber, 2008](#)).

In recent years, there has been a shift toward data-driven and benchmarked car-following models. For instance, [Mo et al. \(2021\)](#) propose a physics-informed deep learning architecture (PIDL-CF) that embeds classical CF structures (IDM/OVM) into neural networks. Their PIDL models predict acceleration across multiple regimes (free, congested, emergency braking) and significantly outperform pure “black-box” networks on the NGSIM data, especially when training data are sparse. [Zhang et al. \(2022\)](#) take another hybrid approach by calibrating an IDM with time-varying parameters (to encode different driving styles) and then using a neural process (latent Gaussian process) to generate realistic trajectories. Their generative CF model achieves high accuracy in behaviour prediction and can synthesise plausible car-following trajectories for both observed and unobserved driving styles. [Hart et al. \(2024\)](#) demonstrate a reinforcement-learning CF model trained in a simulated environment of rare events: unlike models trained solely on natural driving data (which may crash on hard braking), their RL policy remains collision-free, comfortable, and string-stable even under extreme leader behaviour. In fact, the learned RL agent matches or exceeds traditional models’ performance in heavy braking scenarios, highlighting the robustness of learning-based control.

Empirical analyses have been used to calibrate and compare these classical forms. [Brockfeld et al. \(2004\)](#) calibrated ten car-following models on controlled freeway DGPS data and found calibration errors around 12–17% for all models, with inter-driver differences far exceeding inter-model differences. Similarly, [Punzo and Simonelli \(2005\)](#) fit multiple models (Newell, Gipps, IDM, MITSIM, etc.) to highway/freeway trajectories and reported that interpersonal variations were greater than inter-model variations in most cases. [Punzo et al. \(2012\)](#) elaborated this point by proposing rigorous calibration diagnostics for the Gipps model, highlighting limitations of many common calibration settings. [Ossen and Hoogendoorn \(2011\)](#) examined the heterogeneity in car-following behaviour, analysing differences among passenger car drivers, between passenger cars and trucks, and for different types of leading

vehicles. Their findings revealed significant heterogeneity in all cases, underscoring the importance of modelling individual driving preferences.

2.2.2 Lane-changing Models

Lane-changing is a key microscopic driving behaviour involving lateral manoeuvres in response to traffic conditions, road geometry, and driver intentions. Compared to car-following, lane-changing is more complex due to its multidimensional decision-making process, the involvement of multiple vehicles, and the variability of motivations. A wide range of models has been proposed to capture the decision logic, execution timing, and interaction dynamics involved in lane changes.

The categorisation of lane changes based on driver motivation was first introduced by [Gipps \(1986\)](#), who proposed a rule-based model evaluating the possibility, necessity, and desirability of a lane change. Although not explicitly labelled, this framework laid the foundation for distinguishing mandatory lane changes (MLCs) and discretionary lane changes (DLCs). Building on this, [Halati et al. \(1997\)](#) formally defined MLCs as lane changes required for route compliance (e.g., preparing to exit or avoiding a blocked lane), while DLCs are driven by preferences for better driving conditions.

Several studies adopted and extended these distinctions. For example, [Yang and Koutsopoulos \(1996\)](#) implemented a micro-simulation model in which MLCs are triggered by downstream constraints, such as lane blockages or upcoming exits, while DLCs are performed when the leader vehicle is perceived to be moving too slowly. [Ahmed \(1999\)](#) proposed a discrete choice model that decomposes the lane-changing process into three stages: the decision to consider a lane change, the choice of a target lane, and the acceptance of a suitable gap. In his framework, DLCs are only considered when MLC conditions are not active, and a forced merging model was introduced to capture interactions such as courtesy yielding.

To better capture transitions between MLC and DLC motivations, [Toledo et al. \(2003\)](#) proposed a unified utility-based framework in which both types of lane changes are evaluated jointly. The model considers explanatory variables including surrounding vehicle states, route

information, and driving preferences. Subsequent work by [Toledo et al. \(2005\)](#) expanded this approach to incorporate the choice of target lane, acknowledging that multiple sequential lane changes may be necessary. Utility values are assigned to each candidate lane based on speed, density, presence of heavy vehicles, and alignment with future routing. Lane-change execution is governed by a gap acceptance model based on spatial headways to adjacent vehicles. Further developments by [Toledo et al. \(2007\)](#) introduced the concepts of short-term goals (e.g., selecting a target lane) and short-term plans (e.g., identifying a target gap), linked through an acceleration model. This integrated framework captures the interdependence between lane choice, gap selection, and longitudinal control.

[Schakel et al. \(2012\)](#) proposed a lane change model centred around the idea of lane change desire. The desire follows from the trade-off between route, speed, and keep right incentives. As lane change desire increases, drivers become more assertive. Four levels of lane change desire are formulated, from lowest to highest, the lane changer performs: no lane change; lane change only in a free fashion; synchronization with target lane to prepare for lane change; lane change indication to create gap. Relaxation is implemented as drivers accept smaller time headways under large desire.

Alternative modelling approaches have also been explored. The MOBIL model ([Kesting et al., 2007](#)) links lane-changing with longitudinal control by evaluating the acceleration gains and losses for both the ego vehicle and nearby vehicles. Although it incorporates acceleration dynamics into the lane-change decision, the process remains sequential rather than fully coupled. Fuzzy logic and machine learning methods have also been applied to model driver decision-making under uncertainty ([Hou et al., 2012, 2014](#), [Hunt and Lyons, 1994](#)). These approaches often attempt to capture behavioural nuances and adapt to diverse driving styles.

A growing body of work has focused on game-theoretic approaches to model interactions between vehicles during lane changes. [Kita \(1999\)](#) developed one of the earliest such models for merging scenarios, framing the interaction between merging and through vehicles as a non-zero-sum game with complete information. [Liu et al. \(2007\)](#) extended this work by incorporating additional objectives such as speed variation and acceleration lane duration. More recent models have introduced incomplete or imperfect information structures to reflect

real-world uncertainty about other drivers' strategies. These include Bayesian games using Harsanyi's transformation, where player types are sampled from a distribution and strategies are based on beliefs (Ali et al., 2019, Talebpour et al., 2015). For a comprehensive review of game-theoretic lane-change models, see Ji and Levinson (2020).

Capturing the heterogeneity among drivers has often been overlooked in research. Sun and Elefteriadou (2014) collected behavioural data and driver characteristics through both focus group studies and in-vehicle driving tests. The results were categorised and used to model drivers' lane-changing behaviour. The proposed model incorporates a lane-change probability component and a gap acceptance component. A major contribution of the work is the consideration of driver heterogeneity.

Finally, empirical studies have played a critical role in informing model development and validation. Efforts have been made to more accurately model lane-change execution and timing. Traditional models treat lane-changing as a discrete event, but recent studies have approached it as a continuous process. Ji et al. (2023a) developed a hazard-based model that jointly analyses longitudinal and lateral movements during lane changes. The model estimates both the time between lane changes and the duration of the manoeuvre itself, offering a more detailed representation of the entire process. Similar works such as (Ahn and Cassidy, 2007, Treiber et al., 2000, Wang and Coifman, 2008, Zheng et al., 2013) have used naturalistic driving data to explore the frequency, duration, and spatial characteristics of lane changes under varying traffic conditions. These empirical insights continue to shape the evolution of lane-changing models toward greater realism and robustness.

2.2.3 Coupled Car-following and Lane-changing Models

There are clear interdependencies between car-following and lane-changing behaviours (Toledo, 2007). Despite this, only a limited number of studies have attempted to integrate the two processes into a unified modelling framework, largely due to the challenges of joint formulation and calibration. Among the studies that have done so, many rely on data-driven approaches to capture human driving behaviour, often at the cost of interpretability and

physical transparency. This section summarises the existing efforts that model car-following and lane-changing in a coupled manner.

A prominent early effort is the discrete choice model developed by [Toledo et al. \(2007\)](#), which integrates lane-change and acceleration decisions through a hierarchical framework. The model decomposes the overall decision-making process into sub-tasks across multiple layers, including lane-change desirability, target lane selection, gap selection, and acceleration control. Each lower-level decision is conditioned on outcomes from higher levels, allowing for an interdependent structure that reflects driver planning.

[Tomar et al. \(2010\)](#) proposed a machine learning approach using a multilayer perceptron (MLP) to predict vehicle trajectories from the NGSIM dataset. While not a behavioural model in the traditional sense, the MLP was trained to learn patterns in both longitudinal and lateral movements simultaneously. The model was able to achieve accurate short-term trajectory predictions, highlighting the ability of neural networks to implicitly capture coupled driving behaviour.

A more advanced integration was proposed by [Zhang et al. \(2019\)](#), who developed a long short-term memory (LSTM) model to simulate human decision-making in both car-following and lane-changing. The model takes as input the positions of six surrounding vehicles and learns the temporal evolution of the subject vehicle's behaviour. The LSTM architecture enables the model to retain memory of past interactions and encode time-dependent influences on driving decisions. To enhance learning performance, a Hybrid Retraining Constrained (HRC) training method was introduced, which constrains the learning process with auxiliary loss functions.

[Sharath and Velaga \(2020\)](#) developed an algorithm for generating human-like longitudinal and lateral motion for autonomous vehicles. The model builds upon the Intelligent Driver Model (IDM), adding features to incorporate lane-changing and overtaking. The formulation computes an effective gap by considering all surrounding obstacles, and a component of discretionary lateral acceleration is introduced to model discretionary lane changes. Although

TABLE 2.1. Calibration and validation of coupled car-following and lane-changing models of human driving

Paper	Datasets	Calibration procedure	Calibration trajectories	Validation trajectories
Toledo et al. (2007)	Arlington and Southampton	homogeneous	7608	-
Tomar et al. (2010)	NGSIM	homogeneous	50	50
Zhang et al. (2019)	NGSIM	homogeneous	820	783
Sharath and Velaga (2020)	NGSIM	homogeneous	1050	450
Shi et al. (2022)	NGSIM	homogeneous	2250	250

the model was developed for autonomous driving applications, it was calibrated and validated using human driving data.

More recently, [Shi et al. \(2022\)](#) proposed a two-dimensional trajectory prediction model based on deep learning. The model combines an attention mechanism, a bi-directional long short-term memory network (BiLSTM), and a Temporal Convolutional Network (TCN). A switch component determines whether the predicted trajectory includes a lane change. Lane-change types are discretised into three classes: left lane change, no lane change, and right lane change. The calibration and validation procedures of the above integrated models are summarised in Table 2.1. All microscopic models reviewed above are summarised in Table 2.2.

2.2.4 Research Gaps

Despite the progress made in modelling microscopic driving behaviours, several research gaps remain, particularly in relation to the coupling of driving tasks, interaction modelling, behavioural heterogeneity, and the treatment of MLCs.

A significant limitation of existing microscopic models is the continued treatment of car-following and lane-changing as independent processes. Empirical studies have demonstrated that these behaviours are closely interdependent. For instance, instability in car-following movements can induce discretionary lane-changes ([Gao and Levinson, 2023](#)). Similarly, lane-changing decisions can also affect gap sizes, impacting car-following decisions. In addition, rapid lane-changing movements have been shown to contribute to traffic instabilities and oscillations ([Ahn and Cassidy, 2007](#), [Zheng et al., 2013](#)). Despite this, only a limited number

TABLE 2.2. Summary of car-following, lane-changing, and coupled models

Study	Approach
<i>Car-Following Models</i>	
Gipps (1981)	Safety-based model combining desired speed with braking limits.
Bando et al. (1995)	OVM where drivers accelerate toward gap-dependent optimal speed.
Helbing and Tilch (1998)	Adds corrective force to OVM to reduce unrealistic accelerations.
Jiang et al. (2001)	Extends OVM with velocity-difference term for stability.
Treiber et al. (2000)	IDM balancing safe following distance with free-flow desire.
Mo et al. (2021)	Physics-informed deep learning embedding IDM/OVM.
Zhang et al. (2022)	Generative IDM with time-varying parameters and latent GP.
Hart et al. (2024)	RL-based CF trained on rare events; robust and collision-free.
<i>Lane-Changing Models</i>	
Gipps (1986)	Rule-based model of possibility, necessity, and desirability.
Halati et al. (1997)	CORSIM simulator formalising MLC vs DLC.
Yang and Koutsopoulos (1996)	Micro-simulation: MLC from constraints, DLC from slow leaders.
Ahmed (1999)	Discrete choice with three stages: consider, lane, gap.
Toledo et al. (2003)	Utility-based framework unifying MLC and DLC.
Toledo et al. (2005)	Sequential lane choice via utility evaluation.
Toledo et al. (2007)	Integrated goals/plans linking lane choice, gaps, and accel.
Schakel et al. (2012)	Lane-change desire model with four assertiveness levels.
Kesting et al. (2007)	MOBIL model using incentive criterion tied to acceleration.
Hunt and Lyons (1994)	Fuzzy logic for lane-change under uncertainty.
Hou et al. (2014)	Fuzzy logic capturing nuanced driver decisions.
Hou et al. (2012)	Genetic algorithm for adaptive lane-changing.
Kita (1999)	Game-theoretic merging as non-zero-sum game.
Liu et al. (2007)	Game-theoretic merging with speed and lane duration.
Talebpour et al. (2015)	Game-theoretic LC with V2V; Nash equilibrium under information limits.
Ali et al. (2019)	AZHW game model for mandatory LC in connected/traditional environment.
Sun and Elefteriadou (2014)	Empirical model capturing driver heterogeneity in LC.
<i>Coupled CF and LC Models</i>	
Toledo et al. (2007)	Hierarchical discrete choice linking LC and acceleration.
Tomar et al. (2010)	MLP trained on NGSIM for coupled trajectory prediction.
Zhang et al. (2019)	LSTM with HRC learning for joint CF+LC behaviour.
Sharath and Velaga (2020)	IDM extension with effective gap and lateral accel.
Shi et al. (2022)	BiLSTM+TCN+attention predicting LC (left/right/none).

of models have incorporated both behaviours within a unified decision-making structure, instead, most existing models represent lane changes as discrete decisions. In practice, however, lane-changing is a continuous process akin to acceleration and braking, and models that ignore this continuity fail to capture the dynamics of real manoeuvres.

Another important gap is the limited treatment of strategic interactions between drivers. While some models attempt to account for the presence and influence of nearby vehicles, these are often implemented through heuristic rules or via black-box data-driven models that lack explanatory power. Rule-based methods may struggle with adaptability under dynamic,

uncertain conditions, while purely data-driven approaches often fail to generalise beyond observed scenarios. Game-theoretic methods have emerged as a promising framework to address this issue, offering a more principled way to capture interactive, decision-making behaviour. However, their application remains relatively limited in the context of coupled microscopic driving models.

Furthermore, most models assume drivers behave homogeneously, neglecting well-documented variations in driving style and preference. While a few decoupled models have attempted to incorporate driver heterogeneity—for example, by allowing differences in desired headway or speed—such considerations are rarely extended to coupled models involving both car-following and lane-changing. This is despite evidence that human drivers exhibit substantial diversity in risk tolerance, lane-keeping discipline, and merging aggressiveness ([Ossen and Hoogendoorn, 2011](#), [Zhou et al., 2020](#)). Incorporating such heterogeneity is essential for improving model realism.

In the context of MLCs, many existing models implicitly assume that such manoeuvres are absolute, that is, once the conditions for an MLC are met, the driver is expected to perform the lane change without fail. While this assumption may be valid in highly constrained scenarios, such as avoiding an imminent lane blockage, it becomes less appropriate in situations where the lane change serves a longer-term tactical purpose, such as preparing for a downstream turn. In these cases, drivers may delay or even forgo the manoeuvre if current traffic conditions make it difficult or unsafe to change lanes.

Finally, current models rarely quantify the incentive to perform an MLC. The trade-offs involved, such as between maintaining the current lane versus initiating a potentially difficult merge, are typically unaccounted for. Capturing these trade-offs is essential to reflect real-world behaviour more accurately, especially in congested or mixed traffic conditions where drivers must weigh the risks and benefits of proceeding with a lane change under time or spatial constraints.

2.3 Autonomous Vehicle Optimal Control

AV control systems aim to generate safe and efficient driving behaviours through automated decision-making and trajectory planning. These control approaches can vary depending on the traffic environment, with some assuming structured lane-based roads and others designed for more flexible, lane-free scenarios. This section reviews key developments in AV trajectory planning under both frameworks.

2.3.1 Lane-based AV Control

While coupled models of car-following and lane-changing behaviours for human drivers remain limited, two-dimensional driving has been extensively studied in the context of optimal control for autonomous vehicles, particularly on highways.

A prominent line of research adopts hierarchical control architectures that separate high-level manoeuvre selection from low-level trajectory optimisation. [Glaser et al. \(2010\)](#) proposed a two-step trajectory planner for highway driving that first selects feasible longitudinal and lateral manoeuvres based on surrounding traffic and then refines candidate trajectories to optimise safety, comfort, and efficiency. Designed to operate on low-power embedded systems, the approach was shown to generate rule-compliant, collision-free trajectories in real-time. [Zhang et al. \(2013\)](#) extended this paradigm by decomposing complex driving scenarios into lane-keeping and lane-changing manoeuvres. Their method employs a two-stage optimisation framework that generates both static and adaptive trajectories to respond to dynamic traffic environments while balancing performance and comfort. Building on this hierarchical structure, [Chen et al. \(2020\)](#) proposed a hybrid control system for multi-lane highway driving, where a high-level finite-state machine governs discrete manoeuvre decisions, and a low-level model predictive controller computes smooth, coordinated longitudinal and lateral actions.

Other approaches integrate decision-making and motion planning within unified predictive control frameworks that handle both discrete and continuous aspects of driving. [Wang et al. \(2015\)](#) proposed a receding horizon control model based on dynamic game theory

that jointly captures car-following and lane-changing behaviours under both cooperative and non-cooperative scenarios. The framework simultaneously evaluates discrete lane-change decisions and continuous acceleration inputs, enabling the modelling of interactive behaviours in mixed traffic. [Wang et al. \(2021\)](#) introduced a control system structured around four interconnected modules: car-following, lane-changing decisions, trajectory generation, and model predictive trajectory tracking. The model accounts for interactions with nearby vehicles and was validated across diverse field conditions. [Lim et al. \(2021\)](#) proposed a hybrid trajectory planning framework that combines sampling-based lateral planning with optimisation-based longitudinal control for autonomous highway driving. This approach enables real-time adaptation to dynamic environments by generating diverse manoeuvre options laterally while optimising speed profiles longitudinally. The method was implemented in a C++ embedded system and validated through both simulation and real-world driving tests across various on-road scenarios.

Although most studies focus on highway environments, some have extended optimal control methods to two-dimensional turning at intersections. [Ma et al. \(2017\)](#) integrated a modified social force model with behavioural logic and motion constraints to simulate left-turn trajectories at a mixed-flow intersection, while [Bichiou and Rakha \(2019\)](#) applied Pontryagin's minimum principle to derive optimal intersection-crossing trajectories that minimise travel time while satisfying dynamic and static constraints. Compared to traditional control strategies, the model significantly reduces delay, fuel use, and emissions. More recently, [Zhao et al. \(2023\)](#) developed a behaviourally grounded optimal control model to reconstruct two-dimensional turning trajectories during unprotected left turns, validated on high-resolution trajectory datasets.

2.3.2 Lane-free AV Control

In lane-free autonomous driving, both connected and communication-free (unconnected) control strategies have been investigated ([Berahman et al., 2022](#), [Karafyllis et al., 2022b](#), [Karakakou et al., 2022](#), [Levy and Haddad, 2021](#), [Yanumula et al., 2021](#)). In this context, connected vehicles are capable of exchanging information, such as states, trajectories, intended

manoeuvres, or control inputs, while communication-free vehicles rely solely on onboard sensors to perceive and respond to their environment without external coordination.

Several notable connected control strategies have been proposed for lane-free autonomous driving. [Levy and Haddad \(2021\)](#) developed a nonlinear model predictive controller (NMPC) for the centralised control of vehicle groups on a ring road. The objective function seeks to maximise vehicle progress while minimising control effort and energy consumption, subject to constraints that maintain travel within the road boundaries and prevent inter-vehicle collisions. The effectiveness of this approach was further demonstrated using small-scale robotic vehicles in a physical testbed ([Levy and Haddad, 2022](#)).

Similarly, [Yanumula et al. \(2021\)](#) proposed a path planning controller under the assumption of vehicle connectivity. Each vehicle communicates its current state and short-term predicted trajectory to others, and a nonlinear constrained optimal control problem is solved in real-time using a Feasible Direction Algorithm. The framework was later extended to account for nudging behaviour, where vehicles gently influence others' trajectories to improve traffic flow ([Yanumula et al., 2023](#)).

[Berahman et al. \(2022\)](#) proposed a learning-based approach that employs repulsive forces to maintain safe distances between connected vehicles. Control strategies are trained using the Deep Deterministic Policy Gradient (DDPG) algorithm to maintain consistent separation between consecutive vehicles. Beyond repulsive interactions, several studies also explore attractive forces to enable cooperative flocking behaviours ([Chavoshi and Kouvelas, 2021](#), [Dabestani et al., 2024](#)). Notably, [Rostami-Shahrbabaki et al. \(2023\)](#) introduced a flocking-based control strategy in which vehicles are dynamically grouped into clusters based on attractive-repulsive energy functions. A consensus algorithm ensures vehicles within each flock travel at harmonised speeds, while a virtual leader at the upper level guides the collective motion by generating reference trajectories.

In parallel with connected strategies, a variety of communication-free, decentralised controllers have been proposed for lane-free autonomous driving. [Karafyllis et al. \(2022b\)](#) developed a nonlinear cruise controller that integrates Lyapunov functions, barrier functions,

and potential fields to ensure both safety and stability. The potential function introduces repulsive effects to prevent collisions with other vehicles and road boundaries, while the Lyapunov function guarantees convergence to desired motion profiles. The barrier function, in turn, enforces control limits and ensures safety constraints are respected. Building on this work, [Karafyllis et al. \(2022a\)](#) introduced two classes of cruise controllers—viscous and inviscid—derived from kinetic energy formulations inspired by Newtonian and relativistic mechanics. These controllers produce microscopic vehicle dynamics that aggregate into macroscopic fluid-like traffic models governed by conservation and momentum laws, with pressure and viscosity terms directly modulated by control parameters.

Focusing on learning-based approaches, [Karalakou et al. \(2022\)](#) addressed the problem of autonomous driving for a single agent in a lane-free setting. They formulated a set of reward functions aimed at balancing two primary objectives: maintaining a target speed and avoiding collisions. The reward design includes components for longitudinal target tracking, overtaking motivation, and collision risk, quantified using potential fields based on the relative positions and velocities of nearby vehicles. The proposed policy was trained and evaluated using the DDPG algorithm, with comparative analysis across different reward structures to assess effectiveness in various traffic contexts. All microscopic AV controllers reviewed above are summarised in Table 2.3.

2.3.3 Research Gaps

Despite the growing body of work on autonomous vehicle (AV) control, relatively few studies have focused on lane-free traffic environments. Most existing models and control strategies assume lane-based settings, leaving open questions about when and where lane-free traffic may offer efficiency or safety advantages. A systematic comparative analysis between lane-based and lane-free paradigms is currently lacking, and the specific conditions under which lane-free traffic becomes more beneficial remain unexplored. Given that lane-free traffic can be viewed as a generalised form of the traditional lane-based structure, further exploration of this setting may reveal new opportunities to enhance overall traffic performance.

TABLE 2.3. Summary of lane-based and lane-free AV controllers

Study	Approach
<i>Lane-based Control</i>	
Glaser et al. (2010)	Two-step planner: manoeuvre selection then trajectory refinement.
Zhang et al. (2013)	Two-stage optimisation for lane-keeping and lane-changing.
Chen et al. (2020)	FSM–MPC hybrid for highway driving.
Wang et al. (2015)	Game-theoretic receding horizon controller.
Wang et al. (2021)	Modular MPC for car-following and lane-changing.
Lim et al. (2021)	Hybrid lateral sampling and longitudinal optimisation.
Ma et al. (2017)	Social-force model for intersection left turns.
Bichiou and Rakha (2019)	Pontryagin-based optimal intersection crossing.
Zhao et al. (2023)	Behavioural optimal control for unprotected left turns.
<i>Lane-free Control</i>	
Levy and Haddad (2021)	Centralised NMPC for lane-free ring-road traffic.
Levy and Haddad (2022)	Testbed validation of NMPC with robots.
Yanumula et al. (2021)	Connected optimal path planner.
Yanumula et al. (2023)	Extension with trajectory nudging.
Berahman et al. (2022)	DDPG with repulsive-force interactions.
Dabestani et al. (2024)	Flocking via attractive–repulsive rules.
Chavoshi and Kouvelas (2021)	Distributed cooperative flocking.
Rostami-Shahrbabaki et al. (2023)	Consensus-based dynamic clustering.
Karafyllis et al. (2022b)	Lyapunov–barrier cruise controller.
Karafyllis et al. (2022a)	Viscous/inviscid cruise control models.
Karalakov et al. (2022)	DDPG policy balancing speed and safety.

While both connected and communication-free, as well as centralised and decentralised, control frameworks have been studied, communication-free decentralised approaches show particular promise. This is especially relevant in the near to medium term, when full AV adoption and ubiquitous connectivity is unlikely. These decentralised methods allow each AV to independently plan and optimise its trajectory based solely on onboard sensing, trading off global optimality for robustness, adaptability, and scalability. However, current approaches do not yet fully exploit the strategic interactions that can arise between vehicles. The use of game-theoretic reasoning to anticipate and respond to the behaviour of nearby AVs has not been explicitly addressed in most decentralised control studies.

Finally, heterogeneity among AVs has been largely overlooked in existing studies. Most models assume homogeneous behaviour, where all vehicles follow similar decision rules and cost functions. However, this simplification may fail to capture the diverse interactions that emerge in realistic traffic composed of AVs with customisable objectives, capabilities, or

driving styles. Incorporating behavioural heterogeneity could provide new insights into both individual vehicle performance and emergent traffic dynamics, with important implications for controller design and system-wide efficiency.

2.4 Conclusions

This chapter reviewed key developments in microscopic modelling of human driving and AV control. Traditional human driving models have primarily focused on car-following and lane-changing as separate components within lane-based traffic systems. While some efforts have aimed to unify these into integrated two-dimensional frameworks, such models remain limited in number and scope. The lack of comprehensive two-dimensional models restricts our ability to capture complex lateral-longitudinal interactions that are increasingly relevant in more flexible traffic environments.

In contrast, AV control research has grown rapidly, particularly in structured, lane-based environments. Both hierarchical and unified frameworks have been proposed, using techniques such as model predictive control and game theory. More recently, lane-free control approaches have gained attention, offering a generalised alternative to traditional lane-based systems. These include both connected and communication-free decentralised strategies, each with distinct benefits in coordination and scalability. However, gaps remain in understanding when lane-free traffic becomes advantageous, how strategic interactions between AVs can be modelled without communication, and how heterogeneous vehicle behaviours influence traffic dynamics.

Costs of Missing Freeway Exits

3.1 Introduction

As one of the most common microscopic road manoeuvres, lane change has attracted considerable research interest, e.g. (Ji and Levinson, 2020) among others. Poorly executed lane changes can cause flow oscillations and traffic incidents that may incur social and economic costs. More severe impacts are often observed under heavier traffic conditions with the potential of flow breakdowns under extreme cases (Ahn and Cassidy, 2007, Gao and Levinson, 2023, Zheng et al., 2013). Therefore, the development of precise lane change models is a crucial step for understanding the fundamental characteristics of traffic flow.

Based on its motivation, the notion of lane change has traditionally been classified as discretionary or mandatory. Discretionary lane change (DLC) occurs when the driver perceives that the target lane offers superior driving conditions compared to the current one. This encourages the driver to change lanes in aspiration of speed and safety advantages. Mandatory lane change (MLC), on the other hand, applies when the driver must leave the current lane, in order to avoid obstructions downstream (e.g., lane drop) or moving to the appropriate lane in preparation for future turning movements.

Numerous lane change models have been proposed in the literature. However, it is common for these models to tacitly accept that MLCs are absolute, meaning that the possibility of not performing the lane change is often overlooked. While this proposition holds true in scenarios where the driver must make a lane change decision (e.g., changing lanes to avoid a blockage), its validity is questionable in scenarios that are not as restrictive (e.g., changing lanes for

tactical routing). In this chapter, we study what would be the cost for not making a lane change, which consequently reflects the level of mandatoriness of a lane change manoeuvre. We specifically focus on the tactical routing scenario, examining lane changes made to diverge onto a freeway off-ramp. Hence, the study is conducted on the cost of not making the lane change decision at that exiting point and tries to corroborate “how costly would this misstep be”. It should be noted, the lane change we focus on in this chapter is centred around the combined concepts of lane change decision and execution. We assume that a lane change decision directly indicates whether a lane change has been executed and, consequently, whether the vehicle has taken the correct off-ramp. Thus, the problem transforms to examining the costs associated with missing off-ramps.

We start by formulating the exit-missing costs analytically by decomposing the costs into sub-components on the freeway and surface street network levels. Through utilising probability theory and stochastic geometry, the costs can be approximated with simple variables derived from the network structures. The proposed model is then validated using simulation data from the Longitudinal Employer-Household Dynamics (LEHD) Census dataset ([Bureau, 2022a](#)) for the top 50 metropolitan areas in the US.

In addition, we conduct an empirical investigation using two large-scale GPS datasets from the metropolitan region of Minneapolis and St. Paul. Real-life GPS data are considered vital for gaining insight into drivers’ decision-making processes, and as such, it is critical in analysing actual exit-missing costs. The datasets in this study were collected at high sampling rates and include comprehensive GPS trajectories.

Both the analytical and empirical results indicate that the additional costs associated with not executing lane changes to exit freeways are generally on the order of minutes and kilometres. Compared with lane changes performed to avoid obstacles, lane changing for tactical routing has lower mandatoriness. This may partially explain why drivers are sometimes prone to missing their scheduled turns. From a modelling perspective, models should account for the likelihood of drivers missing exits, which could be due to factors such as tiredness, distractions, or miscalculations regarding the ease of lane change manoeuvres. From a design

perspective, this study could be particularly useful in informing lane change decisions at the route choice level.

Findings of this chapter suggest that future research could investigate lane change models that distinguish between three categories: discretionary, expedient (ELC), and mandatory, where the benefits of changing lanes are on the order of seconds/metres, minutes/kilometres, and infinite, respectively. In this chapter, we define ELC as lane changes performed only for tactical routing. All other previously defined MLCs remain in the MLC category. Lane changes made to avoid downstream obstacles are considered MLCs because the cost of not executing such a lane change would, intuitively, be infinite. If no lane change decisions are made, the drivers could be stuck indefinitely. Other previously defined MLCs, such as obeying lane usage indications or yielding to emergency vehicles, remain in the MLC category until further verification of their costs. However, we hypothesise that they might be better classified as ELCs since drivers sometimes choose not to make these lane changes. Drivers often weigh the trade-offs between changing lanes and not changing lanes before making decisions.

3.2 Analytical Costs for Missing Freeway Exits

3.2.1 Mean of Exit-Missing Costs

Our study initially investigates the distance cost incurred due to missed exits. This cost, expressed as the extra travel distance (D_c), can be subdivided into two distinct components: the freeway level cost ($D_{c,F}$) and the surface street level cost ($D_{c,S}$). $D_{c,F}$ is simply the gap length between the alternative and the original exits, whereas $D_{c,S}$ represents the discrepancy between the distance travelled from the alternative exit to the destination and the distance travelled from the original exit to the destination. The corresponding travel distances on the surface street network can be further decomposed to the circuitities (C_a and C_o) multiplying the Euclidean distances ($|l_a|$ and $|l_o|$) from the exits to the destination. Circuituity is defined as the ratio between the route length and the Euclidean distance for an OD pair (Axhausen et al., 2003). Figure 3.1 visualises these relationships.

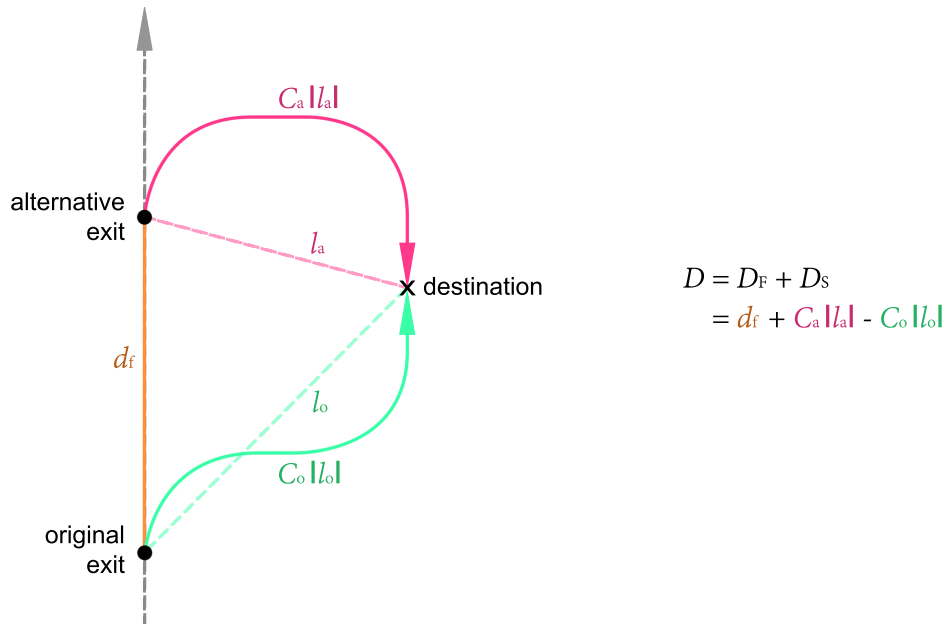


FIGURE 3.1. Diagram illustrating the composition of the costs. In this case, the vehicle moves in an upward direction along the freeway, indicated by the grey dashed line. The pink and green routes denote the paths from the exits on the freeway to the destination (marked as a cross). The orange line corresponds to the section of the freeway located between the two exits. The cost incurred due to missing the original exit and subsequently taking the alternative exit is equal to the combined length of the orange line and the disparity between the lengths of the pink and green lines.

Extending the aforementioned approach to encompass multiple trips over a region, the formulation transforms into determining the average distance cost by calculating the mean of its constituent parts. While the average gap length and the average circuitry can be computed with relative ease, estimating the average Euclidean distance between a destination and its original or alternative exits presents a challenge, as the location of the destination is a random variable, and the distance is contingent upon it. To tackle this issue, we assume that people use a pool of nearby exits to reach their locations and subsequently determine the proportion of trips that belong to each exit. Next, we calculate the average Euclidean distance of the exits, and this value, when considered alongside their proportion of usage, forms the measure of the average distance to the alternative and original exits.

Suppose that freeways can be subdivided into segments at exits. For one arbitrary segment (see Figure 3.2), the percentage of people travelling to the nearest (1st order) and 2nd nearest

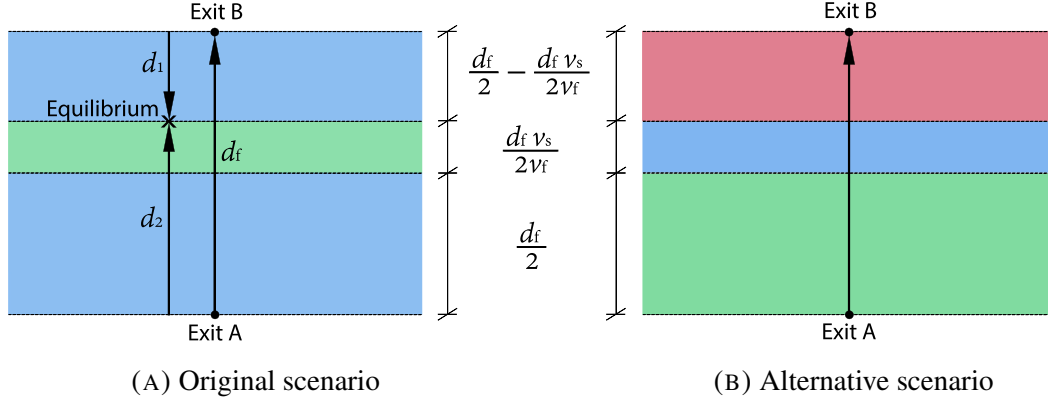


FIGURE 3.2. Diagram showing the contribution of exits to the surrounding destinations. The blue regions indicate that destinations within this area are reached from their nearest exit, the green area corresponds to the second nearest exit, and the red area corresponds to the third nearest exit. Subgraph (a) shows the contribution of exits for the original route while (b) shows the contribution of exits for the alternative route, which we assume that drivers have missed their intended exits.

(2nd order) exits can be approximated as follows: If we assume that individuals are fully rational and choose to travel through the off-ramp that minimises their travel time, then there exists an equilibrium point between exits A and B at which the travel time from both exits is identical. Suppose the velocity on the freeway is higher than that on the surface street network (i.e. $v_s < v_f$). In this case, the location of the equilibrium point can be determined using the following relationship.

$$d_f = d_1 + d_2$$

$$\frac{d_f}{v_f} + \frac{d_1}{v_s} = \frac{d_2}{v_s}$$

where d_f is the length between consecutive exits; d_1 and d_2 are distances to the equilibrium from exit A and exit B depicted in Figure 3.2a; v_f and v_s are speeds on the freeway and surface street networks respectively.

By solving the above equations simultaneously, the length of d_1 and d_2 can be obtained. For the original route (Figure 3.2a), people with destinations below the equilibrium (in d_2)

use exit A and people with destinations above the equilibrium (in d_1) use exit B. Therefore, $1 - v_s/(2v_f)$ proportion of people (in blue areas) exit through their nearest off-ramp and $v_s/(2v_f)$ proportion (in green area) exit through their second nearest off-ramp.

Similarly, considering that the intended exits are now missed (Figure 3.2b), people living in the green area (the bottom segment) who previously travelled using exit A will now use exit B instead, which is their second order exit; people in the blue area (the middle segment) who previously used exit A now use exit B, their nearest exit; and people in the red zone (the top segment), who used exit B for travel now have to use the next downstream exit (not shown on the figure), which would be of order three. Therefore, after missing exits, $v_s/(2v_f)$ proportion of people use their nearest exit, $1/2$ proportion of them use the second nearest exit, and $1/2 - v_s/(2v_f)$ proportion use the third nearest exit.

Assuming that all segments between the exits follow the same allocation throughout the sample region, the representations for the mean $D_{c,F}$ and $D_{c,S}$ are:

$$\langle D_{c,F} \rangle = \frac{\sum_i^m d_{f,i}}{m} = \bar{d}_f \quad (3.1)$$

$$\begin{aligned} \langle D_{c,S} \rangle &= \bar{C} \left[\langle r_1 \rangle \frac{\bar{v}_s}{2\bar{v}_f} + \langle r_2 \rangle \frac{1}{2} + \langle r_3 \rangle \left(\frac{1}{2} - \frac{\bar{v}_s}{2\bar{v}_f} \right) \right] - \bar{C} \left[\langle r_1 \rangle \left(1 - \frac{\bar{v}_s}{2\bar{v}_f} \right) + \langle r_2 \rangle \frac{\bar{v}_s}{2\bar{v}_f} \right] \\ &= \bar{C} \left[\langle r_1 \rangle \left(\frac{\bar{v}_s}{\bar{v}_f} - 1 \right) + \langle r_2 \rangle \left(\frac{1}{2} - \frac{\bar{v}_s}{2\bar{v}_f} \right) + \langle r_3 \rangle \left(\frac{1}{2} - \frac{\bar{v}_s}{2\bar{v}_f} \right) \right] \end{aligned} \quad (3.2)$$

where $d_{f,i}$ represents the distance of the i^{th} gap between two successive exits; m is the total number of gaps between consecutive exits; \bar{C} is the average circuitry of the region; $\langle r_n \rangle$ is the average Euclidean distance between any random destination on the region and its n^{th} closest exit; and \bar{v}_f and \bar{v}_s are the average speeds on freeway and street networks, respectively.

To estimate the average Euclidean distance from any destination to its n^{th} closest exit $\langle r_n \rangle$, one can naively double-integrate the Euclidean norm for x over 0 to infinity and y from 0 to the height of the block. However, the solution does not converge to a finite value and, therefore,

cannot be used for approximation. While it may be possible to numerically determine the range of y values, the resulting error would be too large for the approximation to be valid. Instead, we make the assumption that the spatial distribution of exits across the region follows a two-dimensional homogeneous Poisson point process (HPPP). This assumption allows us to generalise the problem and obtain a solution using probability theory. For a point process $\Psi \subset \mathbb{R}^2$ to be an HPPP, it has to inhere the following properties (Stoyan et al., 2013):

- The number of points of the point process Ψ within any Borel set $B \subset \mathbb{R}^2$ follows a Poisson distribution, i.e.,

$$\mathbb{P}(\iota(B) = n) = \frac{e^{-\Lambda(B)}(\Lambda(B))^n}{n!} \quad (3.3)$$

where $\iota(B)$ is the counting measure, defined as $\iota(B) = \sum_{x_i \in \Psi} \mathbb{1}(x_i \in B)$, and $\Lambda(B)$ is the intensity measure, defined as $\Lambda(B) = \mathbb{E}[\iota(B)]$.

- For M disjoint sets B_1, \dots, B_M , the random variables $\iota(B_1), \dots, \iota(B_M)$ are independent.
- The intensity ξ is a constant.

A useful property that the HPPP exhibits for the derivation of the n^{th} nearest neighbour distribution is motion-invariance (Daley and Vere-Jones, 2008). A point process is said to be motion-invariant if it is both stationary and isotropic. Stationary means that if the observer moves or translates the coordinate system, the distribution of the points do not change. Isotropic, on the other hand, implies that the distribution of points is invariant to the rotation of the axes. Together they suggest that regardless of the realisation of the subset B , the points inside will always follow the same HPPP (Poisson distributed with mean $\xi \ell(B)$, where $\ell(B)$ is the area of B in \mathbb{R}^2). Consequently, the distributions of the n^{th} nearest neighbours are invariant to the underlying distribution of the random destinations, such as homes, offices, and other functional areas, which are known to display heterogeneous spatial patterns. This property enables the modelling of exit-missing costs without making assumptions or evaluations on the spatial distribution of the destinations, which can simplify our analysis.

We now introduce the cumulative distribution function (CDF) and the probability density function (PDF) for the distance R_n to the n^{th} nearest point from the origin. This is an extension to the distance distribution of the nearest neighbour commonly used in wireless communication problems (Baccelli et al., 1997).

LEMMA 1. *The CDF and PDF of the n^{th} nearest-neighbour distance for a motion-invariant HPPP Ψ with intensity ξ are*

$$CDF : F_{R_n}(r_n) = 1 - e^{-\xi\pi r_n^2} \sum_{k=0}^{n-1} \frac{(\xi\pi r_n^2)^k}{k!} \quad (3.4)$$

$$PDF : f_{R_n}(r_n) = \frac{e^{-\xi\pi r_n^2} (\xi\pi r_n^2)^{n-1}}{(n-1)!} 2\xi\pi r_n \quad (3.5)$$

PROOF: By definition, the CDF of R_n is

$$F_{R_n}(r_n) = \mathbb{P}(R_n \leq r_n) = 1 - \mathbb{P}(R_n > r_n)$$

Envisage a disk $B(o, r_n)$ with its centre at the origin and radius r_n . The event that the distance of the n^{th} closest point R_n is greater than r_n implies that the number of points intersecting the disk would be at most $n - 1$. Thus, the above expression can be elaborated as follows

$$\begin{aligned} F_{R_n}(r_n) &= 1 - \mathbb{P}(\iota(B(o, r_n)) \leq n - 1) \\ &= 1 - \sum_{k=0}^{n-1} \mathbb{P}(\iota(B(o, r_n)) = k) \end{aligned}$$

We know from the definition that $\iota(B(o, r_n))$ is Poisson distributed. Therefore, the CDF of R_n is given by

$$F_{R_n}(r_n) = 1 - \sum_{k=0}^{n-1} \frac{e^{-\xi\pi r_n^2} (\xi\pi r_n^2)^k}{k!}$$

Then, the PDF of R_n is simply obtained by differentiating the CDF with respect to r_n . Or else, the PDF can be derived by calculating the product of the probabilities that $n - 1$ points are embedded in the disk and 1 point lies on the disk. Both would yield the same function. \square

From this point, the average distance to the n^{th} closest point can be computed as

$$\begin{aligned} \langle r_n \rangle &= \frac{\int_0^\infty r_n f(r_n) dr_n}{\int_0^\infty f(r_n) dr_n} \\ &= \frac{\int_0^\infty r_n e^{-\xi\pi r_n^2} (\xi\pi r_n^2)^{n-1} 2\pi r_n \xi dr_n}{\int_0^\infty e^{-\xi\pi r_n^2} (\xi\pi r_n^2)^{n-1} 2\pi r_n \xi dr_n} \end{aligned} \quad (3.6)$$

By substituting $\xi\pi r^2$ with u , we can determine the mean cost as follows:

$$\begin{aligned} \langle r_n \rangle &= (\xi\pi)^{-\frac{1}{2}} \frac{\int_0^\infty e^{-u} u^{n-\frac{1}{2}} du}{\int_0^\infty e^{-u} u^{n-1} du} \\ &= (\xi\pi)^{-\frac{1}{2}} \left[\frac{\Gamma(n + \frac{1}{2})}{\Gamma(n)} \right] \end{aligned} \quad (3.7)$$

With the above derivations, the average distance cost for missing exits can be approximated:

$$\begin{aligned} \langle D_c \rangle &= \bar{d}_f + \frac{\bar{C}}{\sqrt{\xi\pi}} \left[\left(\frac{\bar{v}_s}{\bar{v}_f} - 1 \right) \left(\Gamma(1.5) - \frac{\Gamma(2.5)}{2} - \frac{\Gamma(3.5)}{4} \right) \right] \\ &= \bar{d}_f + \frac{11\bar{C}}{32\sqrt{\xi}} \left[1 - \frac{\bar{v}_s}{\bar{v}_f} \right] \end{aligned} \quad (3.8)$$

The time cost can also be approximated using the solution we derived above. By considering the distance cost on the freeway and surface street separately, the approximated time cost is

$$\begin{aligned}
\langle T_c \rangle &= \frac{\langle D_{c,F} \rangle}{v_f} + \frac{\langle D_{c,S} \rangle}{v_s} \\
&= \frac{\bar{d}_f}{v_f} + \frac{\bar{C}}{v_s \sqrt{\xi \pi}} \left[\left(\frac{v_s}{v_f} - 1 \right) \left(\Gamma(1.5) - \frac{\Gamma(2.5)}{2} - \frac{\Gamma(3.5)}{4} \right) \right] \\
&= \frac{\bar{d}_f}{v_f} + \frac{11\bar{C}}{32\sqrt{\xi}} \left[\frac{1}{v_s} - \frac{1}{v_f} \right]
\end{aligned} \tag{3.9}$$

3.2.2 Standard Deviation of Exit-Missing Costs

To fully characterise the exit-missing cost, we also formulate the standard deviation of the costs. The cost functions consist of a number of random variables, including the exit gap (d_f), circuitry (C), freeway speed (v_f), surface street speed (v_s), as well as the Euclidean distance to the three nearest exits (r_1, r_2, r_3). The standard deviation of the costs can be determined by considering the variance of each random variable and their respective covariance. Note, the parameter ξ is implicitly considered in the mean distance to the nearest neighbours (r_1, r_2, r_3).

We approximate the variance of the cost functions using the Delta method (Liu, 2012). Here, we assume all variables are independent, and the covariance terms are zero. Note the Euclidean distance to the three nearest exits (r_1, r_2, r_3) are assumed to be independent because these distances are calculated from any random points in the region where the variables are three independent realisations of the same underlying point process. First, we confirm that both the distance and time cost functions are smooth (i.e. continuous and differentiable), which indicates that they can be approximated with Taylor expansion. To avoid an infinite regress of higher-order terms, we use the first-degree Taylor polynomial around the population mean $\bar{\mu}$. A function G with k independent random variables is expressed as:

$$G(x_1, \dots, x_k) \cong G(\bar{\mu}_1, \dots, \bar{\mu}_k) + (x_1 - \bar{\mu}_1, \dots, x_k - \bar{\mu}_k) \begin{pmatrix} \frac{\partial G}{\partial x_1}(\bar{\mu}_1, \dots, \bar{\mu}_k) \\ \vdots \\ \frac{\partial G}{\partial x_k}(\bar{\mu}_1, \dots, \bar{\mu}_k) \end{pmatrix} \tag{3.10}$$

To use the Delta method for approximating the underlying distribution of function G , we assume all variables are asymptotically normal (i.e., $x_i \sim N(\bar{\mu}_i, \bar{\sigma}_i^2)$ for large samples). This is a reasonable assumption and allows us to avoid other assumptions for the back-transformation. Moreover, this assumption leads to reasonable estimates of mean and variance, as shown later in Section 3.3.2. After rearrangement, $x_i - \bar{\mu}_i \sim N(0, \bar{\sigma}_i^2)$. Therefore:

$$(G(x_1, \dots, x_k) - G(\bar{\mu}_1, \dots, \bar{\mu}_k)) \cong (N(0, \bar{\sigma}_1^2), \dots, N(0, \bar{\sigma}_k^2)) \begin{pmatrix} \frac{\partial G}{\partial x_1}(\bar{\mu}_1, \dots, \bar{\mu}_k) \\ \vdots \\ \frac{\partial G}{\partial x_k}(\bar{\mu}_1, \dots, \bar{\mu}_k) \end{pmatrix} \quad (3.11)$$

Hence, by the properties of the normal distribution:

$$G(x_1, \dots, x_k) \sim N \left(G(\bar{\mu}_1, \dots, \bar{\mu}_k), (\bar{\sigma}_1^2, \dots, \bar{\sigma}_k^2) \begin{pmatrix} \left(\frac{\partial G}{\partial x_1}(\bar{\mu}_1, \dots, \bar{\mu}_k) \right)^2 \\ \vdots \\ \left(\frac{\partial G}{\partial x_k}(\bar{\mu}_1, \dots, \bar{\mu}_k) \right)^2 \end{pmatrix} \right) \quad (3.12)$$

To adapt the above expression to our problem, the partial derivatives of the variables for both the time and distance costs are obtained and summarised in Table 3.1.

TABLE 3.1. Partial derivatives for all random variables in the cost functions.

	distance cost (D_c)	time cost (T_c)
$\frac{\partial}{\partial d_f}$	1	$\frac{1}{v_f}$
$\frac{\partial}{\partial C}$	$r_1 \left(\frac{v_s}{v_f} - 1 \right) + r_2 \left(\frac{1}{2} - \frac{v_s}{2v_f} \right) + r_3 \left(\frac{1}{2} - \frac{v_s}{2v_f} \right)$	$r_1 \left(\frac{1}{v_f} - \frac{1}{v_s} \right) + r_2 \left(\frac{1}{2v_s} - \frac{1}{2v_f} \right) + r_3 \left(\frac{1}{2v_s} - \frac{1}{2v_f} \right)$
$\frac{\partial}{\partial v_f}$	$C \left[-r_1 \frac{v_s}{v_f^2} + r_2 \frac{v_s}{2v_f^2} + r_3 \frac{v_s}{2v_f^2} \right]$	$-\frac{d_f}{v_f^2} + C \left[-r_1 \frac{1}{v_f^2} + r_2 \frac{1}{2v_f^2} + r_3 \frac{1}{2v_f^2} \right]$
$\frac{\partial}{\partial v_s}$	$C \left[\frac{r_1}{v_f} - \frac{r_2}{2v_f} - \frac{r_3}{2v_f} \right]$	$C \left[\frac{r_1}{v_s^2} - \frac{r_2}{2v_s^2} - \frac{r_3}{2v_s^2} \right]$
$\frac{\partial}{\partial r_1}$	$C \left(\frac{v_s}{v_f} - 1 \right)$	$C \left(\frac{1}{v_f} - \frac{1}{v_s} \right)$
$\frac{\partial}{\partial r_2}$	$C \left(\frac{1}{2} - \frac{v_s}{2v_f} \right)$	$C \left(\frac{1}{2v_s} - \frac{1}{2v_f} \right)$
$\frac{\partial}{\partial r_3}$	$C \left(\frac{1}{2} - \frac{v_s}{2v_f} \right)$	$C \left(\frac{1}{2v_s} - \frac{1}{2v_f} \right)$

An estimation of the variance for the exit gaps, circuitry, freeway speed, and surface street speed can all be obtained numerically, whereas the variance ($\bar{\sigma}_{r_n}^2$) of the distance r_n to the

n^{th} nearest exit requires the formulation of its mean (Equation 3.7) and the second statistical moment of the probability density:

$$\begin{aligned}
\bar{\sigma}_{r_n}^2 &= \mathbb{E}(r_n^2) - (\mathbb{E}(r_n))^2 \\
&= \frac{\int_0^\infty r_n^2 e^{-\xi\pi r_n^2} (\xi\pi r_n^2)^{n-1} 2\pi r_n \xi \, dr_n}{\int_0^\infty e^{-\xi\pi r_n^2} (\xi\pi r_n^2)^{n-1} 2\pi r_n \xi \, dr_n} - (\xi\pi)^{-1} \left(\frac{\Gamma(n + \frac{1}{2})}{\Gamma(n)} \right)^2 \\
&= (\xi\pi)^{-1} \left(\frac{\Gamma(n+1)}{\Gamma(n)} \right) - (\xi\pi)^{-1} \left(\frac{\Gamma(n + \frac{1}{2})}{\Gamma(n)} \right)^2 \\
&= (\xi\pi)^{-1} \left[n - \left(\frac{\Gamma(n + \frac{1}{2})}{\Gamma(n)} \right)^2 \right]
\end{aligned} \tag{3.13}$$

Therefore, for a cost function G (either D_c or T_c), the estimated standard deviation is calculated as follows:

$$\hat{\sigma}_G = \sqrt{\sum_{x \in \{d_f, C, v_f, v_s, r_1, r_2, r_3\}} \left(\frac{\partial G(\bar{d}_f, \bar{C}, \bar{v}_f, \bar{v}_s, \langle r_1 \rangle, \langle r_2 \rangle, \langle r_3 \rangle)}{\partial x} \right)^2 \bar{\sigma}_x^2} \tag{3.14}$$

3.3 Simulation Experiments of Exit-Missing Costs in the Top 50 Metropolitan Areas

While we have formulated the exit-missing costs analytically, their performance against the actual values is as yet undetermined. Since real GPS data containing missed exits are rare (we discuss an example in a subsequent section), we evaluate our model with simulated costs for the 50 most populated metro areas in the US. The effectiveness of the model is showcased against a baseline, and various evaluation metrics are calculated. The limitations of the analytical model are discussed in Section 3.5.

3.3.1 Extraction of Exit-Missing Costs

This section demonstrates the computations of exit-missing costs for the top 50 most populated metro areas in the US. All trips involving links outside of the metros are discarded. The trip generation data used in our research are a subset of the 2013 Longitudinal Employer-Household Dynamics (LEHD) dataset termed LEHD Origin-Destination Employment Statistics (LODES), which contains state-wise commuting data at census block levels (Bureau, 2022a). A census block is the smallest areal unit for which detailed social demographic and mode share data are available. A typical census block consists of 250-550 housing units, and for this reason, the number of people commuting between any block pair is usually small (most are 0 or 1).

The shapefiles of these census blocks were acquired from the Census TIGER/line files (Bureau, 2022b). Since not all centroids lie within their respective blocks, the internal points of the polygons are snapped to their nearest nodes in the 2013 Open Street Map (OSM) traffic network. One drawback of the OSM network is the absence of speed limits on some links. This was accounted for by imputing the travel speeds of the missing values using the mean of the remaining links for the particular highway type. The link travel time is then calculated using the speed and the length. It should be noted that we have omitted the impact of traffic congestion in this Section; the travel time calculated is based solely on the free-flow speed. The impact of traffic congestion on exit-missing costs is further evaluated in Section 3.4.2.4.

The method we use to determine the exit-missing costs is by randomly sampling 10,000 OD pairs from the LODES data for each metropolitan area, where the sampling weight follows from the inter-block commuting number (i.e. the number of trips belonging to an OD). For the sake of simplicity, we assume all flows on an OD pair are allocated to the shortest path according to an all-or-nothing assignment. Then for each OD, Dijkstra's algorithm computes the shortest path with travel time as the edge impedance. This route is then regarded as the original route with its final exit stored (if it exists, else remove the OD from the list). The travel costs, namely time and distance, from the upstream point of the final exit to the destination are computed as the original costs. The alternative path assumes the driver has

missed the intended exit and will use the following off-ramp to exit the freeway, that is of saying, the alternative route proceeds along the mainline, and the shortest path algorithm is utilised to determine the most efficient route from the second exit to the destination. The travel costs for this shortest path, along with the extra portion on the freeway, are calculated as the alternative costs. Consequently, the cost of missing an exit is calculated by subtracting the original cost from the alternative cost. Since both routes are computed from the shortest paths, the cost is guaranteed to be positive. Finally, for each metro area, the costs from all the sampled ODs are aggregated across the commuting population, and the mean and standard deviation of the costs are computed. Note that the impacts of traffic flows, variations in speed, and the presence of traffic signals are not accounted for in this section. To overcome the above limitations, we conduct experiments on a real-world GPS dataset in Section 3.4. The LEHD dataset used in this section only contains commuting OD information. Further commuting related analysis are also investigated in Section 3.4. The empirical results demonstrate that the exit-missing costs are consistent with the values computed in this section and thus support the conclusion that costs for missed exits are negligible.

3.3.2 Goodness of Fit with Analytical Costs

We undertake a comparative analysis to showcase the dependencies between predicted and observed costs. Figure 3.3 illustrates the scatter plots of the simulated mean values and the estimated mean values of the distance and time costs of missing the freeway exit. Likewise, residual plots are presented to showcase potential bias in our model. Residuals are calculated by subtracting the predicted values from the actual/simulated values. The scatter plots show a clear linear relationship between the two variables for both time and distance, and the points scatter approximately along the 45-degree line. Las Vegas (LAS), Riverside (RAL), and Salt Lake City (SLC) are exceptions that are computed to be outliers using the Bonferroni outlier test (Fox, 2008). Further looking into the three instances shows that the low exit densities, together with large exit gaps, have contributed to the large discrepancies. As an example, the average exit gap for SLC is 4.380 km, which even exceeds the total predicted distance costs (average exit gap and surface street network cost combined) for most metro areas. This is

due to extremely large exit gaps embedded in the metropolitan regions (maximum gap length of 68.470 km for SLC), which inflate the average gap length. Moreover, the three outliers have the lowest exit densities among all 50 metros. This is due to them having dispersed freeway networks, lowering the exit-area ratio, and further exaggerating the costs. Thus, it is reasonable to conjecture that the models would align more closely in regions where freeway exits exhibit more uniform densities. In urban settings, freeways and their exits tend to be densely distributed, whereas inter-urban freeways and their exits are sparser. When considering the metropolitan region in its entirety, there will be zones with denser exits and others that are less dense. This variability likely accounts for the significant discrepancies observed in the predicted values.

The root-mean-square error (RMSE) and mean absolute error (MAE) are applied as the evaluation metrics for quantifying the performance of the predicted costs. After removing the outliers, namely SLC, RAL, and LAS, the RMSE and MAE for distance costs are 0.524 km and 0.413 km, respectively. While for the time costs, the RMSE and MAE are 0.418 min and 0.333 min. To extrapolate what these figures suggest, we construct a simple baseline for the predictor using the mean of the 50 observed values (3.558 km for distance and 2.307 min for time). The respective RMSE and MAE for distance costs are 0.831 km and 0.716 km, and for time are 0.465 min and 0.392 min. Hence, our model has outperformed the simple baseline model with solely the information on metropolitan road networks and can thus provide valuable insights into the compositions of the costs.

In addition, we conducted a stratified analysis of the simulated costs, distinguishing between trips made to work and those made returning home, and assessed the model's performance in each case. The scatter diagrams for the two scenarios are illustrated in Figure 3.4. Notably, the MAE values indicate that the model performs better for the to-home trips than the to-work trips. This is because our model is formed based on the assumption that the exits are distributed homogeneously, implying that analytical outcomes remain constant irrespective of destination distribution (due to motion-invariance). However, in reality, the exits are not distributed homogeneously, and neither are the destinations. While setting the destinations to homes does not make the exits appear homogeneous, the simulated relative spacings between

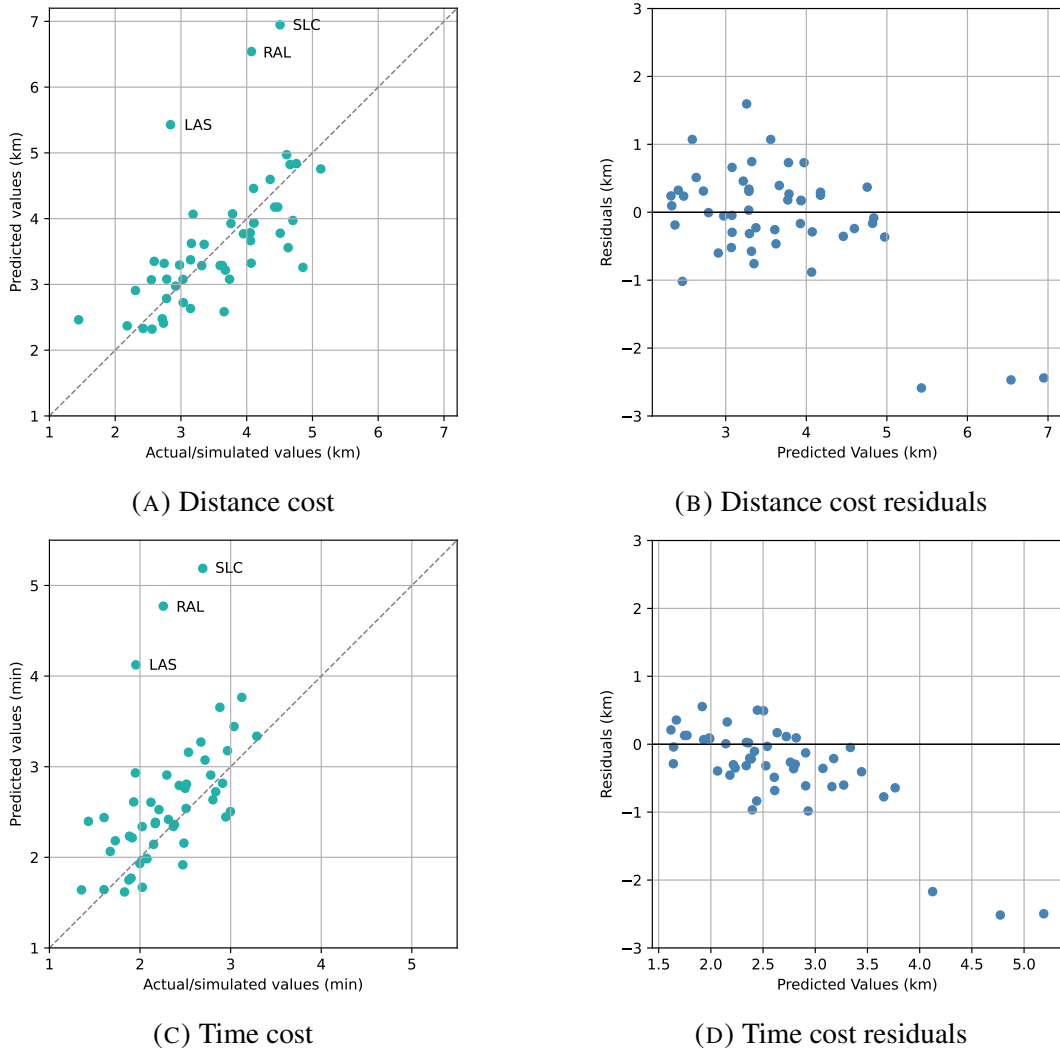


FIGURE 3.3. Scatter plots comparing the actual and predicted mean of costs. Subplots (a) and (c) are scatter plots showing the relationship between the observed and predicted values. The three outliers marked are Las Vegas (LAS), Riverside (RAL), and Salt Lake City (SLC). Subplots (b) and (d) are residual plots showcasing the distribution of residuals against the predictions. The MAE for the exit-missing distance cost is 0.413 km, and is 0.333 min for the exit-missing time cost.

the two would be more homogeneous since homes are more widely located than workplaces, which are more often concentrated in business districts. If the distribution of the destinations is highly clustered or regular in the vicinity of the exit points, then the travel distance between the two sets may be affected by this inhomogeneous nature of the destination set.

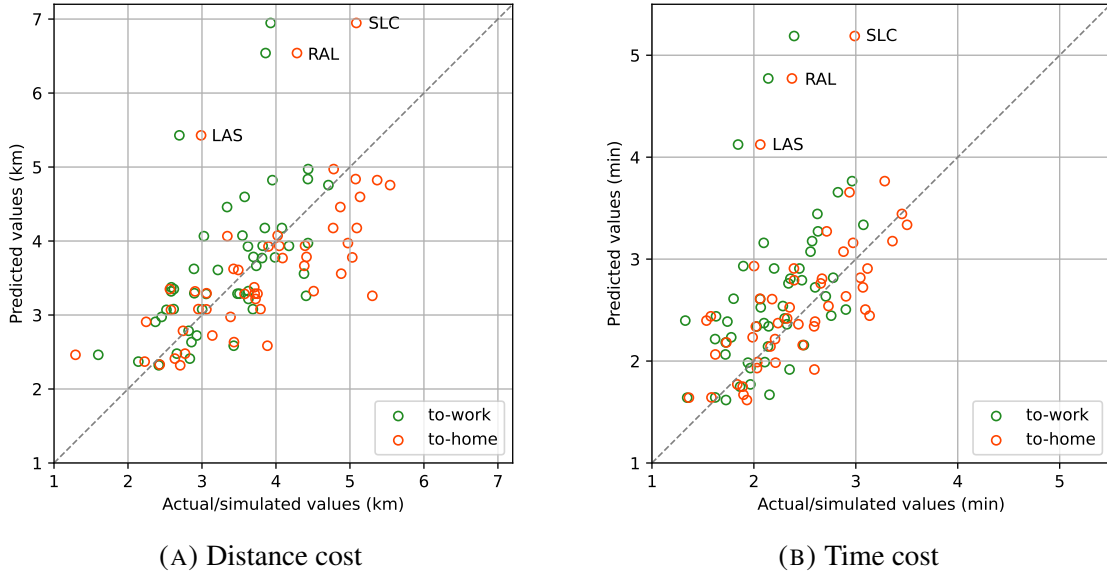


FIGURE 3.4. Scatter plots for the actual and predicted mean of costs for to-work and to-home trips. To-work trips are shown in green and to-home trips are in red. The outliers marked are Las Vegas (LAS), Riverside (RAL), and Salt Lake City (SLC). Subplots (a) and (b) illustrate the scatters of the distance and time costs, respectively. The MAE for the exit-missing distance cost for the respective to-work and to-home trips are 0.438 km and 0.531 km. For the exit-missing time cost, the MAE for the respective to-work and to home trips are 0.414 min and 0.313 min.

The results for the standard deviation of the costs are somewhat less promising (shown in Figure 3.5). Although there is still a positive correlation between the predicted and observed values, the analytical model tends to produce underestimated standard deviations, as shown by the points falling in the lower half of the scatter plot and the residuals clustering at some positive values. The reason behind this discrepancy is twofold. First and foremost, we assume the distribution of exits to be spatially homogeneous, which would down-estimate the variance for the distance to the n^{th} closest exits. Secondly, we assume that the random variables in the formulations are independent, which yields zero covariance between the terms. However, in reality, many of these variables, such as the exit gap length and the distance to the n^{th} nearest exit, are interdependent. As a result, the variance and hence the standard deviation of the costs are undervalued. One outlier has been detected for both the distance and the time costs, which is Salt Lake City. The much higher than rest variance for the exit gap length ($\bar{\sigma}_{d_t}^2 = 64.007 \text{ km}^2$) and the distance to exits ($\bar{\sigma}_{r_1}^2 = 9.483 \text{ km}^2$; $\bar{\sigma}_{r_2}^2 = 10.290 \text{ km}^2$; $\bar{\sigma}_{r_3}^2 = 10.554 \text{ km}^2$)

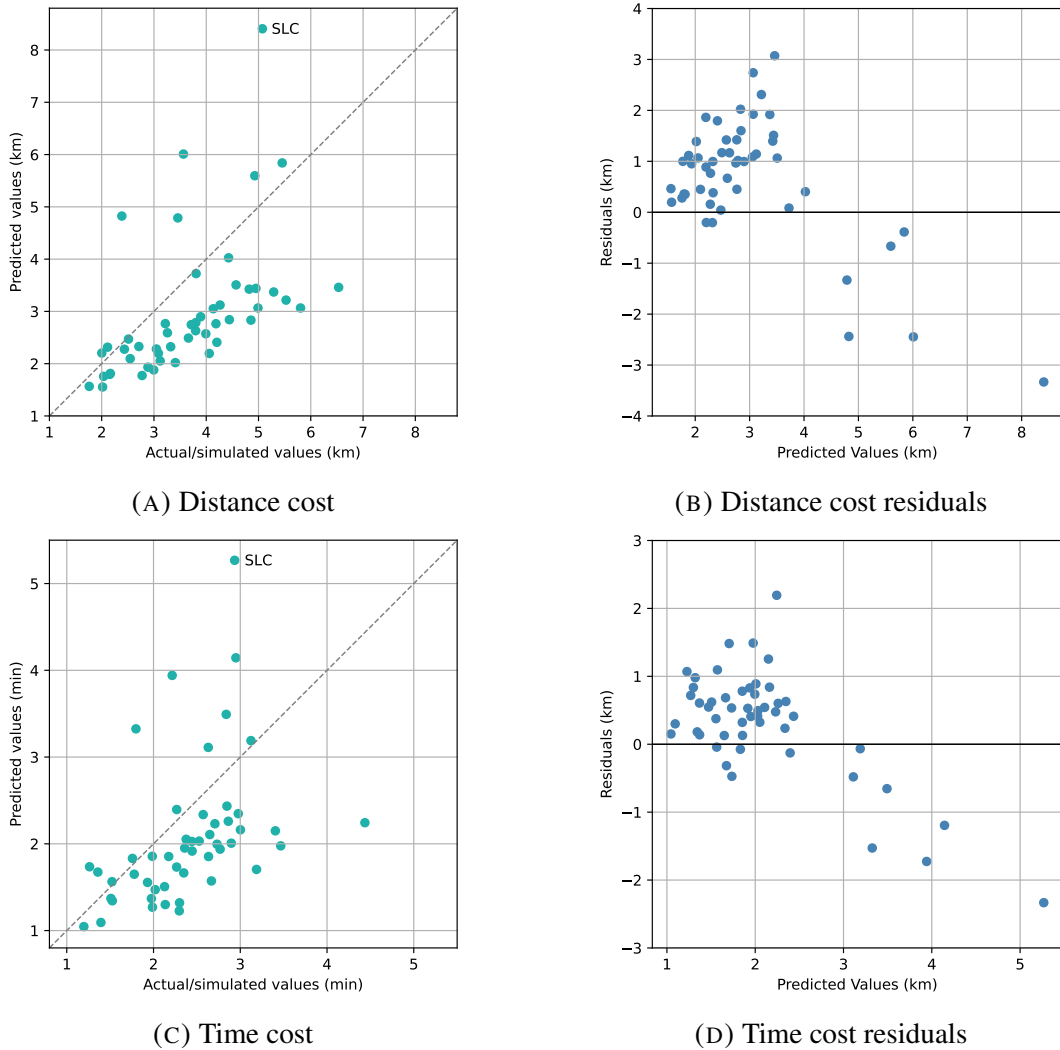


FIGURE 3.5. Scatter plots for the actual and predicted standard deviation of costs. Subplots (a) and (c) are scatter plots showing the relationship between the observed and predicted values. The outlier marked represents Salt Lake City (SLC). Subplots (b) and (d) are residual plots showcasing the distribution of residuals against the predictions. The MAE for the standard deviation of exit-missing distance cost is 1.055 km, and is 0.614 min for the standard deviation of exit-missing time cost.

led to this abnormality. Apart from SLC, the standard deviation of many other metro areas also diverge from the dominant stream. The reasoning is more or less the same: metros with fewer off-ramps have high variation in exit gap length and distance to the nearest exits, which consequently overstates the variance of the functions.

In line with the evaluation of the mean, the RMSE and MAE measures are applied to the predicted standard deviation of the costs. The RMSE and MAE for the standard deviation of the distance costs are 1.257 km and 1.055 km, and are 0.741 min and 0.614 min for time costs. The means of the 50 observed standard deviations are again used as the baseline. The RMSE and MAE for the baseline are computed to be 1.133 km and 0.938 km for distance costs, and 0.628 min and 0.492 min for time costs. This implies that our model underperformed against the baseline. However, the model performance is reasonable, considering that only network variables are fed into the model and that the dispersion of data around the mean is typically harder to predict. Since we have already assessed the impact of the normality assumption on the predictions, it is unlikely that the observed discrepancies stem from this specific assumption. Instead, we suspect the large discrepancy arises from assuming the variables are independent. Without knowledge of the asymptotic distributions of these variables, determining their covariances is challenging. By assuming the variables are independent, only the variance matrix is needed in the calculations. Since we want to find the within-metropolitan variances of the travel costs, it is essential to determine the within-metropolitan covariances of the variables. However, given that the variables were not sampled concurrently and have distinct contexts (e.g., speed was derived from road links, while circuitry was obtained from simulated trips), it is infeasible to compute their covariance. Thus, to quantify the joint variability between these variables, their joint distributions would be necessary, but this information is absent. Consequently, future research should aim to identify and quantify the joint variability between these variables and assess if this contributes to the disparity between predicted and simulated values.

Finally, it is crucial to highlight that the residuals in Figures 3.3b and 3.3d, as well as Figures 3.5b and 3.5d, seem to be non-normally distributed. In fact, the error is decreasing with increasing predicted values. This pattern is likely a result of the inhomogeneous distributions of the exits. Given our HPPP assumption, the analytical model exhibits a strong dependence on the distribution of off-ramps from the networks. The considerable variability in exit density and exit gap length leads the predictions to overshoot the actual values in regions with vast area but sparsely populated off-ramps. This model limitation is elaborated upon in the limitations in Section 3.5.

3.4 Empirical Evidence for Exit Missing Costs

In this section, we analyse GPS datasets to corroborate our prior conclusions that the costs associated with missing off-ramps are not excessively high. Due to the scarcity of real-time vehicle data, our analysis is limited to a single metro area out of the 50 metro areas we have covered. Specifically, we focus our investigation on the Minneapolis-St. Paul metropolitan area. While the LODES dataset only provides commuting OD information, our analysis in this section encompasses both commuting and non-commuting trips, intending to identify any disparities in their cost distributions. The relationship between peak hour congestion and exit-missing costs is likewise assessed. Apart from the evaluation of time and distance costs, we also incorporate the computation of circuitry and a deviation measure, which serves to quantify the spatial dissimilarity between the two routes. These measurements help us explore the consequences of missing freeway off-ramps and are discussed in Appendix A2. Furthermore, in Appendix A3, we conduct a categorical analysis on the exiting behaviours detected in the two datasets. A simple heuristic is introduced to extract trips that possibly missed their intended off-ramps.

Below, we set out a list of hypotheses that are of interest to our study:

- (1) Missing freeway exits will incur additional travel time and distance. However, from the results in Section 3.3.2, we hypothesise that the exit-missing costs are low (only a matter of minutes for time and kilometres for distance);
- (2) The cost distribution of missing an exit for commuting trips and non-commuting trips are not significantly different from the supply/network perspective;
- (3) The costs for missing exits depend on the traffic conditions. In peak hour periods, we presume the costs are more severe, especially for the time cost, since congestion may further delay the travel time;
- (4) There exists a positive relationship between deviation and exit-missing costs. The larger the deviation (spatial dissimilarity) between the original and alternative routes, the higher the additional travel time or distance (discussed in Appendix A2);

- (5) The mean circuitry of alternate routes is anticipated to surpass that of the original routes due to excess rerouting (discussed in Appendix A2);

3.4.1 Method for Empirical Analysis

3.4.1.1 Datasets and Pre-processing

Two datasets were employed in our study: the I-35W dataset and the Travel Behaviour Inventory (TBI) dataset. The I-35W data were collected for part of a traffic behaviour research on the resilience of traffic systems when a detrimental event occurs, in this case before and after the reopening of the I-35W Mississippi River Bridge, which collapsed in 2007 (Zhu et al., 2010). The GPS data were collected for up to 13 weeks in 2008/2009. The participants were required to install either a logging device (Otrek) or a real-time communication device (VMT) on their cars. Out of 190 subjects who participated, 143 records were retrieved and utilised for our study. 47 of them (VMT) were recorded using an installed GPS unit with a frequency of one point per second, while the rest 96 (Otrek) were recorded using a portable GPS unit plugged into an in-vehicle power source, with a frequency of one point every 25 metres.

The TBI dataset was collected in 2010/2011 by the Metropolitan Council of the Twin Cities. GPS devices were deployed to 278 individuals from 250 households, who carried them on their person for all of their travel, including both in and outside of cars. Consequently, prior to undertaking other pre-processing procedures, the TBI data necessitated trip stratification, mode classification, and trip purpose identification. We modified the simple set of rules designed by (Tang and Levinson, 2018) to perform the above three tasks. The mode identification process is based on studies by (Chen et al., 2010, Gong et al., 2012). The trips were first stratified by placing a 300-second threshold on the time gaps of successive GPS points. Then, the car mode was classified with the following rules:

- Average speed of all points > 10 km/hr, which precludes most walking trips.
- Maximum speed of all points > 30 km/hr, which precludes most cycling trips.

- If both the first and last points with speed ≥ 10 km/hr lie within 50 m of bus stops, the trip is considered in bus mode and removed.
- If both the first and last points with speed ≥ 10 km/hr lie within 150 m of rail stations, the trip is considered in rail mode and removed.

Finally, trip purposes are identified by matching the relative location of origin and destination with the known home and work location of the subject. A 500-metre buffer region is used to account for parking and GPS error. The applied process exclusively identifies car trips, wherein these trips are subsequently classified into either commuting or non-commuting categories.

Although the datasets are of high resolution, some data-cleaning steps are still necessary. The data first undergo filtering to eliminate a set of beginning and ending points to address the so-called warm-start-cold-start problem and the excess dwell time at the destination. A trip is updated with its new origin, the first point with distance from the original origin ≥ 50 m, and the updated destination, the first point with distance from the original destination ≤ 50 m. Afterwards, a rough location screening is performed by fitting the trips to the Twin Cities region. In addition, we address potential positional jumps in the data by incorporating a small elliptical zone for each GPS point using their respective coordinate, speed, heading, and a 50 m error buffer. Points outside their neighbours' zones are recognised as off-route points and removed. Abnormal speed changes, defined with an instantaneous acceleration of $\pm 5 \text{ ms}^{-2}$, are also identified and removed from the data. Another common type of speed anomaly, caused by the interference of the satellite signals, usually emerges in the form of a sudden drop of velocity to zero along with linearised trajectory points. These points are likewise removed by tracking the speed change of the vehicles and identifying obstructions, primarily tunnels, in the network. Finally, the missing gaps are filled with interpolated points to construct continuous, connected trips.

Furthermore, we preclude detour routes which may cause misleading outputs. Such excess routing may be undertaken to facilitate the pick-up or drop-off of passengers, resulting in original travel expenses that are well beyond the shortest path costs. Hence, the alternative route computed by the shortest path algorithm can have substantially shorter travel costs than

TABLE 3.2. Summary of the data pre-processing steps. VMT and Otreco together contribute to the I-35W dataset. For each dataset, both the numbers of trips and points are showcased, as well as the original and final data size. The Otreco data do not contain speed attributes; thus, the speed filter is omitted.

	VMT		Otreco		TBI	
	Trips	Points	Trips	Points	Trips	Points
Original	32,730	18,334,840	16,658	7,344,553	3,470	8,721,622
Warm/cold start	-3,906	-3,500,813	-637	-512,852	-198	-3,841,283
Boundary elimination	-730	-1,001,889	-7	-4,654	-192	-969,389
Short trips	-13,630	-413,854	-2,046	-105,458	-117	-6,102
Position jump	0	-937	0	-964	0	-53
Abnormal speed	0	-169,272	-	-	0	-3,879
Fill missing	0	+170,653	0	+698	0	+340,538
Circuitry	-1,653	-2,492,702	-1,191	-542,778	-296	-922,571
Final	12,811	10,926,026	12,777	6,178,545	2,667	3,318,883

the original route, leading to savings in time and distance by missing an exit. In response, we apply the network circuitry to remove trips based on their directness. The top 10th percentile of all trips is removed with an upper bound threshold of 1.84 for the I-35W dataset and 2.25 for the TBI dataset.

Finally, the processed trips are matched with the 2013 OSM network, the earliest and most compatible open-source map regarding the data collection time frame. As one of the test variables for the dataset, the I-35W bridge is removed from the network before the reopening date of 18th September 2008. The map-matching algorithm employed was developed using the hidden Markov model and a set of precomputed shortest paths (Yang and Gidófalvi, 2018). This algorithm takes into account connectivity and adjacency between consecutive points. The links included in the precomputation were within 3 km from the points, thus significantly increasing the computational speed of the matching process when querying from the database.

The above procedure is applied to I-35W and TBI datasets with the carry-over data analysed to test the hypotheses. Table 3.2 summarises trips and points removed in each pre-processing step.

3.4.1.2 Estimation of Exit-Missing Costs

To compute the travel time on specific links, we utilise the TomTom Twin Cities dataset, which contains link-wise speed information from 2010 that has been preserved in databases categorised by the time of day. These data attributes are transferred to the high-resolution OSM network, which had previously been employed for map-matching purposes. Due to the mismatching links for the two road networks, we assign the speed attributes through the largest overlapping proportion between the TomTom network and an auxiliary OSM buffer. The match yields a 95.9% coverage, which implies that only 4.1% of the OSM network is not adapted with a link speed. These remaining links mainly comprise minor local roads, and their speeds are extrapolated using the mean of the three nearest neighbouring edges. To provide a fair comparison between the intended trajectory and the best alternative, we calculate the trip travel times based solely on the network edges matched to the GPS points. It should be noted that the TomTom network takes into account factors such as turns, congestion, and traffic signals implicitly in its speed data. Moreover, we make the assumption that drivers travel at the median (50th percentile) speed for their respective time of day.

The approach of finding the exit-missing costs is similar to that introduced in Section 3.3.1. Based on the GPS dataset and road network, we identify the freeway off-ramps the participants used to exit freeways. To mimic the behaviour of missing an exit, Dijkstra's algorithm is then performed from the source node of the second exit to the final point mapped to the network. Link impedances for shortest path calculations are based on the travel time of each link, which is estimated using the link length divided by the TomTom link speed. The shortest path from the second exit, together with the additional segment on the freeway, constitute the alternative route. Time and distance differences between the original route and its alternative are used to quantify the costs of missing exit.

3.4.2 Empirical Results

This section presents the empirical results from I-35W and TBI datasets and provides a comprehensive discussion of the attributes of exiting activities. Our empirical analysis

TABLE 3.3. Number of exits extracted from datasets. VMT and Otrec together contribute to the I-35W dataset.

	Final		Interim		All	
	Commute	Non-commute	Commute	Non-commute	Commute	Non-commute
VMT	934	5,282	723	3,861	1,657	9,143
Otrec	1,777	4,336	1,486	3,370	3,263	7,706
TBI	177	1,063	137	1,026	314	2,089

indicates that the costs associated with missing exits are consistent with the analytical and simulated values we have obtained. Collectively, these findings suggest that the exit-missing costs are on the order of a few minutes or kilometres. The number of outputted samples from the two empirical datasets are summarised in Table 3.3. A total of 24,002 exiting behaviours are observed, and the spatial distribution of these exits is illustrated in Appendix A1.

3.4.2.1 Costs of Missing Exit

We start by discussing whether missing exits incur additional travel costs for drivers and, if so, how costly missing exits are. Costs are evaluated in terms of additional time and distance from the difference between the original route and its alternative. Computation of the costs is discussed in Section 3.4.1.2. We analyse two scenarios in this study. The first comprises all exits used by the drivers. This approach can be somewhat biased since a human-chosen route (original) is compared with a shortest time path (alternative), which may result in considerably large time/distance savings (negative costs) when an exit is missed. Therefore, we adopt another scenario (named with SP constraint hereafter), which consists of exits with their remaining portion of the trip (from the upstream node of CE to the destination) within 10% of the shortest path travel time. By applying the shortest path filter, we are able to ensure that the original and alternative routes are comparable.

Figure 3.6 presents distributions of the time costs of missing exits for exits in I-35W and TBI datasets. Two exiting scenarios are showcased for the costs. The first includes all exits within each trip, which consists of both final and interim exits. The second scenario involves only the final exits of each trip. From Figure 3.6, density plots without the SP constraint all exhibit

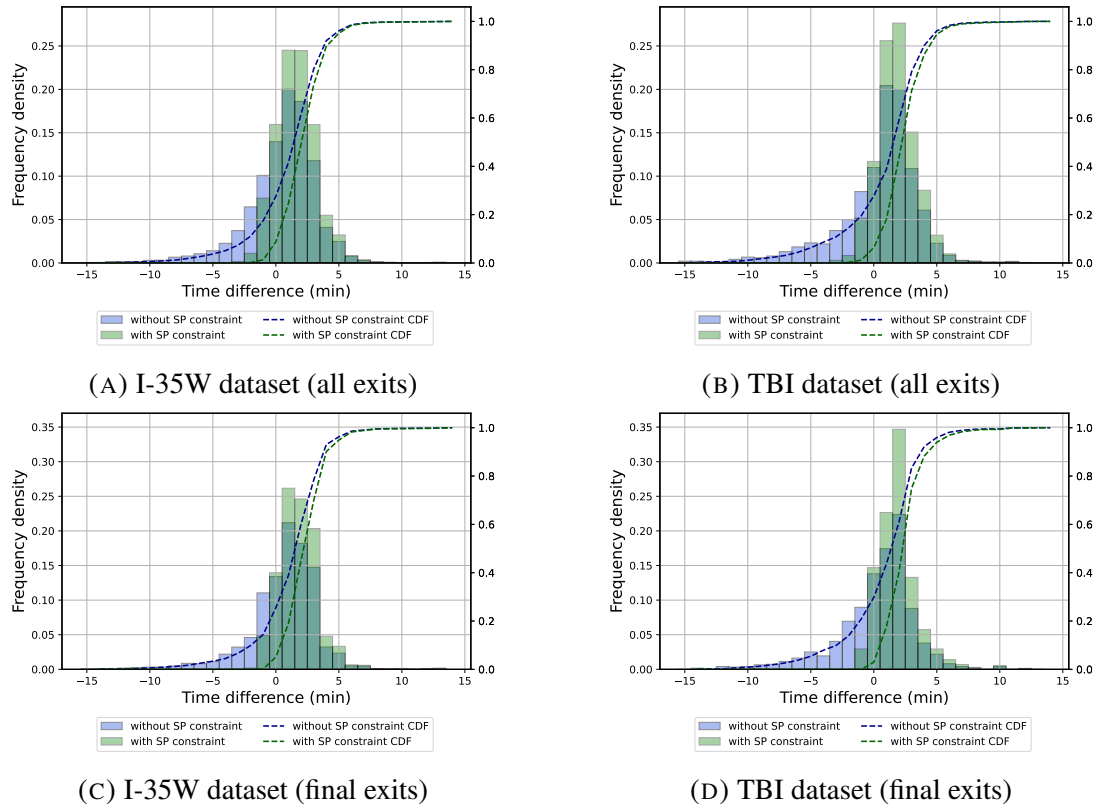


FIGURE 3.6. Distributions of the time cost of missing exits and corresponding CDFs. Cost is calculated by subtracting the travel time of the alternative route from the travel time of the original route. Blue plots correspond to samples without the SP constraint in place. Green plots correspond to samples with the 10% SP constraint.

TABLE 3.4. Summary of the time cost of missing exits. (** $p \leq 0.01$, * $p \leq 0.05$)

	I-35W		TBI	
	without SP constraint	with SP constraint	without SP constraint	with SP constraint
All exits				
Mean (min)	0.479	1.612	0.163	1.833
Median (min)	1	2	1	2
Std (min)	2.931	1.704	4.129	1.683
Min (min)	-24	-8	-49	-31
Max (min)	14	14	12	12
Sample size	21,769	15,727	2,403	1,625
p-value	0.000e-30**	0.000e-30**	0.027*	0.000e-30**
Final exits				
Mean (min)	0.636	1.808	0.275	1.894
Median (min)	1	2	1	2
Std (min)	2.824	1.631	3.161	1.642
Min (min)	-23	-3	-24	-1
Max (min)	14	14	12	12
Sample size	12,329	7,936	1,240	715
p-value	0.000e-30**	0.000e-30**	1.118e-3**	0.000e-30**

left skewness, indicating that a significant proportion of the trips deviate from the shortest path to such an extent that the variation is even greater than that resulting from a missed exit. This has an adverse effect on the accuracy of our analysis and justifies the introduction of the shortest path filter. However, even without the SP constraint, both the mean and the median of all four cases still lie above 0. Table 3.4 summarises the statistics for distributions of the time cost of missing exits and their respective p-values. Here we applied the one-sample t-test with the null hypothesis: the mean of the additional travel time from missing exits is less than or equal to 0. All p-values are statistically significant, which provides evidence for us to reject the null hypothesis in favour of its alternative. Although we confirm missing an exit does incur additional travel time, the magnitude of the increase is generally mild.

Moreover, we compare the empirically derived results with those from our analytical derivations and simulations. To maintain consistency, only the final exits are considered in the comparison. Combining the I-35W and TBI samples yields a mean of 1.815 min and a standard deviation of 1.632 min. Based on the combined mean, there are, respectively, 9% and 14% over-estimations for the analytical and simulated results (1.984 min and 2.077 min). Similarly, the combined standard deviations is compared to the analytical and simulated values (1.472 min and 2.019 min). There is a 10% under-estimation and a 24% over-estimation for the analytical and simulated results. Although the error may seem large, it is worth noting that we are approximating complicated exit-missing costs using analytical and simulated models laden with simplifications and assumptions. Notably, the analytical model employs only five network variables to estimate these costs. Most importantly, the observations from all three approaches show that people tend to spend more time after missing an exit, yet the time costs are, on average, relatively low (approximately 2 min for Twin Cities). Therefore, lane changes intended to exit a freeway may not be strictly required if no opportunity for a safe lane change prior to the exit emerges. Drivers are thus resilient to missing exits.

The extra travel distance incurred by missing an exit is further evaluated. Figure 3.7 presents the travel distance costs of missing exits for the two datasets, while the statistics of the distributions are summarised in Table 3.5. We hypothesise that by missing an exit, one would be penalised with extra travel distance. Thus, for our null hypothesis, we assume that the

mean of the additional travel distance from missing exits is less than or equal to 0. The alternative hypothesis states otherwise. Once again, the p-values favour the alternative, and we conclude that a driver would be subjected to extra travel distance when missing an exit.

Likewise, we evaluate the discrepancies between empirical results and those derived from analytical and simulation methods for the distance cost. By combining the two datasets, the mean and standard deviation are 3.068 km and 2.272 km, respectively. Recall that the mean distance costs from the analytical and simulated means are 2.634 km and 3.145 km, while the standard deviations are 2.020 km and 3.409 km. The combined empirical means are therefore, 14% under-estimated by the analytical value and 3% over-estimated by the simulated value. The standard deviations, on the other hand, are 11% under-estimated and 50% over-estimated by the analytical and simulated methods. From the above analysis, it is clear that while the estimated means serve as reasonable approximations, the standard deviations diverge notably. In particular, the simulated standard deviation tends to surpass the empirical value by a margin. The large difference may be a result of the higher variability in the sampled data. Since the LEHD data cover large amounts of commuting ODs from each metro area, the variability between routes may be substantially different. This will in turn magnify the variance of the exit-missing costs. Nevertheless, the observations underscores that the exit-missing costs are not substantial.

3.4.2.2 Analysis on Relative Proportions

In general, evaluating whether an increase in travel time/distance is relevant for the drivers would require a discussion on what proportion of the total trip this represents. Therefore, we further evaluate the exit missing costs in terms of the relative proportions of the costs to their original trips.

Figure 3.8 illustrates the time costs as a percentage of the original trip length, while Table 3.6 summarizes the statistics. Continuing from the t-tests performed for the absolute costs, in this section, we examine whether the mean percentage costs are also positive. The null hypothesis is set such that the mean values are assumed to be 0. The p-values for all scenarios are significant, indicating that, on average, the percentage time costs are not 0. In particular, for

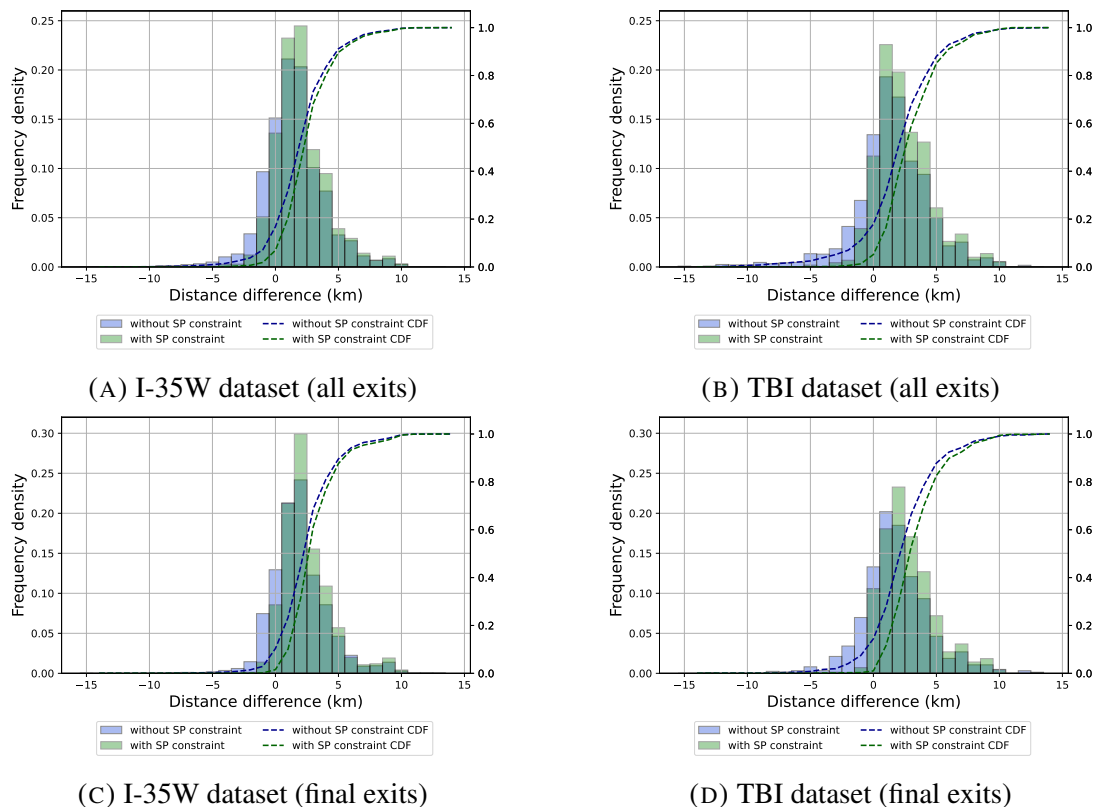


FIGURE 3.7. Distributions of the distance cost of missing exits and corresponding CDFs. Cost is calculated by subtracting the distance of the alternative route from the distance of the original route. Blue plots correspond to samples without the SP constraint in place. Green plots correspond to samples with the 10% SP constraint.

TABLE 3.5. Summary of the distance cost of missing exits. (** $p \leq 0.01$, * $p \leq 0.05$)

	I-35W		TBI		
	without SP constraint	with SP constraint	without SP constraint	with SP constraint	
All exits	Mean (km)	1.958	2.536	1.879	2.978
	Median (km)	1.900	2.218	1.967	2.582
	Std (km)	2.636	2.278	3.811	3.493
	Min (km)	-20.270	-16.548	-43.587	-4.677
	Max (km)	23.802	23.802	23.326	23.326
	Sample size	21,769	15,727	2,403	1,625
	p-value	0.000e-30**	0.000e-30**	0.000e-30**	0.000e-30**
Final exits	Mean (km)	2.452	3.041	2.382	3.370
	Median (km)	2.199	2.522	2.132	2.863
	Std (km)	2.479	2.235	2.949	2.630
	Min (km)	-36.717	-2.872	-12.249	-4.504
	Max (km)	23.697	23.697	23.164	23.164
	Sample size	12,329	7,936	1,240	715
	p-value	0.000e-30**	0.000e-30**	0.000e-30**	0.000e-30**

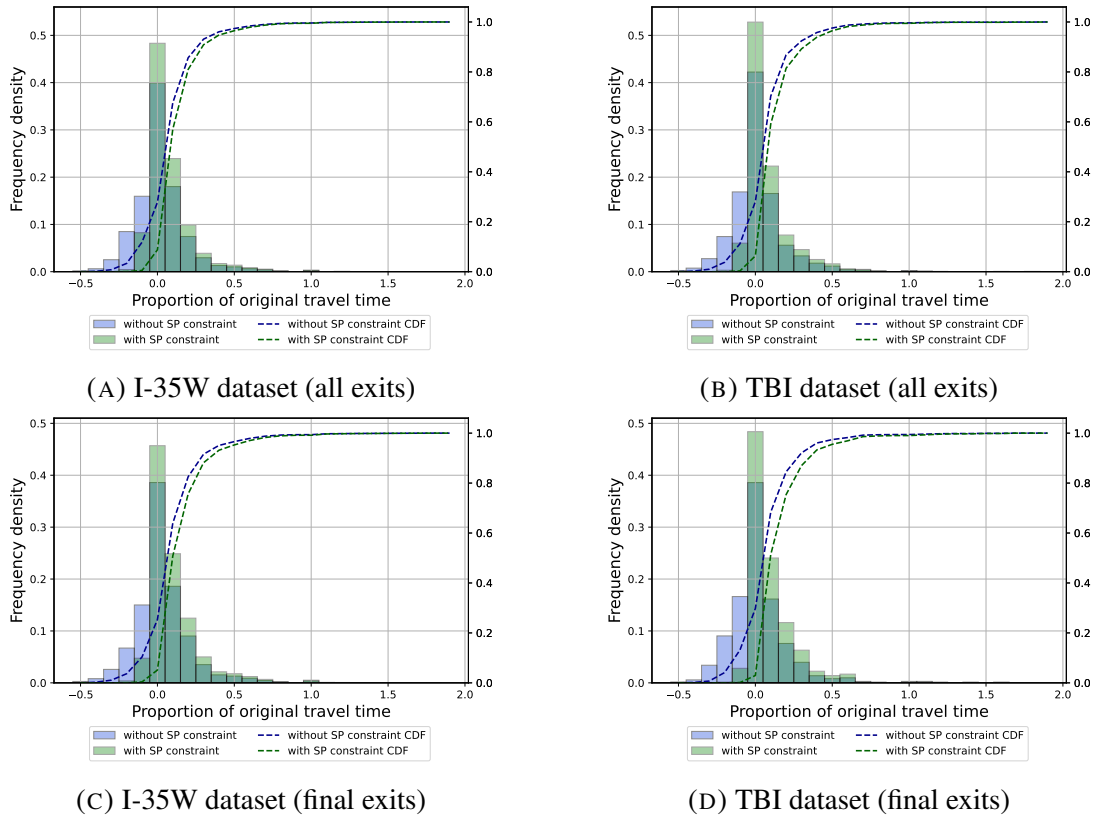


FIGURE 3.8. Distributions of the time cost of missing exits and corresponding CDFs. Costs are in percentage of the original trip length.

TABLE 3.6. Summary of the time cost of missing exits in terms of the proportion of the total trip. (** $p \leq 0.01$, * $p \leq 0.05$)

	I-35W		TBI		
	without SP constraint	with SP constraint	without SP constraint	with SP constraint	
All exits	Mean (%)	0.0637	0.1194	0.0523	0.1201
	Median (%)	0.0465	0.0769	0.0444	0.0769
	Std (%)	0.2058	0.2013	0.2079	0.1594
	Min (%)	-1.7143	-0.8182	-2.7857	-0.1667
	Max (%)	4.6667	4.6667	1.8000	1.8000
	Sample size	21,769	15,727	2,403	1,625
	p-value	0.000e-30**	0.000e-30**	0.000e-30**	0.000e-30**
Final exits	Mean (%)	0.0782	0.1476	0.0592	0.1427
	Median (%)	0.0526	0.0952	0.0400	0.0952
	Std (%)	0.2377	0.2417	0.1836	0.1776
	Min (%)	-1.0000	-0.2500	-0.5385	-0.1250
	Max (%)	4.3333	4.3333	1.6000	1.6000
	Sample size	12,329	7,936	1,240	715
	p-value	0.000e-30**	0.000e-30**	0.000e-30**	0.000e-30**

all our scenarios, the costs are positive. When the SP constraint is in place, the percentage time costs are approximately 12% for all exits and 14% for the final exits. Unsurprisingly,

without the SP constraint, a portion of the distributions shifts to the negative range, and the mean proportional time costs drop.

As for the percentage distance costs, the distributions are illustrated in Figure 3.9, and the statistics are presented in Table 3.7. Again, the null hypothesis is set such that the mean values are 0. All p-values suggest that the distance percentage costs relative to the original trips are positive. The magnitudes of the costs with SP constraints are approximately 17% for all exits and 23% for final exits. The results above suggest that the percentage distance costs are typically larger than the percentage time costs. We speculate that this observation might be due to vehicles having to travel further down the freeway after missing an exit. The freeway section, which has much higher traffic speeds than local streets, then understates the marginal impact on travel time.

Furthermore, both the figures and tables show that the percentage costs are fairly consistent for the two datasets, suggesting a strong dependence of the exit-missing costs to the network in which the trips belong to. From both absolute and percentage costs, it can be seen that the cost tend to be higher for final exits than all exits. The fact that interim exits have higher flexibility for future route choice might have caused this phenomenon. The earlier an exit is missed, we believe the easier it is for the driver to adjust its route to make up for the costs.

It can also be seen that the distributions with the SP constraint have a larger portion in the negative plane. However, instead of skewing to the left, the plot now skews to the right, indicating that trips with negative costs typically have high travel time or distance that forces the percentage costs low. Due to the invalidity of comparing excess routing original trips and shortest path alternatives, analysis hereafter are performed with the SP constraint in place.

3.4.2.3 Analysis on Commuting Effect

Unlike the data obtained from the LEHD dataset, GPS data used in this study consists of both commuting and non-commuting trips, providing an opportunity to examine the impact of trip purposes on exit-missing costs. Here, we only assess the “SP-constrained” scenario. The two-sample independent t-test is applied, assuming equal variances. The null hypothesis states




FIGURE 3.9. Distributions of the distance cost of missing exits and corresponding CDFs. Costs are in percentage of the original trip length.

TABLE 3.7. Summary of the distance cost of missing exits in terms of the proportion of the total trip. (** $p \leq 0.01$, * $p \leq 0.05$)

	I-35W		TBI		
	without SP constraint	with SP constraint	without SP constraint	with SP constraint	
All exits	Mean (%)	0.1372	0.1743	0.1231	0.1736
	Median (%)	0.0898	0.1130	0.0839	0.1166
	Std (%)	0.2270	0.2379	0.2039	0.1939
	Min (%)	-0.7571	-0.7571	-1.8837	-0.1250
	Max (%)	4.4029	4.4029	2.0158	2.0158
	Sample size	21,769	15,727	2,403	1,625
	p-value	0.000e-30**	0.000e-30**	0.000e-30**	0.000e-30**
Final exits	Mean (%)	0.1825	0.2313	0.1540	0.2243
	Median (%)	0.1163	0.2313	0.1022	0.1740
	Std (%)	0.2613	0.2842	0.2052	0.2243
	Min (%)	-0.3694	-0.0497	-0.3184	-0.0719
	Max (%)	4.3433	4.3433	2.0009	2.0009
	Sample size	12,329	7,936	1,240	715
	p-value	0.000e-30**	0.000e-30**	0.000e-30**	0.000e-30**

that the expected values for the costs of missing exits for commuting and non-commuting trips are the same. The alternative states otherwise.

As above, both time and distance costs are examined. Figures 3.10 and 3.11 demonstrate the cost distributions for commuting and non-commuting trips, and their statistics are summarised in Tables 3.8 and 3.9. Interestingly, all p-values for the final exit cases are significant, and in general, the average costs for commuting trips are larger than the average for non-commuting trips. This opposes our initial hypothesis and suggests that the costs of missing an exit are somehow related to the purpose of the trip. The cases involving all exits are not statistically significant due to the greater flexibility after missing an interim exit. In other words, when drivers miss an interim exit, they have more alternative routes available to them than if they missed the final exit. This flexibility allows drivers to more easily counteract the negative effects of missing an exit, which could downplay the impact of different OD locations, leading to seemingly consistent comparison outcomes.

An inspection of the trips containing final exits shows that the average distances between the final exit and the destination for commuting trips (2.992 km) and non-commuting trips (2.970 km) are barely different. Therefore, the flexibility of route choice is similar, and the discrepancy is likely related to the length of rerouting. The deviation factors are computed for both the original paths and the shortest paths (from the exit to the destination). Commuting trips are seen to have a higher average deviation (2.787 km) than non-commuting trips (1.720 km). The difference in the costs thus could not have been caused by a higher sensitivity of the commuting trips towards missing exits due to them travelling more frequently on the shortest path. Therefore, the cost discrepancy is likely related to the network configuration.

3.4.2.4 Analysis on Peak Hour Effect

Considering that traffic congestion may play a crucial role in the composition of exit-missing costs, we examine the effects of travelling during peak hours. Again, we focus solely on the “SP-constrained” scenario. The analysis done in this section divides the trips into peak hour and off-peak hour ones based on the time that a trip was performed. A trip is deemed to have experienced peak hour effects if its entire duration falls within one of the following peak periods: morning peak (7-9 am) or afternoon peak (4-6 pm). These peak hours are considered only on weekdays.

TABLE 3.8. Summary table for the time cost of missing exits for commuting and non-commuting trips. (** $p \leq 0.01$, * $p \leq 0.05$)

		I-35W		TBI	
		All exits	Final exits	All exits	Final exits
Commuting	Mean (min)	1.599	1.957	2.016	2.167
	Median (min)	1	2	2	2
	Std (min)	1.759	1.819	1.785	1.523
	Sample size	3,510	1,688	189	168
Non-commuting	Mean (min)	1.616	1.768	1.808	1.857
	Median (min)	2	2	2	2
	Std (min)	1.688	1.575	1.669	1.653
	Sample size	12,217	6,248	1,436	1,262
p-value		0.595	2.548e-5**	0.111	0.022*

TABLE 3.9. Summary table for the distance cost of missing exits for commuting and non-commuting trips. (** $p \leq 0.01$, * $p \leq 0.05$)

		I-35W		TBI	
		All exits	Final exits	All exits	Final exits
Commuting	Mean (km)	2.472	3.224	3.221	3.763
	Median (km)	2.170	2.436	2.314	3.122
	Std (km)	2.448	2.655	2.743	2.464
	Sample size	3,510	1,688	189	168
Non-commuting	Mean (km)	2.554	2.991	2.947	3.318
	Median (km)	2.231	2.558	2.608	2.760
	Std (km)	2.227	2.105	2.458	2.646
	Sample size	12,217	6,248	1,436	1,262
p-value		0.059	1.444e-4**	0.155	0.039*

Figures 3.12 and 3.13 illustrate the time and distance cost distributions, taking into account the effects of peak hours. Tables 3.10 and 3.11 present the statistics corresponding to the two figures. The t-tests were conducted with the null hypothesis that the costs during peak and off-peak hours are equivalent. Notably, none of the distance costs exhibit a significant p-value, indicating that the distance costs from missing exits during peak and off-peak periods tend not to differ. On the other hand, a majority of the time costs are shown to be statistically significant. Given the consistency in distance costs, the variation in time costs might be explained by the deteriorating traffic condition during the peak hours. This is further corroborated by the higher average costs observed during peak hour instances. Although the p-value is not significant for the final exits of the TBI dataset, the peak hour mean remains considerably higher than

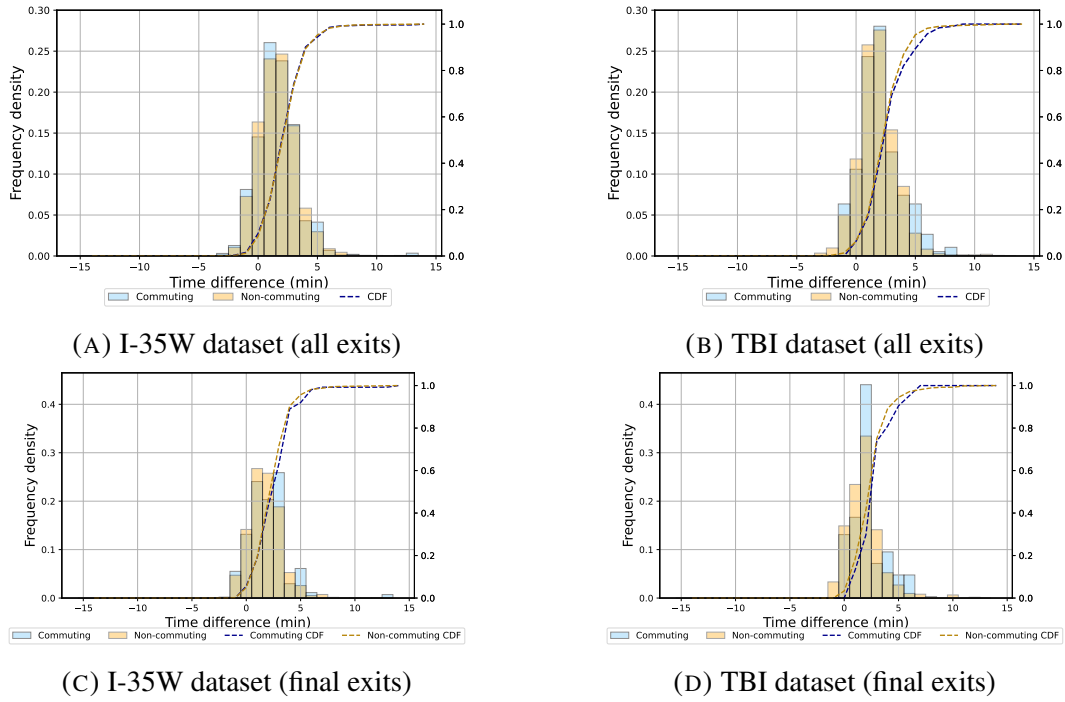


FIGURE 3.10. Distributions for the time cost of missing exits in commuting and non-commuting trips and the corresponding CDFs.

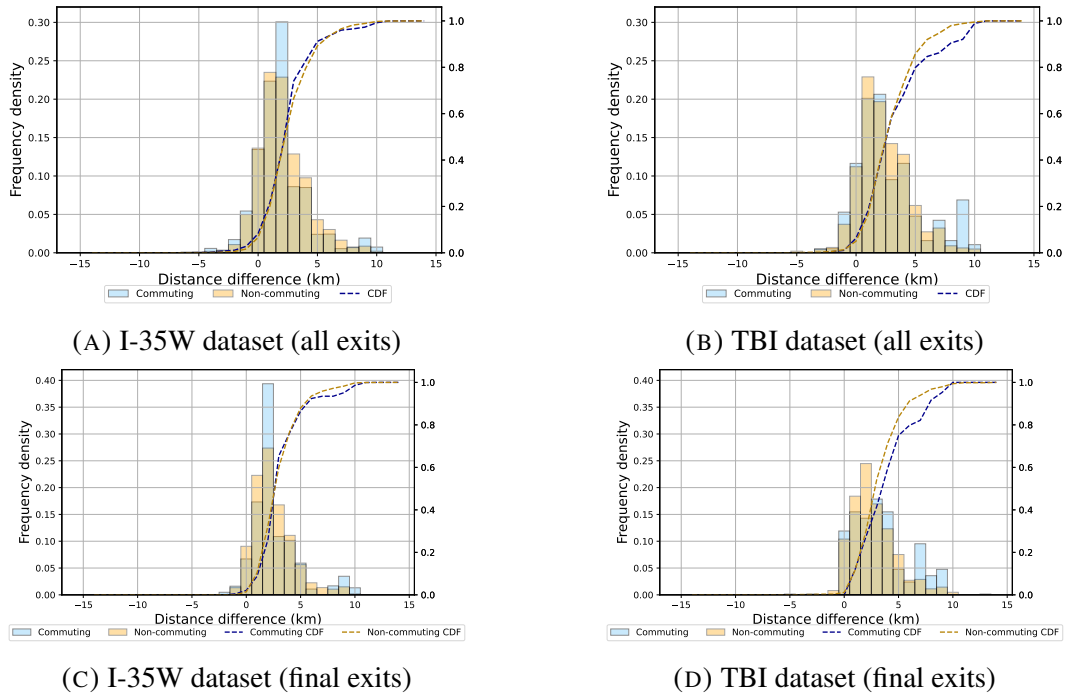


FIGURE 3.11. Distributions for the distance cost of missing exits in commuting and non-commuting trips and the corresponding CDFs.

TABLE 3.10. Summary table for the time cost of missing exits for peak hour and off-peak hour trips. (** $p \leq 0.01$, * $p \leq 0.05$)

		I-35W		TBI	
		All exits	Final exits	All exits	Final exits
Peak hour	Mean (min)	1.677	1.944	2.084	2.1468
	Median (min)	2	2	2	2
	Std (min)	1.756	1.721	1.923	1.7786
	Sample size	4,538	2,374	238	109
Off-peak hour	Mean (min)	1.586	1.751	1.789	1.8482
	Median (min)	1	2	2	2
	Std (min)	1.682	1.588	1.636	1.6130
	Sample size	8,486	5,562	1,387	606
p-value		2.536e-3**	1.421e-6**	0.013*	0.080

TABLE 3.11. Summary table for the distance cost of missing exits for peak hour and off-peak hour trips. (** $p \leq 0.01$, * $p \leq 0.05$)

		I-35W		TBI	
		All exits	Final exits	All exits	Final exits
Peak hour	Mean (km)	2.512	3.057	2.864	3.226
	Median (km)	2.190	2.553	2.767	2.906
	Std (km)	2.311	2.247	2.165	2.085
	Sample size	4,538	2,374	170	109
Off-peak hour	Mean (km)	2.545	3.034	2.992	3.396
	Median (km)	2.220	2.522	2.539	2.845
	Std (km)	2.265	2.230	2.529	2.717
	Sample size	11,189	5,562	1,455	606
p-value		0.411	0.674	0.527	0.534

the off-peak mean. The insignificant p-value might be attributed to a smaller sample size compared to other cases.

From the figures, particularly the CDFs, it is evident that the time cost distributions for peak hour trips are generally more platykurtic. This trend arises because alternative routes that are longer than the original ones, leading to positive exit-missing costs, experience increased time costs due to slower traffic. Similarly, original routes that are longer than their alternative counterparts, resulting in negative costs, also have their travel time amplified, which further decreases the time cost (resulting in more negative values). This rationale underscores our decision to focus solely on the SP-constrained cases. The unconstrained distributions, with a




FIGURE 3.12. Distributions for the time cost of missing exits in peak hour and off-peak hour trips and the corresponding CDFs.

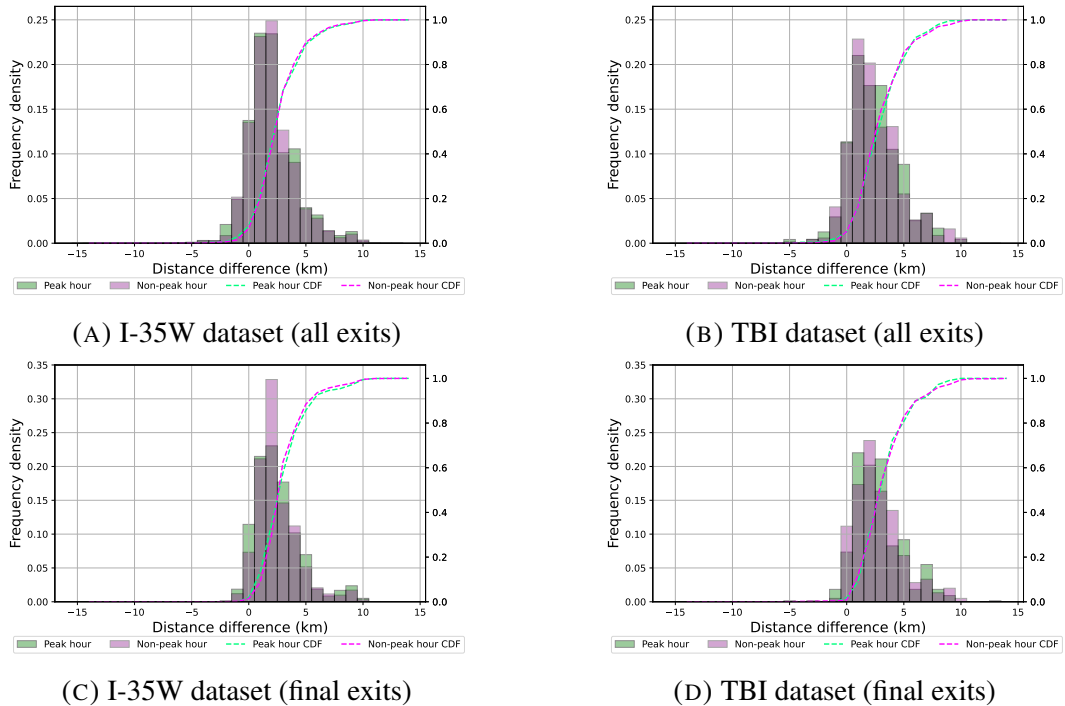


FIGURE 3.13. Distributions for the distance cost of missing exits in peak hour and off-peak hour trips and the corresponding CDFs.

significant portion below 0, would yield more negative costs during peak hours, potentially skewing the results. The SP cases are thus more compelling given that the two routes are more comparable. Thus, the empirical results identify the presence of traffic congestion during peak hours. The lowered speed during peak hours tends to magnify the time cost of missing exits, whereas the distance costs remain relatively unchanged.

3.5 Conclusions and Limitations

This study ascertains the costs of missing exits, thereby exploring the benefits of performing lane changes for exiting freeways. First, an analytical model is proposed, whereby the exit-missing costs are deconstructed into sub-components with their values then estimated. Both the mean and standard deviation of the costs can be approximated. The proposed analytical methodology is then applied to the top 50 metropolitan regions in US. and the outcomes are compared with simulated values. The results demonstrate that the proposed analytical approach is practicable and can provide valuable insights into the nature of costs associated with missed exits. Specifically, networks with lower circuitry, faster freeway speeds, and faster local street speeds often have higher costs associated with missing exits. Furthermore, if a network has sparser exit density and greater distances between exits, the costs for missing an exit tend to be high. However, inhomogeneous exit distributions can cause the prediction to overshoot the actual costs by a margin. Therefore, the model would perform satisfactorily in areas where exits are more homogeneous, such as in urbanised areas. Subsequently, we conduct a thorough analysis of two real-life empirical datasets obtained from the Minneapolis - St. Paul region. The results from both datasets are similar and also correspond with the analytical and simulated outcomes. All three analyses indicate that although missing exits can lead to higher travel costs, the increase is negligible, amounting to only a few minutes or kilometres.

These findings shed light on the costs associated with not executing a lane change to exit off-ramps. When lane changes are performed to avoid obstacles, the costs are essentially infinite, given that the manoeuvre prevents the driver from being indefinitely stuck. In contrast, the

costs tied to not changing lanes for exiting freeways are notably lower. This discrepancy might partially explain why drivers occasionally miss their intended turns. As such, we recommend more lane change models to incorporate the probability of drivers missing exits, which could arise from factors such as fatigue, distractions, or underestimations of the complexity of a lane change manoeuvre. For the design of lane change controllers, our study can assist in balancing the trade-off between executing a lane change and enduring a longer route to the destination.

Given the distinct benefits of making a lane change to exit a freeway versus doing so to avoid blockage, we categorise the mandatoriness of lane changes into three levels: discretionary, expedient (ELC), and mandatory. The respective benefits of these lane changes are on the scale of seconds/metres, minutes/kilometres, and infinite. Based on our findings, we classify lane changes for tactical routing as ELC. Future studies should assess the benefits of other types of MLCs, such as obeying lane usage indications or yielding to emergency vehicles, to determine if they should also be classified as ELCs.

The present study is subject to several limitations, and we will first discuss the ones for the analytical model:

As part of the cost functions, the contribution of exits to the surrounding households (in Section 3.2.1) is formulated based on the assumption that people are rational and fully aware of the travel time to their destinations. In fact, it has been previously demonstrated that drivers are not always rational and often do not travel on the shortest available path ([Dia, 2002](#), [Jiang et al., 2020](#), [Zhu and Levinson, 2015](#)). Therefore, the deterministic proportions we derived may not accurately reflect the true contribution of nearby exits. Since the experiment we conducted in Section 3.3 is based on the shortest path assumption, the potential ramifications of this limitation may not be readily apparent. However, future research should address this shortcoming by incorporating route choice techniques into the model.

It is postulated in the proposed model that exits are characterised by a homogeneous Poisson point process, which represents a simplified approximation of the true distribution. In practice, the design of exit locations would be more dependent on the population of the surrounding

neighbourhood, the spatial distribution of freeways, and the overall accessibility of the region. However, accounting for all possible scenarios would make our model overly complex and would require the computation of input parameters such as the locations of city centres and the density of links on the freeway network.

As demonstrated in Section 3.3, the predicted value exhibits a strong dependence on the number of off-ramps from the street network. In our examples, the high variation in exit density and exit gap length led to the predictions to overshoot the actual values in regions with a larger area and fewer off-ramps. This is essentially a repercussion of the underlying HPPP assumption, which makes the model sensitive to fluctuations in the number of exits in subsets of the regions. Therefore, instead of computing the costs for a large geographical area, our model would perform better for areas where the density of exits is more consistent (e.g. urbanised area instead of the entire metropolitan area).

It should be noted that the findings we have obtained are based on the average costs for an overall network. In dense urban area, with a lot of redundancy, missing an exit would not cost much if there were no congestion. However, in a sparser network, that may not be the case. Examples such as SLC and LAS have sparse networks and are observed to have relatively high costs associated with missing an exit.

Finally, our analytical model fails to take into account supply side influences. As an example, higher flows are observed closer to the city centre than in rural areas. Thus higher weights should be placed on exits in more populated areas where the traffic network is denser and shorter exit-missing costs are expected.

In addition, both the analytical and simulated models have omitted the impact of traffic congestion on the costs. As highlighted in Section 3.4.2.4, heavier traffic is observed to impact exit-missing costs, particularly the time cost. This suggests that further research should be undertaken to incorporate congestion into the cost modelling. However, in the absence of link-wise speed data for different time periods, we will defer this inquiry to future research. Likewise, the fluctuations in traffic flow can influence exit choice of drivers, which could, in turn, impact the exit-missing costs. In particular, a driver may opt for an earlier exit

due to a downstream blockage or choose a later exit due to spillback from the usual off-ramp. In summary, factors such as spillbacks from bottlenecks can greatly affect route choices and consequently, lane-changing decisions and costs. It is challenging to disentangle the effects of congestion from path choice. The impact may be more visible if newer datasets are used, which reflect a wide spread of communication and technology systems where drivers may follow directions from navigation apps. The impact of such path planning behaviour on the costs of missing exits may be explored in future research by comparing old and new travel patterns.

The following are the limitations of our empirical analysis:

- The empirical datasets only cover the Twin Cities region. As a result, it remains unclear whether the same findings can be extrapolated to other regions. It is important to note that the performance of traffic networks is heavily influenced by their structural characteristics (Xie and Levinson, 2007). Therefore, high variability of travel costs may exist for networks with different configurations and road densities.
- The estimation of trips with missed final exits is based solely on a set of simple rules, without knowledge of the actual intentions of drivers. This approach may not accurately capture the complexities of real-world driving behaviours. A more accurate analysis could be achieved through the use of a driving simulation or an in-vehicle dataset that provides detailed vehicle indicator information and tracks driver intent.
- The alternative routes computed are assumed to follow the next downstream exit, which may not be a valid assumption. After missing exits, drivers will likely follow their intuitions to navigate to the destinations, which may not involve the exit directly downstream of the missed one. To address this issue, future research could incorporate data capturing drivers' actual behaviours after missing exits, which would provide a more accurate representation of travel patterns.
- The presence of in-car navigation may also influence the human reaction after an exit is missed. This factor was not directly accounted for in our analysis, since it is unknown which individuals used in-car navigation during the data collection period.

However, we have indirectly considered this by purposefully selecting two datasets that were not collected recently (2008/09 and 2010/11, respectively). At that time, in-car navigation systems were not widely available, and only a fraction of people in the US used such systems regularly ([Leshed et al., 2008](#)). Moreover, Google Maps was released in September 2008, so its influence on the route choice decisions of drivers in the I35-W dataset should be minimal. Regarding the TBI dataset, there might have been a broader adoption of in-car navigation systems by 2010 than in 2008. Nonetheless, based on the shortest path usage analysis in Appendix A3, we believe the impact of GPS navigation on exit choices would not significantly influence the calculated costs. After in-car navigation become wider adopted (in more recent years), we anticipate that there might be changes to the costs of missing exits as people are more likely to travel on the shortest paths, and concern about unfamiliarity routes might diminish.

- By using trajectory data from GPS records, we may have overlooked the influence of human factors on the exit-missing costs, such as variations in the Value of Time (VOT) and drivers' perceived travel costs. VOT differs among individuals and for different trip purposes, making exit-missing costs situation- and individual-dependent. For instance, the costs associated with missing an exit during a leisure trip would differ from those during a trip to a time-sensitive meeting. However, this study focuses solely on objective costs incurred when exits are missed, omitting demand-side factors. To assess the impact of VOT on the cost of missing freeway exits, future studies should investigate the marginal disutility of travel costs for various types of trips and for different driver populations. For research related to commuting VOT, interested readers can refer to ([Carrion and Levinson, 2012](#), [Li et al., 2010](#)).
- Furthermore, the costs perceived by human drivers may differ from the actual costs incurred on the road. Perceived costs can depend on various factors, including the driver's familiarity with the road, the nature of the trip, the driver's expected arrival time, and individual variations. As an example, individuals may encounter discomfort while navigating an unfamiliar route. Missing a frequently used exit

and travelling into an unfamiliar area can significantly impact the urgency of a lane change. In such cases, the lane change may feel “more mandatory” than the objective costs incurred. Extracting this information can be challenging and may require an understanding of the mismatch between drivers’ expected and actual utilities. Future research should explore how mandatory a lane change feels to a driver using datasets beyond trajectory data. Surveys can gauge drivers’ attitudes toward distracted driving and VOT, while driving simulator experiments, or in-vehicle monitoring, coupled with electroencephalogram technologies, may be better suited for investigating the cognitive processes and emotional responses while driving.

As for future research directions, our findings can be applied to model lane changes when dealing with the trade-off between safety and travel time ([Ji and Levinson, 2021](#), [Ji et al., 2023a,b](#)). The analytical model we proposed can be integrated into the cost functions by allowing a certain degree of stochasticity in the merging process. Moreover, our work can provide insight into exiting activities for traffic simulations. Existing microscopic models use rules to model lane changes in preparation for exits. For instance, VISSIM adopts two criteria for exiting: the lane change distance and the emergency stop distance. The lane change distance specifies the distance from which the driver becomes aware of the upcoming exit and will consider it in lane change manoeuvres. The emergency stop distance is the distance from the exit the driver will stop to wait for an acceptable gap to change lanes. On the other hand, AIMSUN uses a look-ahead heuristic, which reduces the chance of a vehicle coming to a stop to exit. However, the same assumption is made that a car will force a gap in the target lane when executing the lane change, which often causes the traffic condition to deteriorate as vehicles build up ([Barceló, 2010](#)). We argue that such an assumption can be relaxed given that the cost of missing an exit is relatively small in most urban cases, and the vehicle will eventually route back to its destination when necessary.

The findings of this study challenge the conventional notion of MLC and suggest a need for researchers to refine existing lane change models.

A Game-theoretic Two-dimensional Driving Model

4.1 Introduction

Microscopic traffic models have seen significant advances in recent years, enabling detailed simulations of individual vehicle behaviours. Existing studies have focused on modelling the longitudinal and lateral movements of vehicles separately (Hou et al., 2014, Kesting et al., 2007, Toledo et al., 2003, Yang and Koutsopoulos, 1996), overlooking the interdependencies between car-following and lane-changing manoeuvres (Toledo, 2007). Lane-changes tend to be treated as discrete processes with many important factors, such as lane-change attempts and durations, neglected (Ji et al., 2023a). Furthermore, only a few have explicitly considered the strategic interactions between human drivers (Ali et al., 2019, Kita, 1999, Liu et al., 2007, Talebpour et al., 2015). That is, the possible actions of other surrounding vehicles are not reasonably accounted for. Finally, existing models tend to treat driver preferences homogeneous despite heterogeneity in human driving (Ji and Levinson, 2021, Ossen and Hoogendoorn, 2011, Zhou et al., 2020).

In light of the previous considerations, this chapter introduces a microscopic trajectory model that addresses the gaps identified above. To account for the interdependence between car-following and lane-changing decisions, our two-dimensional driving model (2DDM) integrates longitudinal acceleration and lateral steering into a single unified objective. The proposed model assumes drivers optimise their actions over a short predictions horizon based on the objective function. To account for the strategic interactions between drivers, we have adopted game theory, which provides a means for drivers to predict actions of surrounding vehicles over the horizon. Considering that individuals perform hierarchical

reasoning processes, we apply level-k game theory to model their strategic interactions while driving. However, due to the short prediction interval, we assume that drivers engage in shallow decision hierarchies, thereby limiting the depth of their reasoning recursions.

The resulting model is then calibrated and validated using a naturalistic trajectory dataset, namely the highD dataset. To capture the heterogeneity between individual drivers, the parameters are calibrated as distributions using a Bayesian calibration procedure. The calibrated distributions are then used to reconstruct the calibrated trajectories as well as two non-calibrated sets of trajectories to validate the model's accuracy and transferability. Results show that the model can reasonably recreate both seen and unseen data.

Lane-changes are typically classified as either mandatory or discretionary. The model in this chapter focuses exclusively on discretionary lane-changes. Vehicles are assumed to travel on a highway segment free from intersections and on/off-ramps, where the goal for each driver is to safely traverse the section in minimal time, performing discretionary lane-changes to achieve this aim.

4.2 Methodology

Our modelling framework employs level-k game theory and receding horizon optimisation. In this framework, each vehicle is modelled to predict potential actions of its opponents (i.e., adjacent vehicles) based on the surrounding environment and uses this information to make decisions that optimise its own objectives. Our methodology relies on three key assumptions:

- **Rationality:** Drivers are boundedly rational and will act to optimise their perceived payoffs based on their beliefs about the reasoning processes of other players.
- **Predictability:** Drivers can foresee a short interval into the future and will utilise this information in their decision-making.
- **Heterogeneity:** Drivers have varying preferences regarding their payoffs.

These assumptions form the foundation of 2DDM. The remainder of this section discusses the formulation of the receding horizon optimisation, the incorporation of level-k game theory, and the numerical method used to solve the entire framework.

4.2.1 Receding Horizon Cost Optimisation

The vehicle's dynamics in 2DDM are modelled using a kinematic bicycle model, a simplified representation that lumps the four wheels of an automobile (two front and two rear) into two central wheels, one at the front axle and one at the rear (Singh et al., 1991). This model describes the vehicle's dynamic state as a four-variable vector $\mathbf{z} = [x, y, v, \psi]'$, where x and y denote coordinates in the global frame, v is the velocity at the vehicle's centre of mass, and ψ indicates the heading in the global frame. The control inputs, expressed as the vector $\mathbf{u} = [a, \delta]'$, correspond to the acceleration a and the front wheel steering angle δ . The velocity and acceleration are applied at the vehicle's centre of mass in the direction of the slip angle β . The dynamics of the vehicle are governed by a set of coupled ordinary differential equations. In this study, we use the discrete-time variant of this dynamic model to periodically update the system ($\mathbf{z}(t + 1) = \mathbf{f}(\mathbf{z}(t), \mathbf{u}(t), \Delta t)$):

$$x(t + 1) = x(t) + v(t) \cos(\psi(t) + \beta(t))\Delta t \quad (4.1a)$$

$$y(t + 1) = y(t) + v(t) \sin(\psi(t) + \beta(t))\Delta t \quad (4.1b)$$

$$v(t + 1) = v(t) + a(t)\Delta t \quad (4.1c)$$

$$\psi(t + 1) = \psi(t) + \frac{v(t)}{l_r} \sin(\beta(t))\Delta t \quad (4.1d)$$

$$\beta(t) = \tan^{-1} \left[\frac{l_r}{l_r + l_f} \tan(\delta(t)) \right] \quad (4.1e)$$

where t represents the discrete time interval index, Δt the length of each time interval in the prediction horizon, and l_f and l_r the distances from the front and rear wheel centres to the vehicle's centre of mass, respectively. In this work, we assume a vehicle's centre of mass coincides with its centroid (i.e. $l_f = l_r = l_B/2$, where l_B is the wheelbase) (Yue et al., 2015). The vehicle's dimensions are defined by its length (l) and width (w). We assume a vehicle

has an overhang ratio of $5/3$, meaning the wheelbase l_B is 0.6 times the vehicle's length l (AASHTO, 2018).

The lane in which the vehicle is travelling (ζ) is determined by the vehicle's y -position relative to the lane markings. Thus, a vehicle is characterised by its index, x -position, y -position, velocity, heading, length, width, and lane. These dynamic and configuration variables are collectively referred to as a vehicle's extended state ($\mathbf{Z} = [id, x, y, v, \psi, l, w, \zeta]'$). Figure 5.1 illustrates a vehicle in its global frame of reference.




FIGURE 4.1. Kinematic bicycle model for an ego vehicle. The vehicle states are defined using the x , y -positions, the velocity, and the heading $[x, y, v, \psi]$. The control actions are the acceleration and steering angle $[a, \delta]$. Velocity and acceleration are applied on the centroid in the direction of the slip angle (β). The vehicle has dimensions $w \times l$ and is travelling in lane ζ .

To realistically model human driving behaviour, we assume that drivers' decisions are influenced by their perceived payoffs, which they aim to optimise over a specific time interval. Essentially, a driver selects a series of optimal actions (\mathbf{u}^*) that minimise their cost function ($\mathcal{L}(\cdot)$) over a finite time horizon. This cost function depends on the initial states of both the ego and opponent vehicles ($\mathbf{Z}(t_0)$ and $\mathbf{Z}^-(t_0)$), as well as their control actions throughout the prediction horizon (\mathbf{u} and \mathbf{u}^-). The method for obtaining the opponents' states and controls is detailed in Section 4.2.2. The optimal action inputs to be applied at the current time step are determined as follows:

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \mathcal{L}(\mathbf{Z}(t_0), \mathbf{Z}^-(t_0), \mathbf{u}, \mathbf{u}^-). \quad (4.2)$$

The cost function, like the system dynamics, is expressed in discrete time, where the running costs ($g(\cdot)$) are summed to measure the cumulative performance of the actions. Thus, finding the optimal actions involves solving the following finite-dimensional optimisation problem:

$$\mathcal{L}(\mathbf{Z}(t_0), \mathbf{Z}^-(t_0), \mathbf{u}, \mathbf{u}^-) = \sum_{t=t_0}^{t_0+N-1} g(\mathbf{Z}(t), \mathbf{Z}^-(t), \mathbf{u}(t), \mathbf{u}^-(t)) \quad (4.3)$$

subject to:

$$\mathbf{z}(t+1) = \mathbf{f}(\mathbf{z}(t), \mathbf{u}(t)), \quad \forall t \in t_0, \dots, t_0 + N - 1 \quad (4.4)$$

$$\mathbf{z}^-(t+1) = \mathbf{f}(\mathbf{z}^-(t), \mathbf{u}^-(t)), \quad \forall t \in t_0, \dots, t_0 + N - 1 \quad (4.5)$$

$$\mathbf{u}(t) \in \mathcal{U}, \quad \forall t \in t_0, \dots, t_0 + N - 1 \quad (4.6)$$

where N denotes the number of intervals over the prediction horizon, $\mathbf{f}(\cdot)$ represents the set of differential equations (Equations 4.1) describing the system dynamics, and \mathcal{U} defines the admissible control space. The admissible control space is defined as a two-dimensional continuous range within fixed bounds:

$$\mathcal{U} = \{(a(t), \delta(t)) \mid a_{\min} \leq a(t) \leq a_{\max}, \delta_{\min} \leq \delta(t) \leq \delta_{\max}\}. \quad (4.7)$$

We set a continuous admissible control space to ensure that the optimisation conducted is smooth and that the trajectories generated are differentiable. Time is discretised within the prediction horizon to ensure a finite number of controls for the computation to be feasible. The running cost is expressed as a weighted summation of individual cost components (g_i), which represent the trade-offs between costs:

$$g(\mathbf{Z}(t), \mathbf{Z}^-(t), \mathbf{u}(t), \mathbf{u}^-(t)) = \sum_i \alpha_i g_i(\mathbf{Z}(t), \mathbf{Z}^-(t), \mathbf{u}(t), \mathbf{u}^-(t)) \quad (4.8)$$

where α_i represents the weighting parameter for each cost component i . Our study considers a total of six cost components.

Speed cost. Drivers want to minimise their travel time, to the extent conditions permit. The **speed** cost component stems from the vehicle speed deviating from a so-called ‘safe speed’. Imagine driving on a level, smooth, straight desert highway with perfect vision and




FIGURE 4.2. The safe velocity (v_{safe}) is defined as the maximum speed at which a vehicle can come to a complete stop without colliding, should an object appear at the sight distance (D_s). Since sight distance is a parameter we later calibrate, the vertical height is ignored.

no other vehicles around. The desire to speed would govern as long as there were nothing else (other vehicles, people, animals, road debris, etc.) in the line of sight. Therefore, we defined the safe speed v_{safe} such that the vehicle can perform a complete stop if an object appears at the furthest visible point (see Figure 4.2). The safe speed (v_{safe}) is then derived through the kinematic equations of motion with consideration of the driver's reaction time. The cost is formulated as a quadratic function of the difference between the current speed and the safe speed, ensuring symmetry in penalisation for over- and under-speeding, and yielding a minimum when the vehicle's speed matches the safe speed:

$$g_{\text{speed}} = (v(t) - v_{\text{safe}})^2 \quad (4.9)$$

$$v_{\text{safe}} = t_r a_{\text{min}} + \sqrt{t_r^2 a_{\text{min}}^2 - 2a_{\text{min}} D_s} \quad (4.10)$$

where t_r denotes the reaction time of the driver, a_{min} the maximum braking, and D_s the sight distance.

Spacing cost. The **spacing** component is defined as a non-decreasing function of the difference between a spacing threshold and the current spacing. Vehicles are assumed to maintain a velocity-dependent distance from the vehicle ahead, balancing the desire to travel faster (captured by the speed cost) against the need to preserve safety (captured by the spacing cost). In this work, we assume that each ego vehicle can have up to two preceding vehicles to account for interactions across lanes during lane-changes. When a vehicle initiates a lane-change, it considers the vehicles ahead in both the original and target lanes as its preceding

vehicles. Similarly, when another vehicle merges in front of the ego vehicle, the ego car treats both the merging vehicle and its original lead vehicle as preceding cars. Preceding vehicles are identified as opponents that are ahead of the ego vehicle in the longitudinal direction and have lateral overlap. Among these, the spacing cost is computed using the preceding vehicle that results in the largest difference between the desired spacing threshold and the current spacing:

$$g_{\text{spacing}} = \max_{j \in \mathbb{J}_p^-} (0, s_j^*(t) - s_j(t)) \quad (4.11)$$

The spacing threshold $s_j^*(t)$ and the current spacing $s_j(t)$ for each preceding vehicle j in the set of preceding cars \mathbb{J}_p^- are calculated as follows:

$$s_j^*(t) = s_0 + t_{\text{safe}}v(t) + \frac{v(t)(v(t) - v_j(t))}{2\sqrt{|a_{\text{max}}a_{\text{min}}|}} \quad (4.12)$$

$$s_j(t) = x_j(t) - x(t) - \frac{1}{2}(l_j + l) \quad (4.13)$$

where $x_j(t)$, $v_j(t)$, and l_j are the x-position, velocity, and length of the j^{th} preceding car. s_0 and t_{safe} respectively denote the minimum standstill gap and the driver's safe time headway, and a_{max} and a_{min} are the maximum and minimum acceleration capacities of the ego vehicle. The IDM model is used to capture the car-following behaviour, as it has been shown to perform well for both passenger cars and trucks in heterogeneous traffic ([Ossen and Hoogendoorn, 2011](#)). It is important to acknowledge that the car-following component, when combined with other cost terms, models the trade-offs between longitudinal acceleration and lateral actions such as lane changes.

Safety cost. The **safety** cost applies to all surrounding opponents of the ego vehicle. It penalises a low time-to-collision (TTC) between the ego vehicle and the surrounding vehicles by considering their relative positions and velocities. To ensure safety and prevent collisions, the ego vehicle must balance all opponents in the game when maintaining a safe distance.

The cost component is defined as the summation of the negative exponential of the time-to-collision between the ego vehicle and each of its opponents.

$$g_{\text{safety}} = \sum_{j \in \mathbb{J}} e^{-TTC_j} \quad (4.14)$$

where TTC_j represents the time-to-collision between the ego car and the j^{th} opponent in set \mathbb{J} , \mathbb{J} is the set of all opponent vehicles to the ego vehicle and the construction of this set is discussed in Section 4.2.2.

The TTC for an ego-opponent pair is calculated by dividing the two-dimensional gap between the two by their closing rate (Hayward, 1971). To determine whether a collision can occur, we need to check if the two vehicles are approaching each other and if their bounding shapes can overlap. We first determine the relative velocity ($\mathbf{v}_{e,j}$) between the two vehicles:

$$\mathbf{v}_{e,j} = \mathbf{v}_e - \mathbf{v}_j \quad (4.15)$$

where \mathbf{v}_e and \mathbf{v}_j are the respective velocity vectors for the ego and opponent vehicles. Let $\mathbf{P} = [x, y]'$. The vehicles are closing in on each other if:

$$(\mathbf{P}_j - \mathbf{P}_e) \cdot \mathbf{v}_{e,j} > 0 \quad (4.16)$$

We then check whether the vehicles overlap with their current velocities. To do this, we project the bounding shapes of the vehicles to the orthogonal plane ($\mathbf{v}_{e,j}^\perp$) of the relative velocity vector. We first make the assumption that a vehicle's shape can be represented by an ellipse, which completely encloses the vehicle bounding box. Ellipses have smooth, continuous boundaries, making the associated mathematical functions easier to deal with and more amenable to our optimisation framework. Rectangle bounding boxes, although more common, have discontinuities in their boundary conditions, complicating the optimisation problem. It is important to note that there can be infinitely many such enclosing ellipses, and we aim to find the one where the ratio of the major to minor axes matches the ratio of the length to the width of the vehicle bounding box. The formula for the ellipse, oriented

according to the local vehicle orientation, is as follows:

$$\frac{2x^2}{l^2} + \frac{2y^2}{w^2} = 1 \quad (4.17)$$

Furthermore, the projection of an ellipse onto a line can be determined by explicitly differentiating the parametric form of the ellipse to find the tangent gradient. The projected length can be determined based on the observation that the gradient at the maximum x would be zero. The following equation calculates the length of the projection (w_{proj}) for an ellipse oriented in the direction ψ onto a line of projection at angle θ :

$$w_{\text{proj}} = 2\sqrt{\frac{l^2}{2} \sin^2\left(\frac{\pi}{2} - \psi + \theta\right) + \frac{w^2}{2} \cos^2\left(\frac{\pi}{2} - \psi + \theta\right)}. \quad (4.18)$$

When projected onto the orthogonal relative velocity plane ($\mathbf{v}_{e,j}^\perp$), the distance between centres of the two vehicles ($\Delta D_{e,j}^\perp$) is:

$$\Delta D_{e,j}^\perp = \left| (\mathbf{P}_e - \mathbf{P}_j) \cdot \frac{\mathbf{v}_{e,j}^\perp}{\|\mathbf{v}_{e,j}^\perp\|} \right| \quad (4.19)$$

where $\mathbf{v}_{e,j} \cdot \mathbf{v}_{e,j}^\perp = 0$. Then, we can use Equations 4.18 and 4.19 to check for overlap in the $\mathbf{v}_{e,j}^\perp$ plane. The vehicles overlap if:

$$\Delta D_{e,j}^\perp < \frac{1}{2}(w_{\text{proj}}^{e,\perp} + w_{\text{proj}}^{j,\perp}) \quad (4.20)$$

where $w_{\text{proj}}^{e,\perp}$ and $w_{\text{proj}}^{j,\perp}$ are the lengths of the projections of the enclosing ellipses for the ego and opponent vehicles onto $\mathbf{v}_{e,j}^\perp$. An ego–opponent pair is considered to be on a collision course if and only if both Inequalities 4.16 and 4.20 are satisfied. Notably, vehicles that are closing in on each other may not necessarily intersect spatially, and conversely, vehicles that spatially overlap based on their instantaneous positions and velocities may not intersect in the future (such overlaps could have occurred in the past). Therefore, to compute the TTC, it is essential that the vehicles are predicted to overlap in the future.

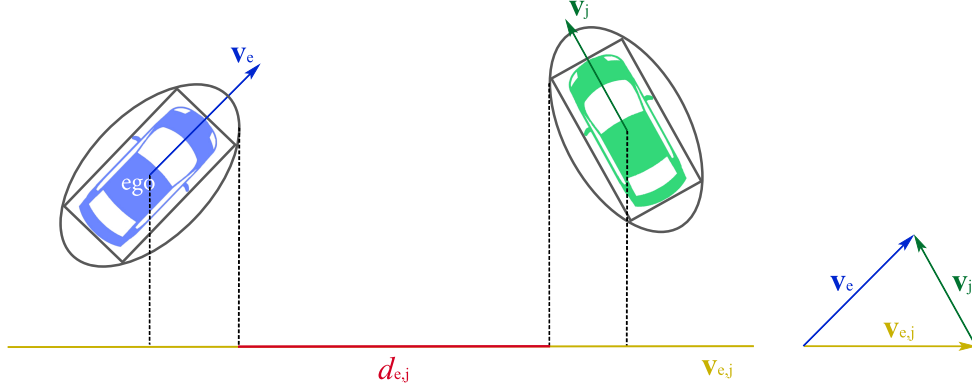


FIGURE 4.3. The elements for computing the time-to-collision between an ego vehicle and one of its opponents.

If both conditions are satisfied, we can calculate the gap ($d_{e,j}$) in the direction of $\mathbf{v}_{e,j}$ similar to how we have computed the overlapping condition 4.20:

$$d_{e,j} = \Delta D_{e,j} - \frac{1}{2}(w_{\text{proj}}^e + w_{\text{proj}}^j) \quad (4.21)$$

Finally, the TTC between the ego car and the j^{th} opponent car is calculated as:

$$TTC_j = \begin{cases} \frac{d_{e,j}}{\|\mathbf{v}_{e,j}\|} & \text{if conditions 4.16 \& 4.20 hold} \\ \infty & \text{otherwise} \end{cases} \quad (4.22)$$

An illustration of the TTC calculation is provided in Figure 4.3.

Steering and acceleration costs. The **acceleration** and the **steering** cost components are adopted to penalise extensive control inputs. These components are mainly used to stabilise the closed-loop system in the receding horizon optimisation. They could also reflect the consideration of comfort in human driving.

$$g_{\text{acceleration}} = \frac{a(t)^2}{|a_{\text{max}} a_{\text{min}}|} \quad (4.23)$$

$$g_{\text{steering}} = \frac{\delta(t)^2}{|\delta_{\text{max}} \delta_{\text{min}}|} \quad (4.24)$$

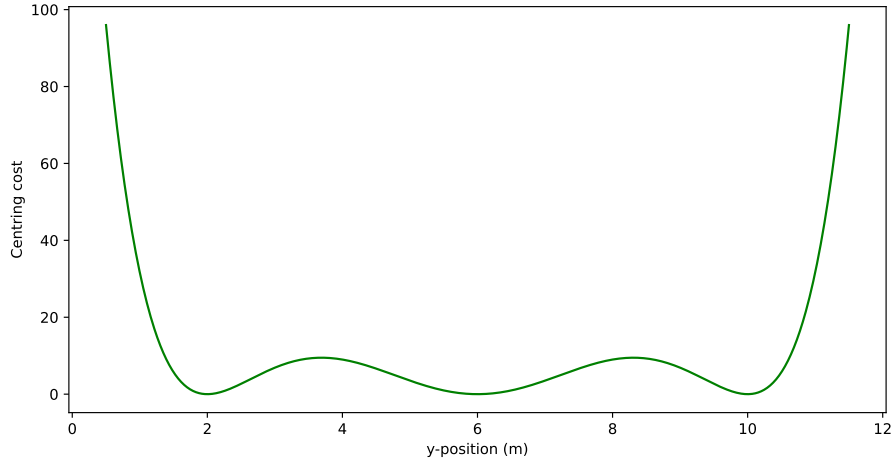


FIGURE 4.4. Lane-centring cost for a road with lane-markings at 0, 4, 8, and 12 m. The minima of the cost function are at 2, 6, and 10 m.

Lane-centring cost. The **lane-centring** cost component penalises the ego vehicle if it deviates from the centre of the lane. This can be considered a lane-keeping cost, ensuring that the vehicle remains centred in the lane unless a lane-change is being made. With this component, vehicles will weigh the trade-off between changing lanes and staying in the current lane, ensuring that lane-changes are not made for marginal utility gains. This cost component is defined as a polynomial, with the lane centres as its minima:

$$g_{\text{centre}} = \prod_{i=1}^{\omega} \left(\frac{y(t) - \frac{1}{2}(\Omega_{i+1} + \Omega_i)}{\frac{1}{2}(\Omega_{i+1} - \Omega_i)} \right)^2 \quad (4.25)$$

where ω denotes the maximum number of lanes and Ω the set of lane marking positions. Figure 4.4 depicts the centring cost for a road with lane-markings at 0, 4, 8, and 12 m.

4.2.2 Level-k Game Theory

To compute the optimal actions over the prediction horizon, it is necessary to consider the interaction of the ego vehicle with surrounding vehicles. For this purpose, we use game theory to replicate the interactive nature of human driving. Humans rely on a combination of visual, auditory, and cognitive cues to perceive the states of surrounding vehicles. In this study, we

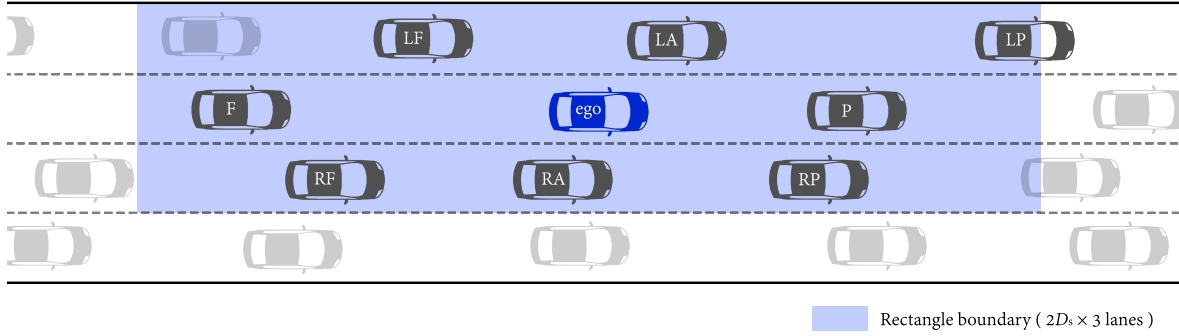


FIGURE 4.5. Game opponent identification by an ego vehicle. The ego vehicle selects 0 to 8 surrounding vehicles to be its opponents. The figure depicts 8 opponents for illustration purpose. Not all 8 need to be selected. The ego vehicle is depicted in dark blue, the dark grey vehicles represent the opponents to the ego vehicle, while the light grey vehicles are the detected surrounding vehicles that are not considered the opponents. The opponent labels: P, F, LP, LA, LF, RP, RA, RF, respectively represent: preceding, following, left-preceding, left-aligned, left-following, right-preceding, right-aligned, right-following. D_s denotes the ego driver's sight distance.

adopt the assumption of certainty equivalence, meaning that we treat the actual states of the opponents as though they were perfectly estimated.

Before detailing the game setup, we first introduce a heuristic used to identify the opponents for an ego vehicle. From the set of all vehicles (\mathbb{J}), the ego vehicle considers a subset (\mathbb{J}^-) that contains the vehicles it will interact with. At each time step, the ego vehicle perceives the surrounding vehicles and selects up to eight of them as opponents for the current iteration. This selection process is illustrated in Figure 4.5. The dark grey vehicles are the opponents to the ego vehicle and constitute set \mathbb{J}^- . The light grey vehicles are those that are detected by the ego vehicle but are not considered its opponents. They belong to set \mathbb{J} . We use the following heuristic to identify the set of opponents (\mathbb{J}^-) from all detected vehicles (\mathbb{J}):

- (1) Define a rectangular box around the ego vehicle. This box can be interpreted as a field of vision/perception of a human driver. The box is centred around the ego vehicle and has a length of D_s and a width of up to 3 lanes (we assume a 3-lane road in our scenarios below).
- (2) Identify all vehicles within this rectangular box as potential game opponents.

- (3) Label the vehicles directly in front of and behind the ego vehicle as P (preceding) and F (following), respectively, if they are present.
- (4) Label any vehicles in the left and right lanes that overlap with the projection of the ego vehicle into these lanes as LA (left-aligned) and RA (right-aligned), respectively, if they are present.
- (5) Label other, non-aligned vehicles that are in front of and behind the ego vehicle's projection in the left lane as LP (left-preceding) and LF (left-following), respectively, if they are present.
- (6) Label vehicles that are in front of and behind the ego vehicle's projection in the right lane as RP (right-preceding) and RF (right-following), respectively, if they are present.
- (7) Add the labels P, F, LA, RA, LP, LF, RP, RF to the set of game opponents \mathbb{J}^* , if they are present.

It is worth noting that a vehicle, especially a long vehicle such as a truck, may have more than one LA or RA vehicle. However, upon inspecting the datasets used for our analysis, we found no such cases. Future studies could analyse the impact of multiple aligned vehicles and provide resolving strategies. After identifying the opponents and their states, the ego vehicle engages with them to make optimal decisions. This interaction is modelled using dynamic game theory, framed by the question: what would be my optimal strategy given my opponents' actions? The ego vehicle then uses this framework to decide its actions based on the predicted actions of the opponents.

There appears to be a hierarchical structure in human reasoning ([Camerer et al., 2004](#), [Stahl and Wilson, 1995](#)). To incorporate this reasoning hierarchy, we leverage 'level-k game theory' (or cognitive hierarchy theory) to predict the actions of the opponents through a set of finite recursions. Level-k game theory has been used to model vehicle interactions in the field of control theory ([Oyler et al., 2016](#)).

In employing a level-k reasoning structure, the ego vehicle is assigned a rationality level K , which represents its level of strategic thinking in the game. Higher levels of K indicate a deeper understanding or anticipation of others' actions. In level-k game theory, players

assume their opponents have cognitive levels one less than themselves, meaning all opponents of the ego vehicle are at level $K - 1$. The opponents of these level $K - 1$ players are at level $K - 2$, and so on. The recursion ends at the base case, level 0, where we assume players have no predictive power and their prediction horizon is zero. Consequently, their objective consists only of the running cost at the current time, leading them to opt for inaction as their best strategy, thus travelling at their currently velocities. All other players, whose levels are higher than 0, believe that their opponents are at level $k - 1$, and they best respond to the strategies they expect from the level $k - 1$ opponents. For an ego vehicle with a rationality level K , its optimal actions (\mathbf{u}^{K*}) can be computed by iteratively calculating the optimal controls for vehicles from level 1 to K . Let $k \in \{1, 2, \dots, K\}$. Then for each k , compute:

$$\mathbf{u}^{k*} = \arg \min_{\mathbf{u}^k} \mathcal{L}(\mathbf{Z}^k(t_0), \mathbf{Z}^{k-1}(t_0), \mathbf{u}^k, \mathbf{u}^{k-1*}) \quad (4.26)$$

where \mathbf{Z}^k denotes the states of level k vehicles, \mathbf{u}^k denotes the actions adopted by the level k vehicles, and $\mathbf{u}^{0*} = \mathbf{0}$. Essentially, to compute the optimal actions of an ego vehicle, we must iteratively determine the optimal actions of all lower-level opponents. This is because the ego vehicle lacks knowledge of the strategies that opponent vehicles will adopt. The level- k reasoning framework addresses this by assuming that each higher-level agent makes decisions based on the anticipated decisions of lower-level agents.

4.2.3 Numerical Solutions

Since our cost function is non-convex, deriving an analytical solution for the optimal actions is infeasible. Instead, we use a numerical method to solve for the global optimum. The Simplicial Homology Global Optimisation (SHGO) method from the SciPy package is employed for this purpose (Endres et al., 2018). SHGO is suitable for low-dimensional, constrained non-linear programming (NLP) problems and is particularly effective for blackbox optimisation. It is a general-purpose algorithm that has proven effective for non-continuous, non-convex, and non-smooth functions.

To simplify the computational complexity, we reduce the dimension of the solution space by specifying a control horizon (n_c) that is shorter than the prediction horizon (N). This

approach is common in receding horizon optimisation problems and has been shown to reduce the computational complexity (Geroliminis et al., 2013). The algorithm optimises control actions only for the duration of the control horizon but still considers the predictions of the system and surrounding objects over the entire prediction horizon, accounting for the future impact of current decisions. A shorter control horizon is often advantageous in fast-changing environments where long-term predictions are less reliable. In our problem, decisions depend not only on the actions of the ego vehicle but also on those of the surrounding vehicles, justifying the use of a shorter control horizon to reduce computational load.

Code to implement the 4th order Runge-Kutta method (RK4) as the numerical integrator was written to solve the system dynamics equations (Equations 4.1). The time interval used for making predictions within the horizon (Δt) is defined separately from the sampling time used for updating decisions (ΔT). A finer prediction interval is employed during optimisation to achieve higher solution granularity.

This completes the final piece of the framework, and the entire model is applied at each discrete time interval to decide the actions of a human-driven vehicle. The complete framework is outlined in Algorithm 3. As the rationality level K increases, the number of cars requiring computation also increases, approximately linearly. Consequently, a smaller K would reduce the computational load.

All computations in this chapter were conducted on a 2021 Apple M1 chip with an 8-core CPU and 16 GB of unified memory with macOS Ventura 13.6.1 operating system. No graphics card was used for numerical computations. The model was developed in Python 3.9.13. Numerical computations were performed using NumPy 1.24.4 and SciPy 1.9.1.

Algorithm 1 Game-Theoretic Receding Horizon Model

Input:

$\mathbf{Z}[0]$: state of the ego vehicles at time 0
 \mathbf{Z}^{oth} : states of all surrounding vehicles for the entire duration
 K : predefined cognitive level of the ego vehicle
 T : number of intervals over the entire duration
 Δt : length of intervals over the prediction horizon
 ΔT : length of intervals over the time axis
 N : number of intervals over prediction horizon
 n_c : number of intervals over control horizon

Output:

\mathbf{Z} : states of the ego vehicle for the entire duration
 \mathbf{u}^* : control actions of the ego vehicle for the entire duration

```

1: Initialise  $t \leftarrow 0$ 
2: while  $t \neq T$  do
3:    $\mathbf{Z}^K[t] \leftarrow$  perceive the state of the ego vehicle and assign a level  $K$ 
4:    $\mathbf{Z}^{\text{oth}}[t] \leftarrow$  perceive the states of all surrounding vehicles
5:   Initialise  $k \leftarrow K$ 
6:   while  $k \neq 0$  do:
7:     for  $player$  in  $\mathbf{Z}^k[t]$  do
8:        $\mathbf{Z}_{player}^{k-1}[t] \leftarrow$   $\mathbf{Z}_{player}^k[t]$ 's opponents from  $\text{Concat}(\mathbf{Z}[t], \mathbf{Z}^{\text{oth}}[t]) \setminus \mathbf{Z}_{player}^k[t]$  (Section 4.2.2)
9:     end for
10:     $k \leftarrow k - 1$ 
11:  end while
12:   $\mathbf{u}^{0*} \leftarrow \mathbf{0}$ 
13:  Initialise  $k \leftarrow 1$ 
14:  while  $k \neq K$  do
15:    for  $player$  in  $\mathbf{Z}^k[t]$  do
16:       $\mathbf{u}_{player}^{k*} \leftarrow \arg \min_{\mathbf{u}_{player}^k} \mathcal{L}(\mathbf{Z}_{player}^k[t], \mathbf{Z}_{player}^{k-1}[t], \mathbf{u}_{player}^k, \mathbf{u}_{player}^{k-1*})$  using SHGO
17:      where  $\mathcal{L}(\mathbf{Z}_{player}^k[t], \mathbf{Z}_{player}^{k-1}[t], \mathbf{u}_{player}^k, \mathbf{u}_{player}^{k-1*})$  is computed as:
18:       $\mathbf{u}_{player}^k \leftarrow \text{CONCATENATE}(\mathbf{u}_{player}^k, [\mathbf{u}_{player}^k[\text{nc} - 1]] \cdot (N - n_c))$ 
19:       $\mathbf{u}_{player}^{k-1*} \leftarrow \text{CONCATENATE}(\mathbf{u}_{player}^{k-1*}, [\mathbf{u}_{player}^{k-1*}[\text{nc} - 1]] \cdot (N - n_c))$ 
20:      Initialise  $\mathcal{L} \leftarrow 0$ 
21:      for  $s$  in  $[t, t + n_p - 1]$  do
22:         $\mathbf{Z}_{player}^k[s + 1] \leftarrow \text{RK4 on } \mathbf{f}(\mathbf{Z}_{player}^k[s], \mathbf{u}_{player}^k[s], \Delta t)$  (Equation 4.1)
23:         $\mathbf{Z}_{player}^{k-1}[s + 1] \leftarrow \text{RK4 on } \mathbf{f}(\mathbf{Z}_{player}^{k-1}[s], \mathbf{u}_{player}^{k-1*}[s], \Delta t)$  (Equation 4.1)
24:         $\mathcal{L} += g(\mathbf{Z}_{player}^k[s + 1], \mathbf{Z}_{player}^{k-1}[s + 1], \mathbf{u}_{player}^k[s], \mathbf{u}_{player}^{k-1*}[s])$  (Equation 4.3)
25:      end for
26:    end for
27:  end while
28:   $\mathbf{u}^*[t] \leftarrow \mathbf{u}^{K*}[0]$ 
29:   $\mathbf{Z}[t + 1] \leftarrow \text{RK4 on } \mathbf{f}(\mathbf{Z}[t], \mathbf{u}^*[t], \Delta T)$  (Equation 4.1)
30:   $t \leftarrow t + 1$ 
31: end while
32: return  $\mathbf{Z}, \mathbf{u}^*$ 

```

4.3 Calibration

4.3.1 Data

In this and the subsequent Section 4.4, we calibrate and validate the proposed model using the highD dataset, a large-scale naturalistic dataset of human trajectories collected on German highways (Krajewski et al., 2018). Vehicles were recorded at 25 fps (one frame per 0.04 second) using a high-resolution camera mounted on a drone. Each frame was annotated using segmentation masks created by a U-Net algorithm, and the tracks were post-processed through a Kalman filter. Each road segment in the dataset is approximately 420 m long without any ramps. Considering the leftward traffic in the dataset may contain mandatory lane-changes, we use only the rightward traffic for our study (Wang and Gao, 2025).

4.3.2 Calibration Procedure

The proposed model includes a total of 19 parameters that can be adjusted for each vehicle trajectory, including a_{\min} , a_{\max} , δ_{\min} , δ_{\max} , α , t_r , D_s , t_{safe} , s_0 , K , Δt , ΔT , N , and n_c . Conducting calibration on all these parameters would be computationally inefficient, so we simplify the model by setting default values for some parameters, as introduced in Table 4.1. This leaves 8 parameters that require calibration.

Instead of calibrating the model for a specific K , we assume the trajectories of the opponents follow their actual behaviour. There are two benefits from this approach. First, the computation and optimisation complexity for the nested optimisation is drastically reduced allowing for faster and more stable convergence. Second, the calibrated parameters are generalised across levels of K as parameters capture underlying behaviour of drivers.

Given the dataset includes both passenger vehicles and trucks, some parameters are defined separately to reflect the different control capacities of these vehicle types. In this study, the vehicles types are pertained only to the ego vehicle. For passenger vehicles, we set acceleration bounds at $[-3.4 \text{ m/s}^2, 3 \text{ m/s}^2]$ and steering bounds at $[-\pi/9 \text{ rad}, \pi/9 \text{ rad}]$.

TABLE 4.1. Parameters of the two-dimensional driving model

Parameter	Passenger vehicle	Truck
Vehicle Characteristics		
Minimum acceleration (a_{\min})	-3.4 m/s^2	-1.5 m/s^2
Maximum acceleration (a_{\max})	3 m/s^2	1 m/s^2
Minimum steering (δ_{\min})	$-\pi/9 \text{ rad}$	$-\pi/12 \text{ rad}$
Maximum steering (δ_{\max})	$\pi/9 \text{ rad}$	$\pi/12 \text{ rad}$
Simulation Parameters		
Horizon discretisation interval (Δt)	0.2 s	0.2 s
Sampling time (ΔT)	0.4 s	0.4 s
Number of prediction steps (N)	5	5
Number of control steps (n_c)	1	1
Driver Characteristics		
Reaction time (t_r)	0.6 s	0.6 s
Minimum standstill gap (s_0)	2 m	2 m
Sight distance (D_s)	to be calibrated	to be calibrated
Safe time headway (t_{safe})	to be calibrated	to be calibrated
Cost function weights (α)	to be calibrated	to be calibrated

For trucks, we assume a smaller action space of $[-1.5 \text{ m/s}^2, 1 \text{ m/s}^2]$ for acceleration and $[-\pi/12 \text{ rad}, \pi/12 \text{ rad}]$ for steering. These bounds are selected with considerations of empirical evidence (AASHTO, 2018, Bokare and Maurya, 2017).

All other parameters are consistent for both passenger cars and trucks. The driver reaction time (t_r) is set to 0.6 s (Johansson and Rumar, 1971). Considering that the safe time headway (T_{safe}) in the IDM model tend to differ between and within vehicle types (Ossen and Hoogendoorn, 2011), we add it to the list of calibration parameters. The minimum standstill gap (s_0) is treated as a constant with a value of 2 m (Kesting and Treiber, 2008). The horizon discretisation interval (Δt) for predictions within the horizon is 0.2 s, and the sampling time (ΔT) for updating the driver's decisions is 0.4 s. The 0.4 s interval is chosen as it is a multiple of the highD frame rate (0.04 s) and balances finer temporal resolution with computational efficiency. The number of prediction steps (N) is set to 5, giving a prediction horizon of 1 s. The number of control steps (n_c) in the optimisation is set to 1, due to the exponentially increasing computational complexity. Although the algorithm does not enumerate all possible

solutions to find the optimal control, a lower control horizon is still preferred due to the minimal impact on performance improvement and its computational burden. The values of the prediction and sampling intervals, along with the number of prediction and control steps, are manually tuned through extensive trial and error, balancing the smoothness of simulated trajectories and computational efficiency. Values reported in related studies are also used as guidelines (Guo et al., 2013, Lazcano et al., 2021, Prokop, 2001).

To calibrate the model, we simulate a trajectory and compare it with the measured trajectory, rather than adjusting to individual points. We take a holistic approach to trajectory calibration to more accurately assess the model’s performance over time. This method captures the dynamics of actual driving more effectively as it considers previous actions in future decision-making. We opt not to use conventional optimisation-based methods for calibration (Keane and Gao, 2021), as our implementation of 2DDM requires heterogeneous calibration across a group of vehicles, and traditional methods could lead to overfitting and a lack of generality. Instead, we employ a technique that calibrates trajectories collectively while recognising the diverse nature of human preferences. This approach involves using a Bayesian inference procedure where a prior is iteratively updated to yield an output posterior distribution. Interested readers can refer to (Von Toussaint, 2011).

Specifically, we perform Bayesian inference by introducing a slight modification to the standard Sequential Monte Carlo (SMC) algorithm (Del Moral et al., 2006). To improve computational efficiency and enable scalable inference on large datasets, we divide the original dataset into overlapping mini-batches and incorporate observations incrementally. This Mini-Batch Sequential Monte Carlo (MB-SMC) approach allows for faster updates by reducing the computational cost per iteration, while still preserving the sequential nature of the inference process. Each mini-batch provides a partial likelihood update, enabling the algorithm to refine the weights gradually and focus resampling on promising regions of the parameter space. The rest of the algorithm remains consistent with the original SMC framework. The MB-SMC algorithm is summarised in Algorithm 2. Hereinafter, we use the term ‘particle’ to refer to a sample of the model parameters used for calibration.

Algorithm 2 Mini-Batch Sequential Monte Carlo

Input:

$highD$: highD data served as observations
 \mathbf{a}_0 : lower limits of the prior distributions
 \mathbf{b}_0 : upper limits of the prior distributions
 σ_{pert} : standard deviations of the perturbation distributions
 γ : weight for exploitation vs exploration
 N_{particle} : number of particles to be generated
 $N_{\text{iteration}}$: number of iterations till termination

Output:

$particles$: optimal particles from the updated posterior distribution

```

1: Initialise  $i \leftarrow 0$ 
2:  $particles \leftarrow \{\mathbf{x}_j\}_{j=1}^{N_{\text{particle}}}$ , where  $\mathbf{x}_j \sim U(\mathbf{a}_0, \mathbf{b}_0)$ 
3:  $MiniBatches \leftarrow \{highD_i\}_{i=1}^{N_{\text{iteration}}}$ , where  $highD_i \subseteq highD$  &  $|highD_i| = 2 \left\lceil \frac{|highD|}{N_{\text{iteration}}} \right\rceil$ 
4: while  $i \neq N_{\text{iteration}}$  do
5:    $i \leftarrow i + 1$ 
6:   Initialise an empty list:  $likelihood \leftarrow []$ 
7:   for  $particle$  in  $particles$  do
8:      $similarity \leftarrow 0$ 
9:     for  $observation$  in  $MiniBatches[i]$  do
10:       $simulation \leftarrow$  simulate the trajectory based on  $particle$  and  $observation$ 
11:       $similarity \leftarrow similarity + D_F(simulation, observation)$ 
12:    end for
13:     $likelihood \leftarrow likelihood + [e^{-\gamma similarity}]$ 
14:  end for
15:   $weights \leftarrow likelihood / \sum_j likelihood$ 
16:   $particles \leftarrow$  sample with replacement from  $particles$  with probabilities  $weights$ 
17:   $particles \leftarrow \{\mathbf{x}_j\}_{j=1}^{N_{\text{particle}}}$ , where  $\mathbf{x}_j \sim Trunc\mathcal{N}(\mathbf{x}_j, \sigma_{\text{pert}}, \mathbf{a}_0, \mathbf{b}_0)$ 
18: end while
19: return  $particles$ 

```

In the MB-SMC method, particles are initialised using a uniform distribution, as we lack specific information about the actual shapes of the underlying distributions. However, we have established bounds for the parameters based on assumptions. The observed range of speeds from the highD dataset is used to approximate the range for the D_s distribution. We have set the limits at 150 m and 250 m, providing a broad span to ensure the initial particles are fairly distributed. The initial values of t_{safe} are between 1 s and 2 s (Kesting and Treiber, 2008). The limits for α are set between 0 and 1, as these weights are relative.

After the initial particles are selected, their performance is evaluated. The similarity between the simulated and observed trajectories is quantified using the discrete Fréchet distance (Eiter and Mannila, 1994). This similarity metric captures the minimal distance necessary to traverse both curves simultaneously from start to finish without backtracking. Unlike more conventional measures such as NRMSE and Theil’s U, the discrete Fréchet distance measures not only the proximity of the points but also how closely the modelled trajectory follows the shape and progression of the actual trajectory, capturing both spatial and temporal variations between the two paths. The formulation of the discrete Fréchet distance is as follows:

$$D(P, Q) = \min_{\rho_P, \tau_Q} \max_{t \in \{1, \dots, T\}} d(p_{\rho_P(t)}, q_{\tau_Q(t)}) \quad (4.27)$$

where $P = [p_1, p_2, \dots, p_n]$ and $Q = [q_1, q_2, \dots, q_m]$ represent two trajectory sequences, $d(p, q)$ is the euclidean distance between points p and q , ρ_P and τ_Q are sequences of indices traversing P and Q , and T is the length of sequences ρ_P and τ_Q , such that $\rho_P(1) = \tau_Q(1) = 1$ and $\rho_P(T) = n$ and $\tau_Q(T) = m$. We also scale the x-positions by multiplying the ratio between section width and section length ($12/420 \sim 0.286$). Otherwise, the optimal parameters would yield trajectories that closely match in the x coordinates while ignoring the closeness in the y coordinates.

The similarity between the trajectories is translated into likelihood using a negative-exponential function, which is then converted into weights through normalisation (done using the softmax function). These weights are used to sample the original particles, forming the basis for the new particles. Sampling allows for repeated draws to maintain the same number of particles. Subsequently, the new particles are perturbed using a truncated normal distribution. The entire process then iterates again with the new particles used to simulate the trajectories, which are compared against the observations. The calibration results are discussed in the following section.

4.3.3 Calibration Results

Since both passenger cars and trucks are available in the highD dataset, we calibrate these two vehicle types using 20% of the vehicles from Track 50. This provides us with a total of




FIGURE 4.6. The average Fréchet distance across all calibrated vehicles throughout the calibration iterations.

214 vehicle trajectories — 162 passenger cars and 52 trucks. Figure 4.6 depicts the change in average Fréchet distance over the calibration iterations. The calibrated particles and their respective best fits are illustrated in Figure 4.7. Table 4.2 summarises the statistics of the estimated distributions. Curve fitting is approached as an optimisation problem by minimising the negative log-likelihood function of a truncated normal distribution (refer to Appendix B1).

The distributions of the sight distance (D_s) appear reasonable. Applying the mean D_s to Equation 4.10, the resulting safe speeds are 35.88 m/s for passenger cars and 25.15 m/s for trucks. The actual mean speeds from Track 50 are respectively 32.07 m/s for passenger cars and 24.33 m/s for trucks. Intuitively, the D_s for trucks should be slightly higher than for passenger cars due to their higher eye level, which might slightly offset the longer braking distance. However, the D_s values should be similar, which is consistent with our findings.

The safe time headway in our calibration is lower than the 1.6 s value used in the original IDM paper (Treiber et al., 2000). This difference may be attributed to variations in the default parameters assumed in 2DDM. Additionally, the calibrated values for trucks are higher than that for passenger cars, which is expected, as trucks generally maintain more conservative gaps due to their longer braking distances.

As the cost components are not normalised, the weights do not indicate the absolute scale of preferences. For instance, the calibration results do not necessarily suggest that the desire for speeding is less significant than other cost considerations. Rather, it indicates that in 2DDM, the weighting for speeding is smaller, potentially because of a larger scale of this component.

TABLE 4.2. Statistics of the truncated normal fits to the calibrated parameters

	Passenger car			Truck		
	Mean	Std	Truncation	Mean	Std	Truncation
Sight distance (D_s)	210.83 m	33.64 m	$[0, \infty)$	225.96 m	22.61 m	$[0, \infty)$
Safe time headway (t_{safe})	0.98 s	0.10 s	$[0, \infty)$	1.19 s	0.08 s	$[0, \infty)$
Speed cost weight (α_{speed})	0.02	0.12	$[0, 1]$	0.05	0.15	$[0, 1]$
Spacing cost weight (α_{spacing})	0.41	0.12	$[0, 1]$	0.51	0.14	$[0, 1]$
Safety cost weight (α_{safe})	0.40	0.11	$[0, 1]$	0.49	0.13	$[0, 1]$
Acceleration cost weight ($\alpha_{\text{acceleration}}$)	0.94	0.16	$[0, 1]$	0.86	0.16	$[0, 1]$
Steering cost weight (α_{steering})	0.88	0.12	$[0, 1]$	0.90	0.16	$[0, 1]$
Centring cost weight (α_{centring})	0.85	0.10	$[0, 1]$	0.80	0.12	$[0, 1]$

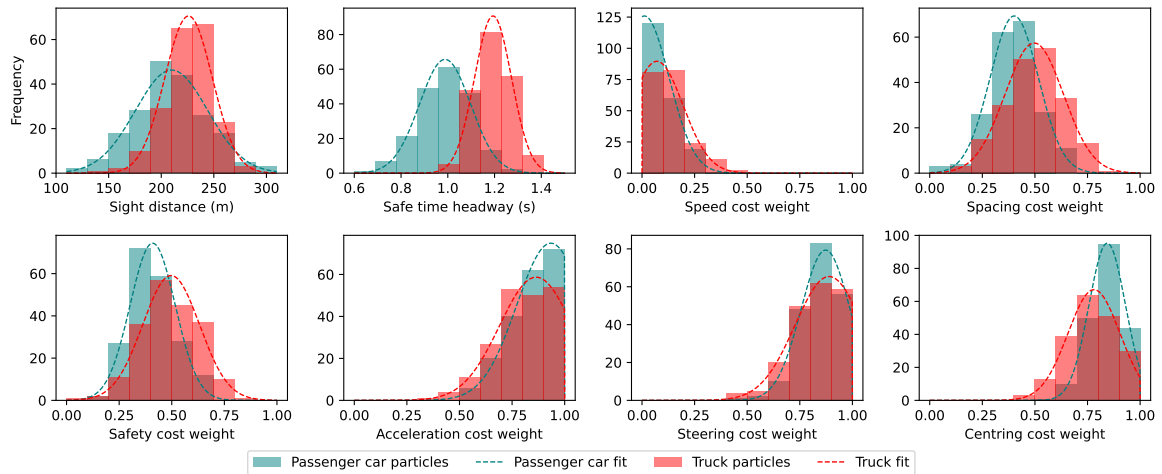


FIGURE 4.7. Calibrated distributions of parameters. The histograms are produced using the calibrated particles from the MB-SMC algorithm, while the dashed lines represent the fitted truncated normal distributions.

We considered normalising the weights, but without a clear scale for each component, this task proved unfeasible. A sigmoid function was applied to each component, but given that the impacts of larger costs, such as exponentially increasing cost with lower TTC, are downscaled, we decided ultimately not to normalise the components.

The mean values of the weights are nevertheless comparable between vehicle types. The standard deviations for the trucks tend to be slightly higher than those for the passenger cars. This might suggest a greater variety of driving behaviour among truck drivers, and/or a higher variation in truck types. Nevertheless, these weights are all relative, so the exact cause of this

discrepancy is not entirely clear. Regardless, the focus of this chapter is on the ability of these parameters to replicate the actual trajectories is tested and discussed in the upcoming section.

4.4 Validation

The highD dataset includes sections of both 2-lane and 3-lane highways. To validate the performance of 2DDM and the tuned parameters, we simulate vehicle trajectories on both 2-lane and 3-lane tracks. Initially, the model is validated on the calibrated Track 50 to demonstrate its performance in a familiar context. The highway section in Track 50 consists of 3 lanes. Subsequently, two other tracks representing different locations and different times are used to assess the transferability of the model and to test its ability to generalise to unseen data.

For each validation track, we simulate trajectories using 2DDM for $K = 1$ and $K = 2$, and the IDM+MOBIL model. This comparison showcases the performance of 2DDM against a model that treats car-following and lane-changing as separate processes. To maintain a fair comparison between the models, we have chosen desired speed, safe time headway, politeness factor, and lane-changing threshold as the calibration parameters. The acceleration and deceleration limits are chosen as $[-3.4 \text{ m/s}^2, 3 \text{ m/s}^2]$ for passenger cars and $[-1.5 \text{ m/s}^2, 1 \text{ m/s}^2]$ for trucks. The IDM+MOBIL model is similarly calibrated using the MB-SMC algorithm. The distributions of its parameters are illustrated in Appendix B3.

Similar to the calibration process, we initiate trajectory simulations by providing the model with the first data point of an observed trajectory. Using this initial point and the positions of surrounding vehicles, the model simulates the entire path. To prevent indefinitely long simulations, we terminate the algorithm once a vehicle reaches 420 m.

Figure 4.8 presents four sets of observed and simulated trajectories obtained from the highD dataset and our proposed model. The chosen observed trajectories contain either zero, one, or two lane-changes. The Fréchet distance is shown in the sub-captions to quantify the level of similarity between the observed and simulated trajectories. Note that calibration is done

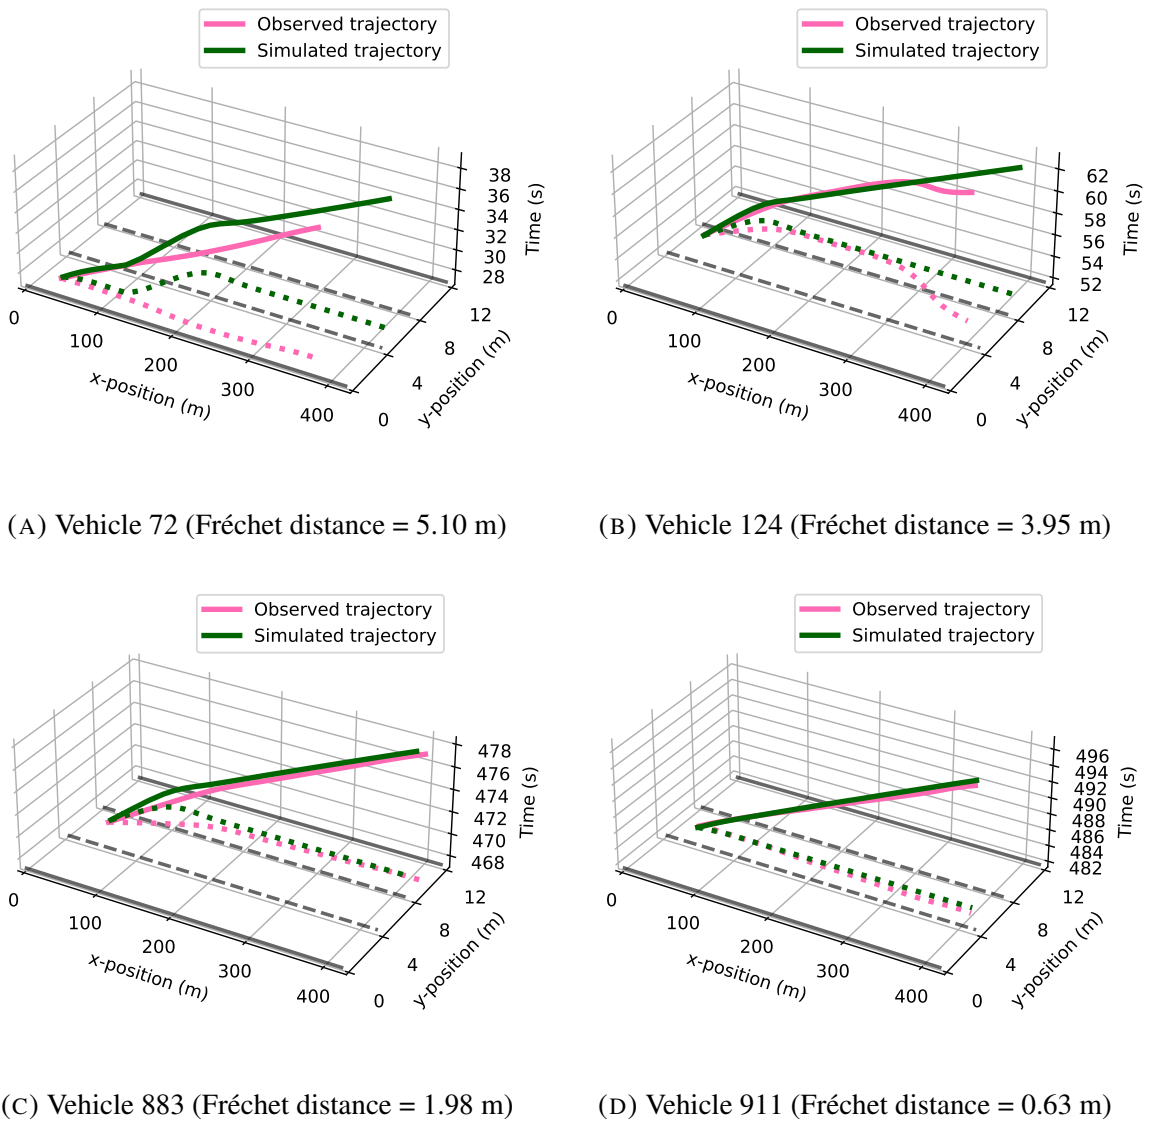


FIGURE 4.8. Four examples from Track 50 showcasing the observed and simulated vehicle trajectories. The observed curves from vehicles 72 and 911 contain no lane-changes, while vehicle 883 contains one lane-change and vehicle 124 contains two lane-changes. Their respective Fréchet distances are presented in the sub-captions. The dashed coloured lines on the x-y planes represent the projected 2D trajectories. The dashed and solid grey lines on the x-y planes represent the road boundaries and lane dividing lines, respectively.

collectively on 20% of the vehicles in Track 50. Therefore, the observations and simulations are not expected to closely match.

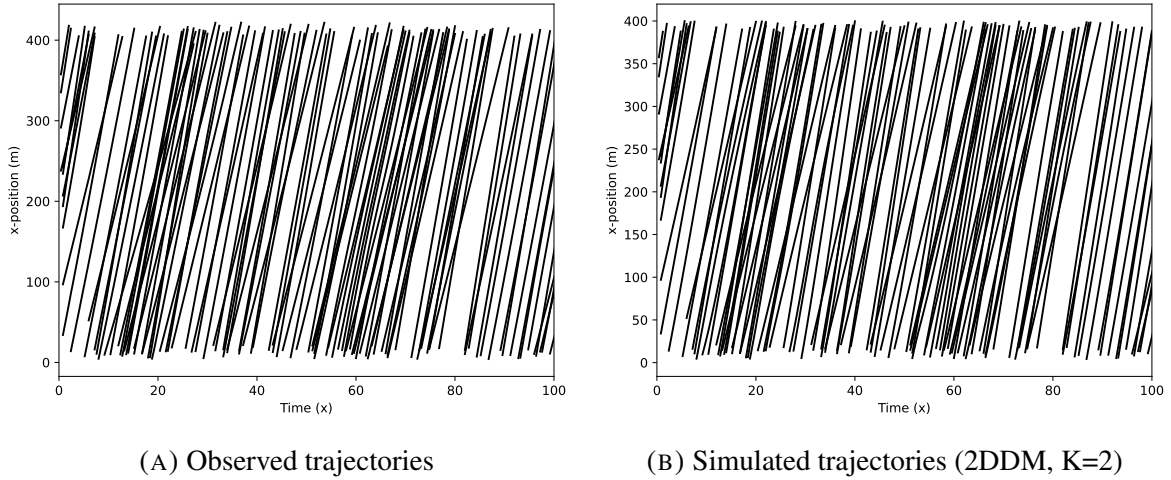


FIGURE 4.9. Space-time diagram for first 100 s in Track 50 showcasing the observed and simulated trajectories.

In addition to the individual trajectories, we also showcase the trajectories of a vehicle fleet for a short time interval. Figure 4.9 illustrates both the observed and the simulated ($K = 2$) trajectories from Track 50. The plots are alike but not identical since vehicles are generated with random sets of parameters drawn from the calibrated distributions.

4.4.1 Validation on the Calibrated Track (Track 50)

As previously mentioned in Section 4.3.3, 20% of vehicle trajectories in Track 50 were used to calibrate the model parameters. We now use all trajectories in the same track to validate the model and its parameters. This track has a total duration of 18.93 minutes and includes 1,093 vehicles. Before initiating the simulation, each vehicle randomly selects a set of 8 parameters from the fitted distributions, as detailed in Table 4.2. Passenger cars and trucks select from their respective distributions. We sample randomly from the 8 distributions, as all parameters are weakly correlated, as seen in Appendix B2. The simulation then runs for the entire duration of the track. This process is repeated 20 times, each iteration using a different set of drawn parameters. The average computation time for Track 50 is 57.67 min for $K = 1$ and 112.03 min for $K = 2$. Considering that the total travel time for all vehicles is 228.65 min, the algorithm computes trajectories for vehicles in real-time. However, the overall computation

TABLE 4.3. Validation results for Track 50. LC denotes lane-change.

	Observed	2DDM, K=1		2DDM, K=2		IDM+MOBIL	
		Mean	Std	Mean	Std	Mean	Std
Generalised flow (veh/hr/lane)	1079.64	1096.05	0.38	1096.05	0.51	1047.55	5.52
Generalised density (veh/km/lane)	9.92	9.91	0.03	9.91	0.03	9.87	0.05
Average speed (m/s)	30.99	31.39	0.10	31.28	0.12	30.45	0.07
Average headway (s)	2.58	2.56	0.01	2.56	0.01	2.63	0.03
LC frequency (LC/hr/lane)	125.79	129.58	10.31	131.03	12.79	136.15	9.69
LC composition (right to mid) (%)	9.24	41.12	2.99	43.08	4.29	29.50	5.16
LC composition (mid to left) (%)	32.77	23.76	4.99	28.75	3.85	29.73	3.38
LC composition (left to mid) (%)	42.86	16.72	2.96	15.67	2.4	14.41	3.01
LC composition (mid to right) (%)	15.13	18.40	2.85	12.50	4.05	26.37	6.58
Computation time (min)	-	57.67	1.17	112.03	2.01	0.51	0.01

time is still high compared to the IDM+MOBIL model which only requires 0.51 min on average using the same computer.


The performance of these simulations is quantified using traffic state variables and compared with their observed counterparts, as shown in Table 4.3. The state variables used include generalised flow, generalised density, average speed, average headway, lane-change frequency, lane-based lane-change composition. Generalised traffic states, as described by (Edie, 1963), are defined over a specific distance and time period. Here “area” (A), refers to the spatio-temporal domain, which is 420 m times 18.93 min. Generalised states are calculated based on A along with the total travelled distance (D_{Total}) and the total time spent (T_{Total}) by all vehicles. The generalised flow is given by $Q_{Generalised} = D_{Total}/A$, while the generalised density is $K_{Generalised} = T_{Total}/A$.

As demonstrated, the simulations of the track have accurately recreated the observations. The absolute percentage errors for all state variables in 2DDM for both $K = 1$ and $K = 2$ are less than 10%. This represents an improvement over the IDM+MOBIL model, which shows lower accuracy in measures such as LC frequency. The overestimation of LC frequency in IDM+MOBIL may stem from its treatment of lane-changing as a discrete event rather than a continuous, two-dimensional process. Specifically, IDM+MOBIL models lane-changes as instantaneous lateral shifts triggered when the acceleration condition in an adjacent lane is superior, without accounting for the dynamic and spatial aspects of lane-changing behaviour.

As seen from Table 4.3, LC compositions show the highest discrepancy from the observed values. The simulated results indicate that vehicles tend to change lanes towards the left. This is reasonable because the default starting points of vehicles in the simulation are generated based on observations from Germany, which adheres to the keep-right directive, and the simulation is over a relatively short distance, resulting in most trucks occupying the right-most lane. Consequently, passenger cars often make left lane-changes to bypass trucks and increase speed. In the observed data, vehicles frequently make lane-changes in the second and third lanes potentially because traffic has already stabilised, with vehicles comfortable at lower speeds already positioned in lane 1, while vehicles desiring higher speeds occupy lanes 2 and 3. As a result, there are no frequent lane-changes made by vehicles in lane 1, while those in lanes 2 and 3 consistently look for gaps to move into.


As shown in Figure 4.10, we also compare the speed profiles of the observed and simulated vehicles. The simulated plot is created by normalising the frequencies of the 20 combined histograms. We compute the Wasserstein distance (D_W) between the simulated and observed plots and the values are stated in the captions of the subplots. Both D_W s and visualisations indicate that the speed profiles generated by 2DDM outperforms the IDM+MOBIL model. Additionally, despite the road section having a speed limit of 33 m/s (120 km/hr), it is observed that drivers frequently exceed this limit, with approximately half of all passenger cars travelling faster than the posted speed. Figure 4.11 presents velocity heatmaps of the freeway section over time.

Since parameters are randomly selected for each vehicle, their behaviour does not tend to replicate the exact observed trajectories. This has been shown in Figure 4.8. The average Fréchet distance for Track 50 is 1.57 for $K = 1$ and 1.59 for $K = 2$. However, we do not seek to align the predicted and actual trajectories perfectly as our objective is to validate the macroscopic pattern of the model.




(A) Observation

(B) 2DDM, $K=1$. D_W : 0.23 for cars, 0.33 for trucks.




(C) 2DDM, $K=2$. D_W : 0.28 for cars, 0.37 for trucks.




(D) IDM+MOBIL. D_W : 0.58 for cars, 0.35 for trucks.

FIGURE 4.10. Comparison between the observed and simulated speed profiles for Track 50. The simulated plot is created by normalising the frequencies of 20 replications. The teal vertical lines indicate the passenger car speed limit of 120 km/hr, while the red lines indicate the truck speed limit of 90 km/hr. The Wasserstein distances (D_W) between the simulated and observed plots are stated in the captions.




(A) Observation

(B) 2DDM, $K=1$



(C) 2DDM, $K=2$



(D) IDM+MOBIL

FIGURE 4.11. Space-time heatmaps for the average velocities for Track 50. The simulated plots are created by normalising the frequencies of 20 replications.

TABLE 4.4. Validation results for Track 01. LC denotes lane-change.

	Observed	2DDM, K=1		2DDM, K=2		IDM+MOBIL	
		Mean	Std	Mean	Std	Mean	Std
Generalised flow (veh/hr/lane)	841.36	849.74	0.52	849.62	0.39	825.47	2.30
Generalised density (veh/km/lane)	7.68	7.64	0.02	7.64	0.03	7.72	0.05
Average speed (m/s)	31.74	31.87	0.09	31.84	0.15	30.54	0.14
Average headway (s)	2.79	2.76	0.01	2.77	0.01	3.07	0.03
LC frequency (LC/hr/lane)	119.84	112.31	10.82	106.52	13.01	99.16	11.59
LC composition (right to left) (%)	41.67	58.46	3.33	58.03	2.71	56.16	3.80
LC composition (left to right) (%)	58.33	41.54	2.98	41.97	2.52	43.84	3.29
Computation time (min)	-	15.27	0.73	29.23	0.82	0.22	0.01

4.4.2 Validation on Non-calibrated Tracks

Apart from Track 50, on which we have conducted calibration, we also analyse the performance of the model on unseen data to test its transferability to other locations and times. For this purpose, we use a 2-lane and a 3-lane track (Tracks 01 and 60). Although these tracks were recorded at different locations and times, their spatial and temporal characteristics should be inherently similar since they were all collected on German highways near Cologne during 2017 and 2018.

Track 01. Track 01 comprises 223 trajectories on a 2-lane highway segment and spans a duration of 15.03 minutes. Compared to Track 50, this track experiences lower traffic flow, and a reduced frequency of lane-changes is observed. Twenty repeated simulations are performed, and the results are aggregated for comparison. The simulated traffic states, alongside the observed values, are displayed in Table 4.4.

The absolute percentage errors are surprisingly low, with most of the values below 10%. Although the errors for some state variables have increased compared to Track 50, the increases are mild. For a model that has been calibrated on a separate dataset, this demonstrates the robustness of 2DDM and its ability to be applied to unseen data. Compared to the IDM+MOBIL model, 2DDM has again achieved lower error across most variables. The composition of the lane-changes remains dissimilar. The original LC composition from lane 1 to 2 are lower than that of lane 2 to 1 while the simulation shows the opposite. Similar to


TABLE 4.5. Validation results for Track 60. LC denotes lane-change.

	Observed	2DDM, K=1		2DDM, K=2		IDM+MOBIL	
		Mean	Std	Mean	Std	Mean	Std
Generalised flow (veh/hr/lane)	851.86	863.20	1.10	863.48	0.66	814.45	4.24
Generalised density (veh/km/lane)	8.37	8.10	0.02	8.10	0.03	7.71	0.04
Average speed (m/s)	29.43	30.51	0.05	30.58	0.13	30.01	0.09
Average headway (s)	3.04	2.99	0.02	2.99	0.01	3.55	0.04
LC frequency (LC/hr/lane)	137.39	128.75	9.59	122.09	12.43	123.5	9.41
LC composition (right to mid) (%)	13.64	28.80	2.20	31.07	2.82	33.24	3.16
LC composition (mid to left) (%)	30.00	32.74	1.80	31.69	2.64	32.97	4.85
LC composition (left to mid) (%)	39.09	27.95	2.87	26.23	1.69	12.86	2.80
LC composition (mid to right) (%)	17.27	10.51	3.29	11.01	2.05	20.93	2.00
Computation time (min)	-	32.78	0.64	65.69	2.00	0.34	0.01

Track 50, lane-changes are made more frequently from the right lane to the left, as vehicles on the right lane attempt to gain speed through lane-changes.


The speed distributions are shown in Figure 4.12. The plots do not match as closely as those previously shown for Track 50, partly because the D_s , which influences speeding tendencies, has not been recalibrated. Unlike Track 50, Track 01 has no posted speed limit, resulting in more platykurtic speed distributions with the maximum recorded speed around 50 m/s (180 km/hr). The observed and simulated speed distributions for trucks are similar, as they are subject to a speed limit of 80 km/hr (22 m/s) even on unrestricted autobahns. Figure 4.13 presents velocity heatmaps of the freeway section over time. The average Fréchet distance for Track 01 is 1.44 for $K = 1$ and 1.48 for $K = 2$. Lower Fréchet distances are observed possibly because manoeuvres on 2-lane segments have fewer options than 3-lane ones. On a 3-lane road, the middle lane can change to both the left and right, introducing extra degrees of freedom.

Track 60. Finally, this model is validated again on a 3-lane track at a different location and time from the calibrated Track 50. The new Track 60 contains 734 vehicles over a duration of 16.01 minutes. The generalised traffic states for both the observed and simulated trajectories are calculated and summarised in Table 4.5.



(A) Observation


(B) 2DDM, $K=1$. D_W : 1.35 for cars, 0.29 for trucks.



(C) 2DDM, $K=2$. D_W : 1.48 for cars, 0.33 for trucks.


(D) IDM+MOBIL. D_W : 1.51 for cars, 0.52 for trucks.

FIGURE 4.12. Comparison between the observed and simulated speed profiles for Track 01. The simulated plot is created by normalising the frequencies of the 20 histograms combined. For this segment, passenger cars are not subject to any speed constraints, while the red lines indicate the truck speed limit of 80 km/hr.



(A) Observation

(B) 2DDM, $K=1$



(C) 2DDM, $K=2$

(D) IDM+MOBIL

FIGURE 4.13. Space-time heatmaps for the average velocities for Track 01. The simulated plots are created by normalising the frequencies of 20 replications.

The absolute percentage errors are again higher than those for Track 50, as the model has not been calibrated on this data. However, the simulated output is still satisfactory, with most performance indicators below 10%. Again, 2DDM performed better than the IDM+MOBIL model for most state variables. There is still a higher tendency for vehicles to change lanes to the left, which differs from the observation.


Regarding the speed profiles, illustrated in Figure 4.14, the match could be further improved by explicitly calibrating the D_s parameter. On this track, which has no posted speed limit, the speeds of passenger vehicles are more evenly distributed than predicted, with the highest speed recorded being approximately 60 m/s (216 km/hr). Figure 4.15 presents velocity heatmaps of the freeway section over time. The average Fréchet distance for Track 60 is 1.64 for $K = 1$ and 1.66 for $K = 2$.

4.4.3 Discussion

In summary, 2DDM can replicate trajectories to produce accurate macroscopic states, with errors for the simulated traffic states generally within 10% of the observed values. The highest errors, lane-changing manoeuvres, tend to be more challenging to capture. The model has demonstrated fairly consistent performance across each of the 20 repeated simulations. The test of transferability shows that the model performs well with unseen data but obviously would achieve even better results with data for which it has been calibrated.


Based on the LC compositions from all three tracks, 2DDM exhibits a larger tendency for vehicles to change lanes to the left rather than to the right. To avoid adding complexity, the current model considers only the essential incentives for making lane-changes, which are speed and safety gains. The effect of lane position on lane-change decisions has been overlooked. Future research could quantify this effect or even incorporate the keep-right directive into the cost formulation of the model.

Interestingly, the $K = 1$ 2DDM tends to slightly outperform $K = 2$ in terms of both the macroscopic measures and the average Fréchet distance. This could indicate that most human



(A) Observation

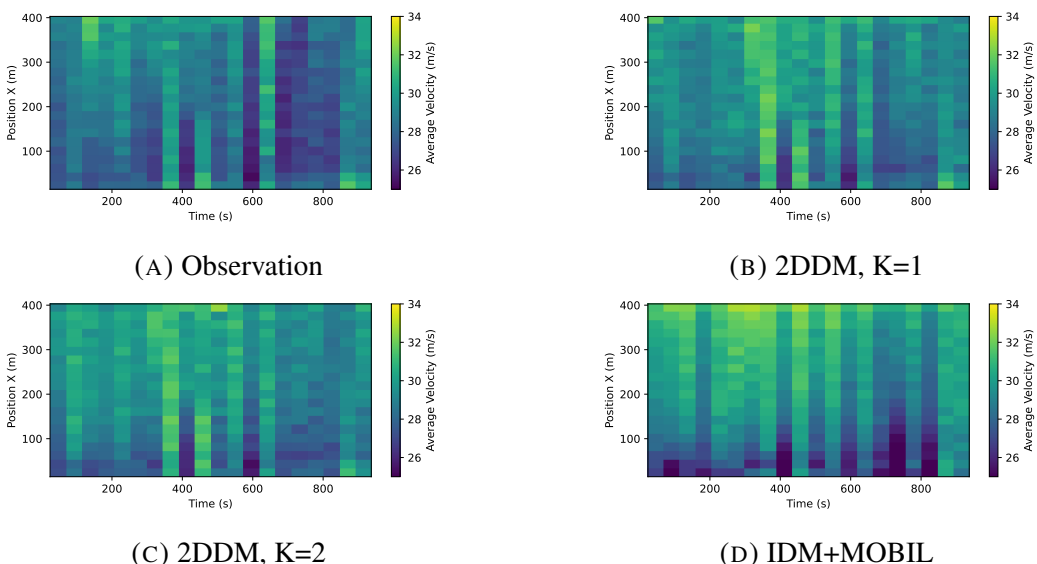
(B) 2DDM, $K=1$. D_W : 1.49 for cars, 0.60 for trucks.



(C) 2DDM, $K=2$. D_W : 1.51 for cars, 0.56 for trucks.

(D) IDM+MOBIL. D_W : 2.11 for cars, 0.80 for trucks.

FIGURE 4.14. Comparison between the observed and simulated speed profiles for Track 60. The simulated plot is created by normalising the frequencies of the 20 histograms combined. For this segment, passenger cars are not subject to any speed constraints, while the red lines indicate the truck speed limit of 80 km/hr.



(A) Observation

(B) 2DDM, $K=1$

(C) 2DDM, $K=2$

(D) IDM+MOBIL

FIGURE 4.15. Space-time heatmaps for the average velocities for Track 01. The simulated plots are created by normalising the frequencies of 20 replications.

drivers don't appear to think strategically when driving. When decisions are updated at a high frequency, people may not anticipate the actions of others.

Based on the average Fréchet distances, 2DDM does not replicate the exact trajectories of the vehicles. For reproducing the microscopic trajectory of a single vehicle, the model can be specifically calibrated using optimisation, although this may lead to overfitting. To understand the impact of changes in calibration parameters on the performance of the model, a sensitivity analysis has been conducted, detailed in Appendix B4. The results indicate that the parameters used are suitable and that the model remains robust under small parameter variations.

Compared to the IDM+MOBIL model, 2DDM consistently outperforms it across most macroscopic measures. This could mainly be attributed to our treatment of integrating longitudinal and lateral motions in a continuous manner. The continuous action space together with the unified objective function allows car-following and lane-changing to be considered at the same time. This means our approach is able to generate smooth, continuous, human-like trajectories. On the other hand, the IDM+MOBIL model assumes vehicles occupy one lane at a time. Lane-changes are thus treated as lateral shifts where the impact of the vehicle on both lanes is ignored.

4.5 Conclusion

The game-theoretical two-dimensional driving model presented in this study integrates the longitudinal and lateral manoeuvres of a vehicle in a continuous manner. Vehicles decide their acceleration and steering actions within a bounded, two-dimensional space. Receding horizon optimisation is applied to simulate real-time prediction and decision-making of drivers. The interactions between the drivers are modelled using level-k game theory to consider the hierarchical reasoning process of humans. Recognising that most existing studies treat human drivers homogeneously, our implementation of 2DDM introduces a degree of stochasticity in people's driving styles. The model is, therefore, heterogeneously defined, treating each vehicle as an individual with its own set of driving preferences. The distributions of these preferences are estimated through calibration with a naturalistic trajectory dataset using

Bayesian inference. The model's performance is validated across three different tracks from the highD dataset, demonstrating its accuracy, consistency, and transferability. The 2DDM is highly flexible. Many of the parameters can be manually or automatically tuned to mimic different driving styles. This study is subject to several limitations. Some can be addressed in future studies:

- In this study, we have assumed that drivers are boundedly rational, meaning they optimise their own objectives based on assumptions of their opponents' reasoning levels. However, people do not behave 'rationally,' even when provided with the same assumptions and information. For example, drivers often overestimate the success rate of making a lane-change or choose not to travel at the highest speed that conditions allow. Future research could quantify the level of randomness in human driving and incorporate stochasticity in human decisions to make actions probabilistic. The optimisation algorithm could be terminated before the optimal actions are found to allow for sub-optimal solutions.
- Another limitation is that the references to passenger cars and trucks pertain only to the ego vehicle. This is because the ego driver does not have access to the parameter sets of surrounding vehicles and must assume they are the same as its own. Future work could explore incorporating behavioural assumptions for other vehicles or updating their parameters through online learning. This would reflect how human drivers adapt their responses to different vehicles by adjusting assumptions about their driving behaviour, such as aggressiveness.
- This study focuses on the modelling of discretionary lane-changes. For the sake of model simplicity and generality, speed and safety gains are assumed to be the major incentives for making lane-changes. A future research direction would be to extend the model to incorporate mandatory lane-change incentives, including considerations for lane drops and future turning manoeuvres. For example, the cost for missing an exit could be directly applied to the objective for making optimal decisions on whether and when to change lanes (Wang et al., 2024). Furthermore, lane-based incentives can be incorporated into the model so that the lane preferences can be

taken into account. The model can be extended to account for the keep-right/left directives by adding a cost into the objective function.

- The calibration and validation of this study were conducted using a single dataset. As a result, the model's performance across different geographical regions and road types remains uncertain. Future research could assess the model's generalisability by evaluating its performance on additional datasets.
- Lastly, future studies could further enhance running time efficiency by compiling the model or reformulating it to enable closed-form analytical solutions that eliminate the need for iterative optimisation.

Recommendations for the applications of the 2DDM are suggested:

- First and foremost, the model can be used to simulate microscopic vehicle behaviours. Validation results have shown that the model is reasonably accurate and consistent in reproducing calibrated human trajectories. The model also performed satisfactorily for unseen data, provided the nature of the data is not significantly different from that which it was calibrated against. The transferability of the model was tested on two other segments recorded in locations and times different to the calibrated track.
- Secondly, the model proposed in this chapter can be adapted to control an AV by mimicking human driving styles. Our 2DDM model allows AVs to adapt to the owner's preferences by calibrating against their historical driving patterns. To use this model as a control framework, several limitations need to be overcome. The online optimisation must be modified to use more efficient algorithms. The objective may need to be transformed or reformulated into linear or convex problems for stability and fast convergence. If possible, analytical solutions to the controls should be derived.
- Third, errors in measurement or prediction must be explicitly accounted for. In this study, we assume certainty equivalence, which is an unrealistic assumption. Both actuators and sensors on AVs are subject to detection errors and delays ([Wang et al., 2023](#)). The robustness of the algorithm must be ascertained before it is applied in practice. The closed-loop stability of the model needs to be tested to ensure the AV

behaves predictably and safely under all expected operational conditions. This can be achieved through theoretical analysis, practical simulations, and real-world trials. Essentially, both internal and external errors need to be accounted for to guarantee the safety, reliability, and performance of the AV.

Decentralised and Communication-free Control of Autonomous Vehicles on Lane-free Freeways

5.1 Introduction

The development of autonomous vehicles (AVs) has seen significant advances in recent years. AVs equipped with advanced sensory systems and efficient actuators offer promising benefits in both travel efficiency and safety ([Van Brummelen et al., 2018](#)). So far, the majority of the research has focused on the development of lane-based control methods, where vehicles travel with lane discipline considerations ([Ali et al., 2019](#), [Balal et al., 2016](#), [Glaser et al., 2010](#), [Ktrakazas et al., 2015](#), [Li et al., 2018](#), [Sankar and Han, 2020](#), [Shi and Li, 2023](#), [Sun and Elefteriadou, 2014](#), [Wang et al., 2015](#), [Yang et al., 2018](#)). Traditionally, lanes have improved both the efficiency and safety of human driving by simplifying the driving task to mainly longitudinal car-following manoeuvres. Lateral lane-changes are considered more complex and riskier, as drivers need to operate in both longitudinal and lateral directions. AVs, on the other hand, can rely on their efficient sensors and fast reactions to navigate complex two-dimensional planes. Recently, there has been increasing interest in the concept of AVs operating on lane-free roads ([Papageorgiou et al., 2021](#), [Troullinos et al., 2021](#)). Unlike traditional roads with designated lanes to guide human drivers, lane-free roads allow AVs to move more flexibly and adaptively. This approach could lead to more efficient use of road space and smoother traffic flow, as AVs optimise their trajectories without the constraints of fixed lanes.

Among the existing studies on lane-free AV control, both connected ([Berahman et al., 2022](#), [Levy and Haddad, 2021](#), [Yanumula et al., 2021](#)) and communication-free (unconnected)

(Karafyllis et al., 2022b, Karalakou et al., 2022) strategies have been explored. In the context of AV systems, connection refers to a vehicle's ability to exchange information with others. The information exchanged can include the current states, predicted trajectories, intended manoeuvres, or control inputs. In contrast, communication-free vehicles rely solely on their onboard sensors to perceive the surrounding environment and nearby vehicles, without sharing or receiving information.

Several notable connected control strategies have been proposed. Levy and Haddad (2021) introduced a nonlinear model predictive controller (NMPC) for centralised control of a group of vehicles on a ring road. The objective function aims to maximise vehicles' progress while reducing control efforts and energy consumption. Constraints are placed to ensure vehicles travel within the road boundary and avoid collisions with each other. The performance of their controller is further validated using reduced-size robots (Levy and Haddad, 2022). Yanumula et al. (2021) proposed a path planning controller that assumes vehicles connected. The current states and short-term future trajectories are communicated between vehicles. These data are used to solve the nonlinear constrained optimal control problem in real-time using a Feasible Direction Algorithm. The authors later expanded their work to consider the effect of vehicle nudging (Yanumula et al., 2023). Similarly, Berahman et al. (2022) introduced repulsive forces to ensure vehicle separations in a connected environment. Deep deterministic gradient policy is used to learn control strategies to maintain a sufficient gap between two consecutive vehicles. Apart from the repulsive forces, attractive forces are likewise studied (Chavoshi and Kouvelas, 2021, Dabestani et al., 2024, Rostami-Shahrbabaki et al., 2023). Rostami-Shahrbabaki et al. (2023) designed a control strategy for vehicle flocking in a connected, lane-free environment. Vehicles are clustered based on the flocking forces realised by energy functions. A consensus algorithm ensures vehicles within the flock are driving at similar speeds. The flock is guided by a virtual leader that generates trajectories at the upper level.

In parallel, a range of communication-free, decentralised controllers have also been developed. Karafyllis et al. (2022b) put forth a nonlinear cruise controller through leveraging Lyapunov functions, barrier functions, and potential functions. Specifically, the potential function ensures collision avoidance with other vehicles and road boundaries by introducing repulsive

effects. Simultaneously, a Lyapunov function is used to guarantee convergence to desired motion profiles, while the barrier function enforces control input limits and safety constraints. The authors also extended their work to design two families of viscous and inviscous cruise controllers (Karafyllis et al., 2022a). These controllers are based on kinetic energy formulations inspired by Newtonian and relativistic mechanics. The resulting microscopic dynamics yield macroscopic fluid-like models, consisting of conservation and momentum equations with controllable pressure and viscosity terms. Karalakou et al. (2022) formulated the problem of single-agent autonomous driving in a lane-free environment and proposed a set of reward functions to address two key objectives: collision avoidance and achieving a target speed. The reward formulation includes terms for longitudinal target tracking, overtaking motivation, collision avoidance, and potential fields to quantify real-time risk based on relative positions and velocities. Using Deep Deterministic Policy Gradient algorithm, they evaluated and compared different reward structures.

While research has been conducted on both connected and unconnected systems, as well as centralised and decentralised control strategies, communication-free decentralised approaches show particular promise. This is especially true in the near to medium term, where full AV penetration remains unlikely. In communication-free decentralised controls, decision-making is distributed to AVs, each independently predicting and optimising its own trajectory. While this may not achieve the full system optimality that centralised methods aspire to, it offers greater robustness and scalability, making it a more realistic option. Moreover, decentralised controllers can still anticipate behaviours of surrounding vehicles using game-theoretic reasoning. The strategic interactions between the AVs are something that have not been explicitly addressed in existing decentralised control studies. In light of the above, this chapter presents a communication-free, decentralised game-theoretic controller that optimises vehicle acceleration and steering in a lane-free environment.

The proposed two-dimensional trajectory controller (2DTC) adopts a receding horizon optimisation framework, where a cost function is minimised over a prediction horizon subject to various constraints. The cost function penalises deviation from the safe velocity, excessive controls, and vehicle sway. Vehicle dynamics are captured using the kinematic bicycle model.

To reflect realistic control capabilities, limits on acceleration and steering are enforced as hard constraints. The presence of other vehicles and road boundaries is handled through inequality constraints that ensure these obstacles are avoided. The resulting problem is a nonlinear optimal control problem (OCP) that is solved online using the continuation/GMRES method, which guarantees convergence within a fixed number of iterations (Ohtsuka, 2004). To model vehicle interactions, Level-k game theory is employed, enabling each vehicle to anticipate the actions of others based on current observations and predefined behavioural assumptions.

The weights of the cost function are tuned through Bayesian optimisation where collision, road departure, and negative heading are treated as constraints. The microscopic performance of 2DTC is evaluated through four simulation experiments featuring various initial conditions and freeway configurations. Results demonstrate smooth and stable trajectories that respect all specified constraints. To assess macroscopic behaviour, large fleets of AVs are simulated on a lane-free freeway with a closing bottleneck. Both K=1 and K=2 reasoning levels are simulated and compared. The resulting fundamental diagrams reveal a higher flow capacity in the lane-free AV system compared to typical lane-based, human-driven traffic. Furthermore, average travel time per vehicle indicates that K=2 outperforms K=1 in terms of efficiency.

5.2 Methodology

5.2.1 Optimal control problem

The OCP is defined as a receding horizon control problem with nonlinear dynamics, cost function, and constraints. Vehicle dynamics are modelled using the kinematic bicycle model, a simplified representation that consolidates the four wheels of an automobile (two front and two rear) into two central wheels, one at the front axle and one at the rear Singh et al. (1991). The state vector of a vehicle at time t is defined by four variables:

$$\mathbf{z}(t) = [x(t), y(t), v(t), \psi(t)]^T \quad (5.1)$$

where $x(t)$ and $y(t)$ denote Cartesian coordinates in the global frame, $v(t)$ is the velocity at the vehicle's centre of mass, and $\psi(t)$ indicates the heading in the global frame. The control inputs are expressed as a vector of two variables:

$$\mathbf{u}(t) = [a(t), \delta(t)]^T \quad (5.2)$$

which correspond to the acceleration $a(t)$ and the front wheel steering angle $\delta(t)$. The velocity and acceleration are applied at the vehicle's centre of mass in the direction of the slip angle $\beta(t)$. The vehicle's motion is governed by a set of nonlinear, coupled ordinary differential equations ($\dot{\mathbf{z}}(t) = \mathbf{f}(\mathbf{z}(t), \mathbf{u}(t))$):

$$\dot{x}(t) = v(t) \cos(\psi(t) + \beta(t)) \quad (5.3a)$$

$$\dot{y}(t) = v(t) \sin(\psi(t) + \beta(t)) \quad (5.3b)$$

$$\dot{v}(t) = a(t) \quad (5.3c)$$

$$\dot{\psi}(t) = \frac{2v(t)}{l_B} \sin(\beta(t)) \quad (5.3d)$$

$$\beta(t) = \tan^{-1} \left[\frac{\tan(\delta(t))}{2} \right] \quad (5.3e)$$

where l_B is the wheelbase, assumed to be 0.6 times the vehicle's length l . In this work, we assume a vehicle's centre of mass coincides with its centroid [Yue et al. \(2015\)](#). The vehicle's dimensions are defined by its length (l) and width (w). Together with the state variables, a vehicle is characterised by its index, x-position, y-position, velocity, heading, length, and width. These dynamic and configuration variables are collectively referred to as a vehicle's extended state, defined as $\mathbf{Z}(t) = [id, x(t), y(t), v(t), \psi(t), l, w]^T$, and hereinafter referred to simply as states. Figure 5.1 illustrates a vehicle in its global frame of reference.

Moreover, the vehicles are subject to physical limits on their control inputs, with the admissible control space defined as a two-dimensional continuous domain bounded by fixed limits:

$$\mathcal{U} = \{(a(t), \delta(t)) \mid a_{\min} \leq a(t) \leq a_{\max}, \delta_{\min} \leq \delta(t) \leq \delta_{\max}\} \quad (5.4)$$




FIGURE 5.1. Kinematic bicycle model for the ego vehicle. The vehicle states are defined using the x , y -positions, the velocity, and the heading $[x, y, v, \psi]$. The control actions are the acceleration and steering angle $[a, \delta]$. Velocity and acceleration are applied on the centroid in the direction of the slip angle (β). The vehicle has dimensions $l \times w$.

The optimal control of an ego vehicle ($\mathbf{u}^*(t)$) is obtained by solving the receding horizon optimisation problem. This is done by minimising the cost function over a prediction horizon subject to the system dynamics, the admissible control limits, and the collision-avoiding inequality constraints:

$$\text{Minimise: } \mathcal{L}(\mathbf{Z}, \mathbf{u}|t_0) = \int_{t_0}^{t_0+T} g(\mathbf{Z}(t), \mathbf{u}(t)) dt$$

$$\text{subject to: } \begin{cases} \dot{\mathbf{Z}}(t) = \mathbf{f}(\mathbf{Z}(t), \mathbf{u}(t)) & \text{ego vehicle dynamics constraint} \\ \dot{\mathbf{Z}}^-(t) = \mathbf{f}(\mathbf{Z}^-(t), \mathbf{u}^-(t)) & \text{opponents dynamics constraint} \\ \mathbf{h}(\mathbf{Z}(t), \mathbf{Z}^-(t)) \leq 0 & \text{collision-avoiding constraint} \\ \mathbf{u}(t) \ \& \ \mathbf{u}^-(t) \in \mathcal{U} & \text{admissible control constraint} \end{cases}$$

where t_0 is the current time, T is the prediction horizon, \mathcal{L} is the cost function, g is the running cost, and \mathbf{h} is a vector of inequality constraints. The vectors \mathbf{Z} and \mathbf{u} represent the state and control inputs of the ego vehicle, respectively. The vectors \mathbf{Z}^- and \mathbf{u}^- denote the state and controls of the opponent vehicles, where $\mathbf{Z}^- = [\mathbf{Z}_j]_{j \in \mathbb{J}^-}$ and $\mathbf{u}^- = [\mathbf{u}_j]_{j \in \mathbb{J}^-}$. Here, \mathbb{J}^- is the set

of opponent vehicles identified by the ego vehicle. The OCP is formulated without terminal conditions.

5.2.2 Continuation/GMRES Method

To solve the above OCP, we discretize the horizon into N finite steps. The problem then becomes:

$$\begin{aligned} \text{Minimise: } \mathcal{L}(\mathbf{Z}, \mathbf{u}|t_0) &= \sum_{\kappa=0}^{N-1} g(\mathbf{Z}(\kappa), \mathbf{u}(\kappa)) \\ \text{subject to: } &\begin{cases} \mathbf{Z}(\kappa + 1) = \mathbf{Z}(\kappa) + \mathbf{f}(\mathbf{Z}(\kappa), \mathbf{u}(\kappa))\Delta t & \text{ego vehicle dynamics constraint} \\ \mathbf{Z}^-(\kappa + 1) = \mathbf{Z}^-(\kappa) + \mathbf{f}(\mathbf{Z}^-(\kappa), \mathbf{u}^-(\kappa))\Delta t & \text{opponents dynamics constraint} \\ \mathbf{h}(\mathbf{Z}(\kappa), \mathbf{Z}^-(\kappa)) \leq 0 & \text{collision-avoiding constraint} \\ \mathbf{u}(\kappa) \ \& \ \mathbf{u}^-(\kappa) \in \mathcal{U} & \text{admissible control constraint} \end{cases} \end{aligned}$$

where κ is the time discretization index and Δt is the step size in time discretization (i.e. $\Delta t := T/N$). For the discrete OCP, $\kappa = 0$ denotes the current time t_0 .

The running cost (g) is defined as a combination of perceived costs. We assume each vehicle has its own preferences regarding the cost components that represent its trade-off between efficiency, safety, and comfort. However, the overall structure of the cost function is the same for all vehicles:

$$g(\mathbf{Z}(\kappa), \mathbf{u}(\kappa)) = \alpha_1 (v(\kappa) - v_{\text{safe}}(\kappa))^2 + \alpha_2 a(\kappa)^2 + \alpha_3 \delta(\kappa)^2 + \alpha_4 \psi(\kappa)^2 - \alpha_{\tilde{\mathbf{u}}} \tilde{\mathbf{u}}(\kappa) \quad (5.5)$$

where $\alpha_1, \alpha_2, \alpha_3,$ and α_4 are weights to the cost components, and $\tilde{\mathbf{u}}(\kappa)$ and $\alpha_{\tilde{\mathbf{u}}}$ are respectively the control dummy variables and their weights. The dummy variables are used to ensure the control inputs follow the admissible control constraints and will be explained later in the section. The safe velocity v_{safe} is determined by:




FIGURE 5.2. A pair of ego-preceding vehicles. The safe velocity is determined based on the relative x-positions and the velocity of the preceding car. Lane boundaries are marked as Ω^+ and Ω^- .

$$v_{\text{safe}}(\kappa) = \min \left(v_d, t_{\text{safe}} a_{\min} + \sqrt{t_{\text{safe}}^2 a_{\min}^2 + v_p(\kappa)^2 - 2a_{\min} \left(x_p(\kappa) - x(\kappa) - \frac{l}{2} - \frac{l_p}{2} \right)} \right) \quad (5.6)$$

where v_d represents the desired velocity, t_{safe} the safe following time of the AV, and x_p , v_p , l_p the x-position, velocity, and length of the preceding car. The preceding vehicle is defined as the closest opponent vehicle located directly ahead of the ego vehicle. A vehicle is considered “directly ahead” if it has a greater x-position than the ego vehicle and its lateral position overlaps with the ego vehicle when projected onto the ego’s front bumper. Figure 5.2 illustrates an ego vehicle following a preceding car on a lane-free road. The cost function we define accounts for four cost components, which are respectively used to penalise the deviation to the safe speed, excessive acceleration/braking, excessive steering, and vehicle sway.

Vehicles avoid causing collisions, assuming other vehicles behave as expected, whether with the road boundaries or with other vehicles. To ensure this, collisions are accounted in inequality constraints defined below:

$$\mathbf{h}(\mathbf{Z}(\kappa), \mathbf{Z}^-(\kappa)) = \begin{bmatrix} \ln(1 + \exp(\Omega^-(x(\kappa)) - y(\kappa) + w/2)) - \epsilon \\ \ln(1 + \exp(y(\kappa) + w/2 - \Omega^+(x(\kappa)))) - \epsilon \\ -\Delta d_j(\kappa) \text{ for } j \in \mathbb{J}^- \end{bmatrix} \leq 0 \quad (5.7)$$

where Ω^+ and Ω^- are the upper and lower road bounds. Softplus functions are used for the road boundary constraints since directly incorporating the linear inequalities would yield

vehicles travelling towards the centre of the road. Δd_j represents the distance between the ego vehicle and the j^{th} opponent in the set of opponents \mathbb{J}^- . This means the size of the vector \mathbf{h} is dependent on the size of the opponent set \mathbb{J}^- . For simplicity, this study assumes a fixed size of 8 surrounding opponents, which means \mathbf{h} has a fixed size of 10. Since the opponent set is updated frequently in the highly dynamic environment, it is difficult to track specific opponents within the vector \mathbf{h} using fixed indices. Instead, we sort the opponents based on their distance to the ego vehicle Δd_j , which means the multipliers would not need to be reinitialised when opponent set changes due to vehicles entering or leaving the detection range. This design enables smooth continuation of the multipliers across time steps.

The distance between an ego-opponent pair is computed by projecting the bounding ellipsoids of the vehicles onto a common axis, and the relative distance is assumed to be the gap between the two projected vehicle outlines. The common axis is the line through the centroids of both vehicles. To determine the relative distance, we first define the ellipsoids that enclose the vehicle bounding boxes. There are infinitely many such ellipses, but we seek the one whose major-to-minor axis ratio matches the vehicle bounding box's length-to-width ratio. The ellipse, aligned with the vehicle's local orientation, is defined by the following formula:

$$\frac{2x^2}{l^2} + \frac{2y^2}{w^2} = 1 \quad (5.8)$$

Furthermore, the projection of an ellipse onto a line can be determined by explicitly differentiating the parametric form of the ellipse to find the tangent gradient. The projected length can be determined based on the observation that the gradient at the maximum x would be zero. The following equation calculates the length of the projection ($w_{\text{proj}}(\kappa)$) for an ellipse oriented in the direction $\psi(\kappa)$ onto a line of projection at angle $\theta(\kappa)$:

$$w_{\text{proj}}(\kappa) = 2\sqrt{\frac{l^2}{2} \sin^2\left(\frac{\pi}{2} - \psi(\kappa) + \theta(\kappa)\right) + \frac{w^2}{2} \cos^2\left(\frac{\pi}{2} - \psi(\kappa) + \theta(\kappa)\right)} \quad (5.9)$$

where $\theta(\kappa)$ is the angle of slope that the line of projection makes with the horizontal axis.

Therefore, the relative distance between two vehicles ($\Delta d_j(\kappa)$) is calculated as:

$$\Delta d_j(\kappa) = \Delta D_j(\kappa) - \frac{w_{\text{proj}}^j(\kappa)}{2} - \frac{w_{\text{proj}}^e(\kappa)}{2} \quad (5.10)$$

where $\Delta D_j(\kappa)$ denotes the distance between the centroids of the two vehicles, and $w_{\text{proj}}^e(\kappa)$ and $w_{\text{proj}}^j(\kappa)$ are respectively the length of projection of the ego vehicle and the j^{th} opponent.

The continuation/GMRES algorithm we use to solve for the optimal trajectory relies on reducing the residual error of the optimisation problem. Since a deviation from the optimal solution is allowed, the complementary slackness condition in the Karush–Kuhn–Tucker conditions may be violated. Therefore, we treat the inequality constraints using the semi-smooth Fischer-Burmeister (FB) method by transforming the inequality constraints into a system of equations that can be solved using standard optimisation techniques. Semi-smoothness implies that the function is locally Lipschitz continuous and have well-defined directional derivatives. The inequality constraint is transformed to:

$$\Xi(\eta(\kappa), -h(\mathbf{Z}(\kappa), \mathbf{Z}^-(\kappa))) = 0 \quad (5.11)$$

where $\eta(\kappa)$ is a Lagrangian multiplier to be treated as an input. The semi-smooth FB function is defined for two variables a and b as:

$$\Xi(a, b) := \sqrt{a^2 + b^2 + \epsilon^2} - (a + b) \quad (5.12)$$

This function has the property that $\Xi(a, b) = 0$ if and only if $a \geq 0$, $b \geq 0$, and $\epsilon = \sqrt{2}ab$. The regularisation term ϵ avoids $\Xi(a, b)$ being undifferentiable when $a = 0$ and $b = 0$ by relaxing the orthogonality $ab = 0$ in the original FB function [Liao-McPherson et al. \(2019\)](#). Note this is deliberate since the complementary slackness condition is relaxed.

Furthermore, the control bounds are also converted to equality constraints by introducing dummy variables $\tilde{a}(\kappa)$ and $\tilde{\delta}(\kappa)$:

$$\mathbf{c}(\mathbf{u}(\kappa), \tilde{\mathbf{u}}(\kappa)) := \left(\mathbf{u}(\kappa) - \frac{\mathbf{u}_{\max} + \mathbf{u}_{\min}}{2} \right)^2 - \left(\frac{\mathbf{u}_{\max} - \mathbf{u}_{\min}}{2} \right)^2 + \tilde{\mathbf{u}}(\kappa)^2 = 0 \quad (5.13)$$

The dummy variables are added to the running cost with weights ($\alpha_{\tilde{\mathbf{u}}}$) to ensure numeric stability at $\tilde{\mathbf{u}}(\kappa) = 0$ and the dummy variables themselves are treated as inputs to be solved. To solve the equality constraint, a Lagrangian multiplier $\mu(\kappa)$ is also added to the list of inputs. The safety constraints and the control bounds are enforced as hard constraints when solving for optimal controls.

We can then define the Hamiltonian as:

$$\mathcal{H}(\mathbf{Z}, \mathbf{Z}^-, \mathbf{u}, \tilde{\mathbf{u}}, \lambda, \eta, \mu) = g(\mathbf{Z}, \mathbf{u}) + \lambda^T \mathbf{f}(\mathbf{Z}, \mathbf{u}) + \eta^T \mathbf{h}(\mathbf{Z}, \mathbf{Z}^-) + \mu^T \mathbf{c}(\mathbf{u}, \tilde{\mathbf{u}}) \quad (5.14)$$

The first-order necessary conditions for the optimal controls, co-states, and multipliers are obtained by the calculus of variation:

$$\frac{\partial \mathcal{H}}{\partial \mathbf{u}}(\mathbf{Z}^*(\kappa), \mathbf{Z}^{*-}(\kappa), \mathbf{u}^*(\kappa), \tilde{\mathbf{u}}^*(\kappa), \lambda^*(\kappa + 1), \eta^*(\kappa), \mu^*(\kappa)) = 0 \quad (5.15)$$

$$\lambda^*(\kappa) = \lambda^*(\kappa + 1) + \frac{\partial \mathcal{H}}{\partial \mathbf{Z}}(\mathbf{Z}^*(\kappa), \mathbf{Z}^{*-}(\kappa), \mathbf{u}^*(\kappa), \tilde{\mathbf{u}}^*(\kappa), \lambda^*(\kappa + 1), \eta^*(\kappa), \mu^*(\kappa)) \Delta t \quad (5.16)$$

$$\lambda^*(N) = \mathbf{0} \quad (5.17)$$

Here, we can define a vector of inputs over the horizon, which includes the controls, dummy variables, and multipliers:

$$\mathbf{U}(t_0) := [\mathbf{u}^*(0), \tilde{\mathbf{u}}^*(0), \eta^*(0), \mu^*(0), \dots, \mathbf{u}^*(N-1), \tilde{\mathbf{u}}^*(N-1), \eta^*(N-1), \mu^*(N-1)] \quad (5.18)$$

Given the input vector and the boundary conditions, the state can be calculated using $\mathbf{Z}(0)$, $\mathbf{Z}^-(0)$, and Equation 5.3. While the costate can be calculated using Equations 5.16 and 5.17. Then, the following must hold if the inputs are optimal:

$$V(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0) := \begin{bmatrix} \frac{\partial \mathcal{H}}{\partial \mathbf{u}}(\mathbf{Z}^*(0), \mathbf{Z}^{*-}(0), \mathbf{u}^*(0), \tilde{\mathbf{u}}^*(0), \lambda^*(1), \eta^*(0), \mu^*(0)) \\ \mathbf{c}(\mathbf{u}^*(0), \tilde{\mathbf{u}}^*(0)) \\ \Xi(\eta^*(0), -h(\mathbf{Z}^*(0), \mathbf{Z}^{*-}(0))) \\ \vdots \\ \frac{\partial \mathcal{H}}{\partial \mathbf{u}}(\mathbf{Z}^*(N-1), \mathbf{Z}^{*-}(N-1), \mathbf{u}^*(N-1), \tilde{\mathbf{u}}^*(N-1), \lambda^*(N), \eta^*(N-1), \mu^*(N-1)) \\ \mathbf{c}(\mathbf{u}^*(N-1), \tilde{\mathbf{u}}^*(N-1)) \\ \Xi(\eta^*(N-1), -h(\mathbf{Z}^*(N-1), \mathbf{Z}^{*-}(N-1))) \end{bmatrix} = 0 \quad (5.19)$$

In other words, Equation 5.19 represents the necessary condition for optimality over the horizon, which can be solved directly using iterative search such as multiple-shooting methods. However, this procedure is inefficient and does not guarantee convergence in a fixed number of steps. Therefore, a continuation method is employed by choosing $\mathbf{U}(0)$ so that $V(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|0) = 0$, then we track $\mathbf{U}(t_0)$ satisfying $V(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0) = 0$ by integrating $\dot{\mathbf{U}}(t_0)$, such that:

$$\dot{V}(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0) = -vV(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0) \quad (5.20)$$

where v is a positive stabilising constant. The stabilising constant is introduced since the optimality condition does not need to be exactly satisfied. Then, we obtain a differentiable equation for $\mathbf{U}(t_0)$:

$$\dot{\mathbf{U}}(t_0) = \frac{\partial V}{\partial \mathbf{U}}^{-1}(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0) \left(-vV(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0) - \frac{\partial V}{\partial \mathbf{Z}}(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0)\dot{\mathbf{Z}}(t_0) - \frac{\partial V}{\partial t}(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0) \right) \quad (5.21)$$

The product $\frac{\partial V}{\partial \mathbf{Z}}(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^-|t_0)\dot{\mathbf{Z}}(t_0)$ can be computed using the finite difference approximation that requires only one evaluation of the vector V . Equation 5.21 is solved using the GMRES

method without conducting matrix inversion on $\frac{\partial V}{\partial \mathbf{U}}(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^- | t_0)$ Shimizu et al. (2009). After obtaining $\dot{\mathbf{U}}(t_0)$, the solution is updated based on the sampling interval $\Delta\tau$:

$$\mathbf{U}(t_0 + \Delta\tau) = \mathbf{U}(t_0) + \dot{\mathbf{U}}(t_0)\Delta\tau \quad (5.22)$$

To obtain the initial solution $\mathbf{U}(0)$ satisfying $V(\mathbf{U}, \mathbf{Z}, \mathbf{Z}^- | 0) = 0$, we solve the following iteratively using the Newton's method for a specified number of iterations Ohtsuka (2004).

$$\begin{bmatrix} \frac{\partial \mathcal{H}}{\partial \mathbf{u}}(\mathbf{Z}(0), \mathbf{Z}^-(0), \mathbf{u}(0), \tilde{\mathbf{u}}(0), 0, \eta(0), \mu(0)) \\ \mathbf{c}(\mathbf{u}(0), \tilde{\mathbf{u}}(0)) \\ \Xi(\eta(0), -h(\mathbf{Z}(0), \mathbf{Z}^-(0))) \end{bmatrix} = 0 \quad (5.23)$$

5.2.3 Level-k Game Theory

Most existing studies overlook the potential actions of other opponents. Our model addresses this limitation by predicting their future states and controls over the prediction horizon. To achieve this, we employ level-k game theory to capture the strategic interactions between vehicles. In a decentralised and connection-free system, each agent aims to optimise its actions by anticipating the behaviour of others. This creates a recursive loop, as agents must predict what others predict about themselves. The process may not converge and raises the question of where to begin. Level k game theory addresses this by defining a base level, level 0, where agents act without modelling others. Level 1 agents then optimise their actions assuming their opponents are level 0 agents. Similarly, level 2 agents respond to level 1 behaviour, and the process continues until it reaches a predefined maximum level K .

In this work, we assume that AVs understand the structure of their opponents' objective functions. Since the exact cost weights of the opponents are unknown, each AV assumes that all other vehicles use the same weights as itself. We also adopt the assumption of certainty equivalence, meaning the ego vehicle treats the actual states of its opponents as if they are perfectly estimated. This is a reasonable approximation in scenarios involving a fleet of AVs.

We further assume that each AV has a detection range of 200 meters (100 meters in each direction) and considers the nearest eight vehicles within this range as opponents.

The ego vehicle is assigned a rationality level K , which reflects its depth of strategic reasoning in the game. A higher value of K corresponds to a greater ability to anticipate the actions of others. In level- k game theory, each player assumes that their opponents operate at one level lower. That is, a level k vehicle assumes all its opponents are at level $k - 1$ and optimises its strategy accordingly. In this study, level 0 vehicles are modelled as travelling at constant velocities, with both acceleration and steering set to zero.

For an ego vehicle with a rationality level K , its optimal actions (\mathbf{u}^{K*}) can be computed by iteratively calculating the optimal controls for vehicles from level 1 to K . Let $k \in \{1, 2, \dots, K\}$. Then for each k , compute:

$$\mathbf{u}^{k*} = \arg \min_{\mathbf{u}^k} \mathcal{L}(\mathbf{Z}^k, \mathbf{u}^k | t_0)$$

$$\text{subject to: } \begin{cases} \mathbf{Z}^k(\kappa + 1) = \mathbf{Z}^k(\kappa) + \mathbf{f}(\mathbf{Z}^k(\kappa), \mathbf{u}^k(\kappa))\Delta t & \text{vehicle dynamics} \\ \mathbf{h}(\mathbf{Z}^k(\kappa), \mathbf{Z}^{k-1}(\kappa)) \leq 0 & \text{collision-avoiding} \\ \mathbf{u}^k(\kappa) \ \& \ \mathbf{u}^{k-1}(\kappa) \in \mathcal{U} & \text{admissible control} \end{cases}$$

where \mathbf{Z}^k denotes the states of level k vehicles, \mathbf{u}^k denotes the actions adopted by the level k vehicles, and $\mathbf{u}^{0*} = \mathbf{0}$. The problem is discretized and solved using the continuation/GMRES method discussed in Section 5.2.2. The entire framework is summarised in the Algorithm 3.

5.2.4 Parameter tuning

To ensure optimal performance by our 2DTC, we tune its parameters using a Bayesian optimisation framework. The variables to be tuned include the four weights in the cost function ($\alpha_1, \alpha_2, \alpha_3, \alpha_4$). Considering that the sub-cost weights can range substantially, we use a log scale for tuning.

Algorithm 3 One Evaluation of the Optimal Control for the Ego Vehicle

Input:

$\mathbf{Z}(t_0)$: states of all vehicles at t_0
 $\mathbf{U}(t_0)$: inputs of all vehicles over the prediction horizon
 K : predefined cognitive level of the ego vehicle
 k_{\max} : maximum number of iterations for the GMRES algorithm
 $\Delta\tau$: sampling time interval
 ϵ_t : finite difference step size
 v : stabilising coefficient
 FORWARD_KINEMATICS($t_0, \mathbf{x}, \mathbf{u}$): computes the state updates
 CAL_V($\mathbf{Z}, \mathbf{Z}_{\text{opp}}, \mathbf{U}, \mathbf{U}_{\text{opp}}, t$): computes V
 GMRES($\mathbf{A}, \mathbf{b}, \text{maxiter}$): solves $\mathbf{A}\mathbf{x} = \mathbf{b}$

Output:

$\mathbf{U}^K[: 2]$: control actions of the ego vehicle

```

1: procedure SOLVE_NMPC( $t_0, \mathbf{Z}(t_0), \mathbf{U}(t_0)$ )
2:   Assign ego vehicle as  $\mathbf{Z}^K(t_0)$ 
3:   Initialise  $k \leftarrow K$ 
4:   while  $k \neq 0$  do:
5:     for  $player$  in  $\mathbf{Z}^k(t_0)$  do
6:        $\mathbf{Z}_{player}^{k-1}(t_0) \leftarrow$  opponents of  $\mathbf{Z}_{player}^k(t_0)$  from  $\mathbf{Z}(t_0)$ 
7:     end for
8:      $k \leftarrow k - 1$ 
9:   end while
10:   $\mathbf{U}^0(t_0) \leftarrow \mathbf{0}$ 
11:  Initialise  $k \leftarrow 1$ 
12:  while  $k \neq K$  do
13:    for  $player$  in  $\mathbf{Z}^k(t_0)$  do
14:       $\dot{\mathbf{U}} \leftarrow \mathbf{0}$ 
15:       $\mathbf{b} \leftarrow$  EVAL_B( $t_0, \mathbf{Z}_{player}^k(t_0), \mathbf{Z}_{player}^{k-1}(t_0), \mathbf{U}_{player}^k(t_0), \mathbf{U}_{player}^{k-1}(t_0), \dot{\mathbf{U}}$ )
16:       $\mathbf{A} \leftarrow$  Linear operator based on EVAL_A( $t_0, \mathbf{Z}_{player}^k(t_0), \mathbf{Z}_{player}^{k-1}(t_0), \mathbf{U}_{player}^k(t_0), \mathbf{U}_{player}^{k-1}(t_0), \dot{\mathbf{U}}$ )
17:       $\dot{\mathbf{U}} +=$  GMRES( $\mathbf{A}, \mathbf{b}, \text{maxiter} = k_{\max}$ )
18:       $\mathbf{U}_{player}^k(t_0) += \Delta\tau \dot{\mathbf{U}}$ 
19:    end for
20:     $k \leftarrow k + 1$ 
21:  end while
22:  return  $\mathbf{U}^K[: 2]$ 
23: end procedure
24:
25: function EVAL_B( $t_0, \mathbf{Z}, \mathbf{Z}_{\text{opp}}, \mathbf{U}, \mathbf{U}_{\text{opp}}, \dot{\mathbf{U}}$ )
26:    $t_1 \leftarrow t_0 + \epsilon_t$ 
27:    $\dot{\mathbf{Z}} \leftarrow$  FORWARD_KINEMATICS( $t_0, \mathbf{Z}, \mathbf{U}[: 2]$ )
28:    $\mathbf{Z}_1 \leftarrow \mathbf{Z} + \epsilon_t \dot{\mathbf{Z}}$ 
29:    $\mathbf{U}_1 \leftarrow \mathbf{U} + \epsilon_t \dot{\mathbf{U}}$ 
30:    $\mathbf{b} = (\frac{1}{\epsilon_t} - v) \text{CAL\_V}(\mathbf{Z}, \mathbf{Z}_{\text{opp}}, \mathbf{U}, \mathbf{U}_{\text{opp}}, t) - \frac{\text{CAL\_V}(\mathbf{Z}_1, \mathbf{Z}_{\text{opp}}, \mathbf{U}_1, \mathbf{U}_{\text{opp}}, t_1)}{\epsilon_t}$ 
31:   return  $\mathbf{b}$ 
32: end function
33:
34: function EVAL_A( $t, \mathbf{Z}, \mathbf{Z}_{\text{opp}}, \mathbf{U}, \mathbf{U}_{\text{opp}}, \dot{\mathbf{U}}$ )
35:    $t_1 \leftarrow t + \epsilon_t$ 
36:    $\mathbf{U}_1 \leftarrow \mathbf{U} + \epsilon_t \dot{\mathbf{U}}$ 
37:    $\mathbf{A} \leftarrow \frac{\text{CAL\_V}(\mathbf{Z}_1, \mathbf{Z}_{\text{opp}}, \mathbf{U}_1, \mathbf{U}_{\text{opp}}, t_1) - \text{CAL\_V}(\mathbf{Z}_1, \mathbf{Z}_{\text{opp}}, \mathbf{U}, \mathbf{U}_{\text{opp}}, t_1)}{\epsilon_t}$ 
38:   return  $\mathbf{A}$ 
39: end function

```

TABLE 5.1. Fixed parameters for the controlled vehicles

Level of ego vehicle in the level-k game (K) [-]	2
Sampling time ($\Delta\tau$) [s]	0.01
FB epsilon (ϵ) [-]	0.01
Dummy variable weights ($\alpha_{\bar{u}}$) [-]	1
Maximum number of iterations in GMRES (k_{\max}) [-]	20
Stabilisation parameter in continuation/GMRES (ν) [-]	10
Finite difference step size (ϵ_t) [-]	1e-6
Prediction horizon length (T) [s]	2
Number of discretization over the prediction horizon (N) [-]	20
Safe following time (t_{safe}) [s]	0.5
Acceleration bounds (a_{\min}, a_{\max}) [m/s ²]	[-3.4, 3]
Steering bounds ($\delta_{\min}, \delta_{\max}$) [rad]	$[-\pi/9, \pi/9]$

The simulation environment is defined for a hypothetical road section with a bottleneck. Vehicles are randomly generated with a minimum headway of 0.5 s. The length, width, and desired velocity of all vehicles are randomly drawn from normal distributions $l \sim \mathcal{N}(5, 0.5)$, $w \sim \mathcal{N}(2, 0.2)$, and $v_d \sim \mathcal{N}(30, 2)$. All vehicles are controlled using our algorithm and the fixed parameters are summarised in Table 5.1. Each simulation is run for a period of 300 s.

The tuning procedure is defined as a constrained optimisation problem where the objective is for the vehicles to travel as close as possible to their desired velocities. Constraints are defined on vehicle collisions, road departure, and negative headings. The optimisation is conducted using an open source Bayesian optimisation library for 20 iterations (Nogueira, 2014). The tuned parameters for the cost function (5.5) are as follows:

$$\alpha^* = [10^{0.52}, 10^{0.29}, 10^{0.33}, 10^{-0.88}] \quad (5.24)$$

5.3 Simulation Results

The proposed 2DTC for $K = 2$ is first tested using simulated scenarios for its microscopic performance. Then, its macroscopic pattern is evaluated by simulating a large fleet of AVs similar to how the model is tuned. All vehicles are simulated based on our proposed model detailed in Section 5.2 with fixed parameters from Table 5.1 and tuned weights

in Equation 5.24. Slight heterogeneity in the cost weights is introduced for macroscopic analysis of the AV fleets. The controller is implemented in C++17 and compiled using Clang 15.0.0 with CMake 3.28.3 as the build system. The Eigen3 library is used for linear algebra operations. All computations are conducted on macOS Sonoma 14.5, equipped with an Apple M3 chip and 18 GB of unified memory.

5.3.1 Microscopic Behaviour


The proposed control framework is evaluated using four simulation scenarios, each lasting 50 s. In every scenario, four vehicles are simulated with different initial conditions. The freeway environment remains the same across the first three scenarios, featuring a road width of 12 m. The fourth scenario has a slight modification with a physical bottleneck that narrows the road width to 8 m. All vehicles are assumed to operate at $K = 2$ and have the same set of cost weights, as specified in Equations 5.24.

The first scenario involves four vehicles positioned sparsely apart. The initial vehicle configurations are summarised in Table 5.3. The simulation results for scenario 1 are shown in Figure 5.4. Vehicles avoid collisions with both one another and the road boundaries. Although they occasionally come close, they maintain safety by moving laterally in opposite directions. When no preceding vehicles are present, they converge to their desired velocities. In contrast, when a slower vehicle is ahead, they exhibit car-following behaviour by reducing speed. For instance, Vehicle 1 initially slows down due to the car-following cost, but after overtaking the slower Vehicle 3, it accelerates to reach its desired velocity of 25 m/s. Similarly, Vehicle 2, which has the highest desired velocity, decelerates in response to the slower Vehicle 4. After briefly reducing their speeds, both Vehicles 1 and 2 successfully overtake their preceding vehicles.


It is worth noting that overtaking does not always require the following car to manoeuvre laterally. At around 38 s, Vehicle 4 yielded to its follower by shifting laterally to the right. This behaviour may reflect Vehicle 4's anticipation of Vehicle 2's potential actions. Since

FIGURE 5.3. Initial vehicle configurations for scenario 1


	Desired speed (v_d)	x-position (x)	y-position (y)	velocity (v)	heading (ψ)	Length (l)	Width (w)
Vehicle 1	25 m/s	0 m	4 m	25 m/s	0 rad	5 m	2 m
Vehicle 2	35 m/s	100 m	7.5 m	30 m/s	0 rad	7 m	3 m
Vehicle 3	20 m/s	150 m	4 m	20 m/s	0 rad	5 m	3 m
Vehicle 4	30 m/s	250 m	6.5 m	20 m/s	0 rad	7 m	3 m




(A) x-position-time diagram




(B) y-position-time diagram




(C) Velocity-time diagram



(D) Heading-time diagram



(E) Acceleration-time diagram



(F) Steering-time diagram

FIGURE 5.4. Change in states and control inputs over the simulated period for scenario 1 (initially sparse vehicles). The four state variables are x-position, y-position, velocity, and heading. The control inputs are acceleration and steering angle.

Vehicle 2 was approaching at a higher speed, it was safer for the leader to steer right and allow the follower to pass.

From Figures 5.4e and 5.4f, all four vehicles have their acceleration and steering within the specified bounds. Vehicles reduce speed when approaching a slower preceding car and steer in opposite directions when overtaking or being overtaken. Overall, the computed controls result in smooth trajectories. An animation of scenario 1 is displayed in Figure C.1 in Appendix C1.

The second set of simulations is conducted on closely initialised vehicles. Once again, four vehicles are used and their initial configurations are summarised in Table 5.5. The simulated states and controls are shown in Figure 5.6. Vehicles consistently avoid one another as they come close, owing to the collision-avoidance constraint. In the first few seconds, Vehicles 1 and 3 overtake the slower Vehicles 2 and 4 by moving laterally to the left. Vehicle 2 initially makes room for Vehicle 1 by shifting to the right. After Vehicle 1 passes, Vehicle 2 also moves left to overtake the slower Vehicle 4 and eventually reaches its desired velocity at around 20 s. The slowest vehicle, Vehicle 4, first shifts left—possibly to yield to Vehicle 2. It then quickly moves right in response to the faster Vehicle 1 approaching from behind.


The velocity–time plot further illustrates that vehicles converge to their desired speeds. When a preceding car is present, a vehicle first reduces its velocity, then gradually accelerates after overtaking. This behaviour is most apparent in Vehicle 2, which maintains a safe distance behind its leader before reaching its desired velocity after bypassing Vehicle 4.

The control inputs are illustrated in Figures 5.6e and 5.6f. Once again, the acceleration and steering for all four vehicles remain within the specified limits. Vehicles reduce speed as they approach slower preceding cars and steer towards opposite directions when coming closer. In this simulation, Vehicles 1 and 2 brake heavily to ensure safety at the start of the frame. The steering is also more aggressive since vehicles are initialised in a complex, unsafe situation. An animation of scenario 2 is displayed in Figure C.2.


The third scenario consists of one vehicle that was initially stopped. The initial configurations are summarised in Table 5.7. The change in states and controls over the 50 s simulation is presented in Figure 5.8. Vehicle 4, initially stationary and positioned ahead of the others,

FIGURE 5.5. Initial vehicle configurations for scenario 2


	Desired speed (v_d)	x-position (x)	y-position (y)	velocity (v)	heading (ψ)	Length (l)	Width (w)
Vehicle 1	30 m/s	0 m	8 m	30 m/s	0 rad	5 m	2 m
Vehicle 2	25 m/s	25 m	6 m	20 m/s	0 rad	7 m	3 m
Vehicle 3	30 m/s	45 m	5.5 m	25 m/s	0 rad	5 m	3 m
Vehicle 4	20 m/s	60 m	3.5 m	15 m/s	0 rad	7 m	3 m




(A) x-position-time diagram




(B) y-position-time diagram




(C) Velocity-time diagram



(D) Heading-time diagram



(E) Acceleration-time diagram



(F) Steering-time diagram

FIGURE 5.6. Change in states and control inputs over the simulated period for scenario 2 (initially close vehicles). The four state variables are x-position, y-position, velocity, and heading. The control inputs are acceleration and steering angle.

gradually accelerates toward its desired velocity of 15 m/s. The remaining vehicles decelerate when approaching Vehicle 4, then steer to the right to bypass it. Vehicle 1, with the highest desired velocity, has not had the chance to reach its desired states as there are always slower vehicles ahead. All other vehicles reach their desired velocities when preceding cars are not hindering their paths.

The plots for control variables are presented in Figures 5.8e and 5.8f. Once again, both acceleration and steering angle are within their designated bounds for all four vehicles. Figure C.3 shows the animation of scenario 3.

Finally, we showcase an extension of scenario 3 by placing a bottleneck that reduces the road width from 12 m to 8 m. The bottleneck is located at 800 m. The initial setup is the same as scenario 3 displayed in Table 5.7. Figure 5.9 illustrates the states and controls for the four vehicles over the simulation. Since the vehicles are initialised in the same way as in Scenario 3, the variables in the first half of the simulation are quite similar. After Vehicles 1, 2, and 3 overtake the slower Vehicle 4, they approach the bottleneck. Owing to the narrowing road boundaries, Vehicle 2 narrowly passes Vehicle 3 just before the road width reduces to 8 m. Toward the end of the simulation, Vehicle 3 begins to yield to the following Vehicle 1.

The control inputs for the vehicles are illustrated in Figures 5.9e and 5.9f. It can be seen that Vehicles 2 and 3 exhibit small lateral shifts while travelling side by side in close proximity. The vehicles constantly adjust their steering to maintain safe distances from one another and from the road boundaries. An animation of scenario 4 is presented in Figure C.4.


5.3.2 Macroscopic Behaviour

In addition to showcasing 2DTC's microscopic performance, we also simulated a fleet of AVs in a lane-free environment and observe their macroscopic behaviour.


Vehicles are generated at the beginning of a freeway section in a stochastic manner. The arrivals are sampled probabilistically at each time step using a normal distribution, with a time-varying inflow rate governed by a sigmoid function. This results in a time-dependent

FIGURE 5.7. Initial vehicle configurations for scenario 3


	Desired speed (v_d)	x-position (x)	y-position (y)	velocity (v)	heading (ψ)	Length (l)	Width (w)
Vehicle 1	30 m/s	0 m	5 m	15 m/s	0 rad	5 m	2 m
Vehicle 2	25 m/s	75 m	6 m	25 m/s	0 rad	7 m	3 m
Vehicle 3	20 m/s	150 m	3 m	20 m/s	0 rad	5 m	3 m
Vehicle 4	15 m/s	200 m	8 m	0 m/s	0 rad	7 m	3 m




(A) x-position-time diagram




(B) y-position-time diagram




(C) Velocity-time diagram



(D) Heading-time diagram



(E) Acceleration-time diagram



(F) Steering-time diagram

FIGURE 5.8. Change in states and control inputs over the simulated period for scenario 3 (one initially stopped vehicle). The four state variables are x-position, y-position, velocity, and heading. The control inputs are acceleration and steering angle.

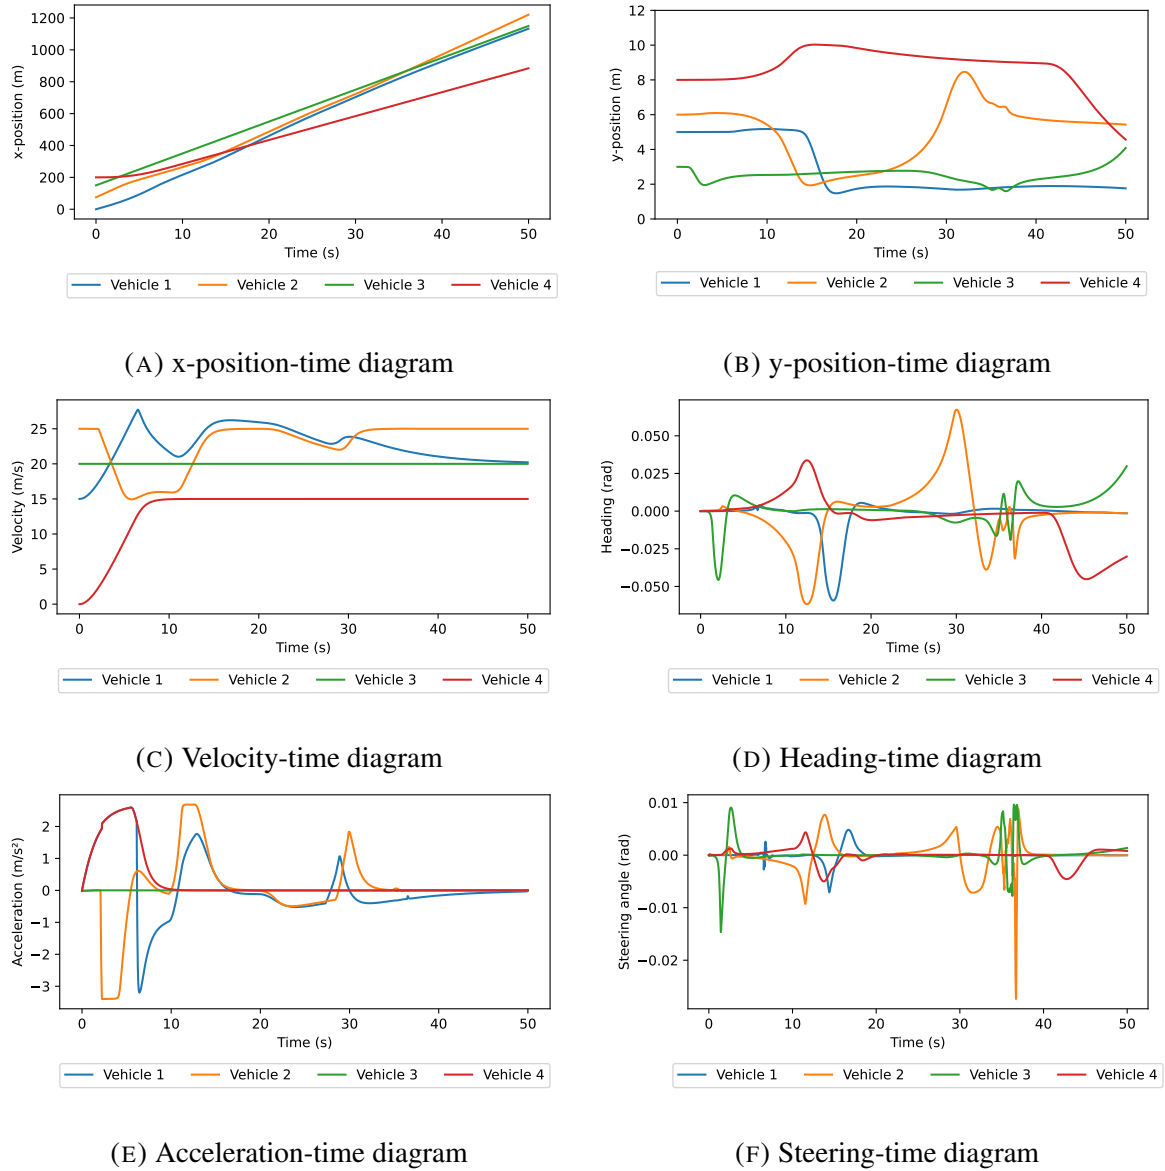


FIGURE 5.9. Change in states and control inputs over the simulated period for scenario 4 (one initially stopped vehicle with a bottleneck). The four state variables are x-position, y-position, velocity, and heading. The control inputs are acceleration and steering angle.

arrival probability that follows the shape of the sigmoid. The flow increases gradually to some maximum threshold. The initial y-position of vehicles are generated from a uniform distribution. To ensure vehicles are not prone to collision when they are first generated, we define a minimum time headway of 0.3 s and an initial speed of 30 m/s. As in Section 5.2.4, the dimensions of the vehicles are drawn from normal distributions $l \sim \mathcal{N}(5, 0.5)$ and

$w \sim \mathcal{N}(2, 0.2)$, and the desired velocity from $v_d \sim \mathcal{N}(30, 2)$. The weights of the cost components are likewise heterogeneous and sampled with the tuned weights α^* as means $([10^{0.52}, 10^{0.29}, 10^{0.33}, 10^{-0.88}])$ and one-tenth the tuned values as the standard deviations $([10^{-0.48}, 10^{-0.71}, 10^{-0.67}, 10^{-1.88}])$. This is based on the assumption that the weights are independent, which is showcased in our previous model on human-driving (Chapter 4).

The simulated freeway section is 500 m. The upstream road width is 12 m, similar to a three-lane freeway. At 350 m, a bottleneck is placed to reduce the road width from 12 m to 8 m by gradually narrowing the left side. Five loop detectors are placed over the freeway section with an interval of 100 m. The configuration of the freeway is illustrated in Figure 5.10a. To ensure vehicles prepare for merging, we add a smooth non-decreasing cost that gradually increases when approaching the bottleneck.

$$g_{\text{merge}} = \frac{\max(0, y(\kappa) - \Omega_r^+)}{1 + \exp(-\alpha_r(x(\kappa) - x_r))} \quad (5.25)$$

where Ω_r^+ denotes the upper road boundary of the narrow section, α_r represents a scaling variable, which we set to 0.01, and x_r is the x-position of the bottleneck.


Traffic flow (Q) and density (ρ) are calculated based on the loop detector recordings using the following equations:

$$Q = \frac{n_{\text{sample}}}{T_{\text{sample}}} \quad (5.26)$$

$$\rho = \frac{T_{\text{occupied}}}{T_{\text{sample}} \bar{l}_{\text{sample}}} \quad (5.27)$$

where T_{sample} denotes the sampling period set to 30 s, n_{sample} is the number of vehicles passing over the loop detector during the sampling period, T_{occupied} is the time that the loop detector is occupied within the sampling period, and \bar{l}_{sample} the average vehicle length of passing vehicles. Note that loop detectors are assumed as thin lines with zero length. Each simulation is run for a period of 5 min. We repeated the simulation multiple times with varying levels of maximum traffic flow then aggregate the results.

We simulated AV fleets with K values of both $K = 2$ and $K = 1$, which are respectively presented in Figures 5.10b and 5.10c. Both sets of plots are similar across various locations.



(A) Configuration of the simulated freeway section and the loop detectors.



(B) FDs generated using $K = 2$.(C) FDs generated using $K = 1$.

FIGURE 5.10. Fundamental diagrams of flux (veh/hr/m) vs spissitude (veh/m²). The FDs are constructed based on loop detectors at different locations specified in Figure 5.10a. The locations of the loop detectors are 50 m, 150 m, 250 m, 350 m, 450 m. The colours of the FDs align with the colours of the loop detectors. The input flow is stochastic and governed by a time-varying sigmoid function.

It can be seen that the bottleneck induces mild congestion in the narrow section. Congestion only occurs near the bottleneck as vehicles accelerate again when their preceding vehicles exit the simulated section.

Based on the fundamental diagrams, the capacity of the lane-free traffic for both $K = 2$ and $K = 1$ is approximately 1,000 veh/hr/m. With an approximation of 4-meter “lanes”, the capacity converts to 4,000 veh/hr/lane. From empirical evidence, the capacity of human-driving vehicles tends to be below 2,500 veh/hr/lane (Hall et al., 1986, Mannering and Washburn, 2020). This showcases a significant increase in capacity for an all AV lane-free scenario. The reasons for this improvement could be twofold. First, AVs having a shorter reaction time than human vehicles can have a shorter car-following distance, improving longitudinal efficiency. Second, due to the lane-free setup, vehicles can effectively fill gaps

in the traffic, improving lateral efficiency. Both reasons result in vehicles flowing in a more compact manner and hence resulting in higher traffic capacity observed.

Although this suggests an improved flow capacity from higher density, the influence of slower vehicles becomes more pronounced due to increased interactions and reduced headways. In particular, slower vehicles can reduce the average speed of surrounding traffic, even in free-flow conditions. With higher heterogeneity in desired speeds, the reduction in overall speed could be significant.

The space-mean speed could be determined easily using the fundamental diagram properties. Based on the slope, the free-flow speed is approximately 28 m/s (~ 100 km/h), which is lower than the average desired velocity and the initial velocity at generation (both at 30 m/s). This discrepancy arises due to the aforementioned reason. Slower vehicles, though fewer in number, impose greater influence on the traffic stream by causing faster-following vehicles to reduce their speed. This effect is due to asymmetric interaction dynamics in car-following behaviour, where slower vehicles act as moving bottlenecks even in uncongested conditions.

Over the 15,301 total vehicle kilometres travelled in simulation, no anomalous behaviour was observed. No vehicles were engaged in collisions nor departed from the road. For $K = 2$, the average computation time for the 0.01 s interval is 0.0024 s and the maximum computation time was 0.0056 s. These values are computed based on 28,070,231 instances of computations conducted. As for $K = 1$, the average computation time is 0.0004 s and the maximum computation time is 0.0005 s. Therefore, the algorithm has been shown to perform computations in real-time with a typical desktop computer described above.

One should note that at $K = 2$, the ego vehicle performs optimisation at most 9 times. 8 for the opponents and 1 for itself. While at $K = 1$, the ego vehicle only optimises once. Therefore, as K increases, the computation time increases (in worst cases) exponentially. With more advanced processors and optimisers, the computation is expected to be even faster and should not be a computational bottleneck when implemented.




FIGURE 5.11. Distributions of the travel times to traverse the 500 m section. All AVs within the fleets are either controlled at $K = 2$ or $K = 1$.

5.3.3 Comparing $K = 1$, $K = 2$, and Mixture

To understand the difference between the $K = 1$ and $K = 2$ traffic, we quantify vehicle efficiency using the time taken to traverse the 500 m freeway section. The distributions of the travel times are illustrated in Figure 5.11. Since each simulation is repeated 20 times, the average and median travel times provide reliable estimates of the actual vehicle efficiency. Vehicles controlled by the $K = 1$ controller had an average travel time of 17.51 s, compared to 17.23 s for those using the $K = 2$ controller. This represents a modest improvement of 0.28 s for each $K = 2$ vehicle. While the difference may appear minor on a per-vehicle basis, it accumulates to approximately 145 s (or 2.43 min) when aggregated over the roughly 520 vehicles generated during a 5-minute simulation. Furthermore, the distributions are slightly right skewed, and the respective medians for $K = 1$ and $K = 2$ are 17.45 s and 17.17 s.

In addition, we have simulated a mixture of both $K = 1$ and $K = 2$ vehicles in the same fleet. Vehicles are generated randomly with a 50% chance of being either $K = 1$ or $K = 2$. Figure 5.12 depicts the distributions of the travel times for $K = 1$ and $K = 2$ AVs. In this case, the average travel times are very similar with a mean of 17.48 s for $K = 1$ and 17.46 s for $K = 2$. This could suggest that with a mixed traffic consisting of AVs at both $K = 1$




FIGURE 5.12. Distributions of the travel times to traverse the 500 m section. The AV fleets are controlled with half of the vehicles at $K = 2$ and the other half at $K = 1$.

and $K = 2$, the efficiency improvement is reduced. Furthermore, the results still suggest that higher levels of anticipation enabled by the $K = 2$ controller allow ego vehicles to better optimise their trajectories, resulting in them achieving more efficient control decisions. The medians for $K = 1$ and $K = 2$ are 17.37 s and 17.33 s, respectively.

5.3.4 Comparing Levels of Heterogeneity in Vehicle Size

We next examine the impact of vehicle size heterogeneity on overall traffic efficiency. To keep computation costs low, the following tests are conducted with $K = 1$. Vehicle dimensions are drawn from normal distributions, and heterogeneity is controlled by varying the standard deviations of the length and width distributions. These standard deviations are expressed as percentages of the mean values, and efficiencies are evaluated at 0%, 10%, and 20%. As noted in Section 5.3.2, vehicles are assumed to have an average length of 5 m and an average width of 2 m. The distributions $[l \sim \mathcal{N}(5, 0), w \sim \mathcal{N}(2, 0)]$, $[l \sim \mathcal{N}(5, 0.5), w \sim \mathcal{N}(2, 0.2)]$, and $[l \sim \mathcal{N}(5, 1), w \sim \mathcal{N}(2, 0.4)]$ correspond to the cases shown in Figure 5.13. The respective average travel times are 17.57 s, 17.51 s, and 17.39 s with medians of 17.48 s, 17.45 s, and 17.36 s. The results indicate that efficiency improves as vehicle size heterogeneity increases.




FIGURE 5.13. Distributions of the travel times to traverse the 500 m section. The AV fleets are controlled at $K = 1$. Vehicle dimensions are generated with coefficients of variation of 0, 10, and 20 percent.

This outcome is expected, since vehicles of different sizes can arrange themselves more effectively than vehicles of uniform dimensions. The effect is likely to be more pronounced in mixed traffic composed of distinct vehicle types, such as autonomous motorcycles, passenger cars, and large trucks, which could collectively optimise space in a lane-free setting.

5.3.5 Comparing Lane-free with Lane-based

Finally, we compare the current lane-free model with a similarly defined lane-based model. Again, the experiments are conducted with $K = 1$. To ensure a fair comparison between the two controllers, we only added a cost component for lane-keeping (lane-centring cost in Chapter 4) while keeping rest of the lane-free OCP unchanged. The lane-based controller is similarly optimised by tuning its parameters using the Bayesian approach as discussed in Section 5.2.4. The travel time distributions for the lane-free and lane-based simulations are displayed in Figure 5.14. The average travel times for lane-free and lane-based fleets are 17.51 s and 17.37 s, while the medians are 17.45 s and 17.28 s.

Interestingly, lane-free traffic being a generalisation of the lane-based traffic, performed worse than its counterpart. This could be due to the AVs needing to place higher emphasis




FIGURE 5.14. Distributions of the travel times to traverse the 500 m section. Both lane-free and lane-based fleets are controlled at $K = 1$.

on avoiding collisions on lane-free roads as vehicles must check a full 360-degree range of potential conflicts. Furthermore, vehicles perform more lateral movements in lane-free environments which could induce oscillations similar to how frequent lane-changes can facilitate stop-and-go waves which can lead to congestion. However, it is important to note that the lane-free concept is best practised when vehicle sizes are heterogeneous (as shown in Section 5.3.4). In a fleet with vehicles having a wider range of sizes, traffic with lanes could be less efficient as all vehicles occupy one lane, despite their sizes. In contrast, in a lane-free environment, vehicles can be packed more efficiently and hence improve the throughput. To test this assumption, we compare lane-free and lane-based AV fleets with a higher size heterogeneity of 30%. As shown in Figure 5.15, the lane-free traffic is slightly more efficient than the lane-based one when both have a larger heterogeneity in vehicle size. In the new simulations with a Coefficient of Variation of 0.3 in vehicle size, the average travel times for lane-free and lane-based fleets are 17.14 s and 17.33 s, while the medians are 17.13 s and 17.25 s.

Figure 5.16 visualises the relationship between vehicle size heterogeneity and traffic efficiency for both lane-free and lane-based fleets. It can be seen that, when the fleet is more heterogeneous, AVs in the lane-free environment achieve larger marginal time savings than those




FIGURE 5.15. Distributions of the travel times to traverse the 500 m section. Both lane-free and lane-based fleets are controlled at $K = 1$. Vehicle size heterogeneity is set with a coefficient of variation of 0.3.




FIGURE 5.16. Relationship between vehicle size heterogeneity and traffic efficiency for lane-free and lane-based fleets.

in the lane-based environment. This is because vehicles occupy road space more efficiently when their sizes vary significantly. Our analysis indicates a trade-off point at a coefficient of variation of around 0.22, beyond which a lane-free configuration becomes more efficient. This finding should be considered when designing mixed traffic systems, where lane-free and lane-based roads coexist to improve overall traffic efficiency.

5.4 Conclusion

In this chapter, we propose a decentralised AV controller that generates optimal trajectories in a lane-free freeway environment. To account for the actions of surrounding opponent vehicles, we introduced level-k game theory to the control framework. By assuming a hierarchical reasoning structure, the controller accounts for the strategic interactions between vehicles in a decentralised setting. To ensure real-time implementation of the controller, the continuation/GMRES method has been used to solve the optimal control problem.

The performance of 2DTC is tested through simulations. Microscopic analysis shows that vehicles can be controlled in a smooth manner and vehicles acknowledge the potential actions of opponents. Macroscopic analysis of the controller demonstrates an improved traffic capacity compared to human-driving traffic. The controller is also shown to operate in real-time and obeys the specified constraints. Furthermore, simulations show that AVs controlled at $K = 2$ outperform ones controlled at $K = 1$ in terms of efficiency. Vehicles that predict the actions of the others can further optimise their own trajectories. This effect is reduced when a mix of $K = 1$ and $K = 2$ AVs are generated. In addition, heterogeneity in vehicle sizes also affects the efficiency. Higher variations in sizes result in vehicles packed more efficiently inducing higher throughput. Finally, the lane-free simulations under the current configuration result in lower efficiency compared to a similarly defined lane-based counterpart. This could be a result of traffic oscillations from excessive lateral manoeuvres or AVs taking slightly longer, non-direct paths.

Several potential improvements and future studies should be acknowledged. First, the present study assumes a pipeline setting with no on- or off-ramps. As Chapter 4, future studies could investigate introducing costs for tactical routing. For instance, the cost for not exiting the freeway could be embedded to help AVs make decisions by integrating microscopic manoeuvres with network-level strategies (Wang et al., 2024). Although we assume a straight pipeline setting of the freeway, the controller could be extended to curved roads by reformulating the optimal control problem using the curvilinear coordinate system (Chu et al.,

2012). An alternative could be to use a hierarchical structure where trajectory planning and tracking are performed separately (Huang et al., 2020).

Another limitation of our work is that we assume an ego vehicle considers its surrounding vehicles to have the same set of cost weights as itself. This may not be a realistic assumption as vehicles have heterogeneous weights. However, since we assume AVs do not communicate, we have simply assumed vehicles consider others the same as themselves. Future research could develop learning methods to update the weights of surrounding vehicles in real-time, similar to how human drivers learn the aggressiveness of surrounding drivers. One potential learning approach could be to utilise Bayesian inference. The AV could start with the prior assumption of opponents having the same weights as itself. Then the weights could be updated as the ego vehicle gathers more observations by playing games with the opponents.

Finally, the proposed 2DTC should be tested for its performance on reduced-scale robots (Xie et al., 2024). Robots, which are essentially scaled-down AVs, could provide valuable information on the real-life applicability of the controller.

Conclusion

6.1 Summary of Main Findings

The microscopic decisions when driving often have far-reaching impacts on traffic as a whole. It is therefore important to generate accurate models that capture the realistic behaviours of human driving and to consider various automation strategies with the emergence of autonomous vehicles (AVs). Considering the trade-offs of operating mandatory lane-changes (MLCs), we analyse the costs of not performing MLCs for exiting freeways. We then model the trajectories of human-driven vehicles by integrating the longitudinal car-following with lateral lane-keeping and lane-changing decisions. Finally, an unconnected, decentralised control strategy is proposed to perform trajectory planning of AVs on a lane-free freeway environment. Through comprehensive analysis, modelling, and the development of novel optimisation approaches, this work contributes to the growing body of research on microscopic modelling and control.

In Chapter 2, we review both microscopic models for human drivers as well as autonomous vehicle controllers. We recognise that although there has been significant effort in developing car-following and lane-changing models, current studies mainly decouple the two dimensions of driving for simplicity. In reality, driving is a two-dimensional continuous process where acceleration and steering decisions are made jointly. Furthermore, the trade-off between the driving objectives is under-explored. For the simpler discretionary lane-changes (DLCs), the weights on different costs can be calibrated in a heterogeneous manner, while for more complex MLCs, the decisions need to be linked to the underlying incentive for performance

such manoeuvres. There is also significant interest in developing autonomous vehicle controllers for both lane-based and lane-free traffic settings. Despite the realism in unconnected, decentralised control systems in the near future, it has been relatively under-explored.

In light of the gaps identified, Chapter 3 explores the incentives for making MLCs for exiting freeways by modelling and analysing the costs for not exiting a freeway. MLCs are often regarded as absolute while we question this widely accepted proposition by evaluating the costs of neglecting an MLC. The core of this study involves a dual approach comprising an analytical model for the costs of missing exits alongside an empirical analysis of two GPS datasets from the Minneapolis - St. Paul metropolitan area. The performance of the analytical model is validated by cross-referencing it against exit-missing costs from the top 50 metropolitan areas in the US. The analytical model developed in this study can be integrated into the trade-off function for deciding a lane-change, enabling drivers to bypass their intended exit when needed. Regarding the empirical study, it was found that while both time and distance costs are associated with missing exits, the magnitudes of these costs are not substantial.

Chapter 4 introduces a microscopic two-dimensional driving model (2DDM) that integrates both longitudinal and lateral decision-making processes into one unified objective, allowing simultaneous consideration of car-following and lane-changing manoeuvres. The proposed model employs receding horizon optimisation to iteratively predict future movements, while the hierarchical reasoning structure inherent in human cognition is captured by leveraging level-k game theory. Drivers are assumed to be heterogeneous and the model calibrates their preferences using the highD dataset. Calibration is conducted using Bayesian inference, which approximates the underlying distributions of the preference weights. The calibrated model effectively reproduces macroscopic traffic patterns with minimal disparities. We demonstrate the model's transferability by validating it with data collected from different locations and times.

Building upon Chapter 4's 2DDM model, Chapter 5 proposes a two-dimensional trajectory controller (2DTC) for AVs operating on a lane-free freeway. Without inter-vehicle communication, each AV independently computes its acceleration and steering in real time. The 2DTC

problem is formulated as a receding horizon OCP with nonlinear dynamics, cost function, and constraints. Optimal controls are obtained by discretizing and solving the OCP using the continuation/GMRES method. Strategic interactions between vehicles are modelled using a level-k game-theoretic framework in a lane-free environment. The performance of the controller is tested through both microscopic and macroscopic analyses. Simulated microscopic scenarios demonstrate that the generated trajectories are smooth and constraint-compliant. Additional bottleneck simulations are used to construct fundamental diagrams, revealing higher flow capacity than typical human-driven traffic. Furthermore, tests are performed across different reasoning depths and levels of vehicle size heterogeneity, and the controller is also compared against an optimised lane-based controller.

6.2 Future Outlook

This thesis has made several novel contributions to the field of microscopic driving modelling and control. It has also opened up numerous avenues for future research by establishing a foundation on which models can be further optimised and controllers applied not only to AV trajectory planning but also to the broader study of lane-free and lane-based traffic systems. These contributions provide a basis for designing more efficient and adaptable transportation systems.

The study on exit-missing costs has shown that the costs for failing to make freeway exit lane-changes are not substantial. This finding suggests that existing microscopic models could relax the assumption that MLCs must always be executed. For instance, when no safe gaps are available, vehicles should be allowed to miss their intended exits. Exit-missing costs could also be incorporated into the trade-offs underlying lane-change decisions. In particular, the proposed analytical cost model can be readily integrated into both Chapter 4's 2DDM and Chapter 5's 2DTC, thereby linking higher-level routing choices with lower-level driving manoeuvres.

The 2DDM provides a strong foundation for generating human-like trajectories, yet several aspects warrant further development to better capture human behaviour. For example, the

current model assumes that drivers make optimal decisions based on their expectations of others. This assumption should be examined more closely by quantifying the levels of bias and randomness inherent in human driving. In reality, humans have limited capacity for fully optimal decision-making. One way to reflect this would be to terminate the solver after a set number of iterations, thereby producing sub-optimal solutions. Moreover, the 2DDM has so far been calibrated and validated on a single dataset. Future studies could expand this evaluation by testing the model on datasets from different regions worldwide and by examining its ability to reproduce microscopic trajectories of individual drivers.

Building on the lane-free controller developed in this study, future work could examine the conditions under which lane-free traffic outperforms lane-based traffic, and vice versa. Since lane-based traffic can be viewed as a special case of lane-free, heuristics could be developed to further enhance overall traffic performance. Such insights would be valuable for the design of mixed traffic systems, where lane-free and lane-based roads coexist to improve both efficiency and safety. Another promising direction is to explore the impact of cognitive hierarchy on traffic performance. For instance, future studies could investigate how different levels of K relate to overall efficiency. Finally, the proposed control strategy should be evaluated experimentally, such as through tests with reduced-scale robotic vehicles.

It is worth noting that both 2DDM and 2DTC can be extended to more complex traffic scenarios, such as intersections and freeway on- and off-ramps. Rather than relying solely on cost-function optimisation to generate trajectories, a bi-level approach could be adopted, where the upper level handles path planning and the lower level executes trajectory control. Furthermore, both 2DDM and 2DTC could benefit from the integration of real-time learning methods within the level- k game framework. In this setting, the assumed weights of surrounding vehicles could be updated dynamically, reflecting how human drivers adapt to others' preferences, such as aggressiveness. For AV control, online learning would enable vehicles to make more informed decisions by continuously acquiring knowledge about the behaviours of other AVs.

APPENDIX A

Appendix for Chapter 3

A1 Exits Analysed




FIGURE A.1. Spatial location of all exits in the datasets

A2 Route Deviation and Circuity

In this appendix section, we compute the circuity for both original and alternative routes to ascertain whether there is a change in circuitousness by missing an exit. To quantify the discrepancy between the original and the alternative routes, we employ the route deviation, which measures the closeness between two trajectories with the same OD. The measure is computed from the square root of the area of the concave hull that is enclosed by the routes (Wang et al., 2022). A higher deviation suggests that the alternative path is more dissimilar to the original path and, by our assumption, may imply higher routing costs. Figure A.2 shows an arbitrary trip and its best alternative. In this case, the alternative readjusts itself back to the original route after exiting at the next downstream ramp. The deviation from missing the exit can be derived from the square root of the shaded area.




FIGURE A.2. Original and alternative routes of trip 497 from sample GPS1019. The deviation of this original-alternative pair is calculated as the square root of the shaded area.

Figure A.3 shows the distributions of deviation for final exits. As shown in the figure, 90% of the exiting samples have deviations less than 3 km, reinforcing that alternative routes may not be very dissimilar to their original pair. Furthermore, it is expected and observed that original routes with the SP constraint deviate less from their alternatives since both paths have optimised travel times.

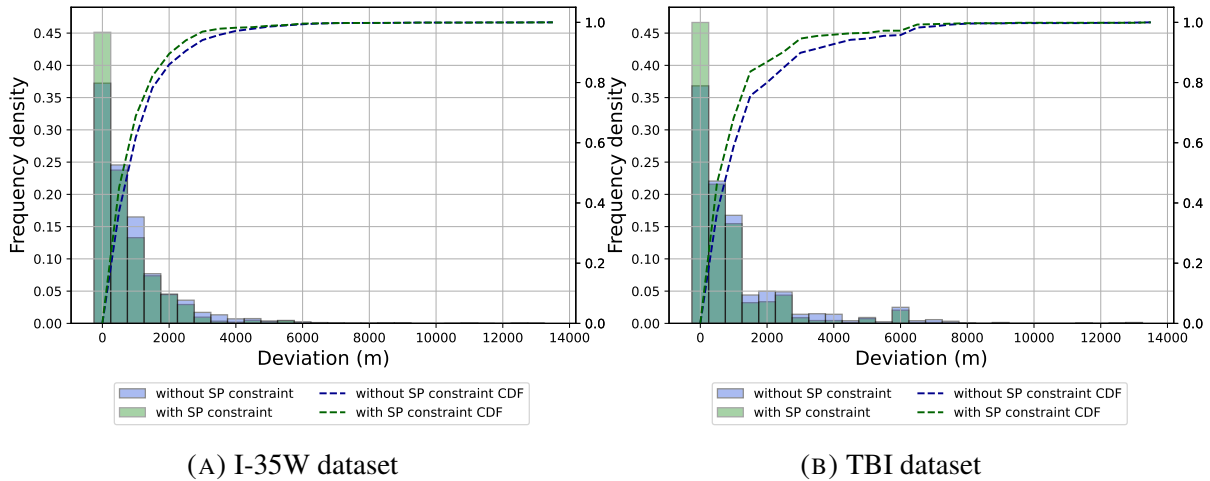


FIGURE A.3. Route deviation distributions and CDFs between original and alternative routes from final exits. Blue plots correspond to samples without the SP constraint in place. Green plots correspond to samples with the 10% SP constraint. Subplots (a) and (b) correspond to the I-35W dataset and the TBI dataset, respectively.

To understand the impact of missing an exit on the directness of the trips, we utilise the network circuitry, which is computed as the ratio between the route length and the Euclidean distance between the origin and the destination. Circuitry is computed for original and alternative routes from the source node of the exit to the final destination of the trip. The outcomes are illustrated in Figure A.4. As can be seen, the densities of the original routes are more concentrated at lower values, while the alternative routes are more platykurtic. Although computed with shortest path assumptions, the alternative routes have a spread-out distribution due to the great variability of the Euclidean distance between the assessed ODs. In general, we assume a shorter distance between the exit source node and the destination will induce a larger circuitry as there is a higher likeliness of taking detours, but this may also depend on the road network topology.




FIGURE A.4. Circuitry distributions and CDFs for original and alternative routes from final exits. The purple plots correspond to the original routes, and the green plots correspond to the alternative routes. Subplots (a) and (c) correspond to the I-35W dataset and subplots (b) and (d) correspond to the TBI dataset.

Moreover, we compare the deviation factors to the change in circuitry and exit missing costs, evaluating our hypothesis that a larger deviation can induce larger time and distance discrepancies. The correlations are assumed to be linear, and the Pearson correlation coefficients are calculated between the variables. As shown in Table A.1, only the costs are strongly correlated, while the deviation has weak correlations with other variables. This suggests the deviation is not suited for estimating the costs of missing an exit. Although counter-intuitive, considering the opposing effect of the negative time/distance costs may explain this outcome. In addition, the deviation is not only related to the closeness of the original and alternative trajectories but also their length. This adds a stochastic nature to the parameter. Therefore, although

the deviation is a fair measure of the spatial variation of two routes, it can be misleading when used as a parallel comparison standard. On the other hand, the circuitry difference is seen to have a moderate positive correlation with the costs, indicating that more circuitous routes correspond to higher travel costs. Unsurprisingly, the travel costs themselves have a strong positive correlation as the average velocities on alternative routes are guaranteed to be positive.

TABLE A.1. The correlation between measurements of missing exits. The factors are: the extra travel time cost, the extra travel distance cost, the circuitry difference, and the routes deviation.

Factors	1	2	3	4
1. Travel time cost	-	0.736***	0.433**	-0.368*
2. Travel distance cost		-	0.523**	-0.207*
3. Circuitry difference			-	-0.187*
4. Route deviation				-

***: Strong correlation ($0.7 < |r| < 1$)

** : Moderate correlation ($0.4 < |r| < 0.7$)

* : Weak correlation ($0.1 < |r| < 0.4$)

A3 Categorical Analysis of Final Exits

In this section, we demonstrate our method of estimating the number of final exits missed by the participants. Here we only assess the final freeway exits since the intention of remaining on the freeways for the interim exits is ambiguous as there is higher flexibility for a driver to choose an interim exit than it is for a final exit. We also neglect non-commuting trips since inconsistent OD pairs can generate distinctive routes with different choice sets of exits. For each driver, we split all commuting trips into to-home and to-work. The I-35W dataset (VMT and Otrek combined) is further stratified into before-bridge-reopen and after-bridge-reopen, ensuring bridge usage consistency.

The final exits under each subcategory are extracted and sorted according to their proportion of appearance, where the most frequent exit is assigned to be the Common Exit (CE). We then retain only the downstream exits of the CE by applying the breadth-first search algorithm on the freeway and ramp network for each distinctive commuting OD pair.

The remaining exits are ranked based on the order of shortest path using a variant of the path deletion K-shortest paths algorithm (Azevedo et al., 1993). The algorithm incrementally eliminates the final exit in the fastest route, which forces the algorithm to search for a new shortest path with a different final exit. The loop stops once the analysed exit is matched with the final exit of the shortest path. The order of the Shortest Exit (SE) is then recorded.

In addition, we identify the Next Exit (NE) downstream of the CE for each OD pair by running a single source Dijkstra search starting at the CE source node. NE can be obtained as the exit edge with the least cost assuming equal link impedances on all edges, excluding the CE. Trips in the same OD pair that contain the NE are then identified and classified as trips with missing exits if and only if their proportion of appearance in their respective OD pair is less than equal to certain threshold. Otherwise, it is assumed to be intentional. In Chapter 4, we set the thresholds to be 25%, 15%, and 5%.

Above we have developed a simple heuristic to categorically analyse the exits based on their frequency of usage, travel time to destination, and likeliness of being missed. The three labels are illustrated in Figure A.5 through an arbitrary commuting OD pair. The CE, in this case, has the highest travel time out of the three options, which in turn emphasises that people are often not rational in their route choice. The implementation of the procedure is showcased in Algorithm 4.

Following the aforementioned procedure, 2,888 final exits from commuting trips are used to analyse the exiting behaviour of drivers. After assigning the labels to all final commuting exits, the counts are aggregated, and the output is presented in the format of a decision tree (Figure A.6). The branches of the trees represent the decisions regarding ordering, while the leaves represent the percentages of the corresponding orders.

Looking top-down from the decision tree, the large majority (roughly 98%) of the final exits are comprised of CEs, while at most, three exits were used by participants during the sampling period. This suggests people tend to minimise variation by sticking to a fixed commuting route.

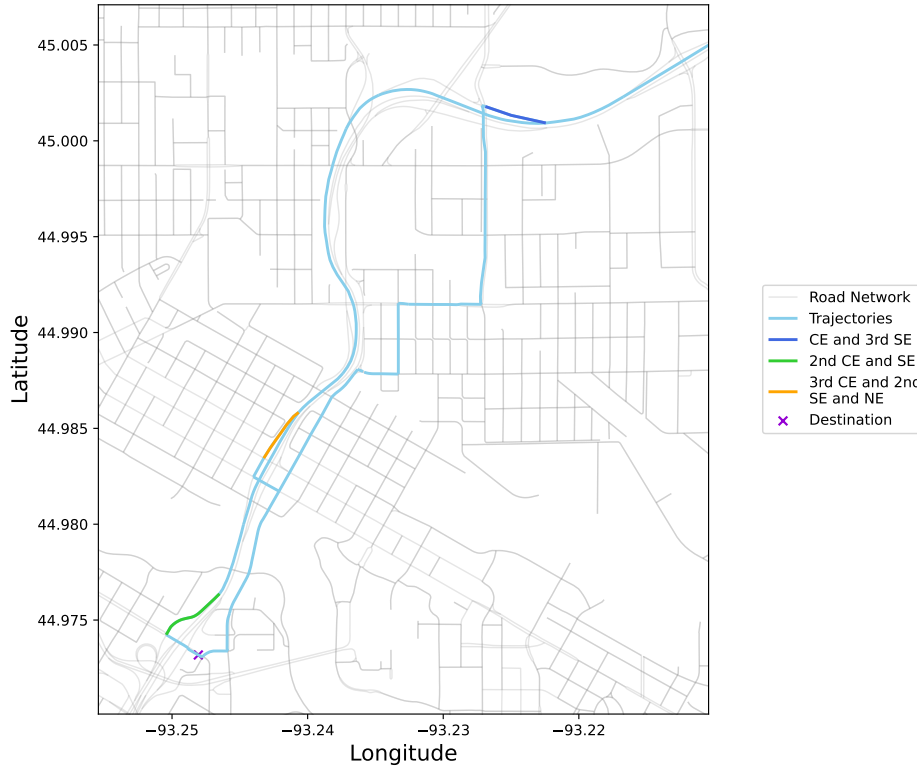


FIGURE A.5. Final exits of to-work trips for sample GPS2038. The subject used three different exits during the sampling period. Common Exit (CE) denotes the ordering of the final exits regarding the frequency of usage; Shortest Exit (SE) corresponds to the ordering of the exits based on their usage in the n^{th} shortest paths (from the source node of CE to the destination marked by the purple cross); Next Exit (NE) is defined as the next exit directly downstream of the CE, which we assume a trip that missed the CE has traversed along.

Furthermore, we see that the CEs are often not the SEs (13.19% for the I-35W dataset and 18.90% for the TBI dataset) even from the point just upstream of the CE. Previous studies (Tang and Levinson, 2018, Zhu and Levinson, 2015) on the same datasets argue the trips used by the participants deviate significantly from the shortest paths, which may provide reasoning as to why a large proportion of CEs do not coincide with SEs. Although consisting of only a tiny proportion, CEs with SE orders as large as 7/8 still suggest that some trips are largely inefficient in terms of travel time. This phenomenon did not seem to improve from 2007/08 (when the I-35W dataset was collected) to 2010/11 (when the TBI dataset was collected). The total percentage of final exits on the shortest path dropped from 85.50% to 79.88%. Despite more people being exposed to advanced navigation systems over the 2-3 years (especially

Algorithm 4 Exit Class Identification**Input:**

OD_{list} : A list of OD vectors with each vector containing different numbers of trips. Each trip (element of a vector) contains different links

Output:

$final_exits[OD]$: For each OD, a table that contains the exits used in that OD pair along with the counts, CE, NE, and SE orders

Definitions:

(U_i, V_i) : Source and target nodes of a link i

G : Road network graph

$G_{freeway}$: A subgraph of the road network graph containing freeway network only

$BFS(graph, start, depth)$: Breadth-first search function

$dijkstra(graph, start, end)$: Dijkstra's shortest path function

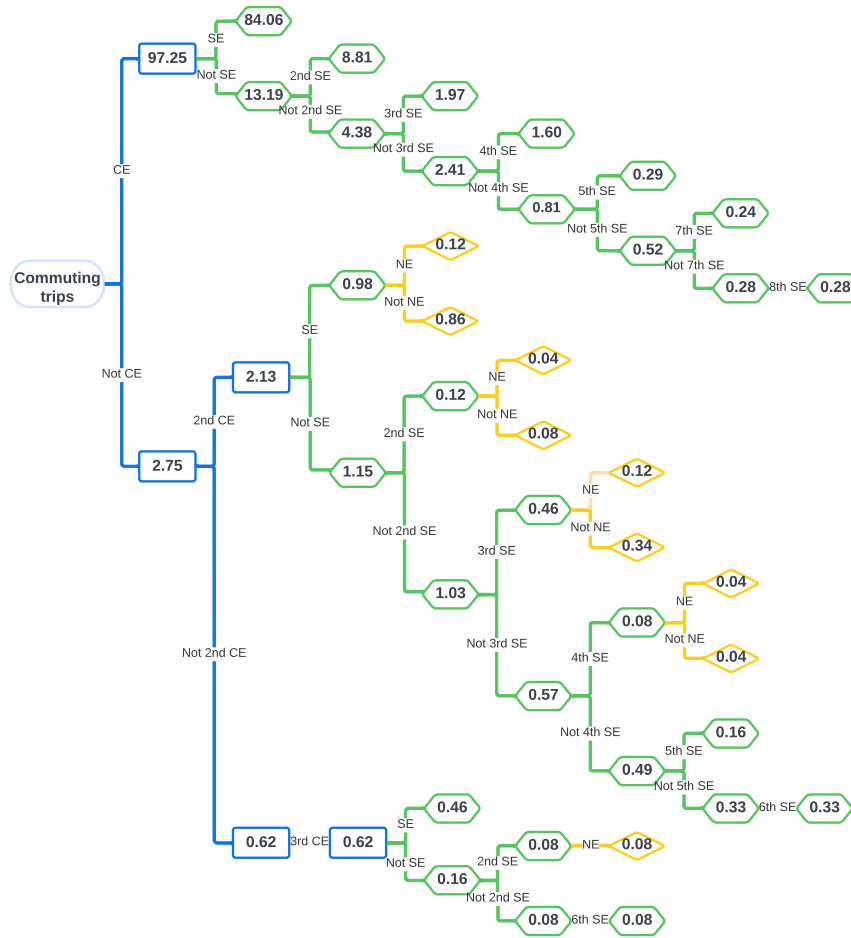
```

1: for each  $OD$  in  $OD_{list}$  do
2:   for each  $trip$  in  $OD$  do
3:     Extract the final exit  $(U_{exit}, V_{exit})$  from trip
4:      $final\_exits[OD] \leftarrow (U_{exit}, V_{exit})$ 
5:   end for
6:    $exit\_counts \leftarrow$  count occurrences of each unique  $(U_{exit}, V_{exit})$  in  $final\_exits$ 
7:    $CE\_exit \leftarrow \text{argmax}(exit\_counts)$ 
8:    $down\_links \leftarrow BFS(graph = G_{freeway}, start = U_{CE\_exit}, depth = 20)$ 
9:    $down\_exits \leftarrow final\_exits \cap down\_links$ 
10:  for  $(U_{exit}, V_{exit})$  in  $down\_exits$  do
11:     $CE[(U_{exit}, V_{exit})] \leftarrow$  ranking of  $(U_{exit}, V_{exit})$  in  $exit\_counts$ 
12:     $NE[(U_{exit}, V_{exit}) | (U_{exit}, V_{exit}) = down\_exits[2]] \leftarrow 1$ 
13:    set  $k \leftarrow 1$ 
14:    while True do
15:       $SP \leftarrow dijkstra(graph = G, start = U_{CE\_exit}, end = V_{destination})$ 
16:      if  $(U_{exit}, V_{exit}) \in SP$  then
17:         $SE[(U_{exit}, V_{exit})] \leftarrow k$ 
18:        break
19:      else
20:         $G \leftarrow G - \{\text{last exit in } SP\}$ 
21:      end if
22:       $k \leftarrow k + 1$ 
23:    end while
24:  end for
25:   $final\_exits[OD] \leftarrow$  sort  $final\_exits[OD]$  based on unique  $(U_{exit}, V_{exit})$ 
26:   $final\_exits[OD] \leftarrow final\_exits[OD] \cup \{exit\_counts, CE, NE, SE\}$ 
27: end for
28: return  $final\_exits$ 


```

after the release of the Google Maps phone app on September 2008), the usage of the shortest path has not seen improvements. This lagging effect on technological advancement possibly indicates that people are reluctant to change their regular routines.

The NEs, on the other hand, are observed to have relatively low ratios, which may result from people intentionally skipping the CE due to other external factors. For instance, the driver may perceive that the CE is congested or wish to fuel their car by taking a different route. Regardless of the cause, the drivers are assumed to fully understand the consequences of skipping the CE and thus are not accounted for missing it. After deducting these trips, only 0.40% in the I-35W dataset and 1.22% in the TBI dataset are regarded as NEs. By applying the 25% threshold, 12 trips remain to be determined to contain missed exits, the 15% threshold yields 3 trips, while 5% yields 2 trips. As expected, the occurrence of missing exits only consists of a small portion of the total trips. It is worth noting that this analysis is independent of the rest of Chapter 3, which focuses on the cost of missing a freeway exit. The extracted missed exits are not used as samples for any other analysis.



(A) I-35W dataset



(B) TBI dataset

FIGURE A.6. Decision tree of exit structures for commuting trips. All numbers on leaves are in percentage. Blue corresponds to the order of CE; green corresponds to the order of SE; orange corresponds to the binary decision of NE.

Appendix for Chapter 4

B1 Fitting of Truncated Normal Distribution

To fit the truncated normal distribution to the parameters, we define an optimisation problem. The optimal parameters $(\bar{\mu}^*, \bar{\sigma}^*, z_a^*, z_b^*)$ are found by minimising the negative log-likelihood function of a truncated normal distribution subjected to the equality constraints that ensure the correct truncation positioning:

$$\begin{aligned} \min \quad & - \sum_{i=1}^N \ln (f(\bar{\mu}, \bar{\sigma}, z_a, z_b; x)) \\ \text{s.t.} \quad & x_a = z_a \bar{\sigma} + \bar{\mu} \end{aligned} \tag{B.1}$$

$$x_b = z_b \bar{\sigma} + \bar{\mu}$$

$$f(\bar{\mu}, \bar{\sigma}, z_a, z_b; x) = \begin{cases} \frac{\phi(\bar{\mu}, \bar{\sigma}; x)}{\Phi(\bar{\mu}, \bar{\sigma}; z_b) - \Phi(\bar{\mu}, \bar{\sigma}; z_a)} & x_a < x < x_b \\ 0 & \text{otherwise} \end{cases} \tag{B.2}$$

where ϕ and Φ are the PDF and CDF of the general normal distribution. Note, z_a and z_b are the z-scores of the respective abscissas, x_a and x_b , at which we wish to truncate the distribution. The optimisation is performed using Sequential Least Squares Programming ([Kraft, 1988](#)).

B2 Correlation of Calibrated Parameters

The Pearson correlation coefficients between all pairs of parameters are estimated based on the calibrated particles. The correlation matrices are depicted in Figure B.1. As can be seen, all parameters are weakly correlated.

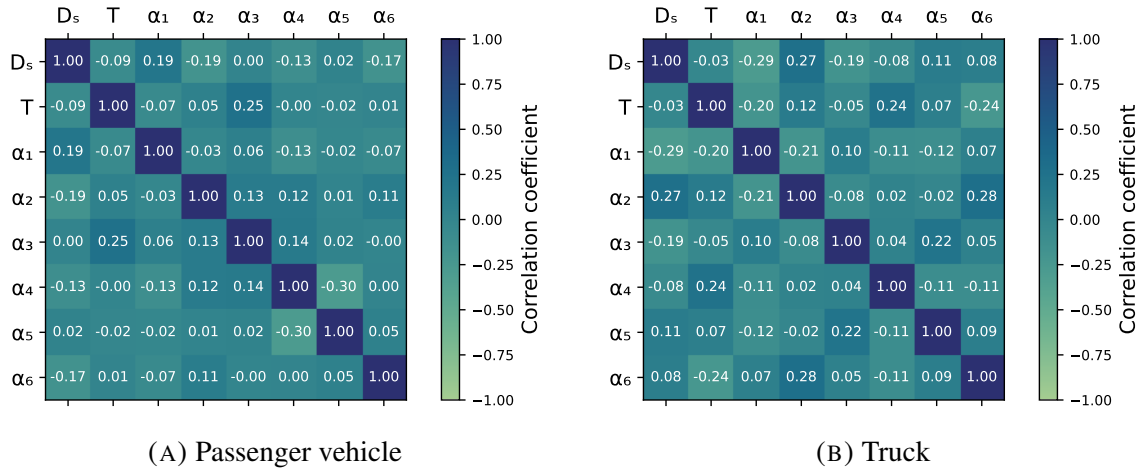


FIGURE B.1. Correlation matrices for calibrated parameters. The corresponding Pearson correlation coefficients are displayed for all pairs of parameters.

B3 Calibrated Distributions of IDM+MOBIL Model

Our proposed model is validated in comparison to another widely used car-following lane-changing model, namely the IDM+MOBIL model. To ensure the models are consistent for comparison, we assume the fixed parameters the same and calibrate the remaining parameters. In Table 4.1, the minimum standstill gap is 2 m. The minimum and maximum accelerations are -3.4 m/s^2 and 3 m/s^2 for passenger cars, and -1.5 m/s^2 and 1 m/s^2 for trucks. The calibration parameters are the desired velocity, safe time headway, politeness, and lane-change threshold. Once again, we use Algorithm 2 to perform the parameter calibration. The calibrated parameter distributions are summarised in Table B.1 and Figure B.2.

TABLE B.1. Statistics of the truncated normal fits to the calibrated parameters

	Passenger car			Truck		
	Mean	Std	Truncation	Mean	Std	Truncation
Desired velocity	35.72 m/s	2.90 m/s	$[0, \infty)$	25.74 m/s	0.97 m/s	$[0, \infty)$
Safe time headway	1.11 s	0.10 s	$[0, \infty)$	1.38 s	0.13 s	$[0, \infty)$
Politeness	0.65	0.11	$[0, \infty)$	0.67	0.13	$[0, \infty)$
Lane-change threshold	0.63 m/s ²	0.13 m/s ²	$[0, 3]$	0.66 m/s ²	0.11 m/s ²	$[0, 1]$

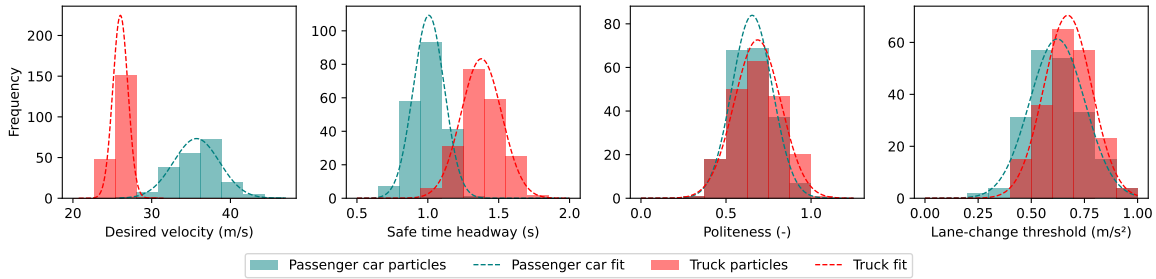


FIGURE B.2. Calibrated distributions of parameters for IDM+MOBIL model. The histograms are produced using the calibrated particles from the MB-SMC algorithm, while the dashed lines represent the fitted truncated normal distributions.

B4 Sensitivity of Parameters

To understand the influence of the calibration parameters on model performance, sensitivity analysis has been conducted. Each parameter is individually increased and decreased by 10% and 20%, while keeping all other parameters fixed. This results in a total of 32 parameter combinations to simulate. Each parameter combination is also repeated 3 times to ensure consistency of the outputs. For simplicity and consistency, the same percentage changes are applied to both passenger cars and trucks. Lane-change frequency is selected as the evaluation metric, as it differs most significantly from observed data while also exhibits the largest standard deviations. The resulting LC frequencies for different parameter combinations are summarised in Table B.2.

TABLE B.2. Sensitivity analysis of calibrated parameters using percentage error of lane-change frequency. The analysis is conducted using the $K = 1$ model and Track 50 from the highD dataset.

	20%	10%	-10%	-20%
Sight distance (D_s)	134.56	129.81	131.59	124.68
Safe headway (T_{safe})	145.66	133.06	119.40	117.01
Speed cost weight (α_{speed})	119.62	129.61	133.88	136.01
Spacing cost weight (α_{spacing})	127.26	124.91	130.71	129.95
Safety cost weight (α_{safe})	128.33	123.42	133.67	132.60
Acceleration cost weight ($\alpha_{\text{acceleration}}$)	117.05	115.55	125.16	135.41
Steering cost weight (α_{steering})	132.81	124.06	127.30	140.71
Centring cost weight (α_{centring})	118.11	121.05	134.13	137.54

The observed LC frequency for Track 50 is 125.79 LC/hr/lane. Using the optimised parameters, the simulated LC frequency for $K = 1$ is 129.58 LC/hr/lane. As shown in Table B.2, some parameter combinations performed better than the optimised parameters. For instance, reducing the sight distance by 20% resulted in an LC frequency of 124.68 LC/hr/lane, which is closer to the observed value. However, it's important to note that lane-change frequency is just one metric. The optimal parameter set was selected based on its overall performance across multiple metrics. For example, although the 80% sight distance setting improves LC prediction, it yields an average speed of 28.99 m/s, which is less accurate compared to the 31.37 m/s predicted using the optimal sight distance.

Furthermore, the relatively small variation in LC frequency across the different parameter settings indicates that the model's behaviour is not highly sensitive to moderate changes in individual parameters. This suggests that the model is robust and likely to generalise well to new traffic conditions and different driver populations.

APPENDIX C

Appendix for Chapter 5

C1 Visualisation and Animations of Simulated Scenarios

C1.1 Scenario 1

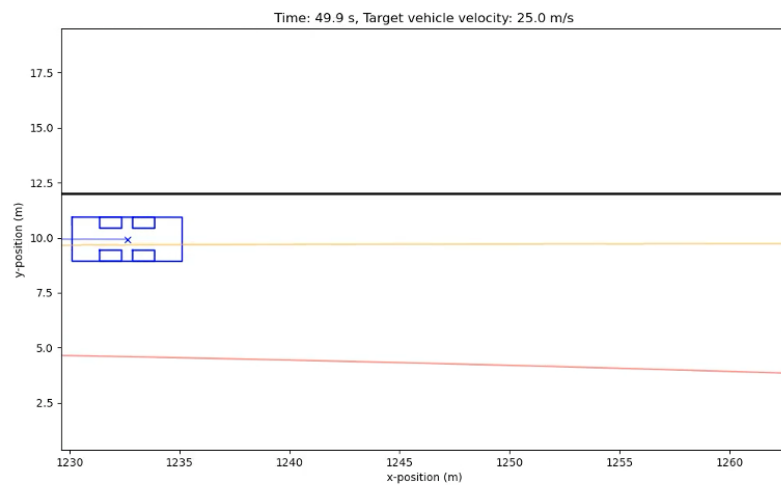


FIGURE C.1. Animation of scenario 1. The camera is focused on Vehicle 1. This video is available in the electronic version of the thesis.

C1.2 Scenario 2

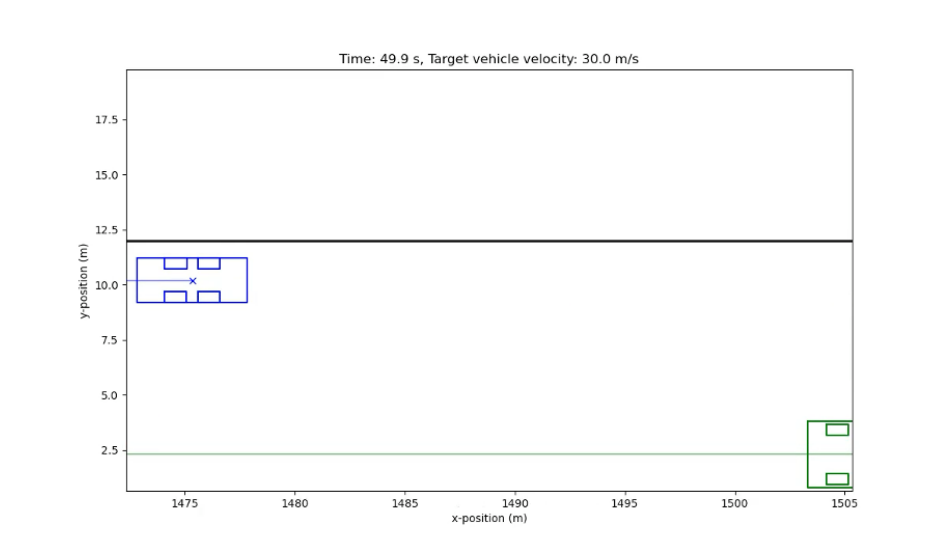


FIGURE C.2. Animation of scenario 2. The camera is focused on Vehicle 1. This video is available in the electronic version of the thesis.

C1.3 Scenario 3




FIGURE C.3. Animation of scenario 3. The camera is focused on Vehicle 1. This video is available in the electronic version of the thesis.

C1.4 Scenario 4

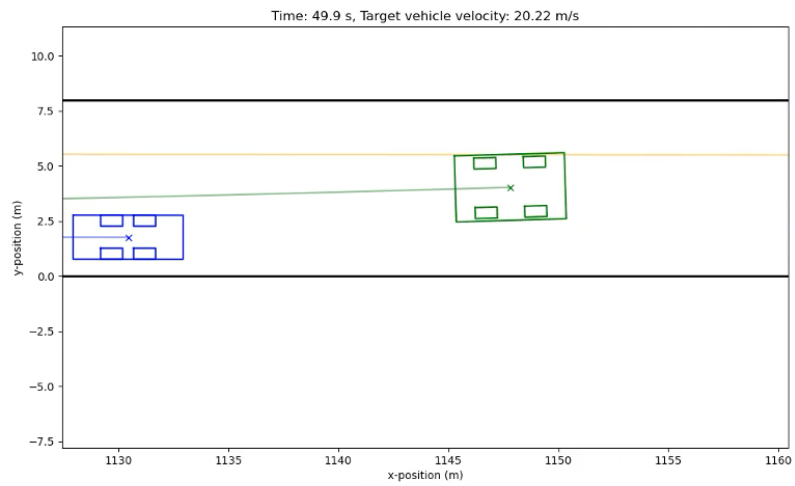


FIGURE C.4. Animation of scenario 4. The camera is focused on Vehicle 1. This video is available in the electronic version of the thesis.

Bibliography

- AASHTO (2018). *A Policy on Geometric Design of Highways and Streets* (7th Edition ed.). Washington, DC, USA: American Association of State Highway and Transportation Officials.
- Ahmed, K. I. (1999). *Modeling drivers' acceleration and lane changing behavior*. Thesis, Massachusetts Institute of Technology. Accepted: 2005-08-19T19:16:47Z.
- Ahn, S. and M. Cassidy (2007, January). Freeway traffic oscillations and vehicle lane-change maneuvers.
- Ali, Y., Z. Zheng, M. M. Haque, and M. Wang (2019). A Game Theory-based Approach for Modelling Mandatory Lane-changing Behaviour in a Connected Environment. *Transportation Research Part C: Emerging Technologies* 106, 220–242.
- Axhausen, K. W., S. Schönfelder, J. Wolf, M. Oliveira, and U. Samaga (2003, August). 80 weeks of GPS-traces: Approaches to enriching the trip information. pp. 29 p. Artwork Size: 29 p. Medium: application/pdf Publisher: ETH Zurich.
- Azevedo, J., M. E. O. Santos Costa, J. J. E. Silvestre Madeira, and E. Q. Vieira Martins (1993, August). An algorithm for the ranking of shortest paths. *European Journal of Operational Research* 69(1), 97–106.
- Baccelli, F., M. Klein, M. Lebourges, and S. Zuyev (1997, June). Stochastic geometry and architecture of communication networks. *Telecommunication Systems* 7(1), 209–227.
- Balal, E., R. L. Cheu, and T. Sarkodie-Gyan (2016, June). A Binary Decision Model for Discretionary Lane Changing Move Based on Fuzzy Inference System. *Transportation Research Part C: Emerging Technologies* 67, 47–61.
- Bando, M., K. Hasebe, A. Nakayama, A. Shibata, and Y. Sugiyama (1995, February). Dynamical model of traffic congestion and numerical simulation. *Physical Review E* 51(2), 1035–1042. Publisher: American Physical Society.

- Barceló, J. (Ed.) (2010). *Fundamentals of Traffic Simulation*, Volume 145 of *International Series in Operations Research & Management Science*. New York, NY: Springer.
- Berachman, M., M. Rostmai-Shahrbabaki, and K. Bogenberger (2022, January). Driving Strategy for Vehicles in Lane-Free Traffic Environment Based on Deep Deterministic Policy Gradient and Artificial Forces. *IFAC-PapersOnLine* 55(14), 14–21.
- Bichiou, Y. and H. A. Rakha (2019, May). Developing an Optimal Intersection Control System for Automated Connected Vehicles. *IEEE Transactions on Intelligent Transportation Systems* 20(5), 1908–1916. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Bokare, P. S. and A. K. Maurya (2017, January). Acceleration-Deceleration Behaviour of Various Vehicle Types. *Transportation Research Procedia* 25, 4733–4749.
- Brockfeld, E., R. D. Kühne, and P. Wagner (2004, January). Calibration and Validation of Microscopic Traffic Flow Models. *Transportation Research Record* 1876(1), 62–70. Publisher: SAGE Publications Inc.
- Bureau, U. C. (2022a). Longitudinal Employer-Household Dynamics.
- Bureau, U. C. (2022b). TIGER/Line Shapefiles. Section: Government.
- Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004, August). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics* 119(3), 861–898.
- Carrion, C. and D. Levinson (2012, May). Value of travel time reliability: A review of current evidence. *Transportation Research Part A: Policy and Practice* 46(4), 720–741.
- Chavoshi, K. and A. Kouvelas (2021, September). Distributed control for laneless and directionless movement of connected and automated vehicles. STRC. Accepted: 2022-03-14T13:54:59Z.
- Chen, C., H. Gong, C. Lawson, and E. Bialostozky (2010, December). Evaluating the feasibility of a passive travel survey collection in a complex urban environment: Lessons learned from the New York City case study. *Transportation Research Part A: Policy and Practice* 44(10), 830–840.
- Chen, K., X. Pei, H. Okuda, M. Zhu, X. Guo, K. Guo, and T. Suzuki (2020, November). A hierarchical hybrid system of integrated longitudinal and lateral control for intelligent vehicles. *ISA Transactions* 106, 200–212.

- Chu, K., M. Lee, and M. Sunwoo (2012, December). Local Path Planning for Off-Road Autonomous Driving With Avoidance of Static Obstacles. *IEEE Transactions on Intelligent Transportation Systems* 13(4), 1599–1616.
- Ciuffo, B., V. Punzo, and M. Montanino (2012, January). Thirty Years of Gipps' Car-Following Model: Applications, Developments, and New Features. *Transportation Research Record* 2315(1), 89–99.
- Dabestani, N., P. Typaldos, I. Papamichail, and M. Papageorgiou (2024, January). Joint Vehicle Path Planning for Interruptible 1-D Snake-like Platoons and 2D Flocks in Lane-free Traffic. *IFAC-PapersOnLine* 58(10), 61–68.
- Daley, D. J. and D. Vere-Jones (2008). *An Introduction to the Theory of Point Processes. Probability and Its Applications*. New York, NY: Springer.
- Del Moral, P., A. Doucet, and A. Jasra (2006, June). Sequential Monte Carlo Samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 68(3), 411–436.
- Dia, H. (2002, October). An agent-based approach to modelling driver route choice behaviour under the influence of real-time information. *Transportation Research Part C: Emerging Technologies* 10(5), 331–349.
- Edie, L. C. (1963). *Discussion of Traffic Stream Measurements and Definitions*. Port of New York Authority. Google-Books-ID: 3uBmYgEACAAJ.
- Eiter, T. and H. Mannila (1994, May). Computing Discrete Frechet Distance.
- Endres, S. C., C. Sandrock, and W. W. Focke (2018, October). A simplicial homology algorithm for Lipschitz optimisation. *Journal of Global Optimization* 72(2), 181–217.
- Fox, J. (2008). *Applied regression analysis and generalized linear models, 2nd ed.* Applied regression analysis and generalized linear models, 2nd ed. Thousand Oaks, CA, US: Sage Publications, Inc. Pages: xxi, 665.
- Gao, Y. and D. Levinson (2023, December). Lane changing and congestion are mutually reinforcing? *Communications in Transportation Research* 3, 100101.
- Geroliminis, N., J. Haddad, and M. Ramezani (2013, March). Optimal Perimeter Control for Two Urban Regions With Macroscopic Fundamental Diagrams: A Model Predictive Approach. *IEEE Transactions on Intelligent Transportation Systems* 14(1), 348–359. Conference Name: IEEE Transactions on Intelligent Transportation Systems.

- Gipps, P. G. (1981, April). A behavioural car-following model for computer simulation. *Transportation Research Part B: Methodological* 15(2), 105–111.
- Gipps, P. G. (1986, October). A model for the structure of lane-changing decisions. *Transportation Research Part B: Methodological* 20(5), 403–414.
- Glaser, S., B. Vanholme, S. Mammar, D. Gruyer, and L. Nouvelière (2010, September). Maneuver-Based Trajectory Planning for Highly Autonomous Vehicles on Real Road With Traffic and Driver Interaction. *IEEE Transactions on Intelligent Transportation Systems* 11(3), 589–606. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Gong, H., C. Chen, E. Bialostozky, and C. T. Lawson (2012, March). A GPS/GIS method for travel mode detection in New York City. *Computers, Environment and Urban Systems* 36(2), 131–139.
- Guo, H. Y., Y. Ji, T. Qu, and H. Chen (2013, January). Understanding and Modeling the Human Driver Behavior Based on MPC. *IFAC Proceedings Volumes* 46(21), 133–138.
- Halati, A., H. Lieu, and S. Walker (1997). CORSIM - CORRIDOR TRAFFIC SIMULATION MODEL.
- Hall, F. L., B. L. Allen, and M. A. Gunter (1986, May). Empirical analysis of freeway flow-density relationships. *Transportation Research Part A: General* 20(3), 197–210.
- Hart, F., O. Okhrin, and M. Treiber (2024, February). Towards robust car-following based on deep reinforcement learning. *Transportation Research Part C: Emerging Technologies* 159, 104486. Publisher: Elsevier BV.
- Hayward, J. C. (1971). Near misses as a measure of safety at urban intersections. Master's thesis, The Pennsylvania State University, Philadelphia, PA. Issue: 384.
- Helbing, D. and B. Tilch (1998, July). Generalized force model of traffic dynamics. *Physical Review E* 58(1), 133–138. Publisher: American Physical Society.
- Hou, Y., P. Edara, and C. Sun (2012, September). A genetic fuzzy system for modeling mandatory lane changing. In *2012 15th International IEEE Conference on Intelligent Transportation Systems*, pp. 1044–1048. ISSN: 2153-0017.

- Hou, Y., P. Edara, and C. Sun (2014, April). Modeling Mandatory Lane Changing Using Bayes Classifier and Decision Trees. *IEEE Transactions on Intelligent Transportation Systems* 15(2), 647–655. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Huang, Y., H. Wang, A. Khajepour, H. Ding, K. Yuan, and Y. Qin (2020, January). A Novel Local Motion Planning Framework for Autonomous Vehicles Based on Resistance Network and Model Predictive Control. *IEEE Transactions on Vehicular Technology* 69(1), 55–66.
- Hunt, J. G. and G. D. Lyons (1994, December). Modelling dual carriageway lane changing using neural networks. *Transportation Research Part C: Emerging Technologies* 2(4), 231–245.
- Ji, A. and D. Levinson (2020, January). A review of game theory models of lane changing. *Transportmetrica A: Transport Science* 16(3), 1628–1647. Publisher: Taylor & Francis
_eprint: <https://doi.org/10.1080/23249935.2020.1770368>.
- Ji, A. and D. Levinson (2021, October). Estimating the Social Gap With a Game Theory Model of Lane Changing. *IEEE Transactions on Intelligent Transportation Systems* 22(10), 6320–6329.
- Ji, A., M. Ramezani, and D. Levinson (2023a, December). Joint modelling of longitudinal and lateral dynamics in lane-changing maneuvers. *Transportmetrica B: Transport Dynamics* 11(1), 996–1025.
- Ji, A., M. Ramezani, and D. Levinson (2023b, April). Pricing lane changes. *Transportation Research Part C: Emerging Technologies* 149, 104062.
- Jiang, G., M. Fosgerau, and H. K. Lo (2020, February). Route choice, travel time variability, and rational inattention. *Transportation Research Part B: Methodological* 132, 188–207.
- Jiang, R., Q. Wu, and Z. Zhu (2001, June). Full velocity difference model for a car-following theory. *Physical Review E* 64(1), 017101. Publisher: American Physical Society.
- Johansson, G. and K. Rumar (1971, February). Drivers' Brake Reaction Times. *Human Factors* 13(1), 23–27. Publisher: SAGE Publications Inc.
- Karafyllis, I., D. Theodosis, and M. Papageorgiou (2022a, September). Constructing artificial traffic fluids by designing cruise controllers. *Systems & Control Letters* 167, 105317.

- Karafyllis, I., D. Theodosis, and M. Papageorgiou (2022b, November). Lyapunov-based two-dimensional cruise control of autonomous vehicles on lane-free roads. *Automatica* 145, 110517.
- Karalakov, A., D. Troullinos, G. Chalkiadakis, and M. Papageorgiou (2022). Deep RL Reward Function Design for Lane-Free Autonomous Driving. In F. Dignum, P. Mathieu, J. M. Corchado, and F. De La Prieta (Eds.), *Advances in Practical Applications of Agents, Multi-Agent Systems, and Complex Systems Simulation. The PAAMS Collection*, Cham, pp. 254–266. Springer International Publishing.
- Katrakazas, C., M. Quddus, W.-H. Chen, and L. Deka (2015, November). Real-time motion planning methods for autonomous on-road driving: State-of-the-art and future research directions. *Transportation Research Part C: Emerging Technologies* 60, 416–442.
- Keane, R. and H. O. Gao (2021, May). Fast Calibration of Car-Following Models to Trajectory Data Using the Adjoint Method. *Transportation Science* 55(3), 592–615.
- Kesting, A. and M. Treiber (2008, January). Calibrating Car-Following Models by Using Trajectory Data: Methodological Study. *Transportation Research Record* 2088(1), 148–156.
- Kesting, A., M. Treiber, and D. Helbing (2007, January). General Lane-Changing Model MOBIL for Car-Following Models. *Transportation Research Record: Journal of the Transportation Research Board* 1999(1), 86–94.
- Kita, H. (1999, April). A merging–giveway interaction model of cars in a merging section: a game theoretic analysis. *Transportation Research Part A: Policy and Practice* 33(3), 305–312.
- Kraft, D. (1988). A software package for sequential quadratic programming. *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt*.
- Krajewski, R., J. Bock, L. Kloeker, and L. Eckstein (2018, November). The highD Dataset: A Drone Dataset of Naturalistic Vehicle Trajectories on German Highways for Validation of Highly Automated Driving Systems. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2118–2125. ISSN: 2153-0017.

- Lazcano, A. M. R., T. Niu, X. C. Akutain, D. Cole, and B. Shyrokau (2021, June). MPC-Based Haptic Shared Steering System: A Driver Modeling Approach for Symbiotic Driving. *IEEE/ASME Transactions on Mechatronics* 26(3), 1201–1211. Conference Name: IEEE/ASME Transactions on Mechatronics.
- Leshed, G., T. Velden, O. Rieger, B. Kot, and P. Sengers (2008, April). In-car gps navigation: engagement with and disengagement from the environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08*, New York, NY, USA, pp. 1675–1684. Association for Computing Machinery.
- Levy, R. and J. Haddad (2021, September). Path and Trajectory Planning for Autonomous Vehicles on Roads without Lanes. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, Indianapolis, IN, USA, pp. 3871–3876. IEEE.
- Levy, R. and J. Haddad (2022, November). Cooperative Path and Trajectory Planning for Autonomous Vehicles on Roads Without Lanes: A Laboratory Experimental Demonstration. *Transportation Research Part C: Emerging Technologies* 144, 103813.
- Li, N., D. W. Oyler, M. Zhang, Y. Yildiz, I. Kolmanovsky, and A. R. Girard (2018, September). Game Theoretic Modeling of Driver and Vehicle Interactions for Verification and Validation of Autonomous Vehicle Control Systems. *IEEE Transactions on Control Systems Technology* 26(5), 1782–1797.
- Li, Z., D. A. Hensher, and J. M. Rose (2010, May). Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. *Transportation Research Part E: Logistics and Transportation Review* 46(3), 384–403.
- Liao-McPherson, D., M. Huang, and I. Kolmanovsky (2019, July). A Regularized and Smoothed Fischer–Burmeister Method for Quadratic Programming With Applications to Model Predictive Control. *IEEE Transactions on Automatic Control* 64(7), 2937–2944.
- Lim, W., S. Lee, M. Sunwoo, and K. Jo (2021, January). Hybrid Trajectory Planning for Autonomous Driving in On-Road Dynamic Scenarios. *IEEE Transactions on Intelligent Transportation Systems* 22(1), 341–355. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Liu, H. X., W. Xin, Z. M. Adams, and J. X. Ban (2007). A Game Theoretical Approach for Modelling Merging and Yielding Behavior at Freeway On-Ramp Sections.

- Liu, X. (2012, June). *Survival Analysis: Models and Applications*. John Wiley & Sons. Google-Books-ID: bEZpGtw39qgC.
- Ma, Z., J. Xie, X. Qi, Y. Xu, and J. Sun (2017). Two-Dimensional Simulation of Turning Behavior in Potential Conflict Area of Mixed-Flow Intersections. *Computer-Aided Civil and Infrastructure Engineering* 32(5), 412–428. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/mice.12266>.
- Mannering, F. L. and S. S. Washburn (2020, July). *Principles of Highway Engineering and Traffic Analysis*. John Wiley & Sons. Google-Books-ID: akn8DwAAQBAJ.
- Mo, Z., R. Shi, and X. Di (2021, September). A physics-informed deep learning paradigm for car-following models. *Transportation Research Part C: Emerging Technologies* 130, 103240. Publisher: Elsevier BV.
- Nash, J. (1951). Non-Cooperative Games. *Annals of Mathematics* 54(2), 286–295. Publisher: [Annals of Mathematics, Trustees of Princeton University on Behalf of the Annals of Mathematics, Mathematics Department, Princeton University].
- Nogueira, F. (2014). Bayesian Optimization: Open source constrained global optimization tool for Python.
- Ohtsuka, T. (2004, April). A Continuation/GMRES Method for Fast Computation of Nonlinear Receding Horizon Control. *Automatica* 40(4), 563–574.
- Ossen, S. and S. P. Hoogendoorn (2011, April). Heterogeneity in car-following behavior: Theory and empirics. *Transportation Research Part C: Emerging Technologies* 19(2), 182–195.
- Oyler, D. W., Y. Yildiz, A. R. Girard, N. I. Li, and I. V. Kolmanovsky (2016, July). A game theoretical model of traffic with multiple interacting drivers for use in autonomous vehicle development. In *2016 American Control Conference (ACC)*, Boston, MA, USA, pp. 1705–1710. IEEE.
- Papageorgiou, M., K.-S. Mountakis, I. Karafyllis, I. Papamichail, and Y. Wang (2021, February). Lane-Free Artificial-Fluid Concept for Vehicular Traffic. *Proceedings of the IEEE* 109(2), 114–121.
- Prokop, G. (2001, January). Modeling Human Vehicle Driving by Model Predictive Online Optimization. *Vehicle System Dynamics* 35(1), 19–53.

- Punzo, V., B. Ciuffo, and M. Montanino (2012, January). Can Results of car-following Model Calibration Based on Trajectory Data be Trusted? *Transportation Research Record* 2315(1), 11–24. Publisher: SAGE Publications Inc.
- Punzo, V. and F. Simonelli (2005, January). Analysis and Comparison of Microscopic Traffic Flow Models with Real Traffic Microscopic Data. *Transportation Research Record: Journal of the Transportation Research Board* 1934(1), 53–63. Publisher: SAGE Publications.
- Rostami-Shahrbabaki, M., S. Weikl, T. Niels, and K. Bogenberger (2023, September). Modeling Vehicle Flocking in Lane-Free Automated Traffic. *Transportation Research Record* 2677(9), 499–512. Publisher: SAGE Publications Inc.
- Sankar, G. S. and K. Han (2020, December). Adaptive Robust Game-Theoretic Decision Making Strategy for Autonomous Vehicles in Highway. *IEEE Transactions on Vehicular Technology* 69(12), 14484–14493.
- Schakel, W. J., V. L. Knoop, and B. van Arem (2012, January). Integrated Lane Change Model with Relaxation and Synchronization. *Transportation Research Record* 2316(1), 47–57. Publisher: SAGE Publications Inc.
- Sharath, M. N. and N. R. Velaga (2020, November). Enhanced intelligent driver model for two-dimensional motion planning in mixed traffic. *Transportation Research Part C: Emerging Technologies* 120, 102780.
- Shi, K., Y. Wu, H. Shi, Y. Zhou, and B. Ran (2022, August). An integrated car-following and lane changing vehicle trajectory prediction algorithm based on a deep neural network. *Physica A: Statistical Mechanics and its Applications* 599, 127303.
- Shi, X. and X. Li (2023, July). Trajectory Planning for an Autonomous Vehicle with Conflicting Moving Objects Along a Fixed Path – An Exact Solution Method. *Transportation Research Part B: Methodological* 173, 228–246.
- Shimizu, Y., T. Ohtsuka, and M. Diehl (2009, May). A real-time algorithm for nonlinear receding horizon control using multiple shooting and continuation/Krylov method. *International Journal of Robust and Nonlinear Control* 19(8), 919–936.
- Singh, S., D. Feng, P. Keller, G. Shaffer, W. F. Shi, D. H. Shin, J. West, and B. X. Wu (1991). A system for fast navigation of autonomous vehicles. Technical report, Carnegie Mellon University.

- Stackelberg, H. v. (2010, November). *Market Structure and Equilibrium*. Springer Science & Business Media. Google-Books-ID: dghH9OH5fDoC.
- Stahl, D. O. and P. W. Wilson (1995, July). On Players Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior* 10(1), 218–254.
- Stoyan, D., W. S. Kendall, S. N. Chiu, and J. Mecke (2013, June). *Stochastic Geometry and Its Applications*. John Wiley & Sons.
- Sun, D. J. and L. Elefteriadou (2014, May). A Driver Behavior-Based Lane-Changing Model for Urban Arterial Streets. *Transportation Science* 48(2), 184–205.
- Talebpour, A., H. S. Mahmassani, and S. H. Hamdar (2015). Modeling Lane-Changing Behavior in a Connected Environment: A Game Theory Approach. *Transportation Research Procedia* 7, 420–440.
- Tang, W. and D. M. Levinson (2018, August). Deviation between Actual and Shortest Travel Time Paths for Commuters. *Journal of Transportation Engineering, Part A: Systems* 144(8), 04018042. Publisher: American Society of Civil Engineers.
- Toledo, T. (2007, January). Driving Behaviour: Models and Challenges. *Transport Reviews* 27(1), 65–84.
- Toledo, T., C. F. Choudhury, and M. E. Ben-Akiva (2005, January). Lane-Changing Model with Explicit Target Lane Choice. *Transportation Research Record* 1934(1), 157–165. Publisher: SAGE Publications Inc.
- Toledo, T., H. N. Koutsopoulos, and M. Ben-Akiva (2007, April). Integrated driving behavior modeling. *Transportation Research Part C: Emerging Technologies* 15(2), 96–112.
- Toledo, T., H. N. Koutsopoulos, and M. E. Ben-Akiva (2003, January). Modeling Integrated Lane-Changing Behavior. *Transportation Research Record* 1857(1), 30–38. Publisher: SAGE Publications Inc.
- Tomar, R. S., S. Verma, and G. Tomar (2010, November). Prediction of Lane Change Trajectories through Neural Network. In *2010 International Conference on Computational Intelligence and Communication Networks*, pp. 249–253.
- Treiber, M., A. Hennecke, and D. Helbing (2000, August). Congested Traffic States in Empirical Observations and Microscopic Simulations. *Physical Review E* 62(2), 1805–1824.

- Troullinos, D., G. Chalkiadakis, D. Manolis, I. Papamichail, and M. Papageorgiou (2021, September). Lane- Free Microscopic Simulation for Connected and Automated Vehicles. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, Indianapolis, IN, USA, pp. 3292–3299. IEEE.
- Van Brummelen, J., M. O’Brien, D. Gruyer, and H. Najjaran (2018, April). Autonomous Vehicle Perception: The Technology of Today and Tomorrow. *Transportation Research Part C: Emerging Technologies* 89, 384–406.
- von Neumann, J., O. Morgenstern, and A. Rubinstein (1944). *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press.
- Von Toussaint, U. (2011, September). Bayesian inference in physics. *Reviews of Modern Physics* 83(3), 943–999.
- Wang, C. and B. Coifman (2008, September). The Effect of Lane-Change Maneuvers on a Simplified Car-Following Theory. *IEEE Transactions on Intelligent Transportation Systems* 9(3), 523–535. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Wang, H., E. Moylan, and D. M. Levinson (2022, June). Prediction of the Deviation between Alternative Routes and Actual Trajectories for Bicyclists. *Findings*.
- Wang, M., S. P. Hoogendoorn, W. Daamen, B. van Arem, and R. Happee (2015, September). Game Theoretic Approach for Predictive Lane-changing and Car-following Control. *Transportation Research Part C: Emerging Technologies* 58, 73–92.
- Wang, Z. and Y. Gao (2025, March). Lane Change Inconsistencies in the highD Dataset. *Findings*. Publisher: Findings Press.
- Wang, Z., P. Keo, and M. Saberi (2023). Real-Time Traffic State Measurement Using Autonomous Vehicles Open Data. *IEEE Open Journal of Intelligent Transportation Systems* 4, 602–610.
- Wang, Z., M. Ramezani, and D. Levinson (2024, August). How mandatory are ‘Mandatory’ lane changes? An analytical and experimental study on the costs of missing freeway exits. *Transportation Research Part B: Methodological* 186, 102994.

- Wang, Z., X. Zhao, Z. Xu, X. Li, and X. Qu (2021). Modeling and field experiments on autonomous vehicle lane changing with surrounding human-driven vehicles. *Computer-Aided Civil and Infrastructure Engineering* 36(7), 877–889.
- Wilson, R. E. (2001, October). An analysis of Gipps's car-following model of highway traffic. *IMA Journal of Applied Mathematics* 66(5), 509–537. Conference Name: IMA Journal of Applied Mathematics.
- Xie, F. and D. Levinson (2007). Measuring the Structure of Road Networks. *Geographical Analysis* 39(3), 336–356. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.2007.00707.x>.
- Xie, Z., M. Ramezani, and D. Levinson (2024, November). Reduced-Scale Mobile Robots for Autonomous Driving Research. *IEEE Transactions on Intelligent Transportation Systems* 25(11), 15367–15387.
- Yang, C. and G. Gidófalvi (2018, March). Fast map matching, an algorithm integrating hidden Markov model with precomputation. *International Journal of Geographical Information Science* 32(3), 547–570.
- Yang, D., S. Zheng, C. Wen, P. J. Jin, and B. Ran (2018, October). A dynamic lane-changing trajectory planning model for automated vehicles. *Transportation Research Part C: Emerging Technologies* 95, 228–247.
- Yang, D., L. Zhu, Y. Liu, D. Wu, and B. Ran (2019, June). A Novel Car-Following Control Model Combining Machine Learning and Kinematics Models for Automated Vehicles. *IEEE Transactions on Intelligent Transportation Systems* 20(6), 1991–2000.
- Yang, Q. and H. N. Koutsopoulos (1996, June). A Microscopic Traffic Simulator for evaluation of dynamic traffic management systems. *Transportation Research Part C: Emerging Technologies* 4(3), 113–129.
- Yanumula, V. K., P. Typaldos, D. Troullinos, M. Malekzadeh, I. Papamichail, and M. Papageorgiou (2021, September). Optimal Path Planning for Connected and Automated Vehicles in Lane-free Traffic. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, Indianapolis, IN, USA, pp. 3545–3552. IEEE.
- Yanumula, V. K., P. Typaldos, D. Troullinos, M. Malekzadeh, I. Papamichail, and M. Papageorgiou (2023, March). Optimal Trajectory Planning for Connected and Automated Vehicles

- in Lane-Free Traffic With Vehicle Nudging. *IEEE Transactions on Intelligent Vehicles* 8(3), 2385–2399.
- Yue, H., L. Zhang, H. Shan, H. Liu, and Y. Liu (2015, October). Estimation of the Vehicle's Centre of Gravity Based on a Braking Model. *Vehicle System Dynamics* 53(10), 1520–1533.
- Zhang, S., W. Deng, Q. Zhao, H. Sun, and B. Litkouhi (2013, October). Dynamic Trajectory Planning for Vehicle Autonomous Driving. In *2013 IEEE International Conference on Systems, Man, and Cybernetics*, pp. 4161–4166. ISSN: 1062-922X.
- Zhang, X., J. Sun, X. Qi, and J. Sun (2019, July). Simultaneous modeling of car-following and lane-changing behaviors using deep learning. *Transportation Research Part C: Emerging Technologies* 104, 287–304.
- Zhang, Y., X. Chen, J. Wang, Z. Zheng, and K. Wu (2022, December). A generative car-following model conditioned on driving styles. *Transportation Research Part C: Emerging Technologies* 145, 103926. Publisher: Elsevier BV.
- Zhao, J., V. L. Knoop, J. Sun, Z. Ma, and M. Wang (2023, September). Unprotected Left-Turn Behavior Model Capturing Path Variations at Intersections. *IEEE Transactions on Intelligent Transportation Systems* 24(9), 9016–9030. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Zheng, Z., S. Ahn, D. Chen, and J. Laval (2013, January). The effects of lane-changing on the immediate follower: Anticipation, relaxation, and change in driver characteristics. *Transportation Research Part C: Emerging Technologies* 26, 367–379.
- Zhou, H., Y. Sun, X. Qin, X. Xu, and R. Yao (2020, March). Modeling discretionary lane-changing behavior on urban streets considering drivers' heterogeneity. *Transportation Letters* 12(3), 213–222.
- Zhu, S. and D. Levinson (2015, August). Do People Use the Shortest Path? An Empirical Test of Wardrop's First Principle. *PLOS ONE* 10(8), e0134322.
- Zhu, S., D. Levinson, H. X. Liu, and K. Harder (2010, December). The traffic and behavioral effects of the I-35W Mississippi River bridge collapse. *Transportation Research Part A: Policy and Practice* 44(10), 771–784.