



Delft University of Technology

Document Version

Final published version

Citation (APA)

Yao, X. (2026). *Driving Heterogeneity in Traffic Flow Theory: An Action-based Framework for Identification, Modelling, and Simulation*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:320b6223-75a8-4a1b-ba26-a54a44ff4d29>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy


Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Driving Heterogeneity in Traffic Flow Theory

An Action-based Framework for Identification, Modelling, and Simulation

Xue Yao



Propositions

accompanying the dissertation

Driving Heterogeneity in Traffic Flow Theory

An Action-Based Framework for Identification, Modelling, and Simulation

by

Xue YAO

1. Driver aggressiveness does not significantly degrade traffic flow performance; rather, heterogeneity does. [Chapter 3&7]
2. Complex human driving behaviour can be understood by breaking it down into structured and interpretable components. [Chapter 4&5]
3. Artificial intelligence reveals what driving data hides, while at the risk of being misleading without expert knowledge to interpret its outputs. [Chapter 5&6]
4. Instead of introducing new models that are only marginally different, the research community should prioritise empirical validation and real-world testing in the future. [Chapter 7]
5. Traffic flow theory is dying.
6. Large language models help overcome technical barriers, but make the knowledge fortress messy.
7. Traffic engineers should reduce heterogeneity wherever possible, while human beings should keep it as much as possible.
8. Most challenges of a PhD are psychological rather than intellectual.
9. Working more than 40 hours a week reduces productivity.
10. Many useful scientific ideas emerge while doing other activities rather than during focused desk work: a PhD candidate should reduce time spent at the desk.

These propositions are regarded as opposable and defensible, and have been approved as such by the promoters prof. dr. ir. S. P. Hoogendoorn, and dr. ir. S. C. Calvert.

Driving Heterogeneity in Traffic Flow Theory

An Action-based Framework for Identification, Modelling, and
Simulation

Xue YAO

Driving Heterogeneity in Traffic Flow Theory

An Action-based Framework for Identification, Modelling, and
Simulation

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus prof. dr. ir. H. Bijl,
Chair of the Board for Doctorates
to be defended publicly on
Wednesday, 21 January 2026 at 15:00

by

Xue YAO

Master of Traffic and Transportation Engineering,
Southwest Jiaotong University, China,
born in Shandong, China.

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus,	Chairperson
Prof. dr. ir. S. P. Hoogendoorn,	Delft University of Technology, promotor
Dr. ir. S. C. Calvert,	Delft University of Technology, copromotor

Independent members:

Prof. dr. ir. J.W.C. van Lint,	Delft University of Technology
Prof. dr. ir. D.A. Abbink,	Delft University of Technology
Prof. dr. D. Work,	Vanderbilt University, USA
Prof. dr. S. Ahn,	University of Wisconsin-Madison, USA
Prof. dr. M. Menendez,	New York University Abu Dhabi, UAE



The research leading to this dissertation is supported by the Automated Driving & Simulation Lab (ADaS) from the Department of Transport & Planning, Delft University of Technology.

TRAIL Thesis Series no. T2026/3, The Netherlands Research School TRAIL

TRAIL
P.O. BOX 5017
2600 GA Delft
The Netherlands
E-mail: info@rsTRAIL.nl

Thesis cover: designed by the author with the assistance of ChatGPT.
ISBN: 978-90-5584-378-7
Copyright © 2026 by Xue YAO

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilised in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission of the author.

Printed in the Netherlands

No journey is ever in vain – every step counts.

Contents

Summary	ix
Samenvatting	xiii
1 Introduction	1
<i>I State-of-the-art for Driving Heterogeneity Analysis</i>	<i>13</i>
2 Driving Heterogeneity Identification using Machine Learning: A Review and Analytical Framework	15
3 Investigation on Car-Following Heterogeneity and Its Impacts on Traffic Flow Performance	39
<i>II An Action-based Framework for Driving Heterogeneity Identification</i>	<i>63</i>
4 Identification of Driving Heterogeneity using Action-chains	65
5 A Novel Framework for Identifying Driving Heterogeneity through Action Patterns	77
<i>III A Pattern-based Approach for Driving Heterogeneity Modelling and Simulation</i>	<i>99</i>
6 Human Driving Pattern Modelling: A Knowledge-Enhanced Deep Learning Approach	101
7 A Pattern-based Framework for Modelling Driving Heterogeneity and Traffic Flow Simulation	125
8 Conclusions and Perspectives	149
Bibliography	157
Acknowledgments	177
About the Author	181
List of Publications	183
TRAIL Thesis Series	185

Summary

Human driving behaviour is inherently heterogeneous, shaped by individual differences in perception, decision-making, risk tolerance, and control strategies. These behavioural variations significantly influence traffic dynamics, affecting traffic safety, flow stability, energy consumption, and emissions. With the emergence of automated vehicles (AVs), understanding and modelling this behavioural variability has become increasingly important, particularly for ensuring seamless human-AV interaction in mixed-traffic environments. To tackle challenges posed by heterogeneity in driving behaviour, this dissertation proposes a data-driven, AI-powered, and interpretable framework for identifying, modelling, and simulating heterogeneous driving behaviours to investigate how driving heterogeneity impacts traffic operations. By leveraging large-scale naturalistic driving datasets and artificial intelligence (AI), this dissertation aims to advance traffic flow theory and the design of more responsive, safer, and sustainable intelligent transportation systems (ITS).

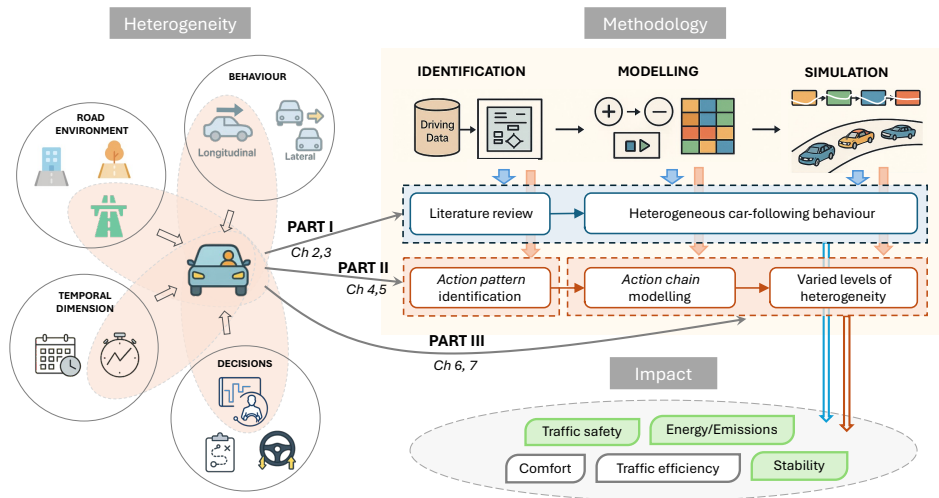


Figure S1: Conceptual framework as an outline of this dissertation.

Figure S1 outlines the conceptual framework developed in this dissertation to investigate driving heterogeneity and its impact on traffic operations. The scope of this dissertation, highlighted by the orange shading, is restricted to *longitudinal driving behaviour* of passenger cars on *highways*, with a focus on *tactical-level decision-making* and *short-term behavioural heterogeneity*. The dissertation is structured into three parts by solving three research questions:

1. What is driving heterogeneity and why does it matter?
2. How to identify driving heterogeneity through behavioural characteristics?
3. How to model and evaluate heterogeneous driving behaviours?

Methodology

PART I provides a comprehensive examination of existing approaches for driving heterogeneity analyses, including a systematic review of machine learning (ML) techniques for heterogeneity identification (*Chapter 2*) and a simulation-based investigation into the impact of car-following heterogeneity on traffic flow performance (*Chapter 3*). Based on the literature review, a four-step ML framework is proposed to structure the identification process, consisting of data collection, feature extraction, model training, and performance evaluation. Semi-supervised learning is used to identify distinct car-following styles, which are subsequently embedded into calibrated car-following models. These models are applied in micro-simulations across 66 traffic scenarios to evaluate how behavioural heterogeneity affects traffic safety and energy efficiency.

PART II introduces a novel *action-based* framework to uncover the underlying mechanisms of driving behaviour and identify heterogeneity more interpretably. This framework is structured around three behavioural units: *Action phases* (primitive driving units), *Action patterns* (group-specific behaviours), and *Action chains* (sequential behavioural structures). These units are extracted using rule-based segmentation and unsupervised learning methods, resulting in interpretable and structured driving patterns. (*Chapter 4&5*).

PART III extends the action-based framework by developing pattern-based modelling and simulation approaches. Longitudinal driving behaviour is modelled as sequences of *Action patterns* using deep learning models, where behavioural knowledge is explicitly integrated to improve predictive performance (*Chapter 6*). A bi-level modelling and simulation approach is then proposed to model heterogeneous driving behaviour, with the upper-level capturing behavioural characteristics of *Action chains*, and the low-level simulating corresponding vehicle dynamics within each identified pattern. This integrated approach is applied in micro-simulations to replicate traffic scenarios with varying levels of behavioural heterogeneity and to examine their impacts on traffic safety, energy efficiency, and flow stability (*Chapter 7*).

Results

The key takeaways and discussions of this dissertation are structured according to its three main research components.

PART I State-of-the-art for driving heterogeneity analysis provides a foundation by examining existing identification methods and revealing the importance of interpretability and contextual adaptability in identifying behavioural variability. Two major outcomes are highlighted:

- ☞ **The ML-based driving heterogeneity identification process can be structured as a four-step framework**, comprising *Trajectory Data Preparation, Traffic Feature Selection, Identification Models of ML, and Performance Evaluation*. The comprehensive

review of ML methodologies emphasises the need to balance accuracy, interpretability, and real-time recognition for effective heterogeneity identification. (*Chapter 2*)

- ☞ **Driver aggressiveness substantially influences traffic safety, fuel consumption, and emissions.** This is because less aggressive drivers can lead to the formation of vehicular platoons, thereby encouraging more aggressive drivers to adopt a milder driving style. Notably, the formation of these platoons is influenced by both the proportion and spatial distribution of less aggressive vehicles, which makes the correlation between aggressive driving and improvements in safety and environmental sustainability more complex. (*Chapter 3*)

PART II An Action-based framework for driving heterogeneity identification introduces a novel action framework that decodes driving behaviour into interpretable units, providing a scalable and explainable foundation for capturing both intra- and inter-driver behavioural heterogeneity. The main results are as follows:

- ☞ **Action phases provide interpretable building blocks for understanding driving variability.** Longitudinal driving behaviour is decomposed into physically meaningful primitives called *Action phases*, extracted from vehicle trajectories. Transitions between phases reveal intra-driver behavioural characteristics, while the constructed *Action phase Library* represents behavioural diversity at the traffic flow level. (*Chapter 4*)
- ☞ **Six calibrated Action patterns capture group-specific behavioural characteristics and enable rigorous labelling.** By clustering *Action phases*, the action-based framework identifies key driving patterns and driving variable importance, enabling interpretable labelling for supervised driving pattern classification. This framework promises advantages in traditional identification methods by capturing subtle behavioural differences within and among drivers. (*Chapter 5*)

PART III A Pattern-based approach for driving heterogeneity modelling and simulation integrates the action-based behavioural constructs into a modelling and simulation framework that enables the realistic simulation of a wide range of traffic scenarios, reflecting both individual behavioural variability and collective heterogeneity at the traffic flow level. The key findings are:

- ☞ **Embedding behavioural knowledge improves deep learning performance in driving pattern prediction.** An attention-based LSTM model (KE-ALSTM) is developed to model sequences of *Action patterns*, incorporating *Action pattern* transitions and duration properties. This integration significantly improves prediction accuracy, demonstrating the benefit of knowledge-enhanced learning. (*Chapter 6*)
- ☞ **The proposed bi-level modelling and simulation framework replicates realistic traffic phenomena driven by behavioural variability.** By capturing the complex effects of heterogeneity on traffic flow dynamics, the framework reveals how greater behavioural variability amplifies instability, inefficient spacing, and unsafe interactions, ultimately degrading traffic safety, energy efficiency, and flow stability. (*Chapter 7*).

Practical relevance and implications

This dissertation demonstrates how artificial intelligence can reveal driving characteristics that are not readily observable, providing insights into underlying mechanisms of behavioural heterogeneity and temporal dynamics that enrich our

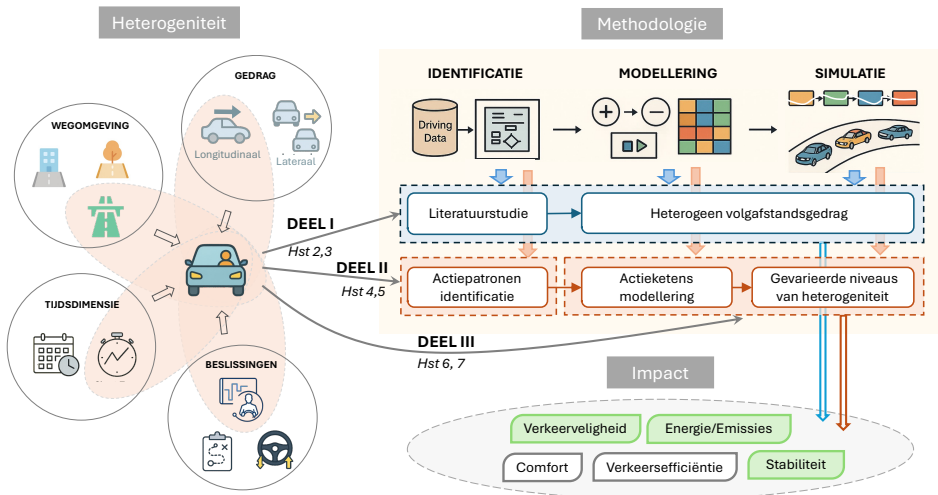
understanding of traffic flow. The developed framework demonstrates how varying levels of behavioural variability introduce disturbances, such as instability, inefficient spacing, and unsafe interactions, into traffic systems. The dissertation also contributes to the integration of domain knowledge, such as behavioural characteristics and temporal patterns, into AI models, which improves their robustness and generalisability, especially in data-scarce or uncertain contexts. The findings of this dissertation support practical applications across three domains:

1. **User-based traffic management:** By simulating traffic dynamics under different behavioural compositions, the framework enables the evaluation and optimisation of control strategies such as traffic signal timing and variable speed limits. These strategies can be adapted to account for the behavioural diversity of road users, enhancing traffic safety, flow stability, and overall network efficiency.
2. **Personalised driver assistance:** The identification and quantification of individual behavioural traits, such as acceleration tendencies, preferred headways, and responsiveness, allow for the development of tailored driver assistance systems, including adaptive cruise control and eco-driving functionalities that adjust control actions in real-time, ultimately improving comfort, safety, and energy efficiency.
3. **Context-aware AV design:** By embedding knowledge of typical human driving patterns and their heterogeneity, the framework contributes to the design of autonomous vehicles that can better interpret, predict, and respond to surrounding human drivers. This behavioural awareness enhances AV performance in mixed-traffic environments, promoting safer and more cooperative interactions.

Together, this dissertation lays a foundation for future mobility systems that are safer, more responsive, and more adaptable to individual and collective driving behaviours in increasingly complex traffic environments.

Samenvatting

Menselijk rijgedrag is van nature heterogeen, gevormd door individuele verschillen in waarneming, besluitvorming, risicotolerantie en controlestrategieën. Deze gedragsvariaties hebben een aanzienlijke invloed op de verkeersdynamiek en beïnvloeden de verkeersveiligheid, doorstromingsstabiliteit, energieverbruik en emissies. Met de opkomst van geautomatiseerde of Autonome voertuigen (AVs) wordt het begrijpen en modelleren van deze gedragsvariabiliteit steeds belangrijker, met name om een soepele interactie tussen mens en AV in gemengd verkeer te waarborgen. Om de uitdagingen van gedragsheterogeniteit aan te pakken, stelt dit proefschrift een data-gedreven, *kunstmatige intelligentie (AI)*-ondersteund en interpreteerbaar raamwerk voor, gericht op het identificeren, modelleren en simuleren van heterogeen rijgedrag, met als doel de impact ervan op verkeersoperaties te onderzoeken. Door gebruik te maken van grootschalige naturalistische rijgegevens en AI beoogt dit proefschrift bij te dragen aan de ontwikkeling van verkeersstroomtheorieën en het ontwerp van responsievere, veiligere en duurzamere intelligente transportsystemen (ITS).



Figuur S1: Conceptueel raamwerk als overzicht van dit proefschrift.

Figuur S1 geeft een overzicht van het conceptueel raamwerk dat in dit proefschrift is ontwikkeld om rijgedrag heterogeniteit en de impact daarvan op verkeersoperaties te onderzoeken. De scope van dit proefschrift, aangegeven met een oranje arcering, richt zich op *longitudinaal rijgedrag* van personenauto's op *snelwegen*, met de nadruk op *tactisch niveau*

van *besluitvorming* en *korte-termijn gedragsheterogeniteit*. Het proefschrift is opgebouwd uit drie delen die elk een onderzoeksvraag beantwoorden:

1. Wat is rijgedragsheterogeniteit en waarom is het belangrijk?
2. Hoe kan gedragsheterogeniteit worden geïdentificeerd via gedragskenmerken?
3. Hoe kan heterogeen rijgedrag worden gemodelleerd en geëvalueerd?

Methodologie

DEEL I biedt een uitgebreide analyse van bestaande benaderingen voor gedragsheterogeniteit, inclusief een systematische review van *machine learning* (ML)-technieken voor gedragsidentificatie (*Hoofdstuk 2*) en een simulatiegebaseerd onderzoek naar de impact van *car-following* heterogeniteit op de prestaties van de verkeersstroom (*Hoofdstuk 3*). Op basis van deze literatuurstudie wordt een vierstaps ML-raamwerk voorgesteld ter structurering van het identificatieproces, bestaande uit: dataverzameling, kenmerkextractie, modeltraining en prestatie-evaluatie. Semigesuperviseerd leren wordt toegepast om verschillende volgstijlen te identificeren, die vervolgens worden ingebed in gekalibreerde *car-following* modellen. Deze modellen worden toegepast in micro-simulaties van 66 verkeersscenario's om te evalueren hoe gedragsheterogeniteit de verkeersveiligheid en het energieverbruik beïnvloedt.

DEEL II introduceert een nieuw *actiegebaseerd* raamwerk om de onderliggende mechanismen van rijgedrag bloot te leggen en gedragsheterogeniteit op een meer interpreteerbare manier te identificeren. Dit raamwerk is opgebouwd uit drie gedragslagen: *Actiefases* (elementaire rijgedragseenheden), *Actiepatronen* (groepsspecifiek gedrag) en *Actieketens* (sequentiële gedragsstructuren). Deze gedragslagen worden geëxtraheerd via regelgebaseerde segmentatie en *unsupervised learning*-methoden, wat resulteert in interpreteerbare en gestructureerde rijgedragspatronen (*Hoofdstuk 4 & 5*).

DEEL III bouwt voort op het actiegebaseerde raamwerk door patroon-gebaseerde modellering en simulatiebenaderingen te ontwikkelen. Longitudinaal rijgedrag wordt gemodelleerd als sequenties van 'Actiepatronen' met behulp van *deep learning*-modellen, waarbij gedragskennis expliciet wordt geïntegreerd om de voorspellingsprestaties te verbeteren (*Hoofdstuk 6*). Vervolgens wordt een tweeledig modelleer- en simulatiekader voorgesteld: het bovenliggende niveau vat de gedragskenmerken van 'Actieketens' samen, terwijl het onderliggende niveau de voertuigdynamica simuleert binnen elk geïdentificeerd patroon. Deze geïntegreerde aanpak wordt toegepast in micro-simulaties om verkeersscenario's met verschillende niveaus van gedragsheterogeniteit te repliceren, en om de effecten ervan op verkeersveiligheid, energie-efficiëntie en doorstromingsstabiliteit te analyseren (*Hoofdstuk 7*).

Resultaten

De belangrijkste bevindingen en discussies van dit proefschrift zijn gestructureerd aan de hand van de drie hoofdonderdelen van het onderzoek.

DEEL I – State-of-the-art in analyse van rijgedrag heterogeniteit legt de basis door bestaande methoden voor gedragsidentificatie te bestuderen en benadrukt het belang van interpreteerbaarheid en contextuele toepasbaarheid bij het herkennen van gedragsvariatie. Twee hoofdbevindingen worden belicht:

☞ **Het ML-gebaseerde identificatieproces van rijgedragheterogeniteit kan worden gestructureerd als een vierstapsraamwerk**, bestaande uit *Trajectdatavoorbereiding, Selectie van verkeerskenmerken, ML-modellen voor identificatie en Prestatie-evaluatie*. De uitgebreide review van ML-methoden onderstreept de noodzaak om nauwkeurigheid, interpreteerbaarheid en real-time herkenning in balans te brengen voor een effectieve identificatie van heterogeniteit. (*Hoofdstuk 2*)

☞ **Rij-agressiviteit beïnvloedt de verkeersveiligheid, het brandstofverbruik en de uitstoot aanzienlijk**. Minder agressieve bestuurders stimuleren de vorming van voertuigenplatoons, wat ertoe kan leiden dat agressievere bestuurders een rustiger rijstijl aannemen. De vorming van deze pelotons wordt echter beïnvloed door zowel de proportie als de ruimtelijke verdeling van minder agressieve voertuigen, waardoor de relatie tussen agressief rijgedrag en verbeteringen in veiligheid en duurzaamheid complexer wordt. (*Hoofdstuk 3*)

DEEL II – Een actiegebaseerd raamwerk voor identificatie van rijgedrag heterogeniteit introduceert een nieuw actiegebaseerd raamwerk dat rijgedrag decodeert in interpreteerbare eenheden en zo een schaalbare en uitlegbare basis biedt voor het vastleggen van zowel intra- als interbestuurdersvariatie. De belangrijkste resultaten zijn als volgt:

☞ **Actiefases vormen interpreteerbaar fundament om rijgedragsvariatie te begrijpen**. Longitudinaal rijgedrag wordt opgesplitst in fysiek betekenisvolle elementen, genaamd *actiefases*, die uit voertuigtrajecten worden geëxtraheerd. Overgangen tussen fases onthullen intra-bestuurderskenmerken, terwijl de samengestelde *Actiefasebibliotheek* gedragsdiversiteit op verkeersstroombniveau representeert. (*Hoofdstuk 4*)

☞ **Zes gekalibreerde actiepatronen vatten groepsspecifieke gedragskenmerken en maken nauwkeurige labelen mogelijk**. Door clustering van actiefases identificeert het raamwerk belangrijke rijpatronen en de bijbehorende gedragsvariabelen, wat interpreteerbare labeling mogelijk maakt voor gesuperviseerde classificatie. Dit raamwerk biedt voordelen ten opzichte van traditionele methoden door subtiele gedragsverschillen binnen en tussen bestuurders vast te leggen. (*Hoofdstuk 5*)

DEEL III – Een patroongebaseerde benadering voor modellering en simulatie van rijgedrag heterogeniteit integreert de actiegebaseerde gedragsstructuren in een modellering- en simulatiekader dat de realistische simulatie mogelijk maakt van een breed scala aan verkeersscenario's, waarin zowel individuele gedragsvariatie als collectieve heterogeniteit op het niveau van de verkeersstroom tot uiting komen. De belangrijkste bevindingen zijn:

☞ **Het integreren van gedragskennis verbetert de prestaties van deep learning bij het voorspellen van rijpatronen**. Een op aandacht gebaseerd LSTM-model (KE-ALSTM) is ontwikkeld om sequenties van *actiepatronen* te modelleren, inclusief de overgangen en duurkenmerken ervan. Deze integratie verhoogt de voorspellingsnauwkeurigheid aanzienlijk en toont het voordeel aan van kennisverrijkte leermethoden. (*Hoofdstuk 6*)

☞ **Het voorgestelde tweeledig modelleer- en simulatiekader bootst realistische verkeersfenomenen na die worden aangedreven door gedragsvariatie**. Door de complexe effecten van heterogeniteit op verkeersdynamiek te modelleren, toont het raamwerk aan hoe grotere gedragsvariabiliteit leidt tot instabiliteit, inefficiënte

afstandscontrole en onveilige interacties, wat uiteindelijk de verkeersveiligheid, energie-efficiëntie en doorstromingsstabiliteit verslechtert. (*Hoofdstuk 7*)

Praktische relevantie en implicaties

Dit proefschrift toont aan hoe kunstmatige intelligentie verborgen rijgedragskenmerken kan blootleggen en inzicht geeft in de onderliggende mechanismen van gedragsheterogeniteit en temporele dynamiek, die onze kennis van verkeersstromen verrijken. Het ontwikkelde raamwerk laat zien hoe verschillende niveaus van gedragsvariatie verstoringen introduceren in verkeerssystemen, zoals instabiliteit, inefficiënte volgafstanden en onveilige interacties. Het proefschrift draagt ook bij aan de integratie van domeinkennis, zoals gedragskenmerken en temporele patronen, in AI-modellen, wat de robuustheid en generaliseerbaarheid van deze modellen verbetert, met name in contexten met beperkte gegevens of hoge onzekerheid. De bevindingen van dit proefschrift ondersteunen praktische toepassingen in drie domeinen:

1. **Gebruikersgericht verkeersmanagement:** Door verkeersdynamiek te simuleren onder verschillende gedragscomposities, maakt het raamwerk evaluatie en optimalisatie van strategieën zoals verkeerslichtregeling en variabele snelheidslimieten mogelijk. Deze strategieën kunnen worden aangepast aan de gedragsdiversiteit van weggebruikers, wat de veiligheid, stabiliteit en efficiëntie ten goede komt.
2. **Gepersonaliseerde rijondersteuning:** De identificatie en kwantificering van individuele gedragskenmerken, zoals acceleratievoorkeuren, gewenste volgafstanden en reactievermogen, maken de ontwikkeling mogelijk van rijhulpsystemen zoals adaptieve cruise control en eco-driving functies die zich in real-time aanpassen. Dit verbetert comfort, veiligheid en energie-efficiëntie.
3. **Contextbewust ontwerp van AV's:** Door kennis van typisch menselijk rijgedrag en de variatie daarin te integreren, draagt het raamwerk bij aan de ontwikkeling van autonome voertuigen die menselijk gedrag beter kunnen interpreteren, voorspellen en erop reageren. Deze gedragsbewustheid verbetert de prestaties van AV's in gemengd verkeer en bevordert veiligere en meer coöperatieve interacties.

Gezamenlijk legt dit proefschrift een fundament voor toekomstige mobiliteitssystemen die veiliger, responsiever en beter afgestemd zijn op zowel individuele als collectieve rijgedragingen binnen steeds complexere verkeersomgevingen.

1

Introduction

This chapter lays the foundation for the thesis, presenting the research topic, objectives, and relevance. Section 1.1 introduces the background, motivation, and theoretical foundations of this thesis. Section 1.2 clarifies the research objectives and scope. Section 1.3 presents the methodological approach adopted in the thesis. Section 1.4 highlights the key scientific and practical contributions. Finally, Section 1.5 outlines the overall structure of the thesis.

1.1 Research background

1.1.1 Driving heterogeneity in traffic flow

Traffic flow is inherently heterogeneous due to variations in vehicle performance, road conditions, and, most importantly, differences in human driving styles, skill levels, and cognitive abilities [1]. Drivers exhibit diverse perceptions, reaction times, risk tolerances, and control strategies, leading to inconsistencies in acceleration, deceleration, car-following, and lane-changing behaviours. Even the same driver may respond differently to similar traffic and environmental situations at different times.

One of the primary concerns arising from driving heterogeneity is traffic safety. Differences in driving styles can create mismatches in driver expectations, increasing the risk of collisions. For instance, aggressive drivers may accelerate and decelerate abruptly, making it difficult for following vehicles, especially those with slower reaction times, to respond appropriately, potentially leading to accidents [2]. A thorough understanding of these behavioural differences is crucial for improving predictive traffic models and effective safety strategies. Beyond safety concerns, fuel consumption and emissions are significantly impacted by driving heterogeneity. Frequent speed fluctuations, aggressive acceleration, and sudden braking have been shown to increase energy consumption and pollutant emissions. Studies indicate that aggressive driving in urban environments can raise fuel consumption and emissions by 30% to 40% compared to calm driving [3]. Conversely, adopting a more consistent driving style with controlled acceleration and deceleration can reduce fuel consumption and emissions by 10% to 25% [4]. Understanding these behavioural variations is essential for designing intelligent traffic control policies, promoting eco-driving strategies, and developing vehicle intervention technologies that contribute to sustainable traffic management.

With advancements in automated driving technology, managing mixed-traffic environments where human-driven vehicles (HDVs) and autonomous vehicles (AVs) coexist presents additional challenges. The heterogeneous nature of HDV behaviour complicates AV decision-making, which often relies on predefined models [5]. Moreover, heterogeneity can also exist among AVs due to variations in control algorithms and system designs across manufacturers. These discrepancies may lead to inefficient vehicle interactions, safety concerns, and reduced trust in automation. As a result, a deeper understanding of driving heterogeneity is critical for refining advanced driver assistance systems (ADAS) and improving AV algorithms to facilitate smoother integration into traffic systems.

To effectively address these challenges, a structured approach to modelling driving heterogeneity is necessary. This involves systematically identifying and modelling variations in driving behaviour and assessing how the complexity of driving heterogeneity

impacts traffic flow performance. By gaining deeper insights into these variations, researchers, technology developers, and mobility service providers can develop practical solutions, such as personalised in-vehicle assistance technologies, adaptive traffic control strategies, and adaptive AV algorithms, to enhance road safety and promote traffic sustainability.

1.1.2 AI-based approach for driving behaviour analysis

The increasing availability of large-scale naturalistic driving datasets has revolutionised the analysis of driving behaviour [6]. These datasets, collected from in-vehicle sensors, GPS trackers, LiDAR, and traffic surveillance systems, provide detailed, high-resolution information on vehicle trajectories, acceleration patterns, and driver decision-making processes [7]. The increasing scale and granularity of driving data have enabled artificial intelligence (AI)-based approaches, particularly machine learning (ML) and deep learning (DL) techniques, to emerge as powerful tools for traffic flow-related research. Unlike traditional models that rely on predefined assumptions, AI methods can automatically learn complex patterns and nonlinear relationships from data. These methods have been widely applied in traffic research to address key challenges such as driver profiling, trajectory prediction [8], and behaviour classification [9]. Machine learning techniques, including decision trees, support vector machines (SVMs), and clustering algorithms, have been used to categorise driving styles [10], detect anomalies [11], and analyse lane-changing patterns [12]. Deep learning models, such as recurrent neural networks (RNNs) and long short-term memory (LSTM) networks, have proven effective in capturing temporal dependencies in driving behaviour, enabling precise trajectory and intention prediction [13]. Additionally, reinforcement learning (RL) has been employed to optimise adaptive traffic control strategies [6] and enhance autonomous vehicle (AV) decision-making [14] in mixed-traffic environments.

Given the high variability and diversity of driving behaviours, data-driven AI techniques offer a promising tool for analysing driving heterogeneity. Leveraging AI-based insights can lead to the development of more accurate and scalable models that better reflect real-world traffic dynamics. These advancements hold practical implications for personalised in-vehicle driver assistance technologies, and adaptive traffic control strategies, ultimately contributing to safer, more efficient, and sustainable transportation systems.

1.1.3 Theoretical foundations of the thesis

This thesis is grounded in two important domains, including traffic flow theory and artificial intelligence (AI). The integration of these fields provides a structured foundation for analysing and modelling driving heterogeneity in a way that balances behavioural realism, computational scalability, and practical relevance.

Traffic flow theory offers a well-established foundation for understanding vehicle interactions and emergent traffic phenomena. Classical microscopic models, such as the Intelligent Driver Model (IDM), Gipps' model, and Wiedemann's model, describe how individual vehicles respond to surrounding traffic conditions based on predefined behavioural rules for car-following, lane-changing, and gap acceptance [15]. These models typically rely on mathematical equations with a set of parameters (e.g., desired time headway,

minimum gap, acceleration) to represent driver behaviour, which are calibrated using empirical data to reproduce observed driving behaviours. To overcome the rigidity of rule-based assumptions, data-driven models have been subsequently developed to learn driving behaviours directly from trajectory data, using methods such as regression, decision trees, and neural networks to capture complex and nonlinear driver responses [16]. However, both traditional and early data-driven models often rely on assumptions of behavioural consistency and homogeneity. In reality, driver behaviour varies substantially across individuals (inter-driver heterogeneity) and within the same individual across different conditions (intra-driver heterogeneity). Addressing this variability is critical for improving the realism and reliability of traffic models.

In this context, machine learning techniques, particularly advances in supervised learning, unsupervised learning, and interpretable AI, provide powerful tools for addressing traffic heterogeneity by uncovering complex behavioural patterns from large-scale naturalistic driving data [17]. For instance, clustering and classification techniques can distinguish driving styles or behavioural modes, while sequence models such as RNNs, and LSTMs capture how driver decisions evolve over time. These insights can be integrated into traffic simulations to better reflect real-world variability. This thesis builds on and extends traffic flow theory by explicitly modelling behavioural heterogeneity using AI techniques, aiming to reduce its negative impacts and enhance traffic performance in terms of safety and environmental sustainability.

1.2 Research objectives and scope

1.2.1 Research objectives

The main objective of this thesis is to advance the empirical understanding and the conceptual, data-driven and simulation-based modelling of driving heterogeneity in traffic flow, which can improve the behavioural realism of traffic simulations and support the development of personalised in-vehicle systems and adaptive traffic management strategies. The objective is to be achieved by developing a data-driven, interpretable framework that leverages AI-based tools to systematically identify, model, and assess the impacts of behavioural variability within and among drivers. This yields three main research questions with several sub-questions, which are as follows:

RQ1. What is driving heterogeneity and why does it matter?

RQ1.1. How is driving heterogeneity characterised and analysed in the literature? (Chapter 2)

RQ1.2. What are the impacts of driving heterogeneity on traffic flow performance? (Chapter 3)

The increasing availability of naturalistic driving data and the growing use of machine learning (ML) have led to a surge in methods for analysing driving behaviour. However, this growing body of literature is fragmented across subdomains, with varying definitions, driving data, and methodological choices, particularly in feature selection and ML model implementation. This lack of consistency poses challenges for comparative evaluation and

practical adoption of identification methods. Therefore, a systematic review and analytical framework are essential to clarify key concepts, methodological trends, and research gaps. Furthermore, understanding how heterogeneity impacts traffic performance indicators, such as safety and sustainability, provides the foundation for developing targeted practical applications. These questions aim to establish a comprehensive understanding of existing methods for driving heterogeneity analyses before moving toward novel approaches.

RQ2. How to identify driving heterogeneity through behavioural characteristics?

RQ2.1. What are the underlying mechanisms of driving heterogeneity? (Chapter 4)

RQ2.2. How can driving heterogeneity be identified from underlying driving characteristics? (Chapter 5)

Existing approaches for driving heterogeneity identification often assign fixed behavioural categories to drivers based on driving trajectory features such as velocity, acceleration, and distance. While intuitive, this approach fails to capture dynamic variations in driving behaviour that arise from temporal changes. Alternatively, high-dimensional representations of driving states obtained from clustering offer greater behavioural characteristics but often lack interpretability and generalisability, particularly when derived from black-box ML methods. There is therefore a need to develop identification methods that strike a balance between behavioural richness and analytical interpretability. These questions explore the underlying mechanisms of driving heterogeneity based on data-driven techniques, with a focus on identifying intra- and inter-driver variability.

RQ3. How to model and evaluate heterogeneous driving behaviours?

RQ3.1. How can heterogeneity in driving behaviour be effectively modelled? (Chapter 6)

RQ3.2. How can different levels of driving heterogeneity be captured in traffic simulation? (Chapter 7)

Although previous micro-simulation studies have demonstrated that driving heterogeneity influences traffic performance, they mostly rely on rough behavioural typologies with personalised behavioural model parameters, such as classified car-following (CCF) models [18]. Such approaches lack the granularity to capture dynamic temporal variability both within and across drivers. Consequently, these approaches fall short of representing the full spectrum of behavioural heterogeneity present in real-world traffic. These questions aim to facilitate a more flexible modelling framework that explicitly incorporates multiple levels of heterogeneity when simulating driving behaviour and effectively emphasises how different levels of driving heterogeneity impact traffic flow performance such as safety and sustainability.

1.2.2 Research scope

Driving heterogeneity arises from various sources, including but not limited to differences in individual driver characteristics, vehicle types, and environmental conditions [19]. This thesis focuses specifically on behavioural heterogeneity within and across drivers by analysing the behaviour of passenger cars in highway settings. To be more specific and a

well-defined examination of driving heterogeneity, the thesis concentrates on three key dimensions of driving behaviours:

1) *Longitudinal driving behaviour*: The research examines longitudinal driving behaviour, which in this thesis is defined by how drivers adjust their speed and maintain headway in response to surrounding traffic. Lateral behaviours, such as lane-changing, merging, and overtaking, are excluded due to different strategic decision-making.

2) *Tactical-level driving behaviour*: This thesis focuses on using real-world vehicle trajectory data to capture driving behaviour at a tactical level ¹. Higher-level strategic decisions (e.g., route choice) and lower-level operational controls (e.g., steering inputs), as well as unobservable human factors such as personality traits or cognitive state, are beyond the scope. This ensures that the analysis in this thesis remains grounded in observable, reproducible driving actions.

3) *Short-term heterogeneity perspective*: The research primarily focuses on short-term behavioural variations, analysing how drivers adjust their behaviour over seconds or minutes in response to immediate traffic conditions. These micro-level variations are essential for understanding diverse interactions between vehicles and their influence on traffic flow performance. In contrast, long-term heterogeneity, which encompasses stable driving traits observed over months or years, is not explicitly considered.

Within this well-defined scope, this thesis aims to develop data-driven methodologies for identifying and modelling driving heterogeneity, providing effective tools that can be applied to traffic flow analysis, and insights that can be used to improve simulation-based studies and ITS applications in practice.

1.3 Research approach

This dissertation study adopts a structured, three-step research approach to address the three main research questions systematically. The approach encompasses both baseline methods and a novel action-based framework for driving heterogeneity identification, modelling, and evaluation. Figure 1.1 provides a comprehensive overview of the various parts of approaches. Each step is guided by one research question.

(i) Investigation of the state-of-the-art (SOTA) methods for driving heterogeneity analysis

The first step addresses **Question 1** by investigating the state-of-the-art (SOTA) methods for driving heterogeneity analysis. It begins with a comprehensive literature review on driving heterogeneity identification, with a particular emphasis on machine learning (ML) methodologies. The review systematically organises the foundational concepts of driving heterogeneity and synthesises state-of-the-art methods into a structured analytical framework. The proposed framework includes data collection and preprocessing, feature engineering, ML modelling techniques, and evaluation strategies, illustrating how different

¹Strategic-level, tactical-level, and operational-level are hierarchical levels of driving behaviour that are defined by Michon [20]. Strategic-level behaviour involves pre-trip decisions such as route planning and activity scheduling. Tactical-level behaviour pertains to short-term manoeuvre planning in response to the immediate driving environment, such as decisions related to car-following and lane-changing. Operational-level behaviour involves moment-to-moment vehicle control, such as acceleration, braking, and steering, to execute the selected manoeuvres [21].

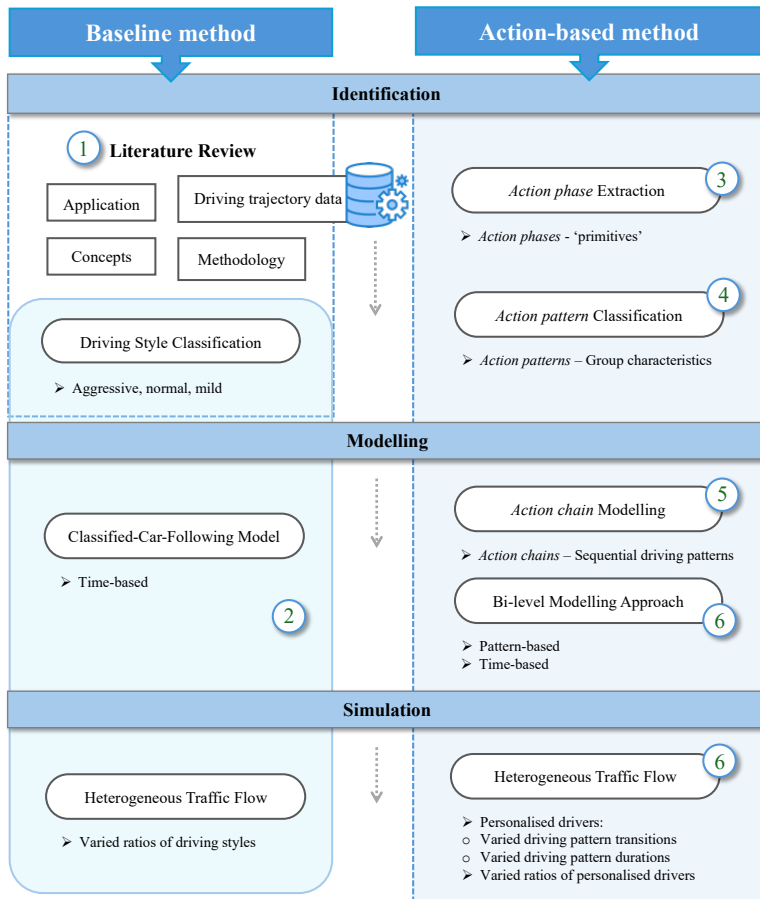


Figure 1.1: An overview of research approach

ML methods have been applied to analyse driving behaviour across various traffic conditions. This directly addresses **Question 1.1**. In addition, the impact of driving heterogeneity on traffic flow performance is examined using micro-simulation. Existing heterogeneity identification methods are integrated into micro-simulations to quantify their influence on traffic safety and sustainability. This directly addresses **Question 1.2**.

This part establishes a strong foundation by critically evaluating current methods and their limitations, thereby motivating the development of the proposed action-based framework.

(ii) Identification of driving heterogeneity from behavioural characteristics

The second part of the research aims to address **Question 2** by introducing a novel action-based framework for identifying and interpreting driving heterogeneity through

structured behavioural units. Three hierarchical concepts are proposed:

- *Action phase*: The most basic unit, characterised by distinct trends in driving variables such as speed, acceleration, and time headway.
- *Action pattern*: A composite behavioural unit capturing consistent driver responses to traffic conditions, built from sequences of *Action phases*.
- *Action chain*: A sequence of action patterns representing the evolving behavioural trajectory of a driver over time.

Drawing inspiration from gene regulation in biology, this framework views driving behaviour as a dynamic and adaptive process influenced by external stimuli—analogue to environmental triggers activating gene expression. In this analogy, *Action patterns* function like gene expressions [22], while *Action chains* represent regulatory sequences influencing behavioural evolution. This action-based framework provides a human-interpretable, information-driven methodology to decode and model heterogeneous driving behaviours. It allows systematic interpretation of how individual drivers respond to varying traffic conditions and how short-term actions accumulate into longer-term behavioural traits. This directly addresses **Questions 2.1 and 2.2**.

(iii) Modelling and evaluation of heterogeneous driving behaviours in traffic flow

The third part of the thesis focuses on addressing **Question 3** by modelling and evaluating heterogeneous driving behaviours in traffic flow. An action-based simulation framework is proposed to model longitudinal driving behaviours based on *Action chains*. This includes modelling both:

- *Action pattern* transitions: How drivers switch between behavioural states, and
- *Action pattern* durations: How long each behavioural state is maintained.

To enhance predictive performance and behavioural realism, transition dynamics and durations are encoded using both graph-based and probabilistic representations. This approach supports accurate driving behaviour modelling and addresses **Question 3.1**.

Furthermore, a bi-level modelling architecture is introduced for traffic simulation, where the high-level process captures the dynamics of *Action pattern* transitions and their corresponding durations, which in turn influences the parameters of the low-level model. For instance, when a driver transitions from *Action pattern A* to *Action pattern B*, the high-level model adjusts the expected acceleration, which then shapes the trajectory generated by the low-level vehicle dynamics model. The low-level modelling process generates driving trajectories by calibrating vehicle dynamics models for each *Action pattern*. This integrated model is implemented in micro-simulation model to evaluate the impacts of driving heterogeneity on traffic performance. Simulations are conducted under various driver compositions and traffic scenarios, considering both personalised driver behaviours (intra-heterogeneity) and population-level behavioural distributions (inter-heterogeneity). This directly addresses research **Question 3.2**.

The entire research workflow illustrated in Figure 1.1 highlights how the proposed action-based framework delivers a more interpretable, human-centric, and temporally adaptive understanding of driving behaviour heterogeneity. By structuring driving actions into meaningful behavioural units, the framework enhances the interpretability and adaptability of ML-based models in real-world traffic scenarios. Moreover, it offers a practical tool for simulating diverse driving behaviours and evaluating their impact on traffic flow performance, with implications for both research and policy.

1.4 Research contributions

1.4.1 Scientific contributions

This thesis contributes to the understanding and modelling of driving heterogeneity and its influence on traffic dynamics. By integrating data-driven techniques, behavioural analysis, and simulation-based methods, the research provides new empirical and theoretical insights and practical tools towards the role of heterogeneous driving behaviours in shaping traffic flow. The main scientific contributions of the thesis are summarised below:

(1) Structured framework for analysing driving heterogeneity

- An analytical framework is proposed to identify driving heterogeneity using machine learning techniques, offering a systematic way for AI-driven driving heterogeneity analysis (*Chapter 2*)
- A simulation-based approach is developed to quantify the effects of car-following heterogeneity on traffic safety and environmental sustainability, demonstrating the importance of incorporating behavioural variability in traffic studies. (*Chapter 3*)

(2) Behavioural pattern discovery from real-world data

- A multi-level analysis framework is introduced to uncover underlying behavioural patterns from real-world driving data, characterising both inter-driver and intra-driver heterogeneity. (*Chapter 4*)
- A pattern-based framework is developed to model and interpret driving heterogeneity through behavioural patterns, providing deeper insights into driving behaviour and its relation to traffic dynamics. (*Chapter 5*)

(3) Integration of domain knowledge into AI-based modelling

- A method for embedding domain knowledge of driving behaviour into data-driven models is proposed, enhancing the relevance and realism of AI-driven modelling. (*Chapters 4 & 5*)
- Knowledge-enhanced AI models are developed that incorporate structured behavioural information, improving the interpretability and predictive accuracy of driving behaviour models. (*Chapter 6*)

(4) Hierarchical simulation framework for heterogeneous traffic flow

- A multi-level simulation framework is developed to model heterogeneous traffic flow by capturing both inter-driver and intra-driver variability, enabling more realistic simulations and supporting better evaluation of traffic heterogeneity. (*Chapter 7*)

1.4.2 Practical relevance

In addition to its scientific advancements, this thesis offers practically applicable methods, tools, and insights that can inform real-world applications in adaptive traffic management, personalised driving assistance systems, and AV design, offering practical contributions that support the development of safer, more efficient, and more intelligent transportation systems. These are relevant to a range of stakeholders, including traffic engineers, policymakers, urban planners, automotive industry professionals, and developers of autonomous vehicle (AV) and traffic simulation technologies. The key practical contributions are summarised below:

(1) Enhanced identification of driving heterogeneity for traffic management

- A structured methodology is developed for identifying, modelling, and evaluating driving heterogeneity, providing researchers and practitioners with practical tools to assess behavioural variability in traffic and develop targeted interventions to improve traffic safety. (*Chapter 2*)
- Simulation-based analyses reveal how car-following heterogeneity affects safety and energy efficiency, offering insights to support sustainable traffic policies and eco-conscious planning. (*Chapter 3*)

(2) AI-driven behavioural modelling for safer and smarter road systems

- AI-based models are developed to improve the classification and prediction of driving behaviours under diverse traffic conditions. These models support practical applications in Advanced Driver Assistance Systems (ADAS) design, autonomous driving systems, and adaptive traffic signal control, enhancing road safety and operational performance. (*Chapter 4 & 5*)
- The proposed method for embedding domain knowledge into machine learning models improves the interpretability and real-world applicability of AI tools, particularly relevant for AV developers and transport system designers. (*Chapter 6*)

(3) Advanced micro-simulation for traffic flow optimisation

- The pattern-based modelling approach introduced in this thesis benefits microscopic traffic simulation platforms by improving behavioural realism, enabling more effective evaluation of traffic control strategies, infrastructure designs, and automation policies. (*Chapter 6*)

- A hierarchical simulation framework enables more realistic assessments of traffic dynamics, supporting the design and testing of personalised traffic management solutions and emerging ITS technologies such as connected and autonomous vehicles. (*Chapter 7*)

1.5 Thesis outline

The structure of this thesis, comprising an introduction at the beginning, three parts in the main body, and a conclusion section at the end, is shown in figure 1.2. The introduction in this Chapter has set out the background and necessity for studying driving heterogeneity. In the main body, Part I presents a state-of-the-art review of driving heterogeneity analysis. It includes a comprehensive literature review, an analytical framework for identifying heterogeneity using machine learning techniques, and an investigation of the impacts of car-following heterogeneity on traffic flow using a micro-simulation approach. Part II introduces a novel action-based framework for identifying driving heterogeneity, including the conceptual development of interpretable behavioural units in various granularity. Based on the action-based framework, Part III develops a hierarchical simulation approach to model heterogeneous driving behaviour and evaluate its impacts on traffic safety, and sustainability. Each chapter is introduced separately on its own title page and, where applicable, the source publications are given that make up the chapter. The thesis concludes with a summary of key findings and perspectives for future research.

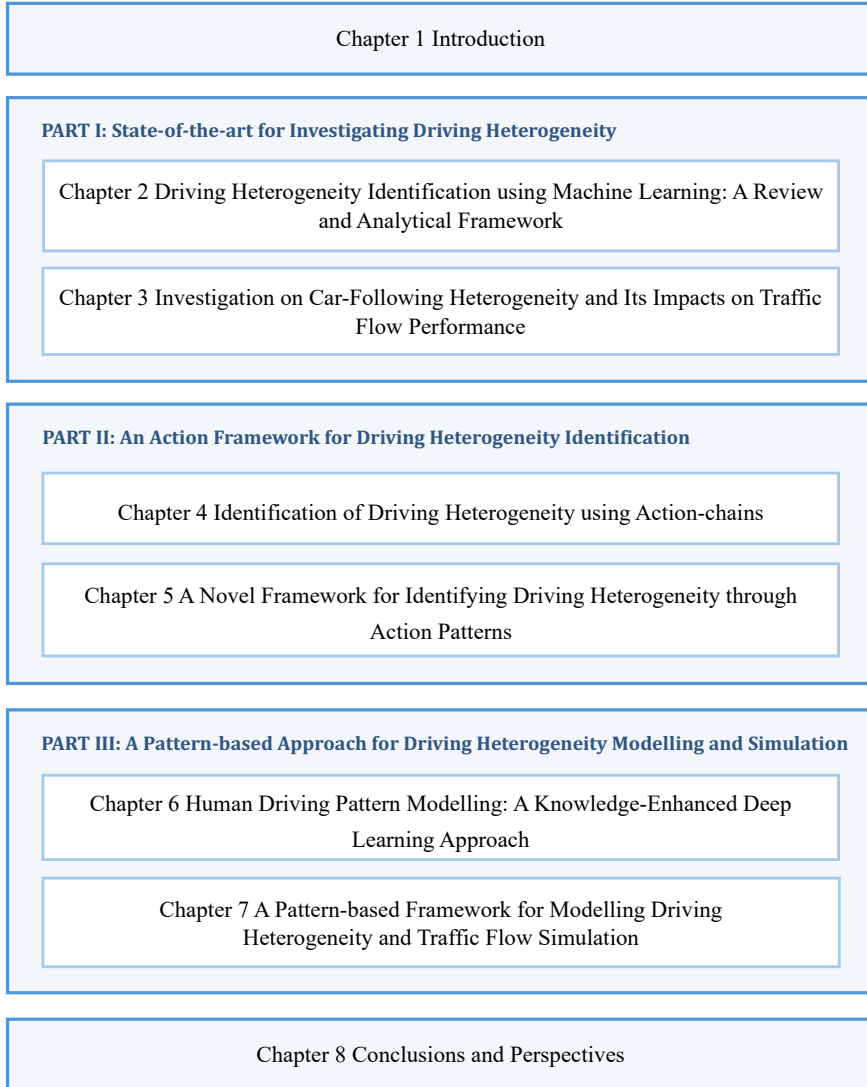


Figure 1.2: Thesis outline

I

State-of-the-art for Driving Heterogeneity Analysis

2

Driving Heterogeneity Identification using Machine Learning: A Review and Analytical Framework

The content of this chapter has been published on

📖 Yao, X., Calvert, S. C., and Hoogendoorn, S. P. (2025). “Driving Heterogeneity Identification Using Machine Learning: A Review and Framework for Analysis.” *Transportation Research Interdisciplinary Perspectives*, 32, 101511.

This chapter presents a systematic review of current Machine Learning (ML)-based methods for driving heterogeneity identification. The review organises key concepts and categorisations of driving heterogeneity, highlights strengths and drawbacks of various methods, and outlines applications of identification analysis. Based on the literature review, a structured framework that guides the ML-based identification process is proposed, including data collection and pre-processing, feature selection, ML model training, and performance evaluation.

2.1 Introduction

Driving behaviour plays a crucial role in shaping traffic flow, influencing road safety, and the overall sustainability of transportation systems [23]. The way a vehicle responds to driver inputs, along with environmental factors and propulsion dynamics, forms the basis of the vehicle-driver interaction. Importantly, drivers exhibit varying behaviours even under identical traffic conditions, a phenomenon known as driving heterogeneity. This variability has been shown to impact traffic performance by increasing crash risk, disrupting traffic flow, and contributing to higher fuel consumption and emissions [24, 25]. For example, delayed reaction times and reduced stimulus sensitivity have been linked to an elevated risk of rear-end collisions [26]. In mixed traffic environments where autonomous vehicles (AVs) and human-driven vehicles (HDVs) coexist, overlooking HDV heterogeneity can result in oversimplified AV behavioural models, thus increasing safety issues [27]. These issues underscore the need for accurate identification and modelling of human driving variability for both simulation and real-world applications.

Identifying driving heterogeneity requires rich data capturing a wide range of driver actions, such as speed, acceleration, and braking, as well as contextual factors like road conditions, traffic density, and weather. These data can come from various sources, including floating car data (FCD) via smartphones and GPS, in-vehicle sensors, or high-resolution imagery from roadside cameras and drones [28]. Based on driving behaviour data, heterogeneity analysis methods generally fall into two categories: subjective approaches (e.g., surveys and questionnaires) and objective approaches, which include rule-based logic, fuzzy systems, and increasingly, machine learning (ML) techniques [29]. With the rise of naturalistic driving datasets, ML techniques have become especially effective in capturing complex behavioural patterns due to their flexibility, high accuracy, and adaptability [30]. Models such as Support Vector Machines (SVM), k-Nearest Neighbours (KNN), and Feedforward Neural Networks (FFNN) have been widely used to classify driving styles, often achieving accuracy rates above 90% [18, 31]. Beyond traditional ML, deep learning models, including Long Short-Term Memory (LSTM) networks, have proven effective in modelling time-dependent behaviours, such as driver responses to external incentives or changes in driving workload [32].

Several previous reviews have explored driving heterogeneity from different angles, such as distinguishing driving styles and manoeuvres [33, 34], improving ADAS and safety systems [35–38], and evaluating vehicle–cloud collaboration [39]. However, these studies often focus on specific subdomains and vary in how they define and categorise driving heterogeneity. There remains a lack of a unified conceptual framework that systematically

organises key definitions, methodologies, and application pathways. In particular, ML-based approaches differ substantially in terms of data processing, feature selection, and model design [31, 40, 41], making it difficult to determine which techniques are most suitable for a given traffic context.

To address these gaps, this paper presents a comprehensive review of ML-based driving heterogeneity identification and proposes a structured framework to support its practical implementation. Our review focuses on longitudinal driving behaviours, especially in highway settings, and covers various traffic and environmental conditions. The goal is to advance data-driven strategies for understanding, classifying, and modelling driving heterogeneity, ultimately enabling personalised driver support, better traffic management, and safer vehicle automation systems. The main contributions of this study are twofold:

(i) **Comprehensive literature review:** We consolidate existing knowledge on driving heterogeneity, organising key concepts, behavioural categories, and ML-based identification methods. This provides a foundation for researchers to navigate and build upon existing work.

(ii) **A framework for analysis:** We propose a structural framework that supports heterogeneity identification through data collection and pre-processing, feature selection, model training, and performance evaluation. The framework incorporates supervised, unsupervised, semi-supervised, and reinforcement learning techniques, with a focus on their strengths, limitations, and suitability for different data types and use cases.

2.2 An overview of driving heterogeneity & identification

This section presents an overview of the review, as illustrated in Figure 2.1. The need for driving heterogeneity identification (WHY) stems from its relevance to real-world applications including traffic management, personalised ADAS, and human-like AV design. The review then organises key concepts of driving heterogeneity (WHAT) in multiple dimensions. Finally, we introduce the methodological process (HOW) for ML-based driving heterogeneity identification. This visual guide provides a structured foundation for the rest of the paper and sets the stage for the proposed analytical framework.

2.2.1 Applications of driving heterogeneity identification

Identifying driving heterogeneity has practical values in improving traffic management, enhancing road safety, and enabling personalised driver support systems. In traffic operations, understanding variations in driver behaviour allows for better predictions of congestion and more effective control strategies, such as adaptive signal timings or alternative route recommendations. It also helps in detecting unusual or high-risk behaviours, which can support real-time interventions such as issuing warnings to nearby vehicles or alerting authorities [42, 43]. In vehicle technology, heterogeneity identification enhances the customisation of Advanced Driver Assistance Systems (ADAS). For example, systems can adapt their feedback based on an individual's driving tendencies, providing earlier alerts to those prone to hard braking or enhanced lane assistance for frequent lane-changers [44, 45]. Similarly, for automated vehicles (AVs), recognising and responding to diverse human driving styles enables AVs to behave more naturally and safely in mixed

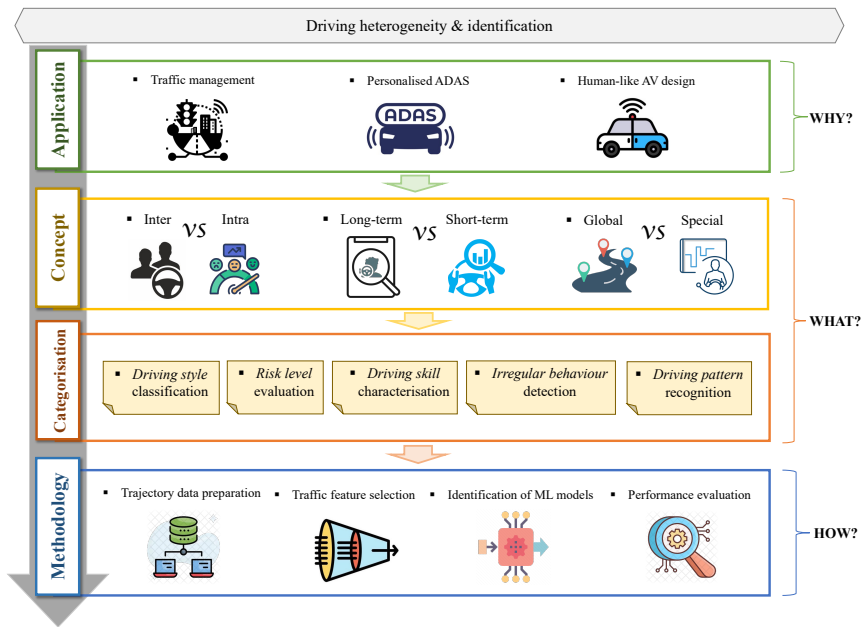


Figure 2.1: An overview of the literature review paper.

traffic environments [36]. These applications influence how we define and study driving heterogeneity by directing attention to specific behavioural differences and shaping the design of identification methods.

2.2.2 Concepts and categorisation of driving heterogeneity

Driving heterogeneity refers to the variability in driver traits, decision-making, and control actions. This variability manifests in both space and time [46], and can be categorised along three main dimensions: (i) **inter-driver** vs. **intra-driver** heterogeneity, (ii) **long-term** vs. **short-term** heterogeneity, and (iii) **global** vs. **special** behavioural patterns.

Inter-driver heterogeneity describes differences among drivers in similar conditions. For instance, some may accelerate more aggressively or maintain smaller headways than others [18, 24]. In contrast, intra-driver heterogeneity refers to how the same driver may behave differently over time or across situations [47, 48]. From a temporal perspective, long-term heterogeneity relates to persistent behavioural tendencies or driving skills developed over months or years, while short-term heterogeneity reflects temporary states like distraction or fatigue during a specific trip [49]. Lastly, global heterogeneity captures overall behaviour over a trip or time period, such as consistent car-following strategies [18], whereas special heterogeneity focuses on specific manoeuvres or behaviours, such as harsh braking or sharp turns [50].

Delineating these concepts provides a fundamental insight into understanding driving heterogeneity, which helps to describe heterogeneity in a human-comprehend manner.

Table 2.1: Categorisation of driving heterogeneity identification

Categorisation	Descriptors	Inter	Intra	Long-term	Short-term	Global	Special
<i>Driving style classification</i>	Aggressive, radical, normal, cautious, etc.	✓	✓	✓	✓	✓	✓
<i>Risk level evaluation</i>	High-moderate-low risk, etc.	✓	✓		✓	✓	✓
<i>Driving skill characterisation</i>	Expert, typical, etc.	✓		✓		✓	
<i>Irregular behaviour detection</i>	Aggressive braking, aggressive acceleration, etc.	✓			✓		✓
<i>Driving pattern recognition</i>	Closing in, closing in, keeping, falling behind, etc.	✓	✓		✓	✓	✓

To operationalise these ideas, researchers identify driving heterogeneity using categories including: *Driving style* (e.g., aggressive, normal, mild), *Risk level* (e.g., safe, risky), *Driving skill* (e.g., novice, expert), *Irregular behaviours* (e.g., harsh braking), and *Driving patterns* (e.g., acceleration phases, lane changes). Table 2.1 summarises these categories, their descriptors, and the type of driving heterogeneity they address. Each plays a distinct role in capturing variability across time, context, and individual differences. For example, driving styles are often associated with inter-driver and long-term global heterogeneity, while risk level and irregular behaviour focus on short-term, situational variability. Additionally, driving pattern recognition provides a flexible tool for capturing both intra-driver variation and broader behavioural tendencies.

Together, these concepts and descriptors for identifying driving heterogeneity not only enhance our understanding of the multifaceted nature of driver behaviour but also facilitate the interpretation and development of ML-based interventions aimed at improving road safety and traffic management.

2.2.3 Methodologies for driving heterogeneity identification

Identifying driving heterogeneity is typically formulated as a classification problem, using behavioural data to distinguish among different drivers or driving patterns. Based on the literature review, three main methodological approaches can be distinguished:

1. Classifying driving behaviours into distinct groups to specific driving profiles.
2. Creating an extensive set for driving states without interpretation.
3. Decomposing complex driving behaviour into simpler, more fundamental patterns with interpretation.

The first approach assigns drivers to predefined categories or clusters, often based on discrete scales (e.g., 2 to 5 groups) or numerical indices (e.g., a score from 0 to 10). For example, drivers may be classified into aggressive (also termed radical), normal (moderate, conventional), or mild (timid, conservative) styles [18, 44, 51–53]. Similarly, driver skill levels have been grouped as novice, typical, or expert [31, 54, 55]. While these

classifications offer clear and interpretable outputs, they are limited in capturing the full range of behavioural diversity due to their coarse granularity. Furthermore, the thresholds used to define these categories are often subjective, potentially introducing bias into the identification process.

Instead of directly classifying or clustering, the second approach creates a broad space of behavioural profiles to represent driving heterogeneity more flexibly. For instance, [56] proposed a high-dimensional style space containing over 20 behavioural types. Similarly, modelled individual behaviour using probabilistic distributions over different driving states rather than discrete groups [9]. By acknowledging more characteristics, this approach can detect a wider range of variations in driving behaviour. However, this extensive categorisation approach might compromise the interpretation of driving profiles, thus limiting its implementation. This highlights the need for methods that can balance complexity and interpretability in behaviour modelling.

The third approach decomposes driving behaviour into simpler, fundamental components, commonly referred to as “primitives”, to analyse heterogeneity at a finer resolution. These primitives represent short, distinct behavioural segments with identifiable characteristics. For example, [40] extracted primitives such as “following behind,” “closing”, “gentle acceleration”, and “aggressive deceleration” to model driving heterogeneity. [57, 58] introduced the concept of “action phases” as basic units to describe transitions in driving behaviour, enabling a clearer interpretation of behaviour dynamics. This approach allows for detailed analysis of intra- and inter-driver variability while maintaining semantic clarity, making it well-suited for both theoretical development and practical applications.

In summary, methodologies for identifying driving heterogeneity differ in their emphasis on interpretability, flexibility, and granularity. Category-based methods are simple and interpretable but limited in detail of behavioural characteristics; continuous profiling offers richer representation but lacks clarity; and pattern decomposition provides high interpretability with fine granularity. Machine learning (ML) techniques are commonly used across these approaches due to their capability to model complex behaviours and handle large datasets. In the next sections, we explore how ML is applied to support and enhance these methodologies.

2.3 A framework for driving heterogeneity analysis

In this section, we introduce the proposed framework for identifying driving heterogeneity using Machine Learning (ML) techniques. The framework is developed based on a comprehensive review of existing literature and consists of four main steps: *Trajectory Data Preparation*, *Traffic Feature Selection*, *Identification Models of ML*, and *Performance Evaluation*, as illustrated in Figure 2.2. The initial step involves collecting, cleaning, and pre-processing raw driving data to ensure it accurately represents real-world driving behaviour. This step is essential for reducing errors and improving the quality of subsequent analysis. The second step, *Traffic Feature Selection*, focuses on identifying relevant variables from the pre-processed data that are most informative for modelling driving behaviour. Effective feature selection enhances model accuracy and reduces computational complexity. The third step, *Identification Models of ML*, applies the selected

features to train ML models capable of detecting and classifying different driving behaviour patterns. This is the core step where heterogeneity is identified using various ML learning algorithms. Finally, *Performance Evaluation* assesses the effectiveness of the ML models. This includes not only traditional accuracy metrics but also interpretability, generalisation to new data, and real-time applicability in the real-world. The following sub-sections describe each step in detail, discussing the methods used in the literature along with their respective advantages and limitations.

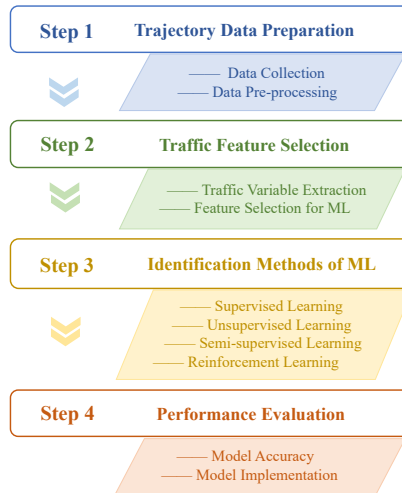


Figure 2.2: A framework for ML-based driving heterogeneity identification.

2.3.1 Step 1: Traffic Data Preparation

Since ML models rely on data to learn and make predictions, the quality of input data is critical to successful driving heterogeneity identification. Therefore, the first step in the framework is preparing trajectory data, which includes data collection and pre-processing. This step ensures that the data used is clean, reliable, and representative of actual driving behaviours.

➤ Data collection

Driving behaviour data is commonly collected using four methods: driving simulator, in-vehicle camera, sensor or hardware, traffic images, and floating car data (FCD, e.g., smartphone-based). These methods vary in controllability, data richness, quality, validity, and cost, as summarised in Table 2.2.

Controllability refers to how much researchers can influence the data collection environment. Driving simulators offer the highest control, allowing for designed experiments under specific conditions [54]. In-vehicle equipment allows to collect driving data in certain traffic scenarios, such as curving sections or ramps [59], while unpredictable situations could happen in real-world data collection, thus with lower controllability

Table 2.2: Comparison analysis of driving trajectory data collection method

Method	Controllability	Data richness	Quality	Validity	Cost
Driving simulator	High	<ul style="list-style-type: none"> • Scale: Small • Repeatability: Easy • Specific manoeuvres: Easy • Driver information: Easy 	Affected by <ul style="list-style-type: none"> • Scenarios design • Observer effect 	Low	High
In-vehicle equipment	Low	<ul style="list-style-type: none"> • Scale: Small • Repeatability: Easy • Specific manoeuvres: Easy • Driver information: Easy 	Affected by <ul style="list-style-type: none"> • Observer effects 	High	Moderate to high
Traffic images	No	<ul style="list-style-type: none"> • Scale: Large • Repeatability: Hard • Specific manoeuvres: Hard • Driver information: Hard 	Affected by <ul style="list-style-type: none"> • Observation errors • Parsing errors 	High	Moderate to high
FCD (Smartphone-based method)	No	<ul style="list-style-type: none"> • Scale: Large • Repeatability: Hard • Specific manoeuvres: Hard • Driver information: Hard 	Affected by <ul style="list-style-type: none"> • Sensor accuracy 	High	Low

than a driving simulator. Traffic images and FCD methods rely entirely on naturalistic driving, making them the least controllable. **Data richness** relates to the diversity and quantity of details available within a specific dataset. Driving simulators and in-vehicle equipment methods can provide driver information and manoeuvre-specific data and allow for repeatable data collection, but are limited in scale due to time and financial constraints. Conversely, traffic images and FCD methods capture large-scale driving behaviour but with less details about driver and manoeuvre information. **Quality** concerns the precision and objectivity of datasets. Data from driving simulators and in-vehicle equipment may suffer from observer effects. For example, drivers know they are observed and might exhibit different driving behaviours compared to driving in a real-world setting, thus reducing the objectivity of collected data [60]. Traffic images and FCD capture real-world behaviours but face issues with sensor accuracy, video resolution, and post-processing errors. **Validity** measures how accurately the data reflects actual driving behaviour. Driving simulators have low validity due to artificial environments, whereas in-vehicle equipment, traffic images, and FCD provide higher validity by capturing real-world driving under diverse conditions. **Cost** refers to the expenses, time, and human effort needed for data collection. Driving simulators are expensive due to equipment and participant costs. In-vehicle systems range from affordable GPS devices to costly telematics. Traffic images require expensive infrastructure and high data-processing costs. FCD, utilising built-in smartphone sensors, is the most cost-effective but raises privacy concerns [61].

Each method offers trade-offs in terms of control, scale, and precision. Researchers can leverage these insights to tailor their data collection strategies effectively, aligning with their research goals and specific research questions. For instance, studies focusing on detailed driver behaviour may prioritise driving simulators or in-vehicle equipment, while traffic management and policy research may favour traffic images or FCD for broader behavioural insights.

Figure 2.3 shows the use of various data collection methods over the years in the reviewed papers. Thanks to the development of technologies such as telematics, GPS




Figure 2.3: Statistics of data collection methods over the year

systems, and 5G, there is an increase in the availability and application of both FCD and in-vehicle equipment methods. The latter, in particular, has seen a trend towards more consistent and long-term usage in studies. As increasing numbers of publicly naturalistic driving datasets have become available since 2018, studies based on traffic image data have increased. Correspondingly, the use of driving simulator experiments and datasets has declined due to more available naturalistic data.

➤ Data pre-processing

Collected driving trajectory data often contains noise and inaccuracies due to sensor errors, video quality limitations, and data extraction inconsistencies [62, 63]. These issues can distort analysis results and misrepresent driving behaviour. To enhance data reliability, various pre-processing techniques are applied to maintain data integrity and optimise input quality for ML-based driving behaviour analysis. Table 2.3 provides a summary of techniques for outlier elimination, filtering, and data synchronisation. Regression-based methods and cubic interpolation are commonly used to detect and correct outliers by either adjusting values based on predictive modelling or estimating missing points from surrounding data [18, 64]. Additionally, filtering techniques, such as the Butterworth filter and Savitzky-Golay filter, are employed to smooth out noise while preserving critical data patterns [65, 66]. To ensure temporal consistency, data synchronisation techniques adjust the sampling rates of datasets. Up-sampling is applied to sparse data to increase resolution and retain essential behavioural details, while down-sampling simplifies large-scale datasets, improving computational efficiency without significant information loss [64, 67].

Overall, each data pre-processing method has its specific purpose when dealing with noise and maintaining data integrity. To ensure that the data is accurately represented to derive meaningful driving characteristics, the selection of data pre-processing methods should be carefully chosen according to the nature of the dataset involved.

2.3.2 Step 2: Traffic Feature Selection

The second step in the proposed framework is traffic feature selection, which aims to reduce dimensionality by selecting relevant variables from the pre-processed dataset. While datasets often include many features, not all of them contribute meaningfully to identifying

Table 2.3: Techniques for pre-processing driving behaviour data

Category	Technique	Characteristic
Outlier elimination	Regression [18]	- Identifies and adjusts anomalies by fitting a predictive model to the data.
	Cubic interpolation [64]	- Fills missing values by estimating them based on nearby data points, preserving dataset smoothness.
Filtering	Butterworth filter [65]	- Smooth response in the passband, preserving the true characteristics of driving data while effectively removing noise
	Savitzky-Golay filter [66]	- Retains data distribution shape for pattern consistency
Data synchronisation	Up-sampling [67]	- Increases the sampling rate in smaller datasets to capture more detailed information while maintaining consistency
	Down-sampling [64]	- Reduces the sampling rate in large-scale datasets, enhancing computational efficiency without significant data loss

driving heterogeneity. Using irrelevant or redundant features can reduce model accuracy and increase computational cost. Therefore, feature selection is essential to improve model performance and interpretability.

➤ **Traffic variable extraction**

There is currently no universally agreed-upon sets of metrics for driving behaviour analysis in literature. According to Abou et al. [34], metrics used in driving studies can be grouped into four categories: vehicle-based, behavioural, physiological, and subjective. Since this study focuses on trajectory-based analysis, we consider only vehicle-based variables, which include vehicle kinematic and dynamic features. Kinematic variables describe the vehicle’s motion, such as speed and acceleration, while dynamic variables reflect the driver’s control inputs, such as braking and throttle use. Both types of variables are widely used to characterise driving behaviour and detect heterogeneity [48, 68].

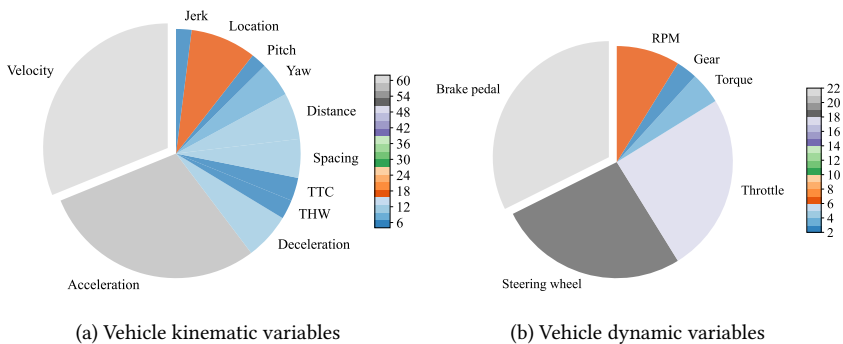


Figure 2.4: Statistics of traffic variables used in literature

Figure 2.4 summarises the use of these variables in existing studies. Kinematic variables are used more frequently than dynamic ones due to their strong correlation with driver

responses to traffic situations. Velocity (used in 61 studies) and acceleration (57) are the most common variables for identifying driving heterogeneity. Among dynamic variables, braking (21), steering wheel angle (18), and throttle position (17) are frequently used because they provide direct information about driving behaviours. The importance of these variables depends on the specific type of heterogeneity being studied. For instance, throttle usage has been found to be a strong indicator of aggressive driving [69], while combining RPM, speed, and acceleration improves driving style classification [70]. Moreover, integrating acceleration and brake events can increase classification accuracy by up to 10% according to Van Ly et al. [71]. Therefore, variable selection should align with the specific behavioural characteristics and heterogeneity concepts under investigation.

➤ Feature selection for ML

Machine learning (ML) models rely on carefully selected traffic features to improve predictive accuracy and computational efficiency. Driving behaviour studies often generate extensive feature sets by incorporating various statistical descriptors such as maximum, minimum, and average values. For instance, one study derived 117 features from three acceleration signals [72], while another extracted 58 brake-event-based features to classify driving behaviour [73]. Given the complexity of high-dimensional datasets, feature selection techniques play a crucial role in refining inputs for ML models. These techniques can be broadly categorised into statistical methods, model-based approaches, and deep learning-based strategies, as summarised in Table 2.4.

Table 2.4: Feature selection techniques used for driving heterogeneity identification

Category	Method	Prons (✓) & Cons (✗)	Reference
Statistical methods	FA	✓Computationally efficient	[45]
	DFT	✓Interpretable results	[31, 46, 74]
	DTW		[44, 74, 75]
	WT	✗Miss feature interactions	[31, 76]
	PCA	✗Questionable assumptions	[18, 52, 77]
Model-based methods	Tree-based	✓Capture feature interactions ✓Yeild better model performance	[78]
	GMM		[79]
	SFFS	✗Computationally intensive ✗Risk of overfitting	[72]
Deep learning-based methods	Autoencoder	✓Handle complex patterns ✓Good for high-dimensional data	[80]
	RNN	✗Computationally expensive ✗Challengable interpretability	[70, 81]

Abbreviations:

FA - Factor Analysis; GMM - Gaussian Mixture Method; SFFS - Sequential Forward Feature Selection - RNN - Recurrent Neural Network.

Statistical methods, including Principal Component Analysis (PCA), Discrete Fourier Transform (DFT) [82], and Dynamic Time Warping (DTW), are widely applied in driving behaviour analysis. These techniques enhance computational efficiency while providing interpretable results. PCA, for example, transforms data into a new coordinate system,

simplifying visualisation and feature ranking [18]. Similarly, DFT analyses signals in the frequency domain to minimise information loss [74]. Some studies integrate multiple techniques, such as combining Wavelet Transform (WT) with PCA [76] or DFT with Discrete Wavelet Transform (DWT) [74], to improve feature selection accuracy. However, these methods often assume linear relationships in data, making them less effective in capturing complex feature interactions. Model-based methods, such as decision trees (DT), random forests (RF), and Gaussian Mixture Models (GMM), offer an alternative by directly assessing feature importance within predictive ML models. These approaches can identify non-linear dependencies between variables, leading to better model performance. However, they can be computationally intensive and susceptible to overfitting, particularly in small datasets. For high-dimensional and complex datasets, deep learning-based methods provide advanced feature selection capabilities. Techniques such as autoencoders and Recurrent Neural Networks (RNNs) effectively capture intricate relationships between features and adapt to evolving driving patterns [80, 81]. While these methods improve feature representation, they are computationally demanding and often lack interpretability. In a word, selecting the appropriate feature selection technique requires balancing computational efficiency, interpretability, and accuracy based on the specific objectives of driving heterogeneity identification.

2.3.3 Step 3: Identification Methods of ML

The third step of the framework focuses on applying machine learning (ML) models to identify driving heterogeneity based on the traffic features selected in *Step 2*. In literature, ML techniques methods used for this purpose are generally classified into four categories: supervised learning (SL), unsupervised learning (USL), semi-supervised learning (SSL), and reinforcement learning (RL), which we elaborate further on below.

➤ Supervised learning methods

Supervised learning (SL) techniques train models to learn the relationship between input features and output labels, allowing them to make predictions or decisions on new, unlabelled data. This process requires training data to be labelled in advance, which is obtained from expert knowledge. Some studies use rule-based strategies to label data by measuring physical variable changes, such as vehicle steering angle or brake/accelerator pedal positions [40]. Threshold values for variables in rule-based labelling are typically determined by data analysts' prior knowledge and sometimes combined with other tools, such as driving style questionnaire (DSQ) [79]. K-means is also often employed for unsupervised learning labelling to group drivers and then label these clusters based on statistical analyses. For instance, Deng et al. [77] classified 30 participants into cautious, moderate, and aggressive drivers based on the PCA and K-means clustering before training models for driving style recognition. In these labelling processes, strict thresholds are usually set to distinguish different groups. A downside of this is that it can result in inaccurate classifications since there aren't always clear-cut boundaries between different driving profiles. To improve this, approaches such as fuzzy logic can be adopted in some studies [75]. Unlike the definitive nature of classical binary logic, fuzzy logic systems bridge inputs to outputs using a set of rules, which allows for establishments of conditions

like $if(A\&B) \Rightarrow C$, where A , B , and C represent different driving profiles. Such an approach avoids the constraints of rigid categories, offering a more accurate classification of driving heterogeneity.

With labelled data, supervised learning classifiers are trained to identify driving heterogeneity. Commonly used traditional ML algorithms include Support Vector Machine (SVM), Random Forest (RF), K-Nearest Neighbors (KNN), and Multi-Layer Perceptron (MLP), among others. Xue et al. [74] employed SVM for driving style recognition and revealed its advantage over RF, KNN, and MLP with an accuracy of 91.7%. More recently, deep learning methods, such as ANNs, CNNs, and RNNs, have demonstrated superior performance in handling complex and high-dimensional driving data. According to Xie et al. [88], the F1-score of CNN-based method is higher than both k-NN and RF-based methods in driving manoeuvre classification. Subsequently, enhanced CNN-based models, such as the adaptive regularised CNN (CNNAR) [51], the parallel Convolutional Neural Network (PCNN) [90], and the Residual Convolutional Network (RCN) [89] were introduced by boosting CNN's capability in driving pattern recognition, showing impressive performance by improving accuracy of 99.3% and reducing training time by two hours. Furthermore, Long Short-Term Memory (LSTM) networks are developed to overcome the vanishing gradient challenges of RNNs, exhibiting high accuracy in driving style classification, and unsafe driving behaviour detection, such as 91% as reported by [91] and 99% according to Saleh et al. [67], which outperform traditional ML models like MLP and DT [92]. Additionally, hybrid deep learning approaches have shown significant improvements in classification accuracy. For example, by taking CNN's ability to extract semantic driving patterns (e.g., turns) from input trajectories and employing an RNN on the sequential driving data to decode interrelationships among driving patterns, the proposed D-CRNN model presented improved accuracies of 8%-31% in driving style identification compared to CNN and RNN models [70].

Table 2.5 summarises SL applications in driving heterogeneity identification, detailing input data types, classifier choices, and output categories. Note that most SL studies focus on kinematic variables, with some integrating dynamic features to capture more driving behaviour features [86, 91]. The output of SL models varies depending on the classification task. Some studies categorise drivers into styles-based groups (e.g., aggressive and normal) or risk-based groups (e.g., high, moderate, and low risk), while others identify specific behaviours such as braking or acceleration patterns [88]. Additionally, some models focus on identifying drivers through driving states [89], while others detect unsafe behaviours like drowsy or distracted driving [87].

➤ Unsupervised learning methods

Unsupervised learning (UL) methods can derive driving profiles by directly examining unlabelled data, which is significantly less labour-intensive and reduces potential labelling biases. Table 2.6 summarises key studies that employ clustering techniques. Clustering methods such as K-means and Fuzzy C-means (FCM) group drivers based on similar driving styles, offering interpretable categorisations such as aggressive, normal, and defensive driving [18, 44]. More advanced clustering techniques like Federated Learning K-means (FL-K-means) and Gaussian Mixture Models (GMM) improve identification

Table 2.5: Identification of driving heterogeneity using supervised learning methods

Paper	Input data ^{*1}		ML models ^{*2}	Output
	Kine.	Dyna.		
[74]	✓		SVM , RF, KNN, MLP	Aggressive, and normal driving
[83]	✓		RF	High-, moderate-, and low-risk driving
[84]	✓	✓	AB , GB, RF, ET, SVM	Aggressive/risky, and normal driving
[72]	✓		KNN	Aggressive, and normal driving
[41]	✓		SVM , ANN, KNN, FL, KNN, RF	Aggressive, calm, and normal driving
[78]	✓		CART, SVM	High-, moderate-, and low-risk driving
[85]	✓	✓	SVM , ANFIS, LRNN, FFNN, LR	Typical and skillful driver
[86]	✓	✓	CNN , LSTM, pretrain-LSTM, SVM	High-, moderate-, and low-risk driving
[87]	✓	✓	2D CNN	Normal, aggressive, distracted, drowsy, and drunk driving
[88]	✓		CNN , KNN, RF	Driving manoeuvres: lane keeping, braking, turning, acceleration, right lane change, and left lane change
[89]	✓	✓	DeepRCN , DeepCNN, DT, RF, SVM, MLP	Driver recognition
[51]	✓		CNNAR , SVM, MLP, KNN, DT	Cautious, moderate, and aggressive driving
[90]	✓		PCNN , CNN, LSTM, HMM, SVM	Aggressive, and non-aggressive driving
[91]	✓		3-CCM-LSTM , 2-CMM-LSTM	Normal, drowsy, and aggressive driving
[67]	✓	✓	stacked-LSTM , MLP, DT	Normal, aggressive, and drowsy driving
[92]	✓		LSTM	Safe and unsafe driving
[70]	✓	✓	D-CRNN , CNN, RNN, ARNet, VRAE, GBDT	Driver recognition
[81]	✓	✓	LSTM-FCN , RF, LSTM, AB, ResNet	Aggressive and non-aggressive driving
[75]	✓		ANFIS	Safe, aggressive, and semi-aggressive driving
[93]	✓		HDC-FFNN , SNN, KNN, SVM, LSTM	Aggressive, and normal driving

*1. Table headings: Kine. - Vehicle kinematic variable; Dyna. - Vehicle dynamic variable;

Dis. - discrete, the output of ML model is distinct groups; Con. - continuous, the output of ML model is continuous trajectories.

*2. ML models: ET - Extra Trees; FL - Fuzzy Logic; CART - Classification and Regression Tree; ANFIS - Adaptive Neuro-fuzzy Inference Systems; 3-CCM - Three-class Classification Model; VRAE - Variational Recurrent Auto-Encoder; FFNN - Feed-forward Neural Networks; SNN - Spiking Neural Networks; LRNN - Layer Recurrent Neural Networks.

precision by handling scattered and sensitive data [45, 94]. Topic models, particularly Latent Dirichlet Allocation (LDA) and its variants (T-LDA, mLDA, mHLDA), segment naturalistic trajectories into behavioural “topics”, enabling fine-grained identification of cautious and radical driving styles [56, 95]. HMM-based approaches, such as sticky HDP-HMM and BP-AR-HMM, capture temporal dependencies by segmenting driving trajectories into hidden states, representing distinct driving states [96, 97]. Similarly, DAA-based methods such as Double Articulation Analyser with Temporal Prediction (DAA-TP) and Nested Pitman–Yor Language Model (NPYLM) can decompose driving sequences into hierarchical structures, mirroring linguistic pattern recognition for driver identification [98, 99]. Other studies employed deep learning models, such as Deep Sparse Autoencoder (DSAE) and Autoencoders with Self-Organized Maps (AESOM) to extract latent driving features, improving real-time recognition of high-risk and moderate-risk driving behaviours [59, 80].

Table 2.6: Identification of driving heterogeneity using unsupervised learning methods

Paper	Input data		ML models* ³	Output
	Kine.	Dyna.		
[60]	✓		K-means	Driving states
[18]	✓		Fuzzy C-means	Aggressive, normal, and mild driving
[44]	✓		SVC	Aggressive, normal, and defensive driving
[94]	✓		IFL K-means, FL K-means, FL-GMM, FFCM, MA K-means	Aggressive, moderate, and calm driving
[45]	✓		HC-GMM, GMM, DBSCAN, HC, K-means	Aggressive and normal driving
[100]	✓		LDA	Trajectory segmentation
[95]	✓		T-LDA, LDA, pLSA	Cautious, normal, radical, very radical driving
[56]	✓		mLDA, mHLDA	Extensive driving styles
[96]	✓		sticky HDP-HMM	Driving primitives
[97]	✓	✓	BP-AR-HMM	Driving states
[98]	✓		DAA-TP	Driving states
[99]	✓		NPYLM	Driver identification
[59]	✓	✓	DSAE	Driving states
[80]	✓		AESOM	Slight, moderate, high risk driving

*3. ML models:

IFL - Improved Federated Learning; FFCM - Federated Fuzzy C-means Method; pLSA - Probabilistic Latent Semantic Analysis; DSAE - Deep Sparse Autoencoder

Similar to SL methods, USL techniques primarily rely on kinematic vehicle data, such as speed, acceleration, and braking patterns, with some studies incorporating dynamic variables like yaw rate and lateral acceleration [59, 97]. Clustering and topic models work well with aggregated statistical features to distinguish different groups (e.g., aggressive, normal, risky), while HMM and deep learning-based models process sequential driving data, enabling better recognition of behavioural transitions [96].

➤ Semi-supervised learning methods

Semi-supervised learning (SSL) methods train classifiers with a small amount of labelled data and a large volume of unlabelled data, making it a promising approach for identifying driving heterogeneity with reduced labelling efforts. For example, semi-supervised SVM (S3VM) has shown better performance than traditional SVM in driving style classification [101]. Advanced SSL models, like the HDP-HSMM, outperform supervised counterparts such as HDP-HMM in extracting driving patterns from high-dimensional data [40]. Similarly, Tri-CatBoost integrates pseudo-labelling through tri-training, achieving higher accuracy than both supervised and unsupervised baselines [52]. Another approach, ARNet, leverages a limited number of labelled samples to guide a regularised autoencoder, improving accuracy by over 3% compared to traditional models. SSL methods are particularly useful when labelled data is scarce. Yet, their performance can be sensitive to noisy labels and data inconsistency. Moreover, their added complexity often increases computational requirements and implementation difficulty.

➤ Reinforcement learning methods

Reinforcement Learning (RL) approaches driving behaviour as a Markov Decision Process (MDP), where agents learn to perform driving actions based on a system of rewards and penalties. Driving heterogeneity is identified by analysing variations in learned behaviours and reward responses. For example, different reward functions can represent different levels of driving aggressiveness [102]. However, the lack of standard reward design and difficulty in validating learned behaviours are key challenges. Inverse Reinforcement Learning (IRL) instead learns the reward functions directly from observed driver behaviour. Techniques such as MLIRL and LogReg-IRL recover style-related parameters from human driving data [103, 104]. More advanced methods, like SIRL [105] and NFACRL [43], further enhance the robustness of driving style identification, particularly in dynamic or risky environments. While IRL offers deeper insights into individual driving preferences, its reliance on high-fidelity simulation environments, complex reward interpretation, and high computational cost limits its scalability in real-world applications.




Figure 2.5: Mechanism of different ML techniques

➤ Summary of ML methods

Each ML technique employs a distinct mechanism, as illustrated in Figure 2.5. From a data processing perspective, SL, USL, and SSL focus on data-driven analysis, including classification and clustering (see Figure 2.5(a-c)). In contrast, RL is primarily concerned

with understanding and modelling the decision-making processes of drivers, see Figure 2.5(d). Each ML technique has strengths and weaknesses and focuses on different aspects of driving heterogeneity, as summarised in Table 2.7. SL methods rely on labelled data and learn from historical labelled examples to classify driving behaviours into specific categories (e.g., aggressive vs. normal driving). They are predominantly used to characterise inter-driver heterogeneity, meaning variations among different drivers. They also frequently identify short-term and global heterogeneity, such as variations in acceleration or braking patterns across different traffic scenarios. USL methods identify patterns by clustering similar driving behaviours or segmenting trajectories based on statistical properties, which are particularly useful for identifying intra-driver heterogeneity (i.e., variations in an individual driver's behaviour over time). They also effectively capture short-term heterogeneity and global driving behaviour trends by analysing large-scale datasets without predefined categories. However, the lack of explicit labels makes their interpretability more challenging. SSL methods leverage both labelled and unlabelled data but can suffer from misclassification issues if the small labelled dataset does not represent the full heterogeneity of the driving population. RL-based models particularly emphasise short-term and global heterogeneity, making them valuable for adaptive driving systems and autonomous vehicle behaviour modelling.

Table 2.7: Comparison of ML Techniques in Driving Heterogeneity Identification

ML technique	Main heterogeneity focus	Strengths	Weaknesses
SL	Inter, short-term, global	High interpretability, effective classification	Requires labelled data, less adaptable
USL	Intra, short-term, global, and special	Captures hidden structures, scalable	Lower interpretability
SSL	Inter, short-term, and global	Reduces need for labels, improves generalisation	Susceptible to misclassification
RL	Inter, short-term, global	Adaptive, useful for AV behaviour modelling	High computational cost, complex reward tuning

2.3.4 Step 4: Performance Evaluation

The final step of the framework is performance evaluation which justifies a model's reliability, ensuring trustworthy and replicable outcomes. This is not only about model accuracy but also interpretability, generalisation, and online processing which displays the model's practical usability in real-world applications for ADAS, traffic management, and autonomous vehicle control.

➤ Model accuracy

Evaluating model accuracy is crucial to finding the best performance for identifying driving heterogeneity. Classification models usually rely on Accuracy, Precision, Recall, F1-score, AUC-ROC, and Cohen's Kappa to measure alignment with ground-truth labels [41, 74, 89, 106–108]. Various models such as SVM, RF, and KNN achieve an accuracy surpassing 85% when used to classify driving skills, some even exhibit better exceeding

95% [54, 84, 109]. Ranking-based metrics like MRR and CRR assess the prioritisation of driving behaviours, crucial for driver feedback systems. MAE and MSE are used for continuous outcome predictions, where MSE minimises large errors while MAE offers straightforward model performance. For unsupervised learning, clustering performance is measured by the Silhouette Score, Davies-Bouldin Index, and Calinski-Harabasz Index, ensuring well-separated clusters [110]. RL models use cumulative reward and convergence time to assess learning efficiency, essential for adaptive ADAS and traffic control.

Studies also employ benchmarking models from state-of-the-art to ensure performance validation, highlighting model strengths, and guiding the selection of optimal algorithms for real-world applications while enabling continuous refinement. This is usually done by using advanced deep learning models and traditional ML algorithms to conduct the same identification tasks, as seen for example with the proposed CNNAR model, which was compared with SVM, MLP, and KNN to demonstrate its superiority in identifying driving heterogeneity [70]. Other studies compare their improved ML models with corresponding foundational counterparts to illustrate enhancements [56, 81].

➤ Model implementation

Interpretability, generalisation, and online processing are crucial factors for real-world applications, influencing how well models can be understood, adapted to diverse driving environments and implemented in real-time driving scenarios. Interpretability ensures that the model's predictions can be understood and trusted by researchers, policymakers, and industry stakeholders. Generalisation determines whether an ML model can maintain its accuracy across different road conditions, weather patterns, and driver populations. And online Processing allows ML models to process incoming traffic data in real-time, making them applicable to adaptive traffic control, driver monitoring systems, and autonomous vehicle decision-making. Each ML approach has distinct strengths and limitations across these three dimensions, as illustrated in Figure 2.6, influencing its suitability for various driving heterogeneity identification tasks.

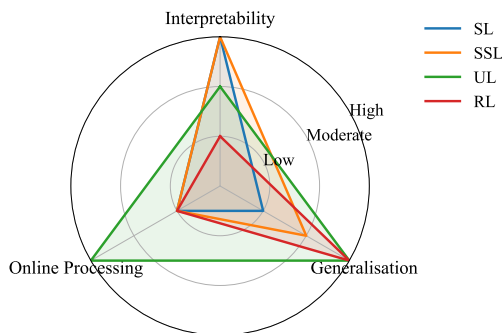


Figure 2.6: Analysis of ML techniques in model implementation

Interpretability: SL and SSL methods offer the highest interpretability, as they classify driving behaviours into well-defined categories, making them useful for safety assessments

and driver feedback. USL and RL methods have lower interpretability because of no label, but their results can be analysed to have human-comprehensive meanings. For instance, interpretable identification results can be provided by conducting statistical analysis on a limited number of clusters and giving semantic meanings [18, 40, 97].

Generalisation: SL struggles with new environments due to reliance on predefined labels, while SSL improves adaptability by leveraging unlabeled data. USL is highly generalisable, as it uncovers hidden driving patterns across diverse conditions. As presented by Ding et al. [9], the proposed GMM-based driving style identification method demonstrated effectiveness when supplemented with new driving data. RL can generalise well but depends on carefully tuned reward functions, which can limit transferability to new scenarios.

Online processing: USL is best suited for real-time applications, as clustering can adapt incrementally without retraining. SL and SSL require batch processing, limiting their real-time use. For instance, incorporating more SVMs can distinguish more driving styles, while this complicates the algorithm and can overburden the computational capacity of on-board vehicle controllers [111]. RL demands extensive training and updates, making it difficult for immediate deployment in adaptive traffic systems.

In total, effective ML models for driving heterogeneity must balance interpretability, generalisation, and online processing for better real-world applications in traffic management, driver profiling, and autonomous driving.

2.3.5 Summary

The proposed framework structures the ML-based identification process using four inherently connected steps. Each step includes multiple methods, offering flexibility in approach selection based on research objectives. By systematically comparing these techniques, the framework facilitates optimised identification, balancing model effectiveness with practical application needs.

2.4 Discussion

In this section, we discuss the main findings of the literature review and the proposed framework for driving heterogeneity analysis. Then we propose challenges and future recommendations towards implementing this framework for real-world applications.

2.4.1 Main findings

Driving heterogeneity identification involves three key aspects including applications, concepts, and methodologies. A clear objective of each aspect ensures reliable identification. The main insights from this study are summarised below.

➤ Driver heterogeneity dimensions

Driving heterogeneity is categorised into three core concepts, with five identification tasks. Inter- and intra- driving heterogeneity refers to differences among drivers and variations within individual drivers' behaviour. Understanding driving heterogeneity from this perspective aids in improving traffic efficiency and road safety.

Driving heterogeneity can be divided into **long-term** and **short-term** from a temporal perspective, which benefits the development of personalised ADAS. Additionally, **Global** and **special** driving heterogeneity are identified based on overall vehicle movements or specific driving events, which is important for AV design. The identification tasks include *driving style* classification, *risk level* evaluation, *driving skill* characterisation, *irregular behaviour* detection, and *driving pattern* recognition. Clearly defining these concepts and identification tasks helps comprehensively understand the essence of driving heterogeneity and enhances the interpretability of ML models.

➤ **ML-based heterogeneity identification**

The ML-based driving heterogeneity identification process can be structured as a four-step framework, consisting of *Trajectory Data Preparation*, *Traffic Feature Selection*, *Identification Models of ML*, and *Performance Evaluation*. Each step includes multiple methods with unique strengths and limitations. Facilitating driving heterogeneity identification requires aligning them with the specific research objectives. Proper data collection, preprocessing, and feature selection are necessary to match the heterogeneity concept being studied. Since ML techniques differ in accuracy, interpretability, and real-time applicability, balancing these factors ensures an effective identification process, improving its practical relevance.

➤ **Applications of the identification framework**

The proposed identification framework is adaptable to various datasets, integrates diverse ML techniques, and supports real-world applications. This flexibility allows for the incorporation of new models and data sources, enhancing identification performance. The framework emphasises interpretability and real-time recognition, encouraging researchers to evaluate methods across multiple dimensions. By guiding the implementation of ML-based driving heterogeneity identification, it supports traffic management, personalized ADAS, and human-like AV designs, ultimately improving transportation systems.

2.4.2 Key issues and future research recommendations

While ML-based methods show strong potential in driving heterogeneity identification, practical deployment, especially in real-time and automated driving systems requires improvements in both data and model design. This section highlights key challenges and outlines future directions to enhance the applicability of the proposed framework in real-world settings.

➤ **Enhancing data quality and availability**

Real-time and high-resolution data collection: ML models are highly sensitive to the quality and granularity of input data [55]. Traditional data sources, including traffic images and basic in-vehicle sensors, often suffer from noise and low sampling resolution. Industry-grade data collection systems, such as those used by Waymo, Mobileye, and Aurora, combine LiDAR, radar, and vision sensors with sub-millisecond-accurate logging

via CAN or Ethernet networks [112, 113]. These multi-sensor platforms enable large-scale collection of naturalistic driving data across diverse environments.

Advanced preprocessing and anomaly detection: Noise filtering and outlier elimination remain essential for improving data reliability. However, commonly used filters such as Savitzky-Golay or Butterworth may smooth out useful behavioural variations. More adaptive preprocessing methods, including machine learning-based anomaly detection and dynamic filtering strategies, are needed to better preserve meaningful heterogeneity in raw data.

Fusion of heterogeneous data sources: Combining data from simulators, traffic cameras, GPS, IMUs, and floating car data can yield a more holistic view of driving behaviour. However, successful integration depends on the development of unified data formats (e.g., ROS, JSON), synchronisation protocols (e.g., timestamp alignment), and cross-platform interoperability standards. Research should focus on frameworks that enable robust data fusion for training consistent ML models.

Balancing dataset use and development: Public datasets like NGSIM [114], HighD [115], KITTI [116], and Lyft5 [117] are widely used due to their accessibility, but they vary in precision and collection context, which may introduce bias or limit generalisability. Future work should aim to balance the reuse of open datasets with the creation of context-specific datasets that better represent traffic conditions.

Ethical data use and privacy protection: The use of personalised driving data raises concerns regarding privacy and data protection. Researchers must prioritise secure data handling practices, including anonymisation techniques (e.g., differential privacy), encryption, and compliance with international regulations such as the GDPR and CCPA [118]. Future work should explore privacy-preserving learning approaches that allow effective model training without compromising individual privacy.

➤ Improving model design and performance

Context-aware and flexible labelling: Traditional labelling strategies often rely on predefined thresholds or global statistics, potentially overlooking subtle, time-sensitive behavioural shifts [50]. Recent industry practices emphasise event-driven and context-aware labelling techniques to capture the complexity of real-world driving behaviours. For instance, AI-powered systems analyse driver gaze, head movements, and body posture to detect distraction and drowsiness, enabling more nuanced labelling of driver states [119]. Future research should continue to explore such adaptive labelling methods to enhance model accuracy.

Feature representation and selection: High-dimensional data with redundant variables can hinder model performance. Attention mechanisms, embedding layers, and spatio-temporal encoders should be employed to capture the most informative features. Domain knowledge can also assist in selecting features that reflect heterogeneity across time, context, and driver type.

Model interpretability vs. accuracy: While deep learning models such as CNNs and LSTMs often achieve high accuracy, they can lack transparency. Techniques like SHAP (SHapley Additive exPlanations) values, attention visualisation, and hybrid rule-based models can improve interpretability without significant performance loss [57]. Enhancing

explainability is essential for increasing trust in ML models, particularly in safety-critical applications.

Cross-disciplinary inspiration: Other domains, such as natural language processing (NLP) and biological modelling of gene expression, offer insights into pattern recognition and behavioural inference. Developing a taxonomy of driving behaviours, similar to biological annotation databases, could support better labelling and model validation. For instance, techniques from NLP sentiment analysis have been adapted to understand driver emotions and intentions, enriching behaviour modelling [99].

Adoption of Emerging AI Techniques: Recent advancements in AI models such as Vision Transformers (ViT), Graph Neural Networks (GNNs), and Large Language Models (LLMs) offer promising capabilities for driving behaviour analysis [120]. These models can capture long-range dependencies, learn contextual semantics, and reduce reliance on large labelled datasets through self-supervised or foundation model approaches. As these techniques gain traction in industry, they are increasingly used to detect nuanced and high-level patterns in driving heterogeneity. When integrated with the proposed framework, these models can significantly enhance the scalability, adaptability, and reliability of driving behaviour identification in complex, real-world traffic scenarios.

➤ Addressing heterogeneity in human-automated traffic

The emergence of automated vehicles in traffic introduces new challenges for driving heterogeneity analysis. Interactions between human-driven vehicles (HDVs) and AVs create more complex behavioural dynamics, requiring the framework to adapt accordingly [27].

Heterogeneity of HDVs in the presence of AVs: Empirical studies show that HDVs often adjust their driving styles, such as maintaining smoother speeds, longer headways, and more cautious manoeuvres, when following AVs [7, 121]. These behavioural shifts are further influenced by the external design of AVs (e.g., visibility of sensors) and the perceived assertiveness or caution in their driving style. Industrial companies such as Waymo and Cruise actively test human-AV interactions in mixed traffic using naturalistic testbeds and closed-loop simulations. Incorporating these behavioural adaptations as variables in heterogeneity models will improve their ability to reflect real-world traffic interactions.

Heterogeneity of human-driven AVs: Drivers of partially automated vehicles (e.g., SAE Level 2–3) exhibit diverse attitudes towards automation, with variation in trust, comfort, and takeover behaviour. Misalignment between AV system behaviour and driver expectations may result in frequent manual overrides, reduced system efficiency, and safety concerns [122, 123]. Industry solutions, such as Ford's and Tesla's driver monitoring systems, track gaze direction, head pose, and hand position to evaluate readiness for control transitions [124]. Future heterogeneity models should incorporate variables such as trust level, takeover frequency, and driver compliance to better model human-in-the-loop behaviours.

Overall, the proposed framework can be extended to address these evolving traffic conditions by integrating new feature variables (e.g., time headway to AVs, takeover intent), incorporating updated datasets from real-world or simulated mixed traffic scenarios, and refining model evaluation criteria to include human-system interaction dynamics. This adaptation will support the design of safer, more personalised, and behaviourally aligned

automated driving systems, thus ensuring better interaction between AVs and a diverse range of human drivers.

2.5 Conclusion

This study provides a comprehensive review of machine learning (ML) techniques for analysing driving behaviour heterogeneity and introduces a structured framework for identifying and interpreting heterogeneity in real-world traffic scenarios. By synthesising key concepts and state-of-the-art methodologies, the proposed framework serves as a systematic guide for data collection, preprocessing, feature selection, modelling, and evaluation, facilitating a more rigorous and interpretable approach to driving heterogeneity identification. Additionally, the review serves as a roadmap for future research, encouraging further exploration of ML applications in traffic analysis with the potential to enhance traffic management, road safety, and vehicle automation. Specifically, a clear conceptualisation of driving heterogeneity deepens our understanding of driver behaviour and lays the foundation for the development of personalised driving assistance systems and human-like autonomous vehicles. By reviewing various ML methodologies and assessing their strengths, limitations, and applicability to different driving contexts, this study emphasises the need to balance accuracy, interpretability, and real-time recognition for effective heterogeneity identification. Recognising individual driving characteristics allows for adaptive traffic control strategies, leading to more intelligent and responsive transportation solutions.

3

Investigation on Car-Following Heterogeneity and Its Impacts on Traffic Flow Performance

The content of this chapter has been published on

📖 Yao, X., Yan, Q., Sun, Z., Calvert, S. C., and Hoogendoorn, S. P. (2024). “Investigation on car-following heterogeneity and its impacts on traffic safety and sustainability” *Transportmetrica A: Transport Science*, 1-25.

This chapter proposes a general framework to investigate car-following heterogeneity and its impacts on traffic safety and sustainability. The framework incorporates rigorous driving style classification using a semi-supervised learning technique and a micro-simulation process that includes 66 fine-grained traffic scenarios exhibiting varying degrees of heterogeneity. Based on two distinct real-world datasets, the impacts of driving heterogeneity are effectively elucidated from the mechanism of underlying characteristics of driving behaviour and traffic flow dynamics.

3

3.1 Introduction

Driving heterogeneity constitutes a vital subject of investigation across various domains, including human-centred vehicle control systems, intelligent transportation systems, road safety, and environmental management. Studies have revealed that the variability in driving behaviours can lead to traffic externalities such as manifesting traffic hysteresis [125], causing more traffic accidents [126], fuel consumption and emissions [127–129]. For example, reaction time and sensitivity to stimuli are directly associated with rear-end collisions, thus contributing to traffic accidents [26, 130]. Also, extreme driving actions, such as over-speeding, excessive acceleration, and sudden stops, are more fuel-intensive [131], thereby increasing emissions and energy consumption [127, 132]. These findings underscore the necessity to investigate the impacts of driving heterogeneity on traffic flow performance and to understand them from the underlying mechanisms of driving behaviour, which can help develop measures to improve traffic safety and environmental sustainability.

According to Ossen et al. [47], driving heterogeneity is defined as the difference between driving behaviours of driver/vehicle combinations under comparable conditions. In literature, driving heterogeneity is usually identified from observed driving behaviours and formulated as a classification problem with the output being categorical (discrete scales vary between two and eight levels) or numerical (score of 0 to 10). Vehicular variables such as velocity, acceleration, and braking have been widely used to identify drivers' different driving styles [57, 133]. Calibration of car-following model parameters is another prevalent way to characterise different driving behaviours [1, 134]. Utilising these trajectory-related variables, computational models such as machine learning (ML) techniques are applied to characterise driving behaviour into distinct groups and infer specific driving profiles. Supervised and unsupervised learning are commonly used ML techniques for this purpose. For example, supervised learning classifiers such as Support Vector Machine (SVM), Random Forest (RF), K Nearest Neighbours (KNN), and Multilayer Perceptron (MLP) were employed to differentiate normal and aggressive driving styles, achieving accuracy as high as 91.7% [74]. A k-means clustering algorithm was implemented to recognise driving profiles into usual, harsh, and eco-driving based on speed and acceleration data [135]. Note that manually labelling of data for supervised learning is usually time-intensive and can introduce bias, and solely relying on unlabelled data in unsupervised learning may lead to unpredictable outcomes. Semi-supervised learning techniques offer a promising solution to overcome these challenges. They train classifiers to identify driving heterogeneity based on both labelled and unlabelled driving data, which can capture more characteristics of driving data

and uncover heterogeneity [136]. Wang et al. classified driving behaviour as aggressive and normal styles using both SVM and S3VM and revealed the superiority of the S3VM-based models over SVM-based models [101].

Based on the identification of driving heterogeneity, many efforts have been made to investigate the impact of driving styles on traffic performance by reproducing heterogeneous driving behaviours in microsimulation. Some studies developed stochastic car-following models by adding time-varying random noise (e.g., white noise) or distributions to deterministic models [137, 138]. For example, Zheng et al. developed a parsimonious enhanced Newell's car-following model incorporating the stochastic reaction time and the fluctuation around the vehicle's desired speed subject to the mean reversion process [139]. The Rakha-Pasumarthy-Adjerid (RPA) car-following model has shown its capability to generate realistic VSP distributions and estimate fuel consumption and emissions [133, 140]. Additionally, both variations in driving styles and vehicle characteristics are introduced in the Microsimulation Free-flow acceleration (MFC) model, which has been utilised to simulate heterogeneous driving behaviour and estimate fuel consumptions and emissions [128]. Based on developed microsimulation methods, the impacts of car-following heterogeneity on traffic flow performance have been revealed. A study reported that promoting more stable driving styles during car-following can potentially mitigate the risk of rear-end collisions [26]. On the contrary, aggressive driving styles are associated with higher levels of speed variability, elevated engine revolutions, and a greater likelihood of road accidents compared to other driving styles [131, 141]. Additionally, aggressive driving increased fuel costs significantly [142], by more than 20% [143] and by 25% in urban areas [144].

Despite the extensive research on heterogeneous driving behaviour modelling and the investigation of the impacts on traffic safety and sustainability, there are still some aspects that need to be explored. Usually, driving behaviours are diverse, and the corresponding driving heterogeneity identification and classified driving behaviour modelling require rigorous analyses, serving as an important preparation for traffic simulation. Furthermore, heterogeneous traffic flow needs more fine-grained traffic scenarios to demonstrate its diversity rather than a small number of representative fixed driving style ratios. Meanwhile, there is a need for nuanced investigation into how these influences happen from the mechanism of underlying characteristics of driving behaviour and traffic flow.

To bridge these research gaps, this chapter proposes a general micro-simulation approach to evaluate the impacts of car-following heterogeneity on traffic safety and sustainability. The novel contributions of this study are threefold: i) A semi-supervised learning method, i.e., multi-classification S3VM, is developed to facilitate rigorous driving style classification and classified car-following behaviour modelling. ii) Heterogeneous traffic flow is refined by 66 distinct traffic scenarios with varying degrees of heterogeneity, which allows a more nuanced examination of the impact of different driving styles and their proportion changes on traffic safety, fuel consumption and emissions. iii) The impacts caused by car-following heterogeneity are elucidated from the mechanism of underlying characteristics of driving behaviour and traffic flow dynamics.

3.2 Methodology

This section outlines the methodology for assessing the impacts of car-following heterogeneity on traffic performance using a micro-simulation approach. The process is illustrated in Figure 3.1. Data plays an important role in ML-based driving heterogeneity identification, thus being prepared as the first step of this methodology. Utilising extracted data of car-following pairs, a multi-class semi-supervised Support Vector Machine (S3VM) is developed to classify drivers into different driving styles. Based on the classification results, 66 refined traffic flow scenarios representing varying degrees of heterogeneity are established in micro-simulation, where various indicators are adopted to estimate traffic safety, and fuel consumption and emissions.

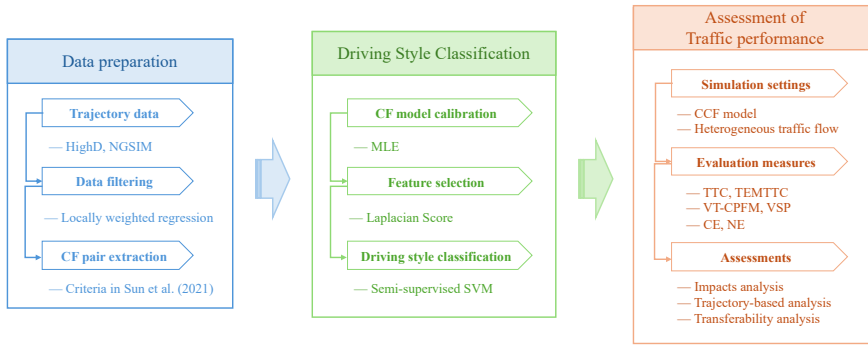


Figure 3.1: Overview of the proposed methodology

3.2.1 Data preparation

The HighD and NGSIM datasets are key and commonly used datasets for highway driving behaviour analysis and are utilised to evaluate the proposed methodology in this study. The HighD dataset captures vehicle trajectories at 25 fps (one frame per 0.04 seconds) through a high-resolution drone-mounted camera, which is notable for its ability to capture a wide range of behaviours and interactions among vehicles from a bird's-eye view [145]. Compared to the NGSIM dataset, HighD has higher granularity with longer recorded duration, driven distance and driven time. Therefore, we adopt HighD as the primary dataset for evaluation and NGSIM for transferability analysis of findings.

It is acknowledged the presence of errors and noise in original trajectory datasets could potentially impact the accuracy of microscopic studies [146]. Thus, we conduct data smoothing and filtering according to reference [18] before the extraction of car-following segments. This car-following pairs extraction process applies several criteria aimed at excluding congestion and free flow conditions, filtering minimum driving during, etc [18]. Finally, 2744 and 1097 car-following trajectory pairs are extracted from HighD and NGSIM datasets, respectively, with each lasting a minimum of 30 seconds.

3.2.2 Driving style classification

➤ Car-following model calibration

Car-following models capture the dynamics of longitudinal interactions between two adjacent vehicles navigating in the same lane without overtaking, aiming to simulate and understand driving behaviours in diverse traffic conditions. From a physics point of view, a car-following model used for microscopic traffic simulation should be as simple as possible [147]. According to Treiber et al. [147], several criteria should be checked when choosing a microscopic traffic model. First, the parameters of traffic models should be intuitive and easy to calibrate, and the corresponding values should be realistic, facilitating the replication of realistic driving behaviour in simulation. Additionally, the model should be capable of representing a variety of traffic conditions, including congestion and hysteresis effects, to ensure simulations accurately mirror real-world traffic dynamics. It's also crucial that the model prevents vehicle collisions and supports efficient numerical simulation.

The time-continuous Intelligent Driver Model (IDM) model is highlighted for its simplicity and effectiveness in simulating realistic acceleration profiles and behaviours in nearly all single-lane traffic scenarios [15], without leading to accidents. This physical model has interpretable parameters, promising advantages compared to machine learning-based models. To this end, IDM is widely used for traffic stability analysis and oscillation analysis in the literature and is capable of satisfactorily reproducing many characteristics of traffic flow [148–151]. The acceleration assumed in the IDM is a continuous function of the velocity v , the gap s , and the velocity difference Δv between the following vehicle and its leading vehicle, as illustrated in Equation 3.1.

$$a(t_{i+1}) = a_0 \left[1 - \left(\frac{v(t_i)}{v_0} \right)^\delta \right] - \left(\frac{s^*(v(t_i), \Delta v)}{s(t_i)} \right)^2 \quad (3.1)$$

$$s^*(v(t_i), \Delta v) = s_0 + \max \left(0, v(t_{i+1}) T + \frac{v(t_i) \Delta v}{2 \sqrt{a_0 b_0}} \right) \quad (3.2)$$

where a_0 is the maximum acceleration/deceleration of the follower, δ is the acceleration index, which is set the value of 4 to reduce additional safety issues [152], v_0 is the desired speed and s_0 is the minimum distance gap. $s^*(v(t_i), \Delta v)$ means the desired gap, which is a function of $v(t_i)$ and Δv as shown in Equation 3.2, in which T is the safety time gap and b_0 is the comfortable deceleration.

Prior research has identified some limitations of the IDM, such as its inability to reproduce high-dimensional oscillations [148]. Huang et al. proposed a stochastic IDM that can qualitatively replicate the concave oscillatory growth pattern [153]. The model adopts white noise to represent driver uncertainty, as shown in Equation 3.3. We utilise this stochastic IDM model to investigate the impacts of car-following heterogeneity on traffic flow performance.

$$v(t_{i+1}) = \max [\min (v(t_i) + a(t_{i+1}) \Delta t + \text{noise}, v_0), 0] \quad (3.3)$$

where $\text{noise} = \text{norm}(0, \sqrt{Q \Delta t})$, Q denotes the noise strength.

Model calibration was conducted at the individual vehicle level in this study, which enables the assignment of specific parameters to each driver indicative of his/her driving behaviour. The calibration of the car-following model involved identifying the most suitable parameter set to minimise the error between simulated and measured trajectories. Thereafter, traffic flow will be simulated using a mixture of drivers with their personalised parameters for car-following behaviour rather than using group average value. Recognising observed limitations of commonly optimised algorithms such as genetic algorithms and random parameter algorithms, Maximum Likelihood Estimation (MLE) is applied for model calibration in this study. Details about the MLE-based calibration procedure were described in [154].

➤ Feature selection

Most car-following models usually have more than five parameters. High-dimensional features easily lead to feature redundancy, and irrelevant features usually negatively affect training efficiency and accuracy. We adopt the Laplacian Score (LS) for dimensionality reduction due to its ability to identify features that represent the data's underlying structure. The idea of LS is to evaluate features according to their locality-preserving power, e.g., two data points are more likely the same type if they are close to each other. In classification problems, the local geometric structure of the data space is more important than the global structure. The Laplacian Score for each feature is calculated based on how much the feature values change within local neighbourhoods. Features that have small variations within neighbourhoods but significant variations between different neighbourhoods are considered more important. For more information about LS, please refer to He et al. [155]. Additionally, other commonly used unsupervised feature selection methods, i.e., Principal Component analysis (PCA) and t-distributed stochastic neighbourhood embedding (t-SNE) are employed for the same feature selection task to verify the effectiveness of LS.

➤ Semi-supervised driving style classification

Based on the selected features, we develop a semi-supervised approach, i.e., a multi-class semi-supervised support vector machine (S3VM), to classify driving style during car-following. S3VM is constructed using a mixture of labelled data (the training set) and unlabelled data (the working set), in which class labels in the training set are assigned to the working set to construct the “best” SVM [155]. The process of multi-class S3VM can be divided into multiple binary classification problems and solved using the “one-vs.-one” or “one-vs.-rest” strategy [156]. For the three-class problem in this paper, the “one-vs.-one” approach can use a smaller size of classifiers and less computational time to get similar results as the “one-vs.-rest” approach. Thus, the “one-vs.-one” approach is adopted in this study. The dataset \mathcal{S} of S3VM consists of labelled and unlabelled data points: $\{\mathcal{S}^{(l)}, \mathcal{S}^{(u)}\} = \{(x_1, y_1^d), \dots, (x_n, y_n^d), x_{n+1}, \dots, x_{n+m}\}$. where $\mathcal{S}^{(l)}$ and $\mathcal{S}^{(u)}$ denote the sets of labelled and unlabeled data, respectively. $y_i^d \in \{-1, 0, 1\}$ indicates whether $x_i \in \mathcal{S}^{(l)}$ belongs to the d -th class, where $d \in D$. D is the total number of classes. Here, the -1, 0, and 1 represent Aggressive, Normal, and Mild driving styles, respectively.

With labelled and unlabelled data, the goal of our classification task is to train a K binary-class classifier: $f^d(x|w^d, b^d) = \langle w^d, x \rangle + b$, where $w^d \in \mathcal{R}^n$ is the desired

hyperplane parameter vector for class d and b^d , is the bias term. To facilitate better performance, we apply a linear kernel and a nonlinear kernel function (i.e., Gaussian radial basis function kernel) separately for S3VMs. According to Gieseke et al. [157], for class d , the optimal solution for f^d can be found with parameter vector $\alpha^d \in \mathcal{R}^{n+m}$,

$$f^d(x|\alpha^d) = \sum_{i=1}^n \alpha_i^d k(x_i, x) \quad (3.4)$$

3.2.3 Assessment of traffic flow performance

In micro-simulation, several indicators are employed to estimate traffic safety, fuel consumption and emissions. Specifically, traffic safety is assessed by Time to Collision (TTC) and Time Exposed Modified Time to Collision (TEMTC). Fuel consumption is evaluated by the Virginia Tech Comprehensive Power-Based Fuel Consumption Model (VT-CPFM) and Vehicle Specific Power (VSP), and emissions of heterogeneous traffic flow are evaluated according to Panis et al. [158].

➤ Measures of traffic safety

Time-to-collision (TTC) is defined as the time when the speed of the objective vehicle is greater than its leading vehicle, the objective vehicle keeps the original driving state and does not take the corresponding deceleration behaviour until the two vehicles collide. It has been widely used in traffic safety evaluations. The formula is shown as follows:

$$TTC_i(t) = \begin{cases} \frac{x_{i-1}(t) - x_i(t) - l_{i-1}}{v_i(t) - v_{i-1}(t)}, & \text{if } \forall v_i(t) > v_{i-1}(t) \\ \infty, & \text{if } \forall v_i(t) \leq v_{i-1}(t) \end{cases} \quad (3.5)$$

where $x_i(t)$ and $v_i(t)$ are the location and speed of vehicle i at time t , respectively. $x_{i-1}(t)$ and $v_{i-1}(t)$ denote the location and speed for vehicle $i-1$; l_{i-1} represents the length of vehicle $i-1$.

A threshold value is selected to distinguish between safe and dangerous traffic conditions in the literature. According to Minderhoud et al. [159], the time-exposed time-to-collision (TET) was proposed based on the TTC, which can be determined by Equations 3.6.

$$TET^* = \sum_{i=2}^N \sum_{t=0}^T \delta_i(t) \cdot \Delta t, \quad \delta_i(t) = \begin{cases} 0, & \text{Otherwise} \\ 1, & 0 \leq TTC_i(t) \leq TTC^* \end{cases} \quad (3.6)$$

where T is the total observation time (denotes simulation time here) and Δt is the interval time. TTC^* represents the safety threshold whose value varies from 1 to 3s [160]; $\delta_i(t)$ is a 0-1 variable. When $TTC_i(t)$ is less than TTC^* , $\delta_i(t)$ is equal to 1; otherwise, it is 0. For $N(i=2 \dots N)$ drivers in the observation section, the total TET^* is calculated by Equation 3.6. The smaller the value of the TET^* , the higher level the of traffic safety. TTC^* is set as 2s in this paper.

The Modified time to collision (MTTC) is another metric that calculates the time required for the following vehicle to collide with a leading vehicle maintaining constant movement

characteristics. Generally, an MTTC below 1.5s is deemed unsafe. TEMTTC represents the accumulated time of unsafe MTTC experienced by each vehicle in the traffic flow [161]. The calculation is expressed as follows:

$$MTTC_i(t) = \frac{\Delta v_i(t) \pm \sqrt{\Delta v_i(t)^2 + 2\Delta a_i(t)[x_{i-1}(t) - x_i(t) - l_{i-1}]}{\Delta a_i(t)} \quad (3.7)$$

$$TEMTTC = \sum_{i=2}^N \sum_{t=0}^T \zeta_i(t) \cdot \Delta t, \quad \zeta_i(t) = \begin{cases} 0, & \text{Otherwise} \\ 1, & MTTC_i(t) < 1.5s \end{cases} \quad (3.8)$$

where $MTTC_i(t)$ is determined by (i) if both of MTTC terms are positive, the minimum of them is considered the final value of MTTC; and (ii) if one is positive while the other is negative, the positive one is considered the final value of MTTC; the Boolean variable $\zeta_i(t)$ takes a value of 1 if the condition is unsafe and 0 otherwise.

➤ Measures of traffic sustainability

The Virginia Tech Comprehensive Power-based Fuel Consumption Model (VT-CPFM) is a well-regarded model designed to estimate the fuel consumption of vehicles based on powertrain dynamics and vehicle-specific characteristics. It can be easily implemented in systems that require the use of a microscopic-level fuel consumption model [162], thus being employed to calculate fuel consumption in this study. The fuel consumption rate $F_{C,i}$ (g/s) is modelled in the form of a quadratic polynomial as:

$$F_{C,i} = \begin{cases} \theta_{0,i} + \theta_{1,i}P_{d,i} + \theta_{2,i}P_{d,i}^2, & \text{if } P_{d,i} \geq 0, \\ \theta_{0,i}, & \text{if } P_{d,i} < 0, \end{cases} \quad (3.9)$$

where $\theta_{(\cdot),i}$ are constant coefficients, and $\theta_{0,i} = 0.54$, $\theta_{1,i} = 0.06$, $\theta_{2,i} = 0.00017$; $P_{d,i}$ (kW) is the power output of the vehicle driveline and it calculated based on the vehicle velocity and acceleration:

$$P_{d,i} = \frac{(ma_i(t) + C_{A,i}v_i(t)^2 + mgf_{r,i})v_i(t)}{\eta_{T,i}} \quad (3.10)$$

where m is the vehicle mass (1500kg); $C_{A,i}$ is the coefficients of aerodynamic drag (0.4kg/m); g is the gravitational acceleration (9.8m/s²); $f_{r,i}$ is the rolling resistance (0.015kg/m); $\eta_{T,i}$ is the mechanical efficiency of the driveline (0.8).

Vehicle specific power model (VSP) refers to the instantaneous power of a vehicle per unit mass and combines the driving characteristics of vehicles such as speed and acceleration, and road characteristics such as the road gradient [163]. It has been widely used for fuel consumption modelling, because both the power for overcoming aerodynamic drag and rolling resistance and for the kinetic and potential energy of the vehicle are taken into account in VSP, by doing this the relationship between VSP and fuel consumption can be explained physically. The calculation of the VSP is shown in Equation 3.11.

$$VSP = 0.132 \cdot v + 1.1 \cdot v \cdot a + 0.0003202 \cdot v^3 \quad (3.11)$$

Vehicle Specific Power (VSP) is usually clustered into bins at certain intervals, and traffic emissions are estimated by average-speed-based VSP distribution within each bin [164]. The categorisation of VSP into 1KW/t intervals is detailed in Equation 3.12.

$$VSP \text{ bin} = n, \forall : VSP \in [n - 0.5, n + 0.5) \quad (3.12)$$

where n is the VSP number.

Then, a nonlinear multivariate regression model is utilised to model instantaneous traffic emission by considering both average speed and other aspects of vehicle operation such as acceleration and deceleration [158], which is expressed by:

$$E_i = \frac{\sum_{t=1}^T \max \left\{ 0, f_1 + f_2 v_i(t) + f_3 v_i(t)^2 + f_4 a_i(t) + f_5 a_i(t)^2 + f_6 v_i(t) a_i(t) \right\}}{T} \quad (3.13)$$

$$E = \frac{\sum_{i=1}^N E_i}{N} \quad (3.14)$$

where E_i is the traffic emission (g/s) of vehicle i ; E is the average traffic emission of the entire traffic flow; the calibrated values of model parameters $f_1, f_2, f_3, f_4, f_5, f_6$ for petrol cars, as outlined in [158], are employed in this study. Specifically, for CO₂ emissions $f_1 = 5.54e^{-1}, f_2 = 1.61e^{-1}, f_3 = -2.89e^{-3}, f_4 = 2.66e^{-1}, f_5 = 5.11e^{-1}, f_6 = 1.83e^{-1}$; for NO_x emissions, model parameters associated with vehicle acceleration. when $a_i(t) \geq -0.5m/s^2$, $f_1 = 6.19e^{-4}, f_2 = 8.00e^{-5}, f_3 = -4.03e^{-6}, f_4 = -4.13e^{-4}, f_5 = 3.80e^{-4}, f_6 = 1.77e^{-4}$; otherwise, $f_1 = 2.17e^{-4}, f_2 = f_3 = f_4 = f_5 = f_6 = 0$.

3.3 Experiments

In this section, classified-car-following models are first established according to driving style classification. Then the micro-simulation setup for estimating heterogeneous traffic performance is introduced. A preliminary experiment is conducted to justify the CCF models' capability to reproduce spatiotemporal traffic flow patterns before formal simulations start.

3.3.1 Classified car-following models establishment

Table 3.1 gives an overview of the bounds of the stochastic IDM model calibration. Following the methodology introduced in Section 3.2, the calibrated behavioural parameters are prepared for the subsequent driving style identification.

Feature selection is then conducted based on the calibrated car-following model parameters and the results of the score array are $[a_0 : -0.0831, b_0 : -0.0638, s_0 : -0.0573, T : -0.0768, v_0 : -0.1237]$. Considering that features with lower Laplacian scores are more important, v_0 emerged as the most important feature, followed by a_0 and others, indicating the greater importance of v_0 and a_0 in describing driving behaviour. Thus, we manually label driving styles for semi-supervised classification according to v_0 and a_0 . Usually, drivers exhibiting higher desired speeds and greater acceleration are considered as a more aggressive driving style [101]. This empirical knowledge enables the identification

Table 3.1: Summary of IDM model parameters and their estimates.

Para. (unit)	Short description	Bounds	Mean	Median	Std
$a_0(m/s^2)$	Max desired acceleration of follower	[0.1,5]	1.18	1.04	0.81
$v_0(m/s)$	Desired speed of follower	[10,40]	29.76	30.04	4.99
$s_0(m)$	Gap at standstill	[0.1,6]	2.22	2.17	0.97
$T(s)$	Desired time headway of follower	[0.1,5]	1.72	1.67	0.94
$b_0(m/s^2)$	Comfortable deceleration of follower	[0.1,5]	1.45	1.33	0.92
Q	Noise strength	[0.01,1]	0.50	0.51	0.28

of drivers with distinct characteristics, forming the foundation for our redefining three driving styles: Aggressive, Normal, and Mild. Specifically, drivers with high designed speeds and accelerations (e.g., above the 75th percentile) were categorised as an Aggressive driving style. Conversely, drivers with lower values in these parameters (e.g., below the 25th percentile) were classified as a Mild driving style. Those with median values are labelled as a Normal driving style. Based on these criteria, a total of 295 driving samples were pre-labelled, which comprise 99, 99, and 97 samples for Aggressive, Normal, and Mild styles, respectively.

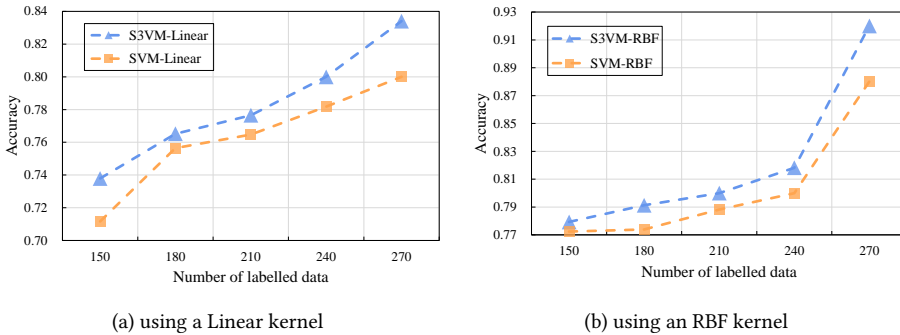


Figure 3.2: Classification accuracy of S3VM and SVM for multi-classification

A linear and a Gaussian radial basis function (RBF) kernel and obtain satisfactory results and reduce computational complexity when using SVMs. All the calibrated model parameters were divided into three disjoint parts: labelled dataset $S^{(l)} = \{x_i\}_{i=1}^n$, unlabeled dataset $S^{(u)} = \{x_j\}_{j=1}^m$, and test dataset $S^{(t)} = \{x_i\}_{i=1}^r$, where $n = 295$, $m = 2449$, $r \leq n$. Insufficient labelled data for training may not take full advantage of unlabeled data to properly reflect the reality of various driving styles. Hence, to further demonstrate and evaluate the S3VM's capability of exploiting the available data, comparative experiments were conducted using different amounts of labelled datasets $S_p^{(l)}$, where $P \in \{150, 180, 210, 240, 270\}$. Classification results of SVM and S3VM are illustrated in Figure 3.2 where the vertical axis represents the multi-classification accuracy and the horizontal axis denotes the varied amount of labelled training data. In both Figure 3.2a and Figure




Figure 3.3: Comparison of fitted distributions of CF model parameters.

3.2b, the blue dotted lines with triangles consistently exceed the yellow dotted lines with squares, indicating S3VM generally outperforms SVM across both kernel types. The overall advantage of S3VM ranges from 1.74%-8.00% when employing a linear kernel, and from 1.82% to 8.60% with an RBF kernel. For instance, with 210 labelled data points, SVM achieved accuracies of 76.47% with a linear kernel and 78.82% with an RBF kernel. In contrast, S3VM attained accuracies of 77.65% and 80% with the respective kernels. Notably, the advantage of S3VM is more pronounced with a larger dataset. For example, with 150 labelled samples using an RBF kernel, the accuracy improvement of S3VM over SVM increases from 77.24% to 77.93%. This advantage becomes more significant with 270 labelled samples, which is with S3VM achieving 92% accuracy compared to 88% for SVM. Classification results with the highest accuracy are used for subsequent analyses.

The Classified Car-following (CCF) model has been proposed to personalise car-following models for heterogeneous drivers [18]. We follow their methodology to establish CCF models for drivers in our micro-simulation. Four widely used parametric distributions, namely Normal, Lognormal, Gamma, and Logistic, are employed to determine the probabilistic distributions of car-following model parameters for each driving style. The best-fitting distribution was chosen by comparing the goodness of fit of all parametric distributions. As shown in Figure 3.3, the Lognormal distribution outperforms the other distributions. Thus, the mean value (μ) of the Lognormal distribution is utilised to determine IDM model parameters for each driving style, as outlined in Table 3.2.

Table 3.2: Classified IDM model parameters.

Driving styles	a_0	v_0	s_0	T	b_0	Q
Aggressive	1.93	33.55	1.91	1.35	1.14	0.373
Normal	1.04	29.45	2.02	1.48	1.04	0.370
Mild	0.44	27.21	2.00	1.36	1.08	0.376


3.3.2 Micro-simulation setup

The micro-simulation is developed by MATLAB R2021b, which was executed on an Apple M1 Pro MacBook. The traffic volume is set as $1600veh/h$, and the vehicles are generated based on the negative exponential distribution $e^{-\lambda t}$. 30 vehicles were counted on the simulated road section.

The proportions of drivers with Aggressive, Mild, and Normal driving styles in the simulated traffic flow are denoted by p_a , p_m , and $p_n = 1 - p_a - p_m$, respectively. To thoroughly assess the effects of various driving styles and their proportional changes on traffic safety and sustainability, 66 fine-grained traffic scenarios are considered. The shares of each driving style (p_a , p_m and p_n) range from 0% to 100% in 10% increments. Acknowledging that drivers with identical driving styles may still exhibit individual differences, we introduce variability in the Car-Following (CF) model parameters in alignment with the values given in Table 3.2. To avoid extreme values resulting from generating parameters directly from distributions, we use Equation 3.15 to vary car-following parameters. By doing this, each driver has their personalised IDM model parameters meanwhile following a particular driving style. It is assumed that drivers maintain a consistent driving style throughout the simulation period [165]. Each traffic scenario is simulated with 100 seeds to enhance the reliability of the results. The final result for each traffic flow scenario is determined by calculating the mean of all repeated simulations under this setting.

$$PM_{ij}^{sim} = 0.9PM_{ij} + 0.1 \times \frac{1}{PM_{ij}\sigma_{ij} \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{\ln(PM_{ij}) - \mu_{ij}}{\sigma_{ij}} \right)^2} \quad (3.15)$$

where PM_{ij}^{sim} denotes the CCF model parameter used in simulation, $i = \{a_0, v_0, s_0, T, b_0, Q\}$, $j = \{Aggressive, Normal, Mild\}$. PM_{ij} follows the values shown in Table 3.2.



(a) Under lower speed driving conditions (10m/s). (b) Under higher speed driving conditions (30m/s).

Figure 3.4: Spatiotemporal patterns of traffic flow under different speed settings.

3.3.3 Preliminary simulation

Preliminary experiments were conducted to examine the traffic flow patterns generated by the CCF models. The spatiotemporal patterns of traffic flow with 20% aggressive and 80% normal drivers are shown in Figure 3.4. Stripes structures are observed, illustrating the propagation, growth, dissipation, and merging of disturbances. In high-speed driving conditions (30m/s) such as the HighD dataset, see Figure 3.4b, fewer disturbances are observed compared to low-speed driving conditions such as the NGSIM dataset (see Figure

3.4a). This is due to the diminished influence of white noise on oscillation evolution on higher-speed traffic flow [148]. These results demonstrate that the stochastic car-following model can effectively replicate typical spatiotemporal traffic flow patterns.

3.4 Results and discussion

Based on the identification of car-following heterogeneity and experimental settings, micro-simulation results are presented in this section. The impacts of variations in driving styles on traffic safety, fuel consumption and emissions are analysed from the mechanism of underlying driving behaviours in heterogeneous traffic flow. And the proposed methodology is then verified by transferability analysis using a different dataset.

3.4.1 Impacts of CF heterogeneity

Statistics of the aforementioned traffic estimation indicators for 66 simulated traffic scenarios are presented in Figure 3.5-3.6. Indicators representing traffic safety, fuel consumption, and emissions are represented by blue, green and red colour sets, respectively. Within each colour set, a darker hue signifies a lower indicator value, such as fewer hazardous incidents or reduced emissions. Figures 3.5a and 3.5b present the results of TTC and TEMTTC, the scenario located in the upper right corner displays the lightest shade of blue, indicating that a traffic flow with 100% aggressive drivers displays a very low level of traffic safety among all heterogeneous traffic scenarios. The shift from light blue towards dark blue is observed from the top right to the bottom left, indicating that a decrease in the number of Aggressive drivers and an increase in the number of Mild drivers improves traffic safety. The safest traffic conditions are achieved with a composition of 100% Mild drivers. Results of fuel consumption and emissions across all simulated traffic scenarios are presented in Figures 3.6a-3.6d. The colour gradient from the darkest in the bottom left corner to the lightest in the top right corner signifies a reduction in traffic sustainability. These results indicate that an increase in the share of Aggressive drivers and a decrease in Mild drivers in traffic flow leads to higher fuel consumption and emissions.

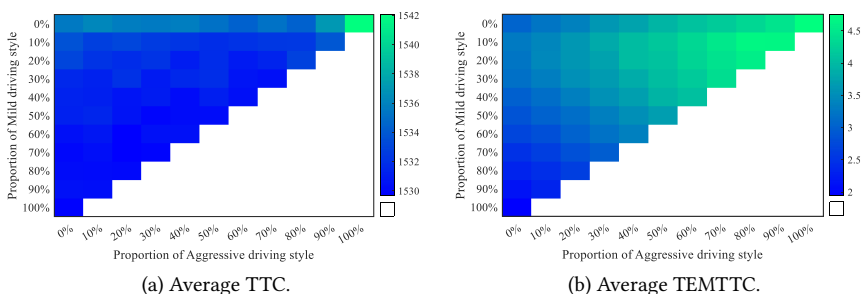


Figure 3.5: Statistics of traffic safety indicators in 66 traffic scenarios.

Overall, this analysis indicates a general trend that traffic safety and sustainability

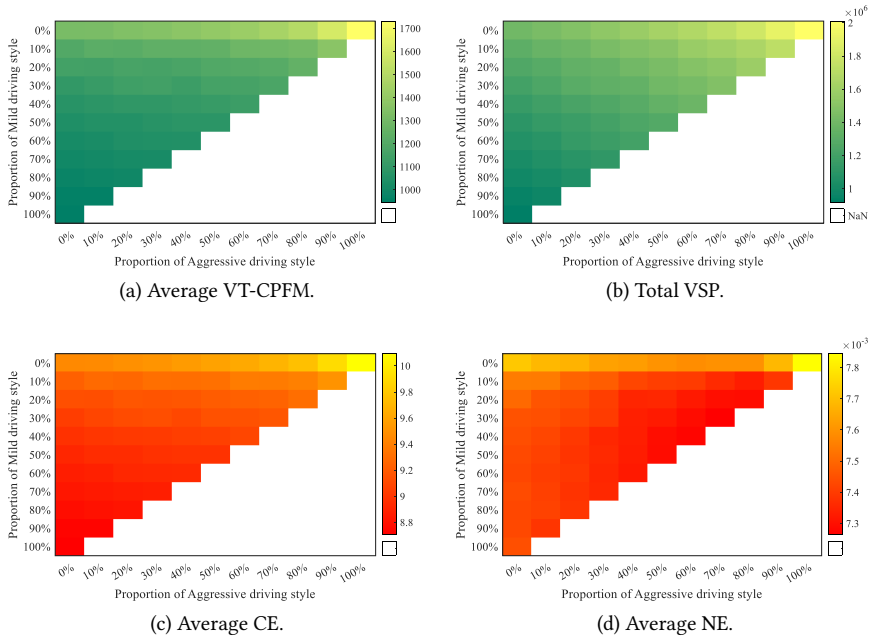


Figure 3.6: Statistics of traffic sustainability indicators in 66 traffic scenarios.

improve as the proportion of milder drivers rises and more aggressive drivers decrease, while the trend does not follow a strictly linear pattern. Take Figure 3.5b as an example, the traffic scenario with 80% Aggressive drivers and 20% Normal drivers is associated with fewer safety issues than one with 80% Aggressive drivers, 10% Normal drivers, and 10% Mild drivers. This observation suggests the complexity of traffic dynamics and a need for a nuanced understanding of how various driving styles and their mixed proportions affect traffic safety and sustainability.

3.4.2 Trajectory-based analysis

To further understand the diversity of safety and sustainability levels caused by CF heterogeneity, we examined individual driving trajectories by diving into two-mixed (with two different driving styles) and three-mixed (with three different driving styles) heterogeneous traffic flows.

➤ Analysis on two-mixed heterogeneous flow

Figure 3.7-3.13 shows the vehicle trajectories when different proportions of Normal drivers are introduced into an Aggressive style traffic flow. Trajectories of Normal and Aggressive drivers are represented by solid red lines and blue dotted lines, respectively. In Figure 3.7a, despite the stochastic nature of the vehicles, they maintain stable and uniform spacing due to their consistent driving style within homogeneous traffic flow, leading to a smoother

traffic flow pattern with little disruption. Similar observations are demonstrated in speed and acceleration diagrams shown in Figure 3.7b and Figure 3.7c. The introduction of Normal drivers disrupts this uniformity and leads to the formation of platoons, and these platoons are formed in traffic flow with Normal-style vehicles as the leader. In a traffic flow with 80% aggressive drivers shown in Figure 3.8a, several platoons are observed which are led by the Normal vehicles, as red lines show. Figure 3.8b shows speed diagrams where light red and light blue represent the vehicle platoons of Normal and Aggressive, respectively. The dark blue line represents the lead vehicle in the Aggressive platoon, which is the first follower of a Normal vehicle. Notice that the speed of the Aggressive follower decreases over time, eventually aligning with the speed of the Normal platoon. This occurs because the Normal leader inhibits the Aggressive follower from maintaining a higher speed. Consequently, other Aggressive vehicles in the platoon, led by this impeded Aggressive vehicle, also reduce their speeds, resulting in the entire platoon adopting a Normal driving style. A similar trend is evident in Figure 3.8c where all platoons ultimately exhibit similar acceleration, despite Aggressive vehicles displaying significant fluctuations in acceleration. This variability of accelerations potentially leads to increased fuel emissions and higher traffic risks, elucidating the lower safety and sustainability caused by aggressive drivers.




Figure 3.7: Vehicle trajectories of traffic flow with 100% Aggressive drivers.




Figure 3.8: Vehicle trajectories of traffic flow with 80% Aggressive drivers and 20% Normal drivers.




Figure 3.9: Quantitative analysis of vehicle trajectories (80% Aggressive and 20% Normal drivers).

For more detailed analysis, trajectories of four exemplified vehicles in a traffic scenario with 80% Aggressive drivers and 20% Normal drivers are visualised in Figure 3.9. The red line with circles denotes a Normal vehicle, and lines with other colours and signs represent Aggressive vehicles. Initially, three vehicles with Aggressive styles exhibited maximum accelerations of 1.92 m/s^2 , 1.81 m/s^2 , and 1.85 m/s^2 , respectively, as the cyan line with asterisks (Aggressive-1), the blue line with plus (Aggressive-2), and the orange line with squares (Aggressive-3), as depicted in Figure 3.9c. The Normal vehicle, denoted by the red line with circles, has the maximum acceleration of 0.98 m/s^2 . Based on originally preset accelerations, speeds of all vehicles increase as the simulation progresses, see Figure 3.9b. During the initial 0-10s, the speeds of Aggressive vehicles surpass that of the Normal vehicle. After 10s, the speeds of Aggressive-1 continue to increase, stabilizing at 33.6 m/s , whereas Aggressive-2 and Aggressive-3 flatten out and show a similar increase trend to that of Normal vehicle, reaching speeds of 27.6 m/s and 27.1 m/s , respectively. To this end, Aggressive-1 can behave as a real Aggressive driving style and distance itself from the following Normal vehicle. In contrast, Aggressive-2 and Aggressive-3 exhibit a Normal driving style due to the blocking of their Normal leader. Here, vehicles that are original with an Aggressive and Normal style are denoted as A_A and N_N, respectively, while Aggressive vehicles that exhibit as Normal due to the block of their leaders are represented as A_N.




Figure 3.10: Vehicle trajectories of traffic flow with 60% Aggressive drivers and 40% Normal drivers.




Figure 3.11: Vehicle trajectories of traffic flow with 40% Aggressive drivers and 60% Normal drivers.




Figure 3.12: Vehicle trajectories of traffic flow with 20% Aggressive drivers and 80% Normal drivers.

In traffic scenarios with a decreasing proportion of Aggressive drivers and an increasing share of Normal drivers, the platoon formation dynamically shifts. In Figure 3.11a-3.12a where there are more Normal drivers than Aggressive drivers, these Aggressive vehicles appear at the end of each platoon. This occurs as Normal vehicles, with their lower acceleration and speed, naturally fall behind their Aggressive leaders and eventually form platoons leading by themselves. Since these traffic flows have a high proportion of Normal drivers, most vehicles in platoons are with an N_N style rather than A_N. Moreover, a homogeneous traffic flow with 100% Normal drivers demonstrates a consistent and smooth pattern, mirroring the uniform behaviour seen in a pure Aggressive traffic flow, see Figure 3.13.

Similar findings are observed through analyses of two-mixed traffic flow scenarios comprising Mild and Normal driving styles. Figure 3.14-3.18 illustrates vehicle trajectories in heterogeneous traffic flow with different proportions of Mild drivers. When Normal drivers are involved in Mild traffic flow, vehicle trajectories are no longer neatly aligned and some platoons are formed, see Figure 3.14. These platoons are formed in traffic flow with Mild-style vehicles as the leader. For example, as the zooming-in trajectory diagram shows in Figure 3.14a, several platoons are led by Mild vehicles, see the green dashed lines. This is because Mild vehicles have smaller acceleration and speed compared to Normal vehicles as settings, which makes a follower Mild vehicle fall behind its Normal leaders.




Figure 3.13: Vehicle trajectories of traffic flow with 100% Normal drivers.

Meanwhile, these Mild vehicles block their Normal followers and force them to behave in a Mild driving style with smaller acceleration and speed as well, see Figure 3.14b and Figure 3.14c. Specifically, the Mild leader notably slows down the Normal followers, eventually causing the Normal vehicles to adopt a mild driving style, as seen in the red region of Figure 3.14b. The overall acceleration of Normal vehicles aligns with a mild driving style, whereas it still exhibits larger fluctuations compared to those in mild platoons, as shown in Figure 3.14c.




Figure 3.14: Vehicle trajectories of traffic flow with 20% Mild drivers and 80% Normal drivers.

When the proportions of Normal drivers decrease and Mild drivers increase, as shown in Figure 3.15-3.17, the formation of platoons in traffic flow changes. In Figure 3.17a where there are more Mild drivers than Normal drivers, these Normal vehicles appear at the end of each platoon. This is because a Mild vehicle which has small acceleration and speed falls behind its Normal leader, forming a platoon leading by itself. Since these traffic flows have a high proportion of Mild drivers, most vehicles in platoons are with an M_M style rather than N_M. Moreover, vehicle trajectories in a homogeneous traffic flow with 100% Mild drivers show a consistent and smooth trend, which is similar to a pure Normal traffic flow, see Figure 3.18.

In summary, in a two-mixed traffic flow scenario, vehicles with lower aggressiveness (Mild in comparison to Normal, or Normal in contrast to Aggressive) tend to lead the




Figure 3.15: Vehicle trajectories of traffic flow with 40% Mild drivers and 60% Normal drivers.




Figure 3.16: Vehicle trajectories of traffic flow with 60% Mild drivers and 40% Normal drivers.




Figure 3.17: Vehicle trajectories of traffic flow with 80% Mild drivers and 20% Normal drivers.

formation of vehicular platoons. This occurs as they naturally fall behind their more aggressive leaders. Meanwhile, vehicles with higher aggressiveness are observed to adopt a driving style associated with lower aggressiveness due to being impeded by their less aggressive leaders. Importantly, the formation of these vehicular platoons is influenced not just by the proportion of less aggressive vehicles within the traffic, but also by their spatial positioning within the traffic flow. These findings explain observations that the negative effects of an aggressive driving style – or conversely, the beneficial impacts of a




Figure 3.18: Vehicle trajectories of traffic flow with 100% Mild drivers.

mild driving style – on traffic safety and sustainability do not escalate strictly linearly with their proportion changes.




Figure 3.19: Traffic flow with 30% Aggressive, 30% Mild drivers, and 40% Normal drivers

➤ **Analysis on three-style mixed traffic flow**

In traffic flows comprising three-mixed driving styles, the dynamics of driving behaviour become increasingly complex. Figure 3.19 shows a heterogeneous traffic flow with 30% Aggressive, 30% Mild drivers, and 40% Normal drivers, in which Normal drivers are represented by red solid lines, and blue and green dash lines denote Aggressive drivers and Mild drivers, respectively. The Normal vehicle falls behind its Aggressive leader and creates a large spacing with its Mild follower. The three vehicles shown as blue, red, and green lines in Figure 3.19b exhibit their original driving styles (Aggressive, Normal, and Mild), largely because they are not impeded by vehicles of lesser aggressiveness ahead of them, enabling them to achieve the desired speed typical of their respective driving styles. This is similarly reflected in their acceleration profiles, as seen in Figure 3.19c, where fluctuations remain minimal due to the lack of hindrance. Beyond these three initial vehicles, drivers with less aggressive styles tend to obstruct those who are more aggressive, compelling them to adopt similar, less aggressive behaviours in terms of acceleration and speed. This results in the formation of vehicle platoons that are predominantly led by drivers of lower aggressiveness, either Mild or Normal. These platoons have similar acceleration but with

different fluctuations, as shown in Figure 3.19c. The original Aggressive platoons exhibit larger acceleration variations, which aligns with the observations of two-mixed traffic flow.

3.4.3 Transferability analysis

To validate the proposed approach and corresponding findings, we conducted a transferability analysis using the NGSIM-I80 dataset. The simulation outcomes are presented in Figure 3.20. The colour shifting in all indicators of traffic safety and sustainability suggests that an increase in the proportion of Aggressive drivers correlates with a reduction in traffic safety, as well as an increase in fuel consumption and emissions. Conversely, the presence of Mild drivers results in a diminishing effect on both traffic safety and environmental issues. Such findings are aligned with the outcomes obtained from the HighD dataset.

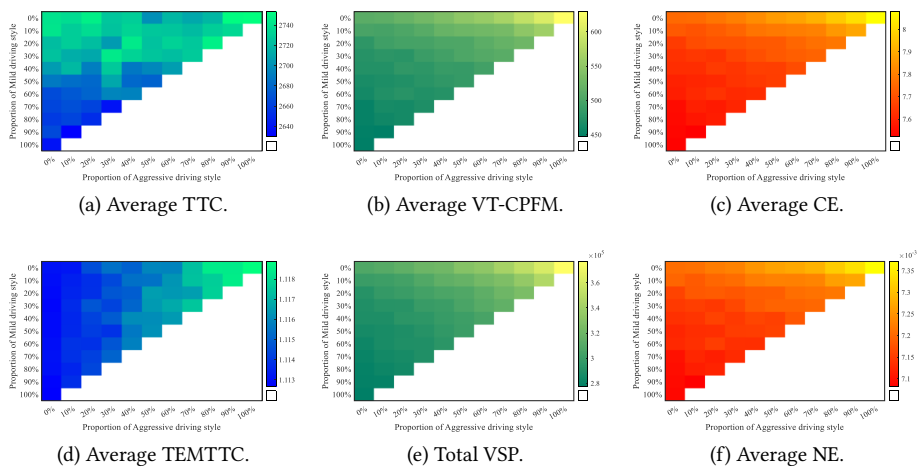


Figure 3.20: Statistics of traffic safety and sustainability indicators in 66 traffic scenarios (using NGSIM dataset).

Several representative traffic scenarios with different driving style proportions are further examined, with results shown in Tables 3.3-3.4. The three proportions represent Aggressive, Normal and Mild driving styles in a certain traffic flow scenario. The traffic scenario with 100% Normal drivers serves as a baseline, with its indicators presented in absolute values. Relative changes in traffic safety, fuel consumption, and emissions for other traffic scenarios are calculated against this baseline, with positive and negative changes indicated in shades of red and blue, respectively; and darker shades signify larger magnitude changes. For example, the traffic scenario with 100% Aggressive drivers showed a notable increase in fuel consumption - specifically, a 31.35% rise according to VT-CPFM compared to the 100% Normal driver traffic scenario. Conversely, scenarios involving 100% Mild drivers demonstrated notable improvements in reducing safety issues, fuel consumption, and emissions. Even though in a three-mixed traffic flow, such as with proportions of 20% Aggressive, 20% Normal, and 60% Mild driving styles, scenarios incorporating Mild drivers

consistently led to decreased safety risks, fuel consumption, and emissions when compared to a purely Normal driving scenario, despite the inclusion of Aggressive drivers.

Comparing the results in Tables 3.3-3.4 reveals that changes in the proportions of driving styles have a more pronounced effect on traffic performance in the NGSIM dataset. For example, a 100% Mild traffic flow in the NGSIM dataset improved emission by 11.86% and reduced fuel consumption by 52.02%, whereas in the HighD dataset, the improvements were only 7.89% and 34.56% HighD dataset, respectively. Similar findings can be observed from other driving style combinations such as 20%, 20%, and 60%. This discrepancy is attributed to the NGSIM dataset being collected under relatively congested traffic conditions, which increase the observable variations in driving behaviours compared to the HighD dataset.

Table 3.3: Impact of heterogeneous driving style on traffic flow with the NGSIM dataset.

Proportions	TTC	TEMTTC	CE	NE	VSP	VT-CPFM
0%,100%,0%	1.51×10^3	10.30	7.57	6.99×10^{-3}	4.49×10^5	4.56×10^2
100%,0%,0%	0.39%	102.99%	6.23%	0.51%	66.37%	34.80%
0%,0%,100%	-1.30%	-27.88%	-11.86%	-7.40%	-52.02%	-50.72%
60%,20%,20%	0.01%	79.46%	-1.73%	-4.51%	17.41%	-7.96%
20%,60%,20%	-0.08%	24.77%	-2.99%	-3.26%	-4.22%	-13.67%
20%,20%,60%	-0.47%	14.38%	-7.75%	-6.17%	-26.42%	-34.37%
30%,30%,40%	-0.13%	35.67%	-4.23%	-4.56%	-5.99%	-19.30%

Table 3.4: Impact of heterogeneous driving style on traffic flow with the HighD dataset.

Proportions	TTC	TEMTTC	CE	NE	VSP	VT-CPFM
0%,100%,0%	1.53×10^3	3.05	9.45	7.72×10^{-3}	1.40×10^6	1.32×10^3
100%,0%,0%	0.42%	56.12%	6.81%	1.64%	43.17%	31.35%
0%,0%,100%	-0.37%	-36.04%	-7.89%	-3.55%	-34.56%	-28.23%
60%,20%,20%	-0.30%	40.57%	-2.34%	-5.13%	5.84%	-8.52%
20%,60%,20%	-0.24%	17.96%	-2.94%	-3.59%	-4.77%	-11.31%
20%,20%,60%	-0.38%	-0.65%	-5.59%	-4.25%	-18.43%	-20.58%
30%,40%,30%	-0.30%	19.86%	-3.38%	-4.13%	-5.38%	-12.74%

3.5 Concluding remarks

This paper proposes a general framework to investigate car-following heterogeneity and its impacts on traffic safety, fuel consumption and emissions. The framework incorporates a rigorous driving style classification using a multi-class S3VM classifier and a micro-simulation process with 66 fine-grained heterogeneous traffic scenarios. The impacts of car-following heterogeneity on traffic flow performance are further elucidated from the mechanism of underlying characteristics of driving behaviour. The key findings of this research are summarised below:

i) S3VM vs. SVM performance: Driving styles classification reveals that S3VM-based classifiers notably outperform traditional SVM classifiers in driving style classification,

with accuracy improvements up to 8.60%. This enhancement is particularly significant when utilising an RBF kernel over a Linear kernel.

ii) Aggressiveness impacts: Less aggressive drivers can lead to the formation of vehicular platoons, thereby encouraging more aggressive drivers to adopt a milder driving style. Importantly, the formation of these platoons is influenced by both the proportion and spatial distribution of less aggressive vehicles, which makes the correlation between less aggressive driving and improvements in safety, fuel consumption, and emissions more complex.

iii) Diversity under traffic conditions: The impact of driving style diversity on traffic performance is more pronounced in congested traffic conditions.

This study promises potential benefits for Intelligent Transportation Systems (ITS) by improving traffic safety and sustainability. The limitation of this study is that only inter-driving heterogeneity is considered, where each driver maintains a consistent driving style throughout the simulation. Future studies can improve the simulation by considering both inter- and intra-driving style heterogeneity. For instance, identifying driving heterogeneity through the underlying mechanisms of driving behaviour using primitive driving patterns [101] or *Action patterns* [166], thereby allowing for varied driving styles over time in micro-simulations. Additionally, external factors such as driving environment and internal factors at driver psychology level are promisingly considered to improve driving style classification.

II

An Action-based Framework for Driving Heterogeneity Identification

4

Identification of Driving Heterogeneity using Action-chains

The content of this chapter has been published on

📖 Yao, X., Calvert, S. C., and Hoogendoorn, S. P. (2023). “Identification of Driving Heterogeneity using Action-chains.” In *IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, Bilbao, Spain, September 2023.

This chapter introduces a comprehensive framework for identifying driving heterogeneity from an action perspective. Driving trajectories are identified into *Action phases* with physical meanings based on rule-based segmentation techniques. The *Action chain* concept is then introduced by implementing *Action phase* transition probability. Evaluating using a naturalistic dataset indicates that this approach effectively identifies driving heterogeneity while providing clear interpretations.

4.1 Introduction

Driving behaviour plays a pivotal role in determining vehicle motion, substantially affecting traffic flow, fuel consumption, and emission. It is widely acknowledged that driving heterogeneity, which is defined as the difference between driving behaviours of driver/vehicle combinations under comparable conditions [24], does exist. Research has shown that this heterogeneity contributes to increased traffic accidents and congestion [18]. Additionally, in mixed automated-human traffic, accurate descriptions and predictions of human-driven vehicle (HDV) behaviour are crucial for the decision-making and control of connected and automated vehicles (CAVs). These have underlined the necessity of a better understanding and identification of the heterogeneity in human driving.

It is well established that driving heterogeneity encompasses both intra-heterogeneity, which refers to driver-independent variability, and inter-driving heterogeneity, which involves differences in driving behaviour among drivers [24, 47]. However, directly measuring or detecting driving heterogeneity is challenging due to its reliance on human cognitive and physiological processes. With the increasing availability of naturalistic driving data, various efforts have been made to comprehensively and quantitatively analyse driving heterogeneity. The identification of driving heterogeneity from observed driving behaviour is typically approached in two ways [167]: 1) Employing techniques to characterise driving behaviour by inferring driving profiles from distinct driving events, and 2) Analysing driving behaviour without explicitly creating driving behaviour profiles.

The former approach addressed the identification of driving heterogeneity as a classification or clustering problem, resulting in categorical output with discrete scales or numerical output with continuous scores. For example, Hoogendoorn et al. [32] developed a method to categorise driver states into low, medium, and high workload categories. Or, clustering techniques have been employed to define a few driving style groups such as aggressive, normal, and mild [18]. However, due to the stochastic and uncertain nature of driving behaviour, these limited groups are insufficient for capturing the diverse characteristics of driving behaviour. Additionally, the criteria used to define these groups are somewhat ambiguous and subjective, posing challenges in effectively eliminating individual biases.

In contrast to employing subjectively defined classes, some research has focused on identifying driving heterogeneity by presenting a driving style space containing a vast array of categories without explicitly establishing driving behaviour profiles. For example, Qi et al. distinguished driving styles based on a space that included over 20 different types [56]. Another study converted car-following sequences into a comprehensive array of primitive driving patterns, and the distributions of these patterns were then utilised

to analyse individual driving styles [40]. This approach allows for the recognition of a greater degree of variability in driving behaviour by encompassing various driving characteristics. However, it is essential to acknowledge that this broader categorisation of driving heterogeneity may lead to reduced clarity of the fundamental driving behaviours and a limited understanding of driving heterogeneity. Consequently, further research in this area is necessary to address these challenges.

To bridge these research gaps, a novel framework is proposed to identify heterogeneity in longitudinal driving behaviour from an "Action-chain" perspective. An "Action-chain" is defined as a series of "Action phases" over time. The contributions of this research are two-fold: i) A rule-based segmentation technique is presented to divide driving trajectories, considering the clear physical meanings of driving behaviour. ii) The concept of "Action phase" and "Action-chain" are first introduced to interpret driving behaviours, based on which a method for evaluating driving heterogeneity is proposed. The effectiveness of the framework was evaluated using real-world datasets, and the results demonstrate that the proposed methods can effectively identify driving heterogeneity at both individual drivers and traffic flow levels, providing clear interpretations. This approach offers valuable insights into understanding driving behaviour by uncovering underlying heterogeneity, which supports the development of accurate and robust driving behaviour and traffic flow models.

4.2 Framework Description

4.2.1 Defining Action Phase and Action-chain

The concept of "action points", which refers to specific moments of change in acceleration during driving [168], serves as the foundation to introduce the concept of "action trend" in this study. While action points capture acceleration or deceleration, they do not fully capture the complexity of driving behaviours. To overcome this limitation, we further propose the concept of "Action phase", which expands the scope by incorporating additional variables to provide more comprehensive information about driving behaviour.

By examining the univariate trajectory of driving behaviour, illustrated by the example of velocity (v) in Figure 4.1, distinct states are obviously observed. Some trajectories exhibit upward trends, others display downward trends, while some maintain a relatively stable range of fluctuations that can be considered as a keeping trend. We refer to these moments of driving behaviour different tendencies as *action trends*, which are segmented by turning points. Specifically, *action trends* are classified as "Increasing (I)", "Decreasing (D)", or "Stable (S)". To further refine the "Stable" trend, it is categorised as "Stable in a high value (H)" or "Stable in a lower value (L)". Thus, the *action trend* space can be represented as $S = \{I, D, H, L\}$, and the driving trajectory shown in Figure 4.1 can be expressed as $S_v = \{D, L, I, D, L\}$.

It is worth noting that while driving behaviour variables often exhibit synchronisation, our definition of *action trends* allows for variations in the temporal changes of different variables. For example, when the velocity state is "Increasing", the acceleration state can be "Increasing", "Decreasing", "Stable", or a combination of them. Consequently, the definition of *action trends* can be extended to other driving behaviour variables, such as acceleration and space headway. Thereafter, the concept of *Action phase* is proposed by encompassing




Figure 4.1: Visualisation of time-series driving trajectory: An example of velocity

multiple variables, and each *Action phase* label consists of multiple *action trend* names. These *action trend* names are estimated using uniform criteria derived from the group level of drivers in a certain traffic flow.

To account for the inherent sequential nature of driving behaviour, it is essential to consider the temporal dependencies between *Actions phases*. Hence, the concept of an *Action-chain* is introduced to represent a sequence of *Action phases* and their relations. The behaviour of a vehicle over time may consist of one or more *Action-chains*, each corresponding to different responses to the environment. With the “Action-chain” structure, driving behaviour over time can be characterised, which provides valuable insights into the underlying patterns and heterogeneity of driving behaviour.

4.2.2 Introducing the Novel Framework

The proposed framework for identifying driving heterogeneity aims to estimate frame-wise driving trajectories and identify driving heterogeneity within specific traffic flow conditions. The entire procedure is illustrated in Figure 4.2, consisting of five main steps: Data Preparation, Trajectory Segmentation, Action phase Extraction, Action-chain Establishment, and Heterogeneity Evaluation. The extraction of Action phase and the establishment of Action-chains involve the preceding steps called Driving Behaviour Interpretation and Action-chain Implementation, respectively.

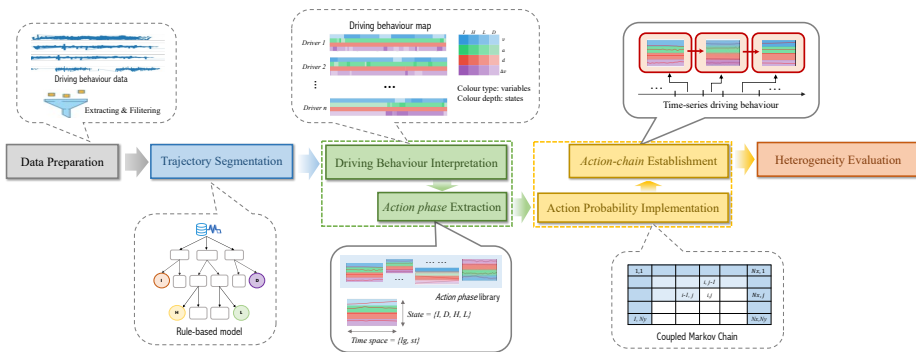


Figure 4.2: A novel framework of identifying driving heterogeneity

Data plays a crucial role in the identification of heterogeneity and serves as a fundamental aspect of the analytical process. After data tracing and preprocessing, the

time-series driving behaviour data are used as input for the segmentation algorithm (**Algorithm 1**). It is represented as x_1, x_2, \dots, x_t , where x_t denotes the driving behaviour variable feature at the t -th frame. The segmentation algorithm (**Algorithm 1**) outputs l_n^m , which represents the *action trend* names of variable m to be recognised for the n -th segment, where $n = 1, 2, \dots, N$. Based on the segmentation results, the driving behaviour of individual drivers can be visualised using driving behaviour maps, which highlight the unique characteristics of each driver. In the driving behaviour map, at the t -th frame, the state of driving behaviour is denoted as $S_t = \{l^1, l^2, \dots, l^m\}$. Subsequently, **Algorithm 2** is designed to detect driving behaviour segments in which all variables have a single *action trend*. The output, denoted as *Action phase* and represented as $S_{n'} = \{l^1, l^2, \dots, l^m\}$, signifies the *Action phase* for the n' -th segment, where $n' \in N'$, N' denotes the total number of *Action phases* for an individual driver. All the output *Action phases* form the "Action phase Library" under a specific traffic flow. The actual size of this Library is generally smaller than the theoretical value m^4 due to the nonexistence of certain state combinations in the real world, in accordance with fundamental driving behaviour theories. The length of the *Action phase* at the n' -th segment, referred to as the time label, is denoted as $\mathcal{T}_{n'}$.

Considering the time-series nature of driving behaviour, an Action phase transition probability algorithm (**Algorithm 3**) is implemented to capture the temporal dependencies between *Action phases*. An *Action phase* and the next *Action phase* obtained through the maximum transition probability constitute an *Action-chain*, representing the most probable driving behaviour adopted by drivers. The *Action-chain* serves as a description of homogeneous driving behaviour and is used to distinguish heterogeneity in driving behaviour. Drivers who deviate more from the *Action-chains* are considered to exhibit greater heterogeneity (conducted by **Algorithm 4**).

4.3 Method Implementation

4.3.1 Rule-based Segmentation

Traditional classification algorithms, such as the K-nearest neighbour method, support vector machines, and Convolutional Neural Networks have been commonly used for classifying driving styles or recognising driving patterns [169]. However, the segments obtained using these algorithms often lack clear interpretability in terms of physical characteristics. In contrast, rule-based segmentation is a relatively simple and interpretable method for dividing driving behaviour trajectories into meaningful segments. Therefore, we propose a rule-based method, referred to as **Algorithm 1** within the framework, to segment driving behaviour trajectories.

Let $V = \{v_1, v_2, \dots, v_m\}$ be a set of driving behaviour variables, such as velocity, acceleration, distance, etc. $P = \{(x_1, y_1), \dots, (x_n, y_n)\}$ represents a set of turning points for a single variable, which are calculated using calculus, specifically the first and second derivatives. **Algorithm 1** consists of the following steps:

1. Data preparation: Load the turning points of the selected variable. Calculate the variable changes Δy and time intervals Δx between neighbouring turning points.
2. Threshold setting: Define threshold values θ_1, θ_2 to differentiate between segments with state Increasing (I), Decreasing (D), or Stable (S). Set γ to determine whether a segment

is too short and should be merged with its neighbouring segments.

3. Initial categorisation: If $\Delta y > \theta_1$, meaning that the variable increases to a certain extent, which cannot be ignored, then label the segment as I. If $\Delta y < \theta_2$, in which case the variable decreases to a non-negligible level, then label it as D. When $\theta_2 < \Delta y < \theta_1$, the variable keeps within a small range of changes and is labelled as S.

4. Merging: For each segment labelled as S, if the time interval $\Delta x_n < \gamma$, and $\Delta x_{n-1} > \gamma, \Delta x_{n+1} > \gamma$, merge the segment with its neighbouring segment $n + 1$.

5. "Stable" refinement: For the updated S segments, calculate the mean value of the variable for each segment. Update the labels S as stable in High (H) or Low (L) values based on the threshold δ .

By implementing the above **Algorithm 1**, each variable in V is assigned *action trend* labels with clear physical meanings. This rule-based method allows for effective segmentation of driving behaviour trajectories based on single variables.

Subsequently, **Algorithm 2** is proposed to extract *Action phases* with simple steps including 1) segmenting the trajectory using turning points of all considered variables, and 2) removing segments shorter than the threshold of drivers' reaction time τ .

4.3.2 Time-series Action Phase Probability Modeling

The length of an *Action phase* can vary and is denoted by the labels "Long (lg)" or "Short (st)" according to a threshold η . Consequently, the time label space for *Action phases* is represented as $\mathcal{T} = \{lg, st\}$. Subsequently, *Action phase* can be further described by a two-dimensional label space $S' = \{S, \mathcal{T}\}$, which is taken as input for the Action phase transition probability model (**Algorithm 3**). Let S'_n and S'_{n+1} represent two adjacent *Action phases*, where $n' \in \{1, 2, \dots, N' - 1\}$. The transition probability between them can be mathematically represented as a function $\mathcal{R}(S'_n, S'_{n+1})$, which captures the underlying characteristics or patterns between S'_n and S'_{n+1} . This transition probability function provides insight into the relationship and dynamics between consecutive *Action phases* in the time-series analysis.

4.3.3 Coupled Markov Chain Theory

Two main approaches are commonly used to implement the transition of driving behaviour segments. The first approach utilises Markov models, including Markov Chains and Hidden Markov Models (HMM) [40], which are easily interpretable and capable of capturing underlying structures. However, when dealing with a large number of hidden layers, HMM may become computationally inefficient and less accurate due to increased complexity. The second approach involves deep learning models such as Recurrent Neural Networks (RNN), Long Short-Term Memory (LSTM), and Gated Recurrent Unit (GRU) Networks [40]. These models can address the complexity limitation of HMM and capture complex relationships between *Action phases*. Nevertheless, they typically require a large amount of training data and are computationally expensive due to their gating mechanisms.

In our case, the Markov Chain method is adopted to implement the Action Probability (**Algorithm 3**). The concept of a coupled chain refers to the collective behaviour of two independent systems, each following the principles of a classical Markov chain [170]. Let's

1,1					$N_x, 1$
			$i, j - 1$		
		$i - 1, j$	i, j		N_x, j
$1, N_y$					N_x, N_y

Figure 4.3: Conditional Markov chain on the states of the future

consider two one-dimensional Markov chains (X_i) and (Y_j) that operate on the state space $\{S_1, S_2, \dots, S_n\}$, with positive transition probabilities defined as

$$\Pr(X_{i+1} = S_k, Y_{j+1} = S_f | X_i = S_l, Y_j = S_m) = p_{lm,kf} \quad (4.1)$$

Here, the (X_i) chain describes the *Action phase* state S' and the (Y_i) chain describes the time label \mathcal{T} . Then the coupled transition probability $p_{lm,kf}$ on the state space $\{S_1, S_2, \dots, S_n\} \times \{S_1, S_2, \dots, S_n\}$ is given by

$$p_{lm,kf} = p_{lk} \cdot p_{mf} \quad (4.2)$$

Two coupled one-dimensional Markov chains can be utilised to construct a two-dimensional spatial stochastic process on a lattice represented by $(Z_{i,j})$. The lattice consists of a two-dimensional domain of cells, as depicted in Figure 4.3. The deep blue cells represent known boundary cells, the light blue cells indicate known cells within the domain (past observations), and the white cells represent unknown cells. The future state used to determine the state of cell (i, j) is cell (N_x, j) , where each cell is identified by its row number i and column number j . Then the conditional probabilities can be expressed as follows [171]:

$$P_{lk}^h = \Pr(X_{i+1} = S_k | X_i = S_l) \quad (4.3)$$

$$P_{mk}^v = \Pr(Y_{j+1} = S_k | Y_j = S_m) \quad (4.4)$$

The stochastic process $(Z_{i,j})$ is obtained by coupling the Markov chains (X_i) and (Y_j) while ensuring that these chains transition to the same states. Therefore, we have:

$$\begin{aligned} & \Pr(Z_{i,j} = S_k | Z_{i-1,j} = S_l, Z_{i,j-1} = S_m) \\ &= C \Pr(X_i = S_k | X_{i-1} = S_l) \Pr(Y_{j-1} = S_m) \end{aligned} \quad (4.5)$$

Here, C is a normalising constant that arises from restricting transitions in the (X_i) and (Y_j) chains to the same states. It is calculated as:

$$C = \left(\sum_{f=1}^n p_{lf}^h \cdot p_{mf}^v \right)^{-1} \quad (4.6)$$

By combining Equation 4.5 and Equation 4.6, the required probability can be expressed as:

$$\begin{aligned}
p_{lm,k} &:= \Pr(Z_{i,j} = S_k | Z_{i-1,j} = S_l, Z_{i,j-1} = S_m) \\
&= \frac{p_{lk}^h \cdot p_{mk}^v}{\sum_f p_{lf}^h \cdot p_{mf}^v}, k = 1, \dots, n
\end{aligned} \tag{4.7}$$

4.4 Data-based Evaluation

4.4.1 Data Preprocessing

In this study, the NGSIM highway dataset, which includes data from I-80 and US-101, was utilised to investigate the heterogeneity of longitudinal driving behaviour based on our proposed framework. A comprehensive preprocessing of the dataset, involving filtering and extraction, was conducted as described by Sun et al. [18]. Especially, drivers with trajectories lasting at least 50 seconds were selected to ensure an adequate amount of data for analysing longitudinal driving behaviour [172]. The final extracted dataset consisted of 123 drivers from the I-80 dataset and 848 drivers from the US-101 dataset.

The driving behaviour variables considered in this study were velocity (v), acceleration (a), distance (d) between the preceding and following vehicles, and their speed difference (Δv). The threshold values used in **Algorithm 1** were determined based on empirical knowledge from literature [173], as summarised in Table 4.1.

Table 4.1: Parameter settings of **Algorithm 1**

Δy /(unit)	θ_1	θ_2	δ	γ	τ	η
$v/(m/s)$	2	-2	20	30	10	50
$a/(m/s^2)$	0.25	-0.25	0.25	30	10	50
$d/(m)$	1	-1	1	30	10	50
$\Delta v/(m/s)$	2	-2	2	30	10	50

4.4.2 Visualisation and Analysis of Action phase

The *action trend* labels for the four driving behaviour variables are obtained using **Algorithm 1**. These results are then visualised, generating unique driving behaviour maps for each driver, as exemplified in Figure 4.4. In the figure, various colours represent different driving behaviour variables, with velocity, acceleration, distance, and speed difference represented in that order. The varying intensity of the same colour indicates different *action trend* names, including Increasing, stable as High, stable as Low, and Decreasing.

In Figure 4.4a, the dominant *action trend* for acceleration is “L”, although instances of “I” and “D” can also be observed. The distance remains relatively stable without frequent *action trend* changes. When comparing the driving behaviour maps of the four drivers shown, driver ID1264 from the I-80 dataset exhibits the fewest *action trend* changes across the four variables. Conversely, drivers ID3 and ID1035 from the US-101 dataset demonstrate a higher frequency of the changes.

Figure 4.4: Visualisation of *actions*: the driving behaviour mapTable 4.2: Statistics of *Action* (Top 10)

I-80		US-101	
<i>Action phase</i>	Frequency	<i>Action phase</i>	Frequency
((L,L,H,H), st)	415	((L,L,H,H), st)	2703
((L,L,H,H), lg)	219	((L,L,H,H), lg)	1661
((L,L,L,H), st)	156	((L,L,L,H), st)	965
((L,L,H,I), st)	68	((L,L,L,H), lg)	672
((L,L,L,H), lg)	65	((D,L,H,H), st)	651
((D,L,H,H), st)	54	((L,L,L,L), st)	480
((L,L,L,L), st)	41	((I,L,H,H), st)	469
((L,L,H,D), st)	39	((D,L,H,H), lg)	419
((D,L,H,I), st)	38	((D,L,H,I), st)	412
((D,I,H,I), st)	30	((L,L,H,I), st)	349

The driving behaviour map offers an intuitive approach to interpreting driving behaviour by visualising the changes in driving behaviour over time. It is also important to note that this visualisation method relies on observation and should be complemented with further quantitative evaluation, which will be carried out in subsequent steps.

The *Action phase* Library for a specific traffic flow was constructed using **Algorithm 2**. The resulting Library consists of 142 *Action phases* for the I-80 dataset and 228 *Action phases* for the US-101 dataset. Table 4.2 presents the top 10 *Action phase* along with their

corresponding frequencies. Notably, both traffic flows exhibit a significant overlap in their high-frequency *Action phases*, and the top three *Action phases* are identical for both datasets. These top *Action phases* include “((L, L, H, H), st)”, “((L, L, H, H), lg)”, and “((L, L, L, H), st)”, which indicate common driving behaviour across the datasets. The reason behind this is that the two datasets were collected during evening and morning rush hours respectively. In these periods, the density of traffic flow is significantly high, with most vehicles exhibiting car-following behaviour and even close to congestion. Due to this, there is limited variability in driving behaviours, resulting in a scarcity of “I” and “D” and a high frequency of “Stable (H and L)”. The high-density traffic flows also provide an explanation for the highest frequency of occurring “L”.

4.4.3 Analysis of Action-chain

The transition probabilities from one *Action phase* to another within the *Action phase Library* were computed using **Algorithm 3**. Some *Action phases* either have no transitions or exhibit very low probabilities of transition. Conversely, other *Action phases* tend to be transitioned to by a greater number of *Action phases*. The results adhere to the fundamental principles of driving behaviour. For example, the *Action phase* “((L, L, L, H), st)” from the US-101 dataset demonstrates higher probabilities of being transitioned. This can be attributed to the fact that the driving data were collected during the morning peak hour when there is typically high traffic flow density, leading drivers to adopt more consistent driving behaviours with lower values.

Overall, each *Action phase* was found to have a following *Action phase* with the highest transition probability, resulting in the formation of an *Action-chain*, as illustrated in the examples provided in Table 4.3. In the I-80 dataset, for instance, the *Action phase* “((D, I, I, H), st)” has a probability of 0.68 to transition to “((L, I, I, H), st)”, which is higher than any other *Action phases*.

Table 4.3: *Action-chain* composed by the highest joint transition probability (JTP)

Dataset	<i>Action phase</i> from	<i>Action phase</i> to	JTP
I-80	((D, D, I, I), st)	((D, L, L, I), st)	0.68
	((D, I, I, H), st)	((L, I, I, H), st)	0.68
	((D, I, D, I), lg)	((L, L, D, H), st)	0.67

	((D, D, H, D), lg)	((L, L, H, H), st)	0.52
US-101	((D, D, D, I), st)	((L, L, L, L), st)	0.64
	((D, H, L, H), st)	((L, L, L, H), st)	0.64
	((I, I, L, D), st)	((L, L, L, H), st)	0.58

	((I, I, D, I), st)	((L, L, D, L), st)	0.32

4.4.4 Evaluation of Driving Heterogeneity

In this study, we assume that the maximum transition probability represents the generally adopted *Action phase* of drivers in a specific traffic flow, indicating the average level of driving behaviour. However, in the real world, drivers often deviate from this general level

due to heterogeneity in their driving behaviours. To quantify the heterogeneity, we define this heterogeneity as the deviation between the actual *Action phase* transition and the *Action-chain*.

The Mean Squared Error (MSE) is a commonly used method for measuring the average squared difference between two sets of data, and it serves as the metric to quantify driving heterogeneity in this context, see Equation 4.8.

$$DH = \frac{1}{N'} \times \sum_{n'=1}^{N'} (P'_{n'} - P_{n'})^2 \quad (4.8)$$

where N' is the total number of *Action phases*, and $P'_{n'}$ and $P_{n'}$ represent the transition probability of actual *Action phase* and the maximum transition probability, respectively. A higher value indicates a greater driving heterogeneity.

Table 4.4: Drivers with the highest heterogeneity

I-80		US-101	
Driver ID	DH	Driver ID	DH
295	0.2172	582	0.1379
535	0.2111	628	0.1360
1174	0.2004	1157	0.1464
-	-	1647	0.1916




Figure 4.5: Driving heterogeneity in a specific traffic flow

4.4.5 Numerical Results and Discussions

The heterogeneity of individual drivers in a specific traffic flow was calculated and subjected to further statistical analysis using the normal distribution. The 3σ rule of thumb is commonly employed in data analysis to identify potential outliers or unusual behaviour. By applying the 3σ principle, drivers with atypical driving behaviour were identified, as

summarised in Table 4.4. These drivers may serve as potential factors contributing to increased traffic flow heterogeneity and negatively affecting traffic performance. It is also noteworthy that the degree of driving heterogeneity in the two traffic flows exhibits the same variance, see Figure 4.5; however, the drivers on US-101 (with $\mu = 0.08$) display overall lower levels of heterogeneity compared to those on I-80 (with $\mu = 0.10$).

4.5 Conclusion

In this study, a novel framework was proposed to address driving heterogeneity in a comprehensive manner. By introducing the concepts of *Action phase* and *Action-chain*, along with specialised algorithms, the framework effectively quantified and explained driving heterogeneity at both the individual driver and traffic flow levels. Real-world datasets were used for evaluation, validating the framework's ability to offer clear interpretations. Although the contributions of novel insights and findings of heterogeneity of driving behaviour, further validation and justification of the methods employed in each step are still required, which is a focus of our ongoing research.

5

A Novel Framework for Identifying Driving Heterogeneity through Action Patterns

The content of this chapter has been published on

📄 Yao, X., Calvert, S. C., and Hoogendoorn, S. P. (2025). “A Novel Framework for Identifying Driving Heterogeneity through Action Patterns.” *IEEE Transactions on Intelligent Transportation Systems*, doi: 10.1109/TITS.2025.3560509.

This chapter proposes a novel framework to identify driving heterogeneity from the underlying characteristics of driving behaviour. The framework includes three processes: *Action phase* extraction, *Action pattern* calibration, and *Action pattern* classification. The framework is then evaluated on a large-scale naturalistic driving dataset to identify *Action patterns*, and the attention mechanism is implemented in LSTM models to enhance both the accuracy and time efficiency of *Action pattern* identification.

5.1 Introduction

Driving behaviour varies from one driver to another, even the same driver can behave differently under identical traffic conditions depending on their current state [47]. This variability in driving behaviour is termed driving heterogeneity. There is strong evidence that driving heterogeneity leads to an increase in traffic accidents and exacerbation of traffic congestion [61, 174]. For instance, inconsistent acceleration and braking patterns of drivers significantly are major contributors to rear-end collisions [175]. Additionally, aggressive drivers may adopt a more moderate driving style when constrained by less aggressive drivers in front of them, which can decrease overall traffic efficiency [176]. To enhance the stability of traffic flow and improve safety, future autonomous vehicle (AV) control systems must account for the diverse driving styles of surrounding human-driven vehicles [27, 129]. Thus, a thorough understanding and identification of driving heterogeneity are crucial for reducing its adverse effects and improving traffic safety and efficiency.

Driving heterogeneity has been studied in the literature from various perspectives, including inter- vs. intra-driver heterogeneity, long-term vs. short-term variability, and global vs. situational patterns [177]. Long-term individual tracking in naturalistic driving data is often unavailable due to privacy concerns and resource constraints. Consequently, much of the driving behaviour analysis has focused on short trajectory segments, which have proven effective in identifying variations in driving characteristics [1]. Key behavioural characteristics, such as acceleration profiles [168] and car-following patterns [40], can be observed within short time windows, and when analysed systematically across large datasets, they reveal stable trends in driving heterogeneity. These short-term driving variations reflect immediate and context-dependent driver responses and are widely used in driving heterogeneity identification such as driving style recognition [18], risk level evaluation [178], and driving skill characterisation [54].

Machine Learning (ML) methods can model complex numerical relationships with accurate results [178], making them widely used in (transportation) research. Recently, the increased availability of naturalistic driving data has boosted using Machine Learning (ML) methods to identify driving heterogeneity. Unsupervised learning techniques such as K-means and fuzzy C-means (FCM) are commonly used to cluster driver behaviours into distinct groups, with driving styles inferred from statistical analyses of the clusters [18]. While this approach allows for direct extraction of behavioural patterns from raw data, it often lacks interpretability, and the discovered clusters may not always correspond to realistic or meaningful driving behaviours. Conversely, supervised learning methods train models to learn the relationship between input features and output labels, enabling them to classify or predict driving behaviours in new, unlabeled data. Various supervised

learning classifiers, such as Support Vector Machine (SVM), k-Nearest Neighbors (KNN), Long Short-Term Memory Network(LSTM) and Convolutional Neural Network(CNN)-based neural network models, have been employed to recognise driving styles, driving skills, and risk levels. Supervised learning models typically require labelled training data, which is often obtained through expert knowledge or a combination of other techniques. For example, some studies use rule-based strategies, where physical driving variables—such as large brake or accelerator pedal inputs—are used to infer aggressive driving styles [40]. Other studies utilise unsupervised learning techniques alongside statistical analysis to categorise drivers for assigning labels. For example, Deng et al. [77] categorised 30 participants into cautious, moderate, and aggressive driving styles using principal component analysis(PCA) and K-means clustering, creating labelled datasets for subsequent classification model training. With properly labelled data, supervised learning models have demonstrated remarkable accuracy in driving heterogeneity identification. SVM models, for instance, have achieved up to 95% accuracy in distinguishing driving patterns and identifying risky drivers [84, 178]. Similarly, deep learning models—such as LSTMs and CNNs—have been developed for driver behaviour classification, achieving recognition accuracies as high as 99% [93].

Labelling approaches often assign fixed driving profiles to drivers by analysing the mean or distribution of variables in a driving trajectory. This type of global trajectory analysis may overlook granular behaviour changes in time, such as specific acceleration changes at short time intervals, consequently missing the full spectrum of driving behaviour [50]. This highlights a necessity for a refined labelling process that uncovers the fundamental characteristics of driving behaviour. Moreover, omitting the temporal dynamics inherent in driving behaviour oversimplifies driving heterogeneity identification tasks, raising concerns about the performance of current supervised models. This underscores the need for advanced models that can handle complex driving data, improving both accuracy and computational efficiency. Attention mechanisms have been successfully used in neural networks to handle complicated classification tasks such as text classification [179] and network traffic classification [180]. This approach offers a promising solution to deal with driving heterogeneity tasks in complex time-series driving data, though its effectiveness in this context remains to be explored.

This chapter proposes a novel framework to identify driving heterogeneity by analysing the underlying characteristics of driving behaviour. The contributions of this research are threefold: Introduce the concepts of *Action phase* and *Action pattern* to decode and interpret driving behaviour, providing a new lens for understanding driving behaviour in a human-comprehensible manner; (ii) propose a novel framework to systematically identify driving heterogeneity, encompassing the processes of *Action phase* extraction, *Action pattern* calibration, and *Action pattern* classification; and (iii) implement an attention mechanism on LSTM models, which significantly improves the *Action pattern* classification accuracy and time efficiency.

5.2 Methodology

This section introduces the overall framework and briefly presents data and experimental setup. The flow diagram of the proposed framework (referred to as the action framework

hereafter) is illustrated in Figure 5.1, and the definitions of the key concepts are summarised in Table 5.1. In the proposed action framework, *Action phase* extraction initially deciphers driving trajectories into “primitives” with semantic meanings. Following this, *Action patterns* are calibrated by categorising *Action phases* and analysing driving characteristics. The outputs of *Action pattern* classification aid in the labelling process for the subsequent *Action pattern* classification, where several supervised learning algorithms equipped with attention mechanisms are utilised.

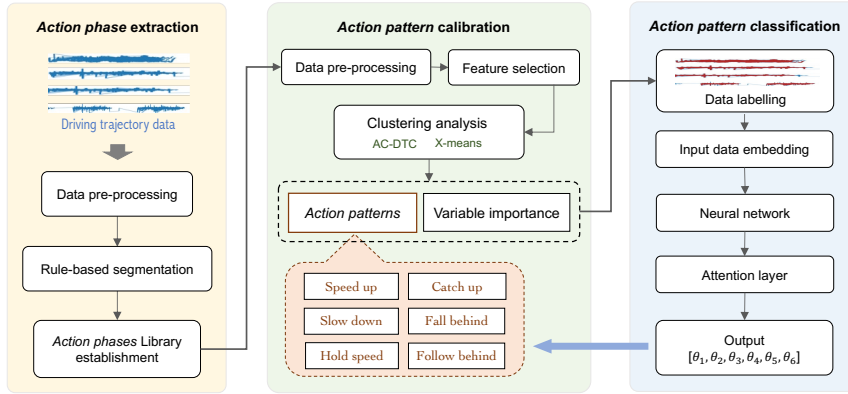


Figure 5.1: The flow diagram of the proposed action framework.

Table 5.1: Brief introduction of concepts.

Name	Meaning
action point [168]	Specific moments of change in acceleration during driving
turning point	Points on the driving variable curve where the concavity changes
<i>Action trend</i>	A distinct change within a driving variable (e.g., speed)
<i>Action phase</i>	Multi-variable driving trajectory segments with each univariate has a single <i>action trend</i> .
<i>Action phase Library</i>	All <i>Action phases</i> extracted from a certain traffic flow condition
<i>Action pattern</i>	Driving behaviours with group-specific driving characteristics

5.2.1 Action phase extraction

Action phases are defined as driving trajectory segments with distinct physical meanings, serving as “primitives” that capture underlying driving characteristics [181]. In literature, “action points” refer to specific moments of acceleration changes [168]. The proposed *Action phases* in this study offer a more comprehensive representation of driving characteristics by incorporating multiple variables, capturing the complexity of driving behaviours that single-variable analyses might overlook.

Action phases are extracted by identifying *action trends* of variables involved in driving behaviour analysis. An *action trend* refers to a distinct change within a driving variable

(e.g., speed), characterised as “Increasing (I)”, “Decreasing (D)”, or “Stable in a high value (H)” or “Stable in a low value (L)”. It should be noted that these action trends are chosen by choice and can also be selected in a different way. *Action phases* represent driving trajectory segments with multiple driving variables, where each univariate exhibits a single *action trend*. Driving variables considered in this study include velocity, acceleration, time headway, and speed difference. Both *action trends* and *Action phases* are identified and extracted using rule-based segmentation methods [181]. All *Action phases* extracted from a given dataset constitute the *Action phase Library* under that traffic condition. Detailed methodology and results of *Action phase* extraction are presented in Section 5.3.

5.2.2 Action pattern calibration

When multiple driving variables are considered, the number of *Action phase* categories increases rapidly as more variables are included. This can lead to minor differences in observed driving behaviour among different *Action phases* and complicate the interpretation of driving behaviour. To address this, we propose the concept of *Action pattern* that represents group-specific driving characteristics. Therefore, the second step of the action framework consolidates similar *Action phases* into a small number of *Action patterns* using unsupervised learning techniques.

Two clustering techniques, gglomerative clustering dynamic tree cut (AC-DTC) and X-means, are employed in this step to facilitate optimal results. Each cluster represents a distinct *Action pattern*, with the characteristics of each pattern obtained through statistical analysis of the *Action phase* data within the cluster. It is assumed that greater dissimilarity of a driving variable among clusters indicates higher importance in distinguishing *Action patterns*. Consequently, driving variable importance is determined. Detailed methodology and results of *Action pattern* calibration are presented in Section 5.4.

5.2.3 Action pattern classification

The third step of the action framework involves training models to classify *Action patterns* with high accuracy and time efficiency. Driving trajectories are first labelled as *Action patterns* using the rule-based method that considers driving variable importance. Then a bidirectional Long Short-Term Memory network integrated with an attention mechanism (ABi-LSTM) is applied for the *Action pattern* classification task, handling multi-variable varied-length data. Baseline models, including basic LSTM, bidirectional LSTM (Bi-LSTM), and attention LSTM (ALSTM), are also employed for the same classification task to evaluate the effectiveness of the ABi-LSTM model in training accuracy and time efficiency. Detailed methodology and results of *Action pattern* classification are presented in Section 5.5.

5.2.4 Data and experimental set up

The Lyft-5 open dataset contains large-scale real-world human driving data with detailed kinematic features, making it well-suited for identifying and analysing distinct characteristics in human driving behaviour. Thus, we utilise the Lyft level-5 as the primary dataset for evaluating the proposed framework in this study. This dataset includes 29k+ HV-following-AV (HV-AV) pairs and 42k+ HV-following-HV (HV-HV) pairs with a total

driving distance of 150k+ km in similar environments [182]. We exclude driving trajectories with a speed of $v = 0\text{m/s}$ to eliminate stopping and restarting behaviour. Additionally, only drivers with trajectories lasting more than 20 seconds were selected to ensure an adequate data volume for analysing longitudinal driving behaviours. A Kalman filter is employed to detect outliers and missing values of driving trajectory data. While larger sizes of data can improve robustness in capturing behavioural variability, they also introduce computational challenges. To balance behavioural detail and computational efficiency, we select 3000 HV-HV pairs for *Action phase* extraction and another set of 3000 HV-HV pairs for *Action pattern* classification. All the extracted *Action phases* are used as input for calibrating *Action patterns* and analysing driving variable importance. In the *Action pattern* classification task, car-following trajectories are first labelled as *Action patterns* according to driving variable importance. These labelled *Action patterns* serve as input for LSTM-based classification models. Detailed experiments are presented in the following sections.

5.3 Action phase extraction

This section details the methodology for *Action phase* extraction, followed by the presentation of experiments and results.

5.3.1 Definition description

A driving trajectory has distinct states, as illustrated by the example of the velocity (v) of an arbitrary vehicle in Figure 5.2. Some periods exhibit upward trends, some display downward trends, and others fluctuate in a small range which can be considered a stable trend. In this study, we refer to these characteristics of a driving trajectory as *action trends*, distinguished by turning points. Specifically, we classify *action trends* into “Increasing (I)”, “Decreasing (D)”, or “Stable (S)”. The “Stable” trend is further categorised as “Stable in a high value (H)” and “Stable in a lower value (L)”. Thus, the *action trend* space can be represented as $S = \{I, D, H, L\}$. For example, the velocity trajectory shown in Figure 5.2 can be expressed as $S_v = \{D, L, I, D, L\}$. Any variables considered in driving behaviour analysis have a corresponding *action trend* space, thus extending the univariate analysis to multivariate analysis. Note that it is possible to define action trends differently, either based on other variables or other characteristics.




Figure 5.2: Visualisation of *action trends*: An example of velocity

Driving variables often exhibit correlated changes. For instance, Makridis et al. [1] indicated a non-linear relationship between speed and longitudinal acceleration. *Action trends* can dissociate temporal changes of different variables. For example, an *action trend* of velocity with “Increasing” might coincide with various acceleration trends of “Increasing”,

“Decreasing”, “Stable”, or a combination of them. This desynchronisation enables *action trends* to address the correlation problem and capture the characteristics of individual variables independently. *Action phases* are extracted by identifying multi-variable driving trajectory segments where each univariate with a single *action trend*.

5.3.2 The procedure of Action phase extraction

Driving trajectory data is represented as x_1, x_2, \dots, x_t , where x_t denotes a driving variable at the t -th frame. The set of driving variables is denoted as $Y = \{v_1, v_2, \dots, v_m\}$, which are velocity, acceleration, time headway, and speed difference in this study. For a given variable v_m , set of turning points is represented as $P_m = (x_1^m, y_1^m), \dots, (x_n^m, y_n^m)$. The procedure *Action phases* extraction is illustrated in Figure 5.3. After data preprocessing, driving trajectory data are fed into a univariate segmentation algorithm. This algorithm assigns an *action trend* l_n^m to each driving segment for variable v_m in the n -th segment, where $n = 1, 2, \dots, N$. Consequently, the *action trends* of multi-variable driving segments at time frame t are denoted as $S_t = \{l^1, l^2, \dots, l^m\}$, highlighting the unique driving characteristics at that moment.




Figure 5.3: The procedure of *Action phase* extraction.

Next, trajectories with assigned *action trends* are fed into another rule-based segmentation algorithm for multivariate segmentation. The outputs of this segmentation process are *Action phases*, represented as $S_{n'} = \{A^1, A^2, \dots, A^{M'}\}$, where n' signifies the n' -th *Action phase*, $n' \in N'$. Here, N' is the total number of *Action phases* identified from an individual driving trajectory. All the extracted *Action phases* collectively form the *Action phase Library* under a specific traffic flow. Details of the univariate and multivariate segmentation processes are provided in Algorithm 1 and 2, respectively. The parameters for both rule-based segmentation algorithms are determined based on uniform criteria that represent collective driving behaviour within the given traffic condition.

5.3.3 Parameter settings

The parameters of rule-based segmentation algorithms are determined based on empirical knowledge from existing literature, as summarised in Table 5.2. Specifically, the threshold γ in Algorithm 1 and τ in Algorithm 2 are both set to 1 second to ensure that the extracted driving trajectory segments exceed the typical human reaction time, effectively capturing meaningful driving phases while preventing excessive fragmentation. This selection is supported by studies indicating that driver reaction times typically range between 0.7 and 1.5 seconds under various driving conditions [183]. Additionally, the thresholds for “stable high/low” in key driving variables—such as speed, time headway, and speed difference—are

Algorithm 1 Univariate Segmentation

```

1: Data Preparation
2: for each variable  $v$  in  $Y$  do
3:   Identify turning points ( $TP$ ) of  $v$ 
4:   Compute  $\Delta y$  and  $\Delta x$  between neighboring  $TP$ s
5: end for
6: Threshold Setting
7: Define thresholds  $\theta_1$  and  $\theta_2$  for  $\Delta y$ 
8: Set threshold  $\gamma$  for  $\Delta x$ 
9: Initial Categorisation
10: for each segment between turning points in each variable do
11:   if  $\Delta y > \theta_1$  then
12:     Label segment as Increasing (I)
13:   else if  $\Delta y < \theta_2$  then
14:     Label segment as Decreasing (D)
15:   else
16:     Label segment as Stable (S)
17:   end if
18: end for
19: Merging Segments
20: for each segment labeled S do
21:   if  $\Delta x$  of segment  $< \gamma$  and  $\Delta x$  of both neighboring segments  $> \gamma$  then
22:     Merge segment with its neighboring segment
23:   end if
24: end for
25: Stable Refinement
26: for each segment labeled S do
27:   Compute mean of variable values in segment
28:   if mean  $> \delta$  then
29:     Refine label to High (H)
30:   else
31:     Refine label to Low (L)
32:   end if
33: end for
34: return action trends for each variable in  $Y$ 

```

Algorithm 2 Multivariate Segmentation

```

1: Identify Turning Points
2: Detect turning points across all variables in  $Y$ 
3: Segment multivariate driving trajectories using detected turning points
4: Discard Short Segments
5: for each multivariate segment do
6:   if segment length  $< \tau$  then
7:     Discard segment
8:   end if
9: end for
10: Extract Action Phases
11: for each driver do
12:   Extract Action phases based on the refined and segmented data
13:   Add extracted Action phases to the Action Phase Library
14: end for
15: return Action Phase Library

```


set based on typical traffic flow conditions in non-congested environments. The acceleration thresholds are derived from driver perception studies, acceleration values reported in driver perception studies [184], ensuring that the segmentation process reflects the acceleration forces that drivers commonly experience and respond to in real-world driving.

Table 5.2: Parameter settings of Algorithm 1.

Δy /(unit)	θ_1	θ_2	δ	$\gamma / 0.1s$	$\tau / 0.1s$
$v/(m/s)$	1.5	-1.5	13	10	10
$a/(m/s^2)$	0.25	-0.25	0.25	10	10
$T/(s)$	0.15	-0.15	1.5	10	10
$\Delta v/(m/s)$	0.8	-0.8	0.8	10	10

5.3.4 Results of Action phase extraction

A total of 18800 *Action phases* with 255 different *action trend* combinations are extracted, constituting an *Action phase Library* of the selected 3000 drivers in Lyft-5 dataset. This *Action phase Library* serves as input data for the subsequent *Action pattern* calibration. Specifically, the extracted *Action phases* vary in duration, ranging from a minimum of 1.1 seconds to a maximum of 13.4 seconds, with an average length of 2.48 seconds. The most frequently observed *Action phases* are “(H, L, L, L)”, (H, L, H, L)”, and “(H, H, L, L)”, reflecting common driving characteristics across the dataset.

Figure 5.4: The differences in *actions trends* with similar stimuli

The frequency of identified *action trends* for each driving variable is presented in Table 5.3. Notably, the most dominant *action trends* are ‘H’ and ‘L’, suggesting that vehicles predominantly maintain steady car-following behaviour under the given traffic condition. To further demonstrate the variability in driver responses under similar stimuli, we illustrate *Action phases* of several drivers experiencing comparable driving conditions in Figure 5.4. Specifically, Figure 5.4(c) shows that these drivers have a time headway with a “Low” trend, indicating that they receive similar stimuli from their preceding vehicles. However, despite this commonality, their driving actions exhibit significant differences. As shown in Figure 5.4(a)-(b), some drivers accelerate (e.g., drivers represented in purple, blue, and grey), while others decelerate (e.g., the driver represented in red and green). These variations highlight the heterogeneity in driving behaviour, which may be influenced by factors such as individual driving styles, risk perception, and latent decision-making tendencies.

This finding underscores the importance of considering multiple driving variables when understanding and identifying driving heterogeneity.

Table 5.3: Frequency of *action name* in each driving variable.

	I	D	H	L
v	3860	2501	7745	4649
a	5764	6226	3483	3327
T	4320	4123	4008	6343
Δv	3661	3729	3383	8027

5.4 Action pattern calibration

This section presents the methodology and results of *Action pattern* calibration.

5.4.1 Clustering techniques

Clustering approaches are generally divided into two distinct categories [185]: hierarchical clustering and partitioning clustering. We employ both techniques to cluster *Action phases*, aiming to achieve optimal results for *Action pattern* calibration.

Agglomerative clustering (AC) is a hierarchical clustering technique that initiates with individual data points as singleton clusters and recursively merges them based on similarity [186]. This method constructs a dendrogram to explore different granularity levels of clustering. Such flexibility in adjusting cluster boundaries is particularly beneficial for *Action phases* data, where the number of clusters is unknown. The dynamic tree cut (DTC) method further refines the hierarchical clustering process by analysing the dendrogram's structure to make context-sensitive decisions on cluster division. The hierarchical nature of AC-DTC effectively captures the temporal dependencies and variations within *Action phases* data, facilitating accurate clustering of similar *Action phases* while reflecting subtle differences. In agglomerative clustering with dynamic tree cut (AC-DTC), the distance among clusters is determined using linkage methods, which represent the hierarchical tree structure by specifying how distances between clusters are calculated [187].

X-means is a partitioning clustering technique that extends the K-means algorithm by introducing a mechanism to automatically determine the optimal number of clusters (k) [188]. It starts with a lower bound for k and iteratively adjusts it, aiming to find an optimal number of clusters by balancing model complexity and fit based on specific criteria. This method maintains a level of scalability and computational efficiency similar to K-means while eliminating the need for pre-defined cluster numbers.

Evaluation indicators provide quantitative measures to guide the clustering process and validate clustering results. Three commonly used indicators - including Silhouette Score (SS), Calinski-Harabasz Index (CHI), and Davies-Bouldin Index (DBI) - are employed for both AC-DTC and X-means clustering in this study. Specifically, the best value of SS is 1 and the worst value is -1. A higher CHI score indicates that clusters are well-separated and dense within clusters. DBI values closer to 0 represent better clustering performance [189].

5.4.2 Experiments and results of calibration

➤ Data preparation

Action phases are composed of multiple driving variables, each with varying durations. The Resampling and Downsampling Method (RDM) is employed to standardise the length of *Action phases* for clustering analysis. Specifically, *Action phases* with lengths shorter than a reference length (e.g., the median length of all *Action phases*) are resampled using Fast Fourier Transform (FFT) and Inverse Fourier Transform (IFFT) [190]. Conversely, *Action phases* exceeding the reference length are downsampled through isometric extraction. To mitigate issues such as the curse of dimensionality and computational burden, Principal Component Analysis (PCA) is utilised to reduce the complexity of *Action phase* data. It consolidates variables into principal components that capture the most significant features. Results in Figure 5.5 reveal that for 96.08% of *Action phases*, the first principal component (PC1) accounted for over 95% of cumulative contributions, as highlighted by the red dotted line. Consequently, PC1s are selected as input for clustering algorithms.




Figure 5.5: Statistics of PC1’s cumulative contributions.

➤ Clustering analysis

Clustering analysis employs both AC-DTC and X-means, with each having ablation studies to facilitate optimal clustering results. Four linkage functions, including “Weighted”, “Average”, “Complete”, and “Ward”, are employed for AC-DTC clustering. The cluster number k for X-means ranged from 4 to 7. The results are presented in Table 5.4. For both clustering techniques, the best-performing settings for each evaluation indicator are highlighted in bold. Note that the “Ward” linkage function, which minimises the variance within clusters to create more homogeneous groups, demonstrated superior performance by identifying six distinct categories of *Action phases*. X-means clustering with $k = 6$ exhibits superior performance across all k values.

Figure 5.6 provides statistics on *action trends* within each cluster. Each block represents a specific *action trend* of a variable, with numbers indicating the percentage of each *action trend* within the six clusters. Darker colours represent higher proportions. For instance, *action trend* “D” for variable v is prevalent in Cluster 3 and relatively rare in other clusters, see Figure 5.6a. Notably, clustering results of both techniques show a high degree of similarity in the distribution of *action trends* for each variable within each cluster, (see

Table 5.4: Results of agglomerative and X-means clustering.

	linkage	SS	CHS	DBS	Clusters
AC-DTC	Average	0.37	4229.43	0.72	8
	Complete	0.28	6286.08	0.90	8
	Weighted	0.31	5823.54	0.92	7
	Ward	0.41	14673.48	0.83	6
X-means	/	0.44	14115.50	0.81	4
	/	0.45	14638.43	0.79	5
	/	0.46	15768.43	0.80	6
	/	0.38	14740.08	0.00	7

numbers in brackets in Figure 5.6. These consistencies justify the six identified categories of *Action phases*. Given the superior performance of X-means clustering over AC-DTC, subsequent analyses of cluster characteristics are based on X-means results.



(a) Agglomerative clustering with dynamic tree cut



(b) X-means clustering

Figure 5.6: Statistics of *action trends* in each cluster. (*After calibration, Clusters 1 to 6 are labelled with the following patterns: “Slow down”, “Speed up”, “Hold speed”, “Follow behind”, “Fall behind”, and “Catch up”).

➤ Recognition of Action patterns and variable importance

The six categories of *Action phases* correspond to six distinct *Action patterns*, recognised by characteristics of *action trends* within their respective clusters. Higher percentages of *action trends* indicate greater importance in reflecting driving characteristics and interpreting driving behaviours. As such, in Figure 5.6b, Cluster 1 is dominated by the *action trends* ‘D’ and ‘L’ for velocity (v) and time headway (T), respectively, indicating a “Slow down” pattern where velocity decreases and time headway remains low. Cluster 2 is characterised by increasing velocity and speed difference, suggesting a “Speed up” pattern. Cluster 3 shows a “Holding speed” pattern with velocity (v) maintaining a high value (H) and acceleration (a) remaining low (L). Clusters 4 and 5 both highlight the primary variables of velocity (v) and time headway (T), with Cluster 4 maintaining constant time headway (L) and Cluster

5 showing increasing time headway (I), indicating a “Follow behind” pattern and a “Fall behind” pattern, respectively. Cluster 6 is characterised by a decreasing time headway and consistently low velocity, indicating a “Catch up” pattern where the following vehicle is closing in their leading vehicle. These six *Action patterns* are detailed in Table 5.5.

Table 5.5: Results of *Action pattern* recognition (according to X-means clustering).

Clusters	Variables	<i>action trend</i>	Observations	<i>Action pattern</i>
1	v, T	D, L	Speed decreasing, time headway stable	Slow down
2	$v, \Delta v$	I, I	Speed increasing, speed difference increasing	Speed up
3	v, a	H, L	Speed stable, small acceleration	Hold speed
4	T, v	L, L	Time headway stable, speed stable	Follow behind
5	T, v	I, L	Time headway increasing, speed stable	Fall behind
6	T, v	D, H	Time headway decreasing, speed stable	Catch up

The importance of a driving variable in distinguishing *Action patterns* is indicated by its dissimilarity across the six clusters. This is quantified using Kullback–Leibler (KL) divergence, denoted as $D_{KL}(P||Q)$, which measures the deviation between two probability distributions P and Q [191]. In this study, *action trends* are quantified using the slope of a given variable: a positive slope indicates an increasing trend, while a negative slope implies a decreasing trend, and sequences of gentle slopes indicate a stable trend. Considering the diverse manifestations of identical *action trends* (e.g., linear increase, convex/concave progression, or slight fluctuations), a “sliding window” method to detect local trends within specific intervals. This approach ensures overall *action trend* representation by averaging trends across these windows.




Figure 5.7: KL divergence of four variables.

Results of the KL divergence analyses are illustrated in Figure 5.7. Significant variance is observed for time headway (T) and velocity (v) across clusters, underscoring their greater variable importance compared to acceleration (a) and speed difference (Δv). This aligns with real-world driving observations, where time headway and velocity are more perceptible and influential in driving behaviour. Recognised *Action patterns* and variable importance provide robust references for labelling driving trajectories in supervised learning driving behaviour analyses, which is then utilised for subsequent *Action pattern* classification.

5.5 Action pattern classification

In this section, we first introduce the methodology for *Action pattern* classification, followed by a presentation of the experiments and results, which include driving trajectory labelling, experimental settings, and results.

5.5.1 Attention-based bidirectional LSTM model

In this paper, we apply a bidirectional Long Short-Term Memory network integrated with an attention mechanism (ABi-LSTM) for *Action pattern* classification. The architecture of this model, illustrated in Figure 5.8, comprises the following four layers:

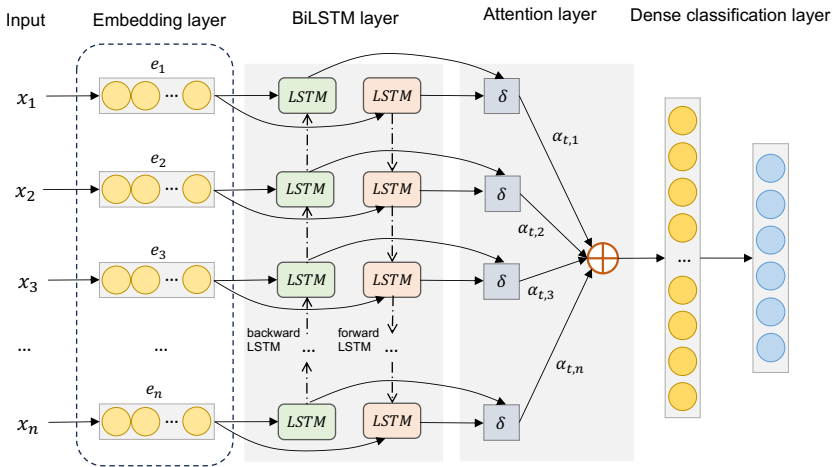


Figure 5.8: ABi-LSTM model architecture.

- **Embedding Layer:** This initial layer transforms each *Action phase* from the input sequence into a compact vector form, facilitating the capture of relationships between different *Action phases*.
- **Bi-LSTM Layer:** The vectorised sequence is then processed by a Bi-LSTM layer which captures contextual information from the entire sequence by analysing embedding data in both forward and backward directions.
- **Attention Layer:** An attention layer then computes context vectors by assigning attention weights to Bi-LSTM's hidden states. This mechanism allows the model to focus on the most relevant parts of the input sequence when classifying *Action phases*, particularly beneficial for handling highly complex sequences.
- **Dense Classification Layer:** The resulting context vectors are finally processed through one or more dense layers, culminating in a softmax output layer to classify *Action patterns* into multiple categories.

5.5.2 Experimental settings

➤ Driving trajectory labelling

As mentioned in Section 5.2.4, 3000 HV-HV pairs are selected from the Lyft-5 dataset for the *Action pattern* classification task. Car-following trajectories are labelled as various *Action pattern* sequences according to the criteria presented in Table 5.6. The labelling criteria are based on the importance of driving variables and *action trends*. Specifically, time headway (T) is the most critical variable, thus prioritised in the labelling process. Driving segments with an increasing trend in T increasing (I) indicates the ego vehicle is distancing from its leading vehicle, labelled as a “Fall behind” pattern. Conversely, a decreasing trend in T is labelled as a “Catch up” pattern. When T is constant, the focus shifts to velocity (v). An increasing trend in v is labelled as a “Speed up” pattern, while a decreasing trend in (v) is labelled as a “Slow down” pattern. Driving segments where both velocity and acceleration show minimal changes are labelled as a “Hold speed” pattern. Otherwise, they are labelled as a “Follow behind” pattern. Examples of the labelling results are shown in Figure 5.9. This labelling approach ensures that all drivers’ trajectories are categorised into distinct *Action patterns*, prepared for the subsequent classification task.

Table 5.6: Rule-based *Action pattern* labelling.

Variable	<i>action trend</i>					
T	I	D	S	S	S	S
v			I	D	S	S
a					I&D	S
Action patter	Fall behind	Catch up	Speed up	Slow down	Follow behind	Hold speed




Figure 5.9: Examples of labelled driving trajectories.

➤ Model training experimental settings

To address the variability in *Action pattern* data lengths, *Action pattern* sequences are reshaped to a consistent length for neural network batch processing. Traditional length reshaping methods such as truncation or padding can distort the original *Action pattern* data - either by omitting critical data or introducing irrelevant zeroes. Therefore, we

employ the PackedSequence technique in PyTorch to manage variable-length sequences [192] effectively. This technique sorts sequences by length and converts them into a PackedSequence object while retaining metadata that indicates each sequence's boundaries, allowing the ABi-LSTM model to process sequences according to their actual lengths without data distortion.

Each variable in *Action pattern* data is independently normalised within a [0, 1] range before being fed to the ABi-LSTM model. The input dimension of the ABi-LSTM model is set as 4, corresponding to the four driving variables considered in this study. Both the batch size and the model's hidden size are set to 64, and the dropout rate is 0.6. Additionally, we use k -fold cross-validation with the k set to 10 to ensure consistent model performance across different subsets of driving trajectory data. The Adam optimizer, known for its efficiency in handling sparse gradients and noisy data, is employed for model training, enhancing the ABi-LSTM model's adaptability to diverse data distributions. Ablation studies are conducted to optimise model performance in *Action pattern* classification, exploring LSTM hidden layers of 64, 128, and 256, and neural network layers ranging from 2 to 5. Learning rates vary from 0.00001 to 0.0001, increasing incrementally by 0.00002.

Benchmark models are utilised to assess the effectiveness of the bi-directional technique and the attention mechanism in *Action pattern* classification, including simple LSTM, bi-directional LSTM (Bi-LSTM), and attention-based LSTM (ALSTM). All hyperparameters are maintained consistent across the models to ensure a fair evaluation.

➤ Performance evaluation metrics

The models for *Action pattern* classification in the experiment are evaluated using four evaluation metrics: accuracy, F1-score, precision, and recall. These parameters are defined as follows.

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + TN + FN} \quad (5.1)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (5.2)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (5.3)$$

$$\text{F1-score} = \frac{2 \times (\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}} \quad (5.4)$$

where True Positives (TP) represent the number of correctly classified positive instances, and True Negatives (TN) denote the number of correctly classified negative instances. False Negatives (FN) refer to the number of positive instances incorrectly classified as negative, and False Positives (FP) signify the number of negative instances incorrectly classified as positive.

5.5.3 Results and discussions

Results of ablation studies on the ABi-LSTM model are illustrated in 5.10. The highest accuracy for both training and testing is achieved at a learning rate of 0.00006 and with

three layers, as depicted in Figure 5.10a and Figure 5.10b, respectively. The training and testing process with a learning rate of 0.00006 and three layers is illustrated in Figure 5.11a, with corresponding numerical results shown in Table 5.7. The training accuracy (see solid red line in Figure 5.11a) increases rapidly during the first 30 epochs, reaching an accuracy around 89%, then stabilises with small increases and finally converges at 200 epochs with an accuracy of 95.68%. This indicates effective learning of the training data by the model. The testing accuracy (see dashed pink line) follows a similar trend as the training process and converges with an accuracy of 94.33%, suggesting the model's good generalisation to unseen data. Both the training loss (solid blue line) and testing loss (dashed cyan line) decrease sharply in the initial 30 epochs and then stabilise at a low value of around 0.3, demonstrating effective error minimisation on both training and testing data. The difference in training and testing loss is around 0.02, indicating that the ABi-LSTM model learns the characteristics of *Action pattern* data without overfitting.

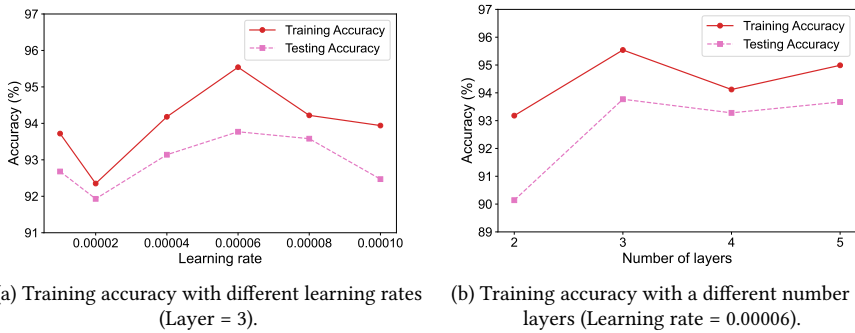


Figure 5.10: Parameter tuning with various learning rates and number of layers.

The training and testing results of baseline models are shown in Figure 5.11 and Table 5.7. In Figure 5.11b, the training and testing accuracies of the Bi-LSTM model steadily increase at the first 50 epochs with an accuracy of around 65%, then increase more slowly, finally converging at 600 epochs with accuracies of 89.88% and 88.97%, respectively. The training loss and testing loss decrease and converge at 0.47 and 0.52, respectively, which are higher than the ABi-LSTM model. The ALSTM model shows rapid increases in training and testing accuracy during the first 15 epochs, reaching around an accuracy of 77%, see Figure 5.11a. Then the accuracies increase slowly and the ALSTM model finally converges at around 200 epochs with training and testing accuracies of 92.41% and 88.51%, respectively. However, the noticeable gap between the training and testing process suggests overfitting, indicating Bi-LSTM models' inferior performance in the *Action pattern* classification task. The LSTM model shown in Figure 5.11d gradually improves training and testing accuracy and finally converges at 700 epochs with accuracies of 86.45% and 83.23%, respectively.

Overall, the ABi-LSTM model achieves the highest accuracy in fewer epochs compared to baseline models, with optimal results of precision (89.39%), recall (87.83%), and F1-score (88.60%), illustrating its remarkable effectiveness in classifying *Action patterns*.

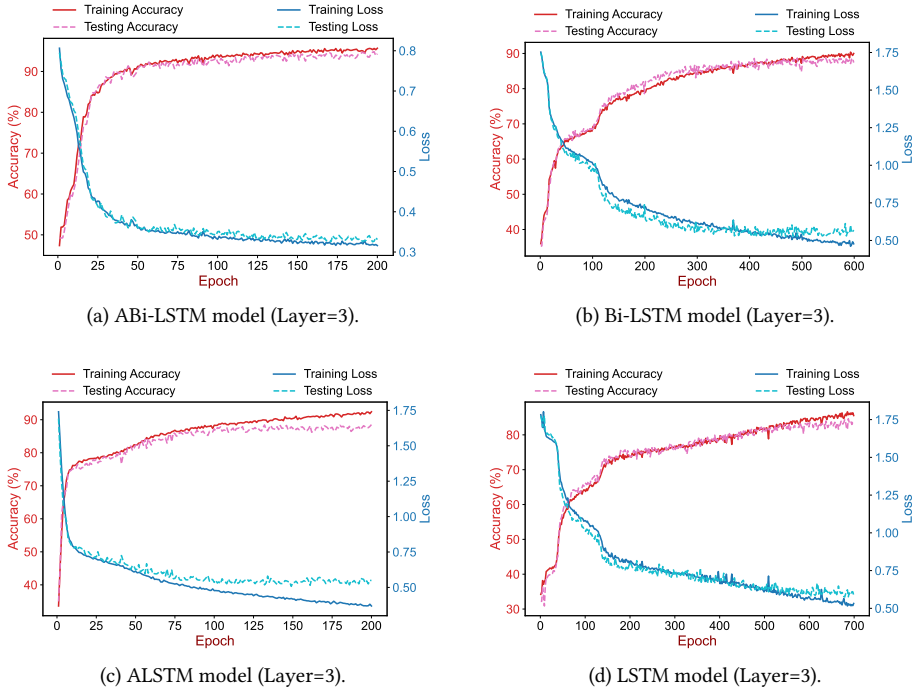


Figure 5.11: Performance of LSTM models and Attention-LSTM methods.

Additionally, models incorporating the attention mechanism (ABi-LSTM and ALSTM) consistently outperform their non-attention counterparts (Bi-LSTM and LSTM). These findings demonstrate the advantages of the attention mechanism in better generalisation by focusing on highly relevant features, thus enhancing classification precision and computational efficiency. The robust performance of *Action pattern* classification highlights the effectiveness of the proposed action framework in analysing complex driving behaviours and identifying driving heterogeneity in a comprehensive way.

Table 5.7: Performance comparison of different models for *Action pattern* classification.

Models	Layers Num.	Epoch	Training acc.(%)	Training loss	Testing acc.(%)	Testing loss	Precision (%)	Recall (%)	F1-score
ABi-LSTM	3	200	95.68	0.31	94.33	0.33	93.89	91.83	92.85
BiLSTM	4	600	89.88	0.47	88.97	0.52	88.27	89.44	88.85
ALSTM	3	200	92.41	0.36	88.51	0.51	91.94	87.10	89.45
LSTM	4	700	86.45	0.52	83.23	0.60	83.41	85.71	84.55

5.6 Discussion

In this section, we summarise the main findings of this study and discuss the limitations along with directions for future research.

5.6.1 Main findings

1) Action phase: The concept of *Action phase* offers a new perspective to decode and interpret the intrinsic characteristics of driving behaviours. The established *Action phase Library* serves as a comprehensive database that encompasses a wide range of driving characteristics within a traffic flow, allowing for more detailed and structured analyses of driving behaviour.

2) Action pattern: Distinct *Action patterns* can be identified from *Action phases* data, representing various group-specific driving characteristics. Six *Action patterns* are delineated in this study, interpreted as “Speed up”, “Slow down”, “Hold speed”, “Catch up”, “Fall behind”, and “Follow behind”. Each pattern represents a unique aspect of driver stimuli-response during driving, highlighting the heterogeneous nature of driving behaviour. This categorisation provides a granular view of driving trajectories, facilitating a more nuanced understanding and interpretation of driving heterogeneity, which can be beneficial for real-world applications such as traffic management and personalised ADAS.

3) Driving variable importance: Time headway (T) and velocity (v) exhibit higher importance than acceleration (a) and speed difference (Δv) in distinguishing *Action patterns*, underscoring their critical roles in reflecting driving characteristics. This aligns with realistic driving observations where time headway and velocity are more perceptible and influential in driving behaviour. Understanding the importance of driving variables can enhance driving behaviour analysis, such as improving the labelling process in supervised learning tasks by prioritising more influential variables.

4) Attention-based LSTMs: Incorporating the attention mechanism significantly enhances the performance of LSTM models in handling the complex *Action pattern* classification task with varied-length, multi-variable time-series data. The ABi-LSTM model improves classification accuracy by 1.39%-8.53% and reduces computational complexity by 66.67%-71.43% compared to baseline counterparts. This finding underscores the effectiveness of advanced techniques in improving accuracy, generalisation, and time efficiency for deep learning models, making them increasingly applicable for real-time driving behaviour analysis.

Overall, these findings highlight the advantages of the proposed action framework in identifying driving heterogeneity: i) the action framework offers a more precise, interpretable, and efficient approach to analyse complex driving behaviours and ii) provides actionable insights for its practical applications, such as enhancing traffic flow analysis, improving vehicle automation systems, and advancing driving behaviour modelling.

5.6.2 Limitations and future work

While the proposed action framework offers promising advantages in understanding and identifying driving heterogeneity, it also has certain limitations that provide valuable directions for future research. These are outlined below.

1) Threshold sensitivity and environmental dependency: Threshold settings are crucial for rule-based *Action phase* extraction, as they influence how driving trajectory data is segmented into meaningful pieces that accurately reflect driving characteristics. This makes the action framework sensitive to threshold selection, where slight adjustments could significantly impact the resulting *Action phases* and the subsequent *Action pattern* calibration. These thresholds are currently derived from expert insights within the literature. However, their effectiveness may vary depending on road characteristics, such as speed limits, roadway design, and traffic infrastructure, which can influence speed and acceleration distributions.

To improve the framework's adaptability across diverse driving environments, future research could explore adaptive thresholding techniques, where threshold values are dynamically optimised based on dataset characteristics rather than predefined fixed values. Additionally, conducting sensitivity analyses across different roadway settings could help refine threshold selection and ensure greater robustness in real-world applications.

2) Selection and interdependence of driving features: Driving variables play an important role in the action framework, as they influence the feature selection process and, consequently, impact the overall identification results. In this study, we utilise four variables (v , a , T , and Δv), which are empirically known to be highly relevant to driving behaviour. The proposed *action trends* treat these variables independently by dissociating variable changes.

Given the correlations among driving variables, future research should explore multivariate feature selection techniques or dimensionality reduction methods to better capture the relationships between variables. Such improvements could further refine the *Action phase* extraction and *Action pattern* calibration processes, improving the framework's ability to distinguish nuanced variations in driving behaviour.

3) Generalisability across different driving contexts: Driving behaviour is inherently influenced by both internal (e.g., driver cognition, risk perception) and external (e.g., road conditions, traffic flow, infrastructure) factors. This study evaluates the proposed action framework based on observed driving variables, and its generalisability across different traffic environments remains further consideration. One potential factor to be considered is the presence of autonomous vehicles (AVs) in the dataset. Interactions between human-driven vehicles (HVs) and AVs may introduce behavioural adaptations that differ from conventional HVs-only environments.

Future studies could investigate how HV-AV interactions influence driving heterogeneity and whether the extracted *Action patterns* remain consistent in AV-integrated traffic. Additionally, testing the framework across datasets from different geographical regions, where variations in driving culture, regulations, and infrastructure exist, can further validate its robustness and applicability.

4) Practical implications for real-world applications: The proposed action framework systematically interprets driving behaviour and can promote implications for various real-world applications. For instance, Advanced Driver Assistance Systems (ADAS) and Adaptive Cruise Control (ACC) could leverage the extracted driving patterns to predict driver behaviour, leading to more adaptive and personalised automation features. Traffic management and control systems could benefit from the framework's ability

to classify heterogeneous driving behaviours, helping optimise traffic signal timing, lane management, or speed harmonisation strategies. Safety assessment models could incorporate *Action patterns* to analyse driving risks, helping improve proactive measures for accident prevention. Future research could further explore how the action framework can be integrated into simulation environments, vehicle automation algorithms, or traffic control systems, ensuring its practical applicability in transportation engineering and intelligent mobility solutions.

5.7 Conclusion

Understanding the heterogeneous nature of driving behaviour is crucial for traffic flow analysis and the design of better road safety measures. Existing methods often lack the granularity and precision required to capture subtle heterogeneity in driving behaviour. This study proposes a novel framework to systematically identify driving heterogeneity by analysing underlying characteristics of driving behaviour. The concept of *Action phase* is introduced to decompose driving trajectories into “primitives” with physical meanings. Then *Action patterns* are calibrated by clustering *Action phases* based on group-specific characteristics. The *Action pattern* calibration process provides a rigorous labelling process for *Action pattern* classification. Evaluation using a large-scale naturalistic driving dataset demonstrates the framework’s effectiveness in capturing driving characteristics and identifying driving heterogeneity. The incorporation of an attention mechanism enhances LSTM models’ performance in terms of both accuracy and time efficiency. This framework improves capturing subtle behavioural differences within and amongst individual drivers, supporting advancements in personalised driving assistance systems and user-based traffic management.

III

A Pattern-based Approach for Driving Heterogeneity Modelling and Simulation

6

Human Driving Pattern Modelling: A Knowledge-Enhanced Deep Learning Approach

The content of this chapter is based on

📖 Yao, X., Qin, Z., Calvert, S. C., and Hoogendoorn, S. P. (2025). "Human Driving Patterns: A Knowledge-Enhanced Deep Learning Approach for Behaviour Modelling." *IEEE Transactions on Intelligent Transportation Systems, R1 Revision*.

This chapter introduces a novel approach to model longitudinal driving behaviour using knowledge-enhanced deep-learning (DL) models. We propose a Knowledge-Enhanced Attention LSTM (KE-ALSTM) model to predict transitions and durations of *Action patterns*. Graph-based and distribution-based knowledge are integrated to improve DL model performance. Evaluation on real-world data demonstrates that KE-ALSTM outperforms baseline models, demonstrating the value of incorporating domain knowledge to enhance deep-learning models in driving behaviour analysis.

6.1 Introduction

Human driving errors, encompassing mistakes of steering, braking or acceleration, as well as distractions caused by mobile devices or human-machine interfaces, frequently lead to unsafe interactions between human-driven vehicles and other road users, causing many traffic accidents [27]. According to the OECD ITF Road Safety Annual Report [193], human factors account for 93–98% of all motor vehicle crashes, underscoring the critical need for accurate modelling and prediction of driving behaviour to improve traffic management and road safety [129]. However, the inherently dynamic and heterogeneous nature of human behaviour, combined with the complexity of driving environments, presents significant challenges for developing robust driving behavioural models.

Over the years, numerous approaches have been proposed to model human driving behaviour [194]. Many efforts have focused on representing longitudinal driving behaviour [25, 195, 196], ranging from physics- and psychology-based formulations to data-driven methods [47]. Among these categories, data-driven models, particularly those leveraging Deep Learning (DL) methods, have gained increasing attention due to their ability to effectively handle high-dimensional, nonlinear relationships beyond the reach of traditional approaches. For example, Khodayari et al. developed an Artificial Neural Network (ANN)-based car-following model that considers drivers' reaction delays, achieving significantly lower prediction errors compared to other car-following models [197]. Similarly, the Bidirectional Long-Short Term Memory Networks and Conditional Random Field (Bi-LSTM-CRF) model proposed by Zhou et al. [198] effectively captures the relationship between driving context and human actions, outperforming state-of-the-art algorithms. Despite the advances, most DL-based studies assume uniform driving behaviour across all drivers, neglecting the heterogeneous nature of human drivers.

To address driving heterogeneity, researchers have clustered drivers into groups with distinct driving styles [18] or introduced models that reflect asymmetric behaviours such as hysteresis and intensity variations [199]. While these approaches capture inter-driver variability, they struggle to represent intra-driver variations that evolve over time or under changing driving conditions [177]. To overcome these limitations, recent efforts have attempted to decompose complex driving behaviour into simpler, interpretable behavioural units to capture more nuanced characteristics of the driving process. Wang et al. decomposed longitudinal driving trajectories into “primitives” and identified 10 different driving patterns such as “closing in”, “follow behind”, and “aggressive acceleration” [40]. Yao et al. utilised rule-based methods to segment car-following trajectories into *Action phases* with semantic meanings, subsequently identifying six *Action patterns* through an

unsupervised learning approach [58]. Such pattern-based representations provide the underlying characteristics of driving heterogeneity, while remaining underutilised in behavioural modelling.

Developing pattern-based behaviour models faces two main challenges. First, trajectories must be segmented and labelled into meaningful patterns to ensure reliable downstream modelling. Second, models must capture not only the regimes of driving patterns but also their transition dynamics and temporal characteristics, i.e., when drivers switch patterns and how long they maintain them. Second, the developed behaviour models must be capable of capturing not only the driving pattern regimes but also their transition dynamics and temporal characteristics, i.e., when drivers switch patterns and how long they maintain them. Our previous study addressed the first challenge by proposing an action-based approach for identifying driving patterns (referred to as *Action patterns*) and established a rigorous labelling process comprising data preparation, feature selection, *Action pattern* calibration, and evaluation of driving variable importance [58]. However, the second challenge remains underexplored in existing literature. The main difficulty lies in handling the complexity of *Action pattern* data, which is inherently multi-dimensional and varies in sequence length. Knowledge-enhanced models, which integrate domain-specific knowledge with data-driven approaches, have demonstrated advantages in vehicle trajectory prediction, such as in autonomous driving contexts where expert knowledge can help to enhance both modelling accuracy and interpretability [200]. Such models offer promising solutions for modelling the contextual and structural complexities of pattern-based driving behaviour.

To address the second challenge, this chapter proposes a novel approach for driving behaviour modelling using driving pattern sequences and knowledge-enhanced DL models. The unique contributions are: (i) Proposing a Knowledge-enhanced attention LSTM (KE-ALSTM) model to capture the underlying characteristics of driving patterns; (ii) Incorporating *Action pattern* transition and duration properties as additional knowledge enhances the performance of DL models in addressing the multi-task pattern-based driving behaviour modelling. Evaluation using a real-world dataset demonstrates that the proposed KE-ALSTM model significantly outperforms benchmark models without knowledge enhancement in both accuracy and computational efficiency.

6.2 Methodology

In this section, we first describe the pattern-based driving behaviour modelling approach, then propose a knowledge-enhanced attention LSTM model (KE-ALSTM) for *Action pattern* modelling.

6.2.1 Problem description

➤ Review of Action phase and Action pattern-based methods

According to Yao et al. [57], *Action phases* are defined as “primitives” to represent underlying characteristics of driving trajectories. This concept introduces multi-variables to capture more complex driving behaviours such as short-term catching up with leading vehicles, which expands “action points” [168] that only consider changes in acceleration. This

expansion is achieved by identifying *action trends* of specific variables (e.g., velocity, acceleration, time headway, and speed difference) involved in driving behaviour analysis. An *action trend* refers to a distinct change of a driving variable within a driving trajectory, characterised as “Increasing (I)”, “Decreasing (D)”, or “Stable (S)”. *Action phases* are driving trajectory segments that involve multiple parallel driving variables, where each univariate exhibits a single *action trend*. Both *action trends* and *Action phases* are identified and extracted using rule-based segmentation methods proposed by Yao et al [57].

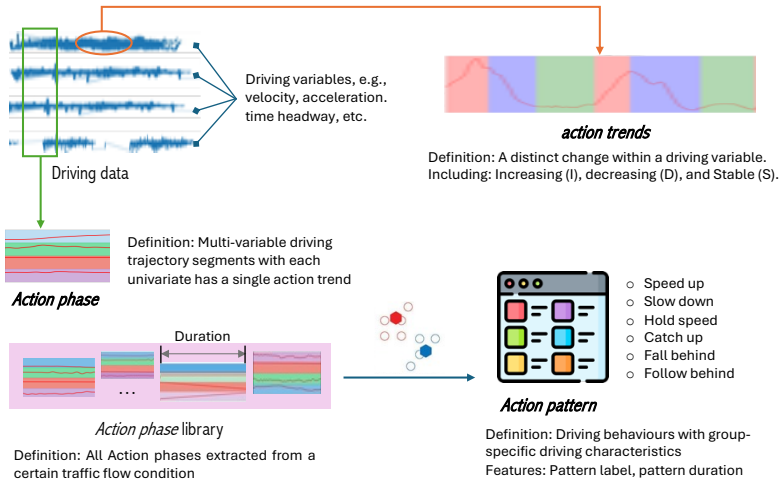


Figure 6.1: Clarification of action-related concepts.

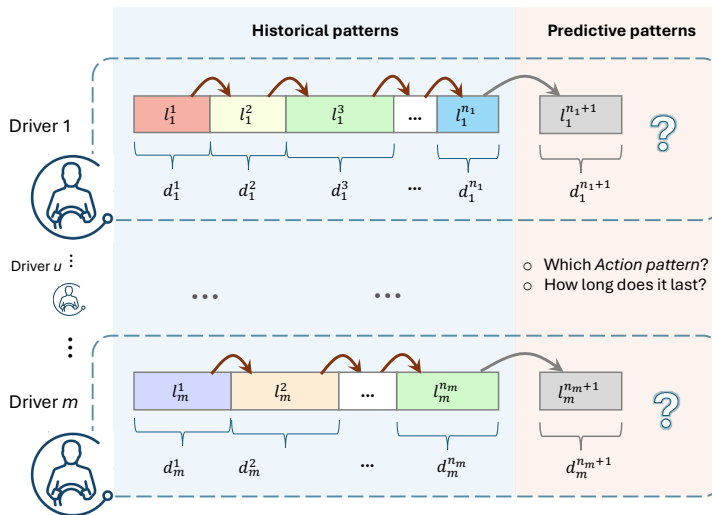
The clarification of action-related concepts is shown in Figure 6.1. *Action phases* encompass multiple driving variables, leading to a rapid expansion in the number of its categories as more variables are considered. For instance, considering four variables results in 3^4 categories, where ‘3’ is the number of *action trends*. Although capturing nuanced driving characteristics, the large number of *Action phases* types can complicate the interpretation of driving behaviour. For example, *Action phase* (I, I, D) - representing velocity increasing, acceleration increasing, and time headway decreasing - may exhibit similar observable behaviour to *Action phase* (I, S, D) where ‘S’ indicates acceleration stable at a positive value. To this end, *Action patterns* are introduced to consolidate similar *Action phases* into a small set of driving patterns that capture group-specific driving characteristics. The calibration of *Action patterns* is performed using an unsupervised learning manner, which involves *Action phase* standardisation, feature selection, and clustering analysis, more details about the calibration process can be found in [58]. Six distinct *Action patterns* are calibrated, which are interpreted as “Speed up”, “Slow down”, “Hold speed”, “Catch up”, “Fall behind”, and “Follow behind”. Each pattern represents a unique aspect of driver stimuli-response during driving.

In our previous work [58], we observed that time headway (T) and velocity (v) exhibit

higher importance than acceleration (a) and speed difference (Δv) in distinguishing *Action patterns*, highlighting their critical roles in characterising driving dynamics. These findings are employed to develop rules for labelling driving trajectories in *Action pattern* prediction. For example, driving segments with time headway increasing (T -I) indicate the ego vehicles are away from their leading vehicles, labelled as “Fall behind”. Conversely, a decreasing trend in T is labelled as a “Catch up” pattern. When the *action trend* of T is keeping stable (S), the focus shifts to the second most important variable, i.e., velocity (v). The rule-based *Action pattern* labelling method is detailed in Table 6.1. Consequently, driving trajectories are segmented and labelled as sequential *Action patterns*, providing input data for the pattern-based modelling task in this study.

Table 6.1: Rule-based *Action pattern* labelling

Variable	<i>action trend</i>					
T	I	D	S	S	S	S
v			I	D	S	S
a					I&D	S
<i>Action pattern</i>	Fall behind	Catch up	Speed up	Slow down	Follow behind	Hold speed

Figure 6.2: An overview of *Action pattern*-based prediction.

➤ Pattern-based driving behaviour modelling

The pattern-based approach to modelling driving behaviour is illustrated in Figure 6.2. For each driver, the *Action pattern* modelling includes two tasks: (i) model the characteristics of sequential transitions of *Action patterns* and their corresponding durations, and (ii) predict

the pattern label and the corresponding duration for the subsequent *Action phase*. Let $l_m^{n_m}$ denote the pattern label of the n_m -th *Action phase* from driver m , and $d_m^{n_m}$ represents the corresponding duration of this pattern. $\{l_m^1, l_m^2, \dots, l_m^{n_m}\}$ and $\{d_m^1, d_m^2, \dots, d_m^{n_m}\}$ signify historical pattern names and corresponding pattern durations of driver m , respectively. Then, $l_m^{n_m+1}$ and $d_m^{n_m+1}$ are pattern label and duration of the $n_m + 1$ -th *Action phase* that needs to be predicted.

6.2.2 Knowledge-enhanced attention LSTM models

The Long Short-Term Memory (LSTM) model serves as the foundational architecture for pattern-based driving behaviour modelling, as it effectively captures long-range dependencies and temporal patterns within sequential data [201], making it well-suited for time-series driving behaviour modelling. Attention-based LSTM (ALSTM) models further enhance the basic LSTM framework by incorporating an attention mechanism, which has shown substantial efficacy in vehicle trajectory prediction [202]. The attention mechanism dynamically focuses on the most relevant portions of input sequences, allowing the LSTM model to prioritise critical time steps and better capture driving characteristics, thereby improving prediction accuracy. In this study, we propose a Knowledge-Enhanced Attention LSTM (KE-ALSTM) model for pattern-based driving behaviour modelling, as illustrated in Figure 6.3.

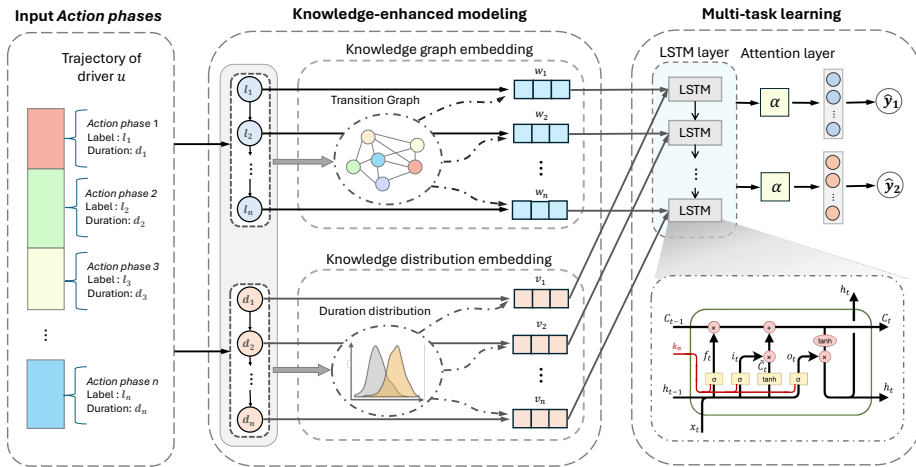


Figure 6.3: Overview of the knowledge-enhanced pattern behaviour modelling framework.

The KE-ALSTM model addresses the multi-task problem of modelling both Action pattern labels and their durations, and leverages both *Action pattern* transition and duration knowledge to enhance prediction accuracy. To achieve this, pattern label sequences and duration sequences are first extracted from input *Action phases*. These sequences are then embedded into equal-length representations, providing a consistent input format that allows the attention LSTM model to effectively handle *Action pattern* data. To enhance the embedding process, a knowledge-enhanced modelling approach is introduced.

This approach encodes *Action pattern* transitions as graph knowledge to represent their interrelationships and models pattern durations as distribution-based knowledge to augment feature representations. By incorporating this domain-specific knowledge, the model gains a deeper understanding of the underlying characteristics of driving patterns, thereby improving prediction accuracy and robustness. The modelling tasks for pattern transitions and durations share common LSTM layers to leverage shared temporal features, but utilise different attention layers and separate output layers to specialise in each task. Details of the knowledge-enhanced embedding and multi-task learning process are presented below.

➤ Driving knowledge modelling

Knowledge-enhanced models incorporate external knowledge, such as domain-specific rules, and expert insights, into machine learning (ML) process, providing additional context or information to improve ML model performance. These models have been increasingly explored to enhance both the accuracy and interpretability of deep learning-based vehicle trajectory prediction (DL-VTP) [200]. Driving-related knowledge can be derived from various sources and extracted using various techniques. For example, traditional ITS and AV research often utilises classic driving features - such as speed difference, relative distance, and time to collision (TTC) - to improve LSTM-based feature encoding for DL-VTP [203]. In this study, statistical methods are employed to encode *Action pattern* transitions as graph knowledge and *Action pattern* duration characteristics as distribution knowledge, enriching the feature representation for more effective driving behaviour modelling.

Action pattern transition knowledge modelling:

Graph-based knowledge represents information as a directed graph $G = (V, E)$, where V is a set of nodes and E is a set of edges that define the relationships between the nodes. In the context of pattern-based driving behaviour modelling, V represents the set of pattern names, and V represents the set of nodes (i.e., pattern names) transitions among these *Action patterns*. The transitions are modelled using a Markov chain, where the transition probabilities are calculated based on the observed frequencies of transitions among *Action patterns*. For a given *Action pattern* P_i at *Action phase* n , each edge $e_{ij} \in E$ from node P_i to node P_j is assigned a weight equal to \Pr which is the probability of transitioning to pattern P_j in the subsequent *Action phase* $n + 1$. \Pr is calculated by Equation 6.1.

$$\Pr(P_{n+1} = P_j | P_n = P_i) = \frac{\psi(P_{n+1} = P_j | P_n = P_i)}{\sum_{k=1}^6 \psi(P_{n+1} = P_k | P_n = P_i)} \quad (6.1)$$

where $\psi(P_{n+1} = P_j | P_n = P_i)$ denotes the frequency of transitions from pattern P_i to pattern P_j in the dataset.

Action pattern duration knowledge modelling:

Distribution-based knowledge captures the statistical properties of the driving trajectories. In this study, durations of *Action patterns* are modelled as distribution knowledge to facilitate the pattern-based prediction. We employ Kernel Density Estimation (KDE), a non-parametric method, to estimate the duration distribution for each *Action*

pattern. KDE offers flexibility by avoiding assumptions about the underlying distribution, which is particularly advantageous given the variability and lack of prior knowledge about duration characteristics across different driving patterns. Compared to parametric approaches such as Gaussian Mixture Models (GMMs), KDE is simpler to implement and does not require model selection procedures (e.g., determining the number of components), making it a practical and effective choice for integrating duration priors into the DL model. Another important consideration is the practicality. The simplicity and ease of implementation make KDE especially suitable for downstream applications, such as integrating into subsequent modelling and simulation research.

Let t_i be the observed duration of *Action pattern* P_i , where $i \in \{0, 1, 2, 3, 4, 5\}$. The KDE for the duration distribution $f_i(t)$ is given by:

$$\hat{f}_i(t) = \frac{1}{zh} \sum_{j=1}^z K\left(\frac{t-t_j}{h}\right) \quad (6.2)$$

where z is the number of duration observations, h is the bandwidth parameter, and K is the kernel function, commonly a Gaussian kernel:

$$k(\mu) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2}\right) \quad (6.3)$$

Accordingly, key parameters, the mean (μ), variance (σ^2), skewness (γ_1), and kurtosis (γ_2) are used as knowledge features for each pattern and are calculated as follows:

$$\begin{cases} \mu = \int_{-\infty}^{+\infty} t \hat{f}_i(t) dt \\ \sigma^2 = \int_{-\infty}^{+\infty} (t-\mu)^2 \hat{f}_i(t) dt \\ \gamma_1 = \frac{\int_{-\infty}^{+\infty} (t-\mu)^3 \hat{f}_i(t) dt}{\sigma^3} \\ \gamma_2 = \frac{\int_{-\infty}^{+\infty} (t-\mu)^4 \hat{f}_i(t) dt}{\sigma^4} - 3 \end{cases} \quad (6.4)$$

where Dr_i is the set of distribution knowledge for the specific pattern i , which is represented as:

$$Dr_i = [\mu \ \sigma^2 \ \gamma_1 \ \gamma_2]_i \quad (6.5)$$

where $i \in \{0, 1, 2, 3, 4, 5\}$, indicating six *Action patterns*.

➤ The structure of the proposed KE-ALSTM model

Figure 6.4 illustrates the structure of the KE_{P+D} -ALSTM model, including sequence feature fusion, pattern knowledge integration, LSTM prediction, duration knowledge integration, attention mechanism and prediction results generation. The pattern label recognition and pattern duration prediction tasks share a common feature extraction module, while each task has its own attention module. This design ensures that the shared feature extraction module learns representations that are useful for both tasks, while the task-specific attention modules focus on capturing information relevant to their respective objectives. Although the attention layers are applied to each task independently, both receive inputs from the

same shared LSTM outputs. This structure allows the model to capture joint temporal dependencies while enabling each task to learn its specific patterns. In this way, the two tasks are not isolated but interact implicitly through shared sequence representations.

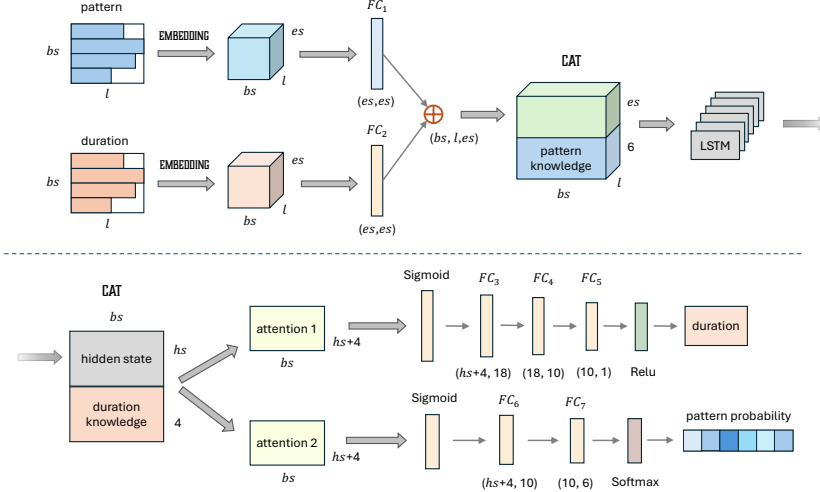


Figure 6.4: The structure of the knowledge-enhanced attention LSTM (KE-ALSTM) model.

Sequence feature fusion: The inputs to the prediction model consist of both pattern label sequences and pattern duration sequences. To address the varying lengths of driver sequences, embedding layers are employed to map these inputs into fixed-length dense vectors, reducing computational complexity while preserving relational information among patterns.

The dimensional changes in this embedding process are described as:

$$\begin{cases} \bar{x}_{(bs,l,es)}^P = \text{Embedding}(x_{(bs,l)}^P) \\ \bar{x}_{(bs,l,es)}^D = \text{Embedding}(x_{(bs,l)}^D) \end{cases} \quad (6.6)$$

where $x_{(bs,l)}^P$ and $x_{(bs,l)}^D$ represent the input of pattern label sequences and pattern duration sequences before embedding, respectively. $\bar{x}_{(bs,l,es)}^P$ and $\bar{x}_{(bs,l,es)}^D$ denote the resulting embeddings, with bs as batch size, l as sequence length, and es as embedding size.

Fully connected layers are then applied for feature extraction, and the resulting pattern label and duration features are merged via element-wise summation:

$$\tilde{x}_{(bs,l,es)} = \text{FC}_1(\bar{x}_{(bs,l,es)}^P) \oplus \text{FC}_2(\bar{x}_{(bs,l,es)}^D) \quad (6.7)$$

Pattern knowledge integration: After feature merging, *Action pattern* transition knowledge is integrated via concatenation (CAT) to preserve the structure of the original features while incorporating external behavioural information. The CAT approach treats the learned features and the appended knowledge as distinct yet jointly learnable

components, enabling flexible adaptation in subsequent layers [204]. Unlike element-wise and attention-based methods, concatenation maintains the interpretability and modularity of each input source.

$$\begin{aligned} \bar{x}_{(b,s,l,es+p)} = \text{CAT}(\bar{x}_{(b,s,l,es)}, \Pr(P_{n+1} = P_j | P_n = P_i)_{(bz,p)}, \\ \text{axis} = 2) \end{aligned} \quad (6.8)$$

where p represents the total number of *Action pattern* types, set to 6 in this study based on empirical observations.

LSTM prediction: The integrated features are fed into the LSTM layers for time-series processing to model sequential *Action pattern* data effectively:

$$h_{n(b,s,hs)} = \text{LSTM}(x_{n(b,s,l,es+p)}, h_{n-1(b,s,hs)}) \quad (6.9)$$

where h_n and x_n denotes the hidden state and input at phase step n , respectively.

Duration knowledge integration: To further enhance prediction accuracy, the final hidden state from the LSTM is augmented with distribution-based knowledge representing Action pattern durations. This augmentation allows the model to account for variations in pattern duration effectively.

$$\bar{h}_{(b,s,hs+4)} = \text{CAT}(h_{(b,s,hs)}, Dr_{(b,s,4)}, \text{axis} = 1) \quad (6.10)$$

where 4 corresponds to the parameters in the KDE distribution, representing the duration characteristics.

Attention mechanism: The augmented features are subsequently passed through two separate attention layers, which process pattern transitions and duration information independently. This design ensures that the model can focus on the most critical aspects of each task. The features are then scored by a linear layer using a linear layer:

$$e_n^{\{1,2\}} = x_n W^{\{1,2\}} + b^{\{1,2\}} \quad (6.11)$$

where W and b are learnable parameters during the model training process.

The attention mechanism computes a set of weights α_n to quantify the importance of sequential *Action patterns* at step n :

$$\alpha_n^{\{1,2\}} = \frac{\exp(e_n^{\{1,2\}})}{\sum_{k=1}^N \exp(e_k^{\{1,2\}})} \quad (6.12)$$

The output of the attention mechanisms is described as:

$$\hat{x}_n^{\{1,2\}} = x_n \odot \alpha_n^{\{1,2\}} \quad (6.13)$$

Thereafter, a *Sigmoid* activation function is applied to prevent extreme output values:

$$\hat{x}_n^{\{1,2\}} = \text{Sigmoid}(\hat{x}_n^{\{1,2\}}) \quad (6.14)$$

Prediction results generation: To facilitate separate predictions for pattern labels and durations, different feature pyramids are utilised to extract features of varying dimensions. Three fully connected layers are used for pattern duration modelling, with Relu activation function to rectify negative outputs:

$$\text{Duration} = \text{Relu}(\text{FC}_5(\text{FC}_4(\text{FC}_3(\hat{x}_n^1)))) \quad (6.15)$$

For pattern label prediction, a Softmax activation function is applied to generate pattern probability after two fully connected layers:

$$\text{Pattern probability} = \text{Softmax}(\text{FC}_7(\text{FC}_6(\hat{x}_n^2))) \quad (6.16)$$

Loss function To optimise the KE_{P+D} -ALSTM model, an adaptive combined-task-wise loss function \mathcal{L} is employed:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{pattern}} + \lambda_2 \mathcal{L}_{\text{duration}} \quad (6.17)$$

where λ_1 and λ_2 are hyperparameters to balance the pattern label and pattern duration prediction tasks as well as order magnitude differences. These two parameters are determined by conducting sensitivity analyses on multiple combinations of λ_1 and λ_2 values, which can be found in Section 6.4.1.

The *Action pattern* label prediction is formulated as a classification problem, where the cross-entropy loss $\mathcal{L}_{\text{pattern}}$ is defined as:

$$\mathcal{L}_{\text{pattern}} = - \sum_{i=1}^C y_i \cdot \log(\hat{y}_i) \quad (6.18)$$

Here, C represents the total number of *Action patterns*. y is a binary indicator (0 or 1) that indicates whether class i is the correct classification for a given observation o , and \hat{y} is the predicted probability of the observation being assigned to class j .

For pattern duration prediction, a confidence interval-based loss function is introduced to penalise deviations from true values. This approach helps accelerate model convergence by imposing heavier penalties on larger errors, as described in Equation 6.19 [205]. The loss weights ω_1 , ω_2 , ω_3 , and ω_4 are assigned to different ranges of standard deviation (σ) to reflect the severity of the error.

$$\mathcal{L}_{\text{duration}} = \begin{cases} \omega_1 \left| \frac{\hat{y} - y}{\sigma} \right|, & \text{if } 0 < \left| \frac{\hat{y} - y}{\sigma} \right| \leq 1 \\ \omega_2 \left| \frac{\hat{y} - y}{\sigma} \right|, & \text{if } 1 < \left| \frac{\hat{y} - y}{\sigma} \right| \leq 2 \\ \omega_3 \left| \frac{\hat{y} - y}{\sigma} \right|, & \text{if } 2 < \left| \frac{\hat{y} - y}{\sigma} \right| \leq 3 \\ \omega_4 \left| \frac{\hat{y} - y}{\sigma} \right|, & \text{otherwise} \end{cases} \quad (6.19)$$

6.3 Experiments

In this section, we introduce the experiment of pattern-based driving behaviour modelling, including the preparation of driving pattern data, the representation of *Action pattern* knowledge, and parameter settings for the knowledge-enhanced attention LSTM model.

6.3.1 Data preparation

The Lyft level-5 open naturalistic dataset [182] is applied in this study. More than 42k HV-following-HV (HV-HV) pairs with a total driving distance of over 150k km are selected, assessed, and enhanced in similar environments in this dataset. We exclude car-following pairs with $v = 0\text{m/s}$ to remove stop-and-go behaviour. Driving trajectories shorter than 20 seconds are removed to ensure sufficient reflection of longitudinal driving behaviours. 3000 car-following pairs are selected for *Action pattern* prediction according to the rule-based labelling method presented in Table 6.1. Examples of labelled driving trajectories are visualised in Figure 6.5. In the purple area, the increase of T exceeds a certain threshold, thus is labelled as a “Fall behind” pattern. When the change in T is minimal, the primary focus shifts to the velocity, as shown in the red line. The increase in velocity is labelled as a “Speed up” pattern, while the decrease in velocity is identified as a “Slow down” pattern, as illustrated in the red and yellow areas, respectively. In this way, all the drivers’ trajectories are labelled as *Action pattern* sequences.

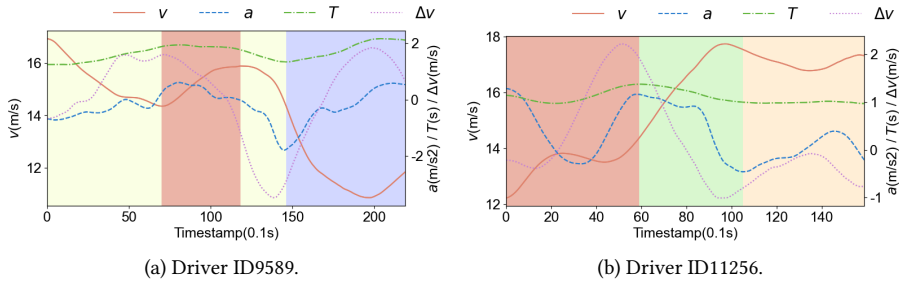


Figure 6.5: Examples of labelled driving trajectories.

6.3.2 Knowledge representation

The transition probabilities between *Action patterns* are calculated using Equation 6.1, with the results shown in Figure 6.6. The arrows indicate the direction of *Action pattern* transitions, and the thickness of the lines represents the corresponding probabilities. Notice that the probability of pattern transitions varies greatly. For instance, some transitions are highly probable, such as “Follow behind” to “Follow behind” with a probability of 0.39, whereas others, such as “Fall behind” to “Hold speed” with a probability of 0.08, occur less frequently. These findings align with real-world observations, where continuous following is common, while holding speed after falling behind is less preferred. Additionally, both “Fall behind” and “Slow down” patterns show a strong tendency to transition to a “Speed up” pattern, reflecting the common behaviour of vehicles accelerating after decelerating or lagging behind. Similarly, a “Catch up” pattern frequently transitions to a “Slow down” pattern, suggesting that vehicles decelerate to maintain a safe distance or speed after catching up with the leading vehicles. The accumulated probability of a given *Action pattern* being transitioned into from all six *Action patterns* is calculated using $\text{Pr}_{b_t}(P_{n+1} =$

$P_j) = \frac{\sum_{i=1}^6 \Pr(P_{n+1}=P_j|P_n=P_i)}{6}$, with results presented in Table 6.2. The results show that the “Follow behind” pattern has the highest cumulative probability, indicating that vehicles often maintain a stable car-following behaviour during driving. These transition probabilities are encoded as graph-based knowledge to enhance the KE-ALSTM model’s ability to predict future *Action patterns*.




Figure 6.6: Transition probabilities from pattern to pattern.

Table 6.2: The probability of *Action patterns* being transitioned

<i>Action pattern j</i>	Fall behind	Catch up	Speed up	Slow down	Follow behind	Hold speed
$\Pr_{b_t}(P_{n+1} = P_j)$	0.1568	0.1379	0.1768	0.1172	0.2849	0.1264

Four statistical measures - mean (μ), variance (σ^2), skewness (γ_1), and kurtosis (γ_2) - capture the KDE distribution from different aspects [206], capturing various characteristics of duration across different *Action patterns*. The mean represents the central value and typical duration of each pattern, revealing which patterns tend to last longer or shorter. Variance quantifies the dispersion of durations around the mean, while skewness measures the asymmetry of the distribution. Kurtosis assesses the tail behaviour of the distribution, indicating the sharpness of the distribution peak and the presence of outliers. The distribution of pattern duration characteristics for each *Action pattern* is shown in Figure 6.7 and Table 6.3. Notice that the “Fall behind” and “Catch up” patterns exhibit similar distributions, with means around 54 and comparable values for variance, skewness, and kurtosis. In contrast, the “Speed up” and “Slow down” patterns have shorter average durations, with means around 47. Notably, the “Slow down” has a high kurtosis (3.82), suggesting a sharper peak and more extreme values compared to “Speed up” (kurtosis of 2.47). The “Follow behind” pattern displays significant right skewness (1.43) and elevated

kurtosis (2.73), indicating more outliers and a sharper peak. Meanwhile, the “Hold speed” pattern has the shortest mean duration (40.12), the lowest variance (663.55), and the highest skewness (1.67) and kurtosis (4.20). This reflects a higher variability in *Action pattern* durations, consistent with real-world driving behaviour, where maintaining a constant speed is often interrupted by frequent acceleration or deceleration. Note that despite the variations across all distributions, the differences are minimal. These statistical descriptors of the KDE distributions offer valuable insights into variability, outliers, and the typical range of values for each pattern. This distribution information is encoded into the KE-ALSTM model to improve its understanding of pattern duration characteristics, thereby enhancing its predictive performance.

Table 6.3: Parameters of KED on *Action patterns*

<i>Action pattern</i>	μ	σ^2	γ_1	γ_2
Fall behind	53.77	1210.97	1.08	1.87
Catch up	54.36	1152.49	0.9	0.93
Speed up	47.84	893.93	1.38	2.47
Slow down	47.2	906.24	1.59	3.82
Follow behind	51.43	1031.59	1.43	2.73
Hold speed	40.12	663.55	1.67	4.20




Figure 6.7: Kernel Density Estimation (KDE) of *Action patterns*.

6.3.3 Model parameter settings

In the KE_{p+D} -ALSTM model, six *Action patterns* are numerically encoded from 0 to 5, and the precision of duration is 0.1 seconds. Given the variability in the number of patterns across different drivers’ trajectories, all input sequences are reshaped to a uniform length

to meet the batch processing requirements of LSTM models. To effectively handle these variable-length pattern sequences without introducing distortion from padding, we employ the PackedSequence technique, following the implementation in [192]. This approach involves two key steps: (i) sorting input sequences by length in descending order, and (ii) converting them into a packed representation that includes both the sequence data and metadata indicating the actual lengths of each sequence. During training, the LSTM processes only the valid time steps in each sequence, skipping padded elements entirely. This not only preserves the integrity of the temporal structure but also significantly improves computational efficiency by reducing unnecessary operations on padded elements. As a result, the model can learn more effectively from sequences of different lengths while minimising computational overhead. After standardising the sequence lengths, the *Action pattern* dataset is split into 60% for training, 20% for validation, and 20% for testing. The training set is exclusively employed during the model’s initial learning phase. The validation set is used throughout the development process for hyperparameter tuning (e.g., number of LSTM layers, hidden layers, etc.) and for implementing early stopping to prevent overfitting. The test set remains entirely unseen during both training and validation. Evaluating the final tuned model on test set provides an unbiased assessment of its generalisation capability to unseen data, which is essential for demonstrating the model’s practical applicability.

The hyperparameter settings for neural networks are summarised in Table 6.4. The loss weights λ_1 and λ_2 are initially set to 1 and are tuned through sensitivity analysis to determine their optimal values. Specifically, the selected values for λ_1 range from 0.1 to 10, and for λ_2 from 0.001 to 1, covering a wide range of absolute values and relative weightings. The ratio $\lambda_1:\lambda_2$ therefore spans a range from 0.1 to 10,000, allowing us to evaluate both symmetric and asymmetric task prioritisation scenarios. The KE_{P+D} -ALSTM model is trained for 100 epochs with a batch size of 32 and a learning rate of 0.008 to ensure convergence. The model employs six LSTM layers with a hidden size of 25 and a dropout rate of 0.5 to mitigate overfitting. All experiments are conducted on a workstation equipped with an NVIDIA A10 GPU, providing the computational resources necessary for training deep learning models efficiently.

Table 6.4: Parameter settings

Hyperparameter	Value	General Parameter	Value
Batch size	32	λ_1	1
Learning rate	0.008	λ_2	1
Training epochs	60	ω_1	1
Dropout rate	0.5	ω_2	1.5
Number of LSTM layers	6	ω_3	2
Hidden size of LSTM	25	ω_4	2.5

Several benchmark models and ablation variants are employed to evaluate the effectiveness of the attention mechanism and knowledge-enhanced approaches. These models include: a standard LSTM model, an attention-based LSTM model (ALSTM), an ALSTM model enhanced with *Action pattern* transition knowledge (KE_P -ALSTM), an

ALSTM model incorporating *Action pattern* duration knowledge (KE_D-ALSTM), a LSTM model integrating both *Action pattern* transition and *Action pattern* duration knowledge (KE_{P+D}-LSTM), and a Transformer model enhanced with both *Action pattern* transition and *Action pattern* duration knowledge (KE_{P+D}-Transformer). All models are tailored for pattern-based driving behaviour modelling and share the same hyperparameter values to ensure a fair performance comparison.

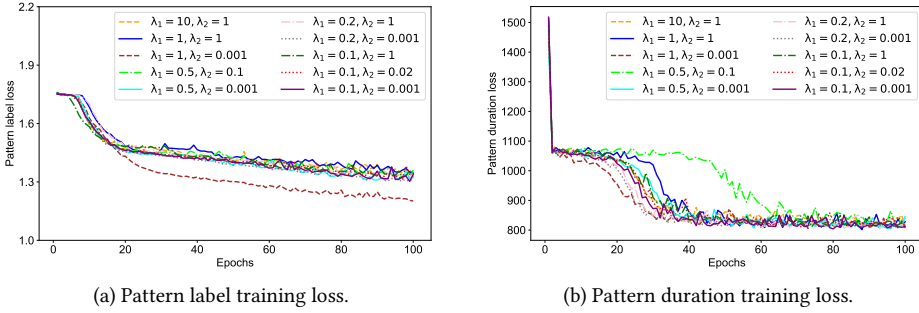
6.4 Results and Discussions

In this section, we present the results of pattern-based driving behaviour modelling using the proposed knowledge-enhanced deep learning models. This starts with the sensitivity analyses for choosing the optimal model hyperparameters, followed by a comparative analysis with benchmark models and ablation variants to demonstrate the effectiveness of incorporating driving knowledge in improving deep learning model performance.

6.4.1 Hyperparameter tuning

To determine the optimal values for the loss weighting parameters λ_1 and λ_2 , we tested a wide range of combinations where λ_1 ranged from 0.1 to 10 and λ_2 from 0.001 to 1. The evaluated configurations include both balanced weights such as $\lambda_1 = \lambda_2 = 1$ and asymmetric settings that emphasise one task over the other, including pattern-prioritised settings (e.g., $\lambda_1 = 10$, $\lambda_2 = 1$), duration-prioritised settings (e.g., $\lambda_1 = 0.1$, $\lambda_2 = 1$), fine-grained adjustments with small values (e.g., $\lambda_2 = 0.001$). The results of the training loss across different combinations of λ_1 and λ_2 are illustrated in Figure 6.8. Both the pattern label loss and pattern duration loss exhibit decreasing trends under these hyperparameter settings, see Figure 6.8a and Figure 6.8b, respectively. While the pattern label loss converges similarly under most configurations, the setting of $\lambda_1 = 1$ and $\lambda_2 = 0.001$ achieves the lowest final loss of pattern prediction and fastest convergence of duration prediction. In contrast, when $\lambda_1 = 0.5$ and $\lambda_2 = 0.1$, the pattern duration loss plateaus and shows minimal reduction within the first 40 epochs. Overall, the $\lambda_1 = 1$, $\lambda_2 = 0.001$ configuration offers the most balanced and efficient learning across both tasks, enabling an empirically supported trade-off between the two objectives. This thorough analysis ensures that the optimal trade-off between the two learning objectives can be empirically determined.

A further sensitivity analysis on the hyperparameters λ_1 and λ_2 was conducted to assess their influence on model performance using the validation set. As presented in Table 6.5, the results reveal that these parameters significantly affect the model's generalisation capability. Among all tested configurations, $\lambda_1 = 1$ and $\lambda_2 = 0.001$ achieves the best overall performance, with a pattern label accuracy of 90.67% and high duration prediction accuracy. This indicates a well-balanced contribution from both loss components. In contrast, configurations such as $\lambda_1 = 10$, $\lambda_2 = 1$ and $\lambda_1 = 0.5$, $\lambda_2 = 0.001$ resulted in reduced performance, especially under stricter duration thresholds, pointing to issues such as overfitting or imbalanced learning. These findings highlight the importance of carefully tuning loss weights based on validation performance. Consequently, the configuration $\lambda_1 = 1$ and $\lambda_2 = 0.001$ was selected as final parameters for prediction models due to its strong convergence behaviour and balanced predictive performance for both tasks.

Figure 6.8: Effect of hyperparameters λ_1 and λ_2 on training loss.Table 6.5: Sensitivity analysis of hyperparameters λ_1 and λ_2 on predictive accuracy.

λ_1	λ_2	Accuracy (%)			
		pattern	duration (3σ)	duration (2σ)	duration (σ)
10	1	76.17	81.00	57.67	19.00
1	1	62.50	98.83	93.00	75.17
1	0.001	90.67	98.67	92.50	76.50
0.5	0.1	76.00	97.50	87.33	66.67
0.5	0.001	90.67	81.00	57.67	19.00
0.2	1	76.50	98.50	90.50	72.33
0.2	0.001	76.67	98.50	90.83	72.83
0.1	1	76.67	96.50	86.50	63.50
0.1	0.02	73.83	97.33	86.67	65.33
0.1	0.001	76.67	97.83	87.50	67.83

To verify that the KE_{P+D} -ALSTM model maintains a balanced prioritisation of both tasks during the training process, we present a comparison of detailed training loss curves for both pattern label and duration tasks. As shown in Figure 6.9, both losses decrease smoothly and stably over epochs, indicating that neither task dominates the learning process. These results demonstrate that the model effectively supports both objectives without requiring dynamic reweighting of the loss components.

Other key hyperparameters, including the number of LSTM layers and the hidden size, were tuned to optimise model performance. As presented in Table 6.6, various configurations were evaluated using the validation set, combining 4, 6, and 8 LSTM layers with hidden sizes of 16, 25, and 32. The results indicate that a configuration of six LSTM layers with a hidden size of 25 yields the best overall performance across both prediction tasks. This setting achieves the highest pattern label prediction accuracy (90.67%) and maintains strong duration prediction accuracy, reaching 92.50% at the 2σ threshold and 76.50% at the 1σ setting. Increasing the hidden size to 32 does not improve performance and results in a slight drop in duration prediction accuracy, suggesting potential overfitting or unnecessary model complexity. Likewise, increasing the number of LSTM layers to 8 negatively impacts




Figure 6.9: Training losses of KE_{P+D} -ALSTM model ($\lambda_1 = 1, \lambda_2 = 0.001$)

generalisation performance, particularly for duration prediction, which drops to 57.67% and 19.00% under the 2σ and 1σ thresholds, respectively. On the other hand, reducing the number of layers or using smaller hidden sizes to some extent improves generalisation for duration prediction but compromises pattern label accuracy. These findings demonstrate that the selected configuration with six LSTM layers with a hidden size of 25 offers the most effective trade-off between model complexity, learning capacity, and generalisation across both prediction tasks.

Table 6.6: Sensitivity analysis of LSTM parameters on predictive accuracy.

Number of LSTM layers	Hidden size of LSTM	Accuracy (%)			
		pattern	duration (3σ)	duration (2σ)	duration (σ)
6	16	76.67	97.83	93.00	76.00
6	25	90.67	98.67	92.50	76.50
6	32	76.67	97.50	87.33	67.17
4	25	79.67	98.67	90.67	74.33
8	25	90.33	81.00	57.67	19.00

6.4.2 Convergence analysis

The convergence behaviour of all models during training is evaluated using loss curves, which reflect the discrepancy between predicted values and ground truth and provide insight into learning efficiency and potential overfitting. Figure 6.10 illustrates the training loss curves for the proposed KE_{P+D} -ALSTM model, its ablation variants (KE_P -ALSTM, KE_D -ALSTM, KE_{P+D} -LSTM), and benchmark models (LSTM, ALSTM, Transformer-based).

All models demonstrate convergence after approximately 100 epochs, with the most rapid loss reductions occurring within the first 30 epochs. Among the benchmark models, LSTM and ALSTM converge to relatively higher loss values (2.5–2.6), indicating limited ability to capture complex driving behaviours. The addition of the attention mechanism in ALSTM does not provide substantial convergence benefits over standard LSTM, suggesting

that attention alone is insufficient for the multi-task pattern-based learning objective. The ablation models offer deeper insights into the effect of knowledge integration. The KE_p -ALSTM model, which incorporates only pattern label transition knowledge, shows moderate improvement in convergence rate but still stabilises at a higher loss level. In contrast, KE_D -ALSTM, which integrates duration knowledge, exhibits a notably steeper decline and lower final loss, indicating that duration knowledge contributes more substantially to training efficiency than pattern knowledge alone. Interestingly, KE_{p+D} -LSTM, which integrates both knowledge types but lacks the attention mechanism, fails to fully capitalise on these enhancements, converging similarly to the baseline models. The proposed KE_{p+D} -ALSTM model achieves the best convergence among all models, with the steepest decline and the lowest final training loss. This demonstrates the synergistic effect of integrating both pattern label and duration knowledge with an attention-enhanced architecture, enabling the model to effectively learn multi-task objectives. Additionally, the KE_{p+D} -Transformer model shows a higher final training loss compared to the LSTM-based models, indicating less effective convergence. This can be attributed to the Transformer's high data requirements and lack of inherent sequential bias.




Figure 6.10: Total training loss.

Overall, the convergence analyses suggest that incorporating both pattern label and duration knowledge enables the model to better capture the sequential characteristics of *Action patterns*, thereby enhancing prediction performance. Additionally, in our modelling task with relatively short driving sequences and a moderately sized dataset, LSTM-based models are better suited to capturing temporal dependencies and pattern characteristics, while the Transformer struggles to generalise effectively.

6.4.3 Accuracy analysis

The accuracy is computed as the ratio of correct predictions to the total number of predictions. For pattern label prediction, a prediction is considered correct when the predicted label exactly matches the ground truth. For pattern duration prediction, accuracy is assessed using confidence intervals, reflecting acceptable error margins in real-world

applications. Given the temporal resolution of *Action patterns* (0.1 seconds), predictions within one standard deviation (σ) of the ground truth are considered highly precise, while predictions within two or three standard deviations (2σ or 3σ) are deemed acceptable.

Table 6.7: Comparison of predictive accuracy

Model	Accuracy (%)			
	pattern	duration (3σ)	duration (2σ)	duration (σ)
LSTM	38.83	97.67	89.33	71.21
ALSTM	27.93	98.50	90.21	72.87
KE _P -ALSTM	76.67	98.00	89.67	70.67
KE _D -ALSTM	52.83	98.83	93.97	73.17
KE _{P+D} -LSTM	21.83	97.67	92.67	74.83
KE _{P+D} -ALSTM	90.67	98.67	92.50	76.50
KE _{P+D} -Transformer	69.83	97.67	89.83	70.50

The results of *Action pattern* prediction across all evaluated models are summarised in Table 6.7. The baseline LSTM and ALSTM models exhibit limited capability in pattern label prediction, achieving accuracies of 38.83% and 27.93%, respectively. While both models perform reasonably well in pattern duration prediction under relaxed tolerance levels, achieving over 97% accuracy at the 3σ threshold, their precision declines at stricter margins, with only 71.21% (LSTM) and 72.87% (ALSTM) accuracy at σ setting. The introduction of behavioural knowledge significantly improves model performance. KE_P-ALSTM, which incorporates pattern transition knowledge, increases pattern label prediction to 76.67%, a 174.51% relative improvement over ALSTM. However, the KE_P-ALSTM does not achieve enhancement on pattern duration prediction. KE_D-ALSTM, which integrates duration knowledge, yields the highest duration prediction accuracy at the less strict threshold settings, reaching up to 98.83% at 3σ margin and 93.97% at 2σ margin. KE_{P+D}-ALSTM, which jointly integrates both transition and duration knowledge, achieves the highest pattern label prediction accuracy of 90.67%, representing an 18.26%–315.35% improvement over all baseline models. It also achieves the best pattern duration prediction at the strictest σ setting with an accuracy of 76.50%, which represents an improvement of 2.23% to 8.51% compared to other baseline models. Note that while the LSTM model achieves superior performance compared to the ALSTM model in pattern label prediction, the KE_{P+D}-ALSTM model outperforms the KE_{P+D}-LSTM model when both pattern transition and duration knowledge are incorporated. This result highlights the effectiveness of the attention mechanism in capturing both *Action pattern* transition and duration characteristics. The KE_{P+D}-Transformer achieves a pattern label prediction accuracy of 69.83%, outperforming the baseline models without knowledge enhancement. However, it does not exhibit superior performance in predicting pattern duration. This aligns with the earlier convergence analysis, suggesting that although the Transformer has strong representational capacity, it appears to underperform in data-scarce scenarios typical of driving behaviour sequences.

Overall, all the tested models demonstrate good performance under less strict duration thresholds (2σ and 3σ), with accuracies exceeding 97%. However, only the

knowledge-enhanced models, particularly KE_D -ALSTM and KE_{P+D} -ALSTM achieve high precision under stricter error settings. These findings underscore the importance of incorporating structured behavioural knowledge into sequential learning frameworks and highlight the KE_{P+D} -ALSTM model as a robust architecture for multi-task pattern-based driving behaviour modelling.

6.4.4 Case study

To further demonstrate the effectiveness of the proposed KE_{P+D} -ALSTM model in *Action pattern* modelling, we analyse the case of driver ID8432, whose driving pattern sequence is shown in Figure 6.11. The pattern sequence for this driver begins with a “Speed up” pattern lasting 3.5 seconds, followed by a “Follow behind” pattern lasting 4.3 seconds. The driver then alternates between “Speed up” and “Slow down” patterns, lasting 3.4 and 2.5 seconds, respectively. The KE_{P+D} -ALSTM model accurately predicts the next pattern, identifying “Catch up” as the most probable pattern, which aligns with the ground truth. Additionally, the model predicts the duration of the “Catch up” pattern as 4.235 seconds, showing a minimal error of 0.035s compared to the ground truth duration of 4.2 seconds. Table 6.8 presents case analyses of additional drivers. The numbers highlighted in red represent *Action pattern* labels and durations to be predicted (i.e., ground truth). Note that the model achieves high accuracy in predicting pattern labels for the majority of the selected drivers and maintains a duration prediction error within 0.6 seconds. The results highlight the effectiveness of the KE_{P+D} -ALSTM model in predicting various pattern labels and durations across driving pattern sequences of varying lengths, showcasing its robustness in modelling *Action patterns* and handling multi-task prediction effectively.




Figure 6.11: The predicted results of driver ID8432.

These results highlight that driving behaviour can be effectively modelled using driving patterns, providing both pattern labels and corresponding durations as outputs. The integration of additional driving pattern knowledge significantly enhances the efficiency and accuracy of deep learning-based prediction models, demonstrating the potential of knowledge-enhanced approaches for improving driving behaviour analysis.

Table 6.8: Predicted *Action patterns* of drivers

Driver ID	Pattern sequences	Predicted value	Error
7029	$l : [0, 0, 4, 1, 1]$	1	0
	$d : [47, 16, 37, 72, 47]$	42.34	0.466s
7224	$l : [3, 2, 1, 1]$	1	0
	$d : [70, 19, 83, 47]$	43.79	0.321s
7351	$l : [3, 1, 4, 4]$	4	0
	$d : [20, 41, 53, 45]$	39.17	0.583s
7491	$l : [2, 5, 5, 5]$	5	0
	$d : [39, 65, 17, 23]$	28.35	0.535s
7643	$l : [5, 0, 0, 3, 3]$	5	1
	$d : [53, 27, 20, 29, 30]$	32.57	0.257s
8382	$l : [2, 1, 4, 2, 2]$	2	0
	$d : [51, 24, 18, 53, 33]$	27.75	0.525s
8570	$l : [0, 5, 0]$	0	0
	$d : [81, 43, 35]$	39.07	0.407s
8682	$l : [0, 0, 1, 3, 3]$	3	0
	$d : [37, 17, 54, 74, 37]$	37.55	0.055s
8808	$l : [4, 3, 3, 2]$	2	0
	$d : [40, 38, 64, 37]$	31.41	0.559s
9081	$l : [4, 0, 2, 2]$	2	0
	$d : [47, 55, 37, 20]$	24.12	0.412s

* l - Action pattern label, d - Action pattern duration (0.1s).

* 0 - Follow behind, 1 - Slow down, 2 - Catch up, 3 - Speed up, 4 - Fall behind, 5 - Hold speed.

6.5 Discussions and Outlook

The proposed knowledge-enhanced deep learning models show significant effectiveness in predicting driving patterns and durations. Despite their advantages, there are limitations that point to promising directions for future research.

6.5.1 Data representation

Despite the structured design of driving pattern modelling, the dataset used in this study exhibits an imbalanced distribution among different *Action patterns*. For example, the “follow behind” pattern accounts for approximately 30.36% of all labelled patterns, while other patterns occur less frequently, ranging from 11.37%–16.61%. Moreover, the length of all *Action patterns* varies considerably, ranging from 1.1s to 13.4s, highlighting substantial temporal variability. This imbalance may introduce bias during training and compromise the model’s ability to accurately capture and predict minority behaviours. Future work could employ strategies such as data augmentation, reconstruction of loss functions, or advanced sampling techniques to achieve a more balanced representation of driving behaviours. Additionally, instead of the current Kernel Density Estimation (KDE) representation, alternative approaches, such as GMMs or deep generative models (e.g., VAEs), could be used to enhance pattern duration representation by modelling distributional uncertainty more flexibly, thereby improving robustness in duration prediction.

6.5.2 Model generalisability

The current model is trained on highway driving data, which inherently lacks diversity in driving contexts. This limitation may constrain the model's generalisability to more complex or heterogeneous environments, such as variations in road geometry, weather conditions, traffic congestion levels, and regional driving norms. Future extensions will focus on incorporating data from diverse driving environments and conditions, including urban, suburban, rural road types, and different traffic flow densities. Additionally, integrating contextual features such as lane configurations, surrounding vehicle interactions, visibility, and environmental variables can improve the model's adaptability and transferability. These enhancements would enable more accurate and context-aware behaviour prediction, facilitating deployment in diverse traffic scenarios.

6.5.3 Model applicability

While the KE-ALSTM model achieves strong performance in predicting behavioural patterns and their durations, it does not explicitly model fine-grained vehicle dynamics within each behavioural segment, such as instantaneous acceleration, deceleration, or lane-level positioning. This limits the model's direct applicability in tasks that require precise trajectory planning or low-level vehicle control, such as cooperative motion planning. Future direction lies in developing hierarchical model architectures, where combining the high-level pattern prediction with low-level vehicle dynamics models. Such integration would allow the model to infer both behavioural intent and the corresponding vehicle response, enabling more complete simulation, planning, and control in mixed-traffic and connected vehicle environments.

6.6 Conclusion

To capture the underlying characteristics of driving trajectories and model longitudinal driving behaviour more comprehensively, this study proposes a novel modelling approach based on driving pattern sequences and knowledge-enhanced deep learning models. Based on a rigorous driving pattern labelling process, a Knowledge-enhanced attention LSTM (KE-ALSTM) model is proposed, which integrates both *Action pattern* label and *Action pattern* duration knowledge, enabling the accurate prediction of driving patterns alongside their corresponding time lengths. Experimental evaluation using real-world datasets demonstrates that the KE-ALSTM model significantly outperforms benchmark models without knowledge enhancement in predicting *Action pattern* sequences. Specifically, the KE-ALSTM achieved accuracy improvements in pattern label prediction ranging from 18.26% to 315.35% and in pattern duration prediction by 2.23%–8.51% at a strict threshold. These results underscore the importance of incorporating additional driving behaviour knowledge to enhance the predictive capability and interpretability of deep learning models. The proposed model captures long-range temporal dependencies and provides robust predictions that support real-time applications such as behaviour-aware planning, traffic flow simulation, and cooperative driving systems. The explicit modelling of pattern duration, often inaccessible to rule-based inference, further enhances its value for downstream decision-making and improves the realism of traffic analysis and control.

7

A Pattern-based Framework for Modelling Driving Heterogeneity and Traffic Flow Simulation

The content of this chapter is based on

📖 Yao, X., Calvert, S. C., and Hoogendoorn, S. P. (2025). “A Pattern-based Framework for Modelling Driving Heterogeneity and Traffic Flow Simulation.” *Under review by a journal.*

This chapter introduces a novel pattern-based modelling and simulation framework to analyse the impacts of driving heterogeneity on traffic flow. A bi-level modelling approach is developed to model high-level behavioural pattern transitions and low-level vehicle dynamics, enabling the capture of both inter- and intra-driver variability in longitudinal driving behaviour in the micro-simulation. Evaluation on real-world data demonstrates its effectiveness in revealing how different levels of driving heterogeneity affect traffic safety, energy efficiency, and traffic stability, offering insights for advanced driver assistance systems and adaptive traffic management strategies.

7.1 Introduction

Heterogeneity in human longitudinal driving behaviour plays a critical role in determining traffic flow dynamics and performance [47]. Empirical evidence and simulation studies have shown that heterogeneity in driving behaviour has been associated with traffic phenomena and externalities, such as the emergence of traffic hysteresis loops [207, 208], increased collision risks [126, 209], and higher fuel consumption and emissions [128, 210]. Advancing our understanding of the mechanisms and manifestations of driving heterogeneity, along with a thorough investigation of how driving heterogeneity impacts traffic flow performance, are therefore essential for developing more robust and adaptive Intelligent Transportation Systems (ITS) that enhance traffic safety, efficiency, and environmental sustainability.

Despite its relevance, studies into the quantification and representation of real-world driving heterogeneity remain limited. The level of heterogeneity present in real-world traffic conditions is not yet well understood [211], and traditional modelling approaches often fail to capture its dynamic, context-sensitive, and multi-dimensional nature. To address this, researchers have used micro-simulation approaches, which involve quantifying driving heterogeneity from empirical data to reproduce heterogeneous driving behaviours in simulated environments. Common methods include clustering observable driving variables such as speed, headway, and acceleration [212, 213], or calibrating individual parameters in car-following models [101]. These approaches typically classify drivers into discrete classes based on their driving styles and assign them distinct parameter sets. For example, the Classified Car-Following (CCF) models proposed by Sun et al. [18] calibrate car-following behaviour separately for different driving styles to reflect inter-driver variability, which has been used to analyse the impacts of aggressiveness in driving on safety and fuel efficiency [25]. However, these category-based models generally assume static inter-driver heterogeneity and are limited in capturing intra-driver variability that evolves over time or in response to situational factors.

To overcome these limitations, recent efforts have attempted to decompose complex driving behaviour into simpler, interpretable behavioural units to capture more nuanced characteristics of the driving process. For instance, Wang et al. [40] identified 75 driving primitives serving as “primitives” to analyse diverse driving styles. Yao et al. [57] introduced the concept of *Action phases* to capture temporally underlying driving characteristics. These methods capture a broader spectrum of driving behaviour; however, the large number of primitives presents challenges in heterogeneity interpretation and model integration. To address this, Yao et al. [58] further grouped *Action phases* into six representative

Action patterns, providing a more structured and interpretable basis for identifying driving heterogeneity. Despite these efforts for driving heterogeneity identification, methods for systematically reproducing heterogeneous behaviours and evaluating their traffic-level impacts remain limited due to the inherent complexity of driving processes.

The present study addresses these gaps by proposing a pattern-based approach to model heterogeneous driving behaviour and analyses its impacts on traffic flow operations. The key contributions are threefold: (i) introducing a novel method to identify and represent driving heterogeneity through structured driving patterns, offering a new lens for behavioural modelling; (ii) developing a bi-level framework that captures both high-level *Action pattern* evolutions and low-level vehicle dynamics; and (iii) developing a micro-simulation approach that enables the investigation of diverse variation of inter- and intra-driver heterogeneity, allowing systematic analysis of its effects on traffic flow stability, safety, and fuel consumption.

7.2 Methodology

In the remainder, driving patterns refer to composite behavioural units capturing consistent driver responses to traffic conditions [58]. In this section, we provide an overview of the proposed pattern-based modelling and simulation framework for evaluating the impact of driving heterogeneity on traffic flow and outline the data preparation process used for its evaluation.

7.2.1 Overview of the pattern-based framework

The proposed framework introduces a pattern-based approach for driving behaviour modelling and simulation to assess the impacts of different levels of heterogeneity on traffic flow performance, offering advantages over traditional modelling methods by capturing temporal dynamic characteristics and both intra- and inter-driver heterogeneity. As illustrated in Figure 6.3, the framework consists of four key components: *Action phase* preparation [57], driving pattern identification [58], driving behaviour modelling, and traffic simulation for performance evaluation.

The process begins by decomposing driving trajectories into *Action phases*, which serve as “primitives” representing fine-grained characteristics of longitudinal driving behaviour. These *Action phases*, first introduced by Yao et al. [57], capture variations in velocity, acceleration, time headway, and speed difference over time and serve as the fundamental units for identifying higher-level driving tendencies. To enhance the interpretability of these “primitives”, the extracted *Action phases* are subsequently clustered into six group-specific driving patterns, referred to as *Action patterns*, as proposed by Yao et al. [58]. These patterns, such as “Follow behind”, “Catch up”, and “Speed up”, etc., serve as interpretable representations of group-specific driving behaviours and provide a compact and semantically meaningful way to understand behavioural dynamics. Specifically, each *Action phase* in driving trajectories can be signed with an *Action pattern*.

A bi-level approach is then introduced for pattern-based driving behaviour modelling, which is the core of the proposed framework. At the high level, this chapter models the temporal evolution of driving behaviour through probabilistic transitions between *Action*

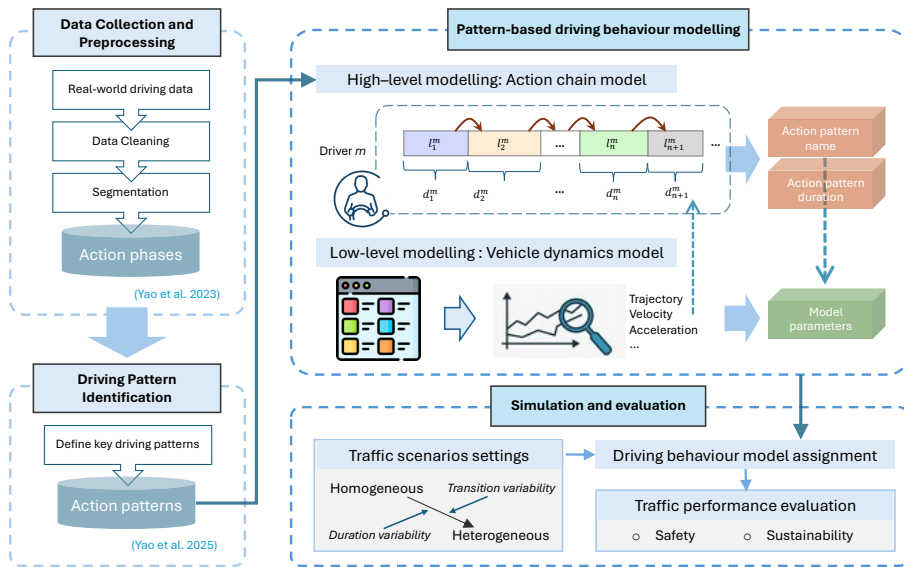


Figure 7.1: An overview of the proposed pattern-based modelling and simulation framework.

patterns, coupled with duration modelling that determines how long an *Action pattern* persists. This level captures the sequence, frequency, and timing of behavioural changes. At the low level, each *Action pattern* is associated with a calibrated vehicle dynamics model that governs acceleration, velocity, and location updates. The two levels interact via a dynamic feedback mechanism, where the high-level model guides the parameterisation of the low-level model, and the outputs from the low-level model, such as driving variable updates, can influence the transitions and durations of high-level patterns. For example, when a driver transitions from a “Follow behind” to a “Speed up” pattern, the parameters of the vehicle dynamics model are adjusted accordingly to generate the corresponding acceleration, velocity, and trajectory segment. This bi-level integration enables the simulation of realistic and diverse driver behaviours, reproducing both the decision-making process and its physical manifestations in terms of vehicle motion.

Finally, the calibrated heterogeneous driving behaviour models are used in micro-simulation, allowing for the simulation of traffic scenarios with varied heterogeneity levels. This enables the assessment of how different levels of driver heterogeneity, both within and across individuals, affect macroscopic traffic performance such as flow dynamics, safety, energy efficiency, and stability. Detailed descriptions of each step are provided in the following sections.

7.2.2 Data preparation

The Lyft Level-5 open dataset is used to evaluate the proposed framework in this study, which provides large-scale, high-resolution naturalistic driving data with diverse driving scenarios, making it suitable for capturing the variability and complexity of human driving

behaviour. We selected 3,000 pairs of human-driven vehicle (HV) following HV data for driving pattern calibration, which includes two key processes:

- *Action phase* extraction [57]: Vehicle trajectories are first segmented into *Action phases* using rule-based and data-driven criteria that consider time-series characteristics of driving variables including velocity (v), acceleration (a), time headway (T), and speed difference (Δv).
- *Action pattern* calibration [58]: The extracted *Action phases* are then categorised into six distinct *Action patterns* which serve as behaviour labels for further modelling. This process includes rigorous clustering based on unsupervised learning techniques (i.e., Agglomerative clustering and X-means clustering) and driving variable importance identification.

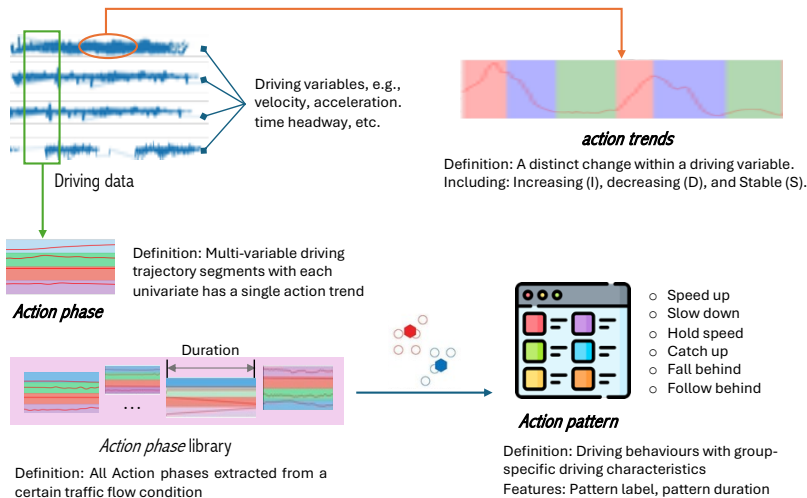


Figure 7.2: Clarification of action-related concepts.

The clarification of action-related concepts is shown in Figure 7.2. More details about the driving pattern preparation process can be found by Yao et al. [57, 58]. Six distinct *Action patterns* are identified, which comprehensively capture the fundamental longitudinal interactions observed in naturalistic driving. These patterns include both individual speed regulation (e.g., “Speed up”, “Slow down”, “Hold speed”) and interactions with other vehicles (e.g., “Catch up”, “Fall behind”, “Follow behind”). The categorisation reflects both self-driven and environment-responsive behaviours, covering a wide range of realistic car-following scenarios. These prepared driving patterns form the behavioural foundation for modelling heterogeneous driving behaviours and simulating diverse traffic scenarios in the later stages of the study. By embedding empirical driving patterns into the simulation, we ensure that the simulated traffic reflects actual driver behaviour, while still allowing exploring different heterogeneity levels through the framework’s parametric structure.

7.3 Pattern-based Driving Behaviour Modelling

In this section, we introduce the modelling structure of the proposed pattern-based driving behaviour framework, which consists of high-level and low-level models. The high-level model describes group-specific characteristics of driving behaviour by modelling discrete driving patterns, while the low-level model captures fine-grained characteristics within these patterns, describing continuous driving behaviours. The two levels interact via a dynamic feedback mechanism, where the high-level model guides the parameterisation of the low-level model, and the outputs from the low-level model, i.e., the updates to driving variables, can influence the transitions and durations of high-level patterns.

7.3.1 High-level modelling: Action chain model

Let $S = [(P_1, D_1), (P_2, D_2), \dots, (P_n, D_n)]$ represent a driver's current sequence of *Action patterns*, where P_n denotes the *Action pattern* name of the n -th phase, and D_n indicates the corresponding *Action pattern* duration. An *Action chain* is defined as a temporal evolution between two consecutive *Action patterns*, denoted by $(P_n, D_n) \rightarrow (P_{n+1}, D_{n+1})$. The high-level modelling aims to capture the sequential characteristics of *Action chains*, providing insights into temporal dependencies in longitudinal driving behaviour.

To capture the stochastic nature of these transitions, we model the evolution of *Action patterns* using a discrete-time Markov process, which assumes that the transition to the next state depends only on the current state. This is well-suited for modelling short-term behavioural decision-making in dynamic contexts such as longitudinal driving behaviours [214]. Let X_n denote a sequence of random variables indexed by discrete phase $n \in \{0, 1, 2, \dots\}$, taking values in a finite, countable state space S . The process $\{X_n\}$ is a Markov chain and satisfies the Markov property:

$$\mathbb{P}(X_{n+1} = \lambda \mid X_0 = x_0, \dots, X_n = x_n) = \mathbb{P}(X_{n+1} = \lambda \mid X_n = x_n) \quad (7.1)$$

The dynamics of a Markov chain are fully characterised by the set of transition probabilities:

$$p_{jk}(n) = \mathbb{P}(X_{n+1} = k \mid X_n = j) \quad (7.2)$$

which denotes the transition probability from state j to state k .

To reflect behavioural variability over time, we adopt an inhomogeneous Markov chain [215], where transition probabilities are phase-dependent. This approach allows the model to capture non-stationary behavioural dynamics influenced by factors such as driving context, traffic conditions, and driver state. Then the time-varying transition probability matrix $\mathbf{P}(n) \in \mathbb{R}^{K \times K}$, where K is the number of distinct *Action patterns*, is defined as:

$$\mathbf{P}(n) = \begin{pmatrix} p_{11}(n) & p_{12}(n) & \cdots & p_{1K}(n) \\ p_{21}(n) & p_{22}(n) & \cdots & p_{2K}(n) \\ \vdots & \vdots & \ddots & \vdots \\ p_{K1}(n) & p_{K2}(n) & \cdots & p_{KK}(n) \end{pmatrix}, \quad (7.3)$$

where each row satisfies $\sum_{j=1}^k p_{jk} = 1 \quad \forall k \in \{1, 2, \dots, K\}$, K is the number of distinct *Action patterns*. $p_{ij}(n)$ denotes the probability of transitioning from state j to state k at phase n .

Action pattern transition probabilities are estimated from empirical observations. To estimate $\eta_{jk}(n)$ from empirical data, let $\eta_{jk}(n)$ be the observed count of transitions from state j to k at phase n . The maximum likelihood estimate (MLE) of the transition probability is computed as:

$$\hat{p}_{jk}(n) = \frac{\eta_{jk}(n)}{\sum_{k=1}^N \eta_{jk}(n)}. \quad (7.4)$$

7.3.2 Low-level modelling: vehicle dynamics model

Based on driving patterns obtained from the high-level model, the low-level modelling process simulates continuous vehicle motions within each *Action pattern*. This is done by adjusting parameters dynamically based on empirically observed relationships among acceleration (a), speed difference (Δv), and distance to the leading vehicle (s). Each *Action pattern* is associated with a dedicated dynamics formulation that captures its unique behavioural characteristics.

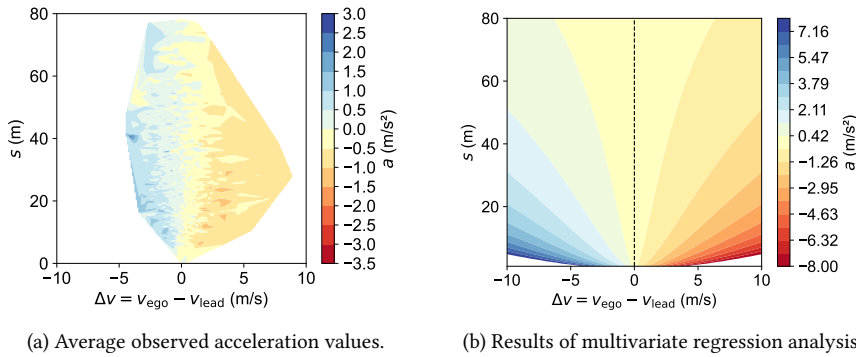


Figure 7.3: Relationship between speed difference, distance, and observed acceleration for speed-related patterns.

For speed-related patterns, including “Speed up”, “Slow down”, and “Hold speed”, empirical analysis in Figure 7.3a shows that the average observed acceleration is predominantly influenced by speed difference and to a lesser extent affected by distance. This aligns with findings by Hoogendoorn et al. [216], which highlight the role of speed difference in the switch of action points. The best-fitting model for acceleration in speed-related patterns is:

$$a = \left(\frac{b_1}{\sqrt{s}} + b_2 \right) \cdot \Delta v + \epsilon, \quad (7.5)$$

the fitted coefficients are $b_1 = 2.28$, $b_2 = 0.19$, $\epsilon \sim \mathcal{N}(0, \sigma)$ represents Gaussian noise. These parameters are determined from the observed driving data analysis, where the fitting results are shown in Figure 7.3b.

For the distance-related patterns, i.e., “Catch up”, “Fall behind”, and “Follow behind”, it is observed that the sign and magnitude of acceleration depend more on the distance than on the speed difference, see Figure 7.4, indicating that distance plays a more significant

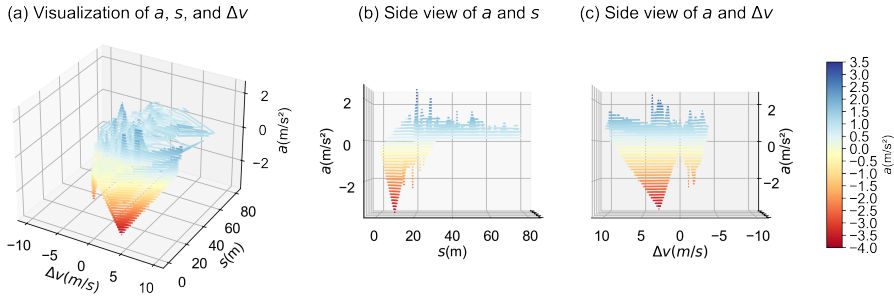


Figure 7.4: Relationship between speed difference, distance, and observed acceleration for distance-related patterns.

role. Accordingly, we define a general form that strongly depends on distance s and less on speed difference Δv :

$$a_{\text{base}} = \underbrace{K (s_{\text{target}} - s)}_{\text{gap control}} + \underbrace{\alpha \Delta v}_{\text{relative speed control}} \quad (7.6)$$

where K and α are parameters for gap control and relative speed control, respectively. $(s_{\text{target}} - s)$ is set to simulate acceleration and deceleration to get certain distance-related *Action patterns*. Instead of forcing a single equation to describe all three distance-related patterns, we calibrate separate formulas for each pattern. The “Catch up” pattern aims to decrease the distance, i.e., if s is too large, accelerate more:

$$a_{\text{catch-up}} = K_{cu} (s_{cu} - s) + \alpha_{cu} \Delta v \quad (7.7)$$

where $K_{cu} = -0.0266$, $s_{cu} = 20$, and $\alpha_{cu} = 1.06$.

Conversely, the “Fall behind” pattern means increasing the distance between the preceding vehicle, i.e., if s is too small, decelerate more:

$$a_{\text{fall-behind}} = K_{fb} (s_{fb} - s) + \alpha_{fb} \Delta v \quad (7.8)$$

where $K_{fb} = 0.00934$, $s_{fb} = 30$, and $\alpha_{fb} = 0.9144$.

The “Follow behind” indicates maintaining a comfortable distance from their leading vehicles:

$$a_{\text{follow-behind}} = K_{fl} (s_{fl} - s) + \alpha_{fl} \Delta v \quad (7.9)$$

where $K_{fl} = -0.0045$, $s_{fl} = 40$, and $\alpha_{fl} = 0.927$.

In reality, drivers cannot change their acceleration instantaneously due to vehicle inertia and human comfort limitations. To mitigate unrealistic discontinuities between *Action patterns*, we introduce a transitional phase when a pattern change is detected. If a pattern switch occurs at time t_χ , acceleration is defined as follows:

$$a_{t_\chi + \omega} = \left(1 - \frac{\omega}{T}\right) a_{t-1} + \frac{\omega}{T} a_t, \text{ for } \omega = 0, 1, \dots, T \quad (7.10)$$

where a_{t-1} is the acceleration before the pattern switch, a_t is the newly computed acceleration based on the new pattern, and T is the number of time steps over the transitional phase which is interpolated.

For each phase n , vehicle speed and location at time t are updated using kinematic equations as follows:

$$v_{t+1} = v_t + f_{a_k}(t) \cdot t, \quad (7.11)$$

$$x_{t+1} = x_t + v_t \cdot t + 0.5 \cdot f_{a_k}(t) \cdot t^2, \quad (7.12)$$

where $t = 0.1, 0.2, \dots$, where $f_{a_k}(t)$ is the acceleration at t determined by the current pattern function k , sustained throughout the duration of phase n . In particular, at the end of each phase, vehicle speed and location are calculated by:

$$v_{\text{after}} = v_{\text{before}} + a_{\text{after}} \cdot t, \quad (7.13)$$

$$x_{\text{after}} = x_{\text{before}} + v_{\text{before}} \cdot t + 0.5 \cdot a_{\text{after}} \cdot t^2, \quad (7.14)$$

This ensures continuity in the speed and trajectory profiles across successive *Action patterns* and enables the generation of smooth, realistic vehicle trajectories that reflect the influence of heterogeneous behavioural patterns.

7.4 Micro-Simulation for Heterogeneous Traffic Flow

In this section, we apply the developed pattern-based driving behaviour model to simulate heterogeneous traffic flow. This includes detailing the micro-simulation setup and specifying the performance indicators used to evaluate traffic safety, energy efficiency, and stability.

7.4.1 Micro-simulation settings

The micro-simulation environment consists of 20 vehicles initially positioned at equal intervals on a single-lane, straight road segment stretching from -1000 to 0 meters. Each vehicle starts with an initial speed of 15 m/s and follows the bi-level pattern-based driving behaviour model developed in Section 7.3. At the high level, each vehicle sequentially executes a series of *Action chains* over a 60-second simulation horizon, with a discrete time resolution of 0.1 seconds. At the low level, vehicle's accelerations are determined by the current activated *Action pattern* and are updated using the vehicle dynamics model introduced in Section 7.3.2, which in turn updates the vehicle's velocity and location. The micro-simulation procedure is illustrated in Figure 7.5, with the core implementation logic is provided as pseudocode in **Algorithm 3** in the Appendix.

To realistically replicate heterogeneous traffic flow conditions, we explicitly model both intra-driver heterogeneity (i.e., behavioural variability within a single driver over time) and inter-driver heterogeneity (i.e., behavioural differences across drivers). These two dimensions of heterogeneity are incorporated through stochastic modelling processes based on empirical data. Using naturalistic driving data from the Lyft Level-5 dataset, we extracted the transition probability matrix P among *Action patterns* and the duration distribution D for each *Action pattern* [217], which are illustrated in Figure 7.6. As shown in Figure 7.6a, some patterns, such as "Follow behind", occur more frequently, indicating a

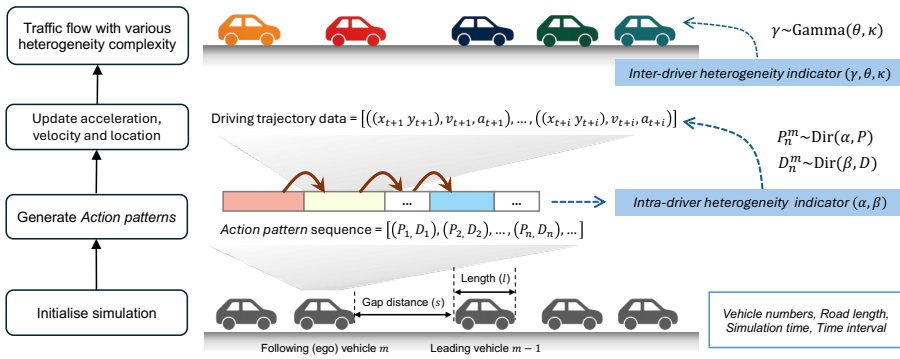


Figure 7.5: Illustration of the micro-simulation process for heterogeneous traffic flow

predominance of steady-state driving behaviours under this traffic condition. Figure 7.6b shows that the temporal consistency of *Action patterns* varies considerably, with some patterns exhibiting shorter durations and higher variability. The transition probability matrix P and the duration distribution D demonstrate the general driving characteristics of drivers. However, drivers are in essence heterogeneous, they do not always keep the same transition probability during driving.

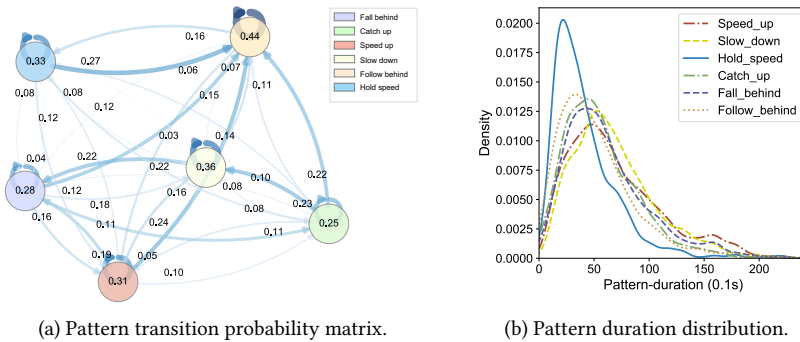


Figure 7.6: Statistics of *Action pattern* transitions and durations.

To model intra-driver heterogeneity, we use the Dirichlet process [218], a nonparametric Bayesian method that is well-suited for capturing behavioural variability and mixture complexity. A Dirichlet distribution $\mathbf{p} \sim \text{Dirichlet}(\alpha_1, \dots, \alpha_K)$ is characterised by the following properties:

$$E[p_j] = \frac{\alpha_j}{\sum_{k=1}^K \alpha_k} \tag{7.15}$$

$$\alpha_0 = \sum_{k=1}^K \alpha_k \tag{7.16}$$

where a large concentration parameter α_0 leads to low variability, with samples concentrated near the mean. conversely, as $\alpha_0 \rightarrow 0^+$, the samples become highly variable. k is a natural number [219]. This property allows us to simulate varying levels of behavioural consistency across *Action phases*.

For each driver m and *Action phase* n , the driver-specific transition probability matrix is sampled as:

$$P_n^{(m)} \sim \text{Dir}(\alpha, P) \quad (7.17)$$

where P is the empirical base pattern transition matrix, and α is a concentration parameter controlling the variability around the base transition matrix. Higher values of α yield more stable transition behaviours, while lower values of α allow for greater variability which causes more frequent and less predictable transitions between *Action patterns*.

Similarly, the pattern duration distribution for each driver m at *Action phase* n is sampled by duration distribution matrix D using:

$$D_n^{(m)} \sim \text{Dir}(\beta, D) \quad (7.18)$$

where D is the base pattern duration distribution and β controls the variability around it. Larger values of β produce more temporally consistent behaviour, whereas smaller β values introduce higher fluctuations in pattern durations.

Together, (α, β) characterises the intra-driver behavioural profile for individual driver m , enabling a flexible representation of temporal and behavioural variability.

To account for heterogeneity across drivers, we introduce a behavioural tendency parameter γ for individual drivers, sampled from a Gamma distribution:

$$\gamma \sim \Gamma(\kappa, \theta), \quad (7.19)$$

where θ represents the prototypical driver profile (i.e., the mean intra-heterogeneity characteristics), and κ is the dispersion parameter controlling variability across the population. A small κ leads to high inter-driver heterogeneity, with drivers exhibiting diverse behavioural tendencies. A large κ yields a more homogeneous population concentrated around the prototypical driver profile. The sampled γ scales the intra-driver concentration parameters (α, β) of each driver, enabling the simulation of heterogeneous traffic flow encompassing a wide spectrum of behavioural types from cautious, stable driving styles to aggressive, unstable ones. The parameters for the micro-simulation are summarised in Table 7.1.

Table 7.1: Parameters and descriptions for heterogeneous traffic flow simulation

Parameter	Description
α	Controlling intra-driver pattern transition variability
β	Controlling intra-driver pattern duration variability
(α, β)	Intra-driver heterogeneity profile of driver m
γ	Inter-driver heterogeneity profile
θ	Mean intra-heterogeneity profile of the population (base driver type)
κ	Dispersion parameter among driver profiles

In this way, unlimited heterogeneous traffic flow scenarios with varying levels of behavioural heterogeneity complexity can be simulated by systematically varying parameters α , β , and κ . Due to the space limit, we show the experiments and results by selecting representative parameter values. The varied concentration parameters α and β are investigated from 10^0 to 10^4 using a hybrid sampling strategy that combines logarithmic scaling (e.g., 10^0 , 10^1 , 10^2 , 10^3 , 10^4) with finer linear steps within each order of magnitude (e.g., 20, 30, ..., 90, 200, 300, ..., 900, etc.). This approach ensured both broad coverage and fine resolution for analysing behavioural variability. For the inter-driver heterogeneity setting, the prototype driver was defined by $(\alpha, \beta) = (5000, 5000)$, and the dispersion parameter κ was varied between 1 and 1000 to scale out different driver's profiles and simulate a wide spectrum of inter-driver heterogeneity in traffic flow.

7.4.2 Assessment indicators

To evaluate the safety and energy efficiency of traffic flow under different heterogeneity settings, two key indicators are used: Time to Collision (TTC) for safety assessment and Vehicle Specific Power (VSP) for energy evaluation.

TTC is widely used to assess the safety margin between a following vehicle and its leader. It is defined as the time when the speed of the objective vehicle is greater than its leading vehicle, the objective vehicle keeps the original driving state and does not take the corresponding deceleration behaviour until the two vehicles collide.

$$TTC_i(t) = \begin{cases} \frac{x_{i-1}(t) - x_i(t) - l_{i-1}}{v_i(t) - v_{i-1}(t)}, & \text{if } \forall v_i(t) > v_{i-1}(t) \\ \infty, & \text{if } \forall v_i(t) \leq v_{i-1}(t) \end{cases} \quad (7.20)$$

where $x_i(t)$ and $v_i(t)$ are the location and speed of vehicle i at time t , respectively. $x_{i-1}(t)$ and $v_{i-1}(t)$ denote the location and speed for vehicle $i-1$; l_{i-1} represents the length of vehicle $i-1$.

To capture safety risks over time, we compute time-exposed time-to-collision (TET), which accumulates the duration where the TTC falls below a critical threshold TTC^* [159]:

$$TET^* = \sum_{i=2}^N \sum_{t=0}^T \delta_i(t) \cdot \Delta t, \quad \delta_i(t) = \begin{cases} 0, & \text{Otherwise} \\ 1, & 0 \leq TTC_i(t) \leq TTC^* \end{cases} \quad (7.21)$$

where T is the total observation time (denotes simulation time here) and Δt is the interval time. TTC^* represents the safety threshold, often assumed to be between 1 and 3s, see [160]. It is set as 2s in this paper. $\delta_i(t)$ is a 0-1 variable. When $TTC_i(t)$ is less than TTC^* , $\delta_i(t)$ is equal to 1; otherwise, it is 0. For $N(i = 2 \dots N)$ drivers in the observation section, the total TET^* is calculated by equation (7.21). Lower values of TET^* indicate safer traffic conditions with fewer safety risks.

VSP represents the instantaneous power demand per unit vehicle mass and is a widely accepted indicator for fuel consumption and emissions [163]. It captures both the power for overcoming aerodynamic drag and rolling resistance, and the kinetic and potential energy of the vehicle are taken into account, by doing this the relationship between VSP and fuel

consumption can be explained physically. The calculation is shown in Eq. 7.22.

$$VSP = 0.132 \cdot v + 1.1 \cdot v \cdot a + 0.0003202 \cdot v^3 \quad (7.22)$$

VSP is usually clustered into bins at certain intervals, and traffic emissions are estimated by average-speed-based VSP distribution within each bin [164]. The categorisation of VSP into 1KW/t intervals is detailed in Eq. 7.23.

$$VSP \text{ bin} = n, \forall : VSP \in [n - 0.5, n + 0.5) \quad (7.23)$$

where n is the VSP number. The distribution of VSP across these bins provides a basis for evaluating energy efficiency under different traffic heterogeneity scenarios.

7.5 Results and Discussions

This section presents the results of applying the proposed pattern-based framework to simulate various traffic flow scenarios with differing levels of behavioural heterogeneity. We analyse both the emerging traffic flow dynamics and the impacts of driving heterogeneity on traffic safety, energy efficiency, and stability.

7.5.1 Analysis of heterogeneous traffic flow dynamics

To understand the impact of behavioural variability on traffic flow performance, it is essential first to demonstrate that the proposed model can reproduce realistic traffic flow dynamics and capture specific interactions between individual vehicles. The simulation results of a heterogeneous traffic flow scenario with moderate intra-driver variability are shown in Figure 7.7, characterised by concentration parameters $\alpha = 500$ and $\beta = 500$. The time-series plots of location, velocity, and acceleration are shown in Figure 7.7a-c, respectively. All vehicles start with a “Follow behind” pattern, representing a stable driving regime. Initially, vehicles maintain smooth trajectories and consistent leader-follower velocities due to small behavioural variability, see the first few seconds in Figure 7.7a. As intra-driver heterogeneity introduces variability in *Action pattern* transitions and durations, behavioural fluctuations within and among drivers emerge, resulting in local traffic disturbances, see inconsistent spacing between vehicles highlighted by the red circle area in Figure 7.7a. To illustrate how such inconsistencies happen, Figure 7.8 demonstrates the trajectory pair of vehicles 9 (leader) and 10 (follower). At approximately 40 seconds, during the #14 phase (see the red-shaded region in Figure 7.8a, the follower unexpectedly switches to a “Speed up” pattern while the leader maintains a “Hold speed” pattern. In the subsequent phase (#15), the leader adopts a “Follow behind” pattern, yet the follower continues accelerating, see Figure 7.8b, resulting in a reduced inter-vehicle distance and increased collision risk. This reflects real-world driving tendencies, particularly under high intra-driver heterogeneity where drivers may misunderstand the intentions of the vehicle ahead or respond impulsively. These findings demonstrate that the proposed pattern-based simulation framework effectively captures realistic traffic phenomena arising from individual behavioural variability. In contrast to traditional deterministic car-following models without considering such inconsistencies, the framework offers

a more nuanced representation of driving behaviour. By incorporating driver-specific differences in behavioural transitions, the framework offers a powerful tool for evaluating the impacts of heterogeneity on traffic performance such as safety and energy efficiency.




Figure 7.7: Traffic dynamics of heterogeneous traffic scenario $\alpha = 500$ and $\beta = 500$.

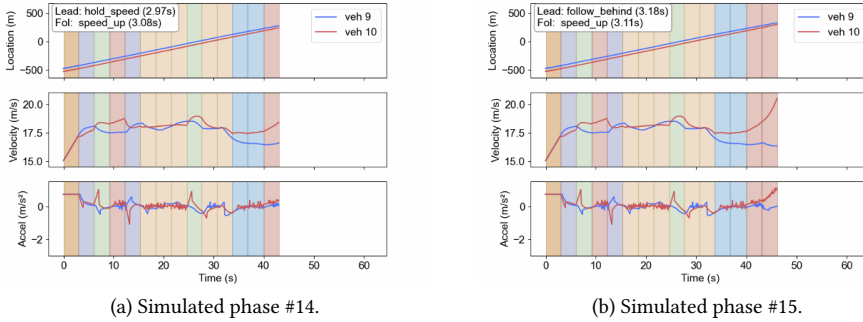


Figure 7.8: Trajectory pair of vehicles #9 (leader) and #10 (follower).

To further investigate the mechanism of intra-driver variability on behavioural dynamics, we examine the evolution of pattern transition probabilities for a single driver across different levels of heterogeneity. The accumulated probability of a given *Action pattern* j being transitioned into from all six *Action patterns* is calculated using $\Pr_{bt}(P_{n+1} = P_j) = \frac{\sum_{i=1}^6 \Pr(P_{n+1}=P_j|P_n=P_i)}{6}$ [217]. Take driver #5 as an example, Figure 7.9 illustrates the distribution of $\Pr_{bt}(P_{n+1} = P_j)$ at each phase during the simulation period. The blue vertical lines denote the range of the probability, and the black horizontal bars represent the median of the distribution. Wider parts of the violin indicate higher probability density. The red stars represent the general probability $\bar{\Pr}_{bt}$ of each *Action pattern* being transitioned from all drivers, which corresponds to the base transition matrix. The general probability $\bar{\Pr}_{bt}$ of each *Action pattern* being transitioned: Speed up(0.18), Slow down(0.12), Hold speed(0.12), Catch up(0.14), Fall behind(0.16), Follow behind(0.28). At low α values e.g., $\alpha = 1$ in Figure 7.9a-b, the transition probabilities deviate substantially from the base transition matrix. This wide spread and multimodal distributions reflect high variability in behavioural decisions across phases, indicating highly fluctuated and

context-sensitive driving characteristics. This corresponds to a manifestation of intra-driver heterogeneity, where driving behaviour may change dramatically depending on the momentary traffic conditions. As α increases, see Figure 7.9c–e, the transition probabilities become more concentrated around the base transition matrix. At $\alpha = 10000$ shown in Figure 7.9f, the sampled transition probabilities exhibit minimal deviation, reflecting a low intra-heterogeneity driver who consistently follows the general behavioural tendencies during driving. These results highlight the effectiveness of the concentration parameter α for regulating intra-driver variability in decision-making and how the proposed simulation framework can represent a wide spectrum of intra-heterogeneity in traffic flow by tuning parameter α .

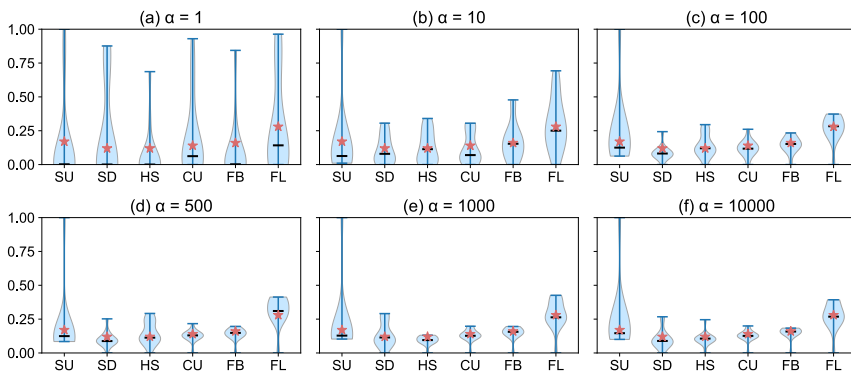


Figure 7.9: Statistics of transition probabilities during simulation, an illustration of Driver #5. [217]

Figure 7.10 demonstrates details of how intra-driver heterogeneity influences traffic flow dynamics by showing the evolution of vehicle trajectories in terms of location, velocity, and acceleration under three α settings: 100 (high variability), 1,000 (moderate variability), and 10,000 (low variability). In the low-variability case where $\alpha = 10000$, see Figure 7.10a, traffic flow remains smooth and coordinated with few and small disturbances: (a) gaps between vehicles are well maintained, (b) velocity profiles remain steady with low dispersion and closely align with their leaders, and (c) acceleration fluctuations are mild. This is because driver-specific transition probabilities are tightly concentrated around the base transition matrix, resulting in vehicles consistently behaving similar *Action patterns* over time. As α decreases to 1000, as shown in Figure 7.10b, variability in pattern transition matrix increases. The traffic flow remains largely coherent, while there are larger deviations in speed and acceleration among drivers. Specifically, some vehicles exhibit abrupt changes in velocity, such as the light purple vehicle and the deep gray vehicle, which diverge from the motion of their leaders, introducing localised disturbances into the traffic flow. When α is further reduced to 100, as shown in Figure 7.10c, the intra-driver variability becomes larger. The velocity profiles become chaotic, with frequent large accelerations and decelerations. Gaps between vehicles become less consistent, and some vehicles lag behind or move ahead unexpectedly. The traffic flow pattern shows clear signs of instability, with dynamic gaps, speed shocks, and increased oscillations. These characteristics correspond to reduced safety

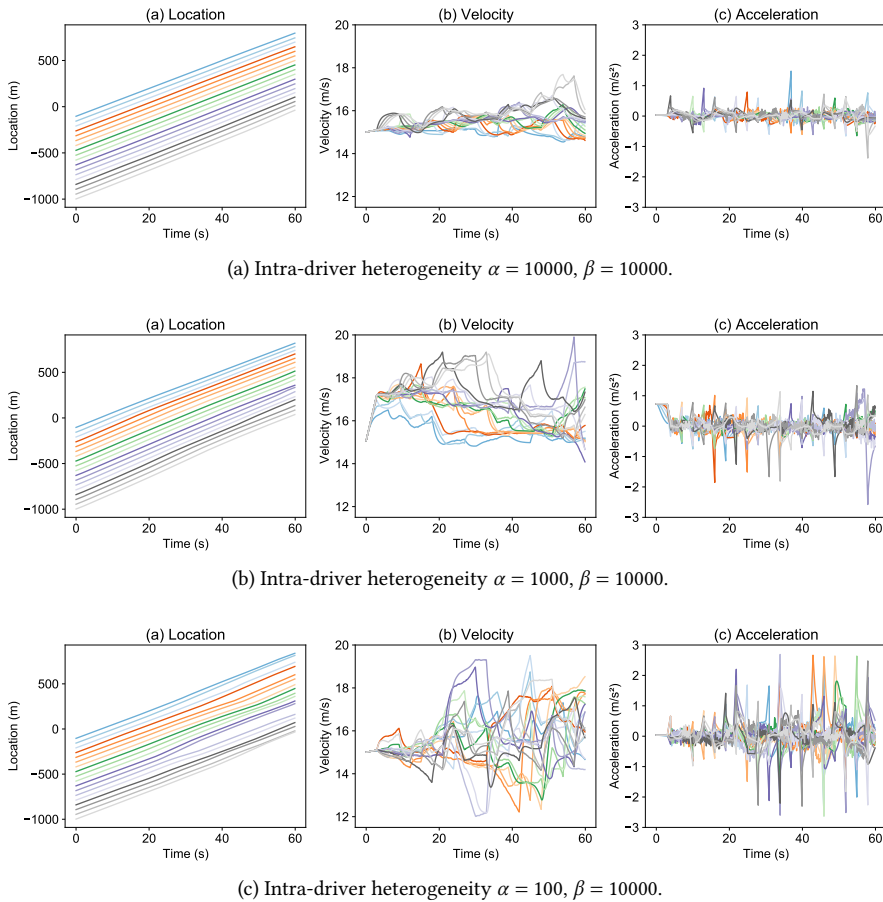


Figure 7.10: Heterogeneous traffic flow scenario with various α settings.

and a higher probability of conflict, highlighting the negative impact of high intra-driver heterogeneity.

The consistency with drivers maintaining selected *Action patterns*, i.e., the duration of behavioural phases, also plays an important role in shaping traffic dynamics. To quantify intra-driver heterogeneity in terms of temporal commitment to driving behaviours, we examine how the duration distribution of each *Action pattern* changes under different values of the concentration parameter β . Figure 7.11 shows the pattern duration distributions for Driver #5 under various values of β , ranging from 1 to 10000. Notice that at low β in Figure 7.11a, the distributions are broad and highly dispersed across patterns. This indicates significant temporal variability in how long *Action patterns* are sustained during simulation, reflecting a high degree of intra-driver heterogeneity in temporal decision-making. As β increases, see Figure 7.11b–f, the distributions become sharper and more concentrated.

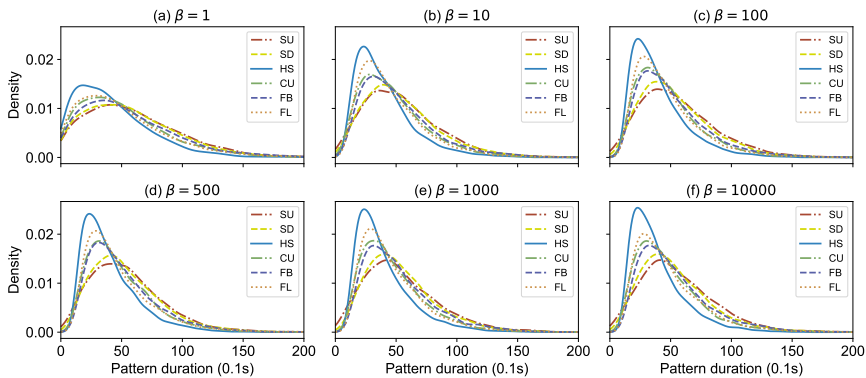


Figure 7.11: Illustration of pattern duration distribution during simulation, an example of Driver #5.

When $\beta = 10000$ is shown in Figure 7.11f, the duration distributions of each *Action pattern* closely align with their base duration distribution, resulting in highly consistent temporal behaviours.

We fix the concentration parameter at $\alpha = 10000$ to ensure a stable pattern transition matrix and vary the duration concentration parameter β across different levels. Figure 7.12 illustrates the results of traffic dynamics for $\beta = 100, 1000$ and 10000 , representing high, moderate, and low duration variability, respectively. When the pattern durations are highly consistent, see Figure 7.12a where $\beta = 10000$, the duration of *Action phases* are highly stable over time and across drivers. This leads to behavioural patterns with consistent following distances, smoother velocity profiles, and less abrupt accelerations and decelerations. At a moderate level of variability where $\beta = 1000$, see Figure 7.12b, the platoon structure is less coherent with some disturbances. Vehicles maintain *Action patterns* for varied intervals, resulting in more gradual changes in velocity (b) and more fluctuations in acceleration (c). In Figure 7.12c where $\beta = 100$, the duration assigned to each *Action pattern* is highly variable, resulting in chaotic velocity profiles and frequently abrupt acceleration and decelerations, leading to visible perturbations in the traffic flow.

In sum, these results highlight the mechanism of disturbance in traffic flow caused by different levels of variability in *Action pattern* transitions and durations, indicating that behavioural variability plays a critical role in shaping traffic flow characteristics. Lower behavioural heterogeneity (high α and β) promotes more stable traffic flow conditions while higher heterogeneity introduces noise into the system, amplifying disturbances and increasing instability, inefficient spacing, and unsafe interactions.

7.5.2 Analysis of traffic safety and energy efficiency

Before assessing how driving heterogeneity influences traffic flow performance, we first examine the role of the dispersion parameter κ , which governs the variability of driver-specific behavioural profiles around a prototype driver. A larger κ leads to high inter-driver heterogeneity, with drivers exhibiting diverse behavioural tendencies. As

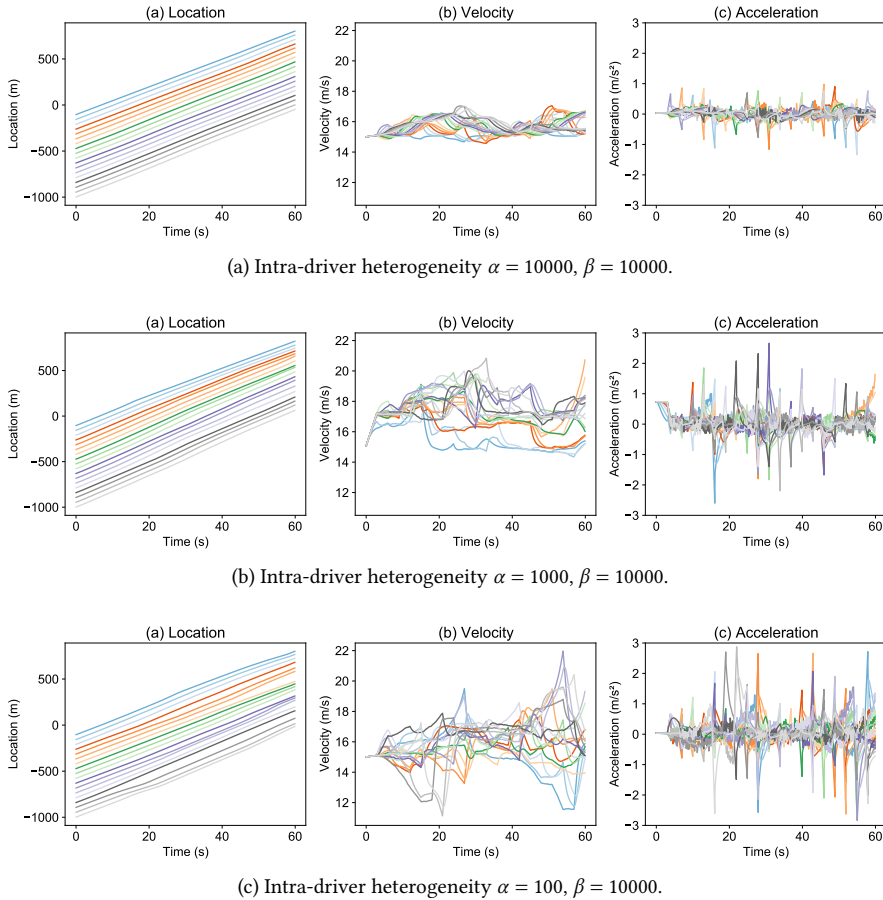


Figure 7.12: Heterogeneous traffic flow scenario with various β settings.

illustrated in Figure 7.13, the prototype driver denoted by the orange star is defined by concentration parameters $\alpha = 5000, \beta = 5000$, representing low intra-driver variability. The blue circles represent individual driver profiles sampled under different κ values from 1 to 1000. In Figure 7.13a where $\kappa = 1$, the driver profiles exhibit wide dispersion across the (α, β) space, exhibiting high variability between each other. This high degree of heterogeneity yields a diverse set of behavioural patterns in the simulated traffic flow. As κ increases, see Figure 7.13b-c, the sampled profiles become increasingly concentrated around the prototype. With large κ values shown in Figure 7.13d, most drivers exhibit nearly identical profile characteristics, resulting in a traffic population with a low level of inter-driver heterogeneity. These observations emphasise that κ serves as an effective parameter for controlling the breadth of inter-driver diversity in the micro-simulation.

We systematically vary κ to simulate a wide spectrum of heterogeneous traffic




Figure 7.13: Illustration of varied κ settings for inter-driver heterogeneity.

flow scenarios, ranging from highly diverse populations ($\kappa = 1$) to low heterogeneous ones ($\kappa = 1000$). This design allows the investigation of how different levels of driving heterogeneity influence key traffic performance indicators: (a) the number of collisions, (b) the number of critical time-to-collision (TET^*), and (c) the total Vehicle Specific Power (VSP). Figure 7.14 illustrates heatmaps of the three key indicators across varying combinations of κ_α (horizontal axis, controlling dispersion in pattern transition variability) and κ_β (vertical axis, controlling dispersion in pattern duration variability). Each grid cell represents the average of 10 simulation runs under a given heterogeneity configuration. In Figure 7.14a and b, both the number of collisions and TET^* events increase substantially when inter-driver heterogeneity is high. This is most evident in the lower-left regions of the heatmaps where κ_α and κ_β are small. Under these conditions, drivers exhibit diverse and uncoordinated behavioural tendencies, leading to inconsistent car-following dynamics and frequent conflict situations. As the values of κ increases, the behavioural profiles of drivers become more consistent, resulting in a reduction in safety-critical events. When κ approaches 1000, safety performance reaches the highest, indicating that consistency in driver characteristics contributes significantly to improved traffic safety. Energy efficiency, as measured by total VSP shown in Figure 7.14c, follows a similar trend. High inter-driver heterogeneity leads to frequent and uncoordinated *Action pattern* changes, resulting in erratic speed profiles, frequent accelerations and decelerations, and ultimately higher fuel consumption, see the lower left area of Figure 7.14c. As κ increases, VSP values generally reduce, reflecting more energy-efficient traffic flow conditions.




Figure 7.14: Statistics of indicators for traffic safety and energy efficiency.

Taken together, the results demonstrate that high inter-driver heterogeneity negatively

affects both safety and energy efficiency. These findings underscore the importance of reducing behavioural heterogeneity within and across drivers, either through vehicle automation, cooperative driving systems, or targeted driver interventions, to promote a safer and more sustainable traffic system.

7.5.3 Analysis of traffic stability

To evaluate the model's capability to capture traffic stability dynamics, we introduce a localised disturbance by enforcing a deceleration manoeuvre on a lead vehicle (Vehicle #5). The vehicle decelerates at a rate of $a = -2\text{m/s}^2$ over a 6-second interval. Figure 7.15 illustrates the spatiotemporal propagation of the disturbance within the vehicle platoon. Notice that the disturbance triggers a velocity drop that is not contained locally but propagates upstream through the platoon. This propagation is evidenced by the alternating red and yellow bands in the trajectory plots, representing oscillations in vehicle speed. To quantify the disturbance propagation rate, we identify the first moment each vehicle's velocity drops below 47km/h , a threshold indicating the onset of congestion flow [220]. These points, marked as black stars in Figure 7.15a, indicate the formation of three distinct congestion wavefronts, rather than a uniform propagation. This deviation is derived from the impacts of driving heterogeneity, particularly in the form of different *Action patterns* derived from individual decision-making tendencies. As shown in Figure 7.15b, the disturbance propagates upstream at a speed of 11.7 m/s , whereas the propagation speeds of the three identified wavefronts (w_1, w_2, w_3) are much higher, ranging from -30.01m/s to -30.65m/s . These values indicate abrupt deceleration waves propagating rapidly through fast-moving vehicles on the highway. This phenomenon is likely driven by heterogeneous driving behaviour, where some drivers initiate braking with relatively short spacing, resulting in a sharp and fast upstream transmission of the disturbance. Additionally, some vehicles, such as vehicles #7 and #11, deviate from the expected upstream propagation trajectory, introducing further inconsistency into the platoon and potentially increasing the risk of collisions.

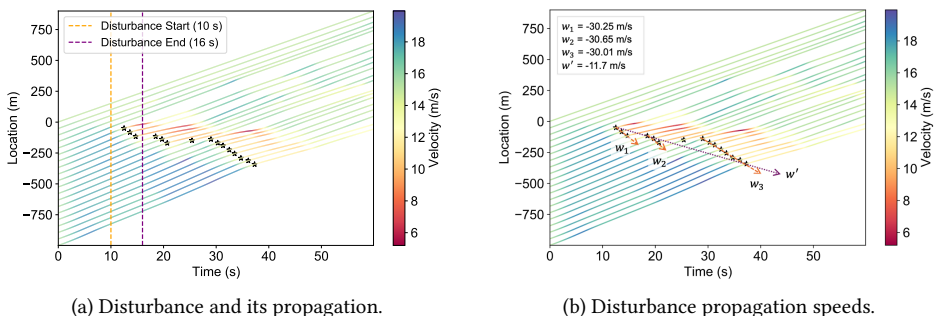


Figure 7.15: Spatiotemporal evolution of a localised disturbance in a vehicle platoon.

Figure 7.16 provides further insight into the disturbance by examining the relationship between relative speed and spacing for the disturbed vehicle (Vehicle #5) and its immediate

follower (Vehicle #6). Different colours represent the six *Action patterns*, and circles and crosses indicate acceleration and deceleration, respectively. In Figure 7.16a, Vehicle #5 exhibits a clear deceleration phase from 10s to 16s, marked by “+”, which increases the spacing from its leader. Although it subsequently begins to accelerate, its speed remains lower than that of the leader for a period, causing the gap to increase further until the vehicle eventually stabilises at a larger spacing than before the disturbance occurred. Figure 7.16b shows Vehicle #6 reacting to the disturbance by decelerating. The decreasing spacing between it and its leader reveals an oscillatory response characteristic of traffic instability. These observations align with empirical Wiedemann car-following dynamics [216], indicating that the proposed behaviour model successfully captures realistic disturbance propagation and response.

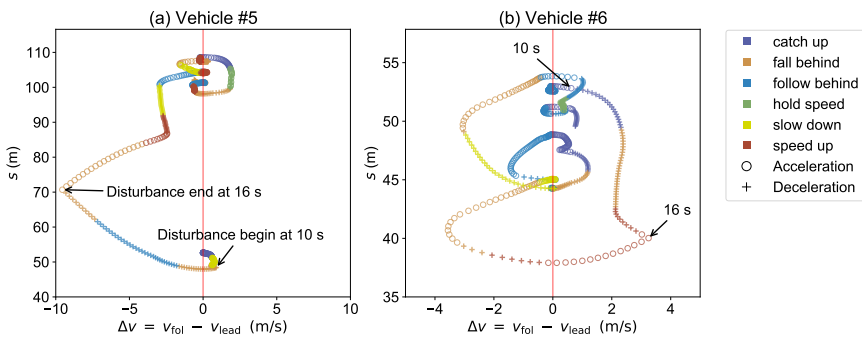


Figure 7.16: Simulated trajectories of the disturbed vehicle and its immediate follower.

Figure 7.17 presents the standard deviation of vehicle velocity along the platoon as a metric for string stability. Two cases are compared: one with a mild disturbance ($a = -2\text{m/s}^2$ denoted by red squares) and one with a larger disturbance ($a = -3\text{m/s}^2$ represented by green dots). Prior to the disturbance, velocity fluctuations remained low and consistent across




Figure 7.17: The standard deviation of the velocity of each vehicle along the platoon.

both cases. However, after the disturbance, velocity variations increase significantly in the downstream vehicles, with the larger disturbance resulting in more severe instability till the end of the platoon. This result further indicates that driving heterogeneity contributes to the amplification of velocity oscillations and overall traffic instability by affecting both the stochasticity and intensity of disturbances.

7.6 Conclusion and Outlook

This paper presents a new pattern-based modelling and simulation framework for investigating the impact of driving behaviour heterogeneity on traffic flow dynamics, safety, energy efficiency, and stability. By decomposing driving behaviour into interpretable sequential *Action patterns*, the framework captures both intra-driver and inter-driver behavioural variability. The bi-level model integrates high-level behavioural sequence modelling through *Action chains* with low-level vehicle dynamics within each driving pattern. The proposed framework enables the simulation of diverse traffic scenarios by varying behavioural heterogeneity within and among individual drivers. Evaluation using real-world naturalistic driving data indicates that the simulations effectively reproduce realistic traffic phenomena such as instability, inefficient spacing, and increased collision risks. The results demonstrate that both transition and duration variability significantly influence traffic flow dynamics and that high levels of driving heterogeneity can degrade safety, energy efficiency, and stability.

We argue that this research provides valuable insights into the behavioural mechanisms underlying traffic instability and inefficiency, with theoretical implications for modelling as well as practical applications in intelligent transportation systems. Specifically, by identifying and quantifying driver-specific behavioural patterns, such as tendencies in acceleration, headway maintenance, and responsiveness, our framework allows for the parametrisation of driver models that can be integrated into algorithms for adaptive cruise control or eco-driving systems. These models can adjust longitudinal control actions in real time based on predicted behavioural responses, enhancing both safety and fuel efficiency. Furthermore, the framework enables the simulation and evaluation of traffic dynamics under varying behavioural compositions, which can support the optimisation of traffic signal timing or variable speed limits by accounting for behavioural heterogeneity in the driver population. In this way, the theoretical insights into behavioural diversity can inform more effective congestion mitigation and risk management strategies in mixed traffic environments.

Based on the promising results of this study, several directions for future investigation can be identified. First, instead of focusing on longitudinal behaviour in single-lane traffic, extending the framework to incorporate lateral dynamics (e.g., lane-changing) in multi-lane environments as well as interactions with diverse road users (e.g., cyclists and pedestrians) would enhance its generalisability and applicability. Second, rather than relying on static, empirically derived matrices to represent intra- and inter-driver heterogeneity, future work could explore dynamically learning pattern transitions and population-level variability using extended datasets and data-driven approaches, such as deep learning techniques. From an application point of view, integrating the proposed framework into

traffic simulation platforms or applying it to enhance in-vehicle driver assistance systems and cooperative traffic control strategies aimed at harmonising driver behaviour presents promising opportunities for future application.

Appendix

Algorithm 3 The procedure of pattern-based micro-simulation for heterogeneous traffic flow

```

1: Initialisation:
2: for  $m = 1$  to  $M$  do
3:   Assign initial Action pattern:  $p_0^{(m)} \leftarrow$  random select an Action pattern.
4:   Assign initial duration:  $d_0^{(m)} \leftarrow$  sample from  $D$ .
5:   Assign pattern-name transition matrix  $P^{(m)}$ :
6:   if  $m \in p_t$  then
7:      $P^{(m)} \sim \text{Dirichlet}(\beta \cdot P)$ 
8:   else
9:      $P^{(m)} \leftarrow P$ 
10:  end if
11:  Assign pattern-duration  $D^{(m)}$ :
12:  if  $m \in p_d$  (proportion with varied durations) then
13:     $D^{(m)} \leftarrow \{D_i \mid i = 1, 2, \dots, k\}$ 
14:  else
15:     $D^{(m)} \leftarrow \bar{D}$ 
16:  end if
17:  Set initial speed  $v_0^{(m)}$  and location  $x_0^{(m)}$ .
18: end for

19: Vehicle generation loop:
20: for  $n = 1$  to  $N$  do
21:   for  $m = 1$  to  $M$  do
22:    Update Action patterns:
23:    if  $m \in p_t$  then
24:       $p_{n+1}^{(m)} \sim P^{(m)}(p_{n+1} \mid p_n^{(m)})$ 
25:       $d_n^{(m)} \sim D^{(m)}$ 
26:    else
27:       $p_{n+1}^{(m)} \sim P(p_{n+1} \mid p_n^{(m)})$ 
28:       $d_n^{(m)} \leftarrow \bar{D}$ 
29:    end if
30:    Obtain speed difference and distance:
31:    if  $m > 1$  then
32:       $\Delta v_n^{(m)} \leftarrow v_n^{(m-1)} - v_n^{(m)}$ 
33:       $s_n^{(m)} \leftarrow x_n^{(m-1)} - x_n^{(m)} - l^{(m-1)}$ 
34:    else
35:       $\Delta v_n^{(m)} \leftarrow 0$ 
36:       $s_n^{(m)} \leftarrow \infty$ 
37:    end if
38:    Compute acceleration:
39:    if  $w_{n+1}^{(m)}$  is speed-related pattern then
40:      
$$a_n^{(m)} \leftarrow \left( \frac{b_1}{\sqrt{s_n^{(m)}}} + b_2 \right) \cdot \Delta v_n^{(m)} + \epsilon$$

41:    else
42:       $a_n^{(m)} \leftarrow K(s_{\text{target}} - s) + \alpha \Delta v$ 
43:    end if
44:    Safety check:
45:     $v_{\text{check}} \leftarrow \left( x_n^{(m-1)} - x_n^{(m)} \right) / TTC^* + v_n^{(m-1)}$ 
46:    Update speed and location:
47:    
$$v_{n+1}^{(m)} \leftarrow \min \left( v_n^{(m)} + a_n^{(m)} \cdot \Delta t_n^{(m)}, v_{\text{check}} \right)$$

48:    
$$a_n^{(m)} \leftarrow \left( v_{n+1}^{(m)} - v_n^{(m)} \right) / \Delta t$$

49:    
$$x_{n+1}^{(m)} \leftarrow x_n^{(m)} + v_n^{(m)} \cdot \Delta t_n^{(m)} + 0.5 \cdot a_n^{(m)} \cdot (\Delta t_n^{(m)})^2$$

50:  end for
51: end for

```

8

Conclusions and Perspectives

This chapter presents conclusions and outlook. The answers to the research questions and key findings of this thesis are presented in Section 8.1. Section 8.2 presents the overall conclusions of this thesis. Section 8.3 gives a comprehensive outlook of this thesis, including both methodological and practical perspectives.

8.1 Key findings

This section presents answers to each individual research question proposed in Section 1.2, followed by the corresponding key findings.

RQ1. What is driving heterogeneity and why does it matter?

RQ1.1. How is driving heterogeneity characterised and analysed in the literature?

This research question is answered in Chapter 2 by presenting a literature review that organises key concepts, methodology, and application of driving heterogeneity. The review introduces three conceptual perspectives and five identification categories. Building on this, a framework is proposed that structures the ML-based methodological identification process, including data preparation, feature selection, ML model implementation and performance evaluation. By reviewing various ML methodologies and assessing their strengths, limitations, and applicability across different driving contexts, this study emphasises the need to balance accuracy, interpretability, and real-time recognition for effective heterogeneity identification.

RQ1.2. What are the impacts of driving heterogeneity on traffic flow performance?

This research question is answered in Chapter 3 by presenting a general framework to investigate car-following heterogeneity and its impacts on traffic safety and energy efficiency. The study indicates that an increase in driving aggressiveness is correlated with higher safety risks and more fuel consumption and emissions. However, the negative impacts are not linearly related to aggressiveness. Elucidation from the underlying characteristics of traffic flow dynamics explains that less aggressive drivers tend to form vehicular platoons, which can induce more aggressive drivers to adopt milder driving styles. The effectiveness of this moderating effect depends on both the proportion and spatial distribution of less aggressive vehicles.

RQ2. How to identify driving heterogeneity through behavioural characteristics?

RQ2.1. What are the underlying mechanisms of driving heterogeneity?

This research question is answered in Chapter 4 by introducing an action approach to understanding driving heterogeneity from both individual driver and traffic flow levels. The proposed fundamental unit, termed *Action phase*, decodes longitudinal driving behaviour with physically meaningful characteristics based on key driving variables. Transitions between *Action phases* reflect driver-specific behavioural traits, while the *Action phase Library*, comprising all extracted phases under specific traffic conditions, captures characteristics of traffic flow level. This study offers interpretable insights into how behavioural diversity determines driver responses to traffic dynamics.

RQ2.2. *How can driving heterogeneity be identified from underlying driving characteristics?*

This research question is answered in Chapter 5 by proposing a novel framework to identify driving heterogeneity based on the action-based approach. The framework calibrates six *Action patterns* by clustering *Action phases* according to group-specific characteristics and reveals driving variable importance. These patterns serve as labels for training a deep learning model, where the integration of an attention mechanism into LSTM-based models significantly improves both classification accuracy and computational efficiency. The proposed framework promises advantages in traditional identification methods by capturing subtle behavioural differences within and among drivers.

RQ3. *How to model and evaluate heterogeneous driving behaviours?*

RQ3.1. *How can heterogeneity in driving behaviour be effectively modelled?*

This research question is answered in Chapter 6 by proposing a pattern-based modelling approach using knowledge-enhanced deep learning models. Driving behaviour is represented as sequences of *Action patterns*, with their transitions and durations modelled by attention-based LSTM models. Incorporating *Action pattern* transition and duration properties into LSTM-based models enhances their performance in driving pattern prediction. This study highlights the effectiveness of incorporating behavioural knowledge to enhance deep-learning models in driving behaviour analysis.

RQ3.2. *How can different levels of driving heterogeneity be captured in traffic simulation?*

This research question is answered in Chapter 7 by proposing a pattern-based modelling and simulation framework. By decomposing driving behaviour into interpretable sequential *Action patterns*, the framework captures both intra-driver and inter-driver behavioural variability. Specifically, the bi-level modelling integrates high-level behavioural sequence modelling through *Action chains* with low-level vehicle dynamics within each driving pattern. The micro-simulation enables the systematic variation of behavioural heterogeneity within and among drivers to replicate diverse traffic scenarios. Results show that the framework reproduces realistic traffic phenomena, such as instability, inefficient spacing, and increased collision risks, and reveals that variability in pattern transitions and durations significantly impacts traffic flow dynamics, safety, energy efficiency, and stability. High levels of driving heterogeneity are shown to negatively affect safety, fuel efficiency, and stability.

8.2 Overall conclusions

This thesis addresses the multifaceted problem of driving heterogeneity by developing a data-driven, AI-powered, and pattern-based framework. By systematically answering the proposed research questions, the thesis advances both methodological understanding of driving heterogeneity and practical applications for improving traffic safety and energy efficiency, by providing:

- A review of existing methods for identifying driving heterogeneity and understanding

its effects on traffic performance.

- A novel action-based methodology to characterise intra- and inter-driver heterogeneity
- An action-based modelling and simulation framework to evaluate the impacts of different levels of heterogeneity on traffic flow dynamics, safety, energy efficiency, and stability.

In the remainder, an overview is provided of the main conclusions drawn from the conducted research:

i) AI-driven analysis improves understanding of driving heterogeneity: This thesis presents a structured roadmap for identifying driving heterogeneity using machine learning techniques. A comprehensive review and synthesis of existing approaches reveals the importance of aligning data preprocessing, feature selection, and model evaluation with behavioural interpretability. The study demonstrates that the choice of learning method (e.g., supervised vs. semi-supervised) significantly influences classification robustness for driving behaviour analysis.

ii) Behavioural mechanisms of heterogeneity can be captured through structured representations: The concepts of *Action phases*, *Action patterns*, and *Action chains* presented in this thesis provide a novel perspective to understand longitudinal driving behaviour. These representations enable the decomposition of complex driving trajectories into interpretable behavioural units, facilitating a more precise analysis of sequential driver responses. Combined with expert knowledge and AI techniques, this approach enables the identification of both individual-specific tendencies and group-level behavioural characteristics.

iii) Driving heterogeneity shapes traffic flow dynamics in measurable ways: The developed bi-level modelling and simulation framework demonstrates how driving heterogeneity impacts traffic flow performance. By incorporating both high-level behavioural patterns and low-level vehicle dynamics, simulations of diverse heterogeneity profiles reveal how variations in driving behaviour contribute to traffic flow disturbances, dispersion in time headways, and increased collision risks.

In addition to these scientific conclusions, the research also offers practical relevance for real-world applications:

i) Personalised driver assistance systems can benefit from behaviour-aware models: The insights into machine learning model selection and interpretability on driving heterogeneity analysis support the development of robust behaviour identification systems. These systems provide a foundation for personalised in-vehicle assistance features, such as adaptive cruise control or eco-driving guidance, that align with individual driving styles and behavioural tendencies, improving safety and comfort.

ii) Action-based modelling enables intelligent traffic control strategies: The identification of behavioural pattern transitions under dynamic traffic conditions enables real-time driver state recognition. This capability can be integrated into intelligent traffic control systems to deliver adaptive interventions tailored to diverse driving behaviours, improving the responsiveness and operational efficiency of the traffic system.

iii) Addressing behavioural variability is key to sustainable and safe mobility:

The results underscore that both intra-driver and inter-driver variability significantly affect traffic stability, safety, and energy efficiency. By accounting for these variations, the proposed framework supports the development of behaviour-aware control strategies, such as adaptive traffic signals, platooning strategies, and lane management policies, contributing to more sustainable and safer mobility systems.

Overall, this thesis concludes that understanding and modelling driving heterogeneity is crucial for both advancing traffic flow theory and designing effective intelligent transportation solutions. The findings provide a foundation for developing infrastructure designs, policy measures, and vehicle-based interventions that mitigate the negative effects of behavioural variability, promoting safer, more efficient, and environmentally sustainable transportation systems.

8.3 Discussion and Outlook

Building on the methodologies and insights developed in this thesis, directions for future research and practical implementation are outlined below.

8.3.1 Future research directions

i) Integration of lateral behaviour and mixed manoeuvres: The current framework primarily addresses heterogeneity in longitudinal driving behaviour in single-lane highway scenarios. However, in real-world traffic, lateral movements such as lane changes, merges, and overtakes occur frequently and are influenced by complex situational and social interactions. The absence of lateral behaviour modelling limits the applicability of the current framework to more complex, multi-lane or urban traffic environments.

Therefore, future work should extend the framework by calibrating lateral *Action patterns* and integrating them into multi-dimensional *Action chains*, which capture both temporal and spatial transitions in driver behaviour. This would enable more realistic and holistic modelling of driving dynamics, particularly in multi-lane highways, merging zones, or urban intersections, thereby enhancing the behavioural realism of both driving behaviour models and traffic simulations.

ii) Contextual generalisability across driving environments: The current framework is designed and evaluated within the context of highway scenarios, which limits its generalisability to other types of road networks such as urban arterials, roundabouts, or signalised intersections. This limitation arises from the use of data collected in relatively uniform highway settings, which impacts the applicability of our findings in more complex and diverse driving environments.

To enhance generalisability, future research should adapt the framework to incorporate diverse driving contexts by including contextual information such as road geometry, traffic density and composition, and weather conditions. Additionally, deep learning techniques such as transfer learning could be employed to fine-tune the model across different environments with minimal retraining, supporting scalable deployment across varied traffic systems.

iii) Multimodal data fusion for human-centric modelling: The current study

focuses on tactical-level driving variables derived from external vehicle kinematics. While effective in capturing observable behavioural patterns, this approach does not account for underlying perceptual or cognitive states that influence driver decision-making, such as driver attention, workload, or stress. This limits the model's ability to explain the causes of behavioural heterogeneity and may reduce its effectiveness in high-risk or cognitively demanding contexts.

Future research could address this by incorporating additional data sources to infer internal driver states, including vision-based inputs (e.g., dashcam footage, eye-tracking), driver biometrics (e.g., eye gaze, head pose), and vehicle control actions (e.g., steering, throttle, brake pressure). Multimodal fusion techniques, such as attention-based architectures or deep sensor fusion, can improve the accuracy and interpretability of behaviour models, helping not only to describe heterogeneity but also to uncover its underlying causes.

iv) Online and adaptive learning mechanisms: The current rule-based approach for extracting *Action phases* relies on static threshold values derived from expert knowledge. While effective in structured settings, these fixed rules may not adapt well to drivers from different traffic conditions, roadway contexts, or cultural backgrounds. This constraint could limit the robustness and flexibility of the framework when deployed in real-world applications.

To overcome this limitation, future research could develop online and adaptive learning mechanisms that dynamically adjust *Action phase* detection thresholds. For example, reinforcement learning (RL) agents could be trained to optimise threshold selection based on feedback from behavioural accuracy or traffic performance indicators. This hybrid approach, combining rule-based interpretability with data-driven adaptability, would improve the framework's generalisation across diverse driving environments and driver types.

v) Modelling interactions and negotiation behaviours: This thesis focuses primarily on individual driver behaviour and the interaction between the leader and the intermediate follower. However, the more complex nature of driving, where negotiation with other road users also plays a central role, particularly in merging, overtaking, or intersection scenarios. The current framework works well for behavioural patterns at the individual level, but it limits the model's applicability in multi-agent traffic environments.

To capture these interactions, future work could incorporate game-theoretic frameworks or multi-agent reinforcement learning (MAREL) to simulate negotiation and adaptation among drivers. These methods can incorporate bounded rationality, trust, and social norms, making them well-suited to modelling complex human-AV interactions and decision-making in mixed traffic scenarios.

vi) Leveraging the capabilities of Large Language Models (LLMs): The current framework focuses on structured, data-driven modelling of driving behaviour based on predefined behavioural features and rule-based or supervised learning techniques. While effective in identifying and simulating heterogeneous behaviours, the framework has limited flexibility in handling unstructured information, generating new behavioural variations, or interacting intuitively with human users. It also relies heavily on manual or rule-based labelling, which can be time-consuming and less scalable.

Large Language Models (LLMs) offer promising opportunities to address these

limitations. First, LLMs can support automated behaviour labelling by interpreting textual descriptions of driving scenarios or behavioural patterns, reducing the need for manual annotation and enabling more scalable data processing. Second, in the context of simulation, LLMs can be used to generate diverse driver profiles or behavioural rules based on natural language inputs (e.g., “an aggressive driver in rainy conditions” or “a defensive driver near intersections”), thereby increasing behavioural diversity and realism in traffic simulations. Furthermore, LLMs enable intuitive human-AI interaction, allowing user-defined instructions to influence system behaviour. For example, future adaptive cruise control systems could be enhanced by LLMs that interpret prompts like “prioritise safety during bad weather” or “minimise fuel consumption on the highway” and adjust control strategies accordingly. These capabilities make LLMs a powerful extension to the current framework, supporting more flexible, interpretable, and user-adaptive modelling of driving behaviour.

8.3.2 Practical recommendations and outlook

i) Personalised in-vehicle assistance systems: The action-based framework can enhance current driver assistance systems by tailoring alerts, suggestions, and control strategies to individual driving characteristics. For example, by recognising a driver’s typical behaviour patterns, such systems can deliver more relevant warnings (e.g., earlier alerts for late-braking drivers) and adjust vehicle responses (e.g., smoother acceleration for comfort-oriented drivers). This personalisation can not only improve traffic safety and comfort but also increase driver acceptance and trust in automated driving technologies.

ii) Autonomous vehicle decision-making in mixed traffic: The heterogeneous nature of human driver behaviour poses challenges for the decision-making of autonomous vehicles in mixed traffic. By embedding human driving pattern recognition and prediction into AV decision-making modules, autonomous systems can better forecast human driving behaviours such as sudden braking and risky time headway. This leads to more socially aware and context-sensitive behaviour planning, which is critical for safe and cooperative driving in mixed-traffic environments.

iii) Traffic management and control: The driving behaviour models and simulation approaches developed in this thesis can be adopted by traffic authorities to evaluate the effects of driver heterogeneity on traffic safety and energy efficiency. This can be conducted by designing adaptive traffic management strategies such as behaviour-aware signal control, variable speed limits, and dynamic lane assignments that respond to real-time variations in driver behaviour. Such approaches can improve throughput, reduce congestion, and enhance road safety, especially in diverse traffic conditions.

iv) Policy design and driving education: Insights into how driver behaviour varies across individuals and contexts can inform more targeted and effective road safety policies. For example, training programs could be designed to correct specific risky behavioural patterns (e.g., aggressive following) observed in certain driver groups. Additionally, infrastructure and road design guidelines can be updated to account for behavioural diversity, for instance, by increasing buffer zones in areas prone to inconsistent driving behaviours.

v) Eco-driving and sustainability strategies: Variability in acceleration and

deceleration patterns directly affects energy use and emissions. By modelling these differences across drivers, it becomes possible to develop personalised eco-driving guidance systems and adaptive cruise control strategies that optimise fuel consumption for individual driving behaviour under varying traffic conditions. These insights support broader sustainability goals by informing the design of vehicle technologies and transport policies that promote energy-efficient driving.

Bibliography

- [1] Michail A Makridis, Aikaterini Anesiadou, Konstantinos Mattas, Georgios Fontaras, and Biagio Ciuffo. Characterising driver heterogeneity within stochastic traffic simulation. *Transportmetrica B: Transport Dynamics*, 11(1):725–743, 2023.
- [2] Anusha Adavikottu, Nagendra R Velaga, and Sabyasachee Mishra. Modelling the effect of aggressive driver behavior on longitudinal performance measures during car-following. *Transportation research part F: traffic psychology and behaviour*, 92: 176–200, 2023.
- [3] Emilia M Szumska and Rafał Jurecki. The effect of aggressive driving on vehicle parameters. *Energies*, 13(24):6675, 2020.
- [4] Kemal Ayyildiz, Federico Cavallaro, Silvio Nocera, and Ralf Willenbrock. Reducing fuel consumption and carbon emissions through eco-drive training. *Transportation Research Part F: Traffic Psychology and Behaviour*, 46:96–110, 2017.
- [5] Ni Dong, Shuming Chen, Yina Wu, Yiheng Feng, and Xiaobo Liu. An enhanced motion planning approach by integrating driving heterogeneity and long-term trajectory prediction for automated driving systems: A highway merging case study. *Transportation research part C: emerging technologies*, 161:104554, 2024.
- [6] Xuesong Wang, Rongjiao Xu, Siyang Zhang, Yifan Zhuang, and Yinhai Wang. Driver distraction detection based on vehicle dynamics using naturalistic driving data. *Transportation research part C: emerging technologies*, 136:103561, 2022.
- [7] Xiao Wen, Zhiyong Cui, and Sisi Jian. Characterizing car-following behaviors of human drivers when following automated vehicles using the real-world dataset. *Accident Analysis & Prevention*, 172:106689, 2022.
- [8] Vinit Katariya, Mohammadreza Baharani, Nichole Morris, Omidreza Shoghli, and Hamed Tabkhi. Deeptrack: Lightweight deep learning for vehicle trajectory prediction in highways. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):18927–18936, 2022.
- [9] Zhezhang Ding, Donghao Xu, Chenfeng Tu, Huijing Zhao, Mathieu Moze, François Aioun, and Franck Guillemard. Driver identification through heterogeneity modeling in car-following sequences. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):17143–17156, 2022.

- [10] Amin Mohammadnazar, Ramin Arvin, and Asad J Khattak. Classifying travelers' driving style using basic safety messages generated by connected vehicles: Application of unsupervised machine learning. *Transportation research part C: emerging technologies*, 122:102917, 2021.
- [11] Mingming Zhang, Chao Chen, Tianyu Wo, Tao Xie, Md Zakirul Alam Bhuiyan, and Xuelian Lin. Safedrive: Online driving anomaly detection from large-scale vehicle data. *IEEE Transactions on Industrial Informatics*, 13(4):2087–2096, 2017.
- [12] Jianqiang Nie, Jian Zhang, Xia Wan, Wanting Ding, and Bin Ran. Modeling of decision-making behavior for discretionary lane-changing execution. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pages 707–712. IEEE, 2016.
- [13] Xulei Liu, Yafei Wang, Kun Jiang, Zhisong Zhou, Kanghyun Nam, and Chengliang Yin. Interactive trajectory prediction using a driving risk map-integrated deep learning method for surrounding vehicles on highways. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):19076–19087, 2022.
- [14] Jiaqi Liu, Donghao Zhou, Peng Hang, Ying Ni, and Jian Sun. Towards socially responsive autonomous vehicles: A reinforcement learning framework with driving priors and coordination awareness. *IEEE Transactions on Intelligent Vehicles*, 9(1): 827–838, 2023.
- [15] Martin Treiber and Arne Kesting. Car-following models based on driving strategies. *Traffic Flow Dynamics: Data, Models and Simulation*, pages 181–204, 2013.
- [16] Vasileia Papathanasopoulou and Constantinos Antoniou. Towards data-driven car-following models. *Transportation Research Part C: Emerging Technologies*, 55: 496–509, 2015.
- [17] Dong-Fan Xie, Zhe-Zhe Fang, Bin Jia, and Zhengbing He. A data-driven lane-changing model based on deep learning. *Transportation research part C: emerging technologies*, 106:41–60, 2019.
- [18] Zhanbo Sun, Xue Yao, Ziyue Qin, Peitong Zhang, and Ze Yang. Modeling car-following heterogeneities by considering leader–follower compositions and driving style differences. *Transportation research record*, 2675(11):851–864, 2021.
- [19] Hani S Mahmassani. 50th anniversary invited article—autonomous vehicles and connected vehicle systems: Flow and operations considerations. *Transportation Science*, 50(4):1140–1162, 2016.
- [20] John A Michon. A critical view of driver behavior models: what do we know, what should we do? In *Human behavior and traffic safety*, pages 485–524. Springer, 1985.
- [21] JWC Van Lint and Simeon C Calvert. A generic multi-level framework for microscopic traffic simulation—theory and an example case in modelling driver distraction. *Transportation Research Part B: Methodological*, 117:63–86, 2018.

-
- [22] Shinya Tasaki, Chris Gaiteri, Sara Mostafavi, and Yanling Wang. Deep learning decodes the principles of differential gene expression. *Nature Machine Intelligence*, 2(7):376–386, 2020.
- [23] Md Nasim Khan and Subasish Das. Advancing traffic safety through the safe system approach: A systematic review. *Accident Analysis & Prevention*, 199:107518, 2024.
- [24] Saskia Ossen, Serge P Hoogendoorn, and Ben GH Gorte. Interdriver differences in car-following: A vehicle trajectory-based study. *Transportation Research Record*, 1965(1):121–129, 2006.
- [25] Xue Yao, Qiruo Yan, Zhanbo Sun, Simeon C Calvert, and Serge P Hoogendoorn. Investigation on car-following heterogeneity and its impacts on traffic safety and sustainability. *Transportmetrica A: Transport Science*, pages 1–25, 2024.
- [26] Junjie Zhang, Yunpeng Wang, and Guangquan Lu. Impact of heterogeneity of car-following behavior on rear-end crash risk. *Accident Analysis & Prevention*, 125:275–289, 2019.
- [27] SC Calvert and Bart van Arem. A generic multi-level framework for microscopic traffic simulation with automated vehicles in mixed traffic. *Transportation Research Part C: Emerging Technologies*, 110:291–311, 2020.
- [28] Jooyoung Lee and Kitae Jang. Characterizing driver behavior using naturalistic driving data. *Accident Analysis & Prevention*, 208:107779, 2024.
- [29] G Priyadharshini and JS Femilda Josephin. A comprehensive review of various data collection approaches, features, and algorithms used for the classification of driving style. In *IOP Conference Series: Materials Science and Engineering*, volume 993, page 012098. IOP Publishing, 2020.
- [30] Paulo Fernandes, Elisabete Ferreira, Eloísa Macedo, and Margarida C Coelho. Unraveling roundabout dynamics: Analysis of driving behavior, vehicle performance, and exhaust emissions. *Transportation Research Part D: Transport and Environment*, 133:104308, 2024.
- [31] Yilu Zhang, William C Lin, and Yuen-Kwok Steve Chin. A pattern-recognition approach for driving skill characterization. *IEEE transactions on intelligent transportation systems*, 11(4):905–916, 2010.
- [32] Raymond Hoogendoorn and Bart Van Arem. Driver workload classification through neural network modeling using physiological indicators. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, pages 2268–2273. IEEE, 2013.
- [33] Soukaina Bouhsissin, Nawal Sael, and Faouzia Benabbou. Driver behavior classification: A systematic literature review. *IEEE Access*, 2023.

- [34] Zouhair Elamrani Abou El Assad, Hajar Mousannif, Hassan Al Moatassime, and Aimad Karkouch. The application of machine learning techniques for driving behavior analysis: A conceptual framework and a systematic literature review. *Engineering Applications of Artificial Intelligence*, 87:103312, 2020.
- [35] Na Lin, Changfu Zong, Masayoshi Tomizuka, Pan Song, Zexing Zhang, Gang Li, et al. An overview on study of identification of driver behavior characteristics for automotive control. *Mathematical Problems in Engineering*, 2014, 2014.
- [36] Clara Marina Martinez, Mira Heucke, Fei-Yue Wang, Bo Gao, and Dongpu Cao. Driving style recognition for intelligent vehicle control and advanced driver assistance: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):666–676, 2017.
- [37] Sinan Kaplan, Mehmet Amac Guvensan, Ali Gokhan Yavuz, and Yasin Karalurt. Driver behavior analysis for safe driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 16(6):3017–3032, 2015.
- [38] Dimitrios I Tselentis and Eleonora Papadimitriou. Driver profile and driving pattern recognition for road safety assessment: Main challenges and future directions. *IEEE Open Journal of Intelligent Transportation Systems*, 2023.
- [39] Peng Mei, Hamid Reza Karimi, Lei Ou, Hehui Xie, Chong Zhan, Guangyuan Li, and Shichun Yang. Driving style classification and recognition methods for connected vehicle control in intelligent transportation systems: A review. *ISA transactions*, 2025.
- [40] Wenshuo Wang, Junqiang Xi, and Ding Zhao. Driving style analysis using primitive driving patterns with bayesian nonparametric approaches. *IEEE Transactions on Intelligent Transportation Systems*, 20(8):2986–2998, 2018.
- [41] Iván Silva and José Eugenio Naranjo. A systematic methodology to evaluate prediction models for driving style classification. *Sensors*, 20(6):1692, 2020.
- [42] Yulin Ma, Zhixiong Li, Yicheng Li, Honghai Li, and Reza Malekian. Driving style estimation by fusing multiple driving behaviors: a case study of freeway in china. *Cluster Computing*, 22(4):8259–8269, 2019.
- [43] Montasir Abbas, Linsen Chong, Bryan Higgs, Alejandra Medina, and CY David Yang. Agent-based evaluation of driver heterogeneous behavior during safety-critical events. In *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 1797–1802. IEEE, 2011.
- [44] Yuxiang Feng, Simon Pickering, Edward Chappell, Pejman Irvani, and Chris Brace. A support vector clustering based approach for driving style classification. *International Journal of Machine Learning and Computing*, 9(3):344–350, 2019.

-
- [45] Junhua Wang, Wenxiang Xu, Ting Fu, and Rui Jiang. Recognition of trip-based aggressive driving: A system integrated with gaussian mixture model structured of factor-analysis, and hierarchical clustering. *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [46] Yajie Zou, Ting Zhu, Yuanchang Xie, Yunlong Zhang, and Yue Zhang. Multivariate analysis of car-following behavior data using a coupled hidden markov model. *Transportation research part C: emerging technologies*, 144:103914, 2022.
- [47] Saskia Ossen and Serge P Hoogendoorn. Heterogeneity in car-following behavior: Theory and empirics. *Transportation research part C: emerging technologies*, 19(2): 182–195, 2011.
- [48] Jeffrey Taylor, Xuesong Zhou, Nagui M Roupail, and Richard J Porter. Method for investigating intradriver heterogeneity using vehicle trajectory data: A dynamic time warping approach. *Transportation Research Part B: Methodological*, 73:59–80, 2015.
- [49] Mozghan Nasr Azadani and Azzedine Boukerche. Driving behavior analysis guidelines for intelligent transportation systems. *IEEE transactions on intelligent transportation systems*, 23(7):6027–6045, 2021.
- [50] Fridulv Sagberg, Selpi, Giulio Francesco Bianchi Piccinini, and Johan Engström. A review of research on driving styles and road safety. *Human factors*, 57(7):1248–1275, 2015.
- [51] Mohammad Mahdi Bejani and Mehdi Ghatee. Convolutional neural network with adaptive regularization to classify driving styles on smartphones. *IEEE transactions on intelligent transportation systems*, 21(2):543–552, 2019.
- [52] Weirong Liu, Kunyuan Deng, Xiaoyong Zhang, Yijun Cheng, Zhiyong Zheng, Fu Jiang, and Jun Peng. A semi-supervised tri-catboost method for driving style recognition. *Symmetry*, 12(3):336, 2020.
- [53] Kaichong Liang, Zhiguo Zhao, Wenchang Li, Jiawei Zhou, and Danshu Yan. Comprehensive identification of driving style based on vehicle’s driving cycle recognition. *IEEE Transactions on Vehicular Technology*, 72(1):312–326, 2022.
- [54] Naiwala P Chandrasiri, Kazunari Nawa, and Akira Ishii. Driving skill classification in curve driving scenes using machine learning. *Journal of Modern Transportation*, 24:196–206, 2016.
- [55] Bing Zhu, Zhipeng Liu, Jian Zhao, Yizhou Chen, and Weiwen Deng. Driver behavior characteristics identification strategies based on bionic intelligent algorithms. *IEEE Transactions on Human-Machine Systems*, 48(6):572–581, 2018.
- [56] Geqi Qi, Jianping Wu, Yang Zhou, Yiman Du, Yuhan Jia, Nick Hounsell, and Neville A Stanton. Recognizing driving styles based on topic models. *Transportation research part D: transport and environment*, 66:13–22, 2019.

- [57] Xue Yao, Simeon C Calvert, and Serge P Hoogendoorn. Identification of driving heterogeneity using action-chains. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pages 6001–6006. IEEE, 2023.
- [58] Xue Yao, Simeon C Calvert, and Serge P Hoogendoorn. A novel framework for identifying driving heterogeneity through action patterns. *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [59] HaiLong Liu, Tadahiro Taniguchi, Yusuke Tanaka, Kazuhito Takenaka, and Takashi Bando. Visualization of driving behavior based on hidden feature extraction by using deep learning. *IEEE Transactions on Intelligent Transportation Systems*, 18(9): 2477–2489, 2017.
- [60] Bryan Higgs and Montasir Abbas. A two-step segmentation algorithm for behavioral clustering of naturalistic driving styles. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, pages 857–862. IEEE, 2013.
- [61] Haluk Eren, Semiha Makinist, Erhan Akin, and Alper Yilmaz. Estimating driving behavior by a smartphone. In *2012 IEEE Intelligent Vehicles Symposium*, pages 234–239. IEEE, 2012.
- [62] Dong-Fan Xie, Tai-Lang Zhu, and Qian Li. Capturing driving behavior heterogeneity based on trajectory data. *International Journal of Modeling, Simulation, and Scientific Computing*, 11(03):2050023, 2020.
- [63] Xiaoyun Chen, Jian Sun, Zian Ma, Jie Sun, and Zuduo Zheng. Investigating the long-and short-term driving characteristics and incorporating them into car-following models. *Transportation research part C: emerging technologies*, 117: 102698, 2020.
- [64] Yongfeng Ma, Wenlu Li, Kun Tang, Ziyu Zhang, and Shuyan Chen. Driving style recognition and comparisons among driving tasks based on driver behavior in the online car-hailing industry. *Accident Analysis & Prevention*, 154:106096, 2021.
- [65] Laurette Guyonvarch, Thierry Hermitte, Frederic Duvivier, Clement Val, and Anne Guillaume. Driving style indicator using udrive nds data. *Traffic injury prevention*, 19(sup1):S189–S191, 2018.
- [66] Nengchao Lyu, Yugang Wang, Chaozhong Wu, Lingfeng Peng, and Alieu Freddie Thomas. Using naturalistic driving data to identify driving style based on longitudinal driving operation conditions. *Journal of intelligent and connected vehicles*, 5(1):17–35, 2022.
- [67] Khaled Saleh, Mohammed Hossny, and Saeid Nahavandi. Driving behavior classification based on sensor data fusion using lstm recurrent neural networks. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6. IEEE, 2017.

-
- [68] Ikki Kim, Taewan Kim, and Keemin Sohn. Identifying driver heterogeneity in car-following based on a random coefficient model. *Transportation Research Part C: Emerging Technologies*, 36:35–44, 2013.
- [69] Bin Shi, Li Xu, Jie Hu, Yun Tang, Hong Jiang, Wuqiang Meng, and Hui Liu. Evaluating driving styles by normalizing driving behavior based on personalized driver modeling. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(12):1502–1508, 2015.
- [70] Sobhan Moosavi, Pravar D Mahajan, Srinivasan Parthasarathy, Colleen Saunders-Chukwu, and Rajiv Ramnath. Driving style representation in convolutional recurrent neural network model of driver identification. *arXiv preprint arXiv:2102.05843*, 2021.
- [71] Minh Van Ly, Sujitha Martin, and Mohan M Trivedi. Driver classification and driving style recognition using inertial sensors. In *2013 IEEE Intelligent Vehicles Symposium (IV)*, pages 1040–1045. IEEE, 2013.
- [72] Vygandas Vaitkus, Paulius Lengvenis, and Gediminas Žylius. Driving style classification using long-term accelerometer information. In *2014 19th International Conference on Methods and Models in Automation and Robotics (MMAR)*, pages 641–644. IEEE, 2014.
- [73] Bernhard Gahr, Benjamin Ryder, Andre Dahlinger, and Felix Wortmann. Driver identification via brake pedal signals—a replication and advancement of existing techniques. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 1415–1420. IEEE, 2018.
- [74] Qingwen Xue, Ke Wang, Jian John Lu, and Yujie Liu. Rapid driving style recognition in car-following using machine learning and vehicle trajectory data. *Journal of advanced transportation*, 2019, 2019.
- [75] Hamid Reza Eftekhari and Mehdi Ghatee. Hybrid of discrete wavelet transform and adaptive neuro fuzzy inference system for overall driving behavior recognition. *Transportation research part F: traffic psychology and behaviour*, 58:782–796, 2018.
- [76] Xinhu Zheng, Pengtao Yang, Dongliang Duan, Xiang Cheng, and Liuqing Yang. Real-time driving style classification based on short-term observations. *IET Communications*, 2022.
- [77] Zejian Deng, Duanfeng Chu, Chaozhong Wu, Shidong Liu, Chen Sun, Teng Liu, and Dongpu Cao. A probabilistic model for driving-style-recognition-enabled driver steering behaviors. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 52(3):1838–1851, 2020.
- [78] Aurenice Cruz Figueira and Ana Paula C Larocca. Proposal of a driver profile classification in relation to risk level in overtaking maneuvers. *Transportation research part F: traffic psychology and behaviour*, 74:375–385, 2020.

- [79] Abdul Wahab, Chai Quek, Chin Keong Tan, and Kazuya Takeda. Driving profile modeling and recognition based on soft computing approach. *IEEE transactions on neural networks*, 20(4):563–582, 2009.
- [80] Jingqiu Guo, Yangzexi Liu, Lanfang Zhang, and Yibing Wang. Driving behaviour style study with a hybrid deep learning framework based on gps data. *Sustainability*, 10(7):2351, 2018.
- [81] Youness Moukafih, Hakim Hafidi, and Mounir Ghogho. Aggressive driving detection using deep learning-based time series classification. In *2019 IEEE International Symposium on INnovations in Intelligent SysTems and Applications (INISTA)*, pages 1–5. IEEE, 2019.
- [82] Xidong Tang. Driving skill recognition: new approaches and their comparison. In *2009 American control conference*, pages 420–425. IEEE, 2009.
- [83] Guofa Li, Shengbo Eben Li, Bo Cheng, and Paul Green. Estimation of driving style in naturalistic highway traffic using maneuver transition probabilities. *Transportation Research Part C: Emerging Technologies*, 74:113–125, 2017.
- [84] Sasan Jafarnejad, German Castignani, and Thomas Engel. Towards a real-time driver identification mechanism based on driving sensing data. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7. IEEE, 2017.
- [85] Fabio Tango and Marco Botta. Real-time detection system of driver distraction using machine learning. *IEEE Transactions on Intelligent Transportation Systems*, 14(2): 894–905, 2013.
- [86] Guofa Li, Fangping Zhu, Xingda Qu, Bo Cheng, Shen Li, and Paul Green. Driving style classification based on driving operational pictures. *IEEE Access*, 7:90180–90189, 2019.
- [87] Mohammad Shahverdy, Mahmood Fathy, Reza Berangi, and Mohammad Sabokrou. Driver behavior detection and classification using deep convolutional neural networks. *Expert Systems with Applications*, 149:113240, 2020.
- [88] Jie Xie, Kai Hu, Guofa Li, and Ya Guo. Cnn-based driving maneuver classification using multi-sliding window fusion. *Expert Systems with Applications*, 169:114442, 2021.
- [89] Najmeddine Abdennour, Tarek Ouni, and N Ben Amor. Driver identification using only the can-bus vehicle data through an rcn deep learning approach. *Robotics and Autonomous Systems*, 136:103707, 2021.
- [90] Zehra Camlica, Jim Quesenberry, Daniel Carballo, and Mark Crowley. Aggressive driver behavior detection using parallel convolutional neural networks on simulated and real driving data. In *2022 9th International Conference on Internet of Things: Systems, Management and Security (IOTSMS)*, pages 1–8. IEEE, 2022.

-
- [91] Moayed A Khodairy and Gibrael Abosamra. Driving behavior classification based on oversampled signals of smartphone embedded sensors using an optimized stacked-lstm neural networks. *IEEE Access*, 9:4957–4972, 2021.
- [92] Jasper S Wijnands, Jason Thompson, Gideon DPA Aschwanden, and Mark Stevenson. Identifying behavioural change among drivers using long short-term memory recurrent neural networks. *Transportation research part F: traffic psychology and behaviour*, 53:34–49, 2018.
- [93] Kenny Schlegel, Florian Mirus, Peer Neubert, and Peter Protzel. Multivariate time series analysis for driving style classification using neural networks and hyperdimensional computing. In *2021 IEEE Intelligent Vehicles Symposium (IV)*, pages 602–609. IEEE, 2021.
- [94] Lin Lu, Yao Lin, Yuan Wen, Jinxiong Zhu, and Shengwu Xiong. Federated clustering for recognizing driving styles from private trajectories. *Engineering Applications of Artificial Intelligence*, 118:105714, 2023.
- [95] Lingyun Xie, Ying Shi, and Zhenwei Li. Driving pattern recognition based on improved lda model. In *2018 5th IEEE International Conference on Cloud Computing and Intelligence Systems (CCIS)*, pages 320–324. IEEE, 2018.
- [96] Wenshuo Wang, Wei Han, Xiaoxiang Na, Jianwei Gong, and Junqiang Xi. A probabilistic approach to measuring driving behavior similarity with driving primitives. *IEEE Transactions on Intelligent Vehicles*, 5(1):127–138, 2019.
- [97] Ryunosuke Hamada, Takatomi Kubo, Kazushi Ikeda, Zujie Zhang, Tomohiro Shibata, Takashi Bando, Kentarou Hitomi, and Masumi Egawa. Modeling and prediction of driving behaviors using a nonparametric bayesian method with ar models. *IEEE Transactions on Intelligent Vehicles*, 1(2):131–138, 2016.
- [98] Tadahiro Taniguchi, Shogo Nagasaka, Kentarou Hitomi, Kazuhito Takenaka, and Takashi Bando. Unsupervised hierarchical modeling of driving behavior and prediction of contextual changing points. *IEEE Transactions on Intelligent Transportation Systems*, 16(4):1746–1760, 2014.
- [99] Tadahiro Taniguchi, Shogo Nagasaka, Kentarou Hitomi, Naiwala P Chandrasiri, Takashi Bando, and Kazuhito Takenaka. Sequence prediction of driving behavior using double articulation analyzer. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 46(9):1300–1313, 2015.
- [100] Asher Bender, Gabriel Agamennoni, James R Ward, Stewart Worrall, and Eduardo M Nebot. An unsupervised approach for inferring driver behavior from naturalistic driving data. *IEEE transactions on intelligent transportation systems*, 16(6):3325–3336, 2015.
- [101] Wenshuo Wang, Junqiang Xi, Alexandre Chong, and Lin Li. Driving style classification using a semisupervised support vector machine. *IEEE Transactions on Human-Machine Systems*, 47(5):650–660, 2017.

- [102] Dimitris M Vlachogiannis, Eleni I Vlahogianni, and John Golias. A reinforcement learning model for personalized driving policies identification. *International journal of transportation science and technology*, 9(4):299–308, 2020.
- [103] Yuande Jiang, Weiwen Deng, Jinsong Wang, and Bing Zhu. Studies on drivers’ driving styles based on inverse reinforcement learning. Technical report, SAE Technical Paper, 2018.
- [104] Daiko Kishikawa and Sachiyo Arai. Estimation of personal driving style via deep inverse reinforcement learning. *Artificial Life and Robotics*, 26(3):338–346, 2021.
- [105] Sascha Rosbach, Vinit James, Simon Großjohann, Silviu Homoceanu, and Stefan Roth. Driving with style: Inverse reinforcement learning in general-purpose planning for automated driving. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2658–2665. IEEE, 2019.
- [106] Nathanael C Fung, Bruce Wallace, Adrian DC Chan, Rafik Goubran, Michelle M Porter, Shawn Marshall, and Frank Knoefel. Driver identification using vehicle acceleration and deceleration events from naturalistic driving of older drivers. In *2017 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, pages 33–38. IEEE, 2017.
- [107] Yongfeng Ma, Ziyu Zhang, Shuyan Chen, Yanan Yu, and Kun Tang. A comparative study of aggressive driving behavior recognition algorithms based on vehicle motion data. *IEEE Access*, 7:8028–8038, 2018.
- [108] Xingjian Zhang, Xiaohua Zhao, and Jian Rong. A study of individual characteristics of driving behavior based on hidden markov model. *Sensors & Transducers*, 167(3):194, 2014.
- [109] Seong Kyung Kwon, Ji Hwan Seo, Jun Young Yun, and Kyoung-Dae Kim. Driving behavior classification and sharing system using cnn-lstm approaches and v2x communication. *Applied Sciences*, 11(21):10420, 2021.
- [110] Bryan Higgs and Montasir Abbas. Segmentation and clustering of car-following behavior: Recognition of driving patterns. *IEEE Transactions on Intelligent Transportation Systems*, 16(1):81–90, 2014.
- [111] Hongqing Chu, Hejian Zhuang, Wenshuo Wang, Xiaoxiang Na, Lulu Guo, Jia Zhang, Bingzhao Gao, and Hong Chen. A review of driving style recognition methods from short-term and long-term perspectives. *IEEE Transactions on Intelligent Vehicles*, 2023.
- [112] Waymo. The waymo driver: A comprehensive sensor suite for autonomous driving. <https://waymo.com/waymo-driver/>, 2024. Accessed: 2024-03-22.
- [113] Mobileye. True redundancy: Two independent perception subsystems. <https://www.mobileye.com/technology/true-redundancy/>, 2024. Accessed: 2024-03-22.

-
- [114] U.S. Department of Transportation, Federal Highway Administration. Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data. <https://ops.fhwa.dot.gov/trafficanalysistools/ngsim.htm>, 2006. Accessed: 2022-03-25.
- [115] Robert Krajewski, Julian Bock, Laurent Kloeker, and Lutz Eckstein. 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)*, pages 2118–2125, 2018. doi: 10.1109/ITSC.2018.8569552.
- [116] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [117] Yuning Wang, Zeyu Han, Yining Xing, Shaobing Xu, and Jianqiang Wang. A survey on datasets for decision-making of autonomous vehicle. *arXiv preprint arXiv:2306.16784*, 2023.
- [118] William Villegas-Ch and Joselin García-Ortiz. Toward a comprehensive framework for ensuring security and privacy in artificial intelligence. *Electronics*, 12(18):3786, 2023.
- [119] Dengfeng Zhao, Yudong Zhong, Zhijun Fu, Junjian Hou, and Mingyuan Zhao. A review for the driving behavior recognition methods based on vehicle multisensor information. *Journal of Advanced Transportation*, 2022(1):7287511, 2022.
- [120] Weichu Xiao, Hongli Liu, Ziji Ma, and Weihong Chen. Attention-based deep neural network for driver behavior recognition. *Future Generation Computer Systems*, 132:152–161, 2022.
- [121] Yiyun Wang, Haneen Farah, Rongjie Yu, Shuhan Qiu, and Bart van Arem. Characterizing behavioral differences of autonomous vehicles and human-driven vehicles at signalized intersections based on waymo open dataset. *Transportation research record*, 2677(11):324–337, 2023.
- [122] John D Lee and Katrina A See. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1):50–80, 2004.
- [123] Zheng Ma and Yiqi Zhang. Drivers trust, acceptance, and takeover behaviors in fully automated vehicles: Effects of automated driving styles and driver’s driving styles. *Accident Analysis & Prevention*, 159:106238, 2021.
- [124] Joonbum Lee, Pnina Gershon, Josh Domeyer, Bruce Mehler, and Bryan Reimer. Understanding hands-off patterns while using tesla autopilot: Individual differences and their impact on hands-off duration. In *THE 9TH INTERNATIONAL CONFERENCE ON DRIVER DISTRACTION AND INATTENTION*, page 26, 2024.

- [125] Danjue Chen, Soyoung Ahn, Jorge Laval, and Zuduo Zheng. On the periodicity of traffic oscillations and capacity drop: The role of driver characteristics. *Transportation Research Part B: Methodological*, 59:117–136, 2014.
- [126] Yao Chen, Ke Wang, and Jian John Lu. Feature selection for driving style and skill clustering using naturalistic driving data and driving behavior questionnaire. *Accident Analysis & Prevention*, 185:107022, 2023.
- [127] Elísabet Björney Lárusdóttir and Gudmundur F Ulfarsson. Effect of driving behavior and vehicle characteristics on energy consumption of road vehicles running on alternative energy sources. *International Journal of Sustainable Transportation*, 9(8): 592–601, 2015.
- [128] Michail Makridis, Ludovic Leclercq, Biagio Ciuffo, Georgios Fontaras, and Konstantinos Mattas. Formalizing the heterogeneity of the vehicle-driver system to reproduce traffic oscillations. *Transportation Research Part C: Emerging Technologies*, 120:102803, 2020.
- [129] Xue Yao, Zhaocheng Du, Zhanbo Sun, Simeon C Calvert, and Ang Ji. Cooperative lane-changing in mixed traffic: a deep reinforcement learning approach. *Transportmetrica A: Transport Science*, pages 1–23, 2024.
- [130] Yanlin Zhang and Alireza Talebpour. Characterizing human–automated vehicle interactions: An investigation into car-following behavior. *Transportation research record*, 2678(5):812–826, 2024.
- [131] Fatima Haque and Mohd Azman Abas. Review of driving behavior towards fuel consumption and road safety. *Jurnal Mekanikal*, 2018.
- [132] Zhanbo Sun, Qiruo Yan, Yafei Liu, Zhijian Fu, and Lei Yang. Fundamental diagram and stability analysis of mixed traffic considering heterogeneous car-following behaviors and platoon factors. *Intelligent Transportation Infrastructure*, 3:liae010, 2024.
- [133] Karim Fadhloun and Hesham Rakha. A novel vehicle dynamics and human behavior car-following model: Model development and preliminary testing. *International journal of transportation science and technology*, 9(1):14–28, 2020.
- [134] Mingfeng Shang and Raphael Stern. Calibrating heterogeneous car-following models for human drivers in oscillatory traffic conditions. In *2020 Forum on Integrated and Sustainable Transportation Systems (FISTS)*, pages 101–106. IEEE, 2020.
- [135] Filippos K Adamidis, Eleni G Mantouka, and Eleni I Vlahogianni. Effects of controlling aggressive driving behavior on network-wide traffic flow and emissions. *International journal of transportation science and technology*, 9(3):263–276, 2020.
- [136] Kristin Bennett and Ayhan Demiriz. Semi-supervised support vector machines. *Advances in Neural Information processing systems*, 11, 1998.

-
- [137] Arne Kesting, Martin Treiber, and Dirk Helbing. Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1928): 4585–4605, 2010.
- [138] Jaime Suarez, Michail Makridis, Aikaterini Anesiadou, Dimitrios Komnos, Biagio Ciuffo, and Georgios Fontaras. Benchmarking the driver acceleration impact on vehicle energy consumption and co2 emissions. *Transportation Research Part D: Transport and Environment*, 107:103282, 2022.
- [139] Shi-Teng Zheng, Rui Jiang, Bin Jia, Junfang Tian, Marouane Bouadi, Michail A Makridis, and Anastasios Kouvelas. A parsimonious enhanced newell’s model for accurate reproduction of driver and traffic dynamics. *Transportation Research Part C: Emerging Technologies*, 154:104276, 2023.
- [140] Jinghui Wang, Hesham A Rakha, and Karim Fadhoun. Validation of the rakha-pasumarthy-adjerid car-following model for vehicle fuel consumption and emission estimation applications. *Transportation Research Part D: Transport and Environment*, 55:246–261, 2017.
- [141] Zhanbo Sun, Qiruo Yan, Yafei Liu, Zhijian Fu, and Lei Yang. Fundamental diagram and stability analysis of mixed traffic considering heterogeneous car-following behaviors and platoon factors. *Intelligent Transportation Infrastructure*, 3:liae010, 07 2024.
- [142] Emilia M. Szumska and Rafał Jurecki. The effect of aggressive driving on vehicle parameters. *Energies*, 13(24), 2020.
- [143] Peter Bakhit, Dalia Said, and Laila Radwan. Impact of acceleration aggressiveness on fuel consumption using comprehensive power based fuel consumption model. *Civ Environ Res*, 7(3):148–156, 2015.
- [144] Georgios Fontaras, Nikiforos-Georgios Zacharof, and Biagio Ciuffo. Fuel consumption and co2 emissions from passenger cars in europe–laboratory versus real-world emissions. *Progress in energy and combustion Science*, 60:97–131, 2017.
- [145] Robert Krajewski, Julian Bock, Laurent Kloeker, and Lutz Eckstein. 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)*, pages 2118–2125. IEEE, 2018.
- [146] Kayvan Aghabayk, Majid Sarvi, and William Young. Including heavy vehicles in a car-following model: modelling, calibrating and validating. *Journal of Advanced Transportation*, 50(7):1432–1446, 2016.
- [147] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Microscopic simulation of congested traffic. In *Traffic and Granular Flow’99: Social, Traffic, and Granular Dynamics*, pages 365–376. Springer, 2000.

- [148] Martin Treiber and Arne Kesting. The intelligent driver model with stochasticity-new insights into traffic flow oscillations. *Transportation Research Procedia*, 23:174–187, 2017.
- [149] Dong Ngoduy. Analytical studies on the instabilities of heterogeneous intelligent traffic flow. *Communications in Nonlinear Science and Numerical Simulation*, 18(10): 2699–2706, 2013.
- [150] Dong Ngoduy. Linear stability of a generalized multi-anticipative car following model with time delays. *Communications in Nonlinear Science and Numerical Simulation*, 22 (1-3):420–426, 2015.
- [151] Mingfeng Shang and Raphael E Stern. Impacts of commercially available adaptive cruise control vehicles on highway stability and throughput. *Transportation Research Part C: Emerging Technologies*, 122:102897, 2021.
- [152] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical review E*, 62(2):1805, 2000.
- [153] Yong-Xian Huang, Rui Jiang, HM Zhang, Mao-Bin Hu, Jun-Fang Tian, Bin Jia, and Zi-You Gao. Experimental study and modeling of car-following behavior under high speed situation. *Transportation research part C: emerging technologies*, 97:194–215, 2018.
- [154] Serge P Hoogendoorn and Raymond Hoogendoorn. Generic calibration framework for joint estimation of car-following models by using microscopic data. *Transportation Research Record*, 2188(1):37–45, 2010.
- [155] Xiaofei He, Deng Cai, and Partha Niyogi. Laplacian score for feature selection. *Advances in neural information processing systems*, 18, 2005.
- [156] Zhanbo Sun and Xuegang Ban. Identifying multiclass vehicles using global positioning system data. *Journal of Intelligent Transportation Systems*, 22(1):1–9, 2018.
- [157] Fabian Gieseke, Antti Airola, Tapio Pahikkala, and Oliver Kramer. Fast and simple gradient-based optimization for semi-supervised support vector machines. *Neurocomputing*, 123:23–32, 2014.
- [158] Luc Int Panis, Steven Broekx, and Ronghui Liu. Modelling instantaneous traffic emission and the influence of traffic speed limits. *Science of The Total Environment*, 371(1):270–285, 2006.
- [159] Michiel M Minderhoud and Piet HL Bovy. Extended time-to-collision measures for road traffic safety assessment. *Accident Analysis & Prevention*, 33(1):89–97, 2001.

-
- [160] Zhibin Li, Seongchae Ahn, Koohong Chung, David R Ragland, Wei Wang, and Jeong Whon Yu. Surrogate safety measure for evaluating rear-end collision risk related to kinematic waves near freeway recurrent bottlenecks. *Accident Analysis & Prevention*, 64:52–61, 2014.
- [161] Weijie Yu, Xuedong Hua, Dong Ngoduy, and Wei Wang. On the assessment of the dynamic platoon and information flow topology on mixed traffic flow under connected environment. *Transportation Research Part C: Emerging Technologies*, 154: 104265, 2023.
- [162] Manjiang Hu, Chongkang Li, Yougang Bian, Hui Zhang, Zhaobo Qin, and Biao Xu. Fuel economy-oriented vehicle platoon control using economic model predictive control. *IEEE Transactions on Intelligent Transportation Systems*, 23(11):20836–20849, 2022.
- [163] Guohua Song and Lei Yu. Estimation of fuel efficiency of road traffic by characterization of vehicle-specific power and speed based on floating car data. *Transportation Research Record*, 2139(1):11–20, 2009.
- [164] H Christopher Frey, Nagui M Roupail, and Haibo Zhai. Speed-and facility-specific emission estimates for on-road light-duty vehicles on the basis of real-world speed profiles. *Transportation Research Record*, 1987(1):128–137, 2006.
- [165] Amin Mohammadnazar, Ramin Arvin, and Asad J. Khattak. Classifying travelers’ driving style using basic safety messages generated by connected vehicles: Application of unsupervised machine learning. *Transportation Research Part C: Emerging Technologies*, 122:102917, 2021.
- [166] Xue Yao, Simeon C Calvert, and Serge P Hoogendoorn. Driving pattern interpretation based on action phases clustering. *arXiv preprint arXiv:2407.17518*, 2024.
- [167] Xue Yao, Simeon C. Calvert, and Serge P. Hoogendoorn. Identification of driving heterogeneity with machine learning: A review. *IEEE Transactions on Intelligent Transportation Systems*, 00:00, 2023, under review.
- [168] Victor L Knoop and Serge P Hoogendoorn. Relation between longitudinal and lateral action points. In *Traffic and Granular Flow’13*, pages 571–576. Springer, 2015.
- [169] Xue Yao, Zhanbo Sun, Qiruo Yan, and C. Simeon Calvert. Performance comparison of multi-class svm and s3vm for driving style classification. *IET Intelligent Transport Systems*, 00:00, 2023, under review.
- [170] Patrick Billingsley. *Probability and measure*. Wiley–Interscience, New York, 3 edition, 1995.
- [171] Amro Elfeki and Michel Dekking. A markov chain model for subsurface characterization: theory and applications. *Mathematical geology*, 33:569–589, 2001.

- [172] Wenshuo Wang, Chang Liu, and Ding Zhao. How much data are enough? a statistical approach with case study on longitudinal driving behavior. *IEEE Transactions on Intelligent Vehicles*, 2(2):85–98, 2017.
- [173] Thomas A Dingus, Vicki L Neale, Sheila G Klauer, Andrew D Petersen, and Robert J Carroll. The development of a naturalistic data collection system to perform critical incident analysis: An investigation of safety and fatigue issues in long-haul trucking. *Accident Analysis & Prevention*, 38(6):1127–1136, 2006.
- [174] Boris S Kerner and Sergey L Klenov. Spatial–temporal patterns in heterogeneous traffic flow with a variety of driver behavioural characteristics and vehicle parameters. *Journal of Physics A: Mathematical and General*, 37(37):8753, 2004.
- [175] Virginia Petraki, Apostolos Ziakopoulos, and George Yannis. Combined impact of road and traffic characteristic on driver behavior using smartphone sensor data. *Accident Analysis & Prevention*, 144:105657, 2020.
- [176] Xue Yao, Qiruo Yan, Zhanbo Sun, C. Calvert, Simeon, and P. Hoogendoorn, Serge. Investigation on car-following heterogeneity and its impacts on traffic safety and sustainability. *Transportmetrica A: Transport Science*, 00:00, 2024.
- [177] Xue Yao, Simeon C Calvert, and Serge P Hoogendoorn. Driving heterogeneity identification using machine learning: A review and framework for analysis. *Transportation Research Interdisciplinary Perspectives*, 32:101511, 2025.
- [178] Dajun Wang, Xin Pei, Li Li, and Danya Yao. Risky driver recognition based on vehicle speed time series. *IEEE Transactions on Human-Machine Systems*, 48(1):63–71, 2017.
- [179] Beakcheol Jang, Myeonghwi Kim, Gaspard Harerimana, Sang-ug Kang, and Jong Wook Kim. Bi-lstm model to increase accuracy in text classification: Combining word2vec cnn and attention mechanism. *Applied Sciences*, 10(17):5841, 2020.
- [180] Haipeng Yao, Chong Liu, Peiyong Zhang, Sheng Wu, Chunxiao Jiang, and Shui Yu. Identification of encrypted traffic through attention mechanism based long short term memory. *IEEE transactions on big data*, 8(1):241–252, 2019.
- [181] Xue Yao, Simeon C. Calvert, and Serge P. Hoogendoorn. Identification of driving heterogeneity using action-chains. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pages 6001–6006, 2023. doi: 10.1109/ITSC57777.2023.10421850.
- [182] Guopeng Li, Yiru Jiao, Victor L Knoop, Simeon C Calvert, and JWC Van Lint. Large car-following data based on lyft level-5 open dataset: Following autonomous vehicles vs. human-driven vehicles. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pages 5818–5823. IEEE, 2023.
- [183] Marc Green. " how long does it take to stop?" methodological analysis of driver perception-brake times. *Transportation human factors*, 2(3):195–216, 2000.

-
- [184] TD Gillespie. Fundamentals of vehicle dynamics. society of automotive engineers (sae). Inc., ISBN-13, pages 978–1560911999, 1992.
- [185] Chris Fraley and Adrian E Raftery. How many clusters? which clustering method? answers via model-based cluster analysis. *The computer journal*, 41(8):578–588, 1998.
- [186] Evan Ackerman. How drive. ai is mastering autonomous driving with deep learning. *IEEE Spectrum Magazine*, 1, 2017.
- [187] Tin T Nguyen, Panchamy Krishnakumari, Simeon C Calvert, Hai L Vu, and Hans Van Lint. Feature extraction and clustering analysis of highway congestion. *Transportation Research Part C: Emerging Technologies*, 100:238–258, 2019.
- [188] Dan Pelleg and Andrew W Moore. X-means: Extending k-means with efficient estimation of the number of clusters. In *Icml*, volume 1, pages 727–734, 2000.
- [189] Xu Wang and Yusheng Xu. An improved index for clustering validation based on silhouette index and calinski-harabasz index. In *IOP Conference Series: Materials Science and Engineering*, volume 569, page 052024. IOP Publishing, 2019.
- [190] Q.H. Liu, N. Nguyen, and X.Y. Tang. Accurate algorithms for nonuniform fast forward and inverse fourier transforms and their applications. In *IGARSS '98. Sensing and Managing the Environment. 1998 IEEE International Geoscience and Remote Sensing Symposium Proceedings. (Cat. No.98CH36174)*, volume 1, pages 288–290 vol.1, 1998. doi: 10.1109/IGARSS.1998.702881.
- [191] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- [192] Sung-Hyun Yoon and Ha-Jin Yu. A simple distortion-free method to handle variable length sequences for recurrent neural networks in text dependent speaker verification. *Applied Sciences*, 10(12):4092, 2020.
- [193] International Transport Forum (ITF). Road safety annual report 2024. Technical report, OECD Publishing, 2024. URL <https://www.itf-oecd.org/sites/default/files/docs/irtad-road-safety-annual-report-2024.pdf>. Accessed: July 18, 2025.
- [194] Vincenzo Punzo, Zuduo Zheng, and Marcello Montanino. About calibration of car-following dynamics of automated and human-driven vehicles: Methodology, guidelines and codes. *Transportation Research Part C: Emerging Technologies*, 128: 103165, 2021.
- [195] Femke van Wageningen-Kessels, Hans Van Lint, Kees Vuik, and Serge Hoogendoorn. Genealogy of traffic flow models. *EURO Journal on Transportation and Logistics*, 4(4): 445–473, 2015.

- [196] Chengyuan Zhang, Wenshuo Wang, and Lijun Sun. Calibrating car-following models via bayesian dynamic regression. *Transportation Research Part C: Emerging Technologies*, 168:104719, 2024.
- [197] Alireza Khodayari, Ali Ghaffari, Reza Kazemi, and Reinhard Brauningl. A modified car-following model based on a neural network model of the human driver effects. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 42(6):1440–1449, 2012.
- [198] Dong Zhou, Hongyi Liu, Huimin Ma, Xiang Wang, Xiaoqin Zhang, and Yuhan Dong. Driving behavior prediction considering cognitive prior and driving context. *IEEE Transactions on Intelligent Transportation Systems*, 22(5):2669–2678, 2021. doi: 10.1109/TITS.2020.2973751.
- [199] Xiuling Huang, Jie Sun, and Jian Sun. A car-following model considering asymmetric driving behavior based on long short-term memory neural networks. *Transportation research part C: emerging technologies*, 95:346–362, 2018.
- [200] Zhezhang Ding and Huijing Zhao. Incorporating driving knowledge in deep learning based vehicle trajectory prediction: A survey. *IEEE Transactions on Intelligent Vehicles*, 8(8):3996–4015, 2023.
- [201] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [202] Lei Lin, Weizi Li, Huikun Bi, and Lingqiao Qin. Vehicle trajectory prediction using lstms with spatial–temporal attention mechanisms. *IEEE Intelligent Transportation Systems Magazine*, 14(2):197–208, 2021.
- [203] Lian Hou, Long Xin, Shengbo Eben Li, Bo Cheng, and Wenjun Wang. Interactive trajectory prediction of surrounding road users for autonomous driving using structural-lstm network. *IEEE Transactions on Intelligent Transportation Systems*, 21(11):4615–4625, 2019.
- [204] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. doi: 10.1109/CVPR.2018.00813.
- [205] Friedrich Pukelsheim. The three sigma rule. *The American Statistician*, 48(2):88–91, 1994.
- [206] Hill Peter D. Kernel estimation of a distribution function. *Communications in Statistics-Theory and Methods*, 14(3):605–620, 1985.
- [207] Saskia Ossen and Serge P Hoogendoorn. Driver heterogeneity in car following and its impact on modeling traffic dynamics. *Transportation Research Record*, 1999(1): 95–103, 2007.

-
- [208] Martin Treiber, Arne Kesting, and Dirk Helbing. Three-phase traffic theory and two-phase models with a fundamental diagram in the light of empirical stylized facts. *Transportation Research Part B: Methodological*, 44(8-9):983–1000, 2010.
- [209] Hui Liu, Meng Li, Hongyang Liu, Jing Wu, and Xiaobo Qu. Impact of driver heterogeneity on traffic flow dynamics and lane-changing behaviors: A naturalistic driving study. *Transportation Research Part C: Emerging Technologies*, 128:103130, 2021. doi: 10.1016/j.trc.2021.103130.
- [210] Xiaobo Niu, Meng Li, Xiaoxu Ma, Wenlong Wang, Chao Xu, and Hui Liu. Eco-approach and departure strategies for connected and automated vehicles at signalized intersections: A survey. *Transportation Research Part C: Emerging Technologies*, 96:121–143, 2018. doi: 10.1016/j.trc.2018.09.009.
- [211] Fatemeh Fakhroosavi, Ramin Saedi, Ali Zockaie, and Alireza Talebpour. Impacts of connected and autonomous vehicles on traffic flow with heterogeneous drivers spatially distributed over large-scale networks. *Transportation research record*, 2674(10):817–830, 2020.
- [212] Lin Zhang, Yujie Bian, Lei Li, and Xiaoxu Ma. Classification of driver lane-changing behavior based on driving style. *Transportation Research Part F: Traffic Psychology and Behaviour*, 58:380–397, 2018. doi: 10.1016/j.trf.2018.06.037.
- [213] Tomer Toledo, Haris N Koutsopoulos, and Moshe E Ben-Akiva. Calibration and validation of microscopic traffic simulation models: Methodology and application. *Transportation Research Part C: Emerging Technologies*, 15(6):396–413, 2007. doi: 10.1016/j.trc.2006.10.003.
- [214] Xiqun Chen, Li Li, and Yi Zhang. A markov model for headway/spacing distribution of road traffic. *IEEE Transactions on Intelligent Transportation Systems*, 11(4):773–785, 2010.
- [215] Emil B Iversen, Juan M Morales, and Henrik Madsen. Optimal charging of an electric vehicle using a markov decision process. *Applied Energy*, 123:1–12, 2014.
- [216] Serge Hoogendoorn, Raymond G Hoogendoorn, and Winnie Daamen. Wiedemann revisited: New trajectory filtering technique and its implications for car-following modeling. *Transportation research record*, 2260(1):152–162, 2011.
- [217] Xue Yao, Ziyi Qin, Simeon C Calvert, and Serge P Hoogendoorn. Human driving patterns: A knowledge-enhanced deep learning approach for behaviour modelling. *IEEE Transactions on Intelligent Transportation Systems*, page Under review, 2025.
- [218] Thomas S Ferguson. A bayesian analysis of some nonparametric problems. *The annals of statistics*, pages 209–230, 1973.
- [219] Shiliang Sun and Xin Xu. Variational inference for infinite mixtures of gaussian processes with applications to traffic flow prediction. *IEEE Transactions on Intelligent Transportation Systems*, 12(2):466–475, 2010.

- [220] Rui Jiang, Cheng-Jie Jin, HM Zhang, Yong-Xian Huang, Jun-Fang Tian, Wei Wang, Mao-Bin Hu, Hao Wang, and Bin Jia. Experimental and empirical investigations of traffic flow instability. *Transportation research part C: emerging technologies*, 94: 83–98, 2018.

Acknowledgments

Before starting the final year of my master's studies, I was considering pursuing a PhD for the next step, but kept doubting myself: Can I? Thinking that sending an email costs nothing, why not try? Four years later, I look back with deep gratitude on that try, and got an answer of YES. Throughout this journey, I have received so much care, encouragement, and support from many people around me. Without all of you, this journey wouldn't be such beautiful and enjoyable.

First and foremost, I would like to express my heartfelt gratitude to my amazing supervision team, Simeon and Serge. Thank you for your patience, insight, trust, and constant guidance throughout this journey. You helped an insecure PhD candidate grow into an independent researcher. I deeply appreciate your individualised supervision style, which gave me the freedom to explore my own research interests, while always providing constructive support when needed. Your influence has shaped not only my academic growth but also my personal development.

To Simeon, my daily supervisor: thank you for your unwavering encouragement and for guiding me through every step of refining my ideas and writing. I was fully aware of how terrible my academic writing was at the beginning, yet you patiently helped me improve step by step. I remember most vividly that we iterated my first two papers more than 20 times each, which was a hard learning process but helped me grow tremendously. You were always there to encourage me when I felt stuck, reminding me that every PhD follows its own growth curve, and that progress may not come daily, weekly, or even monthly, but it always comes. This perspective has been invaluable to my confidence and self-identity. Additionally, your talent and efficiency at work, as well as your passion for life and hobbies, have deeply inspired me and shown me the true vitality of exploration and adventure.

To Serge, my promoter: thank you for your broad vision, encouragement, and continuous support throughout this journey. Your passion for research, deep insights into the field, and forward-looking perspective have made you a true leader and role model for me as a young researcher, guiding me into the wider world of transportation research. As a master's student who read and cited many of your research in my early studies, I cannot express how happy I was when I knew that you would be my promoter during my PhD. Despite your many responsibilities, you always try to make time for your PhD candidates; I can genuinely feel your heartfelt commitment to helping young researchers grow. I am always inspired by your many thought-provoking research ideas and deep passion for work. Especially, I cherish your encouragement during my most difficult moments, when motivation was hard to find. Your guidance, care, and support for my personal and career development have meant more to me than words can express.

My heartfelt thanks also go to Bart for forwarding my very first email to Simeon in September 2021, which made everything possible. What may have been a small act of

kindness for you as an educator had a significant impact on me and ultimately changed my life in the most positive way. I also thank you for the random chats in the corridor, for always being willing to offer advice, and especially for your support of my research visit toward the end of my PhD. I sincerely thank Alexandre for accepting me as a visiting scholar at the Center for Information Technology Research in the Interest of Society and the Banatao Institute (CITRIS) at UC Berkeley. Although this journey has not yet fully started due to unforeseen circumstances, I very much look forward to the inspiring collaboration with you and your wonderful team and to conducting cool research together next year. I am also grateful to my committee members, Hans, David, Dan, Soyung, and Monica, for their time, careful review, and valuable feedback. My deep appreciation goes to Meng for inviting me to join the Executive Team of the IEEE ITSS Benelux Chapter, from which I learned immensely. I also thank Ziran for our enjoyable chats and for your kind support within IEEE. I would like to thank Zhanbo, my master's supervisor, for supporting my early graduation process. Without your support, starting my PhD at TU Delft would not have been so smooth. My deep gratitude also goes to Xiaoyuan, my mentor during my bachelor's studies, whose guidance profoundly shaped my early vision of research and of the world.

I warmly thank the former department head, Serge, the current department head, Oded, and their management team, including Haneen, Gonçalo, Sascha, Winnie, and many others, for creating such a welcoming and supportive working environment. I am especially grateful to Maaike for her constant support and for always doing her best to help us with financial matters, even in difficult situations. My sincere thanks go to Monique, Priscilla, Suzanne, Marije, Dehlaila, and Moreen for being consistently approachable and for all the work you have done for the department. I also thank Peter and Edwin for always helping with equipment matters and for solving problems promptly and efficiently. Thank you to Conchita and Esther for your excellent organization of every TRAIL event, and thank you to everyone contributing to TRAIL for offering insightful courses and valuable networking opportunities for all PhD candidates. I am grateful to many wonderful faculty members in the T&P department. Special thanks go to Yanan for our inspiring discussions and for sharing new ideas and perspectives. Thank you to Irene & Jordi, and Matthew for our chats and especially our special and unforgettable hiking experiences. Thank you to Yufei for organising the summer events that brought us together, to Niels for your constant energy and for bringing me the enjoyable transportation board games, and to Victor, Andreas, and Marco for sharing research ideas and stimulating discussions.

My sincere thanks go to my talented and kind colleagues in the "Dittlab" office. To my "deskmate" Yiru, I feel incredibly lucky to have shared this journey with such a thoughtful, motivated, and talented researcher. The breadth and depth of your thinking, together with your steadfastness and perseverance, make me sincerely wish you succeed in everything you pursue. Together with Guopeng, your dedication to research and your pursuit of excellence have deeply inspired me. Thank you, Saeed, for your care, encouragement, and our many pleasant small talks. Starting our PhDs on the same day made this journey even more special, and I am grateful to have such a "hexagonal warrior" by my side. Thank you, Kexin, Yiyun, and Xiaolin, for our gatherings, heartfelt conversations, and constant support, these moments are among my warmest memories of this PhD journey. Thank you, Zili, for always being supportive; I especially cherish your encouragement when I felt stuck,

and our “special forces” trip was truly amazing and unforgettable. Thank you, Lucas, for showing me how one can pursue many important things at once and do them exceptionally well. I wish you and your family lasting happiness. I also thank Samir, Srinath, Marriko, Tian, Robin, Jonah, Macin, and Wouter for all the joy, inspiration, and happy moments we shared in the office.

I am also grateful to my colleagues from the DAIMoND Lab. Thank you to Mingze, Ziteng & Mengdan, Yuxing, Yanyan, and Xiamei for our shared discussions, exchanges of ideas, and moments of mutual support. I also thank Masha and Saman for organising lab meetings and events that brought everyone together and strengthened our sense of community. My thanks extend to other lab members, Nivarna, Elif, Ting, Vincent, Weiming, Tao, Sherman, Hari, and Alex, for our small but warm conversations, which always made the lab a pleasant place to be. I am also thankful to Sara, Qiaochu, Jiayi, Ziyulong, Zipei, Zixuan, Rina, Gavin, Enshan, Jinyang, Renate, Nina, Jyotsna, Lucas, Mohammad, Ximeng, Arco, Abhinav, Zamzam, Zheng, as well as to all other colleagues in the T&P Department, for the many ways they offered support, encouragement, and inspiration. To my former colleagues, Soyeon, Nagarjun, Konstantinos, Yongqi, Ajan, Yixin, Shiqi, Weining, Zhenjie, Bing, Junchi, Chao, Fei, Yumin, Heqi, Yuxia, Yitao, Dingshan, Zhuotong, Shuang, Chaopeng, Ali, Zahra, Tin, Pancharmy, Alex, and Merve, thank you for the wonderful memories we shared. Time may fade the details, but the warmth of those memories remains.

I deeply thank the many people I met in Delft who have become an important part of my life. My special thanks go to Skirmantas for his constant support, care, and companionship, as well as for the values and passions we share. The countless moments we have spent exploring life together have shaped some of my warmest memories, with many more still to come. My heartfelt thanks go to Shaheen, a very special girl in my life. From our first “accidental” online meeting across two different countries to supporting each other through both happiness and hardship, your care, friendship, and love have been a constant source of warmth and strength. Thank you, Qi, for introducing me to the world of snowboarding and for the many deep conversations, support, and delicious meals we shared. Thank you, Xiliang, for your care, our wonderful trips together, and your ever-fascinating ideas. My thanks also go to my neighbour Qiuju for our heartfelt talks and game nights, where time always seemed to disappear. I also thank the “happy brothers” Lingyu and Yun for our gatherings and mutual support; thank Shengren & Ping, Ran & Mingkun, Sifeng, and Yang for the joyful weekends we shared and the memories we created. My warm thanks go to my lovely friends Bjon & Alena, Sangamesh, Xlander, and Gavin & Rhea for our gatherings, board game nights, and ski trips. My thanks extend to Tianlong & Jing, Zhaochong, Pan, Binbin, Mingkai, Jinbao, Xinhan, Xiaoyu, Hao, Aihui, Haiwei, Ao, Shen, Yiping, Tianqi, Yitao, Xiuli, and Xin for the many enjoyable gatherings and conversations. To my friends from DaiBai snowboarding club, Qixiu, Emma, Sherry & Biebie, Terrance, Wanling, Wei, C’an, Ya, Shiqian, Tianli, Xiaomian, Patrick, Hu, and many others, thank you for the adventure, joy, and shared passion. I also deeply thank my Dutch coach, Frans, and his wife, Helly, for helping me practice Dutch, experience Dutch culture, and for our broad discussions and exchanges of ideas.

I gratefully acknowledge my co-authors and overseas collaborators for their support during this journey. I sincerely thank Chang, Zirui, Zhaocheng, Ziye, Ang, Chuheng, and

Qiruo for their invaluable contributions, insights, and teamwork. I am also thankful to many other talented young researchers I met during my PhD, Mingfeng, Chengyuan, Jing, Yajie, Jingwei, Heye, Zilin, Junyan, for their inspiration and friendship. To my old friends, Liping, Mingqian & Shuxin, Jianguang, Hui, Juncai, Chenghao, Xiao, Ying, Zhihang, Xin & Runzhe, Wenying, Zheng, and Yue, and to my cousin Sen and his family, thank you for your warmth, care, and hospitality whenever I returned to my home country.

Finally, my deepest gratitude goes to my family. I sincerely thank my sister, Mrs. Yao, and my brother-in-law, Mr. Chang, for taking such good care of our parents and for your constant support. To my niece and nephew, Ying and Shuo, you are truly the most precious gift. Above all, I thank my parents, Mr. Yao and Mrs. Xue. You may call yourselves ordinary, but to me you are extraordinary. Your constant support and love shaped who I am today, and you taught me the most important values in life: to be kind, independent, resilient, and so much more.

Please forgive me for not being able to mention every single person who has helped me along the way. Thank you all.

Last but not least, I thank myself for all the passion, persistence, and hard work. Through this beautiful journey, I have become the person I aspired to be, and I now step into the next chapter of my life with confidence, strength, and hope.

*Xue Yao
Delft, August 2025*

About the Author

Xue Yao was born in Weifang, Shandong Province, China, a beautiful city known as “the Capital of Kites”. She obtained her Bachelor’s degree in Traffic Engineering with distinction (Outstanding Thesis) from Shandong University of Technology in 2019. During her undergraduate studies, she visited Japan as a representative of China–Japan friendship in 2017, where she planted the Edo Higan Cherry Tree (*Prunus itosakura*) to commemorate bilateral ties. In 2018, she co-founded a start-up in Zibo, China, focusing on drone technology and services. In 2019, she participated in an exchange program at Sungkyunkwan University in Seoul, South Korea.



In 2019, Xue obtained an exam-free admission to Southwest Jiaotong University in Chengdu, Sichuan Province, where she pursued a master’s degree in Traffic and Transportation Engineering. Her research focused on human driving behaviour modelling and cooperative decision-making for autonomous vehicles. She obtained her Master’s degree in December 2021, half a year ahead of schedule, with a thesis titled “*Cooperative Decision-making and Control in Mixed Traffic by Considering Driving Heterogeneity*”.

In January 2022, Xue joined the Department of Transport and Planning at Delft University of Technology (TU Delft) as a PhD candidate, funded by the TU Delft scholarship. Her research focuses on exploring the application of artificial intelligence (AI) in transport and mobility, leveraging machine learning and data-driven approaches to improve traffic modelling and simulation. In August 2025, she secured a visiting scholar position at the Center for Information Technology Research in the Interest of Society and the Banatao Institute (CITRIS), UC Berkeley, supervised by Prof. Alexandre M. Bayen. In October 2025, Xue began her postdoctoral research at the same department in TU Delft.

Throughout her PhD, Xue actively engaged in professional activities. She has served as Vice Chair of the IEEE ITSS Benelux Chapter since January 2025. In this role, she secured IEEE funds and organised multiple events for students and young professionals, aiming to bridge academia and industry and support professional development.

List of Publications

Journal papers

7. **Yao, X.***, Calvert, S. C., and Hoogendoorn, S. P. (2025). A Pattern-based Framework for Modelling Driving Heterogeneity and Traffic Flow Simulation, *Under review by a journal*.
6. **Yao, X.***, Qin, Z., Calvert, S. C. and Hoogendoorn, S. P. (2025). Human Driving Patterns: A Knowledge-Enhanced Deep Learning Approach for Behaviour Modelling. *IEEE Transactions on Intelligent Transportation Systems, R1 Revision*.
5. **Yao, X.***, Calvert, S. C. and Hoogendoorn, S. P. (2025). A Novel Framework for Understanding and Identifying Driving Heterogeneity Through Action Patterns. *IEEE Transactions on Intelligent Transportation Systems*, doi=10.1109/TITS.2025.3560509.
4. **Yao, X.***, Calvert, S. C. and Hoogendoorn, S. P. (2025). Driving Heterogeneity Identification Using Machine Learning: A Review and Framework for Analysis. *Transportation Research Interdisciplinary Perspectives*, 32, 101511.
3. **Yao, X.***, Yan, Q., Sun, Z., Calvert, S. C., and Hoogendoorn, S. P. (2024). Investigation on Car-Following Heterogeneity and Its Impacts on Traffic Safety and Sustainability. *Transportmetrica A: Transport Science*, 1-25.
2. **Yao, X.**, Du, Z., Sun, Z*, Calvert, S. C., and Ji, A. (2024). Cooperative Lane-Changing in Mixed Traffic: A Deep Reinforcement Learning Approach. *Transportmetrica A: Transport Science*, 1-23.
1. Sun, Z*, **Yao, X.**, Qin, Z., Zhang, P., and Yang, Z. (2021). Modelling Car-Following Heterogeneities by Considering Leader-Follower Compositions and Driving Style Differences. *Transportation Research Record*, 2675 (11), 851-864.

Special papers

1. **X. Yao** and Z. Wang*, [Its PH.D.s], in *IEEE Intelligent Transportation Systems Magazine*, vol. 17, no. 6, pp. 124-126, Nov.-Dec. 2025, doi: 10.1109/MITS.2025.3606735.

Conference contributions

12. **Yao, X.***, Calvert, S. C., and Hoogendoorn, S. P. (2026). Unveiling Heterogeneity in Driving Behaviour: A Pattern-based Approach for Traffic Simulation and Analysis. Oral presentation at *Proceedings of the 105th Annual Meeting of Transportation Research Board (TRB2026)*, Washington, D.C. USA.
11. **Yao, X.***, Qin, Z., Calvert, S. C., and Hoogendoorn, S. P. (2025). Action Pattern Prediction with Knowledge-enhanced Attention LSTM. *Proceedings of the 104th Annual Meeting of Transportation Research Board (TRB2025)*, Washington, D.C. USA.

10. **Yao, X.***, Calvert, S. C., and S. P. Hoogendoorn S. P. (2024). A Novel Framework for Understanding and Identifying Driving Heterogeneity Through Action Patterns. *TRAIL PhD Congress 2024*, Utrecht, the Netherlands.
9. **Yao, X.***, Calvert, S.C., Hoogendoorn, S.P. (2024). Action Pattern Recognition based on Action Phases Clustering. *12th Symposium of the European Association for Research in Transportation (hEART2024)*, Helsinki, Finland.
8. **Yao, X.***, Calvert, S.C., Hoogendoorn, S.P. (2024). Driving Pattern Interpretation based on Action Phases. *Proceedings of the 103th Annual Meeting of Transportation Research Board (TRB2024)*, Washington, D.C. USA.
7. **Yao, X.***, Calvert, S. C., and S. P. Hoogendoorn S. P. (2023). Identification of Driving Heterogeneity with Machine Learning: A Literature Review and Future Perspectives. *TRAIL PhD Congress 2023*, Utrecht, the Netherlands.
6. **Yao, X.***, Calvert, S. C., and S. P. Hoogendoorn S. P. (2023). Identification of Driving Heterogeneity using Action-chains. *IEEE 26th International Conference on Intelligent Transportation Systems (ITSC2023)*, Bilbao, Spain.
5. **Yao, X.***, Hou, S., Hoogendoorn, S. P., and Calvert, S. C. (2023). Performance Comparison of Deep RL Algorithms for Mixed Traffic Cooperative Lane-Changing. *Proceedings of the 102th Transportation Research Board Annual Meeting (TRB2023)*, Washington D.C., USA.
4. **Yao, X.**, Sun, Z.*, Yan, Q. (2023). Performance Comparison of Multi-class SVM and S3VM for Driving Style Classification. *Proceedings of the 102th Transportation Research Board Annual Meeting (TRB2023)*, Washington D.C., USA.
3. **Yao, X.***, Calvert, S. C., and S. P. Hoogendoorn S. P. (2022). Traffic Heterogeneity with Connectivity and Connected Automated Vehicles. *TRAIL PhD Congress 2022*, Utrecht, the Netherlands.
2. **Yao, X.**, Du, Z., Sun, Z.*, (2022). Cooperative Lane-changing in Mixed Traffic: A Deep Reinforcement Learning Approach. *Proceedings of the 101th Transportation Research Board Annual Meeting (TRB2022)*, Washington D.C., USA.
1. Sun, Z*, **Yao, X.** (2021). Modelling Car-Following Heterogeneities by Considering Leader-follower Compositions and Driving Style Differences. *Proceedings of the 100th Transportation Research Board Annual Meeting (TRB2021)*, Washington D.C., USA.

TRAIL Thesis Series

The following list contains the most recent dissertations in the TRAIL Thesis Series. For a complete overview of more than 400 titles, see the TRAIL website: www.rsTRAIL.nl. The TRAIL Thesis Series is a series of the Netherlands TRAIL Research School on transport, infrastructure and logistics.

- Yao, X., *Driving Heterogeneity in Traffic Flow Theory: An action-based framework for identification, modelling, and simulation*, January 2026, TRAIL Thesis Series, the Netherlands
- Versluis, N.D., *Optimising Railway Traffic Management under Radio-Based Distance-to-Go Signalling*, January 2026, TRAIL Thesis Series, the Netherlands
- Jiao, Y., *Proactive Collision Risk Quantification in Multi-directional Traffic Interactions*, January 2026, TRAIL Thesis Series, the Netherlands
- Asadi, M., *Accessibility and Road Safety: Integration of road safety in accessibility evaluation*, November 2025, TRAIL Thesis Series, the Netherlands
- Akse, R., *Understanding and untangling the uncertainty knot: How to catalyse decision-making in mobility innovations*, November 2025, TRAIL Thesis Series, the Netherlands
- Führer, K., *Participatory Decision-making under Deep Uncertainty: Modelling mobility transitions*, November 2025, TRAIL Thesis Series, the Netherlands
- Picco, A., *Monitoring and Feedback in Driving*, T2025/17, October 2025, TRAIL Thesis Series, the Netherlands
- Cebeci, M.S., *Behaviour of Prosumers in Last-mile Logistics: The case of crowdshipping*, T2025/16, September 2025, TRAIL Thesis Series, the Netherlands
- Kuijpers, A., *Enabling Inter-Organizational Collaboration Through Platforms: The role of trust*, T2025/15, September 2025, TRAIL Thesis Series, the Netherlands
- Song, R., *Human-MASS Interaction in Decision-Making for Safety and Efficiency in Mixed Waterborne Transport Systems*, T2025/14, June 2025, TRAIL Thesis Series, the Netherlands
- Destyanto, A.R., *A Method for Evaluating Port Resilience in an Archipelago*, T2025/13, June 2025, TRAIL Thesis Series, the Netherlands

- Karademir, C., *Synchronized Two-echelon Routing Problems: Exact and approximate methods for multimodal city logistics*, T2025/12, May 2025, TRAIL Thesis Series, the Netherlands
- Vial, A., *Eyes in Motion: A new traffic sensing paradigm for pedestrians and cyclists*, T2025/11, May 2025, TRAIL Thesis Series, the Netherlands
- Chen, Q., *Towards Mechanical Intelligence in Soft Robotics: Model-based design of mechanically intelligent structures*, T2025/10, April 2025, TRAIL Thesis Series, the Netherlands
- Eftekhar, Z., *Exploring the Spatial and Temporal Patterns in Travel Demand: A data-driven approach*, T2025/9, June 2025, TRAIL Thesis Series, the Netherlands
- Reddy, N., *Human Driving Behavior when Interacting with Automated Vehicles and the Implications on Traffic Efficiency*, T2025/8, May 2025, TRAIL Thesis Series, the Netherlands
- Durand, A., *Lost in Digitalisation? Navigating public transport in the digital era*, T2025/7, May 2025, TRAIL Thesis Series, the Netherlands
- Dong, Y., *Safe, Efficient, and Socially Compliant Automated Driving in Mixed Traffic: Sensing, Anomaly Detection, Planning and Control*, T2025/6, May 2025, TRAIL Thesis Series, the Netherlands
- Droffelaar, I.S. van, *Simulation-optimization for Fugitive Interception*, T2025/5, May 2025, TRAIL Thesis Series, the Netherlands
- Fan, Q., *Fleet Management Optimisation for Ride-hailing Services: from mixed traffic to fully automated environments*, T2025/4, April 2025, TRAIL Thesis Series, the Netherlands
- Hagen, L. van der, *Machine Learning for Time Slot Management in Grocery Delivery*, T2025/3, March 2025, TRAIL Thesis Series, the Netherlands
- Schilt, I.M. van, *Reconstructing Illicit Supply Chains with Sparse Data: a simulation approach*, T2025/2, January 2025, TRAIL Thesis Series, the Netherlands
- Ruijter, A.J.F. de, *Two-Sided Dynamics in Ridesourcing Markets*, T2025/1, January 2025, TRAIL Thesis Series, the Netherlands
- Fang, P., *Development of an Effective Modelling Method for the Local Mechanical Analysis of Submarine Power Cables*, T2024/17, December 2024, TRAIL Thesis Series, the Netherlands
- Zattoni Scroccaro, P., *Inverse Optimization Theory and Applications to Routing Problems*, T2024/16, October 2024, TRAIL Thesis Series, the Netherlands

- Kapousizis, G., *Smart Connected Bicycles: User acceptance and experience, willingness to pay and road safety implications*, T2024/15, November 2024, TRAIL Thesis Series, the Netherlands
- Lyu, X., *Collaboration for Resilient and Decarbonized Maritime and Port Operations*, T2024/14, November 2024, TRAIL Thesis Series, the Netherlands
- Nicolet, A., *Choice-Driven Methods for Decision-Making in Intermodal Transport: Behavioral heterogeneity and supply-demand interactions*, T2024/13, November 2024, TRAIL Thesis Series, the Netherlands



Summary

Human driving behaviour is inherently heterogeneous, shaping traffic dynamics and affecting traffic safety, efficiency and sustainability. This dissertation develops an interpretable, AI-driven framework to identify, model, and simulate heterogeneous driving behaviour using naturalistic data. By analysing action phases, patterns, and behavioural sequences, it reveals how behavioural variability influences traffic performance and supports improved traffic management, personalised driver assistance, and human-aware autonomous vehicle design.

About the Author

Xue Yao conducted her PhD research in the Department of Transport & Planning at Delft University of Technology. Her research focuses on exploring the application of artificial intelligence (AI) for transport and mobility to improve safety, efficiency and sustainability.

TRAIL Research School ISBN 978-90-5584-378-7



Radboud University



rijksuniversiteit
 groningen



UNIVERSITY OF TWENTE. TU/e

Technische Universiteit
 Eindhoven
 University of Technology