

Communication-Enabled Interactions in Highway Traffic
A joint driver model for merging

Siebinga, O.

DOI

[10.4233/uuid:ba814d00-4060-4544-9690-8eace5ddcffa](https://doi.org/10.4233/uuid:ba814d00-4060-4544-9690-8eace5ddcffa)

Publication date

2024

Document Version

Final published version

Citation (APA)

Siebinga, O. (2024). *Communication-Enabled Interactions in Highway Traffic: A joint driver model for merging*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:ba814d00-4060-4544-9690-8eace5ddcffa>

Important note

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

Copyright

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.

Communication-Enabled Interactions in Highway Traffic

A joint driver model for merging

Communication-Enabled Interactions in Highway Traffic

A joint driver model for merging

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op vrijdag 17 mei 2024 om 12:30 uur.

door

Olger SIEBINGA

Werktuigbouwkundig Ingenieur,
Technische Universiteit Delft, Nederland,
geboren te Heerenveen, Nederland.

Dit proefschrift is goedgekeurd door de promotor en copromotor.

Samenstelling promotiecommissie:

Rector Magnificus,
Prof. dr. ir. D.A. Abbink
Dr. A. Zgonnikov

voorzitter
Technische Universiteit Delft, promotor
Technische Universiteit Delft, copromotor

Onafhankelijke leden:

Prof. dr. M.P. Hagenzieker
Prof. dr. ir. S.P. Hoogendoorn
Prof. dr. G. Markkula
Prof. dr. N. Ruswinkel
Dr. E. Vinkhuyzen

Technische Universiteit Delft
Technische Universiteit Delft
University of Leeds, UK
University of Lubeck, Germany
King's College London, UK



Dit onderzoek is gefinancierd door Nissan Motor Company, Ltd. en een toeslag van de Rijksdienst voor Ondernemend Nederland (TKI2012P01)

Druk: Ridderprint
Layout: Olger Siebinga, based on the Latex template:
"TUD-Dissertation" (2020, K.P. Hart)
Vormgeving: Rixt Siebinga | rixt-art.com.

Copyright © 2024 by O. Siebinga

ISBN 978-94-6384-577-9

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

Contents

The background is a vibrant orange with a subtle, grainy texture. In the bottom-left corner, there is a white diagonal line that extends towards the center. Scattered in the lower-left quadrant are several blue dots of varying sizes, some appearing as small specks and others as larger, more prominent circles.

Summary	ix
Samenvatting	xv
1 Introduction	1
2 A human factors approach to validating driver models for interaction-aware automated vehicles	18
3 Uncovering variability in human driving behavior through automatic extraction of similar traffic scenes from large naturalistic datasets	43
4 Modelling communication-enabled traffic interactions	57
5 Interactive merging behaviour in a coupled driving simulator: Experimental framework and case study	86
6 Human merging behaviour in a coupled driving simulator: How do we resolve conflicts?	97
7 A model of dyadic merging interactions explains human drivers' behaviour from input signals to decisions	119
8 Conclusion	145
Appendices	165
A TraViA: a Traffic data Visualization and Annotation tool in Python	167
Acknowledgements	175
Curriculum Vitæ	178
List of Publications	179

Summary

Automated driving technologies hold many potential benefits for society, such as safer, more accessible, and more environmentally friendly road transportation. However, many challenges remain before Automated Vehicles (AVs) can be widely used on public roads. One of these important open problems is the interaction between AVs and human-driven vehicles in traffic. Interactions between vehicles during lane changing and merging on highways are the specific focus of this thesis.

In these merging interactions, two or more vehicles negotiate a safe outcome while adhering to spatio-temporal constraints such as relative speeds and distances. All drivers (human or automated) adapt their control behaviour continuously to prevent collisions. Besides that, human drivers use their vehicle's kinematics to communicate their intent. During a merging interaction, drivers make joint decisions, such as who goes first and exhibit joint behaviour, such as the size of the gap they keep between the vehicles. However, these joint outcomes of an interaction stem from individual decisions and behaviours, such as whether to yield and how and when to accelerate. Therefore, the interaction between multiple vehicles on the highway can best be described as a joint driver effort and best modelled with a **joint driver model**: a model that describes multiple drivers' joint and individual behaviours.

Such a joint driver model should cover the joint system of drivers and their individual contributions on multiple levels. High-level decisions are needed to capture the most likely outcome and individual contributions to that outcome (e.g., who yields or accelerates). Gap-keeping and velocity behaviour are used by human drivers for communicative actions and need to be covered by a model to understand those. However, such a joint driver model capturing the dynamics of traffic interactions during merging is missing. Current models used in autonomous vehicles do not always generalise well to real-world behaviour (Chapter 2). Other driver behaviour models only regard a single driver responding to their environment and often have a limited scope. For example, they only cover the decision to yield or merge or the acceleration input while following another car, but not both. Combining multiple models of individual drivers to describe the joint behaviour will not necessarily cover the dynamics between the drivers since these models are only designed to respond unilaterally.

Current interactions between AVs and human drivers are often over-conservative and sometimes awkward. This can potentially decrease the acceptance of the AV's behaviour and impact safety and travel efficiency. A potential solution to this problem is to provide AVs with knowledge of human driving behaviour. This knowledge, in the form of a driver model, could inform the AV on two topics: first, how the other drivers in the interaction are most likely to act, and second, how human-like the behaviour of the AV itself is. This can help the AV to display safe and acceptable behaviour in interactions. However, current AVs use models of individual drivers that are assumed to only respond to the AVs's behaviour. To facilitate a step towards making interaction-aware autonomous vehicles a reality **this thesis aims to increase the fundamental understanding of merging and lane-changing interactions and capture this knowledge in a joint driver model**. This thesis includes work in three pillars: naturalistic driving behaviour, model theory, and controlled experiments.

Naturalistic Driving Behaviour The work on naturalistic driving behaviour considers existing approaches to interaction-aware control: Controllers for AVs that incorporate a model of human driving behaviour to reason about the future actions of

other drivers. Although such controllers were previously demonstrated in simulated environments, showing that they generate safe behaviour, their interactive capabilities were never evaluated, and the used driver models were never validated. Chapter 2 presents a method to validate driver models used in interaction-aware controllers. In a case study, an inverse-reinforcement-learning-based driver model generalises poorly to real-world driving behaviour. This chapter uses a naturalistic (real-world) dataset recorded on a German highway: the HighD dataset and a specifically developed visualisation software package named TraViA (a **Traffic Visualisation and Annotation tool**) (Appendix A).

Chapter 3 presents a method to extract similar scenes from this large dataset to investigate the variability in human responses to similar scenarios. This investigation showed that interactive scenarios have multiple possible outcomes on both an *operational* level (how a manoeuvre is performed) and *tactical* level (which manoeuvre is chosen). For example, when approaching a slower-moving vehicle on the highway, a driver can decide to slow down or change lanes (tactical variability). Still, if they slow down, the amount of braking can vary (operational variability). Chapter 3 shows that driver responses to similar scenarios vary on both levels and it provides a method to uncover and quantify this variability.

Model Theory Chapter 4 focuses on the theory behind driver models for interactions and discusses the limitations of the current approaches. One limitation is that many existing driver models only describe a single driver. They assume this driver responds to their environment, but the environment (including all other drivers) does not respond to the modelled driver. We named this the one-way interaction assumption because there is only one responsive driver. This assumption prevents a model from describing dynamic interactions. A much-used approach to joint driver modelling uses game theory, developed initially to model one-shot decisions. This does not align with the dynamic negotiations observed in traffic interactions. Furthermore, game theory assumes humans behave as rational utility maximisers who do not communicate. Chapter 4 presents evidence from the literature that humans do not continuously (i.e., on every timestep in a simulation) optimise their behaviour and that communication is an important aspect of traffic interactions.

This theoretical evaluation led to the development of a new model framework to describe the joint driver dynamics in traffic interactions, presented in Chapter 4. The framework explicitly incorporates the communication between multiple drivers; therefore, we named it the Communication-Enabled Interaction (CEI) model framework. The CEI framework assumes drivers have a deterministic plan for their future movements. They communicate this plan to others through implicit or explicit communication; at the same time, they receive similar communication from these other drivers. Based on the received communication, drivers form a probabilistic belief of the other driver's future movements. The probabilistic belief combined with the deterministic plan results in a level of perceived risk. If the perceived risk exceeds a personal threshold, the modelled driver updates their plan to get the risk under control. These triggered plan updates are based on the concept of satisficing: the notion that humans do not have the mental capacity or the time to optimise their behaviour constantly but instead search for a solution that suffices and satisfies. In a case study, an instance of a CEI-model implemented to describe a merging scenario showed plausible joint driver behaviours based on explainable model parameters. Furthermore, human-like gap-keeping behaviour emerged in a car-following scenario.

Controlled Experiments Controlled experiments have been widely used to investigate the effects of controlled parameters on the behaviour of individual drivers. However, there are not many controlled experiments involving multiple drivers. One of the complicating factors of including multiple drivers in a single experiment is the large number of signals that must be measured and processed per driver (e.g., steering inputs and pedal inputs). Furthermore, analytic tools and valid metrics to analyse the data are missing because previous experiments used metrics and tools specifically developed for individual drivers. One example of such a metric is time-to-collision, which is used to investigate car-following behaviour. Chapter 5 takes a step towards controlled experiments with two drivers in a simulator. It presents a simplified merging scenario, with fewer signals for each driver, that can be used in a coupled driving simulator. In this scenario, two human drivers resolve a merging conflict in a top-down view simulation. Furthermore, three analytic tools are specially developed to analyse the drivers' behaviour, including a metric to capture the duration of the conflict: the Conflict Resolution Time (CRT).

Chapter 6 presents the empirical analysis of an experiment with the simplified merging scenario. It reveals important aspects of driver behaviour in interactions. Firstly, humans use intermittent piecewise-constant control for the acceleration of their vehicles. Secondly, the experiment showed that the high-level outcome of a merging interaction (i.e., which vehicle merges first) mostly depends on the kinematics of the vehicles at the start of the interaction, not on differences between the drivers. These effects are quantified in a statistical model. Thereby, Chapter 5 helps to increase the fundamental understanding of interactive driver behaviour during merges. Finally, the effects of vehicle kinematics on CRT are also quantified in Chapter 6.

The empirical analysis provided insights that inspired improvements to the model presented in the case study in Chapter 4. This updated joint driver model is presented in Chapter 7. It uses intermittent piece-wise constant control analogues to the observed behaviour in Chapter 6. A new implementation of the belief module reflects the safety margins observed in human merging interactions. Finally, dynamic risk-thresholds provide an incentive to act based on relative kinematics (e.g., a following vehicle in car following is more inclined to act than the leading vehicle). The model in Chapter 7 can qualitatively and quantitatively describe the driver behaviour from the experiment in Chapter 6 on multiple levels. High-level decisions of individual drivers (e.g., to yield or not) lead to accurate joint outcomes (e.g., which driver merges first?). Human-like input signals (i.e., velocity profiles) accurately describe the individual contributions to the joint safety margins (i.e., the gap between the vehicles). Finally, the model reproduces typical qualitative interactions between drivers in the experiment, such as miscommunications when both drivers initially take the same action.

The three overarching conclusions of this thesis are: 1) Different drivers respond with different tactical and operational behaviours to similar interactive situations; therefore, driver models should capture operational and tactical variability, which should be assessed independently. 2) An important requirement for a joint driver model is the ability to capture that drivers do not continuously (rationally) optimize their acceleration inputs; instead, they use intermittent piece-wise constant control – as was empirically observed in a simplified merging scenario. 3) With communication-enabled, risk-based intermittent control, the proposed CEI-model can describe abstract merging interactions between two drivers, including their

decisions (who goes first), safety margins over time, and underlying individual contributions and control inputs (brake/accelerate). The most important limitation of this thesis is that all the modelling was done in a simplified merging scenario. This scenario involved just two drivers that only have longitudinal control. One of the major open challenges on the road to leveraging the results of this thesis is extending the model to scenarios with more vehicles and full control. The conclusions, limitations, possible future work and potential model applications are further discussed in Chapter 8.

To conclude, I believe the work presented in this thesis has yielded valuable knowledge about human lane-changing and merging interactions and how to model them. The studies presented here have not only identified the shortcomings of existing interactive autonomous driving approaches based on naturalistic traffic data but also proposed a solution to modelling traffic interactions that shows much potential in a simplified merging scenario. Therefore, I believe the work could be a considerable step towards equipping autonomous vehicles with knowledge of how drivers interact in traffic.

Samenvatting

Technologie voor geautomatiseerd rijden biedt veel potentiële voordelen voor de samenleving, zoals veiliger, toegankelijker en milieuvriendelijker vervoer. Maar een aantal uitdagingen moeten worden aangepakt voordat geautomatiseerde of Autonome Voertuigen (AV's) breed kunnen worden ingezet op openbare wegen. Eén van deze belangrijke open problemen is de interacties in het verkeer tussen AV's en door mensen bestuurde auto's. In dit proefschrift staan zulke interacties tussen voertuigen op snelwegen centraal, specifiek tijdens het wisselen van rijstrook en bij het invoegen.

Bij die interacties moeten de bestuurders van twee of meer voertuigen samen tot een veilige uitkomst komen. Hierbij moeten ze rekening houden met de beschikbare ruimte en tijd, zoals met relatieve snelheden en afstanden. Alle bestuurders (menselijk of geautomatiseerd) passen hun gedrag continu aan om botsingen te voorkomen. Bovendien gebruiken mensen de positie en snelheid van hun auto om aan anderen te communiceren wat ze van plan zijn. Bij het invoegen nemen bestuurders samen beslissingen, zoals wie voor gaat, en bepalen ze samen de veiligheidsmarge die ze aanhouden. Deze gezamenlijke aspecten van een interactie komen echter voort uit individuele beslissingen en gedrag, zoals de beslissing om wel of geen voorrang te verlenen en het gedrag hoe hard iemand remt. De interactie tussen meerdere voertuigen op de snelweg is daarom een gezamenlijke inspanning van individuele bestuurders en kan dus het best worden gemodelleerd met een **gezamenlijk model**: een model dat het gezamenlijke en individuele gedrag van meerdere bestuurders beschrijft.

Om waardevol te kunnen zijn voor een AV moet een model van menselijk rijgedrag de bestuurders gezamenlijk en individueel kunnen beschrijven op meerdere niveaus. Gezamenlijke beslissingen beschrijven de uitkomst (wie gaat er eerst) waaraan individuele beslissingen bijdragen, bijvoorbeeld de beslissing om af te remmen en ruimte te maken voor een andere invoegende bestuurder. Daarnaast gebruiken bestuurders hun snelheid en afstand ten opzichte van anderen om te communiceren, een model moet dit kunnen beschrijven om van waarde te zijn in een AV. Maar zo'n model wat de gezamenlijke dynamiek van bestuurders tijdens het invoegen kan beschrijven ontbreekt in de literatuur. Modellen die worden gebruikt in autonome voertuigen beschrijven gedrag op de weg niet altijd goed (Hoofdstuk 2). Andere modellen in de literatuur beschrijven vaak maar één enkele bestuurder die alleen reageert op zijn of haar omgeving. Daarnaast beschrijven deze modellen vaak het gedrag op één enkel niveau, ze beschrijven bijvoorbeeld alleen de beslissing om voorrang te verlenen of alleen de versnellingen tijdens het volgen van een andere auto. Een combinatie van meerdere van deze individuele modellen beschrijft niet automatisch het gezamenlijke gedrag omdat deze modellen ontworpen zijn om eenzijdige reacties te beschrijven.

Huidige AV's gedragen zich vaak conservatief in interactie met menselijke bestuurders, hierdoor kan de interactie onnatuurlijk zijn. Deze onnatuurlijke interacties kunnen de acceptatie van autonoom gedrag en de veiligheid en efficiëntie in het verkeer negatief beïnvloeden. Om interacties met mensen te verbeteren worden AV's in sommige gevallen voorzien van kennis van menselijk rijgedrag. Deze kennis, in de vorm van een model, helpt een AV op twee manieren. Ten eerste door te voorspellen wat de andere bestuurders waarschijnlijk zullen doen in de nabije toekomst. Daarnaast kan het model worden gebruikt om te beoordelen hoe menselijk het gedrag van de AV zelf is. Deze informatie kan helpen om de AV zich veilig en acceptabel te laten gedragen in interacties op de weg. Echter, de modellen die in AV's worden gebruikt zijn modellen van enkele bestuurders die

eenzijdig reageren op de AV. Om natuurlijke interacties tussen AV's en mensen dichterbij te brengen **heeft dit proefschrift als doel om de fundamentele kennis van interactie tussen bestuurders tijdens het invoegen en wisselen van rijbaan te vergroten en deze kennis te vangen in een model van meerdere bestuurders**. Het werk in dit proefschrift kan worden onderverdeeld in drie pijlers: natuurlijk rijgedrag, modellen, en gecontroleerde experimenten.

Natuurlijk Rijgedrag In het eerste deel over natuurlijk rijgedrag worden bestaande technologieën voor interactieve AV's beschreven en gevalideerd. Deze AV's gebruiken een model van menselijk rijgedrag om de toekomstige acties van andere bestuurders te voorspellen. In simulaties is aangetoond dat deze AV's veilig gedrag genereren, maar de gebruikte modellen van rijgedrag werden nooit gevalideerd voor gebruik buiten deze simulaties. In Hoofdstuk 2 wordt een methode gepresenteerd om met behulp van data die is opgenomen op de weg modellen van rijgedrag te valideren die worden gebruikt in interactieve AV's. Deze methode wordt vervolgens gebruikt om een model op basis van *inverse reinforcement learning* te valideren. Dit model wordt in de literatuur gebruikt in simulaties, maar blijkt niet goed te werken voor gedrag op de weg. In dit hoofdstuk wordt gebruik gemaakt van een dataset die is opgenomen op een Duitse snelweg: de HighD-dataset. Daarvoor is specifieke visualisatie-software ontwikkeld: TraViA (een **Traffic Visualisatie en Annotatietool**) (Bijlage A).

In Hoofdstuk 3 wordt een methode ontwikkeld om vergelijkbare scènes in deze grote dataset te vinden. Die scènes kunnen vervolgens worden gebruikt om de variabiliteit in menselijke reacties op die scène te onderzoeken. Deze reacties blijken meerdere mogelijke uitkomsten te hebben op zowel een *operationeel* niveau (hoe een manoeuvre wordt uitgevoerd) als op een *tactisch* niveau (welke manoeuvre wordt gekozen). Zo kan een bestuurder achter een langzamer rijdend voertuig op de snelweg beslissen om af te remmen of van baan te wisselen (tactische variabiliteit). Maar als de bestuurder besluit om af te remmen kan dit op meerdere manieren (operationele variabiliteit). Hoofdstuk 3 laat zien dat er variabiliteit bestaat op beide niveaus als bestuurders regeren op vergelijkbare scènes. In een casestudie wordt ook een methode gedemonstreerd om die variabiliteit te kwantificeren.

Modellen In Hoofdstuk 4 worden recente ontwikkelingen in modellen van menselijk rijgedrag voor interacties geanalyseerd. Veel van deze modellen beschrijven slechts één enkele bestuurder; ze gaan ervan uit dat deze bestuurder reageert op zijn of haar omgeving, maar dat de omgeving (inclusief alle andere bestuurders) niet reageert op de gemodelleerde bestuurder. We noemen dit de *aanname van eenzijdige interactie* omdat er slechts één bestuurder reageert op de ander. Door deze aanname kan een model de dynamiek van interacties tussen meerdere bestuurders niet beschrijven. Een veelgebruikt alternatief voor interactie modellen maakt gebruik van speltheorie (*game theory*). Deze theorie is oorspronkelijk ontwikkeld om beslissingen bij het spelen van een spel te modelleren en om de strategie te vinden die de kans op winst maximaliseert. Deze eenmalige beslissingen hebben weinig overeenkomsten met de continue dynamische aanpassingen in het gedrag van weggebruikers. Daarnaast maak speltheorie de aannames dat mensen zich rationeel gedragen en niet communiceren. Echter is uit de literatuur bekend dat mensen hun gedrag niet voortdurend optimaliseren en dat communicatie een belangrijk aspect is in interacties in het verkeer.

Op basis van deze evaluatie van recente literatuur wordt in Hoofdstuk 4 een nieuw ontwerp-kader ontwikkeld voor modellen die dynamische interacties tussen be-

stuurders kunnen beschrijven. Dit kader is gebaseerd op het feit dat communicatie tussen bestuurders een belangrijke rol speelt in het verkeer, vandaar de naam: het *Communication-Enabled Interaction* (CEI) model-kader. Het CEI-kader gaat ervan uit dat bestuurders een deterministisch plan hebben voor hun acties in de nabije toekomst. Ze communiceren dit plan naar anderen doormiddel van impliciete of expliciete communicatie; ook ontvangen ze communicatie van deze andere bestuurders. Op basis van de ontvangen communicatie vormen bestuurders een probabilistisch geloof over de wat de andere bestuurder in de nabije toekomst gaat doen. De combinatie van dit plan en het geloof vormen de basis voor een perceptie van risico. Als het waargenomen risico hoger is dan een persoonlijke norm werkt de bestuurder zijn of haar plan bij om het risico onder controle te krijgen. Deze plan-updates die getriggerd worden door risico perceptie zijn gebaseerd op het concept van 'satisficing': het idee dat mensen niet is staat zijn om continu hun gedrag te optimaliseren, maar in plaats daarvan op zoek gaan naar een oplossing die 'goed genoeg' is. In een casestudie met een invoegscenario laat een CEI-model plausibel gedrag zien van twee bestuurders. Dit op basis van modelparameters met een duidelijke functie. Bovendien hield het model in een scenario waar twee auto's elkaar volgen op een menselijke manier afstand tot anderen terwijl het daar niet specifiek voor ontwikkeld is.

Gecontroleerde Experimenten Met gecontroleerde experimenten worden de effecten van (gecontroleerde) parameters op het gedrag van individuele bestuurders onderzocht, dit wordt echter weinig gedaan voor interacties tussen meerdere bestuurders. Meerdere bestuurders laten deelnemen aan één experiment is complex door het grote aantal signalen dat moet worden gemeten en verwerkt per bestuurder (bijvoorbeeld de stuur- en gaspedaal-hoeken). Bovendien ontbreken analytische tools (zoals geschikte statistieken) om de gegevens te analyseren. Veel eerdere experimenten gebruiken statistieken en tools die specifiek zijn ontwikkeld voor individuele bestuurders; bijvoorbeeld de statistiek *time-to-collision*, die wordt gebruikt om het gedrag bij het volgen van auto's te onderzoeken. In Hoofdstuk 5 wordt een gecontroleerd experimenten met twee bestuurders in één simulator ontwikkeld. Onderdeel van dit experiment is een vereenvoudigd invoegscenario met minder signalen per bestuurder. Dit scenario kan worden gebruikt in een gekoppelde rijnsimulator met twee bestuurders en een simulatie in vogelperspectief. Daarnaast worden er in dit hoofdstuk drie analytische tools ontwikkeld om het gedrag van de bestuurders te analyseren, waaronder een statistiek om de duur van het conflict te meten: de *Conflict Resolution Time* (CRT).

In Hoofdstuk 6 wordt een empirische analyse gemaakt op basis van de resultaten van het experiment met het vereenvoudigde invoeg-scenario. In deze analyse komen belangrijke aspecten van het gedrag van bestuurders in interacties naar voren. Ten eerste gebruiken mensen discontinue constante inputs voor de acceleratie van hun voertuigen. Daarnaast laat het experiment zien dat de uitkomst van de interactie (welke auto eerst gaat) voornamelijk afhangt van de kinematica van de voertuigen aan het begin van de interactie, niet van individuele verschillen tussen de bestuurders. Hoofdstuk 5 draagt bij aan het fundamentele begrip van interactief gedrag tijdens het invoegingen door deze effecten te kwantificeren in een statistisch model. Ook de effecten van voertuigkinematica op de *Conflict Resolution Time* worden gekwantificeerd in Hoofdstuk 6.

Met deze inzichten uit de empirische analyse kon het model uit de casestudy in Hoofdstuk 4 worden verbeterd. Dit verbeterde gezamenlijke model wordt gepresenteerd in Hoofdstuk 7. De nieuwe versie van het model maakt gebruik van

discontinue constante inputs, zoals is waargenomen in Hoofdstuk 6. De *belief*-module is vernieuwd om de afstanden tussen voertuigen (veiligheidsmarges) te kunnen reproduceren. Tenslotte heeft dit model een ingebouwde stimulans voor individuele bestuurders om te handelen op basis van relatieve kinematica. Hierdoor kan het model bijvoorbeeld beschrijven dat een bestuurder die een auto volgt eerder geneigd is om actie te ondernemen om de afstand tussen de twee te vergroten dan de bestuurder van het voorste voertuig. Het model in Hoofdstuk 7 kan het individuele en gezamenlijke rijgedrag uit het experiment in Hoofdstuk 6 kwalitatief en kwantitatief beschrijven op meerdere niveaus. Beslissingen van individuele bestuurders (e.g., het wel of niet voorrang verlenen) leiden tot gezamenlijke uitkomsten (e.g., welke bestuurder eerst gaat). Snelheidsprofielen met menselijke karakteristieken beschrijven nauwkeurig de individuele bijdragen aan de gezamenlijke veiligheidsmarges (e.g., de ruimte tussen de voertuigen). Ten slotte reproduceert het model typische kwalitatieve interacties tussen bestuurders die zijn waargenomen in het experiment, zoals een miscommunicatie als beide bestuurders in eerste instantie dezelfde actie ondernemen.

Dit proefschrift heeft drie overkoepelende conclusies: 1) Verschillende bestuurders reageren tactisch en operationeel verschillend op vergelijkbare interactieve situaties; daarom zouden modellen van menselijk rijgedrag operationele en tactische variabiliteit moeten vastleggen, daarnaast moeten ze hierop onafhankelijk worden beoordeeld. 2) Een belangrijk aspect voor een gezamenlijk model van rijgedrag is dat bestuurders hun acceleratie-inputs niet voortdurend (rationeel) optimaliseren; in plaats daarvan gebruiken ze discontinue constante inputs – zoals empirisch waargenomen in een vereenvoudigd invoegscenario. 3) Met het modelleren van communicatie en op risico gebaseerde discontinue constante inputs kan het CEI-model abstracte interacties tussen twee bestuurders beschrijven, inclusief hun beslissingen (wie gaat er eerst), veiligheidsmarges over tijd, en het onderliggende individuele gedrag (remmen/accelereren). De belangrijkste beperking van dit proefschrift is dat het model en het experiment beide een vereenvoudigd invoegscenario gebruiken. In dit scenario zijn er slechts twee bestuurders die alleen de snelheid van hun auto kunnen controleren; ze kunnen niet sturen. Om de resultaten van dit proefschrift toe te kunnen passen in dagelijks verkeer is het uitbreiden van het model naar scenario's met meer voertuigen en volledige controle noodzakelijk. De conclusies, beperkingen, mogelijkheden voor verder onderzoek, en de potentiële toepassingen voor het model worden verder uiteengezet in Hoofdstuk 8.

Samenvattend heeft het werk in dit proefschrift een bijdrage geleverd aan de kennis over, en het modelleren van, menselijke gedrag bij interacties tijdens het wisselen van rijstrook en het invoegen. De studies in de hoofdstukken van dit proefschrift hebben niet alleen de tekortkomingen van modellen in bestaande interactieve AV's geïdentificeerd op basis van natuurlijk rijgedrag, maar ook een potentiële oplossing voorgesteld voor het modelleren van verkeersinteracties. Dit nieuwe model heeft in een vereenvoudigd invoegscenario veelbelovende resultaten laten zien. Daarom hoop ik dat we met deze studies een stap hebben gezet in de richting van het ontwikkelen van geautomatiseerd rijgedrag op basis van kennis van interacties in het verkeer.

1

Introduction

We live in a time most interesting for researchers and engineers working on automated vehicles (AVs); the development of AVs is booming. Impressive works on AV perception, path planning, and decision-making are published in videos and research articles almost weekly. This high research output can be partially explained by the available (financial) resources and the great public interest in the topic. But it cannot be denied that a major factor in this activity is the large number of challenging open problems that still exist for automated driving. One of these open problems is how autonomous and automated vehicles should handle interactive driving scenarios such as merging on a highway.

While access to commercial autonomous vehicles for scientific research is limited, online videos provide examples of the difficulties AVs have with interactive scenarios. A video on YouTube (Figure 1.1) shows how a white AV tries to merge onto a highway with heavy traffic (Figure 1.1-1). A grey human-driven vehicle drives directly behind the white AV in the merging lane (Figure 1.1-2). This vehicle merges in quickly without any problems. The AV, however, keeps driving in the merging lane until it reaches the end of the lane, where it comes to a full stop (Figure 1.1-3). Even when the vehicle that recorded the video leaves a large gap as an offer to the AV to merge, it remains stationary (Figure 1.1-4).

This anecdotal evidence shows that this AV has yet to reach a level of driving comparable to human drivers in a merging scenario. The AV fails to merge in on its own and does not take the offered gap, showing that it does not understand the communicative action of the camera car. In 2023, Brown et al. systematically reviewed online videos where AVs interact with human-driven cars [1]. They show multiple examples of awkward interactions between AVs and human traffic participants. In one example, the passenger of an AV even opens the car window to apologise to pedestrians it was interacting with: "Sorry! Self-driving car." In this case, the passenger and the pedestrians understand the interaction in a way the AV misses. Brown et al. concluded that: "To build self-driving cars that can integrate in traffic, [...] understanding traffic has a fundamental role to play in debates and design of self-driving vehicles. [...] The challenge then is not one of computability but of understanding social interaction."

These traffic interactions are difficult to "understand" for AVs because of the inherent risk of colliding with other vehicles and how they try to minimize this risk. While human drivers share the road with automated vehicles, this collision risk must be managed in interaction (and in conjunction) with human road users. However, these human drivers are accustomed to traffic interactions with other humans, and much is still unknown about how human drivers interact with each other in traffic [2]. During an interaction, all drivers (human or automated) can adapt their control behaviour continuously to communicate their intent [1]–[3]. Drivers use their vehicles' kinematics –velocity, position, and the gaps with respect to others– to communicate their intent [2]. They make individual decisions and exhibit individual behaviours, such as whether to merge now or to wait and how and when to accelerate. These individual behaviours lead to joint outcomes, such as who goes first in a merge, and joint continuous behaviour, such as the size of safety margins between the vehicles. Without a good understanding of joint and individual human interactive behaviour, or in the words of Brown et al. without "understanding social interaction", it will be challenging to design and validate safe and acceptable automated behaviour in interactions.

The common solution to enable interactions in automated vehicles is using Interaction-Aware controllers (IACs, e.g., [4]–[6]). These controllers incorporate




Figure 1.1: Stills from a video showing an autonomous vehicle merging onto a busy highway. Orange arrows show the positions of the (white) autonomous vehicle. The green arrow shows points out the merging human-driven vehicle. source: <https://www.youtube.com/watch?v=NMGYv0HpmOI>

an understanding of human driver interactions in the form of a driver model (see Figure 1.2). IACs use these models to predict the future trajectories of other vehicles and make driving decisions based on these predictions. However, such predictions are uncertain; therefore, driving decisions become a trade-off between safety and speed. The AV can either take more risk (i.e., be more assertive) and proceed quicker or be more conservative but slower. Because AVs are often designed to be safe above everything else, this can lead to a situation where the AV cannot find a plan it regards as safe enough to proceed. In this case, it will decide to stop moving. This issue is known as the freezing robot problem and is directly related to the accuracy of the driver model used [7].

In less extreme situations, the AV might not come to a full stop but can still behave too conservatively from a human perspective. Un-human-like behaviour in interactions could be difficult to understand and predict from the perspective of other drivers on the road. This could cause annoyance or even unsafe situations. For example, consider an intersection with two vehicles, where the vehicle with the right-of-way seems to stop and yield (for similar real-world examples, see [1], [2]). This will be annoying for the other driver. It could even be dangerous if both vehicles believe the other will yield. Thus, unnatural traffic interactions could lead to unsafe behaviour and decreased acceptability of AV behaviour by both the passengers of the AV and the other drivers it is interacting with. To prevent situations like this, the models that inform automated driving technology should accurately capture the higher-level outcomes and decisions in an interaction, such as which vehicle goes first, and the lower-level communicative signals drivers use to communicate their intent. Combined, these form the underlying principles of driver behaviour in traffic interactions.

To limit the scope, the focus of this thesis will be on highway interactions, more specifically on merging and lane changing. These interactions contain commu-

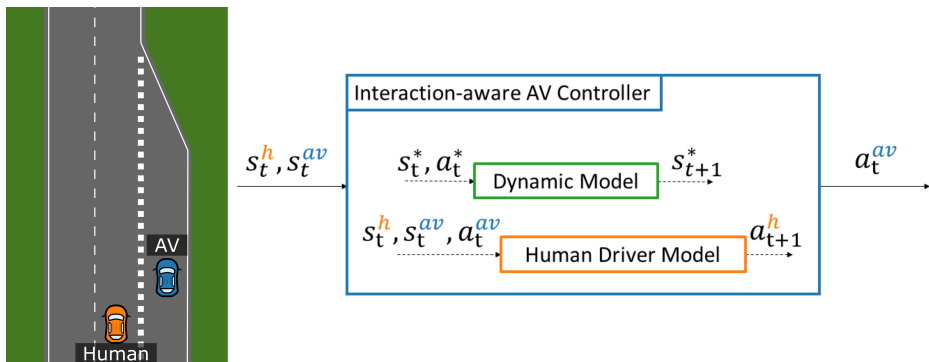


Figure 1.2: A schematic depiction of the inner workings of Interaction-Aware Controllers (from Chapter 2). Interaction-aware controllers for Autonomous Vehicles (AVs) are developed to let AVs interact with human-driven vehicles in interactive scenarios, such as the highway merging scenario on the left. These controllers use two internal models to predict the future state (s) of the vehicles and human response to their actions (a).

nicative signals and higher-level decisions, such as yielding described earlier (see Figure 1.1). Furthermore, they occur in a relatively structured environment without vulnerable road users such as pedestrians and bicyclists. Finally, these interactions can occur with just two drivers but can include multiple vehicles and are thus scalable. Therefore, they form a suitable scenario to study these vehicle-vehicle interactions.

Figure 1.3 (panel A) shows an example of such a merging interaction from a dataset containing vehicle trajectories recorded on German highways: the HighD dataset [8]. In this example, the green vehicle wants to merge onto the highway. It shows how this manoeuvre plays out by showing the positions and velocities of the green and purple vehicles over time (Panel B). The behaviour of the interacting (green and purple) vehicles can be broken down into three levels: Decisions, safety margin, and control inputs (Panel C). Input behaviour can only be seen from the perspective of a single vehicle, but the other two levels can also be viewed from the perspective of the joint behaviour of the two vehicles.

The green vehicle has a head start at the beginning of the interaction but travels at a significantly lower velocity. However, the joint outcome (i.e. decision) is that the green vehicle merges in front of the purple vehicle (i.e., Green goes first). This can be seen in the top plot (panel B) at the time when Green crosses the lane marker (marker I). Both vehicles individually contribute to this decision: Green speeds up while Purple slows down. These individual decisions (e.g., to yield or not, marked with 1) contribute to the joint decision that Green goes first. The joint safety margin at the merge is depicted by the gap plot, which shows the distance between the vehicles at the time of the merge (marker II). However, the velocity plot reveals (at marker 2) that the green vehicle provides an extra individual contribution to this safety margin: Green speeds up even further just before the merge, while purple stops decelerating. This is purely to increase the gap; the decision that green would go first was already made at marker 1. These individual contributions to the gap can also be used as a means of communication [2]. The velocity plot also provides an intuitive and general insight into the characteristics of control inputs. In this example, it can be seen that the green vehicles mostly maintain a constant velocity at the start and end of the interaction, with a linear acceleration phase in

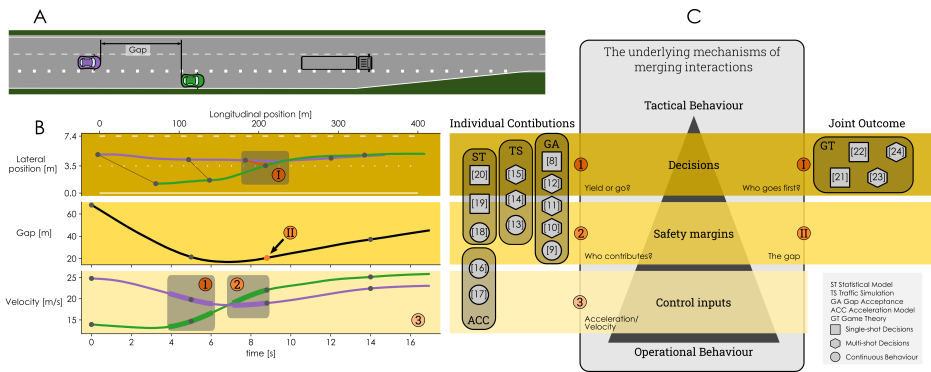


Figure 1.3: An introduction to modelling of interactive highway scenarios. Panel A shows a schematic representation of a typical interactive scenario: the green vehicle wants to merge onto the Highway. It has a position advantage but a significantly lower velocity than the Purple vehicle. This example is taken from the HighD dataset [8]. Panel B shows the individual position traces, the gap between the vehicles, and the individual velocity traces. Markers indicate individual and joint perspectives on the three levels of behaviour indicated in panel C. Panel C shows three levels of behaviour in between Michon's *operational* and *tactical* behaviour [26]. It also shows four types of driver models, each with examples from literature, that capture part of the interactive behaviour. A hypothetical model that would capture all three levels on the individual and joint levels is more likely to have captured the true underlying mechanisms of driver interactions.

between. The velocity plot is preferred to obtain insights into the input behaviour over an acceleration plot because the velocities are easier to understand intuitively.

There have been many efforts to model highway traffic interactions during merges and lane-changing. These models can be classified into five model classes (panel C): Gap acceptance, traffic simulation, statistical, acceleration-based and game-theoretic models. This overview includes references to examples of driver models within each of these classes. Control approaches are not included in this overview because even though they can be inspired by human behaviour, they do not aim to reproduce human behaviour as closely as possible (including its negative aspects). Each model class in the overview targets (and is evaluated on) a specific part of the individual and joint levels of behaviour just described.

Gap acceptance (GA in Figure 1.3) models (e.g. [9]–[13]) describe the decisions made by the individual drivers that want to merge onto the highway (e.g., the green vehicle) by evaluating available gaps between other vehicles in heavy traffic conditions. Drivers decide if a gap is large enough to merge into based on a personal threshold: the preferred individual safety margin. Merging models that are used in traffic simulations (TS in Figure 1.3) often rely on the same gap acceptance theory [14]–[16] and are sometimes combined with acceleration-based models that describe the lower-level inputs of drivers [17], [18]. Statistical models (ST in Figure 1.3) describe the probability that a certain vehicle will merge or change lanes. Some include desired safety margins [19], [20], while others do not [21]. Finally, game-theoretic (GT in Figure 1.3) models describe the high-level joint outcome (who goes first) and decision-making in interactions by considering multiple drivers in a single model [22]–[25], i.e., a **joint driver model**.

Panel C also shows that the fundamental understanding of how the levels are connected is missing; no driver model yet covers all aspects of a merging interaction between specific drivers. The example in panel B highlights how the individual

contributions of the vehicles lead to a joint outcome (e.g., how two individual decisions lead to Green going first). This indicates that to understand merging interactions fully, they should be regarded as a joint system of drivers with individual contributions. Studying the (modelled) behaviour of one of the two drivers while disregarding or fixing the others' behaviour disregards the interactive aspects of this system. This could mean (unintentionally) dividing a system into multiple parts that inherently belong together, changing the dynamics. Therefore, I argue that merging interactions can be best described and modelled as a joint system of drivers. In this thesis, I refer to this approach as a **joint driver model**: a model that describes multiple drivers' joint and individual behaviours. A joint driver model that accurately captures the aspects of merging interactions at these three levels is also more likely to have captured the underlying mechanisms of interactive driving behaviour in general.

Such a model would have a large potential for practical applications since it could be used to actively inform AVs and help them make safe and acceptable decisions. It could provide a more fundamental understanding of the important aspects of human driving interactions, such as how drivers form beliefs about others based on communicative actions and how these beliefs influence their future actions. A fundamental understanding of how humans handle merging and lane-changing situations can, in turn, help design automated behaviours that are understood and accepted by humans. Therefore, **the aim of this thesis is to increase the fundamental understanding of merging and lane-changing interactions and capture this knowledge in a joint driver model.**

1.1. The three pillars of this thesis

Three pillars form the foundation for the work in this thesis to increase the fundamental understanding of driver behaviour in merging and lane-changing interactions: Natural(istic) driving, controlled experiments, and model theory (Figure 1.4). Naturalistic driving considers driver behaviour in the real world; this is often studied with datasets recorded in real traffic, commonly referred to as naturalistic datasets. Controlled experiments are usually conducted in driving simulators or on test tracks and aim to uncover the underlying principles of human driving behaviour. Model theory covers the design of driver behaviour models and the assumptions about human behaviour made in these models. In this section, I will provide background in these three areas and discuss the (previously) open questions that led to the work in this thesis.

1.1.1. Natural(istic) behaviour

The main motivation behind investigating human traffic interactions in this thesis is to improve the interactive capabilities of automated vehicles in real-world traffic. Therefore, it seems logical to start our investigation there: with natural driver behaviour. Multiple datasets were recently published that were recorded on real roads, for example, in Germany [8], Greece [27], China [28], and the US [28], [29]. These datasets provide the opportunity to investigate natural traffic interactions. These naturalistic datasets contain a large range of behaviours. Specifically in interactions, multiple levels of driver behaviour are important to consider: drivers can respond to, for example, a merging vehicle, by deciding what to do (e.g., change lanes) and how to do it (e.g., fast or slow). These levels have previously been described by Michon [26], who distinguished between operational, tactical and strategic behaviour. Tactical behaviour is the choice of manoeuvre (i.e., what

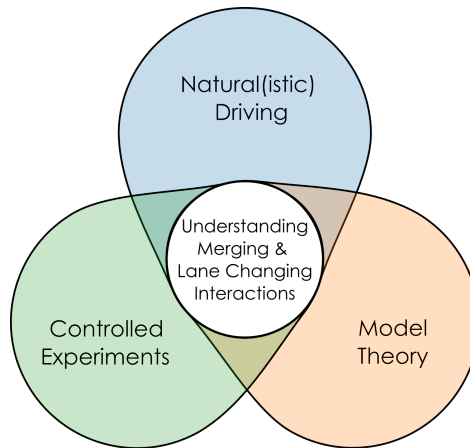


Figure 1.4: An overview of the three pillars used in this thesis to increase the fundamental understanding of merging and lane-changing interactions and capture this knowledge in a joint driver model.

to do), and operational behaviour regards how a this is executed (i.e., how to do it). Strategic behaviour is a higher level that concerns decisions such as route choice and is therefore not relevant to the traffic interactions in this thesis.

A large amount of work has been devoted to studying traffic behaviour in the real world using naturalistic datasets [19], [20], [30]–[32]. Most of these empirical studies were motivated to better understand the overall traffic flow [19], [20], [31]. Therefore, these studies focus on the operational behaviour of the whole population within specific tactical behaviours (e.g., lane changes or car following), or on the distribution between tactical behaviours. For example, these studies aim to find the probability that vehicles merge under specific conditions (e.g. [19]) or try to describe the distribution of the safety margins that are kept when drivers are following another car (e.g. [33]). This makes these studies very suitable for understanding (and improving) the overall traffic system, but individual vehicle interactions are obscured. Specifically, **it is unknown how to uncover the operational and tactical variability in responses of individual drivers facing the same (real-world) scenario.**

These naturalistic datasets have not been widely used in research on interaction-aware autonomous driving. State-of-the-art interaction-aware controllers (from literature) have never been evaluated in real-world traffic interactions [4]–[6]. These controllers have been developed and tested in simulated (top-down-view) environments. Deploying them in real vehicles or even real-world traffic to evaluate their interactive capabilities could potentially be dangerous. The available datasets could be promising environments to gain more insight into the potential of these IACs. Therefore, an open question is: **how can we leverage the available naturalistic data for a preliminary evaluation of interaction-aware controllers?** The answer to this question could also be a step towards answering another open question: how well will current interaction-aware controllers perform in real traffic? However, this is beyond the scope of this thesis.

1.1.2. Modelling theory

Driving behaviour has been a research topic ever since the automobile was invented, and this research has yielded many models. To name just a few to show-

case the depth and width of this field: In 1938, Gibson and Crooks reasoned that driving behaviour and specifically choices about what trajectory to use can be explained by a "field of safe travel" [34]. In 1976, Lee presented a model that linked the braking actions of drivers directly to the capability of the human visual system to estimate relative velocities from the increasing or decreasing size of objects in the field of view (called looming) [35]. The intelligent driver model was proposed in 2000 by Treiber et al. to describe car-following behaviour and has been used in many studies since [18]. Finally, and more recently, in 2020, Kolekar et al. proposed a unified risk-based driver model that can describe driver behaviour in 7 different (non-interactive) scenarios [36].

In this wide variety of driver models, only a limited number of models target traffic interactions and regard all drivers in the interaction. Some others focus on interactive scenarios but only describe a single driver within this scenario. Four types of driver models that aim to describe such merging and lane-changing scenarios were identified earlier (see Figure 1.3): Gap acceptance, traffic simulation, statistical, and game theoretic models.

Gap acceptance models (e.g. [9]–[13]) describe the decisions made by the individual merging drivers by evaluating available gaps between other vehicles in heavy traffic conditions. Drivers decide if a gap is large enough to merge into based on a personal threshold. Merging models used in traffic simulations often rely on the same gap acceptance theory [14]–[16]. In these simulations, the models of the high-level decision to accept a gap are complemented with car-following models (e.g., the intelligent driver model (IDM) [18]). This way, the control behaviour before and after the merging decision is included in the modelled behaviour. Thus, two models are stitched together to describe a full merging manoeuvre of a vehicle. This might describe the population well in traffic simulations; however, it remains unknown how well this combination of models describes individual interactions.

Statistical models are based on the previously discussed naturalistic traffic data. These models describe the probability that a certain vehicle will merge or change lanes. Some include desired safety margins [19], [20], while others do not [21]. As with the gap acceptance models, these models only describe high-level decisions for individual vehicles. This raises the question of whether models of individual vehicles are the best way to describe interactions that inherently play out between multiple vehicles.

Game-theoretic models describe the interaction by considering multiple drivers in a single model. Game theory describes the high-level outcome and decision making in the merging process [22]–[25] (for a more extensive review of game-theoretic merging and lane-changing models see [37]). Game theoretic models have in common that they assume humans to be rational utility-maximizing agents that do not communicate. Are these assumptions valid when modelling the merging behaviour of drivers on the highway?

In conclusion, many approaches to driver modelling have been explored in the past. These approaches all have their strengths and limitations, leaving us with the open question: **How to model the joint dynamics of multiple drivers in merging interactions?**

1.1.3. Controlled experiments

Besides naturalistic data, controlled experiments are a valuable data source for studying driver behaviour. These experiments are used to investigate how specific

controlled variables influence behaviour. They can take place in actual vehicles in a controlled environment, such as on a test track (e.g., [38], [39]) or in the lab in driving simulators (e.g., [40]–[42]). Many aspects of driving behaviour have been studied, such as steering behaviour [41], braking behaviour [39], [42], and risk perception [38], [40]. However, in all these cases, only a single driver participated in the experiment.

Studying the behaviour of a single driver in a controlled experiment entails recording many different signals. This makes the analysis of such experiments a time-consuming effort. An often-used solution to this problem is to capture the behaviour in metrics that describe the aspects of behaviour we are interested in. For example, in car following, the safety margin is usually described in terms of time to collision (TTC) (e.g., [43]). This is the time until two vehicles collide, given that they keep driving at their current velocities. This metric captures the distance between the vehicles and the relative velocity between the vehicles. However, this metric is only valid when vehicles follow each other in a single lane. What metrics accurately describe the joint behaviour in interactive scenarios where multiple tactical behaviours (such as lane changing and overtaking) are possible is currently unknown.

This problem of the lack of meaningful metrics to analyse becomes larger when multiple drivers are involved. Including more drivers in the experiment means recording more input signals, which could hinder the analysis of interactive behaviour in driving simulators. An experiment with a scenario specifically designed to reflect traffic interactions yet limit the input signals per driver is needed to enable such studies. However, **it is unclear what scenario and metrics to use in a controlled experiment to investigate merging interactions in traffic.**

With such an experiment, driver behaviour in a coupled simulator can be investigated. This will answer questions such as: what is the input strategy drivers use during interactions? What individual contributions lead to a specific high-level outcome (e.g., who goes first in a merging interaction)? How do drivers respond to differences in relative velocity and position? Or, to put it more generally: **How do interacting drivers behave in terms of decisions, safety margins, and control inputs?**

1.2. Structure and approach

All the bold questions posed in the previous section will be answered in the following chapters of this thesis. Figure 1.5 shows an overview of which chapter contributes to which pillar. Each chapter, its aim, and main findings will be briefly discussed here.

In Chapter 2 a framework is proposed to validate driver models used for interaction-aware controllers. This framework provides a method to leverage naturalistic data for this validation. A case study shows that a state-of-the-art inverse-reinforcement-learning-based driver model, similar to those used in interaction-aware controllers, does not capture human driving behaviour, motivating the research into new models for traffic interactions.

Chapter 3 investigates how to uncover the variability in naturalistic human driving behaviour. It proposes a method to extract similar traffic scenes from large naturalistic datasets automatically. A case study shows that there is variability present in tactical and operational behaviour. Appendix A presents a software package called TraViA (for Traffic Visualisation and Annotation) which was specially developed for the work in chapters 2 & 3.

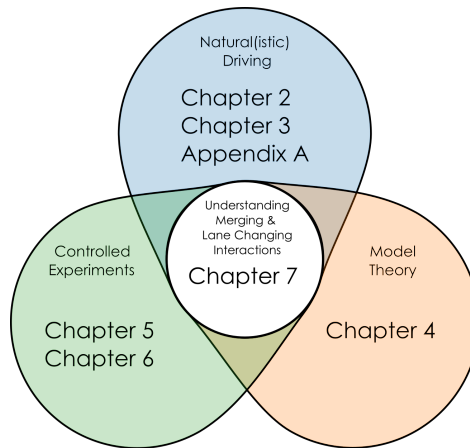


Figure 1.5: Contributions of the chapters in this thesis to the final goal: "to increase the fundamental understanding of merging and lane-changing interactions and capture this knowledge in a joint driver model"

In Chapter 4, the Communication-Enabled Interaction model (CEI-model) is introduced. This is a model of traffic interactions that includes communication between drivers and attributes drivers' actions to risk perception following from a probabilistic belief about the other driver's future trajectory. This chapter discusses the theoretical background for modelling traffic interactions, which led to the development of the CEI framework. The chapter also includes simulations of the model that show plausible human-like behaviours in a merging interaction. Furthermore, human-like gap-keeping behaviour in a car-following scenario emerged from the model.

In Chapter 5 an experiment in a controlled environment is designed, including a simplified merging scenario and novel analysis tools. This simplified merging scenario is used to investigate model and human behaviour from this chapter onward. The chapter introduces the experiment used to gather data on human-human merging interactions in a coupled top-down-view driving simulator. It also describes three analysis tools, specifically developed to analyse the behaviour of a pair of drivers. One of these tools is the Conflict Resolution Time (CRT), a metric that describes the time it took the drivers to resolve the conflict.

Chapter 6 provides an empirical analysis of the data gathered in the merging experiment. This includes a qualitative analysis of individual human behaviour and statistical models of both individual and pair-wise behaviour. This chapter aims to answer the question: How do humans behave in interactive merging? The conclusions of this chapter were used to improve the CEI model presented in Chapter 4.

Chapter 7 presents an improved version of the CEI model. This chapter also validates this final model, using the data gathered from the experiment in Chapters 5 and 6. The results show that the model can accurately describe individual and joint driver behaviour in terms of high-level decisions, safety margins, and control inputs. It also discusses the new insights the model gave into how drivers solve merging conflicts.

To conclude the thesis, the overarching conclusions are drawn and discussed in Chapter 8. This final chapter also includes a discussion on the results, potential applications of the model, and an outline of future work.

1.3. Remarks on publications

The following chapters have been previously published or submitted for publication and were included in the thesis unaltered:

- Chapter 2 was published as:

O. Siebinga, A. Zgonnikov, and D. Abbink, "A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles", *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 4, pp. 1–21, Dec. 2022, ISSN: 2573-9522. DOI: [10.1145/3538705](https://doi.org/10.1145/3538705). [Online]. Available: <https://dl.acm.org/doi/10.1145/3538705>
- Chapter 3 was published as:

O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Uncovering Variability in Human Driving Behavior Through Automatic Extraction of Similar Traffic Scenes from Large Naturalistic Datasets", in *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, IEEE, Oct. 2023, pp. 4790–4796. DOI: [10.1109/SMC53992.2023.10393913](https://doi.org/10.1109/SMC53992.2023.10393913). eprint: [2206.13386](https://arxiv.org/abs/2206.13386). [Online]. Available: <https://ieeexplore.ieee.org/document/10393913/>
- Chapter 4 was published as:

O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Modelling communication-enabled traffic interactions", *Royal Society Open Science*, vol. 10, no. 5, May 2023, ISSN: 2054-5703. DOI: [10.1098/rsos.230537](https://doi.org/10.1098/rsos.230537). [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsos.230537>
- Chapter 5 was published as:

O. Siebinga, A. Zgonnikov, and D. Abbink, "Interactive merging behavior in a coupled driving simulator: Experimental framework and case study", *Human Factors in Transportation*, vol. 60, pp. 516–525, 2022. DOI: [10.54941/ahfe1002485](https://doi.org/10.54941/ahfe1002485)
- Chapter 6 was published as:

O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Human Merging Behavior in a Coupled Driving Simulator: How Do We Resolve Conflicts?", *IEEE Open Journal of Intelligent Transportation Systems*, vol. 5, no. October 2023, pp. 103–114, 2024, ISSN: 2687-7813. DOI: [10.1109/OJITS.2024.3349635](https://doi.org/10.1109/OJITS.2024.3349635). arXiv: [2308.04842](https://arxiv.org/abs/2308.04842). [Online]. Available: <http://arxiv.org/abs/2308.04842%20https://ieeexplore.ieee.org/document/10380755/>
- Chapter 7 has been submitted for publication in a journal and is available at:

O. Siebinga, A. Zgonnikov, and D. Abbink, "A merging interaction model explains human drivers' behaviour from input signals to decisions", pp. 1–18, Dec. 2023. arXiv: [2312.09776](https://arxiv.org/abs/2312.09776). [Online]. Available: <https://arxiv.org/abs/2312.09776>

Because these chapters have been published or submitted to different journals and conferences, they differ in their use of British or American English. Chapters 2, 3, and 5 are written in American English while Chapters 4, 6, and 7 use British English. All other texts are also in British English. I choose to include the chapters unaltered to preserve their standalone readability. This means that some information and some figures are repeated throughout the thesis.

Bibliography

- [1] B. Brown, M. Broth, and E. Vinkhuyzen, "The Halting problem: Video analysis of self-driving cars in traffic", *Conference on Human Factors in Computing Systems - Proceedings*, 2023. DOI: [10.1145/3544548.3581045](https://doi.org/10.1145/3544548.3581045).
- [2] B. Brown, E. Laurier, and E. Vinkhuyzen, "Designing Motion: Lessons for Self-driving and Robotic Motion from Human Traffic Interaction", *Proceedings of the ACM on Human-Computer Interaction*, vol. 7, no. GROUP, pp. 1–21, 2023. DOI: [10.1145/3567555](https://doi.org/10.1145/3567555).
- [3] D. Dey and J. Terken, "Pedestrian interaction with vehicles: Roles of explicit and implicit communication", *AutomotiveUI 2017 - 9th International ACM Conference on Automotive User Interfaces and Interactive Vehicular Applications, Proceedings*, pp. 109–113, 2017. DOI: [10.1145/3122986.3123009](https://doi.org/10.1145/3122986.3123009).
- [4] D. Sadigh, N. Landolfi, S. S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state", *Autonomous Robots*, vol. 42, no. 7, pp. 1405–1426, Oct. 2018, ISSN: 0929-5593. DOI: [10.1007/s10514-018-9746-1](https://doi.org/10.1007/s10514-018-9746-1). [Online]. Available: <https://doi.org/10.1007/s10514-018-9746-1><http://link.springer.com/10.1007/s10514-018-9746-1>.
- [5] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [6] E. Ward, N. Evestedt, D. Axehill, and J. Folkesson, "Probabilistic Model for Interaction Aware Planning in Merge Scenarios", *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 2, pp. 1–1, 2017, ISSN: 2379-8904. DOI: [10.1109/TIV.2017.2730588](https://doi.org/10.1109/TIV.2017.2730588). [Online]. Available: <http://ieeexplore.ieee.org/document/7987753/>.
- [7] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds", *IEEE/RJSJ 2010 International Conference on Intelligent Robots and Systems, IROS 2010 - Conference Proceedings*, pp. 797–803, 2010. DOI: [10.1109/IROS.2010.5654369](https://doi.org/10.1109/IROS.2010.5654369).
- [8] R. Krajewski, J. Bock, L. Kloeker, and L. 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)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 2118–2125, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569552](https://doi.org/10.1109/ITSC.2018.8569552). arXiv: [1810.05642](https://arxiv.org/abs/1810.05642). [Online]. Available: <http://ieeexplore.ieee.org/document/8569552/>.
- [9] A. Kondyli and L. Eleftheriadou, "Modeling driver behavior at freeway-ramp merges", *Transportation Research Record*, no. 2249, pp. 29–37, 2011, ISSN: 03611981. DOI: [10.3141/2249-05](https://doi.org/10.3141/2249-05).
- [10] J. A. Laval and L. Leclercq, "Microscopic modeling of the relaxation phenomenon using a macroscopic lane-changing model", *Transportation Research Part B: Methodological*, vol. 42, no. 6, pp. 511–522, 2008, ISSN: 01912615. DOI: [10.1016/j.trb.2007.10.004](https://doi.org/10.1016/j.trb.2007.10.004).
- [11] C. F. Choudhury, "Modeling Driving Decision with Latent Plans (PhD Dissertation)", no. 2002, 2007.
- [12] R. M. Michaels and J. Fazio, "Driver behavior model of merging", *Transportation Research Record*, no. 1213, pp. 4–10, 1989, ISSN: 03611981.
- [13] Kazi Iftekhhar Ahmed, "Modeling Drivers' Acceleration and Lane Changing Behavior", Ph.D. dissertation, 1999.
- [14] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model MOBIL for car-following models", *Transportation Research Record*, no. 1999, pp. 86–94, 2007, ISSN: 03611981. DOI: [10.3141/1999-10](https://doi.org/10.3141/1999-10).

- [15] Q. Yang and H. N. Koutsopoulos, "A microscopic traffic simulator for evaluation of dynamic traffic management systems", *Transportation Research Part C: Emerging Technologies*, vol. 4, no. 3 PART C, pp. 113–129, 1996, ISSN: 0968090X. DOI: [10.1016/S0968-090X\(96\)00006-X](https://doi.org/10.1016/S0968-090X(96)00006-X).
- [16] P. Hidas, "Modelling lane changing and merging in microscopic traffic simulation", *Transportation Research Part C: Emerging Technologies*, vol. 10, no. 5-6, pp. 351–371, 2002, ISSN: 0968090X. DOI: [10.1016/S0968-090X\(02\)00026-8](https://doi.org/10.1016/S0968-090X(02)00026-8).
- [17] X. Wan, P. J. Jin, F. Yang, J. Zhang, and B. Ran, "Modeling Vehicle Interactions during Merge in Congested Weaving Section of Freeway Ramp", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2421, no. 1, pp. 82–92, 2014, ISSN: 0361-1981. DOI: [10.3141/2421-10](https://doi.org/10.3141/2421-10).
- [18] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, vol. 62, no. 2, pp. 1805–1824, 2000, ISSN: 1063651X. DOI: [10.1103/PhysRevE.62.1805](https://doi.org/10.1103/PhysRevE.62.1805). arXiv: [0002177](https://arxiv.org/abs/0002177) [cond-mat].
- [19] W. Daamen, M. Looft, and S. P. Hoogendoorn, "Empirical analysis of merging behavior at freeway on-ramp", *Transportation Research Record*, no. 2188, pp. 108–118, 2010, ISSN: 03611981. DOI: [10.3141/2188-12](https://doi.org/10.3141/2188-12).
- [20] F. Marczak, W. Daamen, and C. Buisson, "Merging behaviour: Empirical comparison between two sites and new theory development", *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 530–546, 2013, ISSN: 0968090X. DOI: [10.1016/j.trc.2013.07.007](https://doi.org/10.1016/j.trc.2013.07.007). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2013.07.007>.
- [21] C. Dong, J. M. Dolan, and B. Litkouhi, "Smooth Behavioral Estimation for Ramp Merging Control in Autonomous Driving", *IEEE Intelligent Vehicles Symposium, Proceedings*, vol. 2018-June, no. Iv, pp. 1692–1697, 2018. DOI: [10.1109/IVS.2018.8500576](https://doi.org/10.1109/IVS.2018.8500576).
- [22] H. Kita, "A merging-giveway interaction model of cars in a merging section: A game theoretic analysis", *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3-4, pp. 305–312, 1999, ISSN: 09658564. DOI: [10.1016/S0965-8564\(98\)00039-1](https://doi.org/10.1016/S0965-8564(98)00039-1).
- [23] H. Liu, W. Xin, Z. Adam, and J. Ban, "A game theoretical approach for modelling merging and yielding behaviour at freeway on-ramp section", *Transportation and Traffic Theory*, no. January, pp. 1–15, 2007. [Online]. Available: [http://www.ce.umn.edu/\\$%5Csim\\$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf](http://www.ce.umn.edu/$%5Csim$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf).
- [24] S. Coskun, Q. Zhang, and R. Langari, "Receding Horizon Markov Game Autonomous Driving Strategy", in *2019 American Control Conference (ACC)*, vol. 2019-July, IEEE, Jul. 2019, pp. 1367–1374, ISBN: 978-1-5386-7926-5. DOI: [10.23919/ACC.2019.8815251](https://doi.org/10.23919/ACC.2019.8815251). [Online]. Available: <https://ieeexplore.ieee.org/document/8815251/>.
- [25] N. Li, D. W. Oyler, M. Zhang, Y. Yildiz, I. Kolmanovsky, and A. R. Girard, "Game theoretic modeling of driver and vehicle interactions for verification and validation of autonomous vehicle control systems", *IEEE Transactions on Control Systems Technology*, vol. 26, no. 5, pp. 1782–1797, 2018, ISSN: 10636536. DOI: [10.1109/TCST.2017.2723574](https://doi.org/10.1109/TCST.2017.2723574). arXiv: [1608.08589](https://arxiv.org/abs/1608.08589).
- [26] J. A. Michon, "A Critical View of Driver Behavior Models: What Do We Know, What Should We Do?", in *Human Behavior and Traffic Safety*, Boston, MA: Springer US, 1985, pp. 485–524, ISBN: 0306422255. DOI: [10.1007/978-1-4613-2173-6_19](https://doi.org/10.1007/978-1-4613-2173-6_19). [Online]. Available: http://link.springer.com/10.1007/978-1-4613-2173-6_19.
- [27] E. Barmponakis and N. Geroliminis, "On the new era of urban traffic monitoring with massive drone data: The pNEUMA large-scale field experiment", *Transportation Research Part C: Emerging Technologies*, vol. 111, no. November 2019, pp. 50–71, 2020, ISSN: 0968090X. DOI: [10.1016/j.trc.2019.11.023](https://doi.org/10.1016/j.trc.2019.11.023). [Online]. Available: <https://doi.org/10.1016/j.trc.2019.11.023>.
- [28] W. Zhan, L. Sun, D. Wang, *et al.*, "INTERACTION Dataset: An INTERNATIONAL, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps", Sep. 2019. arXiv: [1910.03088](https://arxiv.org/abs/1910.03088). [Online]. Available: <http://arxiv.org/abs/1910.03088>.

- [29] U.S. Department of Transportation Federal Highway Administration, *Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data. [Dataset]*, 2016. [Online]. Available: <https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajectory/8ect-6jqj>.
- [30] A. R. Srinivasan, M. Hasan, Y.-S. Lin, et al., *Comparing merging behaviors observed in naturalistic data with behaviors generated by a machine learned model*, 2021. arXiv: 2104.10496 [cs.LG].
- [31] L. Klitzke, K. Gimm, C. Koch, and F. Koster, "Extraction and Analysis of Highway On-Ramp Merging Scenarios from Naturalistic Trajectory Data", in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, IEEE, Oct. 2022, pp. 654–660, ISBN: 978-1-6654-6880-0. DOI: 10.1109/ITSC55140.2022.9922191. arXiv: 2104.05661. [Online]. Available: <https://arxiv.org/abs/2104.05661v2%20https://ieeexplore.ieee.org/document/9922191/>.
- [32] H. Wang, W. Wang, S. Yuan, X. Li, and L. Sun, "On Social Interactions of Merging Behaviors at Highway On-Ramps in Congested Traffic", *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 11 237–11 248, 2022, ISSN: 15580016. DOI: 10.1109/TITS.2021.3102407. arXiv: 2008.06156.
- [33] C. Thiemann, M. Treiber, and A. Kesting, "Estimating acceleration and lane-changing dynamics from next generation simulation trajectory data", *Transportation Research Record*, no. 2088, pp. 90–101, 2008, ISSN: 03611981. DOI: 10.3141/2088-10.
- [34] J. J. Gibson and L. E. Crooks, "A Theoretical Field-Analysis of Automobile-Driving", *The American Journal of Psychology*, vol. 51, no. 3, pp. 453–471, 1938.
- [35] D. N. Lee, "A Theory of Visual Control of Braking Based on Information about Time-to-Collision", *Perception*, vol. 5, no. 4, pp. 437–459, Dec. 1976, ISSN: 0301-0066. DOI: 10.1068/p050437. [Online]. Available: <http://journals.sagepub.com/doi/10.1068/p050437>.
- [36] S. Kolekar, J. de Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model", *Nature Communications*, vol. 11, no. 1, p. 4850, Dec. 2020, ISSN: 2041-1723. DOI: 10.1038/s41467-020-18353-4. [Online]. Available: <http://dx.doi.org/10.1038/s41467-020-18353-4%20https://www.nature.com/articles/s41467-020-18353-4>.
- [37] A. Ji and D. Levinson, "A review of game theory models of lane changing", *Transportmetrica A: Transport Science*, vol. 9935, no. May, pp. 1–19, 2020, ISSN: 2324-9935. DOI: 10.1080/23249935.2020.1770368. [Online]. Available: <https://doi.org/10.1080/23249935.2020.1770368>.
- [38] S. Kolekar, B. Petermeijer, E. Boer, J. de Winter, and D. Abbink, "A risk field-based metric correlates with driver's perceived risk in manual and automated driving: A test-track study", *Transportation Research Part C: Emerging Technologies*, vol. 133, no. October, p. 103 428, 2021, ISSN: 0968090X. DOI: 10.1016/j.trc.2021.103428. [Online]. Available: <https://doi.org/10.1016/j.trc.2021.103428>.
- [39] J. Potzy, M. Feuerbach, and K. Bengler, "Communication strategies for automated merging in dense traffic", *IEEE Intelligent Vehicles Symposium, Proceedings*, vol. 2019-June, no. lv, pp. 2291–2298, 2019. DOI: 10.1109/IVS.2019.8813835.
- [40] S. Kolekar, J. de Winter, and D. Abbink, "Which parts of the road guide obstacle avoidance? Quantifying the driver's risk field", *Applied Ergonomics*, vol. 89, no. July, p. 103 196, Nov. 2020, ISSN: 00036870. DOI: 10.1016/j.apergo.2020.103196. [Online]. Available: <https://doi.org/10.1016/j.apergo.2020.103196%20https://linkinghub.elsevier.com/retrieve/pii/S0003687018307373>.
- [41] S. Barendsward, L. V. Breugel, B. Schelfaut, et al., "Effect of velocity and curve radius on driver steering behaviour before curve entry", *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, vol. 2019-October, pp. 3866–3871, 2019, ISSN: 1062922X. DOI: 10.1109/SMC.2019.8914263.
- [42] U. Durrani, C. Lee, and D. Shah, "Predicting driver reaction time and deceleration: Comparison of perception-reaction thresholds and evidence accumulation framework", *Accident Analysis and Prevention*, vol. 149, no. November 2020, p. 105 889, 2021, ISSN: 00014575. DOI: 10.1016/j.aap.2020.105889. [Online]. Available: <https://doi.org/10.1016/j.aap.2020.105889>.

- [43] T. Kondoh, T. Yamamura, S. Kitazaki, N. Kuge, and E. R. Boer, "Identification of Visual Cues and Quantification of Drivers' Perception of Proximity Risk to the Lead Vehicle in Car-Following Situations", *Journal of Mechanical Systems for Transportation and Logistics*, vol. 1, no. 2, pp. 170–180, 2008, ISSN: 1882-1782. DOI: [10.1299/jmtl.1.170](https://doi.org/10.1299/jmtl.1.170). [Online]. Available: http://www.jstage.jst.go.jp/article/jmtl/1/2/1%5C_2%5C_170/%5C_article.
- [44] O. Siebinga, A. Zgonnikov, and D. Abbink, "A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles", *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 4, pp. 1–21, Dec. 2022, ISSN: 2573-9522. DOI: [10.1145/3538705](https://doi.org/10.1145/3538705). [Online]. Available: <https://dl.acm.org/doi/10.1145/3538705>.
- [45] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Uncovering Variability in Human Driving Behavior Through Automatic Extraction of Similar Traffic Scenes from Large Naturalistic Datasets", in *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, IEEE, Oct. 2023, pp. 4790–4796. DOI: [10.1109/SMC53992.2023.10393913](https://doi.org/10.1109/SMC53992.2023.10393913). eprint: [2206.13386](https://arxiv.org/abs/2206.13386). [Online]. Available: <https://ieeexplore.ieee.org/document/10393913/>.
- [46] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Modelling communication-enabled traffic interactions", *Royal Society Open Science*, vol. 10, no. 5, May 2023, ISSN: 2054-5703. DOI: [10.1098/rsos.230537](https://doi.org/10.1098/rsos.230537). [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsos.230537>.
- [47] O. Siebinga, A. Zgonnikov, and D. Abbink, "Interactive merging behavior in a coupled driving simulator: Experimental framework and case study", *Human Factors in Transportation*, vol. 60, pp. 516–525, 2022. DOI: [10.54941/ahfe1002485](https://doi.org/10.54941/ahfe1002485).
- [48] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Human Merging Behavior in a Coupled Driving Simulator: How Do We Resolve Conflicts?", *IEEE Open Journal of Intelligent Transportation Systems*, vol. 5, no. October 2023, pp. 103–114, 2024, ISSN: 2687-7813. DOI: [10.1109/OJITS.2024.3349635](https://doi.org/10.1109/OJITS.2024.3349635). arXiv: [2308.04842](https://arxiv.org/abs/2308.04842). [Online]. Available: <http://arxiv.org/abs/2308.04842> <https://ieeexplore.ieee.org/document/10380755/>.
- [49] O. Siebinga, A. Zgonnikov, and D. Abbink, "A merging interaction model explains human drivers' behaviour from input signals to decisions", pp. 1–18, Dec. 2023. arXiv: [2312.09776](https://arxiv.org/abs/2312.09776). [Online]. Available: <https://arxiv.org/abs/2312.09776>.

2

A human factors approach to
validating driver models for
interaction-aware automated
vehicles

A major challenge for autonomous vehicles is interacting with other traffic participants safely and smoothly. A promising approach to handle such traffic interactions is equipping autonomous vehicles with interaction-aware controllers (IACs). These controllers predict how surrounding human drivers will respond to the autonomous vehicle's actions, based on a driver model. However, the predictive validity of driver models used in IACs is rarely validated, which can limit the interactive capabilities of IACs outside the simple simulated environments in which they are demonstrated. In this paper, we argue that besides evaluating the interactive capabilities of IACs, their underlying driver models should be validated on natural human driving behaviour. We propose a workflow for this validation that includes scenario-based data extraction and a two-stage (tactical/operational) evaluation procedure based on human factors literature. We demonstrate this workflow in a case study on an inverse-reinforcement-learning-based driver model replicated from an existing IAC. This model only showed the correct tactical behaviour in 40% of the predictions. The model's operational behaviour was inconsistent with observed human behaviour. The case study illustrates that a principled evaluation workflow is useful and needed. We believe that our workflow will support the development of appropriate driver models for future automated vehicles.

2.1. Introduction

One of the great technological and societal promises of the 21st century is the autonomous vehicle (AV) [1]–[3]. This technology has been under development in laboratories and under controlled conditions for decades and is now transitioning to the real world. However, a major challenge for real-world implementation of AV technologies is enabling AVs to handle complex interactions with human road users. AV controllers have recently been proposed that aim to address this challenge through *interaction-aware* controllers (IACs) [4]–[18]. IACs incorporate a model of human driver behaviour in the controller, to predict how another driver is likely to respond to the AV's behaviour. Based on this prediction and its own reward function (e.g., incorporating safety, comfort, etc.), the IAC finds the optimal action for the AV (Figure 2.1). However, up to now the interactive capabilities of these controllers have only been demonstrated in simplified simulated environments (e.g. top-down view computer simulations). Whether the state-of-the-art IACs are capable of predicting naturalistic driver behaviour and interacting with humans in real traffic remains an open question.

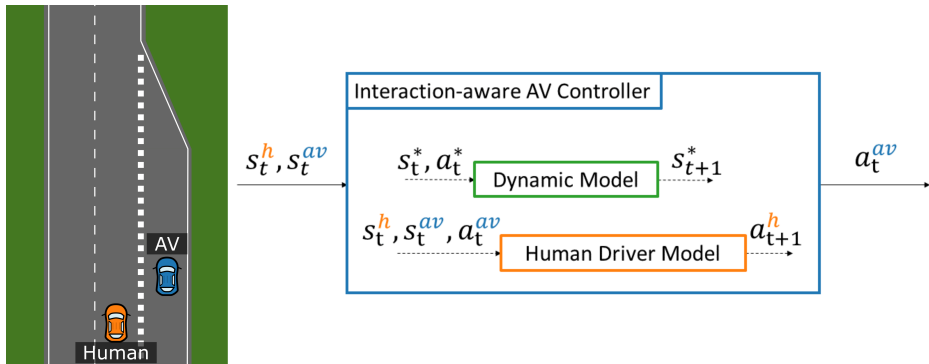


Figure 2.1: A high-level diagram of a typical interaction-aware controller (IAC) for autonomous vehicles (AVs). Such a controller operates in situations where the states and actions of a human-driven vehicle (superscript h) and an AV (superscript av) influence each other, e.g. the merging situation depicted in the left panel. Future states and actions are denoted with subscript $t + 1$, all other states and actions are at time t . An IAC determines the optimal action a for the AV based on the current state s of both the AV and the human. To find this optimal action, IACs make use of at least two prediction models: a dynamic model to predict future states (s_{t+1}^*) based on current states (s_t^*) and actions (a_t^*) (the superscript $*$ denotes it can either be used for the AV or the human), and a human driver model to predict the actions surrounding human drivers will take in response to the AV's action. Both the dynamic and human behaviour predictions are evaluated to find the optimal action for the AV, this is usually done with a reward function that incorporates aspects like safety and comfort. The validation of the human driver model is the focus of this work.

Although demonstrating a proposed controller in a simulated traffic environment is a necessary first step to show its potential, it does not provide sufficient evidence on how well the controller will generalize to real-world environments. In this work, we take the position that before implementing an IAC in vehicles to validate its behaviour in the real world, its underlying driver model should be validated on natural human driving behaviour. If the model fails to predict real-world behaviour accurately, the controller will act on false predictions which can lead to annoying or even unsafe situations. Such driver model validation can therefore provide an early indication of the IAC validity without much of the cost associated with implementing and testing it in real traffic interactions. However, driver model validation is currently not a part of the mainstream approach to IAC validation (see

e.g. [4], [5], [7]), and a principled framework for such validation is missing from the literature.

The contribution of our work lies in proposing and demonstrating a human-factors-based evaluation workflow, in order to help IAC designers in the process of selecting appropriate driver models. The proposed workflow validates driver models using empirical data obtained from *naturalistic* (real-world) traffic interactions, acknowledging two levels of driving behaviour [19]: *tactical* choices and *operational* safety margins. Tactical behaviour refers to *which* manoeuvres are executed (e.g., a lane change or car following) and operational behaviour describes *how* they are executed (e.g., in terms of safety margins). To demonstrate the potential of this workflow, we perform a case study that shows that an inverse-reinforcement-learning-based model, replicated from a model used in a previously developed IAC [4], does not generalize to real-world data. Even though we do not quantify the implications of these results for any specific IAC, they still underline the importance of using validated driver models in AV controllers.

2.2. Validating driver models for interaction-aware controllers

2.2.1. Why validate?

Part of the reason why model validation is necessary is that the simulated environments in which IACs are evaluated are not sufficient to assume safe generalization to the real world. A particular aspect of the evaluation is the human response to the AV's actions. Two approaches to generate this response are used. Some studies [5]–[12], [16]–[18] simulate human driver responses using driver models. However, many of the driver models used for this purpose are also not validated on natural human driver behaviour, which could indicate a discrepancy between the simulation and natural behaviour. Other studies [4], [13]–[15] use real-time responses of a human test subject in an abstract top-down view computer simulation, much like a video game. The gap between such abstract test environments and real-world driving is large, e.g. due to the absence of risk perception [20], motion cues, and visual looming [21]. So, again we can expect the participants' responses to differ from driver responses in real-world traffic. This means that both approaches can only provide very limited evidence for generalization of the demonstrated interactive capabilities of the IAC to the real world.

To show that the IAC's behaviour does generalize to the real world, one could propose to implement the IAC in a real vehicle and demonstrate its workings in a natural environment. However, deploying a proof-of-concept IAC in the real world might result in unsafe situations even under highly controlled conditions. This raises ethical concerns about such real-world testing. Another possibility would be to use real-time human responses and minimize the mismatch between the simulation environment and the real world, e.g. by using a high-fidelity driving simulator. However, such experiments are expensive and time-consuming, and human behaviour even in realistic driving simulators can still differ from behaviour in real traffic [20], [22]. For this reason, we advocate a complementary approach: validating the driver model on naturalistic traffic data before implementing it in an IAC. The combination of the model validation on real-world data and demonstrating the IAC's interactive capabilities in a (simplified) simulated environment provides a firm ground for the further implementation and testing of the IAC in real vehicles.

To the best of our knowledge, validation on naturalistic driving data for use in IACs has not been performed for two of the most commonly used driver models




Figure 2.2: The proposed driver model validation workflow for interaction-aware autonomous vehicle controllers. The workflow consists of three steps. In the first step, a suitable dataset is selected to perform the validation of the driver model on. From this selected dataset, specific situations are automatically extracted. The actual validation of the model takes place in the last two steps. A distinction is made based on the level of behaviour. First, the tactical behaviour is validated in step 2. This step reveals to what extent the driver model shows tactical behaviour that is consistent with human behaviour in the dataset. Behaviour inconsistent with human data, e.g. collisions, is not regarded in the final step. The third step evaluates the operational behaviour of the model based on human factors literature. This is done for every tactical behaviour separately. The final conclusion of the validation should be based on the combined results of steps 2 and 3.

proposed for IACs. These models are the intelligent driver model IDM [23] (used in [6], [7], [14] to predict driver behaviour and in [12], [15], [17] to simulate other drivers' responses) and the expected-utility-maximizing model (used e.g. in [4], [5] to predict other drivers' behaviour) that uses a reward function learned from human demonstrations with inverse reinforcement learning (IRL). Although the reward function in this model is learned from naturalistic driving data, none of the studies which proposed IACs based on an IRL-based model have validated the resulting model with respect to its ability to capture human behaviour.

2.2.2. How to validate?

We propose a three-step evaluation workflow (Figure 2.2) that incorporates important aspects of driver model validation: evaluation against naturalistic data on both the tactical and operational levels.

Step 1: Select naturalistic data

When validating a driver model for an IAC, we propose that the model is compared against human behaviour data recorded in a natural environment, i.e. a naturalistic driving dataset. There are increasingly many naturalistic datasets available, but which dataset should one choose? And once the dataset is chosen, should all the data in the dataset be used uniformly for model validation?

When selecting a naturalistic dataset, one should be aware whether the data recording was done with obtrusive or unobtrusive methods. Obtrusive methods are methods where the driver is aware their behaviour is being recorded (e.g., the SHRP2 dataset [24]). As a result, the driver might have changed their behaviour e.g., to conform to the expectations of the researchers. Other datasets are gathered *without* the drivers knowing that their behaviour is being recorded, typically with drones and cameras (several open-access datasets are available e.g., [25]–[27]). Because of the possibility of adapted behaviour in obtrusive naturalistic datasets, unobtrusive datasets are preferable for model validation.

When a suitable dataset is chosen, specific parts of the data need to be selected to perform the validation on. Data recorded in the real world often contains many different scenarios, e.g. different locations, vehicle types, and manoeuvres. Using all this data to validate a driver model would be intractable because humans behave differently in different scenarios. Instead, comparable scenarios can be selected from the dataset to be evaluated together. These scenarios should fit the intended environment of the IAC. At the same time, one should avoid hand-picking scenarios, or selecting them on low-level characteristics (e.g. only include

vehicles that reach a certain velocity) because this will reduce the variability in the data and thus negate the purpose of the validation, to show that the model generalizes to real-world behaviours. Instead, scenarios should be selected on higher-level similarities, e.g. include *all* lane changes or *all* unprotected left turns. Open-source software is available that includes examples of how to extract such scenarios automatically e.g., [28].

Behaviour validation

After selecting relevant scenarios, the model can be trained and validated. Validation of models of human behaviour is often difficult because there are many aspects that determine if the model's behaviour resembles human behaviour. In most cases, the difference cannot be captured by a single metric. For example: when validating a driver model in a lane-changing scenario, it could be tempting to use a distance-based error-metric to describe the goodness-of-fit. However, an event like a collision with a vehicle in an adjacent lane can, in some cases, be described by a small lateral distance error with respect to a human-driven trajectory. If only this distance error would be examined when validating the model, it would seem to perform well, but in reality, the model predicts that a human would collide with another vehicle. The collision is missed in the single-metric validation procedure, and the (wrong) conclusion would be that the model describes human behaviour with only a small error margin.

This example illustrates that a distinction should be made between *what* behaviour is executed (e.g., car following, crashing, or lane changing) and *how* it is executed (i.e. specific trajectories and safety margins with respect to lane boundaries and traffic participants). This bears resemblance to the common distinction in driving behaviour [19] of tactical and operational behaviour (note that strategic behaviour, e.g. route selection, is not covered by the models in IACs). In this distinction, the manoeuvres executed by the driver, like a lane change, are tactical behaviour. The manner in which they are executed, e.g., expressed in accelerations or dynamics of the gaps with respect to other vehicles, is called operational behaviour. Making this distinction in driver model validation is especially relevant for driver models used in IACs because these models are mostly designed to incorporate multiple tactical behaviours. This is in contrast to traditional driver models that were more often designed to only represent one specific tactical behaviour.

For many tactical behaviours, the corresponding operational behaviour has been studied in human-factors experiments (e.g., for car following [29]–[34]). These studies provide the important metrics of human operational behaviour, given a specific tactical behaviour. Making the same distinction during the validation allows one to leverage the existing human factors literature, enabling researchers without in-depth human-factors expertise to validate their models.

Determining what tactical behaviour is executed by the model and if it matches human behaviour is something that can be done without any expert knowledge. For instance, it is straightforward to specify if a lane change is made, and to compare if the model performs a lane change in the same situation where a human does. Once the tactical behaviour is determined, the metrics specifying the operational behaviour can be defined based on the relevant human-factors literature. This will require obtaining some knowledge on the subject, but with a properly specified tactical behaviour, a brief, non-exhaustive driver-behaviour literature survey would be enough for a researcher to make a motivated choice of the metrics characterizing the corresponding operational behaviour.

Because making a distinction between tactical and operational behaviour is relevant for IACs and makes the validation process easier, we propose a sequential two-stage validation process. The first stage (step 2 in the workflow of Figure 2.2) is to validate the model's behaviour on a tactical level, providing a quick and straightforward distinction between behaviour that clearly resembles or does not resemble the observed human driving behaviour in the same circumstances. The second stage (step 3 in Figure 2.2) examines the tactical behaviours separately on the operational level.

Step 2: Tactical validation

The purpose of the tactical validation step is two-fold. First, it serves to determine which of the model's responses are consistent with human behaviour and which are not. A valid driver model does not predict tactical responses inconsistent with human behaviour, therefore we will refer to such responses as undesirable tactical behaviour. Desirable behaviours on the other hand, are all tactical responses that can be observed in naturalistic human driving data. Second, this step will categorize the model's responses so its desirable behaviours can be validated in the operational validation step according to the right criteria. Undesirable behaviour can be disregarded during the operational validation step because it does not matter how the model performs a behaviour that is undesirable in the first place. To achieve this, a mutually exclusive set of possible tactical behaviours exhibited by the model should be defined. The distinction between these tactical behaviours should be based on simple rules (or inclusion and exclusion criteria) such that all exhibited model behaviour falls in one and only one tactical category. Which and how many of these categories to include depends on the outcome of the literature survey discussed earlier. All behaviours in one category should be validated on the same operational characteristics, which should be taken into account when determining the categories.

Step 3: Operational validation

For the operational validation step, human-factors literature provides signals and metrics that best describe human behaviour for specific tactical behaviour. This operational validation step can compare individual trajectories or averaged metrics between human and model behaviour as long as the metrics and signals are chosen appropriately and the tactical behaviours are regarded separately. Examples of such metrics are metrics that relate to the dynamics of the behaviour, e.g. the gap between vehicles, or to the properties of the manoeuvre, e.g. the duration of a lane change. Human-factors literature can also provide methods on how to compare the signals and metrics. For example, in [29] figures are presented that relate phase diagrams in car following to responsive actions of human drivers, such plotting methods can also be used for model validation.

The validation conclusion

The final conclusion of the validation procedure should be based on both the tactical and operational behaviour displayed by the model. The model should display desirable tactical behaviour in a way that resembles how humans perform the same behaviour on an operational level. But because the eventual goal is to incorporate the driver model in an IAC, the controller's ability to safely operate while using the model's predictions can be seen as the most important factor in the final conclusion.

When a driver model shows behaviour that deviates from human behaviour to a large extent, but the controller that implements the model can still safely operate with these errors, it can still be concluded that the model is "good enough" for use in the IAC. To draw such a conclusion, the maximal acceptable difference between the model's output and human behaviour has to be defined. This should be done for every IAC separately due to differences in IACs, scenarios, and regarded tactical behaviours. The maximal acceptable difference can for example be based on an evaluation that shows that the controller can still reliably execute safe and acceptable interactive behaviour when confronted with predictions that have this maximal deviation from future human behaviour.

However, even if an IAC is robust to inaccurate predictions of the driver model, we argue that it is still important to validate the model and report the magnitude of the deviation from human behaviour. This improves the re-usability of the proposed model for other IACs and provides a basis for a re-evaluation of the model when extending or improving the IAC.

2.3. Case study: Methods

To demonstrate the proposed workflow we use it to validate an inverse reinforcement learning (IRL) based model replicated from a study that proposed one of the first IACs for autonomous vehicles [4]. The choice to validate an IRL-based model was made because this increasingly popular type of model describes dynamic human behaviour in multiple scenarios and has not been validated previously. The two IACs with IRL-based driver models discussed earlier [4], [5] use similar implementations of such a model. However, only the work by Sadigh et al. [4] provides enough detail, in the form of mathematical description and open-source code, to replicate the used IRL-based model. For that reason, the model used by Sadigh et al. is used as a reference for this case study.

2.3.1. Model implementation

IRL-based driver models assume that human behaviour is "driven" by an underlying reward function. A parameterized reward function is assumed and inverse reinforcement learning is used to infer the parameters directly from human demonstrations (see [35]–[37]). This reward function with the learned parameters can be used in an agent to generate individual predictions of human behaviour. Driver models based on IRL use a utility-maximizing rational agent for this purpose. Throughout this paper, we refer to this method of generating predictions combined with a specific assumed reward function as the model. We refer to instances of the model with a specific set of parameters as an agent. In IRL-based driver models, the used reward function consists of a linear combination of features, each with its own weight:

$$R^h(s, a) = \sum \theta_i^h \phi_i(s, a). \quad (2.1)$$

In this formula, R^h denotes the reward of a specific human, s is the state (at time t) and a is the action sequence the human will take. This action sequence is subject to a finite planning horizon. ϕ_i denotes the i_{th} feature and θ_i represents the corresponding weight, which is learned by IRL from demonstrations produced by a human driver h . Note that the features ϕ_i in equation (2.1) are designed beforehand and do not vary over humans, demonstrations, or situations. The weights θ_i are learned from the demonstrations and vary over humans. These weights

are learned by maximizing the log-likelihood of an observed demonstration with respect to the weights, given the assumed features.

2

2.3.2. Assumed reward function

The reward function R^h used for the IRL-based model in this work was replicated from [4] and consists of four features for: maintaining a preferred velocity, lane-keeping, staying on the road, and collision avoidance. The collision avoidance feature is modeled by a two-dimensional Gaussian function, based on distances between the centers of vehicles. Because the human demonstrations we use for the case study were recorded on highways, the heading angles of the vehicles take very low values and are therefore neglected for collision avoidance. They are assumed to be equal to the road heading (this is a deviation from the model used in [4]). The lane-keeping and road boundary features are both Gaussian functions of the lateral road axis, they are constant over the longitudinal axis of the road. The velocity feature is the squared error with respect to the desired velocity. Since the exact desired velocity is not known for the human drivers that provide the demonstrations, and the legal speed limits that could be used for this purpose are not always provided with the data, the maximum recorded velocity of a vehicle is taken as the driver's desired velocity. The full reward function is given in equation (2.2).

$$R^h(x, y, v_x) = \theta_{\text{vel}}^h \phi_{\text{vel}}(v_x) + \theta_{\text{lane}}^h \phi_{\text{lane}}(y) + \theta_{\text{bounds}}^h \phi_{\text{bounds}}(y) + \theta_{\text{collision}}^h \phi_{\text{collision}}(x, y), \quad (2.2)$$

where

$$\begin{aligned} \phi_{\text{vel}}(v_x) &= (v_x - v_d)^2 \\ \phi_{\text{lane}}(y) &= e^{-c(y_{lc} - y)^2} \\ \phi_{\text{bounds}}(y) &= e^{-c(y_{rb} - y)^2} \\ \phi_{\text{collision}}(x, y) &= \frac{1}{\sigma_x \sqrt{2\pi}} e^{-(1/2)((x - x_o)^2 / \sigma_x^2)} \frac{1}{\sigma_y \sqrt{2\pi}} e^{-(1/2)((y - y_o)^2 / \sigma_y^2)} \end{aligned}$$

In these formulae, x and y denote the longitudinal and lateral position as defined in Figure 2.4, lc and rb denote the lane center and road boundaries respectively, where the road boundaries are defined at half a lane width outside the outermost marking. v represents velocity and subscript o denotes the other vehicle. The constants c , σ_x and σ_y are used to shape the features. A visual representation of the reward function, excluding the velocity feature, can be found in Figure 2.3. The constants that shape these features were determined with a grid search on the first 15 demonstrations of the used dataset. Initial guesses of the parameters were based on a visual comparison of the heat map to the road image. Variations around these initial guesses were estimated based on dimensions of the lanes. (For example $\min(\sigma_y) = 1.4 \text{ m}$, thus 95.4% of the lateral influence on collision prevention lies within 2.8 m distance between vehicle centers. With a lane width of 4 m and a 2 m wide vehicle, this means the lane marking has to be crossed before the collision prevention starts contributing to the reward. Thus, the lower bounds of our parameter grid are close to the smallest plausible parameter values.) We used the following sets in the grid search: $c = \{0.14, 0.18, 0.22\}$, $\sigma_x = \{5.0, 10.0, 15.0, 20.0\}$, $\sigma_y = \{1.4, 1.8, 2.2\}$, where the **bold** value the selected value. Each parameter combination in the grid was evaluated based on the resulting number of desired tactical




Figure 2.3: A heat map of the reward function (equation (2.2)) used for the IRL-based driver model where the black block indicates the ego vehicle. Warmer colours indicate low reward, cooler colours indicate high reward. The feature for velocity is not shown here because it does not depend on the position. The dimensions and positions of the features displayed here are assumed to be constant over different humans. The weights, represented here by the colours, differ between humans and are learned from demonstrations.

behaviours by the agent (see Section 2.2.2, Step 2 for the definition of desired behaviour). The parameter sets $c, \sigma_x, \sigma_y = 0.14, 15.0, 1.4$ and $c, \sigma_x, \sigma_y = 0.14, 20.0, 1.4$ had the maximum number of desired tactical behaviours in this grid search, we chose to select the combination containing our initial guess.

2.3.3. Using the proposed workflow

Here we will discuss the use of the proposed workflow (Figure 2.2) to validate the IRL-based model with the reward function as shown in equation (2.2) step by step.

Step 1: select data

The first step of the proposed workflow is to select a naturalistic dataset. Among multiple naturalistic driving datasets that are openly available, in this case study we considered three datasets: the NGSIM dataset [26], the pNEUMA dataset [27], and the HighD dataset [25]. Of these three the NGSIM data has larger uncertainties in the trajectories because it was recorded with fixed-base cameras instead of drones. The pNEUMA dataset was recorded in an urban environment, this does not match the environment of the regarded IAC [4] which focuses on multi-lane scenarios (e.g. a highway) with human behaviour that mostly consists of actions to prevent collisions, like lane changing. The HighD dataset contains high-precision data recorded in a multi-lane environment. It also contains dynamic behaviour such as lane changes to prevent collisions. For these reasons, we will use the HighD dataset. This dataset consists of 60 separate recordings, recorded in 6 different locations in Germany. All recordings were made on highways using drones equipped with cameras; from these recordings, trajectory data was automatically extracted [25]. Every recording is of a fixed stretch of highway, the average length of these recorded stretches is 416 m, the average duration of a single-vehicle track is 14.34 s.

To visualize the data and the resulting agent behaviour we used TraViA [28], an open-source visualization and annotation tool for trajectory datasets. TraViA can visualize all mentioned datasets and we extended it to train and visualize the IRL-based model. The extension code is available online [38]. An example frame of the HighD dataset visualization can be found in Figure 2.4.

From this dataset, we automatically select suitable scenarios for training and validating the model. These scenarios should fit the intended use of the IAC [4]: in our case study, we assume the goal of the IAC is to interact with human drivers who perform lane changes. This means we could consider two distinct behaviours in the HighD dataset: lane changing and merging. A merging lane is only present in 3 of the 60 HighD recordings. For this reason, we will use human lane-changing



Figure 2.4: An example frame of the HighD dataset [25] as visualized using TraViA [28]. The frame includes a stretch of a highway in Germany, where vehicles drive on the right side of the road and where, in some of the cases, there are no legal speed limits. The orange shapes represent regular cars, the green shapes are trucks. All vehicles have a vehicle-ID shown in white. The white arrows display the coordinate frame and the yellow marking shows a visualization of the gap between two vehicles as used in the metrics for step 3 of the validation workflow.

manoeuvres for validation. For consistency, the three recordings with a merging lane were not considered.

As mentioned before, the features in the reward function consider collision avoidance, lane-keeping, staying on the road, and maintaining a preferred velocity. This means that not all lane changes can be explained with this model. Lane changes to the right are not covered because they are "driven" by a need to adhere to (socially acceptable) traffic rules that are not incorporated in the reward function (In Germany, it is obligatory to drive in the rightmost lane if it is free. So a lane change to the right is most often performed simply because that lane is free, not to avoid a collision. It can therefore not be explained by the used reward function). Therefore, only single lane changes to a left lane are considered for training and validation. The highD dataset includes the number of lane changes for every trajectory (based on lane crossings) and the current lane number at every frame. We automatically extracted all used trajectories based on these metrics.

Step 2: tactical validation

The next step is to define a set of tactical behaviour categories. There are only a limited number of possible tactical behaviours on a highway without an exit lane, we will consider four possibilities: car following, lane changing, colliding, and crossing the road boundaries. Lane-keeping is not regarded as a separate behaviour since all vehicles on a highway essentially follow another vehicle. In this set car following and lane changing are regarded as desirable behaviours, and colliding and going off-road are considered undesirable.

Besides defining the behaviour categories, we established a procedure to place the trajectories produced by the agent in one of these categories based on a hierarchy in tactical behaviours. First, if an agent collided with another vehicle, this is labeled as "collision". If the agent did not collide, a check is done to see if the center of the vehicle stayed within the outer road boundaries; if not, the tactical behaviour is labeled "off-road". Agents that did not fall in one of the two categories above are checked for lane changes; if there is one, the tactical behaviour is labeled "lane change". And finally, agents that showed none of these three behaviours are placed in the "car following" category. All of these checks are implemented in the software and are performed automatically for all agents by checking for overlap with other vehicles and evaluating the vehicle's center position for every time step.

The used hierarchy is based on the idea that a predicted collision has the highest impact for IACs. If a model predicts a collision, an IAC will act to avoid this, independent of the fact that the model predicts a lane change first. Vehicles leaving

the road will also have a big impact on IAC behaviour because it reduces the number of vehicles to consider and thus changes the scene. However, the IAC will not take drastic actions to avoid this, therefore it comes second in the hierarchy. Only if none of these undesirable behaviours are executed by the model, lane changes are relevant. Finally, all other behaviours within a single lane is grouped as car following. A more fine-grained distinction could have been made here by including behaviours such as nudging or aborted lane changes. But before considering those more sophisticated behaviours, we chose to evaluate if and how the model displays car following in general.

To evaluate if the model's tactical performance is adequate for use in an IAC, a maximum acceptable deviation from human behaviour needs to be specified. Because no IAC implementation is used in this case study, we cannot specify such a threshold here.

Step 3: operational validation

The last step is to determine how to evaluate the model's operational behaviour for the cases where the tactical behaviour falls in one of the desirable categories. We have defined two desirable categories: lane changes and car following. Earlier studies investigated human car-following behaviour and risk perception using inverse time-to-collision vs time gap plots [34], [39], these metrics were also used to evaluate human lane changes before [40].

Time-to-collision is defined as the time it will take until a vehicle collides with the preceding vehicle given that they both continue at their current velocity,

$$\text{TTC} = \frac{x_{\text{gap}}}{v_{\text{rel}}}. \quad (2.3)$$

The time gap is the time it will take a vehicle to close the current gap with the preceding vehicle, given its current velocity,

$$t_{\text{gap}} = \frac{x_{\text{gap}}}{v_{\text{agent}}}. \quad (2.4)$$

In these equations, TTC is time-to-collision, v_{rel} is the relative velocity of the agent and the preceding vehicles, and x_{gap} is the distance gap between the vehicles. This distance gap is visualized in Figure 2.4. Finally, v_{agent} is the longitudinal velocity of the agent vehicle.

Both TTC and time gap are available in the HighD dataset for human behaviour; for the agent behaviour, the metrics are calculated using the equations (2.3) & (2.4). Again, quantifying an acceptable error margin can only be done for a specific controller. Because we don't demonstrate a controller, we can only show the difference between the model and human behaviour, but in this case study, we cannot quantify if this is acceptable for any specific IAC.

2.3.4. Model training

The optimization procedure to find the weights that fit a human demonstration best is the same as used by Sadigh et al. [4]. The negated log-likelihood function as proposed by [41] is minimized with respect to the weights. To keep this tractable, the human demonstration is divided into sections with the same number of frames as the control horizon used in the agent ($N = 5$). All data frames are used, so the time step is $\frac{1}{25}$ s and the planning horizon is $\frac{1}{5}$ s. The log-likelihood functions of the parts of the demonstration are summed and the summed negated log-likelihood

is minimized. We assume that every lane-change trajectory in the dataset comes from a different human, an agent is trained separately for every trajectory, this resulted in 3279 trained agents. Demonstrations on which the optimization procedure fails (i.e., no minimum of the negated log-likelihood function could be found) were discarded (2302 demonstrations, 41%).

Because highway data is used, the velocities of the vehicles are high (mean = $29.7m/s$) and heading angles are small. The heading angles of the vehicles are ignored in the dataset. For this reason, the dynamics of the vehicles are modeled as point masses. Because the trajectories are extracted from videos, no direct acceleration data was recorded. Acceleration data is available from the HighD dataset, but this has been reconstructed from velocity data. For this reason, the humans in the demonstrations are assumed to have direct control over the longitudinal and lateral velocities. Making the state and action vectors both 2-dimensional containing respectively an x, y -position and -velocity. This assumption is justified because the goal of the model is to learn the reward function, not the dynamics of human control.

2.3.5. Validation of agent behaviour

To validate the agent's behaviour, we evaluate the response of every agent individually in the same scenario that was used to train the agent. A dedicated test-set is not required contrary to most machine learning approaches because the log-likelihood optimization proposed by Levine and Koltun accounts for sub-optimal demonstrations by humans. This means that the learned reward function does not need to be fully optimized in the human-driven demonstration, but the agent will fully optimize the reward function. So the agent might display different behaviour than the human in the same situation and thus this situation can be re-used for validation.

For evaluation, the agent will be placed in the same initial position and its behaviour is recorded for the same duration as the demonstration trajectory. Because the agent learned its reward function from this exact situation, this is a best-case scenario for the model. This approach also has the advantage that we can directly compare the agent's behaviour to the human demonstration it was trained on.

As for the IRL training, heading angles are neglected and the dynamics of the vehicle are assumed to be point mass dynamics. To approximate the states and actions of real drivers, the agents are assumed to have direct control over the linear accelerations of the vehicle. This results in a 4-dimensional state vector per vehicle, containing both x, y -position and -velocity, and a 2-dimensional action vector containing the x, y - accelerations. The agent is a utility-maximizing rational agent, so it will select an action a in state s , that maximizes its summed reward function R over a time horizon $N = 5$. Again, the time-step is equal to the frame rate of the HighD dataset ($\frac{1}{25}$ s). The agent has full knowledge about the future trajectories of all adjacent vehicles. As for the choice of situation, this can be regarded as a best-case scenario for the agent, since it has a perfect prediction system to predict other human behaviour.

The direct control over lateral accelerations, combined with the point mass dynamics, can result in trajectories that are not subject to normal vehicle dynamic constraints. To approximate normal vehicle dynamics, the agent's actions (x, y - accelerations) are constrained to the maximal values of these

accelerations found in the HighD dataset. The x -acceleration is constrained between $(-6.63, 20.06) \text{ m/s}^2$ and the y -acceleration between $(-1.63, 1.63) \text{ m/s}^2$.

2.4. Case study: Results

From the first 57 recordings in the HighD dataset, all 5581 single lane changes to the left lane were automatically detected. These lane changes served as human demonstrations for the IRL-based driver model. Out of these 5581 demonstrations, 3279 resulted in a set of weights after the inverse reinforcement learning procedure. For the other 2302 demonstrations, the IRL procedure failed to converge.

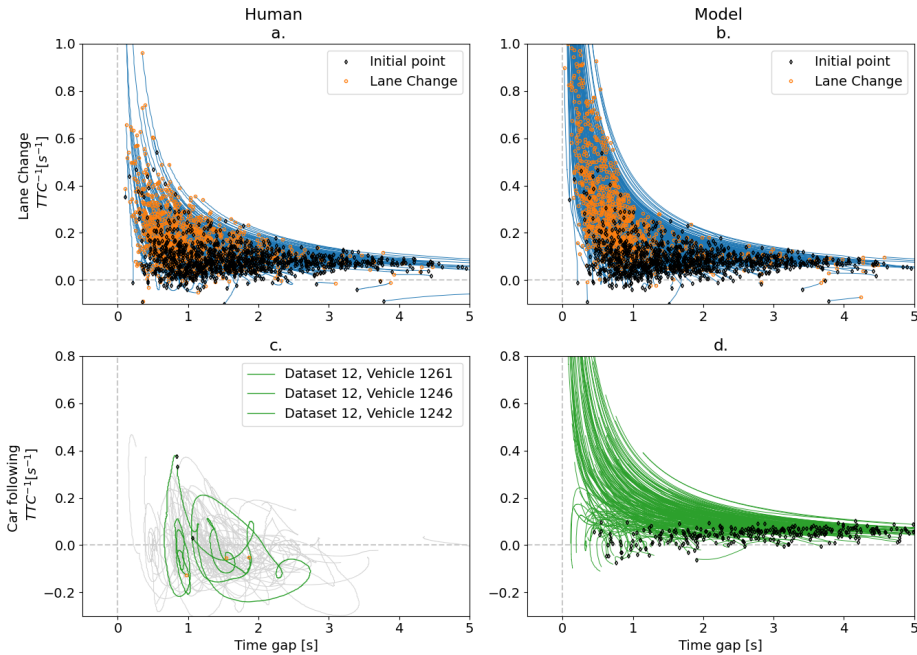


Figure 2.5: Inverse TTC vs time gap plots of human demonstrations (panels a and c) and IRL agent behaviour (panels b and d) in lane changes and car following. Panel a shows human behaviour in the used demonstrations. Since these demonstrations do not contain any car following, panel c shows 55 illustrative examples of car-following behaviour selected from other trajectories in the dataset, three are highlighted for clarity. Black diamonds indicate the initial position, this is the first frame in which a vehicle appears in the HighD dataset. In panels a and b, orange dots indicate a lane change, corresponding to the frame in which the center of a vehicle crosses the center-line between lanes. From panels a and b we conclude that the model's lane change behaviour has human-like dynamics in general, however, the model makes lane changes at substantially higher inverse-TTC (lower TTC) compared to humans. From panels c and d we conclude that the model's car-following behaviour does not resemble human car-following behaviour.

In practice, the failure of the IRL procedure means that the likelihood function adopted from [41] becomes intractable. This function contains the logarithm of the determinant of the Hessian matrix ($\log |\mathbf{H}|$). When this determinant becomes negative, the optimization fails. We found that this can happen when the optimization algorithm assigns a positive value to θ_{vel}^h (i.e., when deviating from the desired velocity is rewarded instead of punished). Note that weights are not re-

stricted to be either positive or negative by the IRL procedure. IRL learns if a feature represents a reward or penalty for the human demonstration.

To examine if this was the cause of the high rate of failures in our training procedure, we estimated the Jacobian used in the optimization procedure for the initial values of θ . For 97% of the demonstrations where IRL failed, the Jacobian value for θ_{vel}^h was negative and had a magnitude at least 10 times larger than all other Jacobian values. This did not happen in demonstrations where IRL succeeded (0 out of 100 randomly selected cases). Which indicates that the optimization algorithm attempted to use positive weights for θ_{vel}^h as they have a high likelihood to explain the demonstration in the failed cases.

This could mean that features in the reward function that are based on the deviation from a maximum observed (or allowed) velocity are not suitable for use on real-world traffic conditions. On the other hand, the IRL procedure might not have failed in these cases if θ_{vel}^h was restricted to always be negative (or more generally, if weights are restricted to represent either rewards or penalties). Further investigation to answer these questions is left for future work. We discarded the demonstrations where training failed and continued the attempt to validate the IRL-based driver model using the training data for which the model converged.

The 3279 agents that trained successfully were placed in the same scenario they were trained on to examine to what extent they show human-like behaviour on a tactical and operational level. We would like to remind the reader that, combined with the fact that all agents had access to perfect predictions of all surrounding vehicles, this constituted a 'best-case scenario' for the model.

	number of agents	percentage of agents	percentage of humans
Lane-change	1318	40.2%	100.0%
Collision	875	26.7%	0.0%
Car following	593	18.1%	0.0%
Off-road	493	15.0%	0.0%
Total	3279	100%	100%

Table 2.1: Tactical behaviour as shown by the IRL agents and in the human-driven demonstrations the agents were trained on.

Tactical behaviour

On a tactical level, we have defined four possible behaviours to categorize the resulting agent behaviour: car following, lane changing, colliding, and crossing the road boundaries. Only in 40.2% of the cases the model showed the same tactical behaviour as the human demonstration, a lane change (Table 2.1). In more than 41% of the cases, the model either collided or went on an off-road adventure. This behaviour was not present in the chosen subset of the human data, so we conclude that model behaviour is inconsistent with human behaviour.

Operational behaviour

We then compared the operational behaviour of the model to the operational behaviour in the human demonstrations using the inverse time to collision vs time gap plots (Figure 2.5). Trajectories with multiple preceding vehicles show jumps in these plots due to suddenly changing values, for that reason those trajectories

were omitted. Agents and humans that perform a lane change when the preceding vehicle is out of sight are also omitted since no inverse TTC and time gap data can be calculated for them for the final frames. All car-following trajectories are cropped to the point where the preceding vehicle gets out of sight.

The plots on the left side of Figure 2.5 (a & c) show human operational driving behaviour. In the case of lane changing (2.5-a.), the inverse TTC increases while the time gap decreases, until the point where the center lane-marking is crossed, depicted with an orange circle. In the case of car-following (2.5-c), humans oscillate around a preferred equilibrium point.

The model's behaviour for the same manoeuvres can be seen on the right side of Figure 2.5. The model's lane-changing behaviour (2.5-b) has human-like dynamics in general (as in 2.5-a), however, the model makes lane changes at substantially higher inverse TTC (lower TTC) compared to humans. Also, the time gap at the moment of the lane change is on average smaller than for the human demonstrations. To further illustrate the differences in the lane change dynamics, we investigated the distributions of inverse TTC and time gap at the moment of lane change (Figure 2.6). This shows substantial differences between the estimated distributions. We performed a paired t-test to check for significant differences, both the inverse TTC ($t(1075) = -7.61$, $p = 6.1e-14 < 0.001$, Cohen $d=0.302$) and time gap ($t(1075) = 13.49$, $p = 2.0e-38 < 0.001$, Cohen $d=0.234$) values at the moment of lane change differ significantly between the model and human demonstrations. So for lane-changing behaviour, we conclude that the IRL-based model does not resemble human behaviour on an operational level.

When comparing the agent's car-following behaviour (2.5-d) with the human's car-following behaviour (2.5-c), there are no oscillations around an equilibrium point for most agents. The general shape resembles that of a human lane-changing manoeuvre (2.5-a) without crossing the center lane-marking. From this, we conclude that if the model shows car-following behaviour, it does not do that in a way that resembles human oscillatory car-following behaviour but instead it tailgates the preceding vehicle.




Figure 2.6: Estimated distributions of inverse time to collision and time gap at the moment of the lane change. The orange distributions represent the model's behaviour and the blue distributions represent the human demonstrations. The mean values for inverse TTC are 0.19 s^{-1} for human lane changes and 0.47 s^{-1} for the model. The mean values for time gap are 1.05 s for human behaviour and 0.85 s for the model behaviour.

Reason for the agents' behaviour

Why do the IRL-based agents show behaviour that is so different from human behaviour, even though their reward function was learned from human demonstrations? We randomly selected several agent trajectories for manual examination using the TraViA [28] traffic visualization tool to answer this question. Examples of these trajectories can also be found as videos in the supplementary materials. From these manual evaluations, two main causes were identified that explain why the behaviour of the agents does not represent human behaviour: the model's assumptions and the IRL fitting procedure.

To start with the cases where the model's assumptions cannot explain the desired behaviour. Consider a demonstration where a human is merging in a slow-moving and crowded left lane to overtake a truck farther ahead in the right lane. This might be beneficial in the long run because the truck can be overtaken, but such behaviour is unlikely to be beneficial within the short planning horizon of the model, especially because the distance-based collision features promote staying away from other vehicles. This issue is similar to the previously identified problem that lane changes to a right lane cannot be explained by the currently assumed reward function. In both of these cases, no matter the learned weights, the assumed reward function will not lead to the desired behaviour within the planning horizon. In other cases, the approach of learning the weights from a demonstration using an assumed reward function can be identified as the cause of the problem. Many agents that collided learned their weights from a demonstration where the human moves into the area influenced by the collision feature (see Figure 2.7 for an example). Because the dimensions of this collision feature are fixed and only the weights are learned in the IRL procedure, the resulting collision weight will be low, i.e. a low collision weight is the only way to explain the human moving into this area. When this low collision weight is used in the agent to generate behaviour, the agent will not perform a lane change, because moving into the collision-feature area will always decrease the reward. Instead, the agent will stay in its lane. When it approaches the preceding vehicle, it will collide, because of the low collision-prevention weight.

The underlying problem here is that the assumed reward function cannot describe the human's demonstrated behaviour properly. Suppose using such a flawed reward function with hand-picked weights. In that case, one would expect prediction errors on the operational level, because the timing of the lane change is determined by the distance-based collision feature. In this case however, the IRL procedure exaggerates the effects of the flawed reward function by learning weights that result in more collisions. So even though the problem lies in the flawed reward function and not the IRL procedure itself, the combination of the IRL-procedure and reward function might not only limit the performance of the model, it can actively make it worse.

2.5. Discussion

In this work, we have proposed a validation workflow for driver models in interaction-aware AV controllers. We illustrated its utility through a case study of validating an inverse reinforcement learning-based driver model replicated from literature [4] using naturalistic highway driving data extracted from the HighD dataset [25]. Our validation workflow (Figure 2.2) incorporates the automatic extraction of comparable lane change scenarios (5581) on which the IRL model was trained (step 1). The validation of the model was then performed in two




Figure 2.7: An example of a demonstration where the assumed anti-collision feature does not describe the human’s demonstrated behaviour. In this figure, the black shape represents the position of the human-driven demonstration vehicle during a lane change manoeuvre, the white arrow indicates its direction of motion. Only the collision feature is visualized with warmer colours indicating higher cost. In this example, the demonstrating vehicle is moving from a low-cost area (right lane) to a high-cost area (left lane). The only way to explain this behaviour with the assumed features is to assign a low weight to the collision feature. Apparently, the demonstrating human does not care so much about moving into the higher cost area in the left lane, other features must be more important. When these learned weights are then used in a utility-maximizing agent in the same situation, it will not make the demonstrated lane change. Instead, it will stay in the right lane with a lower penalty and finally collide with the preceding vehicle (969), because collision prevention has a low weight.

related stages. First, we examined the tactical behaviour of the model (step 2). Even though no collisions or off-road driving were present in the training data, the model produced such behaviour in more than 41% of the cases (Table 2.1). Second, we analyzed the operational behaviour of the model in the 59% remaining trajectories (step 3). This analysis revealed that even though the dynamics of the model’s lane changes are similar to humans (Figure 2.5a,b), the model performed the lane changes with significantly smaller safety margins (Figure 2.6). Furthermore, the dynamics of the model’s car following behaviour was largely inconsistent with human behaviour (Figure 2.5c,d).

In conclusion, despite training the IRL-based model on data of real-world driving behaviour, our 3-step evaluation workflow exposed how the model is not able to produce realistic behaviour in the same scenarios. This case study illustrated that despite promising results in simple IAC demonstrations, the models used for human behaviour prediction in IACs can deviate substantially from actual human behaviour, which can have serious ramifications for generalization of IACs to real-world environments. Our results highlight the importance of validating the models used in interaction-aware motion planning for autonomous vehicles, and suggest an easy-to-use framework to aid researchers in doing so.

Practical applicability of the validation workflow

The case study of validating an IRL-based driver model illustrated the practical applicability of the proposed validation workflow (Figure 2.2). In the first step of the workflow, the case study showed the feasibility of automatically extracting data from an open-access naturalistic dataset. Even after narrowing down the extracted data to select specific scenarios (in our case, lane changes), the data were sufficiently rich to serve as training data for the IRL model. Note that multiple other datasets were available for consideration (e.g., NGSim [26] & PNeuma [27]) to further enlarge the data and/or use scenarios other than lane changes.

The second and third steps of the workflow propose a two-stage evaluation approach, separated into *tactical* and *operational* driver behaviour. The case study illustrated why this two-stage validation is useful and necessary. On the tactical level, the large number of collisions and off-road driving would have been hard to identify in a one-stage metric-based validation (e.g. mean square-root error in [42]). On the operational level, the evaluation illustrated that the differences

between car following and lane changing in human behaviour were not reflected in the model's behaviour. This would have been impossible to identify without first examining the tactical behaviour.

The results of the case study also underline the importance of validating driver models for IACs in general. The discrepancy between the driver model and human behaviour suggests that an IAC using this model might not safely generalize to real-world scenarios. The case study shows that models that do not actually capture human behaviour are not just a hypothetical issue, but a practical concern for IACs developed for autonomous vehicles.

Implications for interaction-aware controllers

The results of the IRL-based model validation have implications for IACs that would use this model to predict other drivers' responses. Wrong predictions on the tactical level can lead to dangerous situations. If an AV decides to accelerate based on an inaccurate prediction that a vehicle in an adjacent lane will stay there, a dangerous situation might occur when the other vehicle moves in front of the AV. The same holds for inaccurate predictions on an operational level. For example, the model will close the gaps to a (very) high inverse TTC (low TTC) compared to human drivers. This can lead to over-conservative AV behaviour because the controller overestimates the aggressiveness of the human. The full extent of these implications needs to be further examined in future work.

Related work and generalizability

To the best of our knowledge, our work is the first attempt to validate a driver model used in interaction-aware controllers on both the tactical and operational levels. The work on which this model was based [4] does not use naturalistic data and reports no validation attempt of the behaviour model. Another related study [5] does use naturalistic data (NGSim) to train the IRL-based model, but also does not report any validation of the trained model. In the supplementary material (available at [42]) Schwarting et al. do report the mean squared errors of their model for merging scenarios. However, given the complexity of human behaviour in traffic interactions, such one-dimensional averaged error metrics provide only rudimentary information on how well the model captures human behaviour.

Driver model validation using naturalistic data has been performed for other use cases than IACs. In [43] five car-following models are validated for use in microscopic traffic simulations on naturalistic data collected in Shanghai. Their validation method could also be useful when designing IACs, and our validation method could as well be used to validate models developed for applications other than IACs. However, we argue that because our method includes the tactical and operational validation steps, it is more suitable to validate models displaying multiple higher-level behaviours.

Besides the IAC literature, there have been other driver modelling attempts using IRL. However, IRL-based models can differ substantially from each other in terms of the used reward function. Naumann et al. studied the suitability of different cost functions for different driving scenarios [44] and showed that there are substantial differences. The other modelling attempts that use IRL differ from our work precisely in the sense that they target another scenario (e.g. [45] who regards curve negotiation) or use different reward function features (e.g. [46] who uses velocity-based features for risk perception). That means that the results of those works should be regarded as validations of different models, despite the fact that they are also based on IRL. This observation leads to two conclusions. First, other

models that use different features should be validated as new models even when they also use IRL. Second, the choice of features for the reward function impacts the performance of the model, which might provide an opportunity to improve models that underperform.

The reasons why the IRL-based models perform poorly in our case study will most likely generalize to other IRL-based models that use similar distance-based collision features in the reward function. The results show that such distance-based features do not capture the essence of human driving behaviour. Only changing the shape or dimension of a position-based feature will not solve this. Instead, we advocate that the metrics used in the reward function features should be based on human factors literature for the targeted tactical behaviours, as was done in the operational validation of the model; e.g. the distance-based collision feature could be replaced with a TTC-based feature (similar to the previously mentioned model in [46]).

Other validation attempts of human driver models that do not specifically target IACs and do not use IRL also exist, one especially related to our work is [47]. In that work, Srinivasan et al. compare the trajectories generated with a deep-learning-based model to naturalistic driving data. As in our work, the comparison is based on an in-depth analysis of the resulting trajectories instead of one-dimensional metrics. They show that, also for deep-learning-based driver models, validation should be grounded on a low-level comparison of trajectories, not just high-level metrics. They do however not provide a generalized framework for performing such validations as we do with our proposed workflow.

Limitations and recommendations

This work has three main limitations. First, we used only a single demonstration of a lane change to train the IRL, which might explain part of the discrepancy between human and model behaviour in the results. However, providing the system with more training data might only slightly improve the model's performance. In the case study we identified the causes of the observed problems to be the features of the reward function, not the weights. Adding more training data could result in weights that better fit a specific driver. But it will not negate the problem with the features used in the reward function.

Second, it should be noted that the planning horizon of the model is very short due to the combination of a low number of frames and a high frame rate ($N = 5$ at $\frac{1}{25}$ Hz). The number of frames within the planning horizon was chosen based on the previous work [4] and to keep the IRL procedure tractable. The frame rate was directly adopted from the HighD dataset for simplicity, both for reproduction purposes and to not introduce any extra assumptions when down-sampling the data. Increasing the planning horizon and examining the model's behaviour under those conditions is left for future work.

Finally, our case study only attempts to validate the model, it does not quantify the implications of the outcome for use of the model in an IAC. Therefore, we are unable to say which aspects of the model's behaviour would be tolerable when used in an IAC or which aspects have major consequences. Quantifying the implications of the mismatch between the model's, and naturalistic human behaviour is left for future work. Answering such a question is an interesting topic of research on its own, a perspective on how to approach such an evaluation can be found in [48].

Future work should also focus on validating more driver models for use in AV controllers, e.g. the Intelligent Driver Model [23] mentioned in the introduction is used in many simulations and demonstrations to model individual human behaviour for IACs and should be validated for such use. Future work on IRL-based driver models could focus on redesigning the used reward function such that it better captures similarities between human drivers by using human-factors literature as a starting point. Besides that, the IRL-based model used here could be extended to take the uncertainty in human behaviour into account. Either the uncertainty over the learned rewards could be targeted by learning multiple reward functions (as is done in [44], [49]) instead of only single parameters and selecting the best fit, or stochasticity could be added when selecting the actions to relax the assumption of humans being utility maximizers (also done in [49]). However, such changes to the model could complicate the implementation in an IAC.

2.6. Conclusions

In this paper, we argued for validation of the driver models used in interaction-aware controllers. We proposed an evaluation workflow for such validation, illustrated through a concrete case study. Based on the findings in our paper, we conclude the following:

- The proposed workflow allowed for a detailed evaluation of a driver model replicated from literature, based on an open-source dataset from which 3279 human-driven lane changes in moderately heavy highway traffic could be extracted. After training the model on each lane change, it did not reproduce adequate behaviour when exposed to the same conditions. It generated crashes and road departures in 41.7% of the cases (inadequate tactical behaviour). For the remaining cases, unrealistic safety margins were observed (inadequate operational behaviour). These unrealistic predictions show that models that do not capture realistic human behaviour are a practical concern for implementing IACs in future autonomous vehicles.
- During the case study, the proposed workflow proved to be practically applicable, providing a structured basis for model validation in two stages:
 - First, validating the tactical behaviour illustrated to what extent high-level choices are correctly predicted (e.g., that a lane change occurs, rather than staying behind the lead vehicle see Table 2.1).
 - Second, correct tactical behaviours produced by a model should be validated in additional detail, by evaluating to what extent the behaviour is executed in a way that resembles the timing and spatio-temporal safety margins acceptable to human drivers (see Figure 2.5).
 - In these two stages, different tactical behaviours should be evaluated based on different operational criteria because differences in human operational behaviour were observed for different tactical behaviours (see Figure 2.5).

Acknowledgements

The authors thank Daniel Góngora for his valuable feedback and contributions to the article.

Bibliography

- [1] C. D. Harper, C. T. Hendrickson, S. Mangones, and C. Samaras, "Estimating potential increases in travel with autonomous vehicles for the non-driving, elderly and people with travel-restrictive medical conditions", *Transportation Research Part C: Emerging Technologies*, vol. 72, pp. 1–9, Nov. 2016, ISSN: 0968090X. DOI: [10.1016/j.trc.2016.09.003](https://doi.org/10.1016/j.trc.2016.09.003). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2016.09.003> [20https://linkinghub.elsevier.com/retrieve/pii/S0968090X16301590](https://linkinghub.elsevier.com/retrieve/pii/S0968090X16301590).
- [2] L. M. Clements and K. M. Kockelman, "Economic Effects of Automated Vehicles", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2606, no. 1, pp. 106–114, Jan. 2017, ISSN: 0361-1981. DOI: [10.3141/2606-14](https://doi.org/10.3141/2606-14). [Online]. Available: <http://journals.sagepub.com/doi/10.3141/2606-14>.
- [3] S. Pettigrew, "Why public health should embrace the autonomous car", *Australian and New Zealand Journal of Public Health*, vol. 41, no. 1, pp. 5–7, Feb. 2017, ISSN: 13260200. DOI: [10.1111/1753-6405.12588](https://doi.org/10.1111/1753-6405.12588). [Online]. Available: <http://doi.wiley.com/10.1111/1753-6405.12588>.
- [4] D. Sadigh, N. Landolfi, S. S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state", *Autonomous Robots*, vol. 42, no. 7, pp. 1405–1426, Oct. 2018, ISSN: 0929-5593. DOI: [10.1007/s10514-018-9746-1](https://doi.org/10.1007/s10514-018-9746-1). [Online]. Available: <https://doi.org/10.1007/s10514-018-9746-1> [20http://link.springer.com/10.1007/s10514-018-9746-1](http://link.springer.com/10.1007/s10514-018-9746-1).
- [5] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [6] N. Evestedt, E. Ward, J. Folkesson, and D. Axehill, "Interaction aware trajectory planning for merge scenarios in congested traffic situations", *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, pp. 465–472, 2016. DOI: [10.1109/ITSC.2016.7795596](https://doi.org/10.1109/ITSC.2016.7795596).
- [7] E. Ward, N. Evestedt, D. Axehill, and J. Folkesson, "Probabilistic Model for Interaction Aware Planning in Merge Scenarios", *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 2, pp. 1–1, 2017, ISSN: 2379-8904. DOI: [10.1109/TIV.2017.2730588](https://doi.org/10.1109/TIV.2017.2730588). [Online]. Available: <http://ieeexplore.ieee.org/document/7987753/>.
- [8] W. Liu, S.-W. Kim, S. Pendleton, and M. H. Ang, "Situation-aware decision making for autonomous driving on urban road using online POMDP", in *2015 IEEE Intelligent Vehicles Symposium (IV)*, vol. 2015-Augus, IEEE, Jun. 2015, pp. 1126–1133, ISBN: 978-1-4673-7266-4. DOI: [10.1109/IVS.2015.7225835](https://doi.org/10.1109/IVS.2015.7225835). [Online]. Available: <http://ieeexplore.ieee.org/document/7225835/>.
- [9] F. Meng, J. Su, C. Liu, and W.-H. Chen, "Dynamic decision making in lane change: Game theory with receding horizon", in *2016 UKACC 11th International Conference on Control (CONTROL)*, IEEE, Aug. 2016, pp. 1–6, ISBN: 978-1-4673-9891-6. DOI: [10.1109/CONTROL.2016.7737643](https://doi.org/10.1109/CONTROL.2016.7737643). [Online]. Available: <http://ieeexplore.ieee.org/document/7737643/>.
- [10] D. Lenz, T. Kessler, and A. Knoll, "Tactical cooperative planning for autonomous highway driving using Monte-Carlo Tree Search", in *2016 IEEE Intelligent Vehicles Symposium (IV)*, vol. 2016-Augus, IEEE, Jun. 2016, pp. 447–453, ISBN: 978-1-5090-1821-5. DOI: [10.1109/IVS.2016.7535424](https://doi.org/10.1109/IVS.2016.7535424). [Online]. Available: <http://ieeexplore.ieee.org/document/7535424/>.
- [11] R. Tian, S. Li, N. Li, I. Kolmanovsky, A. Girard, and Y. Yildiz, "Adaptive Game-Theoretic Decision Making for Autonomous Vehicle Control at Roundabouts", in *2018 IEEE Conference on Decision and Control (CDC)*, vol. 2018-Decem, IEEE, Dec. 2018, pp. 321–326, ISBN: 978-1-5386-1395-5. DOI: [10.1109/CDC.2018.8619275](https://doi.org/10.1109/CDC.2018.8619275). arXiv: [1810.00829](https://arxiv.org/abs/1810.00829). [Online]. Available: <https://ieeexplore.ieee.org/document/8619275/>.


- [12] Q. Zhang, D. Filev, H. E. Tseng, S. Szwabowski, and R. Langari, "Addressing Mandatory Lane Change Problem with Game Theoretic Model Predictive Control and Fuzzy Markov Chain", in *2018 Annual American Control Conference (ACC)*, vol. 2018-June, IEEE, Jun. 2018, pp. 4764–4771, ISBN: 978-1-5386-5428-6. DOI: [10.23919/ACC.2018.8431530](https://doi.org/10.23919/ACC.2018.8431530). [Online]. Available: <https://ieeexplore.ieee.org/document/8431530/>.
- [13] H. Yu, H. E. Tseng, and R. Langari, "A human-like game theory-based controller for automatic lane changing", *Transportation Research Part C: Emerging Technologies*, vol. 88, no. October 2017, pp. 140–158, Mar. 2018, ISSN: 0968090X. DOI: [10.1016/j.trc.2018.01.016](https://doi.org/10.1016/j.trc.2018.01.016). [Online]. Available: <https://doi.org/10.1016/j.trc.2018.01.016%20https://linkinghub.elsevier.com/retrieve/pii/S0968090X18300494>.
- [14] C. Hubmann, J. Schulz, G. Xu, D. Althoff, and C. Stiller, "A Belief State Planner for Interactive Merge Maneuvers in Congested Traffic", in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 1617–1624, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569729](https://doi.org/10.1109/ITSC.2018.8569729). [Online]. Available: <https://ieeexplore.ieee.org/document/8569729/>.
- [15] S. Coskun, Q. Zhang, and R. Langari, "Receding Horizon Markov Game Autonomous Driving Strategy", in *2019 American Control Conference (ACC)*, vol. 2019-July, IEEE, Jul. 2019, pp. 1367–1374, ISBN: 978-1-5386-7926-5. DOI: [10.23919/ACC.2019.8815251](https://doi.org/10.23919/ACC.2019.8815251). [Online]. Available: <https://ieeexplore.ieee.org/document/8815251/>.
- [16] M. Garzón and A. Spalanzani, "Game theoretic decision making for autonomous vehicles' merge manoeuvre in high traffic scenarios", *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019*, pp. 3448–3453, 2019. DOI: [10.1109/ITSC.2019.8917314](https://doi.org/10.1109/ITSC.2019.8917314).
- [17] D. Isele, "Interactive Decision Making for Autonomous Vehicles in Dense Traffic", in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, IEEE, Oct. 2019, pp. 3981–3986, ISBN: 978-1-5386-7024-8. DOI: [10.1109/ITSC.2019.8916982](https://doi.org/10.1109/ITSC.2019.8916982). arXiv: [1909.12914](https://arxiv.org/abs/1909.12914). [Online]. Available: <https://ieeexplore.ieee.org/document/8916982/>.
- [18] P. Hang, C. Lv, Y. Xing, C. Huang, and Z. Hu, "Human-like decision making for autonomous driving: A noncooperative game theoretic approach", *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 4, pp. 2076–2087, 2021. DOI: [10.1109/TITS.2020.3036984](https://doi.org/10.1109/TITS.2020.3036984).
- [19] J. A. Michon, "A Critical View of Driver Behavior Models: What Do We Know, What Should We Do?", in *Human Behavior and Traffic Safety*, Boston, MA: Springer US, 1985, pp. 485–524, ISBN: 0306422255. DOI: [10.1007/978-1-4613-2173-6_19](https://doi.org/10.1007/978-1-4613-2173-6_19). [Online]. Available: http://link.springer.com/10.1007/978-1-4613-2173-6_19.
- [20] T. A. Ranney, "Psychological fidelity: Perception of risk", in *Handbook of Driving Simulation for Engineering, Medicine, and Psychology*, D. L. Fisher, M. Rizzo, J. Caird, and J. D. Lee, Eds., 2011, ch. 9, pp. 9–1–9–13, ISBN: 9781420061000.
- [21] D. N. Lee, "A Theory of Visual Control of Braking Based on Information about Time-to-Collision", *Perception*, vol. 5, no. 4, pp. 437–459, Dec. 1976, ISSN: 0301-0066. DOI: [10.1068/p050437](https://doi.org/10.1068/p050437). [Online]. Available: <http://journals.sagepub.com/doi/10.1068/p050437>.
- [22] J. Greenberg and M. Blommer, "Physical Fidelity of Driving Simulators", in *Handbook of Driving Simulation for Engineering, Medicine, and Psychology*, 2011, CRC Press, Apr. 2011, pp. 7–1–7–24, ISBN: 9781420061017. DOI: [10.1201/b10836-8](https://doi.org/10.1201/b10836-8). [Online]. Available: <http://www.crcnetbase.com/doi/10.1201/b10836-8>.
- [23] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, vol. 62, no. 2, pp. 1805–1824, 2000, ISSN: 1063651X. DOI: [10.1103/PhysRevE.62.1805](https://doi.org/10.1103/PhysRevE.62.1805). arXiv: [0002177 \[cond-mat\]](https://arxiv.org/abs/0002177).
- [24] J. F. Antin, S. Lee, M. A. Perez, T. A. Dingus, J. M. Hankey, and A. Brach, "Second strategic highway research program naturalistic driving study methods", *Safety Science*, vol. 119, no. January, pp. 2–10, Nov. 2019, ISSN: 09257535. DOI: [10.1016/j.ssci.2019.01.016](https://doi.org/10.1016/j.ssci.2019.01.016). [Online]. Available: <https://doi.org/10.1016/j.ssci.2019.01.016%20https://linkinghub.elsevier.com/retrieve/pii/S0925753518301012>.

- [25] R. Krajewski, J. Bock, L. Kloecker, and L. 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)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 2118–2125, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569552](https://doi.org/10.1109/ITSC.2018.8569552). arXiv: [1810.05642](https://arxiv.org/abs/1810.05642). [Online]. Available: <https://ieeexplore.ieee.org/document/8569552/>.
- [26] U.S. Department of Transportation Federal Highway Administration, *Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data. [Dataset]*, 2016. [Online]. Available: <https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajectory/8ect-6jqj>.
- [27] E. Bampounakis and N. Geroliminis, "On the new era of urban traffic monitoring with massive drone data: The pNEUMA large-scale field experiment", *Transportation Research Part C: Emerging Technologies*, vol. 111, no. November 2019, pp. 50–71, 2020, ISSN: 0968090X. DOI: [10.1016/j.trc.2019.11.023](https://doi.org/10.1016/j.trc.2019.11.023). [Online]. Available: <https://doi.org/10.1016/j.trc.2019.11.023>.
- [28] O. Siebinga, "TraViA: a Traffic data Visualization and Annotation tool in Python", *Journal of Open Source Software*, vol. 6, no. 65, p. 3607, Sep. 2021, ISSN: 2475-9066. DOI: [10.21105/joss.03607](https://doi.org/10.21105/joss.03607). [Online]. Available: <https://joss.theoj.org/papers/10.21105/joss.03607>.
- [29] M. Saifuzzaman and Z. Zheng, "Incorporating human-factors in car-following models: A review of recent developments and research needs", *Transportation Research Part C: Emerging Technologies*, vol. 48, pp. 379–403, Nov. 2014, ISSN: 0968090X. DOI: [10.1016/j.trc.2014.09.008](https://doi.org/10.1016/j.trc.2014.09.008). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2014.09.008><https://linkinghub.elsevier.com/retrieve/pii/S0968090X14002551>.
- [30] S. Ossen and S. P. Hoogendoorn, "Heterogeneity in car-following behavior: Theory and empirics", *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 2, pp. 182–195, 2011, ISSN: 0968090X. DOI: [10.1016/j.trc.2010.05.006](https://doi.org/10.1016/j.trc.2010.05.006). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2010.05.006>.
- [31] S. Hoogendoorn, S. Ossen, and M. Schreuder, "Empirics of Multianticipative Car-Following Behavior", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1965, no. 1965, pp. 112–120, Jan. 2006, ISSN: 0361-1981. DOI: [10.3141/1965-12](https://doi.org/10.3141/1965-12). [Online]. Available: <http://trrjournalonline.trb.org/doi/10.3141/1965-12>.
- [32] R. Jiang, M. B. Hu, H. M. Zhang, Z. Y. Gao, B. Jia, and Q. S. Wu, "On some experimental features of car-following behavior and how to model them", *Transportation Research Part B: Methodological*, vol. 80, pp. 338–354, 2015, ISSN: 01912615. DOI: [10.1016/j.trb.2015.08.003](https://doi.org/10.1016/j.trb.2015.08.003). [Online]. Available: <http://dx.doi.org/10.1016/j.trb.2015.08.003>.
- [33] M. Mulder, M. Mulder, M. M. Van Paassen, and D. A. Abbink, "Effects of lead vehicle speed and separation distance on driver car-following behavior", *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, vol. 1, no. 4, pp. 399–404, 2005, ISSN: 1062922X. DOI: [10.1109/icsmc.2005.1571179](https://doi.org/10.1109/icsmc.2005.1571179).
- [34] T. Kondoh, T. Yamamura, S. Kitazaki, N. Kuge, and E. R. Boer, "Identification of Visual Cues and Quantification of Drivers' Perception of Proximity Risk to the Lead Vehicle in Car-Following Situations", *Journal of Mechanical Systems for Transportation and Logistics*, vol. 1, no. 2, pp. 170–180, 2008, ISSN: 1882-1782. DOI: [10.1299/jmtl.1.170](https://doi.org/10.1299/jmtl.1.170). [Online]. Available: http://www.jstage.jst.go.jp/article/jmtl/1/2/1%5C_2%5C_170/%5C_article.
- [35] A. Ng and S. Russell, "Algorithms for inverse reinforcement learning", *Proceedings of the Seventeenth International Conference on Machine Learning*, vol. 0, pp. 663–670, 2000, ISSN: 00029645. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3). [Online]. Available: <http://www-cs.stanford.edu/people/ang/papers/icml00-irl.pdf>.
- [36] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning", *Proceedings, Twenty-First International Conference on Machine Learning, ICML 2004*, pp. 1–8, 2004. DOI: [10.1145/1015330.1015430](https://doi.org/10.1145/1015330.1015430).

- [37] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning", in *Proceedings of the National Conference on Artificial Intelligence*, vol. 3, 2008, pp. 1433–1438, ISBN: 9781577353683.
- [38] O. Siebinga, *IRL Model Validation - TraViA extension code*, <https://github.com/tud-hri/irlmodelvalidation>, 2021.
- [39] M. Mulder, M. M. Van Paassen, M. Mulder, J. J. Pauwelussen, and D. A. Abbink, "Haptic car-following support with deceleration control", *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, no. October, pp. 1686–1691, 2009, ISSN: 1062922X. DOI: [10.1109/ICSMC.2009.5346803](https://doi.org/10.1109/ICSMC.2009.5346803).
- [40] Ruina Dang, Fang Zhang, Jianqiang Wang, Shichun Yi, and Keqiang Li, "Analysis of Chinese driver's lane change characteristic based on real vehicle tests in highway", in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, vol. 100084, IEEE, Oct. 2013, pp. 1917–1922, ISBN: 978-1-4799-2914-6. DOI: [10.1109/ITSC.2013.6728509](https://doi.org/10.1109/ITSC.2013.6728509). [Online]. Available: <http://ieeexplore.ieee.org/document/6728509/>.
- [41] S. Levine and V. Koltun, "Continuous Inverse Optimal Control with Locally Optimal Examples", *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, vol. 1, pp. 41–48, Jun. 2012. arXiv: [1206.4617](https://arxiv.org/abs/1206.4617). [Online]. Available: <http://arxiv.org/abs/1206.4617>.
- [42] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, *Social behavior for autonomous vehicles - Supporting Information*, 2019. [Online]. Available: <https://www.pnas.org/content/suppl/2019/11/22/1820676116.DCSupplemental> (visited on 06/19/2021).
- [43] M. Zhu, X. Wang, A. Tarko, and S. Fang, "Modeling car-following behavior on urban expressways in Shanghai: A naturalistic driving study", *Transportation Research Part C: Emerging Technologies*, vol. 93, no. September 2017, pp. 425–445, Aug. 2018, ISSN: 0968090X. DOI: [10.1016/j.trc.2018.06.009](https://doi.org/10.1016/j.trc.2018.06.009). [Online]. Available: <https://doi.org/10.1016/j.trc.2018.06.009> <https://linkinghub.elsevier.com/retrieve/pii/S0968090X18308635>.
- [44] M. Naumann, L. Sun, W. Zhan, and M. Tomizuka, "Analyzing the Suitability of Cost Functions for Explaining and Imitating Human Driving Behavior based on Inverse Reinforcement Learning", in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2020, pp. 5481–5487, ISBN: 978-1-7281-7395-5. DOI: [10.1109/ICRA40945.2020.9196795](https://doi.org/10.1109/ICRA40945.2020.9196795). [Online]. Available: <https://ieeexplore.ieee.org/document/9196795/>.
- [45] S. Rosbach, V. James, S. Grosjohann, S. Homoceanu, and S. 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)*, IEEE, Nov. 2019, pp. 2658–2665, ISBN: 978-1-7281-4004-9. DOI: [10.1109/IROS40897.2019.8968205](https://doi.org/10.1109/IROS40897.2019.8968205). arXiv: [1905.00229](https://arxiv.org/abs/1905.00229). [Online]. Available: <https://ieeexplore.ieee.org/document/8968205/>.
- [46] Z. Huang, J. Wu, and C. Lv, "Driving behavior modeling using naturalistic human driving data with inverse reinforcement learning", *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–13, 2021. DOI: [10.1109/TITS.2021.3088935](https://doi.org/10.1109/TITS.2021.3088935).
- [47] A. R. Srinivasan, M. Hasan, Y.-S. Lin, et al., *Comparing merging behaviors observed in naturalistic data with behaviors generated by a machine learned model*, 2021. arXiv: [2104.10496](https://arxiv.org/abs/2104.10496) [cs.LG].
- [48] G. Markkula and M. Dogar, "How accurate models of human behavior are needed for human-robot interaction? For automated driving?", Feb. 2022. arXiv: [2202.06123](https://arxiv.org/abs/2202.06123). [Online]. Available: <https://arxiv.org/abs/2202.06123>.
- [49] L. Sun, Z. Wu, H. Ma, and M. Tomizuka, "Expressing Diverse Human Driving Behavior with Probabilistic Rewards and Online Inference", in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2020, pp. 2020–2026, ISBN: 978-1-7281-6212-6. DOI: [10.1109/IROS45743.2020.9341371](https://doi.org/10.1109/IROS45743.2020.9341371). arXiv: [2008.08812](https://arxiv.org/abs/2008.08812). [Online]. Available: <https://ieeexplore.ieee.org/document/9341371/>.

3

Uncovering variability in human driving behavior through automatic extraction of similar traffic scenes from large naturalistic datasets



Recently, multiple naturalistic traffic datasets of human-driven trajectories have been published (e.g., highD, NGSim, and pNEUMA). These datasets have been used in studies that investigate variability in human driving behavior, for example for scenario-based validation of autonomous vehicle (AV) behavior, modeling driver behavior, or validating driver models. Thus far, these studies focused on the variability on an operational level (e.g., velocity profiles during a lane change), not on a tactical level (i.e., to change lanes or not). Investigating the variability on both levels is necessary to develop driver models and AVs that include multiple tactical behaviors. To expose multi-level variability, the human responses to the same traffic scene could be investigated. However, no method exists to automatically extract similar scenes from datasets. Here, we present a four-step extraction method that uses the Hausdorff distance, a mathematical distance metric for sets. We performed a case study on the highD dataset that showed that the method is practically applicable. The human responses to the selected scenes exposed the variability on both the tactical and operational levels. With this new method, the variability in operational *and* tactical human behavior can be investigated, without the need for costly and time-consuming driving-simulator experiments.

3.1. Introduction

In recent years, multiple open-access naturalistic datasets have been published. Some of these datasets are constructed by first recording videos of traffic with mounted cameras (e.g., NGSIM [1]) or drones (e.g., highD [2] and pNEUMA [3]). Image recognition techniques are then used to extract trajectory data from these videos. Such datasets contain trajectories for all vehicles that pass through a specific area.

Researchers have used these datasets for multiple purposes, among which: scenario-based validation of autonomous vehicle (AV) behavior (see [4] for a review), modeling and predicting driver behavior (e.g., [5]–[7]), and validating driver models to be used in autonomous vehicles (e.g., [8]). In all these applications, the *variability* (the range of human behaviors that can be expected in response to a given situation, sometimes referred to as uncertainty) in driver behavior is relevant.

Currently, variability is mostly regarded on the level of *operational* driving behavior (e.g., [5], [6], [9]–[11]). *Operational* driving behavior considers the execution of a maneuver [12], for example a lane change. However, variability does also exist on the *tactical* level, that is, in the choice of maneuver when a driver responds to a traffic scene [12]. For instance, some drivers might respond to a slower-moving vehicle in their lane by overtaking it, while others will brake in the same situation. Understanding variability on both the operational and the tactical level is important for assessing the human-likeness and acceptability of AV behavior, and also for validating human driver models used in AVs [8]. The reason is that both these applications must consider all possible tactical behaviors under given conditions. Traditional driver models on the other hand, mostly target a specific tactical behavior (e.g., car following in the Intelligent Driver Model [13]), thus for their application, only operational variability is relevant. When designing driver models that describe multiple tactical behaviors, the variability in tactical behavior also needs to be understood.

To study driver behavior variability to its full extent, similar traffic scenes have to be (automatically) extracted from the previously mentioned datasets to compare the human responses to these scenes. However, most automatic extraction methods select traffic *scenarios* (see [4] for a review), not traffic *scenes*. According to Ulbrich et al. [14], a scenario "describes the temporal development between several scenes in a sequence of scenes..", where "a [traffic] scene describes a snapshot of the environment including the scenery and dynamic elements.." (dynamic elements in the discussed datasets are (human-driven) vehicles). These definitions show that (extracted) traffic scenarios include part of a trajectory. Similar trajectories describe the same tactical behavior in most cases. Thus, selecting similar traffic scenarios implicitly means selecting similar tactical responses. Some other approaches even explicitly extract data corresponding to a pre-specified tactical behavior (e.g. lane changes in [6], [8]).

These existing approaches expose the variability in the operational execution of a given tactical maneuver but disregard the variability in tactical behavior. Furthermore, including trajectories as part of the automatically extracted data conflates the initial traffic scene (i.e., what a driver is responding to) and the driver's response itself. This makes it more difficult to investigate the full extent of variability in human responses to a specific initial traffic scene.

A method to automatically extract similar traffic scenes from naturalistic datasets would support studies into driver behavior variability on both operational and tac-




Figure 3.1: The four steps of the proposed method to find similar scenes in a naturalistic traffic dataset. In the first step, the scene of interest (of which similar instances should be found) is defined by manually selecting an example of this scene from the naturalistic dataset. The second step consists of extracting the traffic context from this example and converting it to a mathematical set of points — the context set. We use the scaling parameter λ in this conversion because we target highway traffic, but this is optional. The distance between this set and all other sets present in the data is calculated in step three using the Hausdorff distance. Optional pre-filtering steps can be used here to reduce the computational load of the method. Finally, in step four the N scenes with the shortest distance to the example can be selected.

tical levels. With such a method, the trajectories in response to an initial scene can be studied, both in terms of operational and tactical characteristics. However, to the best of our knowledge, all the methods that have been proposed to automatically extract traffic scenarios from the data cannot extract traffic scenes. This paper proposes a method to automatically extract similar traffic scenes from large naturalistic datasets (for a schematic overview, see Figure 3.1). We introduce the concept of *traffic context* to specify the part of the initial traffic scene that is relevant for comparing human responses. We define *traffic context* as all positions and velocities of all surrounding vehicles at a given time. Compared to the complete traffic scene, the traffic context excludes scenery and the state of the ego vehicle. Therefore, the traffic context represents the aspects of the scene the ego-vehicle is responding to.

We use a distance metric to express the difference between traffic contexts. With such a measurable distance, it becomes possible to automatically find scenes that are similar to a manually-selected example. However, such a distance metric for traffic scenes does not readily exist. Instead of developing a completely new distance metric, we use existing mathematical constructs and methods. This allows us to leverage existing literature and software implementations in our work. We propose to convert the traffic context to a mathematical set where each element in the set represents a vehicle in the traffic context. Once converted to a set, we use the Hausdorff distance for mathematical sets [15] to represent the distance with respect to other traffic sets.

An implementation of our method is provided on GitHub¹ as an extension of the traffic visualization software TraViA [16]. We validated the method in a case study using the highD dataset, where we show that this method is practically applicable and provides insight into the operational and tactical variability of driver behaviour.

3.2. Proposed Method

Our proposed method consists of four steps (Figure 3.1). These steps are briefly introduced in this section and explained in more detail when applied in the case study. In the first step, one should manually select an example of the scene of interest from the dataset. This example represents the traffic scene of interest of which multiple instances should be found. The traffic context from this example is converted to a mathematical set (the *context set*) in the second step. After that, the Hausdorff distance is used to determine the distance between the traffic

¹github.com/tud-hri/hausdorffsceneextraction

context in the selected scene and all other scenes in the dataset. Finally, one selects the N contexts with the shortest distance to the example. The resulting scenes are the scenes with traffic contexts most similar to the example across the whole dataset.




Figure 3.2: A visual example of how to calculate the Hausdorff distance from the blue set A to the orange set B in a 2-dimensional space. First, for every point in the blue set, find the closest (Euclidean) distance to any point in the orange set. These are shown here as solid lines. Then, select the longest of these minimum (solid-line) distances, this is the Hausdorff distance. This distance is here shown as the green solid line.

The Hausdorff distance is a distance metric for mathematical sets proposed by Felix Hausdorff in 1914 [15]. It can be used to express the distance between two non-empty compact sets. For two sets A and B , the Hausdorff distance represents the maximum distance between any point in set A and the closest point in set B (Figure 3.2). It can be used to describe the similarity between two sets, even if these sets have a different number of points. This practically means that it can be used to compare scenes with a different number of vehicles.

More extensive (and formal) explanations of the Hausdorff distance and how it can be calculated can be found in the literature (e.g., [17]) and online (e.g., [18]). Mathematically, the (directed) Hausdorff distance between two sets A and B is defined as:

$$h(A, B) = \max_{a \in A} \{ \min_{b \in B} \{ d(a, b) \} \} \quad (3.1)$$

Where we use the Euclidean distance between points a and b for distance $d(a, b)$. The directed Hausdorff distance from set A to B is not equal to the distance from B to A (i.e., $h(A, B) \neq h(B, A)$). Therefore we use the general Hausdorff distance H :

$$H(A, B) = \max \{ h(A, B), h(B, A) \} \quad (3.2)$$

3.3. Case Study: Methods

In this case study, we make use of the highD dataset [2] to show the potential of our proposed method. The highD dataset consists of traffic data recorded in Germany at 6 different highway locations. The dataset is made up of 60 independent recordings. To visualize the data and generate the images used in this work, we used the TraViA visualization software [16]. The proposed method's source code is publicly available as an extension to TraViA [19].

3.3.1. Step 1: Select an example

The first step of the proposed method is to select an example of the scene, to which the variability in responses is the topic of research. We will refer to this scene as *the traffic scene of interest*. This example should be at a specific point in time, seen from the perspective of a selected ego vehicle. For the highD dataset, this means that the example can be fully defined by a combination of three numbers: a dataset id, a vehicle id, and a frame number. For this case study, we have



Figure 3.3: The hand-picked example of the traffic scene of interest, as it is used in step 1 of our case study (highD dataset 1, frame 379, ego vehicle id 21). The blue vehicle is the ego vehicle, with the red arrow indicating the driving direction. The orange vehicles (id 25,26,20) denote other traffic participants, making up the traffic context, all driving in the same direction on the two-lane highway. In this example, the driver could decide to stay behind the vehicle it is currently following (id 20), but could also decide to accelerate and overtake that vehicle. The white arrows depicting x and y show the ego vehicle's reference frame. The gap between vehicle 25 and the ego vehicle is 95 meters. The numbers in red denote the longitudinal velocities of the surrounding vehicles.

selected the example as depicted in Figure 3.3. This scene can be found in dataset 1, frame 379 with ego vehicle id 21.

This example was selected because the driver of the ego vehicle (blue, id 21) can respond to this scene in multiple (tactical) ways as illustrated in Figure 3.3: the driver could decide to stay behind the vehicle it's currently following (id 20), but could also decide to accelerate and overtake that vehicle. The headway between the following vehicles (ids 25 and 26) and the ego vehicle is large enough (95 m and 128.7 m) to allow the ego vehicle to change lanes, but small enough to expect some effect of their presence on the ego vehicle's behavior.

3.3.2. Step 2: Extract the context set

The second step of our proposed method is to convert the traffic context to a mathematical set of 4-dimensional points. We will refer to this set as the context set consisting of context points. There is one context point for every surrounding vehicle. The context set can contain any number of context points, depending on the number of surrounding vehicles that are assumed to be part of the traffic context. In our case study, we used the definitions provided by highD to determine the vehicles that make up the traffic context. In the highD dataset, 8 positions for surrounding vehicles are reported for every ego vehicle. We assume these surrounding vehicles make up the traffic context.

To convert the state of the surrounding vehicles to context points, the 2-dimensional position of each vehicle is expressed in the ego vehicle's reference frame. The 2-dimensional absolute velocities (expressed along the same axes) of each vehicle are then concatenated to the relative positions. The result is a 4-dimensional point per surrounding vehicle. In mathematical form, a single context point representing a single surrounding vehicle can be expressed as

$$p_i = [x_i, y_i, v_{x,i}, v_{y,i}], \quad (3.3)$$

where i denotes the i^{th} surrounding vehicle, x and y denote the center x -position and y -position of the vehicle relative to the ego vehicle and v_x and v_y the velocities in x and y direction.

One potential problem with defining the context points as is done in Equation 3.3 is that it regards differences in longitudinal and lateral positions equally. However, on highways such as in the highD dataset, a small (e.g. 4 m) lateral difference in position can mean a vehicle is driving in another lane. This would substantially change the scene. While the same difference in longitudinal position would have a much smaller effect. To account for this difference, we introduce the parameter λ to scale the lateral dimension of the context points. This parameter should be estimated to reflect the relative importance of longitudinal and lateral position

and velocity differences. With the new parameter λ , the definition of the context points becomes

$$p_i = [x_i, \lambda y_i, v_{x,i}, \lambda v_{y,i}]. \quad (3.4)$$

In our case study, we have assumed $\lambda = 10.0$. This corresponds to the notion that a 1 meter change in the lateral position of a surrounding vehicle is equally important as a 10 meter change in the longitudinal position.

3.3.3. Step 3: Apply the Hausdorff distance

Now that the traffic context has been represented as a mathematical set, we can use the Hausdorff distance to compare different context sets. This step of the proposed method requires the Hausdorff distance to be calculated between the context set of the selected example and the context sets for all possible combinations of frame number and vehicle id in the dataset. For the highD dataset, there are 39.7×10^6 such combinations. Because this is a very large number of distances to be calculated, we will reduce it by filtering the relevant vehicles before calculating the distances.

When searching for scenes with similar traffic contexts, an important aspect is the lane in which the ego vehicle is driving. This determines where surrounding vehicles can be present and in which directions the ego vehicle can change lanes (e.g., a vehicle driving in the center lane can go both left and right, but a vehicle driving in the left lane can only change lanes to the right). For that reason we only consider vehicles driving in the same lane as the ego vehicle in the selected example. We consider 4 possible lanes: the left lane, the center lane, the right lane, and the merging lane. We determine the lane in the selected example (e.g., for Figure 3.3: the right lane) and only use the vehicle frame combinations from the dataset where the vehicle drives in the same lane. For our case study, this leaves 12,515,286 vehicle-frame combinations.

If the resulting number of distances to be calculated is still too large after applying this filter, one could consider down-sampling the frames. Depending on the specific frame rate of a chosen dataset, one could assume that the traffic context does not substantially differ within a certain number of frames and therefore only look at a subset of all frames. This would reduce the number of distances to be calculated even further. In our case study, this was not necessary because the resulting number of required distance calculations proved to be feasible. The calculation time on a high-end desktop (9th Gen Intel i9 8-core) was 3 hours, on a consumer-grade laptop (8th Gen Intel i7 4-core) it took 5 hours.

3.3.4. Step 4: obtain scenes

When all Hausdorff distances are calculated, the scenes in the dataset that are closest to the example can easily be obtained by selecting the N shortest distances. The only caveat here is that consecutive data frames are very similar, which results in groups with the same vehicle id and many consecutive frame numbers having very similar (short) Hausdorff distances to the example of the scene of interest. This problem can be accounted for by sorting all results based on the shortest distance only keeping the highest entry for every vehicle. Selecting the top N entries from the resulting table yields the final result.




Figure 3.4: The spread of the context points representing the results of the case study obtained after the final step. The top plot shows the positions of surrounding vehicles relative to the ego vehicle (see Figure 3.3 for the frame definition) where the stars represent the scene of interest. The ego vehicle always drives in the right-most lane. The bottom plot shows the absolute velocities of the surrounding vehicles (denoted by the stars). The ego vehicle's velocity is not regarded as part of the traffic context, so it differs for all scenes and is not depicted here. The stars represent the context set extracted from the selected example (Figure 3.3). The blue dots represent the 250 closest context sets that were automatically extracted from the highD dataset.

3.4. Case Study: Results

We used the proposed method in a case study to extract 250 scenes with a similar traffic context from the highD dataset. The hand-picked example of the scene of interest in step 1 is illustrated in Figure 3.3. The proposed method resulted in 250 scenes, where the traffic context is closest to this example. The spread of the resulting context points is shown in Figure 3.4. Of the 250 found scenes depicted in Figure 3.4, 233 contain precisely 3 surrounding vehicles, the same number as in the scene depicted in Figure 3.3. The other 17 scenes contain 4 surrounding vehicles. Figure 3.4 shows that the proposed method for automatically selecting scenes from a dataset succeeds in selecting context sets that are similar to the traffic context of the scene of interest. Note that the three clusters in this figure are not three independent distributions. The Hausdorff distance between sets can be interpreted as a trade-off between the points in a set. If one point is far away from the example, the other two need to be closer to result in a short Hausdorff distance. Therefore, the points within one set cannot be seen as samples from independent distributions.

Figure 3.4 also shows that the resulting spread is larger in the longitudinal direction than in the lateral direction. For example, in longitudinal positions, the maximum difference between the found sets and the selected example is approximately 25 m where the maximum lateral deviation is approximately 2 m. These values correspond to the used λ value of 10.

The variability in the results does depend on the amount of data and the scene of interest. The proposed method finds the closest available sets, so if the example represents a more common scene or the dataset to search is larger, lower variability in the found context sets can be expected. The variability in results can also be reduced by selecting fewer context sets i.e. select the $N = 100$ closest set instead of the $N = 250$, but this is a trade-off with the power of the resulting variability estimation.

Among other use cases, these results can be used for research targeting the variability in human responses to similar traffic contexts. To illustrate the utility of the results, Figure 3.5 shows these human responses. 26 drivers (10%) responded to this scenario by changing lanes within 3 seconds, 62 drivers (25 %) slowed down by




Figure 3.5: The variability in driver responses (driven trajectories) as they evolve from the 250 traffic scenes with similar traffic context automatically extracted from the highD dataset (represented in Figure 3.4). The lateral positions are normalized such that 0 m indicates the center of the original (right-most) driving lane. Lane widths differ slightly within the highD dataset but are approximately 4 m . The horizontal grey bars show the range in which all lane markings fall. Blue dots depict the drivers' initial positions and grey lines depict their individual trajectories - with markers for 1 (orange), 2 (green), and 3 (red) seconds. Distributions on longitudinal and lateral vehicle behavior are estimated and shown on the top and right sides of the figure. The figure illustrates tactical variability: some vehicles ($n=19$) make a lane change (the red dots with substantial positive lateral positions) while others keep car following in the original lane. Operational variability can also be observed in position and velocity for both the lane-changers and the car-followers.

more than 1 % of their initial velocity, and the other 162 drivers (65 %) did not slow down or change lanes within the 3 second period.

The figure also shows kernel density estimations of the longitudinal and lateral distributions for multiple points in time. These estimated distributions could be used to validate driver models that make predictions in the form of distributions. The figure illustrates two potential benefits of the proposed method: the method can be used to extract scenes to which humans respond with different tactical behaviors, and distributions of human behavior can be estimated from the responses to these selected scenes.

3.5. Discussion

In this paper, we propose a novel method to automatically extract similar traffic scenes from large naturalistic datasets. In a case study on the highD dataset, we showed that the proposed method is practically applicable and provides insightful results that expose the operational *and* tactical variability in human responses to similar traffic scenes. Therefore, our proposed method can be a valuable tool for the development of autonomous vehicles and traffic systems that incorporate human responses in their control decisions. Also, the case study showed that humans respond to similar traffic scenes with different tactical behaviors (some change lanes while others stay in their initial lane).

One approach closely related to our method is that of clustering scenarios. As discussed in the introduction, obtaining similar scenarios serves a different use case than extracting similar scenes. However, clustering requires a distance metric which makes it comparable to our method. There are two specific trajectory clustering methods that bear resemblance to our approach. In [20], the same distance metric is used as in our approach: the Hausdorff distance. However, in their approach, it is used to determine the distance between two trajectories

by regarding the waypoints as a set while we convert the traffic context to a mathematical set. In [21], another distance metric for scenes is proposed based on a grid around the ego vehicle and the longitudinal distances to other vehicles. Although similar scenes can indeed be found using only longitudinal distance, our method based on the Hausdorff distance is more complete because it also takes into account the lateral positions and longitudinal and lateral velocities of the surrounding vehicles.

Using naturalistic traffic datasets is not the only way to investigate variability in human responses to the same scene, driving-simulator experiments are a well-established alternative. In a driving simulator, multiple participants can be subjected to exactly the same scene with the same traffic context. However, naturalistic data should be used for some applications, for instance when validating human driver models for autonomous vehicles [8]. In other cases, a large diversity of drivers might be needed (e.g., when interested in behavior across the population). This would make driving-simulator experiments time-consuming and expensive. For those reasons, our proposed method based on naturalistic data to study human responses to similar traffic contexts is a valuable new approach.

The proposed approach has four main limitations, some of which can be addressed in future work. First, there is no measure to determine how similar two traffic scenes are from a human perspective. This means that the magnitude of similarities and differences between the selected scenes, and thus the method's performance, cannot be quantified. The best way to construct such a measure would be to collect similarity ratings from humans by letting them experience selected pairs of scenes from the dataset.

Second, the dimensions of vehicles are not taken into account for the traffic context. This could be addressed in a post-processing step if these dimensions are deemed important to answer one's research question. This might, however, limit the amount of extracted scenes. Third, the initial velocity of the ego vehicle is ignored. This was done purposefully because we argue that the initial velocity of the ego vehicle is part of the human response, not of the traffic context. This is a limitation when the resulting data is used to validate driver models that do take this information into account. Including the ego vehicle's velocity could be done by adding the ego vehicle as an extra context point to the context set at position (0.0, 0.0).

Finally, we presented no systematic approach to determine the parameter λ . The main reason is that the relative importance of the longitudinal and lateral positions of other vehicles can depend on a number of factors, such as the environment, vehicles' velocities, road dimensions, and the targeted scene. We would like to point out that we needed the λ parameters because we used highway data. In other environments with a single lane per direction, such as in inner-city traffic, the λ parameter would not be needed.

To investigate the sensitivity of the λ parameter we repeated our selection procedure with the λ -value increased and decreased by 10% ($\lambda = 9.0$ and $\lambda = 11.0$). We evaluated the results by comparing the found dataset and vehicle IDs between the original and new λ values. The results for the lowered λ deviated for 15 vehicles (6%) and for the increased λ for 12 vehicles (5%). We thus conclude that our proposed method is not extremely sensitive to the choice of λ , and that it can be safely estimated manually.

Besides the limitations, there are some possibilities for extending the proposed method. It could be extended to include multiple types of traffic users (e.g.,

vehicles, pedestrians, and cyclists). To do this, every group of traffic users should be converted to an individual set. The distances of all sets can then be summed in step 3 to find the closest scenes. Besides that, contextual information could be regarded in the post-processing step. This would allow for the inclusion of factors such as weather or lighting conditions (if this information is available with the data).

Furthermore, we used the highD dataset in our case study, but the method itself is suitable for use with all trajectory data (e.g., with the pNeuma or NGSim datasets). The main advantage of the highD dataset is that includes information about the surrounding vehicles. This is a pre-processing step that has to be performed on other datasets before they can be used with the proposed method.

In this paper, we have shown a case study on a single example from the highD dataset. Although we believe it is an illustrative example, it only shows the results of our method for a single scene. To verify if the method is generalizable, we repeated the procedure for other example scenes. However, due to the page limit, we did not share those results here. To aid in the reproduction of these results and to enable the replication (including the generation of figures) on other scenes from the highD dataset, we openly share the source code of our method [19]. Future studies can also use this code to systematically investigate the use of our method for different applications and other traffic datasets.

3.6. Conclusion

We conclude that:

- When a set of comparable traffic scenes needs to be extracted from a large naturalistic dataset (e.g. for human factors analyses), our proposed methodology offers an automated and repeatable approach. We demonstrate our method on a HighD dataset, showing our method could find 250 comparable traffic scenes for a handpicked car-following scenario with three other vehicles surrounding the ego vehicle.
- With the extracted scenes, the variability in human responses be investigated, independent of the executed maneuver, and without the need for costly and time-consuming driving-simulator experiments.
- Our case study illustrates how the trajectories evolving from similar initial conditions (of 250 comparable traffic scenes) can be analyzed to show variability in operational and tactical driver behavior.

Bibliography

- [1] U.S. Department of Transportation Federal Highway Administration, *Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data*. [Dataset], 2016. [Online]. Available: <https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajector/8ect-6jqj>.
- [2] R. Krajewski, J. Bock, L. Kloeker, and L. 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)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 2118–2125, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569552](https://doi.org/10.1109/ITSC.2018.8569552). arXiv: [1810.05642](https://arxiv.org/abs/1810.05642). [Online]. Available: <https://ieeexplore.ieee.org/document/8569552/>.
- [3] E. Bampounakis and N. Geroliminis, "On the new era of urban traffic monitoring with massive drone data: The pNEUMA large-scale field experiment", *Transportation Research Part C: Emerging Technologies*, vol. 111, no. November 2019, pp. 50–71, 2020, ISSN: 0968090X. DOI: [10.1016/j.trc.2019.11.023](https://doi.org/10.1016/j.trc.2019.11.023). [Online]. Available: <https://doi.org/10.1016/j.trc.2019.11.023>.
- [4] S. Riedmaier, T. Ponn, D. Ludwig, B. Schick, and F. Diermeyer, "Survey on Scenario-Based Safety Assessment of Automated Vehicles", *IEEE Access*, vol. 8, pp. 87 456–87 477, 2020, ISSN: 2169-3536. DOI: [10.1109/ACCESS.2020.2993730](https://doi.org/10.1109/ACCESS.2020.2993730). [Online]. Available: <https://ieeexplore.ieee.org/document/9090897/>.
- [5] V. Mahajan, C. Katrakazas, and C. Antoniou, "Prediction of Lane-Changing Maneuvers with Automatic Labeling and Deep Learning", *Transportation Research Record*, vol. 2674, no. 7, pp. 336–347, 2020, ISSN: 21694052. DOI: [10.1177/0361198120922210](https://doi.org/10.1177/0361198120922210).
- [6] R. Krajewski, T. Moers, D. Nergler, and L. Eckstein, "Data-Driven Maneuver Modeling using Generative Adversarial Networks and Variational Autoencoders for Safety Validation of Highly Automated Vehicles", in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 2383–2390, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569971](https://doi.org/10.1109/ITSC.2018.8569971). [Online]. Available: <https://ieeexplore.ieee.org/document/8569971/>.
- [7] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [8] O. Siebinga, A. Zgonnikov, and D. Abbink, "A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles", *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 4, pp. 1–21, Dec. 2022, ISSN: 2573-9522. DOI: [10.1145/3538705](https://doi.org/10.1145/3538705). [Online]. Available: <https://dl.acm.org/doi/10.1145/3538705>.
- [9] S. Ossen and S. P. Hoogendoorn, "Heterogeneity in car-following behavior: Theory and empirics", *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 2, pp. 182–195, 2011, ISSN: 0968090X. DOI: [10.1016/j.trc.2010.05.006](https://doi.org/10.1016/j.trc.2010.05.006). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2010.05.006>.
- [10] V. Kurtc, "Studying Car-Following Dynamics on the Basis of the HighD Dataset", *Transportation Research Record*, vol. 2674, no. 8, pp. 813–822, 2020, ISSN: 21694052. DOI: [10.1177/0361198120925063](https://doi.org/10.1177/0361198120925063).
- [11] C. Thiemann, M. Treiber, and A. Kesting, "Estimating acceleration and lane-changing dynamics from next generation simulation trajectory data", *Transportation Research Record*, no. 2088, pp. 90–101, 2008, ISSN: 03611981. DOI: [10.3141/2088-10](https://doi.org/10.3141/2088-10).
- [12] J. A. Michon, "A Critical View of Driver Behavior Models: What Do We Know, What Should We Do?", in *Human Behavior and Traffic Safety*, Boston, MA: Springer US, 1985, pp. 485–524, ISBN: 0306422255. DOI: [10.1007/978-1-4613-2173-6_19](https://doi.org/10.1007/978-1-4613-2173-6_19). [Online]. Available: http://link.springer.com/10.1007/978-1-4613-2173-6_19.
- [13] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, vol. 62, no. 2, pp. 1805–1824, 2000, ISSN: 1063651X. DOI: [10.1103/PhysRevE.62.1805](https://doi.org/10.1103/PhysRevE.62.1805). arXiv: [0002177](https://arxiv.org/abs/0002177) [cond-mat].

- [14] S. Ulbrich, S. Grossjohann, C. Appelt, K. Homeier, J. Rieken, and M. Maurer, "Structuring Cooperative Behavior Planning Implementations for Automated Driving", *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2015-Octob, pp. 2159–2165, 2015. DOI: [10.1109/ITSC.2015.349](https://doi.org/10.1109/ITSC.2015.349).
- [15] F. Hausdorff, *Grundzuge der Mengenlehre*. 1914.
- [16] O. Siebinga, "TraViA: a Traffic data Visualization and Annotation tool in Python", *Journal of Open Source Software*, vol. 6, no. 65, p. 3607, Sep. 2021, ISSN: 2475-9066. DOI: [10.21105/joss.03607](https://doi.org/10.21105/joss.03607). [Online]. Available: <https://joss.theoj.org/papers/10.21105/joss.03607>.
- [17] W. Rucklidge, "The Hausdorff Distance", in *Efficient visual recognition using the Hausdorff distance*, ser. Lecture Notes in Computer Science, W. Rucklidge, Ed., vol. 1173, Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, ch. 2, pp. 27–42, ISBN: 978-3-540-61993-2. DOI: [10.1007/BFb0015091](https://doi.org/10.1007/BFb0015091). [Online]. Available: <http://link.springer.com/10.1007/BFb0015091>.
- [18] Wikipedia, *Hausdorff distance* - accessed on 25-2-2022, https://en.wikipedia.org/wiki/Hausdorff_distance, 2022. (visited on 02/25/2022).
- [19] O. Siebinga, *Hausdorff Scene Extraction*, 2022. [Online]. Available: <https://github.com/tud-hri/hausdorffsceneextraction>.
- [20] S. Atev, G. Miller, and N. P. Papanikolopoulos, "Clustering of vehicle trajectories", *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 3, pp. 647–657, 2010, ISSN: 15249050. DOI: [10.1109/TITS.2010.2048101](https://doi.org/10.1109/TITS.2010.2048101).
- [21] J. Kerber, S. Wagner, K. Groh, et al., "Clustering of the Scenario Space for the Assessment of Automated Driving", *IEEE Intelligent Vehicles Symposium, Proceedings*, no. Iv, pp. 578–583, 2020. DOI: [10.1109/IV47402.2020.9304646](https://doi.org/10.1109/IV47402.2020.9304646).

4

Modelling
communication-enabled traffic
interactions

A major challenge for autonomous vehicles is handling interactions with human-driven vehicles, for example in highway merging. A better understanding and computational modelling of human interactive behaviour could help address this challenge. However, existing modelling approaches predominantly neglect communication between drivers and assume that one modelled driver in the interaction responds to the other, but does not actively influence their behaviour. Here we argue that addressing these two limitations is crucial for the accurate modelling of interactions. We propose a new computational framework addressing these limitations. Similar to game-theoretic approaches, we model a joint interactive system rather than an isolated driver who only responds to their environment. Contrary to game theory, our framework explicitly incorporates communication between two drivers, and bounded rationality in each driver's behaviours. We demonstrate our model's potential in a simplified merging scenario of two vehicles, illustrating that it generates plausible interactive behaviour (e.g., aggressive and conservative merging). Furthermore, human-like gap-keeping behaviour emerged in a car-following scenario directly from risk perception without the explicit implementation of time or distance gaps in the model's decision-making. These results suggest that our framework is a promising approach to interaction modelling that can support the development of interaction-aware autonomous vehicles.

4.1. Introduction

Autonomous and automated vehicles (AVs) hold the potential to help address major societal challenges related to mobility and sustainability. However, one of the major open problems in autonomous vehicle development is safely and acceptably dealing with driving scenarios that require *reciprocal* interaction with human road users. In these interactions, such as in highway merging or intersection negotiation, both vehicles influence and respond to the actions of each other. It entails quick and sometimes iterative negotiations, based on communication (see e.g., [1]–[3]) that can either be implicit (vehicle motions) or explicit (e.g., honking, signalling). The continuous dynamics of a reciprocal interaction govern safety, priority (who goes first, who gives way), and acceptance (by passengers and other road users). For example, drivers can be misunderstood or cause annoyance by being too conservative or aggressive (interfering with or ignoring others' communication). Therefore, fundamental knowledge about continuous human reciprocal interactions is necessary to develop and evaluate safe and acceptable AV behaviour for these scenarios. However, this fundamental knowledge about the dynamics of interactions is currently lacking. We advocate using a modelling approach for human reciprocal traffic interactions to develop the fundamental understanding that in the future can help design better AV behaviour.

Modelling is a common way of gaining an understanding of human driving behaviour. But it has so far mostly been done with a focus on single-driver behaviour, either in single-vehicle (e.g., [4], [5]), or multi-vehicle scenarios such as car following [6], [7], lane changing [8], [9], and gap acceptance [10], [11]. Most multi-vehicle approaches assume that the modelled driver responds to other traffic participants, but that they don't respond in turn. For example, car-following models assume that the following driver responds to the leading vehicle, but this leading vehicle does not change its behaviour based on the follower's actions. We call this the *one-way interaction* assumption because the influence on behaviour is unidirectional. This assumption disentangles the behaviours of the multiple drivers and thereby enables the researchers to better understand and model the behaviour of the driver of interest. The scope of these models is thus deliberately restricted to a single driver. This one-way interaction assumption is justified for car-following models and the likes, but not for interactive driving scenarios such as merging or intersection negotiations, which are inherently reciprocal. Simply joining two one-way interaction models to describe an interaction will neglect the drivers' beliefs about the other's future actions and their expected influence on them. Furthermore, it also neglects the presence and effects of communication between the drivers. Therefore, we argue that the scope of an interaction model should include all participants to begin with. It should be a model of a joint interactive system.

The current mainstream approach to modelling joint interactive systems in traffic (as opposed to individual drivers) is using game theory. Game theory was developed as a framework to describe reciprocal interactions between players in abstract games. It has been used extensively to model traffic interactions. The first model of human merging behaviour based on game theory was proposed in 1999 by Kita [12]. In 2007, Liu et al. improved the game theoretical approach by removing the assumption of constant velocity [13]. After that, many works followed (e.g. [14]–[17]). However, applying game theory to model dynamics between two drivers is not trivial, because game theory makes three strong assumptions about these players.

First, the assumption that all players rationally maximize some utility function. Empirical evidence has shown that even in simple economic games [18], but also in driving behaviour [19] and traffic interactions [20], this assumption does not hold for human players. Second, game theory does not allow communication between the players, an aspect known to be important in interactive driving scenarios [3]. Third, the majority of game-theory-based interaction models use a set of discrete actions for the drivers. Although this is useful to describe the higher-level tactical [21] decisions of drivers accurately (for example the decision to yield or merge), it does not describe the lower-level operational [21] dynamics of the interaction (e.g. changes in velocity or trajectory). Therefore, these approaches are not sufficiently detailed for developing safe and acceptable AV behaviour. Combined, these three limitations motivate the need for an alternative approach to modelling reciprocal traffic interactions that allows for communication, bounded rationality, and continuous dynamic actions.

To address this gap, we propose a framework for Communication-Enabled-Interaction (CEI) modelling. It can be used to create model implementations, of which we provide one example in a case study¹. Our modelling framework relaxes the common (game-theoretic) assumptions that drivers are rational agents and have full information about the strategies of other drivers. It is based on the notion that all drivers have a plan they want to execute and a belief about what other drivers are going to do. Combined, this plan and belief result in a perceived risk for every driver. The drivers are assumed to act to keep this risk below their individual threshold. The key insight of the framework is that the beliefs about others are updated based on communication between the agents. In a simulation case study, we show that an implementation of a CEI-model produces plausible behaviour of two interacting drivers in a simplified merging scenario. Besides that, human-like gap-keeping behaviour emerges directly from the notion of risk perception. These results show that the proposed modelling framework provides a promising new approach for modelling human-human driver interactions.

4.2. Communication-Enabled Interaction (CEI) Modelling

We propose a framework to model reciprocal human-human traffic interactions between two drivers. This framework captures the joint system of both drivers rather than a single driver responding to their environment. The framework explicitly includes (implicit/explicit) communication between the drivers which facilitates the joint interaction. Each of the two drivers is described by four components: a notion of that driver's perceived risk, a deterministic plan for the driver's own control behaviour (e.g. accelerating/decelerating), a means of communication, and a probabilistic belief about the future behaviour of the other driver (Figure 4.1). The general framework we present here only defines loose requirements for how these four components should be operationalized. When implementing the model for a specific scenario or use case, these components can be operationalized based on existing literature (e.g., from the fields of human behaviour modelling, traffic communication, intent inference, or vehicle path planning). This means the model framework allows the incorporation of different methods to operationalize each of the four components, without having to fully redesign it. In this section, we will discuss the four components and our reasoning behind their functionality and

¹The software implementation of the presented model and its simulation environment are available online at [22]. The data discussed in the results section can be found at [23].

requirements. The assumptions and requirements that need to be taken into account when implementing a model based on this framework will also be discussed per component. In Section 4.3, we will illustrate how each component can be implemented in an example implementation for a simplified merging scenario.

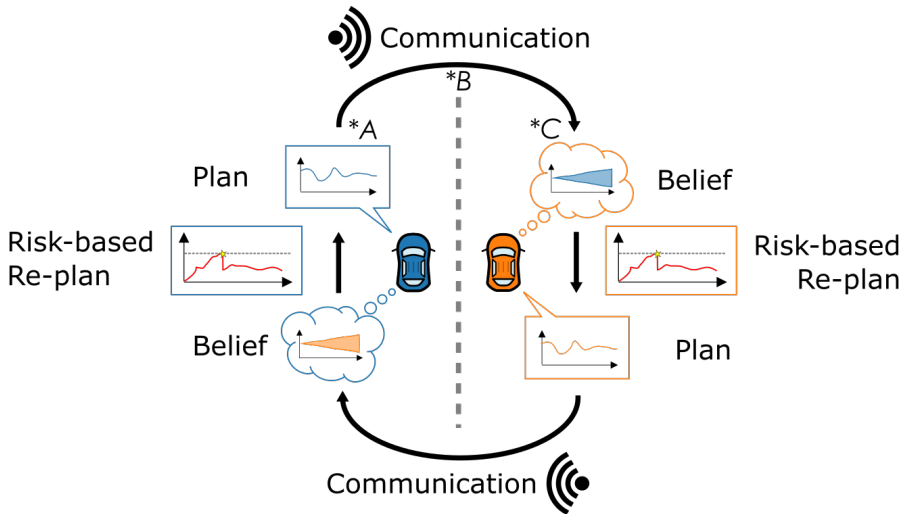


Figure 4.1: An overview of the proposed Communication-Enabled-Interaction (CEI) modelling framework. This framework is designed to capture the reciprocal interaction between two drivers, rather than the one-way interaction behaviour of one driver with respect to another. Each driver has a *plan* for their own behaviour. Plan updates are triggered based on a *risk threshold* and a *risk estimate* arising from a *belief* of how the other driver will move over time. Each driver *communicates* their plan (intention) either implicitly (e.g., through vehicle motion), or explicitly (e.g., through light signals) to the other driver. This communication links one driver's plans to the belief of the other and can be divided into three components denoted *A, *B, and *C. *A represents the mapping of a driver's plan to its communication, *B represents the means of communication, and *C denotes the belief update of the other driver based on the received communication.

4.2.1. Framework components

Risk-based re-plan

Recent research in non-interactive scenarios has shown that a quantification of perceived risk can explain driver adaptations in speed and lateral position and can be used to accurately predict future trajectories [4], [24], [25]. In our framework, we combine this notion of risk-based decision-making in driving with Simon's ideas of *bounded rationality* [26] and *satisficing* [27]. *Bounded rationality* implies that humans are not capable of fully optimizing their behaviour all the time. *Satisficing* (a portmanteau of *satisfy* and *suffice*) is an example of bounded rationality in which humans are assumed to not continuously search for an optimal solution. Instead, they are *satisfied* with a "good enough" solution that *suffices*. We reason that the only solutions that suffice and satisfy in a driving interaction are the ones that are *subjectively safe enough*. To formalize these ideas, and combine them in a framework, we hypothesize that drivers act to keep their perceived risk below their personal risk threshold.

Using such a threshold incorporates Simon's ideas in two ways. First, it defines what solutions are subjectively safe enough. Second, it limits (or bounds) the cognitive capacities (or effort) required from the driver because it allows the driver to only

rethink their plan when the situation has changed and the current plan does not suffice or satisfy them anymore. This is what we call a *risk-based re-plan* (Figure 4.1). By incorporating these ideas, we step away from the fundamental assumption of game theory that humans are rational utility maximizers and move towards a formulation that allows for team effort and mutual goals.

In summary, our framework assumes every driver evaluates the risk of their current deterministic plan, given their probabilistic belief about what other drivers are planning to do. Risk perception can be based on a number of factors, such as high velocity, high acceleration, or the probability of a collision. This evaluation happens continuously, but drivers will only perform a re-plan if the perceived risk exceeds their threshold. This should result in drivers with a low risk-threshold adapting their plan in an early stage of the interaction to reduce the estimated risk. At the same time, drivers with a high risk-threshold will instead continue their current plan and take advantage of the fact that the risk of the situation is lowered by the other driver. Intuitively this can be explained as the driver with the higher risk-threshold being more aggressive.

Plan

The second component in our framework is the *plan*. We assume that drivers have a deterministic plan about the actions they will take in the immediate future. In the framework, this plan takes the form of a deterministic set of waypoints over a limited time horizon. This time horizon should be long enough to include (part of) the interaction.

The construction of this plan (i.e., the planning algorithm) should only consider features that are not related to risk and safety (e.g., desired velocity or comfort), as the perceived risk is constantly evaluated separately to determine if the current plan still suffices and satisfies. This evaluation is done taking into account both the plan and the belief. When re-planning, the risk threshold should be used as a constraint in the planning algorithm. As long as such a constraint can be imposed, the plan can be constructed using any suitable path-planning algorithm.

Communication

One of the key concepts of the framework is that drivers actively communicate their plan to other drivers. This assumption is based on field studies on human-human traffic interactions that confirm that traffic participants actively communicate their plan both explicitly and implicitly to others (e.g. [3]). Experiments on other (non-driving) tasks that require team effort have shown that humans use their movement actions to coordinate with their team member [28] (which is a form of implicit communication). The assumption of communication can also be effectively used to model human behaviour in those tasks [29]. Finally, in simulation, communication can be beneficial for controlling co-bots that navigate among humans [30], resulting in fewer dead-lock situations. In summary, previous research suggests that humans communicate in traffic and that the assumption of communication can be used both for the effective modelling of human teamwork behaviour and the effective control of robots.

In the CEI modelling framework, communication links the *plan* of one driver to the *belief* of the other driver. In practice, this means that three aspects of communication need to be designed when implementing a CEI-model. First, one needs to determine the mode of communication; What signals are used to communicate? These signals can be explicit (e.g., turn indicators) or implicit (e.g., velocity,

heading angle, or acceleration). Second, a mapping from a plan to its communication is required. This can be as simple as just executing the plan, but one could come up with more elaborate mappings based on traffic communication studies such as slowing down, purely to communicate that the other driver can go first (for an example of modelling such exaggerated trajectories in a bottle grasping task, see [29]). Finally, a mapping from communication to belief is needed, this mapping specifies how a probabilistic belief is updated based on the received communication.

Belief

Both drivers are assumed to have probabilistic beliefs about what the other driver will do in the near future. This belief consists of a number of points over a time horizon. Each of these belief points is represented by a probability distribution over positions for the other driver for that specific time in the future (Figure 4.1). This assumption is based on the intuition that human drivers have a general but uncertain idea about what other drivers are planning to do, a concept that has been successfully applied in other modelling frameworks such as belief-desire-intention programming (based on [31]) and (Bayesian) theory of mind [32].

When implementing the belief part of the CEI-model, the only requirement is that the chosen probability distribution can be updated using new information (coming from the observed communication). In practice, this means that most parametric probability distributions are suitable because they can be updated with methods such as Bayesian updates.

4.3. Case Study: an example of an implementation

To demonstrate the feasibility of the proposed model framework and to investigate the effects of design choices (parameters) on model behaviour, we have implemented a CEI-model for a simplified merging scenario. In this case study, we show that even with simple components the model framework can produce plausible, human-like interactive behaviour. At the same time, it is not the purpose of this case study to quantitatively assess the model's consistency with human behaviour. Such an assessment using fine-grained data on the interactive behaviour of two drivers requires a detailed investigation and is therefore left for future work.

4.3.1. Simplified merging scenario

For this case study, we used a simplified symmetric merging scenario (Figure 4.2). In this scenario, two vehicles approach a merge point on a predefined track. The model can directly control the acceleration of the vehicles, but there is no steering involved. The vehicles have a rectangular bounding box for collision detection. The heading of the vehicles is pre-defined and always corresponds to the heading of the road. At the merge point, the heading of the vehicles changes instantly. The vehicles in the simplified scenarios are subject to a negative acceleration due to resistance and drag. The net acceleration (a^{net}) is the applied input (a^{in}) minus the negative acceleration a^r (a function of the vehicle's velocity v):

$$a^{net}(v) = a^{in} - a^r(v), \text{ where} \quad (4.1)$$

$$a^r(v) = \alpha v^2 + \beta. \quad (4.2)$$

Parameters α and β define the magnitude of the drag and constant resistance ($\alpha = 0.0005$ and $\beta = 0.1$). Besides the resistance, the vehicles have a maximum




Figure 4.2: A top-down view of the simplified merging scenario as used in the case study, rotated 90 degrees clockwise. Vehicles follow pre-defined paths (road centres) that merge at a pre-defined merge point. Vehicles have a two-dimensional body ($4.5 \text{ m} \times 1.8 \text{ m}$) and their headings change instantly at the merge point. The model controls the accelerations of the vehicles directly. The dimensions of the track are defined by two parameters. Distance l_a (25 m) denotes the distance between the start points of the vehicles. Distance l_b (50 m) is the distance to travel from the start point until the merge point, and from the merge point until the end of the track.

acceleration $a^{max} = 2.5 \frac{\text{m}}{\text{s}^2}$, which is the same for positive and negative accelerations. The velocity of the vehicles is restricted to non-negative values. The simulation updates all dynamics at a rate of 20 Hz .

4.3.2. Plan

The planning part of the model consists of a path-planning algorithm that minimizes the following cost function:

$$c = \sum^N (v_n - v^d)^2 + (a_n^{in})^2. \quad (4.3)$$

Where n denotes the time-step and v the vehicle's velocity. This cost function includes terms for minimizing the squared input a^{in} and for travelling at a desired velocity v^d . The path is planned at the same frequency as the simulation (20 Hz) and is subject to a time horizon of 4 s ($N = \frac{4}{0.05} = 80$).

A visual example of the plan, belief, and risk perception is shown in Figure 4.3. When initially planning the path, the cost function of Equation 4.3 is minimized, so an optimal path is found with respect to comfort and speed (Figure 4.3-A). If, at the next time step, the current plan still satisfies (i.e., the risk threshold is not exceeded), the current plan is continued. We assume that maintaining velocity at the final time step is the practical equivalent of maintaining the current plan.

When the risk threshold is exceeded, the cost function is minimized again to find a new plan (Figure 4.3-C). This time the minimization is subject to a risk constraint. Based on the ideas of satisficing, we hypothesise that humans do not spend unlimited effort to find an optimal plan, but instead search for a new solution that satisfies and suffices. We hypothesize that re-planning is easiest (i.e., requires the least cognitive effort) if the new plan is close to the previous plan (i.e., uses the same strategy). Therefore, the re-planning optimization is executed with the old plan as the initial condition. When using a gradient descent algorithm, this will result in a solution that is close to the previous plan while the risk constraint is met. For example, if the current plan is to decelerate and pass behind the other driver, the most likely outcome of the re-planning will be to decelerate even more and increase the gap. This will lower the perceived risk while using the current strategy.




Figure 4.3: An example to illustrate the plan and the belief of the model. a) shows four (of 80) deterministic plan points along the one-dimensional track. These are the planned centre positions of the own vehicle at four points in time. The distributions represent the believed centre position of the other vehicle at the same four (of 16) points in time, where colours denote the points in time. b) shows these plan and belief points after a single belief update. This update increased the certainty of the belief about the other vehicle's position. The belief is updated at every time step. c) shows the risk evaluation for one of the points. To evaluate the risk, the probability of a collision (p_c) is evaluated by calculating the probability that the other vehicle will be within the bounds of collision for the given planned position. This risk evaluation is done at every time step for all belief points. If the maximum perceived risk value exceeds the upper risk-threshold, a re-plan is triggered. This re-plan uses the perceived risk as a constraint for the optimization. To lower the risk, the planned position could be moved in the direction of the black arrow.

If the optimization with the current plan as the initial condition does not succeed, three other initial conditions are considered: full braking at all time steps, no acceleration input at all time steps, and full acceleration at all time steps. The candidate plan with the lowest cost is used as the initial condition for a second re-plan. This can result in a change of strategy, but only if the current strategy is not feasible anymore. For example, when the driver was decelerating but decelerating even more will not reduce the risk enough, it will investigate if acceleration will reduce the risk and change its strategy if needed.

4.3.3. Belief

The belief is kept as a sequence of probability distributions over positions for the other vehicle, each at a specific point in time (Figure 4.3-A). This sequence of belief points uses the same time horizon as the planning part of the model (4 s) but contains fewer points for simplicity. Belief points are kept at a 4 Hz frequency (this number was based on an initial evaluation of the model), resulting in a sequence of $4 \cdot 4 = 16$ points. Each belief point is represented by a Gaussian distribution.

The Gaussian distributions are initialized by combining the initial velocity and position of the other vehicle with the maximum bounds of acceleration. To initialize a belief point, the mean of the Gaussian is set to the position that corresponds to the other driver maintaining its current velocity. To calculate the standard deviation, an upper and lower position bound (ub and lb) are used. These are calculated by predicting the position of the other vehicle if it would apply the maximum

and minimum possible acceleration continuously. The standard deviation is then calculated as the difference between the bounds and the mean divided by 3 ($\sigma = \frac{ub-\mu}{3}$). The factor $\frac{1}{3}$ is based on the fact that 99.73% of the area under a normal distribution corresponds to $\mu + / - 3\sigma$. Once the simulation time is equal to the timestamp corresponding to the first belief point, this point is removed from the sequence and a new point is initialized.

4.3.4. Communication

Human communication during driving is a complex topic on which a lot of research has been done. Thus, there is much potential for including complex communication models based on empirical evidence in a CEI-model. However, for this initial investigation of the modelling framework, we used a simple implicit communication model that does not include any explicit communication signals (e.g., turn indicators). We only use velocity and position as communication signals. These two values are assumed to be constantly observed by the other driver without any errors or noise.

When sending communication, the drivers do not use a mapping from their current plan to the actions they take. Instead, they just take the next action from their plan. When receiving communication, drivers use a constant velocity model combined with bounds of comfortable acceleration to update their belief. All belief points are updated every time step using Bayesian updating.

Updating the belief

For Bayesian updating, the previous belief point serves as the prior distribution, and the resulting posterior is adopted as the updated belief point (Figure 4.3-B). The likelihood is constructed using the constant-velocity model. We assume the likelihood to be a Gaussian distribution where the standard deviation is constant and known. This means the likelihood and prior form a conjugate pair, meaning that the posterior will also be a Gaussian distribution of which the μ and σ^2 have a closed-form solution. The likelihood function for the belief point at time t is defined as follows:

$$\mathcal{N}\left(\mu = \frac{p}{t}, \sigma^2 = \left(\frac{a_c t}{6}\right)^2\right) \quad (4.4)$$

In this equation, p denotes a position sampled from the prior (the previous belief point), t denotes the time corresponding to the belief point, and a_c is the maximum comfortable acceleration ($a_c = 1.0 \frac{m}{s}$). The same value is used for positive and negative accelerations, thus the distribution is symmetrical. The likelihood function describes the probability of observing a velocity v (now) given a sampled predicted position p (at time t) from the prior belief. The mean μ corresponds to constant velocity, and σ is determined based on the assumption that 99.73% of the distribution falls within the bounds of comfortable acceleration.

With this likelihood function, the posterior has a closed-form solution. We denote the prior as $\mathcal{N}(\mu_0, \sigma_0^2)$ and the posterior as $\mathcal{N}(\mu_1, \sigma_1^2)$. When updating with a single data point v , the solution for the posterior becomes²:

²For a complete derivation of this closed-form solution, see the supplementary material.

$$\mu_1 = \frac{\mu_0 \sigma^2 + v \sigma^2 \frac{1}{t}}{\sigma^2 + \sigma_0^2 \frac{1}{t^2}} \quad (4.5)$$

$$\sigma_1^2 = \frac{\sigma^2 \sigma_0^2}{\sigma^2 + \sigma_0^2 \frac{1}{t^2}} \quad (4.6)$$

4

4.3.5. Risk

The risk perceived by the drivers is assumed to be proportional to the probability of a collision. Other aspects (i.e., high velocity and high acceleration) are assumed not to contribute to the perceived risk for simplicity. To estimate the probability of a collision, we define the concept of *bounds of collision* (Figure 4.3-C). These are the extreme positions of the other vehicle that would result in a collision, given the position of the own vehicle. These bounds are calculated for every point in the driver's plan. For example, if we know the driver will be at position x at time t , we can use the vehicles' dimensions to calculate that a collision will occur if and only if the other vehicle is at a position between $x + c_1$ and $x - c_2$ at the same time; these are the bounds of collision. The believed probability that the other vehicle will be within these bounds at that time can be calculated using the belief about the other vehicle's position. This probability is then equal to the probability of a collision at that time.

The perceived risk for a complete plan is determined by taking the maximum risk over all belief points. A re-plan is triggered if the perceived risk exceeds an upper threshold ρ_u . Only using the upper threshold, however, poses a potential problem when the merging conflict is resolved because after that there will be no triggers to re-plan anymore. This might cause vehicles to stall or drive very slowly for no reason. We avoid this by extending the risk module with a lower risk threshold ρ_l and a saturation time τ . If the perceived risk is lower than ρ_l and the last update was longer than τ ago, a re-plan is also triggered. When a re-plan optimization is performed, the perceived risk is constrained to be lower than the average of the two thresholds. For the implementation of this constraint, the instant heading change at the merge point in the track posed a problem. Therefore, a linear approximation of the bounds of collision is used.

4.3.6. Investigated scenarios

In total, every driver in the model has four parameters that determine their behaviour: a desired velocity v_d , an upper risk-threshold ρ_u , a lower risk-threshold ρ_l , and a saturation time τ . Besides these parameters, the initial velocity and position (v_0 and x_0) of the drivers can also be adjusted. Both drivers always start from the beginning of the track. In the case study, we investigate the effect of these parameters and the effect of differences in the initial condition in four scenarios (Table 4.1).

The first two scenarios (A & B) manipulate the initial and desired velocity of the right driver while keeping the parameters of the left driver fixed; the drivers here have the same risk thresholds. In scenario A, the drivers are not expected to be on a collision course if they would stick to their desired velocity, but in scenario B, they are.

Scenarios C & D focus on the risk thresholds. Scenario C investigates the effect of a difference in risk thresholds between drivers. Scenario D investigates the sensitivity

Table 4.1: Parameters of the investigated scenarios. Underlined values denote deviations from the default values. ρ_l and ρ_u denote the lower and upper risk-thresholds. v_0 and v_d are the initial and desired velocity respectively. x_0 denotes the initial position of the vehicle along the track.

	Side	ρ_l	ρ_u	v_0 [m/s]	v_d [m/s]	x_0 [m]
Condition A: No expected collision	left	0.2	0.5	10.0	10.0	0.0
	right	0.2	0.5	<u>9.0</u>	<u>9.0</u>	0.0
Condition B: On a collision course	left	0.2	0.5	10.0	10.0	0.0
	right	0.2	0.5	<u>9.0</u>	<u>9.0</u>	<u>1.2</u>
Condition C: High and low thresholds	left	0.2	<u>0.4</u>	10.0	10.0	0.0
	right	<u>0.3</u>	<u>0.6</u>	10.0	10.0	0.0
Condition D: Threshold sensitivity	left	<u>0.3</u>	<u>0.4</u>	10.0	10.0	0.0
	right	<u>0.3</u>	<u>0.6</u>	10.0	10.0	0.0

of model behaviour to variations of these thresholds in one of the drivers. The saturation time τ only affects the behaviour after the conflict is resolved, therefore it is kept constant at 2.0 s for all scenarios.

4.4. Results

4.4.1. Scenario A: No expected collision

Scenario A serves as a baseline scenario. Here, both drivers have an initial velocity that is equal to their desired velocity, but that differs from the velocity of the other driver (Table 4.1). If they would keep their initial (desired) velocity up until the merge point, no collision would occur. The left driver would pass the merge point first with a small distance gap of 0.2 m. Therefore we would expect a rational optimizing model (that does not explicitly include human-like gap-keeping) to maintain the desired velocity all the way. A behaviour expected from human drivers, on the other hand, is to increase this small safety margin. In an empirical study [33], it was found that human drivers in the Netherlands merged on three different highway locations with mean headways of 12.6, 13.4, & 36.1 m for velocities below 60 km/h = 16.7 m/s, and standard deviations of respectively 10.3, 12.8 & 18.2 (the headway is defined as the gap plus the leading vehicle length).

In the modelled outcome of scenario A (Figure 4.4), the left driver reached the merge point first. They accelerated slightly to increase the safety margin at the merge point, after that, they returned to their preferred velocity. The headway when the second vehicle reached the merge point was 6.4 m. This corresponds to the expected human behaviour, and can not be modelled with utility-maximization unless utility is explicitly awarded for keeping a gap. The right driver did not take any action in this scenario. The reason for that is highlighted in the risk perception plot. The left driver's risk increases earlier because it expects to reach the merge point earlier. This increase causes the left driver to take action to lower the risk, while the right driver can continue their plan without exceeding their risk threshold. The right driver's perceived risk also decreases as soon as the left driver takes action; they perceive that the conflict was resolved by the left driver.

4.4.2. Scenario B: On a collision course

In scenario B, the drivers have the same desired and initial velocities as in scenario A. However, the right vehicle starts with a 1.2 m head-start. Therefore, the projected positions of the two vehicles at the merge point overlap by 1.0 m. Thus,




Figure 4.4: Model behaviour in scenario A (no expected collision). Line colours correspond to the vehicle colours in Figure 4.2. a) Positions of the left and right vehicles over time. The x positions of the vehicles are plotted with an offset to prevent the lines from overlapping after the merge point. The grey dots and dashed lines indicate vehicle positions at equal time stamps with an interval of 1.0 s. b) Velocities of the vehicles over time. The stars indicate the moment when the simulated driver performed a re-plan because the upper risk-threshold was exceeded, and a circle denotes a re-plan because the risk fell below the lower threshold. These re-plans are only triggered if the last re-plan was longer than τ ago. c) Accelerations of the vehicles over time. d) Perceived risk of both simulated drivers. In the case of a re-plan, the perceived risk after the re-plan is shown. The dashed horizontal lines in the lowest plots indicate the risk thresholds of the drivers. In this scenario, the drivers increased the small projected gap, even though they were initially not on a collision course. The simulated drivers behaved in a way to increase the initially narrow safety margin.

if neither driver deviates from their desired velocity, this scenario will result in a collision. We would therefore expect that this scenario requires more severe action to be resolved than scenario A, but we do expect the model to avoid a collision.

The modelled outcome of scenario B (Figure 4.5) shows that this scenario indeed requires more effort from both drivers to resolve the conflict compared to scenario A. Both drivers start braking until the left driver decides they can only reduce the risk of a collision by accelerating. This can be explained by the fact that the left driver has a slightly higher velocity at this point compared to the right driver. The right driver sticks to their plan and keeps decelerating until the risk drops below the lower threshold and the saturation time has passed, only then they accelerate again. This behaviour results in a safety margin between the vehicles that is not explicitly included in the reward function. Because the left driver is the first to accelerate, they reach the merge point first. This explainable interactive behaviour combined with the collision-free outcome can be regarded as a plausible human-like interaction.




Figure 4.5: Model behaviour in scenario B. The simulated drivers prevent a collision by slowing down. Initially, they both slow down, but after approximately one second, the left (initially faster) driver speeds up and reaches the merge point first. For details of the notation, see the caption of Figure 4.4.

4.4.3. Summary Scenarios A and B

In scenario A, the driver with the higher preferred velocity that approached the merge point first also passed the merge point first. But the distance gap between the vehicle was enlarged by the drivers. This corresponds to what we expected from human drivers. If the drivers approach the merge point with an expected collision (scenario B), however, the drivers take more drastic action but still manage to resolve the conflict by interacting with each other.

4.4.4. Scenario C: High and low thresholds

Scenario C represents a case where the simulated drivers of both vehicles have the same initial conditions and desired velocities, but different risk thresholds. Compared to the previous scenarios, the right driver has higher risk thresholds while the left driver has lower thresholds. The left driver, having lower thresholds, is expected to act early in the interaction to reduce their perceived risk. In terms of human behaviour, this would correspond to risk-averse, conservative driving. The right driver (high thresholds meaning higher tolerance to risk) is expected to react to a potential conflict at a later point and therefore to keep their velocity at the desired level longer. We expect that the right driver reaches the merge point first, and deviates less from their desired velocity compared to the left driver.

The modelled outcome of scenario C (Figure 4.6) is as expected: the left driver reached their upper threshold first and started to decelerate to reduce the perceived risk. In terms of human driving, this can be seen as more conservative behaviour. The right driver reacts later because their risk threshold is exceeded




Figure 4.6: Model behaviour in scenario C. The right driver maintains their initial velocity longer. After briefly decelerating, they accelerate and reach the merge point first. For details of the notation, see the caption of Figure 4.4.

at a later moment. They briefly decelerate, but quickly start to accelerate to reduce the risk since the left driver already decelerated. This results in the right driver reaching the merge point first and deviating less from their desired velocity than the left driver. This corresponds to the intuition that lower sensitivity to risk (i.e. higher risk thresholds) could be associated with more aggressive behaviour.

4.4.5. Scenario D: Threshold sensitivity

Scenario D investigates the sensitivity of the modelled drivers' behaviour to variations in the lower risk threshold. This scenario is the same as scenario C, with the only exception that the left driver has a slightly higher value for ρ_l (lower risk threshold). We, therefore, expect a very similar outcome in scenarios C and D. The only expected difference is that the left driver in scenario D re-plans more frequently because the risk for the new plan is constrained to the average of the two risk thresholds. With a smaller difference between ρ_l and ρ_u , the absolute risk decrease at the re-plan points is smaller. This should cause the perceived risk to reach the upper threshold quicker and thus result in more frequent re-plan events. However, the model simulation results show major differences between scenarios C and D (Figures 4.6 & 4.7). As expected, the smaller difference between the left driver's lower and upper risk-threshold resulted in more plan updates. But unexpectedly, this more frequent re-planning resulted in the left driver starting to accelerate and reaching the merge point first. To keep their perceived risk under control, the left driver deviated from their desired velocity to a larger extent than the right driver. This observation can be explained by the fact that high velocities

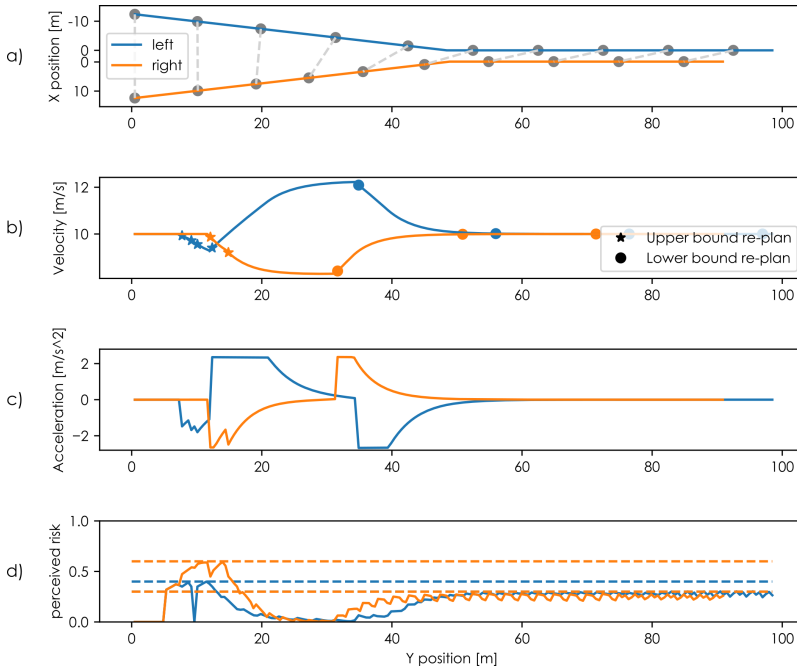


Figure 4.7: Model behaviour in scenario D. The slight change in p_l for the left driver (in comparison to scenario C) resulted in a major change in high-level outcome. Instead of the right driver, the left driver now reaches the merge point first. For details of the notation, see the caption of Figure 4.4.

and accelerations do not contribute to risk. The left driver takes whatever action is needed to keep the probability of a collision below their threshold (in this case, high acceleration and high velocity). The slight change in risk thresholds and more frequent re-plans resulted in one of the re-plans initially failing. This triggered a change in the left driver's high-level strategy, they accelerated instead of braked, and this heavily influenced the outcome.

4.4.6. Summary scenarios C and D

In scenario C, the driver with the higher risk thresholds (the right driver) passed the merge point first. This driver changed their plan at a later moment compared to the other driver. In terms of human behaviour, this can be explained as being more aggressive. The effect of slight changes to the lower threshold was shown to be substantial in scenario D. A small change resulted in a different interaction strategy, making the theoretically more "conservative" left driver arrive at the intersection first. This more conservative driver used high velocities and accelerations to lower their perceived risk even though high velocities would be interpreted by many human drivers as high-risk behaviour. The reason for this seemingly counter-intuitive model behaviour is that the high velocities and accelerations on their own do not contribute to the perceived risk of these modelled drivers.

4.4.7. Emergent gap-keeping behaviour for car following

Although the main focus of our model is on the interactive behaviour of the drivers when approaching the merging point, it also provides insight into their behaviour

after the merging conflict is resolved. Specifically, in the four scenarios above, we found that the simulated drivers continued maintaining a gap on the straight section after the merge point. This behaviour was not explicitly programmed and the planner has no cost associated with short time or small distance gaps (a feature frequently used in human driver models [34], [35]). Instead, these distance gaps appear to emerge from the combination of risk perception and a probabilistic belief about the plan of the other driver.

To further investigate this effect, we investigated a scenario without a merging point. In this scenario, the drivers drive behind each other on a straight stretch of road (400 m). We used the default parameters from Table 4.1, except for the velocity parameters. The leading vehicle has lower desired and initial velocities (9 m/s) compared to the following vehicle (10 m/s). Figure 4.8 shows that a steady-state gap emerges after approximately 100 meters. In this scenario, the leading driver mostly acts to reduce the risk and prevent a collision.

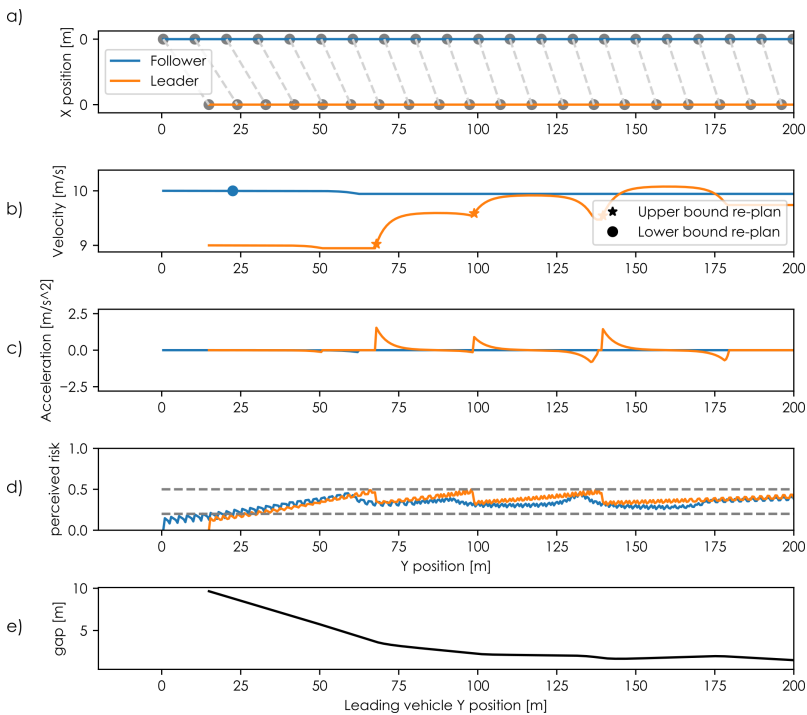


Figure 4.8: Model behaviour in the straight road scenario. For details of the notation, see the caption of Figure 4.4. The bottom panel shows the gap between the vehicles as a function of the leading vehicle position. In this scenario, the blue (following) vehicle has a higher preferred velocity compared to the orange (leading) vehicle. The x-axes have been cropped to the first 200 meters of the 400 meter track.

Although the fact that the leading, not the following, driver mostly acts to maintain this gap is not uncommon for human drivers and has been observed under some conditions [36], it is not the most common behaviour for reducing the risk during car following [37]. We identified two causes for this model behaviour. First, the belief and risk perception in the model are purely symmetrical. There is no difference in perceived risk between drivers that are in front of or behind another driver, nor is there any difference in believed probability that the other driver will accelerate

or decelerate. In natural traffic this simplification will not hold. This should be accounted for when extending the model for use in those scenarios. Second, the risk thresholds of both drivers are equal in this example. It can be expected that in other situations, even under the previously mentioned assumption, the driver with the lower risk-threshold will act to maintain the gap, as was seen in scenario C. This can be either the leading or the following driver, as was observed in human behavior [36], [37].

We investigated the effect of absolute velocities on the resulting steady-state distance gap, where we take the average gap over the final second of simulation as the steady-state gap. We simulated the model behaviour in this scenario for different velocities, every time with a 10 % velocity difference between the drivers, and an initial time gap of 1 s. We found that the emerging steady-state gap increased linearly with increasing velocities (Figure 4.9). This corresponds to human behaviour: the same linear relationship has been previously observed in a study on human gap-keeping behaviour on highways with low speeds [38].




Figure 4.9: Steady-state gap sizes (averaged over the last second) on a straight road where the following vehicle has a higher preferred velocity. The velocity difference between the vehicle is 10% and the initial time gap is 1 s.

Our model explains this relationship between velocity and distance gap as follows: The leading driver (orange) is unsure about the future plan of the following driver (blue). It could be possible that the blue driver will accelerate in the near future; In this case, a collision can occur. Because the orange driver keeps its risk below a threshold, it will keep a distance from the blue driver to make sure that its plan does not overlap too much with the possible future positions of the blue driver. Higher velocities, with the same maximum comfortable acceleration, result in a high standard deviation in the belief points. This causes the gap size to increase with velocity.

The mentioned study [38] also showed that humans keep larger gaps (approximately 12 m to 23 m for the same velocity range) compared to our model. We, therefore, conclude that the model qualitatively captures the underlying risk-mitigation mechanism in human car-following behaviour, but needs to be further explored to investigate if fitting the model parameters to human data would also allow it to capture the magnitude of the gap characteristic of human drivers.

4.5. Discussion

In this work, we have proposed a modelling framework for reciprocal human-human interactions in traffic. We illustrated the utility of the framework by implementing a concrete model based on the framework, targeted at interactive behaviour in a simplified merging situation. We investigated the model's behaviour in four scenarios, one where the drivers are not on a collision course, one where they are, and two where we investigated the effects of the model parameters. The model captures the actions of two drivers who 1) successfully resolve merging conflicts without collisions, 2) increase safety margins that are clearly too small (a 20 *cm* gap) for human drivers, and 3) exhibit individual conservative and aggressive behaviour, based on physically meaningful model parameters: their risk thresholds. In all scenarios, the model behaves in a plausible way that corresponds to intuitions about human interactive behaviour in merging conflicts.

Furthermore, from the model's underlying principle (the notion of risk combined with the probabilistic belief about the other driver's plan) plausible behaviour emerged outside of the situations we developed and tuned the model for. Specifically, a realistic gap-keeping behaviour emerged, where the drivers keep larger distance gaps at higher velocities, as humans do [38]. This behaviour was observed even though no distance or time gap-related costs are incorporated in the model. These results show that the proposed model framework is a promising novel approach for modelling reciprocal multi-agent interactions in traffic.

Modelling interactions in traffic has both practical and fundamental applications. In practice, a modelling framework like the one we propose could aid the development of autonomous vehicle controllers that aim to increase acceptability and safety in interactive scenarios. More fundamentally, such modelling, even when limited to an isolated traffic scenario, could contribute to gaining fundamental knowledge of human behaviour by highlighting the cognitive mechanisms humans use when interacting with each other. Our novel framework addresses the limitations of existing modelling and control approaches, among which game-theoretic models and interaction-aware controllers, because it explicitly incorporates communication and reciprocal interaction. Furthermore, our model framework does not make strong assumptions about human behaviour, such as the assumption that humans are rational utility maximizers. We hope that the initial exploration of the model framework presented here can spark a new strain of interaction modelling research.

Similar approaches

Among existing approaches to modelling traffic interactions, by far the most explored one is game theory. For example, for an extensive review of game-theory-based lane-changing models, see [39]. What is similar to our framework, is that game theory aims at modelling joint interactive systems instead of modelling only one driver responding to another (for examples, see [12]–[17]). What is different, is that our approach is not limited by two main assumptions (rationality, and lack of communication), and –for the majority of game-theoretic approaches– a focus on decision-making without describing operational behaviour. Finally, and more conceptually, game-theoretic models implicitly approach traffic interactions as a competition, while in our framework the agents have a joint primary objective (interaction safety) that makes the interaction a cooperative effort.

In contrast with game theory, our approach explicitly incorporates the drivers' ability to communicate their plan to other drivers, implicitly or explicitly. Although there are similarities with game theory, for example, our case study uses the same

modality of communication as many game-theoretical approaches, position and velocity observations (e.g., [13], [40], for an overview, see [39]), there are two fundamental distinctions in how we approach communication.

First, the communication in our framework allows drivers to construct and update a belief about the other vehicle's plan without the need for any prior information about the other driver. This is a fundamental contradiction with game theory, where players are assumed to know each other's utility functions (at least partially) beforehand. Therefore, in game theory, communication is not necessary because players can reason about what the other player is going to do to maximize their utility given the current state. Thus, the observations of position and velocity only serve to determine the state of the world. While in our model, position and velocity are used to convey information about the intention of drivers.

Second, in game theory, observations are not "remembered". They only serve to determine the *current* state, which is enough to reason about the other players' actions. Previous states are irrelevant. This is also known as the Markov condition or assumption. While in our work, the history of communication is kept in the belief about the other driver's intentions. Thus, the belief about a driver's future actions is based on its recent behaviour, not only on the current state. Some approaches combine game theory with an online estimation of the other player's utility function, thereby indirectly basing the belief about future actions (which directly depends on the utility function) on recent behaviour (e.g., [35], [41]). However, in these approaches, the conveyed information is not regarded as intentional communication. Furthermore, these approaches only estimate part (e.g., a single parameter) of the utility function online, the rest is assumed to be known a priori.

Another modelling concept that bears resemblance to our approach is that of *Belief-Desire-Intent (BDI) modelling*. BDI modelling is based on the philosophical work of Bratman [31] and models single agents that have a belief, a desire (goal), and an intent (plan). Many implementations of BDI models have been proposed for different applications [42]. The BDI framework and our CEI framework share the concepts that agents construct a (probabilistic) belief about other agents and the world, and then make a plan based on that belief to reach a final goal. The BDI framework, however, was not indented to account for interactions. It is primarily a model framework for individual agents that perform individual tasks. It therefore also does not incorporate communication but instead updates its beliefs based on changes that occurred in the world.

Finally, an important concept that can be complementary to the CEI-model framework, and bears resemblance to the BDI framework is the concept of *Theory of Mind (ToM)* [43] (for examples of applications to human-robot interaction, see [44], [45]). ToM is a psychological concept that assumes humans have an internal model of the beliefs, goals, and intentions of other humans in an interaction. Thereby, having the ability to reason about what other humans wants, and how they will try to achieve that goal. This idea that humans understand the mechanisms behind the actions and beliefs of others could be used in an implementation of our proposed CEI-model framework, which, in principle, only requires humans to form a basic belief about the future movements of others. As an example, the implementation of the CEI-model in the case study assumes drivers predict where the other driver is going, not why they are doing that. A complete ToM model could extend this belief about future actions of the other, with beliefs about their beliefs and goals. Implementing a CEI-based model with an internal ToM model is an interesting avenue for future research.

Besides these different types of modelling approaches, recently a great deal of effort was put into approaches for controlling (autonomous) vehicles in merging scenarios (e.g. [35], [41], [46]). Although the underlying techniques (such as finding a policy by optimizing some utility function) can be similar, the goal of these approaches is very different. While modelling approaches (such as ours) aim to best describe human behaviour. Control approaches aim to find a safe and optimal solution to a control problem. Game theory can therefore be very suitable for use in control approaches (as was done in [35], [41], [47]).

Two recent works on modelling come close in scope to this work. In 2022, Markkula et al. proposed a modelling approach for individual agents in a driver-pedestrian interaction rather than multiple agents in a driver-driver interaction [48]. Using different versions of a model that incorporates a variety of concepts from psychology, with varying levels of complexity, they conclude that "modelling of human road user interaction is a formidable challenge". Similar to our work, their findings suggest that the problem cannot be solved with simple rational models. Besides that, accounting for specific, previously unexplained, phenomena observed in human interactive behaviour could only be done using complex cognitive models. These conclusions resonate with our argument that the development of new model frameworks that go beyond game theory and the assumption of one-way interaction is a necessary step to improve our understanding of human traffic interactions.

Secondly, in 2014, Wan et al. also proposed an approach to model vehicle-vehicle interactions on merging ramps [49]. As in our work, they specifically address the reciprocal influence vehicles have on each other. Their (and our) work, therefore, differs from traditional driver models that usually describe a single driver responding to – but not influencing – other traffic. Another similarity between our proposed framework and the work by Wan et al. is that we both explicitly consider communication between vehicles. However, the model proposed by Wan et al. specifically targets congested traffic and uses different mathematical models for vehicles that have different roles in the interaction (i.e., they determine who will lead, follow, and merge a priori). Wan et al. also do not consider individual differences between drivers.

Framework extensions

Although we have only demonstrated our proposed model framework for a simple merging scenario with two vehicles, it could easily be extended to more vehicles or to traffic interactions with other types of participants. The underlying reason is that while we put the model's bounding box around the complete interaction, the drivers within the model are strictly separated; the only component connecting the two drivers is communication (Figure 4.1). This has two main advantages. First, communication in our framework is based on observable signals (e.g., turn indicators or velocity). This means that sending and receiving communication can easily be shared between multiple drivers, i.e., the communication is broadcast to all surrounding road users rather than sent directly to one of them. For that reason, the model framework can be extended to any number of drivers without requiring a redesign. Second, because the drivers are separated, it is possible to swap one of the drivers in the model with another type of agent, for example, a pedestrian. This would require adding the agent type to the observed communication, but since this is also an observable feature, it would not make the model more complex. One could even go as far as replacing one of the agents in the model with a non-model agent altogether. This could, for example, be used to let a real human

interact with the model in a driving simulator (this would require an optimized model implementation capable of running in real-time). This in turn would allow for the possibility of human drivers subjectively evaluating the ability of the model to describe natural interactions. Alternatively, a model could be used to evaluate autonomous vehicle controllers by letting the model interact with such a controller. Another potential extension useful for AV development is integrating the model into an AV controller to help it make decisions with an online evaluation of potential outcomes of an interaction.

We believe our model could also be adapted to other types of human-human interaction tasks. An example of such a task is cooperative bottle reaching, for which a communication model was developed in [29]. The task in [29] is similar to our task in that it constitutes a joint effort for which communication and action take place along the same channel (velocity/acceleration in our case). The main difference between our model framework and the communication model in [29] is that we target the interaction dynamics, in which we assume communication plays an important role, instead of targeting to model the communication as a stand-alone feature.

Limitations and future Work

Both the specific model implementation and the general modelling framework have important limitations. To start with the former, the model used for the simplified merging scenario uses very simplistic implementations for all components. The plan is based on desired velocity and acceleration alone. The beliefs are one-dimensional and assumed to be Gaussian distributions. The communication is assumed to be perfect (continuous without any noise), and only based on implicit cues. And finally, the risk is only based on collision avoidance, not influenced by high velocities or accelerations. In future implementations of the model, these limitations need to be addressed, and more realistic (and complex) model components should be investigated. However, it is important to first identify which of these limitations (if any) play a role in the model's ability to accurately reproduce human-human interactions. This could be done by comparing the model to data on human-human interactions gathered in a driving simulator experiment.

Another limitation of the current model implementation lies in the updates of the belief function. The assumption that the likelihood function (used for the Bayesian updates) has a known and fixed standard deviation results in the fact that every update reduces the standard deviation of the posterior, even if the new information contradicts the current belief. This is counter-intuitive: contradicting information (incoming through communication) should increase the variability of the belief, not decrease it. Put differently, if another person or driver sends unclear communication about what they are going to do by alternating between accelerating and braking, one should keep all options open, not decrease the standard deviation of the predicted position after a couple of seconds while shifting the mean around on every time step. How to properly address this limitation remains an open question.

Finally, the model's satisficing-based decision-making can result in unstable outcomes for high-conflict scenarios. When re-planning, the drivers in the model will first search for a new solution close to the previous solution. For example, if the previous plan was to brake, the driver will first explore if braking harder will satisfy the new constraint. Only if this optimization fails, the driver will explore other strategies (i.e., acceleration) to lower the perceived risk. This drastic change in high-level behaviour is thus triggered by the first optimization failing. Therefore,

slight numerical or temporal differences in this optimization can lead to different high-level outcomes, especially for situations that are highly symmetrical (e.g., when drivers have very similar parameters and none of the vehicles has a clear kinematic advantage). This was already observed in scenario D, where a slight change in model parameters caused a different outcome, but a similar outcome change could also result from changes in the type of numerical optimization solver or its parameters. One way of addressing this sensitivity is to make the model stochastic: introducing variability in the model's behaviour will make the outcome in high-conflict scenarios inherently stochastic and therefore could help to make it less sensitive to small external perturbations.

Adding stochasticity also addresses the main limitation of the overall framework, which is that currently, the framework is fully deterministic: with the exact same parameters (for model and solver), the model will always produce the same behaviour. This is inconsistent with the substantial behavioural variability that humans exhibit in traffic [50]. We see multiple possible ways of introducing stochasticity in the framework to account for this. To name two: adding stochasticity could be done in the receiving of communication (translating perceptual information to an updated belief) by using evidence accumulation mechanisms [10] or additive noise, or by including noise directly in the risk perception. However, more work is needed to determine the best approach.

A second limitation of the overall framework concerns improvements and redesigns of the model. Although the different components in the framework are separated, which should allow for easy redesign of parts of the model, they do depend on each other. This could mean that when redesigning one aspect of the model, a redesign of another aspect is inevitable. As an example, in the case study, we used velocity and position as the means of communication. These values are directly used in the belief update. However, if we would change the communication component of the model, the belief and its update also need to be changed. This is an important consideration when starting a redesign of the model since this could be the case for more components.

Finally, event-based triggering of the re-plan based on perceived risk results in an uneven computational requirement from the model: some time steps may take significantly more time to compute than others. A result of this is that our current implementation of the model cannot run in real-time. Instead, we used offline simulation for the case study. This could pose a problem when an experiment needs to be performed where the model interacts directly with a human.

Although the presented case study shows promising results, there is much future work to be done on the proposed framework. In addition to accounting for stochasticity in human behaviour and optimizing the runtime performance of the model, a necessary next step is to compare the model to human-human interactive behaviour. However, even validating single-driver models that do not incorporate interactions is already a complex task [51], therefore comparing our model to human-human interaction data requires a separate detailed investigation.

4.6. Conclusion

In this paper, we proposed a novel modelling framework to model human-human driving interactions. The key insight underlying this framework is the focus on the joint behaviour of the drivers during the interaction, rather than the isolated behaviour of a single driver. The framework explicitly includes communication be-

tween drivers and mutual influences (reciprocal interaction). We implemented the model for a simplified merging scenario and investigated its behaviour in four scenarios. We conclude the following:

- The model avoids impending collisions via plausible driver-driver interactive behaviours;
- Changing the risk threshold parameters per driver results in changes in behaviour that can be interpreted as more aggressive or conservative;
- Velocity-depended gap-keeping behaviour emerges from the combination of risk-based planning and a probabilistic belief about other drivers' plans. With this behaviour, the model shows a fundamental aspect of human driving behaviour, without it being explicitly programmed;
- The proposed model framework is a promising novel approach for modelling two-way multi-agent interactions in traffic.

Bibliography

- [1] J. Potzy, M. Feuerbach, and K. Bengler, "Communication strategies for automated merging in dense traffic", *IEEE Intelligent Vehicles Symposium, Proceedings*, vol. 2019-June, no. lv, pp. 2291–2298, 2019. DOI: [10.1109/IVS.2019.8813835](https://doi.org/10.1109/IVS.2019.8813835).
- [2] S. Ulbrich, S. Grossjohann, C. Appelt, K. Homeier, J. Rieken, and M. Maurer, "Structuring Cooperative Behavior Planning Implementations for Automated Driving", *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2015-October, pp. 2159–2165, 2015. DOI: [10.1109/ITSC.2015.349](https://doi.org/10.1109/ITSC.2015.349).
- [3] Y. M. Lee, R. Madigan, O. Giles, *et al.*, "Road users rarely use explicit communication when interacting in today's traffic: implications for automated vehicles", *Cognition, Technology & Work*, vol. 23, no. 2, pp. 367–380, May 2021, ISSN: 1435-5558. DOI: [10.1007/s10111-020-00635-y](https://doi.org/10.1007/s10111-020-00635-y). [Online]. Available: <https://doi.org/10.1007/s10111-020-00635-y> <https://link.springer.com/10.1007/s10111-020-00635-y>.
- [4] S. Kolekar, J. de Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model", *Nature Communications*, vol. 11, no. 1, p. 4850, Dec. 2020, ISSN: 2041-1723. DOI: [10.1038/s41467-020-18353-4](https://doi.org/10.1038/s41467-020-18353-4). [Online]. Available: <http://dx.doi.org/10.1038/s41467-020-18353-4> <https://www.nature.com/articles/s41467-020-18353-4>.
- [5] D. D. Salvucci and R. Gray, "A two-point visual control model of steering", *Perception*, vol. 33, no. 10, pp. 1233–1248, Oct. 2004, ISSN: 03010066. DOI: [10.1068/p5343](https://doi.org/10.1068/p5343). [Online]. Available: <http://journals.sagepub.com/doi/10.1068/p5343>.
- [6] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, vol. 62, no. 2, pp. 1805–1824, 2000, ISSN: 1063651X. DOI: [10.1103/PhysRevE.62.1805](https://doi.org/10.1103/PhysRevE.62.1805). arXiv: [0002177 \[cond-mat\]](https://arxiv.org/abs/0002177).
- [7] A. Kesting, M. Treiber, and D. 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*, vol. 368, no. 1928, pp. 4585–4605, 2010, ISSN: 1364503X. DOI: [10.1098/rsta.2010.0084](https://doi.org/10.1098/rsta.2010.0084). arXiv: [0912.3613](https://arxiv.org/abs/0912.3613).
- [8] M. Rahman, M. Chowdhury, Y. Xie, and Y. He, "Review of Microscopic Lane-Changing Models and Future Research Opportunities", *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1942–1956, Dec. 2013, ISSN: 1524-9050. DOI: [10.1109/TITS.2013.2272074](https://doi.org/10.1109/TITS.2013.2272074). [Online]. Available: <http://ieeexplore.ieee.org/document/6570532/>.
- [9] D. D. Salvucci and A. Liu, "The time course of a lane change: Driver control and eye-movement behavior", *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 5, no. 2, pp. 123–132, 2002, ISSN: 13698478. DOI: [10.1016/S1369-8478\(02\)00011-6](https://doi.org/10.1016/S1369-8478(02)00011-6).
- [10] A. Zgonnikov, D. Abbink, and G. Markkula, "Should i stay or should i go? cognitive modeling of left-turn gap acceptance decisions in human drivers", *Human Factors*, vol. 0, no. 0, p. 00187208221144561, 0, PMID: 36534014. DOI: [10.1177/00187208221144561](https://doi.org/10.1177/00187208221144561). eprint: <https://doi.org/10.1177/00187208221144561>. [Online]. Available: <https://doi.org/10.1177/00187208221144561>.
- [11] K. Tian, G. Markkula, C. Wei, *et al.*, "Explaining unsafe pedestrian road crossing behaviours using a Psychophysics-based gap acceptance model", *Safety Science*, vol. 154, no. 5, p. 105837, 2022, ISSN: 18791042. DOI: [10.1016/j.ssci.2022.105837](https://doi.org/10.1016/j.ssci.2022.105837). [Online]. Available: <https://doi.org/10.1016/j.ssci.2022.105837>.
- [12] H. Kita, "A merging-giveway interaction model of cars in a merging section: A game theoretic analysis", *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3-4, pp. 305–312, 1999, ISSN: 09658564. DOI: [10.1016/S0965-8564\(98\)00039-1](https://doi.org/10.1016/S0965-8564(98)00039-1).
- [13] H. Liu, W. Xin, Z. Adam, and J. Ban, "A game theoretical approach for modelling merging and yielding behaviour at freeway on-ramp section", *Transportation and Traffic Theory*, no. January, pp. 1–15, 2007. [Online]. Available: [http://www.ce.umn.edu/\\$%5Csim\\$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf](http://www.ce.umn.edu/$%5Csim$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf).


- [14] F. Meng, J. Su, C. Liu, and W.-H. Chen, "Dynamic decision making in lane change: Game theory with receding horizon", in *2016 UKACC 11th International Conference on Control (CONTROL)*, IEEE, Aug. 2016, pp. 1–6, ISBN: 978-1-4673-9891-6. DOI: [10.1109/CONTROL.2016.7737643](https://doi.org/10.1109/CONTROL.2016.7737643). [Online]. Available: <http://ieeexplore.ieee.org/document/7737643/>.
- [15] R. Tian, S. Li, N. Li, I. Kolmanovsky, A. Girard, and Y. Yildiz, "Adaptive Game-Theoretic Decision Making for Autonomous Vehicle Control at Roundabouts", in *2018 IEEE Conference on Decision and Control (CDC)*, vol. 2018-Decem, IEEE, Dec. 2018, pp. 321–326, ISBN: 978-1-5386-1395-5. DOI: [10.1109/CDC.2018.8619275](https://doi.org/10.1109/CDC.2018.8619275). arXiv: [1810.00829](https://arxiv.org/abs/1810.00829). [Online]. Available: <https://ieeexplore.ieee.org/document/8619275/>.
- [16] Q. Zhang, D. Filev, H. E. Tseng, S. Szwabowski, and R. Langari, "Addressing Mandatory Lane Change Problem with Game Theoretic Model Predictive Control and Fuzzy Markov Chain", in *2018 Annual American Control Conference (ACC)*, vol. 2018-June, IEEE, Jun. 2018, pp. 4764–4771, ISBN: 978-1-5386-5428-6. DOI: [10.23919/ACC.2018.8431530](https://doi.org/10.23919/ACC.2018.8431530). [Online]. Available: <https://ieeexplore.ieee.org/document/8431530/>.
- [17] S. Coskun, Q. Zhang, and R. Langari, "Receding Horizon Markov Game Autonomous Driving Strategy", in *2019 American Control Conference (ACC)*, vol. 2019-July, IEEE, Jul. 2019, pp. 1367–1374, ISBN: 978-1-5386-7926-5. DOI: [10.23919/ACC.2019.8815251](https://doi.org/10.23919/ACC.2019.8815251). [Online]. Available: <https://ieeexplore.ieee.org/document/8815251/>.
- [18] C. F. Camerer, *Behavioral game theory: Experiments in strategic interaction*. 2003, ISBN: 0691090394. DOI: [10.1016/j.socec.2003.10.009](https://doi.org/10.1016/j.socec.2003.10.009).
- [19] M. Schmidt-Daffy, "Prospect balancing theory: Bounded rationality of drivers' speed choice", *Accident Analysis and Prevention*, vol. 63, pp. 49–64, 2014, ISSN: 00014575. DOI: [10.1016/j.aap.2013.10.028](https://doi.org/10.1016/j.aap.2013.10.028). [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2013.10.028>.
- [20] A. H. Kalantari, Y. Yang, N. Merat, and G. Markkula, "Modelling vehicle-pedestrian interactions at unsignalised locations : Road users may not play the Nash equilibrium", *PsyArXiv*, pp. 1–34, 2023. DOI: [10.31234/osf.io/axseu](https://doi.org/10.31234/osf.io/axseu). [Online]. Available: <https://psyarxiv.com/axseu/>.
- [21] J. A. Michon, "A Critical View of Driver Behavior Models: What Do We Know, What Should We Do?", in *Human Behavior and Traffic Safety*, Boston, MA: Springer US, 1985, pp. 485–524, ISBN: 0306422255. DOI: [10.1007/978-1-4613-2173-6_19](https://doi.org/10.1007/978-1-4613-2173-6_19). [Online]. Available: http://link.springer.com/10.1007/978-1-4613-2173-6_19.
- [22] O. Siebinga, *CEI-model repository*, <https://github.com/tud-hri/cei-model>, 2022. DOI: [10.4121/21196357](https://doi.org/10.4121/21196357).
- [23] O. Siebinga, A. Zgonnikov, and D. A. Abbink, *Data underlying the publication: Modelling communication-enabled traffic interactions*, 2022. DOI: [10.4121/20749069](https://doi.org/10.4121/20749069).
- [24] S. Kolekar, J. de Winter, and D. Abbink, "Which parts of the road guide obstacle avoidance? Quantifying the driver's risk field", *Applied Ergonomics*, vol. 89, no. July, p. 103196, Nov. 2020, ISSN: 00036870. DOI: [10.1016/j.apergo.2020.103196](https://doi.org/10.1016/j.apergo.2020.103196). [Online]. Available: <https://doi.org/10.1016/j.apergo.2020.103196>
<https://linkinghub.elsevier.com/retrieve/pii/S0003687018307373>.
- [25] E. Jensen, M. Luster, H. Yoon, B. Pitts, and S. Sankaranarayanan, "Mathematical Models of Human Drivers Using Artificial Risk Fields", in *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2022-October, May 2022, pp. 909–916, ISBN: 9781665468800. DOI: [10.1109/ITSC55140.2022.9922389](https://doi.org/10.1109/ITSC55140.2022.9922389). arXiv: [2205.12722](https://arxiv.org/abs/2205.12722). [Online]. Available: <https://arxiv.org/abs/2205.12722v1>
<http://arxiv.org/abs/2205.12722>.
- [26] H. A. Simon, "A Behavioral Model of Rational Choice", *The Quarterly Journal of Economics*, vol. 69, no. 1, p. 99, 1955, ISSN: 00335533. DOI: [10.2307/1884852](https://doi.org/10.2307/1884852).
- [27] H. A. Simon, "Rational choice and the structure of the environment.", *Psychological Review*, vol. 63, no. 2, pp. 129–138, 1956, ISSN: 1939-1471. DOI: [10.1037/h0042769](https://doi.org/10.1037/h0042769). [Online]. Available: <http://doi.apa.org/getdoi.cfm?doi=10.1037/xge0000013%20http://doi.apa.org/getdoi.cfm?doi=10.1037/h0042769>.

- [28] L. M. Sacheli, E. Tidoni, E. F. Pavone, S. M. Aglioti, and M. Candidi, "Kinematics fingerprints of leader and follower role-taking during cooperative joint actions", *Experimental Brain Research*, vol. 226, no. 4, pp. 473–486, May 2013, ISSN: 0014-4819. DOI: [10.1007/s00221-013-3459-7](https://doi.org/10.1007/s00221-013-3459-7). [Online]. Available: <http://link.springer.com/10.1007/s00221-013-3459-7>.
- [29] G. Pezzulo, F. Donnarumma, and H. Dindo, "Human Sensorimotor Communication: A Theory of Signaling in Online Social Interactions", *PLoS ONE*, vol. 8, no. 11, J. Daunizeau, Ed., e79876, Nov. 2013, ISSN: 1932-6203. DOI: [10.1371/journal.pone.0079876](https://doi.org/10.1371/journal.pone.0079876). [Online]. Available: <https://dx.plos.org/10.1371/journal.pone.0079876>.
- [30] M. Dadvar, K. Majd, E. Oikonomou, G. Fainekos, and S. Srivastava, "Joint Communication and Motion Planning for Cobots", in *Proceedings - IEEE International Conference on Robotics and Automation*, Sep. 2022, pp. 4771–4777, ISBN: 9781728196817. DOI: [10.1109/ICRA46639.2022.9812261](https://doi.org/10.1109/ICRA46639.2022.9812261). arXiv: [2109.14004](https://arxiv.org/abs/2109.14004). [Online]. Available: <http://arxiv.org/abs/2109.14004>.
- [31] M. E. Bratman, *Intention, plans, and practical reason*. Cambridge, Mass: Harvard University Press, 1987, ISBN: 0674458192.
- [32] C. Baker, R. Saxe, J. Tenenbaum, and C. L. Baker, "Bayesian Theory of Mind: Modeling Joint Belief-Desire Attribution Publication Date Bayesian Theory of Mind: Modeling Joint Belief-Desire Attribution", *Proceedings of the annual meeting of the cognitive science society*, no. 33, 2011, ISSN: 1069-7977.
- [33] W. Daamen, M. Loot, and S. P. Hoogendoorn, "Empirical analysis of merging behavior at freeway on-ramp", *Transportation Research Record*, no. 2188, pp. 108–118, 2010, ISSN: 03611981. DOI: [10.3141/2188-12](https://doi.org/10.3141/2188-12).
- [34] M. Naumann, L. Sun, W. Zhan, and M. Tomizuka, "Analyzing the Suitability of Cost Functions for Explaining and Imitating Human Driving Behavior based on Inverse Reinforcement Learning", in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2020, pp. 5481–5487, ISBN: 978-1-7281-7395-5. DOI: [10.1109/ICRA40945.2020.9196795](https://doi.org/10.1109/ICRA40945.2020.9196795). [Online]. Available: <https://ieeexplore.ieee.org/document/9196795/>.
- [35] D. Sadigh, N. Landolfi, S. S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state", *Autonomous Robots*, vol. 42, no. 7, pp. 1405–1426, Oct. 2018, ISSN: 0929-5593. DOI: [10.1007/s10514-018-9746-1](https://doi.org/10.1007/s10514-018-9746-1). [Online]. Available: <https://doi.org/10.1007/s10514-018-9746-1> <http://link.springer.com/10.1007/s10514-018-9746-1>.
- [36] Y. Xu, S. Bao, and A. K. Pradhan, "Modeling drivers' reaction when being tailgated: A Random Forests Method", *Journal of Safety Research*, vol. 78, pp. 28–35, 2021, ISSN: 00224375. DOI: [10.1016/j.jsr.2021.05.004](https://doi.org/10.1016/j.jsr.2021.05.004). [Online]. Available: <https://doi.org/10.1016/j.jsr.2021.05.004>.
- [37] T. Kondoh, T. Yamamura, S. Kitazaki, N. Kuge, and E. R. Boer, "Identification of Visual Cues and Quantification of Drivers' Perception of Proximity Risk to the Lead Vehicle in Car-Following Situations", *Journal of Mechanical Systems for Transportation and Logistics*, vol. 1, no. 2, pp. 170–180, 2008, ISSN: 1882-1782. DOI: [10.1299/jmtl.1.170](https://doi.org/10.1299/jmtl.1.170). [Online]. Available: http://www.jstage.jst.go.jp/article/jmtl/1/2/1%5C_2%5C_170/%5C_article.
- [38] J. Piao and M. McDonald, "Low speed car following behaviour from floating vehicle data", *IEEE Intelligent Vehicles Symposium, Proceedings*, pp. 462–467, 2003. DOI: [10.1109/IVS.2003.1212955](https://doi.org/10.1109/IVS.2003.1212955).
- [39] A. Ji and D. Levinson, "A review of game theory models of lane changing", *Transportmetrica A: Transport Science*, vol. 9935, no. May, pp. 1–19, 2020, ISSN: 2324-9935. DOI: [10.1080/23249935.2020.1770368](https://doi.org/10.1080/23249935.2020.1770368). [Online]. Available: <https://doi.org/10.1080/23249935.2020.1770368>.
- [40] N. Li, D. W. Oyler, M. Zhang, Y. Yildiz, I. Kolmanovsky, and A. R. Girard, "Game theoretic modeling of driver and vehicle interactions for verification and validation of autonomous vehicle control systems", *IEEE Transactions on Control Systems Technology*, vol. 26, no. 5, pp. 1782–1797, 2018, ISSN: 10636536. DOI: [10.1109/TCST.2017.2723574](https://doi.org/10.1109/TCST.2017.2723574). arXiv: [1608.08589](https://arxiv.org/abs/1608.08589).

- [41] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [42] L. de Silva, F. Meneguzzi, and B. Logan, "BDI agent architectures: A survey", *IJCAI International Joint Conference on Artificial Intelligence*, vol. 2021-January, pp. 4914–4921, 2020, ISSN: 10450823. DOI: [10.24963/ijcai.2020/684](https://doi.org/10.24963/ijcai.2020/684).
- [43] D. Premack and G. Woodruff, "Premack and Woodruff : Chimpanzee theory of mind", *Behavioral and Brain Sciences*, vol. 4, pp. 515–526, 1978.
- [44] B. Scassellati, "Theory of mind for a humanoid robot", *Autonomous Robots*, vol. 12, no. 1, pp. 13–24, 2002, ISSN: 09295593. DOI: [10.1023/A:1013298507114](https://doi.org/10.1023/A:1013298507114).
- [45] S. Devin and R. Alami, "An implemented theory of mind to improve human-robot shared plans execution", *ACM/IEEE International Conference on Human-Robot Interaction*, vol. 2016-April, pp. 319–326, 2016, ISSN: 21672148. DOI: [10.1109/HRI.2016.7451768](https://doi.org/10.1109/HRI.2016.7451768).
- [46] T. Nishi, P. Doshi, and D. Prokhorov, "Merging in congested freeway traffic using multipolicy decision making and passive actor-critic learning", *IEEE Transactions on Intelligent Vehicles*, vol. 4, no. 2, pp. 287–297, 2019, ISSN: 23798858. DOI: [10.1109/TIV.2019.2904417](https://doi.org/10.1109/TIV.2019.2904417).
- [47] H. Yu, H. E. Tseng, and R. Langari, "A human-like game theory-based controller for automatic lane changing", *Transportation Research Part C: Emerging Technologies*, vol. 88, no. October 2017, pp. 140–158, 2018, ISSN: 0968090X. DOI: [10.1016/j.trc.2018.01.016](https://doi.org/10.1016/j.trc.2018.01.016). [Online]. Available: <https://doi.org/10.1016/j.trc.2018.01.016>.
- [48] G. Markkula, Y.-s. Lin, A. R. Srinivasan, *et al.*, "Explaining human interactions on the road requires large-scale integration of psychological theory", *PsyArXiv*, 2022.
- [49] X. Wan, P. J. Jin, F. Yang, J. Zhang, and B. Ran, "Modeling Vehicle Interactions during Merge in Congested Weaving Section of Freeway Ramp", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2421, no. 1, pp. 82–92, 2014, ISSN: 0361-1981. DOI: [10.3141/2421-10](https://doi.org/10.3141/2421-10).
- [50] V. Kurtc, "Studying Car-Following Dynamics on the Basis of the HighD Dataset", *Transportation Research Record*, vol. 2674, no. 8, pp. 813–822, 2020, ISSN: 21694052. DOI: [10.1177/0361198120925063](https://doi.org/10.1177/0361198120925063).
- [51] O. Siebinga, A. Zgonnikov, and D. Abbink, "A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles", *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 4, pp. 1–21, Dec. 2022, ISSN: 2573-9522. DOI: [10.1145/3538705](https://doi.org/10.1145/3538705). [Online]. Available: <https://dl.acm.org/doi/10.1145/3538705>.

5

Interactive merging behaviour
in a coupled driving simulator:
Experimental framework and
case study



Human highway-merging behaviour is an important aspect when developing autonomous vehicles (AVs) that can safely and successfully interact with other road users. To design safe and acceptable human-AV interactions, the underlying mechanisms in human-human interactive behaviour need to be understood. Exposing and understanding these mechanisms can be done using controlled driving simulator experiments. However, until now, such human-factors merging experiments have focused on aspects of the behaviour of a single driver (e.g., gap acceptance) instead of on the dynamics of the interaction. Furthermore, existing experimental scenarios and data analysis tools (i.e., concepts like time-to-collision) are insufficient to analyze human-human interactive merging behaviour. To help facilitate human-factors research on merging interactions, we propose an experimental framework consisting of a general simplified merging scenario and a set of three analysis tools: (1) a visual representation that captures the combined behaviour of two participants and the safety margins they maintain in a single plot; (2) a signal (over time) that describes the level of conflict; and (3) a metric that describes the amount of time that was required to solve the merging conflict, called the conflict resolution time. In a case study with 18 participants, we used the proposed framework and analysis tools in a top-down view driving simulator where two human participants can interact. The results show that the proposed scenario can expose diverse behaviours for different conditions. We demonstrate that our novel visual representation, conflict resolution time, and conflict signal are valuable tools when comparing human behaviour between conditions. Therefore, with its simplified merging scenario and analysis tools, the proposed experimental framework can be a valuable asset when developing driver models that describe interactive merging behaviour and when designing AVs that interact with humans.

5.1. Introduction

One of the main present-day challenges in the development of autonomous vehicles (AVs) is enabling them to interact with human-driven vehicles safely, efficiently, and in a manner acceptable for occupants and other road users. To reach this goal, a deep understanding of human behaviour in interactive scenarios is required. One example of such a scenario is merging on a highway.

While many studies have been performed to understand and model human behaviour in non-interactive scenarios (e.g., car following [1] or lane changing [2]), there are only a limited number of studies concerning interactive merging behaviour. Some of these studies simulate behaviour and decision-making to investigate interactions. Mostly using game theory to investigate higher-level traffic phenomena (e.g., [3]). Because the behaviour is simulated and not recorded, these studies do not provide insight into the dynamics of driving interactions. Other studies investigate merging using naturalistic traffic datasets (e.g., [4]). These have the disadvantage that they cannot capture the variability within (and between) pairs of interacting vehicles, because there are no (known) repeated trials per pair. Thus, when aiming to understand the dynamics and variability of interactive driving behaviour, controlled human factors experiments are needed.

Existing human-factors experiments on merging behaviour, have so far focused mainly on the behaviour of one of the participants of the interaction while using generated behaviour for other traffic (e.g., [5]). But to fully understand the interaction, the joint behaviour of both drivers and their mutual influence should be considered. Therefore, controlled experiments are needed that are designed specifically to expose these underlying mechanisms of interactive behaviour.

For such an experiment, a merging scenario should be designed that evokes human interactive behaviour and that can be repeated and analyzed systematically. Additionally, meaningful signals and metrics should be defined that provide insight into how the merging conflict is resolved by a pair of merging drivers (e.g., comparable to time-to-collision for car following). Because experiments thus far focused on the behaviour of single drivers, both this simplified merging scenario and these meaningful signals and metrics are lacking.

This work addresses that gap, by proposing an experimental framework for human-factors experiments in coupled driving simulators. We propose a simplified merging scenario that reduces the action space compared to natural merging scenarios. This enables a first principled analysis for which we propose three novel tools: 1) a visual representation of the pair-wise behaviour, 2) a signal describing the time-varying level of conflict, and 3) a novel metric named conflict resolution time (CRT). The three analysis tools provide insight into the combined behaviour of two drivers and how they solve the merging conflict. In a case study, we show the practical applicability of the experimental framework and outline its potential for measuring and modelling human-human merging interactions.

5.2. Merging scenario

We consider a simplified merging scenario as illustrated in Figure 5.1. In this scenario, two vehicles of equal dimensions approach a predefined merge point. The acceleration of each vehicle can be controlled by a participant, who is instructed to maintain their initial velocity, but prevent collisions. No steering is considered: the headings of the vehicles are always equal to the heading of the road, and at the merge point, the headings of the vehicles change instantaneously.




Figure 5.1: The proposed simplified merging scenario: a screenshot of the top-down view driving simulator used in the case study, rotated 90 degrees clockwise. The left and the right vehicle approach a merge point where their lanes merge into one instantaneously. Trees and roadside markers give participants a visual cue for velocity. Each section of the track has a track length of 50 meters.

These simplifications reduce the vehicles' action and position spaces to single dimensions (longitudinal velocity and travelled distance along the track respectively). Note that with two participants, the combined state of the two vehicles and their possible actions are both two-dimensional. Besides the simplifications of action and positions, the proposed merging scenario also simplifies environmental factors. In this scenario, there is no right of way, therefore, the gathered data is symmetrical.

The proposed merging track consists of three sections of equal track length (50 m): the tunnel, the approach (capturing the actual merging behaviour), and a subsequent car-following section. In the tunnel section, participants cannot control their vehicle but can only observe the two vehicles, they gain control once both vehicles have exited the tunnel. This moment marks the explicit start of the interaction.

5.3. Case study

To illustrate the utility of our simplified merging scenario, we performed a case study. All obtained results combined with material to further detail the experimental protocol and methods can be found online [6]. The software used in the experiment can be found on GitHub [7].

Eighteen drivers volunteered to participate in our experiment (6 female, 12 male, mean age: 25, std: 2.6), they were divided into 9 pairs of two. The participants were seated at separate tables and divided by a black screen to prevent them from seeing each other (Figure 5.2). Participants were instructed to remain seated, use one foot on the gas or brake pedal, keep both hands on the steering wheel (which was only used for feedback, not for steering), and to avoid making sounds. Finally, participants were told that this is a scientific experiment, not a game or a race, and that no vehicle had the right of way. When the vehicles in the simulation collided, the participants got a time penalty of 20 seconds.

To investigate the impact of initial conditions on merging behaviour, the experiment consisted of 11 experimental conditions in which two variables were varied: the initial relative velocity of the vehicles (the average velocity was always 10 m/s), and the projected headway when the first vehicle reaches the merge point (assuming that the vehicles maintain their velocity). Both are defined to be positive when the value for the left vehicle is larger. The conditions were labelled with

a combination of numbers. First, the projected headway at the merge point in meters (-4, -2, 0, 2, or 4 meters), and second, the relative velocity of the vehicles (-8, 0, or 8 decimeters/second). The 11 conditions used were (-4_-8), (-4_0), (-4_8), (-2_-8), (0_8) (0_0) (0_-8), (2_-8), (4_-8), (4_0), and (4_8). Each condition was repeated 10 times in random order for every pair of participants. Five additional trials from random conditions were used at the start of the experiment as training runs.

5.4. Results and analysis tools

Besides the simplified merging scenario, we propose three analysis tools that provide insight into interactive behaviour: 1) a visual representation of the pair-wise behaviour, 2) a signal describing the time-varying level of conflict, and 3) a novel metric named conflict resolution time (CRT). Using these tools, an overview figure that provides insight into the conflict resolution behaviour of the pair of participants was made for every trial. Figure 5.3 shows a representative example of this overview, figures for all other trials can be found online [6]). We will use the example in Figure 5.3 to introduce our proposed analysis tools in the following sections.



Figure 5.2: The experimental setup as seen from the view of one participant. The participant sees a top-down view of the vehicle they can control with the gas and brake pedal. The steering wheel provides vibration feedback if the participant deviates from the designated velocity. Visual velocity feedback is provided through the speed dial in the lower part of the screen.

average travelled distance of the two vehicles (we expand on this later in this section). The resulting trajectory in the “headway – average distance” plane represents the joint dynamics of the two vehicles (Figure 5.4).

In Figure 5.3, we first illustrated the positions and velocities of the individual participants (panels A and B). Both these plots illustrate the dynamics of individual vehicle motion but provide little insight into how and when the conflict was resolved, besides the fact that the right vehicle reached the merge point first. Therefore, we propose to use the visualization in panel C of Figure 5.3.

5.4.1. Interaction visualization during merging and car-following

A meaningful visualization of a traffic interaction should capture both the state and possible actions of the involved vehicles, as well as the safety margins. Consider car following for example. The commonly used plot of the distance gap over time shows the current state (gap), action (the slope of the trajectory represents the relative velocity), and safety margin (if the gap is 0, a collision occurs). We propose to extend this gap-plot in two ways to make it applicable to the interactive merging scenario.

First, we represent the behaviour of the two drivers by plotting the headway; the distance between the front bumpers of the vehicles. We define the headway as positive if the left vehicle is ahead. And second, instead of against time, we plot the headway against the

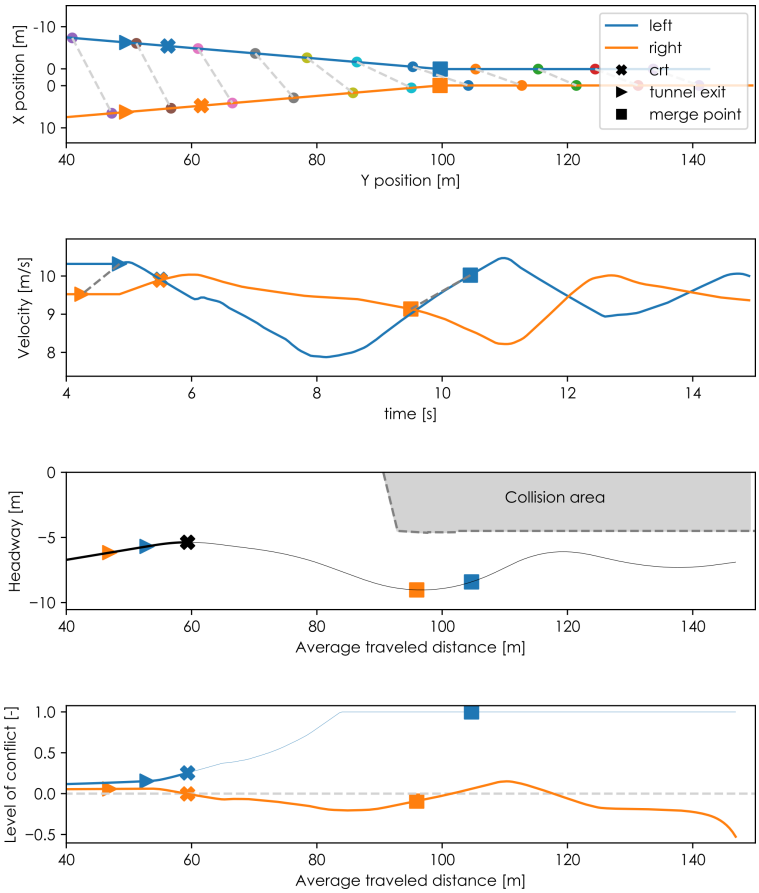


Figure 5.3: Dynamics of a representative case study trial (pair 1, trial 1, condition -2_8). Panel A shows individual positions, the connected markers represent positions at the same points in time showing the left car exits the tunnel later, and arrives later at the merge point. Panel B shows individual velocities over time. Panels C and D show the “headway – average distance” trajectory and the level-of-conflict signal, respectively. These will be introduced in the following sections. The conflict resolution time (CRT) is indicated by the crosses, this metric will be explained in the second last section. Finally, triangles indicate the moment when a vehicle exits the tunnel and squares denote the moment a vehicle reached the merge point. The average traveled distance is average position (distance along the track) of the two vehicles.

In the car-following gap plot, the imminent collision is indicated simply by the gap approaching zero (or the headway approaching the length of the vehicle), in our interactive merging scenario no collision can occur during the approach, yet a conflict can still be present (i.e. the vehicles are heading towards a collision). To visualize this approach conflict, we define the collision area in the “headway – average distance” plane (grey area in Figure 5.4). This area is a block during the car following section and does not exist during the approach. The shape and size of it at the merge point depend on the dimensions of the track and the vehicles (the dark part in Figure 5.4). If the joint trajectory of the vehicles goes inside the collision area a collision occurs. The exact boundaries of the area can




Figure 5.4: The proposed visualization of the interactive behavior, this in an extended view of Figure 5.3 panel C. The darker part of the collision area is influenced by the merge point, the light part represents the car-following part. The joint trajectory in the "headway – average distance" plane shows the dynamics of how the conflict is resolved and when it is resolved (when the line no longer heads for the collision area). Additionally, the safety margin is shown as the headway between the trajectory and the collision area (this is the gap).

5

be calculated by minimizing the headway while constraining the overlap between the vehicles to be 0, for all average travelled distances.

We opted for visualizing the headway against the average distance and not time because the existence of the described collision area depends on the section of the track the vehicles are in. Therefore, representing it over time makes it dependent on the vehicles' positions at a certain time. This would result in a different visual representation of the collision area for every repeated trial of the experiment. Thus, it would make it impossible to compare trials in a single plot. Conversely, plotting the headway over the average distance along the track anchors the collision area (and track sections) and enables the visual comparison of experiment trials.

In the context of the overview figure with case study results (Figure 5.3-C), our proposed visual representation provides additional insight compared to the position and velocity plots (panels A and B). The "headway-average distance" trajectory shows the initial conditions (i.e., the situation at the tunnel exit), the chosen solution (the trajectory bends down so the right vehicle went first), and the safety margin (the gap between the vehicles at the merge point is observable as the vertical distance between the line and the collision block: approximately 4 m).

5.4.2. Level-of-conflict signal

The visual representation of the interaction also inspired us to propose a signal describing the level of conflict (Figure 5.3-D). It quantifies the amount of effort needed to resolve the conflict as well as the safety margins after the conflict is resolved. The signal is calculated using three checkpoints along the track (illustrated as orange lines in Figure 5.5). These checkpoints are situated at the end of the track, at the merge point, and at the collision threshold (the first position on the track where a collision can occur). This collision threshold is located before the merge point because there is a possibility of a side-to-side collision on the approach.

There are always two solutions to the merge conflict: either the left vehicle merges first or the right one does. At every checkpoint, the minimum safe headway can be determined for each solution (this is positive when left merges first and vice versa). These minimum safe headways at the checkpoints can be represented as points on the boundary of the collision area in the "headway – average distance" plane (orange markers in Figure 5.5).

For every point on the headway trajectory, we can now calculate the angle between the current slope of the trajectory (blue arrow in Figure 5.5) and vectors to-




Figure 5.5: The construction of the level-of-conflict signals from the “headway – average distance” plane representation. The blue marker and arrow represent a point and slope of the headway trajectory. The orange lines denote the checkpoints along the track, orange markers indicate the minimum safe headway there, for both the left and right-first solutions. The green lines represent the minimum and maximum slope needed to clear the collision area. They thus represent the deviation needed from the current trajectory, to safely reach the left-first (top) or right-first (bottom) solution.

wards all 6 boundary points (orange markers in Figure 5.5) (for the implementation, see [7]). The minimum and maximum values of these 6 angles can be visualized as lines from the current point on the trajectory to a boundary point, which precisely clears the collision area (green lines in Figure 5.5).

Note that the slope of the headway trajectory is directly related to the relative velocity of the vehicles, these angles thus represent changes in relative velocity. Therefore, they quantify the deviation from the current relative velocity that is required to safely reach one of the two possible solutions (either the left goes first and the trajectory goes above the collision area, or vice versa).

The minimum and maximum angles for every point on the trajectory make up two conflict signals, one for the left-first solution (maximum) and one for the right-first solution (minimum). These signals are inversely related: if one solution becomes easier, the other becomes harder. Right-first solutions always require a clockwise rotation of the slope because the trajectory will go under the collision area. Therefore, we multiply the right-first (minimum) angles with -1. This ensures that the conflict signal is always positive if the vehicles are currently on a collision course, and negative if the conflict is resolved. Finally, we normalize the calculated angles by dividing them by the maximum possible angle. The headway can maximally increase with a factor of 2 over the average travelled distance thus the maximum absolute angle is $\text{atan}(2)$. The result is a conflict signal that ranges from -1 (maximum safety margin) to 1 (maximum conflict).

In the context of the overview figure with case study results (Figure 5.3-D), the proposed signals illustrate the conflict dynamics to a larger extent than the proposed “headway – average distance” representation (panel C) alone. The level of conflict for the right-first solution is initially lower, suggesting that it is easier for the two participants to resolve the conflict in a way where the right vehicle goes first. This is indeed what happened in this trial: the conflict was resolved when the level of conflict for the right-first solution reached 0. When the left-first conflict signal reached 1, that solution was not reachable anymore.

Furthermore, the conflict signals highlight what happened during the rest of the experiment. In the car following section, the right-first conflict signal becomes positive indicating that the vehicles are on a collision course again. This can also be seen in the headway trajectory, which is heading towards the collision area at that point. While this would have been visible in a time-to-collision plot, it is not clear in the raw position and velocity data (panels A and B).

5.4.3. Conflict resolution time

Besides having visual representations of a complete trial of the experiment, it is vital to have a metric that captures the dynamics of the conflict resolution in a single number. Such a metric can be used to compare different experimental conditions and human pairs. We propose to use the conflict resolution time (CRT) as this metric. We define CRT as the time between the start of the interaction (tunnel exit) and the first moment the vehicles are no longer on a collision course (assuming constant velocities).




Figure 5.6: the conflict resolution time (CRT) values for all trials of the experiment, presented per condition.

To calculate the CRT, we use the same three checkpoints along the track that were used for the conflict signal (the orange lines in Figure 5.5). For every time step, we calculate if continuing at the current velocity will result in a collision at any of the checkpoints (for the implementation, see [7]). At the first point in time where no collision would occur, we assume the conflict is resolved.

To investigate differences between the conditions, we analyzed CRT over all participants (Figure 5.6). We found that in some conditions, the conflict was solved faster than in others, which can be interpreted as the conflict being easier to solve. The condition 0_0 for example, where neither vehicle had an advantage in terms of projected headway or velocity, had the highest median CRT. This decreased for conditions where one of the vehicles did have an advantage. The CRT was lowest for the conditions where the vehicles had equal velocity, but where one had a projected headway advantage.

5.5. Discussion and conclusion

In this work, we proposed an experimental framework for investigating the merging behaviour of a pair of human drivers in a coupled driving simulator. Additionally, we proposed three novel analytic tools to quantify essential characteristics of merging behaviour: the “headway – average distance” trajectory, the level-of-conflict signals, and the conflict resolution time (CRT). The results of our case study show that our proposed visual representation and level-of-conflict signals provide additional insight in individual trials compared to basic velocity and position traces. Our proposed CRT metric can be used to expose aggregate differ-

ences between conditions. Together, these analysis tools can help to meaningfully compare joint human-human behaviour in merging interactions between trials and between conditions.

There are three main limitations to the proposed framework and the setup of our case study. One major limitation is that it is currently unknown how human behaviour in this simplified scenario exactly relates to natural merging behaviour. Simplifying the scenario in a controlled environment is a necessary first step for obtaining insight into the merging behaviour. However, future work should focus on extending the controlled environment to a more natural 3-dimensional space that includes environmental factors such as right-of-way.

Second, the proposed analysis tools are not yet suitable for use in a 3-dimensional environment because they are all related to the dimensions of the track and the vehicles. If steering control is added to the vehicles, and the merge point is converted to a merge line, the fixed collision area in the interaction “headway – average distance” plane is no longer valid. Thus, how to extend the proposed analysis tools to a 3-dimensional environment remains an open question.

The final limitation is specific to our case study. In the case study, the same pair of participants performed 110 trials of the experiment. Participants were aware that these trials all involved the same opponent. This could have led to participants anticipating the driving style of the other driver after several trials, something which is not possible in natural driving. Future studies could include multiple participants at the same time, with random pairing of participants at the start of each trial to account for this.

Despite these limitations, we believe our experimental framework can be a valuable asset in future studies of human-human interactive merging behaviour. It can therefore support the development and validation of human behaviour models, advanced driver assistance systems, and autonomous vehicles that interact with humans.

Bibliography

- [1] Y. Li and D. Sun, "Microscopic car-following model for the traffic flow: The state of the art", *Journal of Control Theory and Applications*, vol. 10, no. 2, pp. 133–143, 2012, ISSN: 0974-5572. DOI: [10.1007/s11768-012-9221-z](https://doi.org/10.1007/s11768-012-9221-z).
- [2] M. Rahman, M. Chowdhury, Y. Xie, and Y. He, "Review of Microscopic Lane-Changing Models and Future Research Opportunities", *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1942–1956, Dec. 2013, ISSN: 1524-9050. DOI: [10.1109/TITS.2013.2272074](https://doi.org/10.1109/TITS.2013.2272074). [Online]. Available: <http://ieeexplore.ieee.org/document/6570532/>.
- [3] H. Kita, "A merging-giveway interaction model of cars in a merging section: A game theoretic analysis", *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3-4, pp. 305–312, 1999, ISSN: 09658564. DOI: [10.1016/s0965-8564\(98\)00039-1](https://doi.org/10.1016/s0965-8564(98)00039-1).
- [4] H. Wang, W. Wang, S. Yuan, and X. Li, "Uncovering Interpretable Internal States of Merging Tasks at Highway On-Ramps for Autonomous Driving Decision-Making", pp. 1–12, 2021. arXiv: [2102.07530](https://arxiv.org/abs/2102.07530). [Online]. Available: <http://arxiv.org/abs/2102.07530>.
- [5] A. Calvi and M. R. De Blasiis, "Driver Behavior on Acceleration Lanes", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2248, no. 1, pp. 96–103, Jan. 2011, ISSN: 0361-1981. DOI: [10.3141/2248-13](https://doi.org/10.3141/2248-13). [Online]. Available: <http://journals.sagepub.com/doi/10.3141/2248-13>.
- [6] O. Siebinga, A. Zgonnikov, and D. A. (Abbink, "Data underlying the publication: Interactive merging behavior in a coupled driving simulator: Experimental framework and case study", 2022. DOI: [10.4121/19550377.v1](https://doi.org/10.4121/19550377.v1). [Online]. Available: https://data.4tu.nl/articles/dataset/Data_underlying_the_publication_Interactive_merging_behavior_in_a_coupled_driving_simulator_Experimental_framework_and_case_study/19550377.
- [7] O. Siebinga, *simple-merging-experiment [Software]*, <https://github.com/tud-hri/simple-merging-experiment>, 2022.

6

Human merging behaviour in a
coupled driving simulator:
How do we resolve conflicts?

Traffic interactions between merging and highway vehicles are a major topic of research, yielding many empirical studies and models of driver behaviour. Most of these studies on merging use naturalistic data. Although this provides insight into human gap acceptance and traffic flow effects, it obscures the operational inputs of interacting drivers. Besides that, researchers have no control over the vehicle kinematics (i.e., positions and velocities) at the start of the interactions. Therefore the relationship between initial kinematics and the outcome of the interaction is difficult to investigate. To address these gaps, we conducted an experiment in a coupled driving simulator with a simplified, top-down view, merging scenario with two vehicles. We found that kinematics can explain the outcome (i.e., which driver merges first) and the duration of the merging conflict. Furthermore, our results show that drivers use key decision moments combined with constant acceleration inputs (intermittent piecewise-constant control) during merging. This indicates that they do not continuously optimise their expected utility. Therefore, these results advocate the development of interaction models based on intermittent piecewise-constant control. We hope our work can contribute to this development and to the fundamental knowledge of interactive driver behaviour.

6.1. Introduction

Interactions between vehicles, such as in highway merging, play a major role in everyday traffic. Therefore, driving behaviour in these interactions is an essential aspect of many transportation technologies. Empirical data and microscopic traffic models of human driving behaviour are thus essential tools for transportation engineers. These models and data are used in the design and safety assessment of highway on-ramps [1], [2] and urban intersections [3]. Microscopic traffic models can be used to evaluate traffic management systems [4]. And finally, autonomous vehicle designers are interested in these interactions to develop socially acceptable and human-like autonomous behaviour [5], [6]. Particularly for the last use case, a good understanding of the individual negotiations and the continuous reciprocal actions of the drivers during interactions is essential.

Many recent studies have investigated interactive merging behaviour by modelling this behaviour or by conducting empirical investigations. Most of these studies use naturalistic data, i.e., data recorded in real-world scenarios. For example, Daamen et al. [7] and Marczak et al. [8] performed empirical analysis on traffic data which they recorded with helicopters. Wang et al. [9] and Srinivasan et al. [10] used existing open datasets to evaluate driver behaviour on merge ramps. Others have modelled interactive driver behaviour using naturalistic data to gain insights, e.g., using game theory [11]–[13], acceleration models comparable to car-following models [14], or machine-learned models [10], [15].




Figure 6.1: The simplified merging scenario used in the experiment. Two vehicles approach a pre-defined merge point at which their lanes merge seen from a participant's into one. The track consists of three sections of equal length (50 m, total view. The other participants could observe both vehicles, but not control their vehicles. During used an identical setup. the approach, the participants could control the acceleration of their vehicles The participants could not to resolve the merging conflict. During the car-following section, the vehicles see each other. follow each other in the same lane.

Figure 6.2: The experimental setup as used in the experiment. The other participants could observe both vehicles, but not control their vehicles. During used an identical setup. the participants could not to resolve the merging conflict. During the car-following section, the vehicles see each other.

The usage of naturalistic data has the advantage that real-world behaviour can be studied. However, this approach has two main drawbacks. First, it is challenging to investigate the interacting drivers' operational behaviour and control inputs from naturalistic data. The previously mentioned studies use naturalistic data that was recorded with cameras on helicopters, quad-copters, or high buildings. Therefore, only sequential positions are recorded. Velocities and accelerations are reconstructed from this position data and control inputs are not included. Other naturalistic datasets that are recorded from within a vehicle (e.g., [16]–[18]). They do contain these signals for the ego vehicle. However, these datasets do not provide the same signals for the surrounding vehicles, which complicates the study of interactions.

The second drawback is that although the kinematic differences between situations can be observed, they cannot be controlled. This makes it difficult to investigate the relationship between the initial kinematics of the vehicles and the outcome of the merging conflict (e.g., who merges first and who yields). To gain a deeper understanding of individual reciprocal interactions, controlled experiments are needed.

However, only a very limited number of studies in a controlled environment (i.e., in a driving simulator) targeted interactions during merging (i.e., excluding studies of autonomous control strategies, gap acceptance, or traffic flow). Stoll et al. investigated human decision-making in merging scenarios based on videos of a controlled simulation [19]. Participants had to select their preferred reaction (e.g., accelerate or decelerate) after watching videos of vehicles they were "interacting with". Shimojo et al. used a driving simulator to investigate how the merging behaviour of drivers is affected by their perception of other drivers [20]. They used predetermined controls for one of the vehicles in the interaction, to influence this perception in a controlled way. In both experiments, the behaviour of one of the drivers was predetermined. Thus, there was no interaction or dynamic negotiation between two human drivers. We conclude that the existing literature misses studies that investigate the reciprocal merging interactions between at least two human drivers in a controlled environment.

To address this gap, we conduct an experiment in a top-down view, coupled driving simulator in which we investigate reciprocal merging interactions between two human drivers. We investigate the operational behaviour of the drivers in terms of inputs (acceleration and velocity profiles). Furthermore, we examine the influence of different initial kinematics (both position and velocity) on the outcome of the interaction. Both on a high level in terms of which driver merges first, and in more detail through the metric Conflict Resolution Time (CRT) [21]. The focus of our work is on the dynamics of interactive behaviour. We hope this experiment advances the fundamental knowledge about vehicle-vehicle interactions in traffic and contributes to the development of interaction-aware intelligent transportation systems.

6.2. Methods

We conducted an experiment in a coupled, top-down view driving simulator with 9 pairs of participants (6 female, 12 male, mean age: 25, std: 2.6). All participants met their "opponent" before the experiment and most participant pairs knew each other before the experiment. The details of this experiment (including Figures 6.1 and 6.2), and the analysis tools we developed to gain insight into the merging behaviour, have been previously published in [21]. This experiment was approved by TU Delft's Human Research Ethics Committee (HREC). All participants gave their consent before participating in the experiment.

The experiment regarded a symmetric simplified merging scenario (Figure 6.1) in which participants could control the acceleration of their vehicle using the gas and brake pedal of a steering-wheel game controller (Logitech Driving Force GT). The headings of the vehicles were always equal to the heading of the road, so no steering was involved. Participants could see the simulation on a computer screen (Figure 6.2). However, they could not see the other participant, who was seated in the same room behind a curtain. To prevent auditory communication, participants wore noise-cancelling headsets (Sony WH-1000XM3) with ambient music. All gathered data, the information letter we provided to participants before the

experiment, and the informed consent form we used were published in the 4TU data repository [22]. The software needed to reproduce the experiment can be found on GitHub¹. Interactive plots of all our results can be found in the online supplementary materials².

To investigate the effects of the initial vehicle kinematics on the outcome of the merging conflict we varied the initial positions and initial velocities of the vehicles. Participants were instructed to "maintain their initial velocity yet prevent a collision". To ensure a merging conflict, all conditions were chosen such that if both drivers would maintain their initial velocity, they would collide. Furthermore, participants were instructed to "remain seated, use one foot on the gas or brake pedal, keep both hands on the steering wheel, and not to communicate by making sounds or noise." Finally, participants were told that "this is a scientific experiment, not a game or a race" and that "no vehicle has the right of way."

The participants received visual feedback on their computer screens. Their visuals were randomly mirrored such that they appeared to approach the merge point from the left or the right side randomly. While in the experimenter's view, and in all results discussed here, we refer to the same participant in a pair as the left or right driver. If participants deviated from their initial velocity, their steering wheel provided vibration feedback, increasing with the deviation and with a dead band around the initial velocity. The vibration was implemented to facilitate speed perception. If the vehicles collided, the participants got a time penalty of 20 seconds. This was longer than the duration of a single trial, which took approximately 16 seconds. During this time, the experiment was paused and the participants had to wait and watch an animation on the screen. This increased the total duration of the experiment and therefore provided an incentive not to collide.

The vehicles started in a tunnel where participants could observe the initial velocities of both vehicles, but they could not control their vehicles yet. The drivers gained control when both vehicles exited the tunnel. The tunnel, with its different background colour, served merely as a visual representation of the possibility of controlling the vehicles. This section of the track had two purposes. First, it ensured that drivers could perceive their velocity relative to the velocity of the other vehicle before starting the interaction. In a pilot study, we tested a setup where a driver could only see their own vehicle in the tunnel. However, in this pilot, drivers accelerated directly at the tunnel exit to anticipate the appearance of another vehicle, which is a unilateral decision and thus prevents an interaction. Therefore, we decided to make both vehicles visible in the tunnel. Second, the tunnel exit marked an unambiguous moment when the interaction started (i.e., the start of the interaction).

The vehicles' initial kinematics were varied to create 11 experimental conditions. We investigated both differences in velocity and headway (distance from front bumper to front bumper). For the differences in headway, we used the projected headway at the merge point as the underlying metric to design the conditions and determine the initial positions for a given velocity difference. The projected headway is the headway at the merge point if both drivers would maintain their initial velocity. We chose this metric because it does not depend on track dimensions or a snapshot of the vehicle state at an arbitrary point along the track (e.g., at the tunnel exit). A combination of relative velocity and projected headway fully defines the positions and velocities of both vehicles at the start of the experiment

¹<https://github.com/tud-hri/simple-merging-experiment>

²<https://tud-hri.github.io/simple-merging-experiment>

and at the tunnel exit because the drivers have no control over their vehicles in the tunnel (see Figure 6.4).

To visualise the differences between conditions, we plotted them in a 2D projected-headway - relative-velocity plane (Figure 6.3). This figure shows the conflict space. If the projected headway is larger than the vehicle length, there is no conflict. These areas are shown in grey on the left and right side of Figure 6.3. The figure also shows in which areas we expected the right or the left driver to have an advantage. This expectation was based on a (shorter) pilot experiment with the same experimental setup but different kinematic conditions.




Figure 6.3: The experimental conditions in their two-dimensional space. The x-axis shows the projected headway at the merge point if both drivers would keep their initial velocity. If the headway is larger than the vehicle length (4.5 m) there is no projected collision, this is indicated by the grey areas on the left and right side. The y-axis shows the initial velocity differences. Positive values mean that the left vehicle is (projected to be) ahead or moving faster. The diagonal darker area divides the space into areas where the left or right driver has the advantage of passing the merge point first. This line was estimated by interpolating the results of a pilot experiment (with different kinematic conditions) to find a 50% distribution between left and right going first. Note that this does not simply divide the plane into areas where one driver has the velocity or projected headway advantage.

We used this expectation to design and spread the conditions evenly over the conflict space. The diagonal darker area represents the area in which the (kinematic) advantage changes from the left to the right driver. We decided not to investigate this area but to (first) focus on driver behaviour in cases where the outcome is more distinct. Our aim here is to gain insight into the interactions and negotiations between the two drivers in these situations. However, we did include a baseline condition where neither driver has a position or velocity advantage. With these conditions, we aim to obtain a quantitative description of the most likely outcome (who merges first) based on the initial kinematics. We used the Python package Pymer4 [23] for all statistical models in this work.

We named the conditions based on the two dimensions that define them: the projected headway in meters and the velocity difference in decimetres per second. Positive numbers indicate that the left driver has an advantage. For example, in condition **-2_8**, the right driver has a projected headway advantage of 2 m, but the left driver drives 0.8 m/s faster. For more visual examples of conditions and their names, see Figure 6.4. In our experiment, every condition was repeated 10 times in a random order for every pair of participants.




Figure 6.4: Three visualisations of experimental conditions. The figures show the relative positions of the vehicles and the start point, tunnel exit, and merge point. These merge point positions would occur if both vehicles would maintain their initial velocity. In most conditions, the slower vehicle has a position advantage at the tunnel exit. The exceptions are conditions 4_8 and -4_-8 , where the vehicles exit the tunnel at the same time.

6

We used the Conflict Resolution Time (CRT) [21] to analyse the conflict resolution behaviour of the pairs of participants. The CRT denotes the time from the start of the interaction until the first moment at which the vehicles are no longer on a collision course (assuming constant velocity). To calculate the CRT, we post-process the data and determine for every time step if a collision would occur on the remaining track if both vehicles would continue their velocity. The time between the tunnel exit and the first moment where no collision would occur is the CRT. Drivers had limited time to resolve the conflict after exiting the tunnel; they reached the merge point (where they would collide if they take no action) in 4.9 seconds on average. CRT is a measure of the amount of time needed to resolve the conflict and, therefore, can be used as a measure of the difficulty of the merging conflict.

6.3. Results

Table 6.1: The number of observed collisions per condition. The total number of trials per condition was 90. Most collisions occurred between 4 and 6 seconds after the vehicles exited the tunnels.

Condition	-4_-8	-4_0	-4_8	-2_8	0_8	0_0	0_-8	2_-8	4_-8	4_0	4_8
Collisions	3	1	3	2	4	5	3	2	1	2	2

We structure our investigation of driver conflict resolution behaviour into two parts. First, we present the analysis of the joint behaviour of two drivers, to analyse the outcome of the conflict (who gives way) and how quickly each pair of drivers resolved the merging conflict. Metrics that capture the joint behaviour for each pair under different conditions include a percentage of who merged first, as well as the Conflict Resolution Time (CRT). Second, we investigate the contributions of each individual driver in a pair to resolve the conflict. This includes the actions the individual drivers took in terms of accelerations and the resulting velocity profiles.

6.3.1. Joint behaviour

Who merged first?

The high-level outcome of a merging conflict can be summarised by which driver reached the merge point first, except for the trials where the vehicles collided. However, collisions were rare across all conditions (Table 6.1). We plot the proportion of left and right vehicles that went first as a function of initial conditions in




Figure 6.5: An overview of the high-level outcome per condition: which driver went first? Every condition was repeated 10 times for all 9 participant pairs. Therefore, the total number of trials per condition is 90. The markers show the measured data as the percentage of the left driver merging first, with the vertical line representing the 95% binomial proportion confidence intervals. Collisions were omitted from these results (see Table 6.1). The lines and shaded areas represent the (population) predictions of the mixed-effects logistic regression model (Table 6.2) with the 95% confidence interval.

Table 6.2: Mixed-effects logistic regression model describing the effect of projected headway and relative velocity on which driver reached the merge point first. Collisions were excluded, the left vehicle going first was labelled as 1, right first as 0. The model includes a random intercept for participant pairs to account for between-pair differences.

	Estimate	SE	Z	P-value	Confidence interval	
					0.025	0.975
Intercept	-0.32	0.212	-1.50	1.326×10^{-1}	-0.73	0.10
Projected headway	1.15	0.080	14.4	6.966×10^{-47}	0.99	1.31
Relative velocity	-3.4138	0.321	-10.6	2.858×10^{-26}	-4.04	-2.78

Figure 6.5. In the "neutral" 0_0 condition this proportion is almost evenly distributed. For the other 10 conditions with kinematic differences between the drivers, 5 conditions show a consistent outcome over all pairs and trials. This indicates that the outcome in these conditions is entirely defined by kinematics, with no variation between participant pairs. In one other condition (2_−8), only a single trial deviated from the outcome norm. Four conditions (−4_−8, 4_8, 0_−8, and 0_8) show a large majority of the outcomes where a particular driver merges first and a minority of the other driver merging first.

To investigate the relationship between the initial conditions (i.e. the kinematics at the start of each scenario) and the outcome (which driver merges first), we fitted a mixed-effects logistic regression model to the data. The model parameters are shown in Table 6.2, and the model outcome is visualised in Figures 6.5 and 6.6.

Table 6.3: Fixed effects estimates of the random intercept values per pair for the mixed-effects logistic regression model (Table 6.2).

Participant Pair	1	2	3	4	5	6	7	8	9
Intercept	-0.54	-0.42	-1.17	0.06	-0.13	-0.51	0.16	-0.22	-0.13

These results show that increasing the projected headway advantage increases the chances of a driver merging first ($z = 14.4, p < 10^{-46}$). The relative velocity on the other hand has a negative effect on the probability the driver merges first ($z = -10.6, p < 10^{-25}$). This means that, for equal projected headways, a driver with a higher initial velocity tends to merge behind the driver with a lower initial velocity. The explanation for this is that drivers with a higher initial velocity exit the tunnel later than the slower vehicle in most conditions (Figure 6.4). An important side-note to these effects is that we found these in a symmetric scenario with no right of way for either of the drivers.

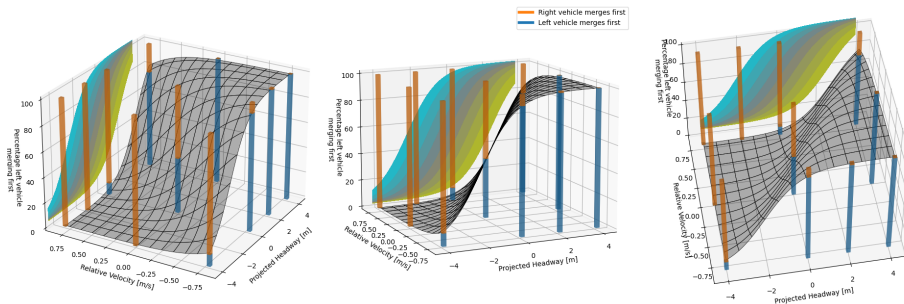


Figure 6.6: A 3-dimensional visualisation of a (population) prediction of the logistic regression model on the data. All three subplots show the same data for different angles. The model predictions are shown as the black surface and the background projections. The coloured bars show the data from the experiment. The x and y -axis represent the condition kinematics. The z -axis shows the percentage of trials where the left driver merged first. Collisions were excluded from this data (see Table 6.1). An interactive version of this plot can be found in the online supplementary materials.

The population level intercept had a negative estimated value that is not significant ($z = -1.5, p = 0.13$). This could be explained by the fact that the intercept explains a bias in the data towards the left or the right driver. This effect is clearest in the neutral condition (0_0), where we found that the right driver merged first in a small majority of the cases. Table 6.3 shows the estimated intercept values for the individual participant pairs. We expect that with more participants, the bias on the population level will disappear and the intercept value will approach 0. To visualise at which locations in the conflict space the left or right driver is more likely to merge first, we have created a top-down view heat map of the regression model. This heat map is shown in Figure 6.7 and closely resembles Figure 6.3.

Conflict Resolution Time

Besides how the conflict was resolved (which driver merged first) we investigated how quickly the conflict was resolved by examining the Conflict Resolution Time (CRT). This is a measure of the time it took the drivers to resolve the conflict and therefore resembles the difficulty of the conflict in a specific trial. Figure 6.8 shows the CRT distributions we found for all experimental conditions. The median CRT is highest for the neutral condition 0_0. In this condition, no driver has a headway or velocity advantage. Drivers have to negotiate a solution without a "most-likely" candidate solution. The lowest median CRT was found for the conditions where one driver only had a projected headway difference but the velocity was the same for both drivers. The conditions with velocities differences but no projected headway difference had high median CRTs. Thus conflicts where one driver has

a pure projected headway advantage are easier to resolve than conflicts where one driver has a pure velocity advantage.




Figure 6.7: A heat map of a logistic regression model prediction for the driver that will merge first. The conditions where data was gathered are marked with black squares.




Figure 6.8: Distribution of the Conflict Resolution Time (CRT) for all conditions. The CRT is the time from the moment at which the drivers gain control until the first moment when they are no longer on a collision course (assuming constant velocities). The coloured horizontal bars indicate the average time at which the first vehicle reached the merge point in that condition. A figure that shows the same CRT distribution placed in the 2-dimension conflict space on the locations of the corresponding conditions is available in the online supplementary material.




Figure 6.9: Distribution of the Conflict Resolution Time (CRT) from the perspective of the first merging driver. In this plot, positive numbers for headway and velocity differences indicate an advantage for the driver that merged first in that trial. This results in a different number of trials per box (see the labels at the top of the figure). The lines and shaded areas visualise predictions of the mixed effects model (Table 6.4) and its 95% confidence interval.

But besides these high-level observations, Figure 6.8 reveals no clear relationship between the initial kinematics and the CRT of the merging conflicts. We expected that the high-level outcome of the conflict (who merged first) might partly explain

Table 6.4: Mixed-effects linear regression model analysing the Conflict Resolution Time (CRT) as a function of the kinematic conditions. Positive headways and relative velocities indicate an advantage for the driver who merged first. Collisions were excluded.

	Estimate	SE	T-stat	P-value	Confidence interval	
					0.025	0.975
Intercept	1.61	0.107	15.2	4.97×10^{-10}	1.41	1.83
Projected headway	-0.25	0.016	-15.3	2.16×10^{-47}	-0.28	-0.22
Relative velocity	0.40	0.080	5.02	6.07×10^{-7}	0.25	0.56
Relative velocity : projected headway	-0.14	0.023	-6.09	1.68×10^{-9}	-0.18	-0.09

the CRT of that trial. More concretely, we expected trials where the driver with the kinematic advantage went first, to be resolved more quickly than trials where the driver with a disadvantage went first. To investigate this, we analysed CRT as a function of the kinematic advantage from the perspective of the first merging driver (Figure 6.9, Table 6.4). The projected headway and velocity differences in this figure are positive if the first merging driver had the advantage. We found that trials with a larger headway advantage for the driver that merged first had a lower CRT ($t = -15.3$, $p < 10^{-46}$). Trials with a velocity advantage for the first merging driver had a higher CRT ($t = 5.02$, $p < 10^{-6}$). Moreover, we found that the association between the CRT and the projected headway advantage was stronger for larger velocity advantage ($t = -6.09$, $p < 10^{-8}$). One important side note is that drivers with a higher initial velocity have a headway disadvantage in the approach section, i.e., they are approaching the merge point behind the other driver.

6.3.2. Individual behaviour

To gain insight into the operational behaviour of the drivers, we investigated the aggregated velocity traces of all drivers (Figure 6.10). We choose to show the velocity traces for the neutral condition (0_0) here because this condition has the widest variety of solutions (in terms of who merges first). Because of this spread, this velocity plot is easier to read than the same plot for other conditions. However, the key aspects identified in this plot are representative of the other conditions (for the raw data, including plots, see [22]. Interactive versions of these plots are available in the supplementary material).

One of the striking characteristics of the velocity traces in Figure 6.10 are the triangular patterns that can be observed in many traces. Such triangular-shaped velocity patterns indicate two things. First, it shows that drivers use blocks of constant acceleration (step inputs on gas/brake) to control their vehicle during an interaction. Second, in between these step inputs, or straight lines in the velocity trace, the input changes rapidly, causing a sharp angle in the velocity trace. This indicates that drivers select an input level and stick to that until something triggers a new decision resulting in a new input level. We refer to this combination as *intermittent piecewise-constant control*, where intermittent refers to the observed decision moments, and piecewise-constant to the constant acceleration levels in between.

With this intermittent piecewise-constant control, drivers use key decision moments at which they determine a plan. After this decision, they stick with this plan until something triggers a new decision. Therefore, Figure 6.10 provides evidence that drivers do not continuously optimise their acceleration input while interacting in




Figure 6.10: Velocity traces of the left and right drivers for all trials in the neutral condition 0_0, from the tunnel exit up until the merge point. The trials of a representative pair are highlighted to provide more insight into individual traces. The markers at the end of the trials indicate the final outcome of the trial. These plots show that drivers use triangular velocity patterns while interacting. These triangular patterns indicate that drivers use blocks of constant acceleration input with key decision moments in between. Interactive versions of these plots for all conditions are available in the online supplementary material.




Figure 6.11: The outcome of the merging conflict plotted versus the initial acceleration input at tunnel exit for the left (x-axis) and right (y-axis) drivers for all conditions.

traffic. Thus, the assumption of continuous utility maximisation that is made in many models of driver behaviour (e.g. [6], [24]–[27]) does not hold for these interactions. Another aspect shown in Figure 6.10 is that in many cases, the drivers immediately accelerate or decelerate at the moment they gain control. This indicates, that even in this purely symmetrical condition, drivers exit the tunnel with an intended solution in mind (i.e., they plan to go first or yield). To further investigate if drivers start the interaction with a mutual solution in mind, and if this solution is also reached, we plotted the outcome of the merging conflict versus the initial drivers' actions in Figure 6.11.

Figure 6.11 shows that in the majority of the interactions that do not end in a collision, the drivers initially cooperate. In most interactions that end in the left vehicle reaching the merge point first, the left driver's initial input was to accelerate and the right driver's initial input was to decelerate. This indicates two things. First, it shows that if drivers share the same perspective and observations of a merging situation, they form compatible ideas about who will merge first before they even start interacting (in that trial), i.e., drivers use a shared mental model [28]. Second, even though there are cases where the conflict is resolved by only one of the drivers (i.e., where the other driver's input is 0), in most cases, both drivers initially act simultaneously to prevent a collision.

6

6.4. Discussion

In this paper, we investigated the conflict-resolving behaviour of pairs of drivers in a simplified merging scenario. Our four most important findings are: 1) both the relative velocity and projected headway have a significant effect on which driver merges first; 2) the time it takes drivers to resolve the conflict (CRT) can be explained by the kinematics from the perspective of the driver that merges first; 3) drivers used a shared mental model about which driver merges first based on observations before the start of the interaction; and 4) drivers use intermittent piecewise-constant control to resolve the conflict, suggesting they do not constantly optimise some utility function. Rather the observed control behaviour is in line with satisficing (see [29]): in our experiment drivers seem to search for a plan that is good enough and stick to that plan until it no longer suffices. At this key decision moment, they re-plan to find a new input that is good enough, and act accordingly.

6.4.1. Relation to the existing literature

Our study indicated for the first time that both the relative velocity and the projected headway significantly influence which driver merges first. When drivers are on a collision course, a velocity advantage decreases the probability of a vehicle merging first while a projected headway advantage increases that probability. Earlier studies mostly used naturalistic data, where these kinematics can not be controlled (e.g., [7]–[9], [30]), or reduced the analysis of kinematics to one dimension by studying time to arrival (e.g., [10]).

The finding that humans do not constantly optimise their behaviour corresponds to previous findings in simple economic games [31], velocity choice for isolated drivers [32], and high-level skill switching (between manual braking and using cruise control) during driving [33]. The key-decision moments with constant inputs in-between have previously been observed in individual truck driver behaviour in real traffic [34], and in steering behaviour in high-fidelity driving simulators [35].

However, our results are the first to show that these operational aspects of human driving are also present in merging interactions in a controlled experiment.

Previous empirical studies on merging behaviour used naturalistic data [7]–[10], [30], in which these operational aspects are not included. Most of these studies focus on evaluating gap acceptance behaviour and were inspired by an interest in the effects of merging behaviour on traffic flow [7], [8], [30]. Among the existing studies of naturalistic merging conflicts, two –in particular– had a goal similar to ours: to understand the dynamics of drivers' conflict-resolving behaviour.

Wang et al. [9] studied social interactions on congested highways in the INTERACTION dataset [36]. They divided merges based on the positions of the vehicles at (what they define as) the start of the interaction. They label situations based on the through-lane vehicle initially being ahead or behind the merging vehicle. Through-lane drivers who overtake a merging car before they merge were labelled "rude", while drivers who let the merging vehicle merge in front of them were labelled "courteous". Thereby, the authors attribute the outcome of who goes first purely to driving style. However, our results indicate that the outcome (who goes first) strongly depends on the vehicles' kinematic states at the start of the interaction. In our experiment, both the relative velocity and the size of the initial gap are important indicators of who merges first; we did not find substantial individual differences based on driving style. Although driving styles play an important role in real traffic, we interpret our results as a call for cautiousness when referring to drivers as rude or courteous purely based on the fact that they overtake each other.

Srinivasan et al. [10] used naturalistic data to evaluate a machine-learned model of human merging behaviour. They concluded that this machine-learned model can successfully predict the trajectories shown by drivers in scenarios where one of the vehicles has a large kinematic advantage. Compared to our work, they reduced the kinematic differences to a single dimension: time-to-arrival. A 0.0 s time-to-arrival difference corresponds to a 0.0 m projected headway in our work, but other time-to-arrival differences can be obtained with multiple combinations of projected headway and relative velocity. Our results show that these both have a significant impact on the outcome of the conflict in terms of the driver that merges first (Table 6.2) and on the CRT (Table 6.4). An important difference between our work and [10] is that we only regarded situations where the drivers are on a collision course from the start of the interaction while [10] regards large(r) kinematic differences. Nevertheless, we advocate using both relative velocity and projected headway for the kinematic analysis, because they have different effects on the outcome of the interaction. Besides that, we expect no major implications for machine-learned models of human behaviour based on our results.

6.4.2. Implications

However, when regarding approaches that are not purely data-driven, our results could have major implications for models and control strategies. Many driver models make the assumption that humans behave as rational utility maximisers (e.g., [12], [24], [27], [37]). And because these models make this assumption, many control strategies for autonomous vehicles in mixed traffic were proposed that make the same assumption (e.g., [5], [6], [38]–[41]).

Roughly, two kinds of rational utility maximisation are used in driver models. First, there are the models that regard merging as a single high-level decision about who merges first, such as Kita already proposed in 1999 [37]. Second, there are

models that assume drivers continuously optimise some reward function to determine their current input (Naumann et al. showed many examples of reward functions used for this approach in 2020 [27]). Our results have major implications for both assumptions.

For models that regard merging as a single decision, our exploration of different kinematic conditions provides valuable insights into driver behaviour. Our results confirm that the vehicles' kinematics at the start of the interaction have a major impact on which driver merges first. This is in line with the model proposed by Kita [37]. However, our results also show that the individual differences in outcomes between pairs of drivers are restricted to a limited range of kinematic scenarios. In most scenarios, the same driver merges first for all driver pairs. This would indicate that modelling the decision of who merges first based on individual preferences (differences in reward function) is only valuable for a limited set of conditions where the kinematic differences are small.

For models that assume continuous optimisation, our results have more far-reaching implications. The aggregated velocity plot (Figure 6.10) shows that drivers do not continuously optimise, but re-plan at specific decision moments. This indicates that the assumptions that drivers **continuously** either: optimise, approximately optimise (up to a threshold), or noisily optimise their inputs are not consistent with driver behaviour. Instead, drivers seem to be triggered to change their behaviour at a certain point (at which they might partially optimise to find a new plan). Besides the key-decision moments, Figure 6.10 also shows piecewise-linear velocity patterns. This indicates that the assumption that drivers aim to minimise a squared difference between their current and desired velocity (as used in many models, e.g., [24], [27]) is also inconsistent with driver behaviour because that would lead to non-linear velocity profiles.

In general, our findings imply that the mathematical convenience related to main assumptions in game-theoretic models comes at a serious cost to their descriptive power. Thus, although game-theoretic approaches can be very valuable to determine optimal control decisions between rational agents (e.g., in vehicle-to-vehicle communication approaches [42]–[44]), we advise caution in applying them to predicting driver behaviour (either in driver models or in AV control).

6.4.3. Recommendations, limitations, and future work

Therefore, we interpret our results as an encouragement to develop new types of traffic interaction models that do allow for intermittent piecewise-constant control in operational behaviour. Siebinga et al. previously proposed a model framework that could describe intermittent control in traffic interactions [45]. But there are other (existing) lines of research that also hold potential for application to interactive scenarios, such as evidence accumulation models (e.g., [46]–[48]). Besides the intermittent control, new interaction models should use piecewise-constant acceleration as control inputs. Furthermore, they should be able to describe the most likely outcomes for different initial kinematics, independent of individual driver differences (Figure 6.5).

Although our work might provide inspiration for the development of novel interaction models, it also has some limitations. The main limitation is the simplification of the merging scenario. We started our investigation into interaction dynamics in driving simulators with a simplified symmetric merging scenario in a top-down view simulator. We chose to use this simplified scenario because of the complexity of real-world merging. Our scenario does not include lateral control (i.e., steering),

right-of-way or surrounding traffic. This allowed us to focus on the longitudinal control dynamics of two drivers who are on a collision course. We chose to use the top-down-view simulator because setting up a high-fidelity coupled simulator is a complicated and costly endeavour. Driver behaviour in our simulator has not been validated on naturalistic data yet. A detailed investigation into the relationship between the behaviour in our simulator and real-world merging is left for future work.

Such validation is complex because, as discussed in the introduction, the available naturalistic datasets lack insight into control inputs of multiple vehicles and control over kinematics. Some datasets do include many (uncontrolled) kinematic examples (e.g., [49]), which could allow validation of high-level outcomes (i.e., who goes first). Others include the control behaviour of individual drivers (e.g. [16]) enabling validation of the low-level control behaviour observed. However, it is likely that a custom naturalistic dataset needs to be collected for full validation of the simulator. A possible intermediate step could be an experiment with real vehicles on a test track. This would allow for a controlled but real test environment.

But nonetheless, we are confident that our conclusions will generalise to real-world driving. To explain why, we will discuss the three major differences between our simulated scenario and real-world driving and their potential effects on the results. The first major difference between our scenario and the real world is the absence of traffic rules and customs such as the right of way. These rules govern who can go first. Therefore, they probably don't greatly affect the operational behaviour we found. However, their absence may have affected the results regarding who goes first, conflict resolution time, and initial actions. Because the traffic rules effect is absent, the kinematic effects we found in our experiment may have been exaggerated. However, there is no reason to assume that the effects we found will not be present in a situation with traffic rules.

The second major difference is the simplification of the control inputs to acceleration and deceleration only. This design choice decreased the possible actions a driver can take as well as the difficulty of the task. This will have reduced the variability in our results. The same holds for the third (and maybe largest) major difference with driving on real roads: the top-down view of the situation. This perspective makes it easier for participants to estimate relative velocities and distances. Such a decrease in the inaccuracies in human perception could decrease the variability in the results, increasing the statistical power of our model. However, these factors are unlikely to affect the nature of the acceleration inputs (intermittent piece-wise constant control). The fact that the same input behaviour was previously found in real traffic [34] strengthens our belief that the operational driver behaviour in our simulator resembles that of the real world.

Finally, the interactions in the experiment were not as risky and anonymous as real highway interactions. Participants knowingly executed 110 merging manoeuvres against the same opponent with a name and a face, while the consequences of a collision were not as severe in real life. This could have influenced the outcome because the participants could have learned the other driver's behaviour. However, we found no evidence of learning effects for any participant pair beyond the familiarisation trials (plots can be found with the online supplementary materials). Furthermore, participants could have changed how risky they behaved. The decreased severity of a collision could have caused more risky behaviour (which would explain the large number of collisions we observed) while the identifiable opponent could have led to more courteous behaviour. However, because we

only draw conclusions on the operational behaviour and the differences in behaviour between conditions, it is unlikely that this influenced our conclusions. In the future, an experiment with more than two drivers and an experimental setup with random pairing could be used to verify this.

Another limitation of our work lies in the experimental construct of "the start of the interaction". In our experiment, we control this moment by giving participants control over their vehicle at a certain point in time when we are sure they have had enough time to observe the kinematics of the situation. This provided us with the opportunity to investigate a situation where both drivers observe and start acting at the same time. However, in real traffic, this is mostly not the case. There will be differences in when drivers see each other and consequently in when they act. How to extend metrics such as CRT, and thus how to leverage some of our findings in the real world, is not trivial. More work is needed to thoroughly investigate this.

6

6.5. Conclusion

In this paper, we investigated how drivers resolved merging conflicts in a coupled, top-down view driving simulator. We used a simplified merging scenario that only includes longitudinal control. We investigated driver behaviour under initial conditions with varying relative velocities and projected headways. We used mixed effects regression models, the concept of Conflict Resolution Time (CRT), and aggregated velocity plots to gain insight into driver behaviour. For the experimental conditions studied, we conclude:

- Drivers used intermittent control (modifying acceleration only at key decision moments) to resolve merging conflicts. This suggests that drivers do not behave as continuous rational utility maximisers in merging interactions.
- Drivers use piecewise-constant acceleration control (blocks of continuous acceleration) resulting in triangular velocity patterns to control their vehicle.
- Relative velocity and projected headway are good predictors of which driver is most likely to merge first. They have different effects and are thus both needed for a reliable prediction (instead of reducing the kinematics to a single time-to-arrival value).
- We used a metric to describe the amount of time the drivers need to resolve a merging conflict (CRT). We found CRT is associated with the outcome of the interaction combined with the initial kinematic differences (projected headway and relative velocity).
- Conditions where one driver has a pure projected headway advantage are resolved faster than conditions with a pure velocity advantage.
- Drivers used shared mental models and observations before the start of the interaction to determine which driver will merge first.

Acknowledgements

We thank Alexis Derumigny for the advice on statistical modelling.

Bibliography

- [1] Y. Hassan, M. Sarhan, and M. Salehi, "Probabilistic model for design of freeway acceleration speed-change lanes", *Transportation Research Record*, no. 2309, pp. 3–11, 2012, ISSN: 03611981. DOI: [10.3141/2309-01](https://doi.org/10.3141/2309-01).
- [2] D. Lord and J. A. Bonneson, "Calibration of predictive models for estimating safety of ramp design configurations", *Transportation Research Record*, no. 1908, pp. 88–95, 2005, ISSN: 03611981. DOI: [10.3141/1908-11](https://doi.org/10.3141/1908-11).
- [3] C. Wang, C. Xu, J. Xia, Z. Qian, and L. Lu, "A combined use of microscopic traffic simulation and extreme value methods for traffic safety evaluation", *Transportation Research Part C: Emerging Technologies*, vol. 90, no. March, pp. 281–291, 2018, ISSN: 0968090X. DOI: [10.1016/j.trc.2018.03.011](https://doi.org/10.1016/j.trc.2018.03.011).
- [4] Q. Yang and H. N. Koutsopoulos, "A microscopic traffic simulator for evaluation of dynamic traffic management systems", *Transportation Research Part C: Emerging Technologies*, vol. 4, no. 3 PART C, pp. 113–129, 1996, ISSN: 0968090X. DOI: [10.1016/S0968-090X\(96\)00006-X](https://doi.org/10.1016/S0968-090X(96)00006-X).
- [5] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [6] D. Sadigh, N. Landolfi, S. S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state", *Autonomous Robots*, vol. 42, no. 7, pp. 1405–1426, Oct. 2018, ISSN: 0929-5593. DOI: [10.1007/s10514-018-9746-1](https://doi.org/10.1007/s10514-018-9746-1). [Online]. Available: <https://doi.org/10.1007/s10514-018-9746-1> <http://link.springer.com/10.1007/s10514-018-9746-1>.
- [7] W. Daamen, M. Loot, and S. P. Hoogendoorn, "Empirical analysis of merging behavior at freeway on-ramp", *Transportation Research Record*, no. 2188, pp. 108–118, 2010, ISSN: 03611981. DOI: [10.3141/2188-12](https://doi.org/10.3141/2188-12).
- [8] F. Marczak, W. Daamen, and C. Buisson, "Merging behaviour: Empirical comparison between two sites and new theory development", *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 530–546, 2013, ISSN: 0968090X. DOI: [10.1016/j.trc.2013.07.007](https://doi.org/10.1016/j.trc.2013.07.007). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2013.07.007>.
- [9] H. Wang, W. Wang, S. Yuan, X. Li, and L. Sun, "On Social Interactions of Merging Behaviors at Highway On-Ramps in Congested Traffic", *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 11 237–11 248, 2022, ISSN: 15580016. DOI: [10.1109/TITS.2021.3102407](https://doi.org/10.1109/TITS.2021.3102407). arXiv: [2008.06156](https://arxiv.org/abs/2008.06156).
- [10] A. R. Srinivasan, M. Hasan, Y.-S. Lin, *et al.*, *Comparing merging behaviors observed in naturalistic data with behaviors generated by a machine learned model*, 2021. arXiv: [2104.10496](https://arxiv.org/abs/2104.10496) [cs.LG].
- [11] A. Ji and D. Levinson, "A review of game theory models of lane changing", *Transportmetrica A: Transport Science*, vol. 9935, no. May, pp. 1–19, 2020, ISSN: 2324-9935. DOI: [10.1080/23249935.2020.1770368](https://doi.org/10.1080/23249935.2020.1770368). [Online]. Available: <https://doi.org/10.1080/23249935.2020.1770368>.
- [12] H. Liu, W. Xin, Z. Adam, and J. Ban, "A game theoretical approach for modelling merging and yielding behaviour at freeway on-ramp section", *Transportation and Traffic Theory*, no. January, pp. 1–15, 2007. [Online]. Available: [http://www.ce.umn.edu/\\$%5Csim\\$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf](http://www.ce.umn.edu/$%5Csim$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf).
- [13] K. Kang and H. A. Rakha, "Game Theoretical Approach to Model Decision Making for Merging Maneuvers at Freeway On-Ramps", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2623, no. 1, pp. 19–28, Jan. 2017, ISSN: 0361-1981. DOI: [10.3141/2623-03](https://doi.org/10.3141/2623-03). [Online]. Available: <http://journals.sagepub.com/doi/10.3141/2623-03>.


- [14] X. Wan, P. J. Jin, F. Yang, J. Zhang, and B. Ran, "Modeling Vehicle Interactions during Merge in Congested Weaving Section of Freeway Ramp", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2421, no. 1, pp. 82–92, 2014, ISSN: 0361-1981. DOI: [10.3141/2421-10](https://doi.org/10.3141/2421-10).
- [15] C. Dong, J. M. Dolan, and B. Litkouhi, "Smooth Behavioral Estimation for Ramp Merging Control in Autonomous Driving", *IEEE Intelligent Vehicles Symposium, Proceedings*, vol. 2018-June, no. lv, pp. 1692–1697, 2018. DOI: [10.1109/IVS.2018.8500576](https://doi.org/10.1109/IVS.2018.8500576).
- [16] J. F. Antin, S. Lee, M. A. Perez, T. A. Dingus, J. M. Hankey, and A. Brach, "Second strategic highway research program naturalistic driving study methods", *Safety Science*, vol. 119, no. January, pp. 2–10, Nov. 2019, ISSN: 09257535. DOI: [10.1016/j.ssci.2019.01.016](https://doi.org/10.1016/j.ssci.2019.01.016). [Online]. Available: <https://doi.org/10.1016/j.ssci.2019.01.016><https://linkinghub.elsevier.com/retrieve/pii/S0925753518301012>.
- [17] R. Eenink, Y. Barnard, M. Baumann, X. Augros, and F. Utesch, "UDRIVE: The European naturalistic driving study", *Transportation Research Arena*, vol. 32, no. 2, pp. 1–10, 2014. [Online]. Available: <http://eprints.whiterose.ac.uk/93078/1/Paper%20-%20UDRIVE%20the%20European%20naturalistic%20driving%20study.pdf>.
- [18] T. A. Dingus, S. Klauer, V. L. Neale, et al., "The hundred-car naturalistic driving study. Phase 2, results of the hundred-car field experiment", United States. Office of Advanced Vehicle Safety Research ; Virginia. Dept. of Transportation, Tech. Rep. April, 2006.
- [19] T. Stoll, L. Weihrauch, and M. Baumann, "After you: Merging at Highway On-Ramps", *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 64, no. 1, pp. 1105–1109, 2020, ISSN: 2169-5067. DOI: [10.1177/1071181320641266](https://doi.org/10.1177/1071181320641266).
- [20] A. Shimojo, Y. Ninomiya, K. Miwa, et al., "How impressions of other drivers affect one's behavior when merging lanes", *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 89, no. June, pp. 236–248, Aug. 2022, ISSN: 13698478. DOI: [10.1016/j.trf.2022.06.007](https://doi.org/10.1016/j.trf.2022.06.007). [Online]. Available: <https://doi.org/10.1016/j.trf.2022.06.007><https://linkinghub.elsevier.com/retrieve/pii/S1369847822001309>.
- [21] O. Siebinga, A. Zgonnikov, and D. Abbink, "Interactive merging behavior in a coupled driving simulator: Experimental framework and case study", *Human Factors in Transportation*, vol. 60, pp. 516–525, 2022. DOI: [10.54941/ahfe1002485](https://doi.org/10.54941/ahfe1002485).
- [22] O. Siebinga, A. Zgonnikov, and D. A. (Abbink, "Data underlying the publication: Interactive merging behavior in a coupled driving simulator: Experimental framework and case study", 2022. DOI: [10.4121/19550377.v1](https://doi.org/10.4121/19550377.v1). [Online]. Available: https://data.4tu.nl/articles/dataset/Data_underlying_the_publication_Interactive_merging_behavior_in_a_coupled_driving_simulator_Experimental_framework_and_case_study/19550377.
- [23] E. Jolly, "Pymer4: Connecting r and python for linear mixed modeling", *Journal of Open Source Software*, vol. 3, no. 31, p. 862, 2018. DOI: [10.21105/joss.00862](https://doi.org/10.21105/joss.00862). [Online]. Available: <https://doi.org/10.21105/joss.00862>.
- [24] O. Siebinga, A. Zgonnikov, and D. Abbink, "A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles", *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 4, pp. 1–21, Dec. 2022, ISSN: 2573-9522. DOI: [10.1145/3538705](https://doi.org/10.1145/3538705). [Online]. Available: <https://dl.acm.org/doi/10.1145/3538705>.
- [25] A. Ji and D. Levinson, "Estimating the Social Gap With a Game Theory Model of Lane Changing", *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–10, 2020, ISSN: 1524-9050. DOI: [10.1109/tits.2020.2991242](https://doi.org/10.1109/tits.2020.2991242).
- [26] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and Decision-Making for Autonomous Vehicles", *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 187–210, 2018, ISSN: 2573-5144. DOI: [10.1146/annurev-control-060117-105157](https://doi.org/10.1146/annurev-control-060117-105157).

- [27] M. Naumann, L. Sun, W. Zhan, and M. Tomizuka, "Analyzing the Suitability of Cost Functions for Explaining and Imitating Human Driving Behavior based on Inverse Reinforcement Learning", in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2020, pp. 5481–5487, ISBN: 978-1-7281-7395-5. DOI: [10.1109/ICRA40945.2020.9196795](https://doi.org/10.1109/ICRA40945.2020.9196795). [Online]. Available: <https://ieeexplore.ieee.org/document/9196795/>.
- [28] C. M. Jonker, M. Birna van Riemsdijk, and B. Vermeulen, "Shared mental models: a conceptual analysis.", *COIN 2010 International Workshops*, no. Section 2, pp. 132–151, 2010.
- [29] H. A. Simon, "Rational choice and the structure of the environment.", *Psychological Review*, vol. 63, no. 2, pp. 129–138, 1956, ISSN: 1939-1471. DOI: [10.1037/h0042769](https://doi.org/10.1037/h0042769). [Online]. Available: <http://doi.apa.org/getdoi.cfm?doi=10.1037/xge0000013%20http://doi.apa.org/getdoi.cfm?doi=10.1037/h0042769>.
- [30] L. Klitzke, K. Gimm, C. Koch, and F. Koster, "Extraction and Analysis of Highway On-Ramp Merging Scenarios from Naturalistic Trajectory Data", in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, IEEE, Oct. 2022, pp. 654–660, ISBN: 978-1-6654-6880-0. DOI: [10.1109/ITSC55140.2022.9922191](https://doi.org/10.1109/ITSC55140.2022.9922191). arXiv: [2104.05661](https://arxiv.org/abs/2104.05661). [Online]. Available: <https://arxiv.org/abs/2104.05661v2%20https://ieeexplore.ieee.org/document/9922191/>.
- [31] C. F. Camerer, *Behavioral game theory: Experiments in strategic interaction*. 2003, ISBN: 0691090394. DOI: [10.1016/j.socec.2003.10.009](https://doi.org/10.1016/j.socec.2003.10.009).
- [32] M. Schmidt-Daffy, "Prospect balancing theory: Bounded rationality of drivers' speed choice", *Accident Analysis and Prevention*, vol. 63, pp. 49–64, 2014, ISSN: 00014575. DOI: [10.1016/j.aap.2013.10.028](https://doi.org/10.1016/j.aap.2013.10.028). [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2013.10.028>.
- [33] M. A. Goodrich, W. C. Stirling, and E. R. Boer, "Satisficing revisited", *Minds and Machines*, vol. 10, no. 1, pp. 79–110, 2000, ISSN: 09246495. DOI: [10.1023/A:1008325423033](https://doi.org/10.1023/A:1008325423033).
- [34] G. Markkula, "Modeling driver control behavior in both routine and near-accident driving", *Proceedings of the Human Factors and Ergonomics Society*, vol. 2014-Janua, pp. 879–883, 2014, ISSN: 10711813. DOI: [10.1177/1541931214581185](https://doi.org/10.1177/1541931214581185).
- [35] G. Markkula, E. Boer, R. Romano, and N. Merat, "Sustained sensorimotor control as intermittent decisions about prediction errors: computational framework and application to ground vehicle steering", *Biological Cybernetics*, vol. 112, no. 3, pp. 181–207, 2018, ISSN: 14320770. DOI: [10.1007/s00422-017-0743-9](https://doi.org/10.1007/s00422-017-0743-9). arXiv: [1703.03030](https://arxiv.org/abs/1703.03030). [Online]. Available: <https://doi.org/10.1007/s00422-017-0743-9>.
- [36] W. Zhan, L. Sun, D. Wang, et al., "INTERACTION Dataset: An INTERnational, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps", Sep. 2019. arXiv: [1910.03088](https://arxiv.org/abs/1910.03088). [Online]. Available: <http://arxiv.org/abs/1910.03088>.
- [37] H. Kita, "A merging-giveway interaction model of cars in a merging section: A game theoretic analysis", *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3-4, pp. 305–312, 1999, ISSN: 09658564. DOI: [10.1016/s0965-8564\(98\)00039-1](https://doi.org/10.1016/s0965-8564(98)00039-1).
- [38] J. F. Fisac, E. Bronstein, E. Stefansson, D. Sadigh, S. S. Sastry, and A. D. Dragan, "Hierarchical game-theoretic planning for autonomous vehicles", *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2019-May, pp. 9590–9596, 2019, ISSN: 10504729. DOI: [10.1109/ICRA.2019.8794007](https://doi.org/10.1109/ICRA.2019.8794007). arXiv: [1810.05766](https://arxiv.org/abs/1810.05766).
- [39] S. Coskun, Q. Zhang, and R. Langari, "Receding Horizon Markov Game Autonomous Driving Strategy", in *2019 American Control Conference (ACC)*, vol. 2019-July, IEEE, Jul. 2019, pp. 1367–1374, ISBN: 978-1-5386-7926-5. DOI: [10.23919/ACC.2019.8815251](https://doi.org/10.23919/ACC.2019.8815251). [Online]. Available: <https://ieeexplore.ieee.org/document/8815251/>.
- [40] M. Garzón and A. Spalanzani, "Game theoretic decision making for autonomous vehicles' merge manoeuvre in high traffic scenarios", *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019*, pp. 3448–3453, 2019. DOI: [10.1109/ITSC.2019.8917314](https://doi.org/10.1109/ITSC.2019.8917314).

- [41] D. Isele, "Interactive Decision Making for Autonomous Vehicles in Dense Traffic", in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, IEEE, Oct. 2019, pp. 3981–3986, ISBN: 978-1-5386-7024-8. DOI: [10.1109/ITSC.2019.8916982](https://doi.org/10.1109/ITSC.2019.8916982). arXiv: [1909.12914](https://arxiv.org/abs/1909.12914). [Online]. Available: <https://ieeexplore.ieee.org/document/8916982/>.
- [42] A. Talebpour, H. S. Mahmassani, and S. H. Hamdar, "Modeling lane-changing behavior in a connected environment: A game theory approach", *Transportation Research Part C: Emerging Technologies*, vol. 59, pp. 216–232, 2015, ISSN: 0968090X. DOI: [10.1016/j.trc.2015.07.007](https://doi.org/10.1016/j.trc.2015.07.007). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2015.07.007>.
- [43] M. Elhenawy, A. A. Elbery, A. A. Hassan, and H. A. Rakha, "An Intersection Game-Theory-Based Traffic Control Algorithm in a Connected Vehicle Environment", *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2015-October, no. September, pp. 343–347, 2015. DOI: [10.1109/ITSC.2015.65](https://doi.org/10.1109/ITSC.2015.65).
- [44] L. Banjanovic-Mehmedovic, E. Halilovic, I. Bosankic, M. Kantardzic, and S. Kasapovic, "Autonomous Vehicle-to-Vehicle (V2V) Decision Making in Roundabout using Game Theory", *International Journal of Advanced Computer Science and Applications*, vol. 7, no. 8, 2016, ISSN: 2158107X. DOI: [10.14569/ijacsa.2016.070840](https://doi.org/10.14569/ijacsa.2016.070840).
- [45] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Modelling communication-enabled traffic interactions", *Royal Society Open Science*, vol. 10, no. 5, May 2023, ISSN: 2054-5703. DOI: [10.1098/rsos.230537](https://doi.org/10.1098/rsos.230537). [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsos.230537>.
- [46] G. Markkula, Y.-s. Lin, A. R. Srinivasan, et al., "Explaining human interactions on the road by large-scale integration of computational psychological theory", *PNAS Nexus*, vol. 2, no. 6, S. Gavrillets, Ed., pp. 1–13, May 2023, ISSN: 2752-6542. DOI: [10.1093/pnasnexus/pgad163](https://doi.org/10.1093/pnasnexus/pgad163). [Online]. Available: <https://doi.org/10.1093/pnasnexus/pgad163%20https://academic.oup.com/pnasnexus/article/doi/10.1093/pnasnexus/pgad163/7202259>.
- [47] A. Zgonnikov, D. Abbink, and G. Markkula, "Should I Stay or Should I Go? Cognitive Modeling of Left-Turn Gap Acceptance Decisions in Human Drivers", *Human Factors*, 2022, ISSN: 15478181. DOI: [10.1177/00187208221144561](https://doi.org/10.1177/00187208221144561).
- [48] U. Durrani, C. Lee, and D. Shah, "Predicting driver reaction time and deceleration: Comparison of perception-reaction thresholds and evidence accumulation framework", *Accident Analysis and Prevention*, vol. 149, no. November 2020, p. 105889, 2021, ISSN: 00014575. DOI: [10.1016/j.aap.2020.105889](https://doi.org/10.1016/j.aap.2020.105889). [Online]. Available: <https://doi.org/10.1016/j.aap.2020.105889>.
- [49] T. Moers, L. Vater, R. Krajewski, J. Bock, A. Zlocki, and L. Eckstein, "The exid dataset: A real-world trajectory dataset of highly interactive highway scenarios in germany", in *2022 IEEE Intelligent Vehicles Symposium (IV)*, 2022, pp. 958–964. DOI: [10.1109/IV51971.2022.9827305](https://doi.org/10.1109/IV51971.2022.9827305).

7

A model of dyadic merging interactions explains human drivers' behaviour from input signals to decisions



One of the bottlenecks of automated driving technologies is safe and socially acceptable interactions with human-driven vehicles, for example during merging. Driver models that provide accurate predictions of joint and individual driver behaviour of high-level decisions, safety margins, and low-level control inputs are required to improve the interactive capabilities of automated driving. Existing driver models typically focus on one of these aspects. Unified models capturing all aspects are missing which hinders understanding of the principles that govern human traffic interactions. This in turn limits the ability of automated vehicles to resolve merging interactions. Here, we present a communication-enabled interaction model based on risk perception with the potential to capture merging interactions on all three levels. Our model accurately describes human behaviour in a simplified merging scenario, addressing both individual actions (such as velocity adjustments) and joint actions (such as the order of merging). Contrary to other interaction models, our model does not assume humans are rational and explicitly accounts for communication between drivers. Our results demonstrate that communication and risk-based decision-making explain observed human interactions on multiple levels. This explanation improves our understanding of the underlying mechanisms of human traffic interactions and poses a step towards interaction-aware automated driving.

7.1. Introduction

Automated driving holds many potential benefits for society [1]–[3], but, safe and efficient interactions between Automated Vehicles (AVs) and human-driven vehicles remain an open problem [4]. Such interactions frequently occur in everyday traffic: at intersections, on roundabouts, and on highways. This work will focus on merging interactions on highways as they are especially intricate due to the high speeds and multiple available options to resolve a conflict (Figure 7.1-A).

A potential solution to handling such interactions in AVs is through interaction-aware controllers (e.g., [5], [6]). These controllers assume that human drivers unilaterally respond to the AV's behaviour and use a model to predict these responses [7]. However, real-world merging interactions are inherently *reciprocal* [8]: a driver does not only respond to another driver but also influences their behaviour through implicit (or even explicit) communication [4], [9]. Individual control inputs and decisions of two or more drivers in a merging situation lead to a *joint interaction* outcome on multiple levels (Figure 7.1-C): high-level decision-making (negotiating who goes first), acceptable safety margins, and required individual control inputs. This makes real-world merging behaviour complex to understand and model, both from an individual and joint perspective (for a real-world example, see Figure 7.1-B). Interaction-aware AVs should use a model that captures this complexity, which is currently lacking.

Previous work has shown that in merging interactions, *high-level individual decisions* are made to yield or to go [4], [10], which lead to universal joint outcomes in terms of who merges first based on the kinematics of a merging scenario [11]. The *safety margins* (e.g., gaps between two vehicles) are the result of joint behaviour [12], but at the same time, these gaps are used by individual drivers to communicate their intent [9]. *Low-level control inputs* (i.e., acceleration, velocity, and position) are used by individual drivers to communicate [9], [13], thereby playing an essential role in both the *outcome* of the interaction and the *human perception* of other vehicles' behaviour. These interrelated aspects of individual and joint behaviours are not well understood, and a driver model capturing all aspects is missing.

Our work builds on related work in merging and lane-changing models (merging is often considered a special type of lane change [14]), in which we identified five classes: 1) gap acceptance; 2) traffic simulation; 3) statistical; 4) acceleration; and 5) game theoretic models (Figure 7.1-C). Gap acceptance models (e.g. [15]–[19]) describe the decisions made by the individual merging drivers by evaluating available gaps (safety margins) against a personal minimal acceptable gap size. Traffic simulation models often rely on the same gap acceptance theory for making high-level decisions [20]–[22], and are complemented with acceleration models (e.g., the intelligent driver model (IDM) [12]) to include the control behaviour before and after the merging decision. These acceleration models describe individual accelerations and do not include interactive or communicative behaviour. We found one acceleration-based model that describes individual drivers' control inputs and safety margins during interactions [23]. Statistical models provide a probability that a certain vehicle will merge or change lanes based on naturalistic traffic data. Some include desired safety margins [24], [25], while others do not [26]. Finally, game-theoretic models describe the high-level outcome and decision-making of multiple drivers in a single model (e.g., [10], [27]–[29], see [30] for a review). Game theoretic models assume humans to be

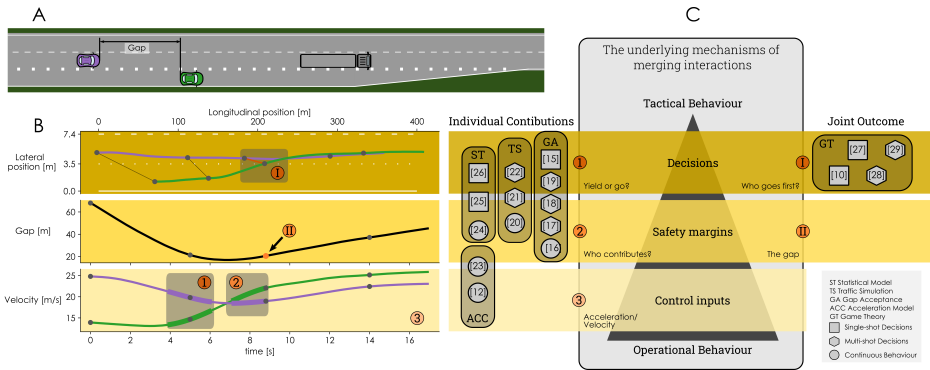


Figure 7.1: Models of highway interactions and the aspects of interaction they describe in a merging scenario. **A:** a typical interactive merging scenario taken from the HighD dataset [31] (dataset 60, vehicles 458 and 468). In this scenario, the driver of the green vehicle wants to merge onto the highway. This vehicle has a position advantage but a significantly lower velocity compared to the purple vehicle. **B:** the vehicles' position traces, the gap between the vehicles, and the individual velocity traces. In this example, the joint high-level decision is indicated at I: Green merges ahead of Purple (i.e., Green goes first). Both vehicles individually contribute to this decision by accelerating and decelerating respectively (at 1). After the decision has been made, Green keeps accelerating and thereby individually contributes to maintaining a safety margin while purple stops decelerating (2). The gap between the vehicles, when Green crosses the lane marker (II), denotes the joint safety margin. Finally, the underlying characteristics of the individual vehicle control inputs are depicted here as the total velocity traces (3). We evaluate velocity traces instead of raw accelerations because they are easier to perceive for other drivers and provide more insight into the trend of the driver's actions since they are less noisy. These individual and joint perspectives on the three levels of behaviour are also indicated in panel C. **C:** the three levels of behaviour in between Michon's *operational* and *tactical* behaviour [32]. It also shows five modelling strategies for merging interactions, each with examples from the literature. The icons indicate if the models describe a single decision at the start of an interaction (one-shot), repeated decisions (multi-shot), or continuous behaviour. Every modelling strategy captures part of the overall interactive behaviour, but none covers all five aspects. We postulate that a model capturing all three levels of individual and joint behaviour simultaneously is likely to have captured the underlying mechanisms of merging behaviour.

rational utility-maximizing agents that do not communicate. However, it is known that these assumptions do not hold for merging drivers [8].

To achieve predictable, legible [33], acceptable, and safe automated behaviour, we need to provide interaction-aware AVs with a driver model that covers all three levels. The decision level is important because if automated driving violates the underlying behavioural norms of human drivers on this level (e.g., it claims the right of way) its behaviour will be unacceptable to passengers and other drivers [4]. Automated behaviour should adhere to acceptable safety margins (on a joint level) and understand how individual drivers keep these margins. This way, AVs will show behaviour that is not just safe but is also perceived as safe. Finally, understanding the subtleties of gaps, positions, and velocities can help AVs understand the communication from other drivers and act accordingly. Since merging is a reciprocal interaction, such a model should capture the joint behaviour, not just that of a single driver responding to their environment [8].

Besides direct applications in interaction-aware AVs, a complete model of merging interactions could also prove to be a valuable step towards theories and a better fundamental understanding of human interactive capacities and behaviour in general [34]. A joint driver model could be used to understand and investigate how drivers perceive the behaviour of others, how they communicate, and how they negotiate a safe solution in general traffic interactions. This fundamental understanding is needed to design automation that can interact in a natural manner.

In a review of automated interactive traffic behaviour, Brown et al. stated: "Designing systems that can understand and react to such [implicit traffic] communication will rely upon developing an understanding of that communication beyond statistical regularity" [4]. Black-box trajectory prediction models (e.g., [35]–[37]) do not provide insight into the underlying mechanisms of merging interactions and will thus not suffice for this purpose. If a model of merging interactions succeeds in capturing the underlying mechanisms of merging interactions, which is likely if it captures behaviour across multiple levels, it could generalise to other interactive traffic scenarios, helping us gain insight into the fundamentals of interactive human driving behaviour.

The main contribution of this manuscript is a novel computational model for a simplified merging scenario with human drivers, based on the Communication-Enabled Interaction (CEI) framework [8]. The model assumes that drivers have a deterministic plan for the near future and form a probabilistic belief about another driver's intentions based on implicit communication. The plan and belief result in a perception of risk. If this risk exceeds a personal threshold, a driver updates their plan to get the risk under control. We validate our model on empirical data collected in a top-down view driving simulator with pairs of drivers [11], [38]. Our model accurately describes the (qualitative and quantitative) control input characteristics, safety margins and high-level decisions (i.e., who goes first?) of human drivers. It captures differences in individual contributions of drivers and joint behaviour. Finally, our model does not assume human rationality and explicitly incorporates communication between drivers as one of the fundamental aspects of interactions, making it the first merging interaction model to avoid these common game-theoretic assumptions.

7

7.2. Results

Simplified merging scenario and experiment To study and model merging interactions between two human drivers, we used data previously gathered in an experiment using a simplified merging scenario (Figure 7.2-A) in a coupled, top-down view driving simulator [11], [38]. The scenario simplifies merging by simulating two roads (or lanes) that merge into one at a single point. The vehicles start in a tunnel where the drivers can only observe both vehicles travelling at constant velocity to facilitate velocity perception before the interaction. Once both vehicles have exited the tunnel, the drivers gain control over the accelerations of their vehicles (steering is not possible) to resolve a merging conflict. The drivers (9 pairs of participants) were instructed to maintain their initial velocity yet prevent a collision. The experiment used 11 different experimental conditions that would end in a collision if the vehicles kept their initial velocity (10 repetitions per condition). A condition consisted of a combination of initial relative velocity ($-0.8, 0.0, \text{ or } 0.8 \text{ m/s}$) and the projected headway at the merge point if both vehicles would continue their initial velocity ($-4, -2, 0, 2, \text{ or } 4 \text{ m}$). The names of the conditions denote *projected headway_relative velocity* (e.g., 4_-8). Positive numbers indicate an advantage for the left driver. We refer to an individual driver as either the left or right driver throughout all trials, based on their physical location during the experiment. The scenario was completely symmetrical (i.e., there was no right of way). Note that from the drivers' view, they randomly perceived approaching the merge point from the left or right side of the track – to account for potential biases due to traffic rules. They were seated in the same room but could not see each other or communicate in any other way than via vehicle motion. For a more extensive description of the

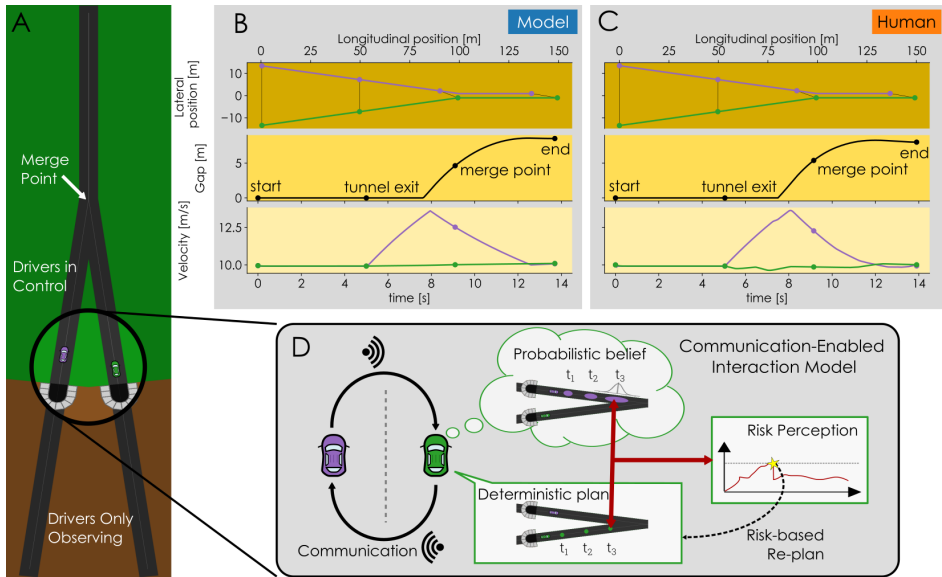


Figure 7.2: **A**: the simplified merging scenario used both in model simulations and the experiment in a coupled top-down-view driving simulator [11]. Two vehicles start on different roads, which they follow to a single merge point of those roads. They start in a tunnel where the driver (i.e., either human participants or the model) can observe both vehicles and their initial velocities but have no control yet. After exiting the tunnel, the drivers can now control the vehicles' acceleration (no steering control is needed or available). Beyond the merge point is a short road where the vehicles follow each other. **B and C**: typical examples of human and model interactions in this scenario (participant pair 3, the model behaviour resulted from a fit on the human behaviour across all conditions and all trials; the model trial shown here was not fitted specifically to this human trial). **D**: The Communication-Enabled-Interaction (CEI) framework. For clarity, the panel only visualises the three model components for the green vehicle, but the model is symmetrical, so the purple vehicle has the same three components. Each driver has a *deterministic plan* for their own behaviour and a *probabilistic belief* of the positions of the other driver in the near future. Combined, the plan and belief result in a continuous perception of risk. If this perceived risk exceeds a *risk threshold* the driver alters their plan to return the risk under the threshold. Each driver *communicates* their plan (intention) implicitly (e.g., through vehicle motion) to the other driver, who bases their belief on the received communication. Thus, this communication links one driver's plan to the belief of the other driver.

experimental protocol, see [11], [38]. Figure 7.2-C shows a typical trial outcome with human participants.

Model and Communication-Enabled Interaction (CEI) framework We created a novel model based on the Communication-Enabled Interaction (CEI) framework in [8] to describe the joint behaviour of a pair of merging drivers¹. At the core of the CEI framework lies the idea that drivers communicate their plans (intentions) to others using implicit or explicit communication. Empirical evidence has shown that this kind of communication plays an important role in traffic interactions (e.g. [9], [39]), which to date has been absent in interactive driving models. We assume that drivers form a probabilistic belief about the other driver's future movements (intent) based on this communication. Combined with their own deterministic plan, this belief underlies a driver's perceived risk. We assume that if the risk exceeds the driver's individual risk threshold, they will unilaterally alter their plan to get the risk

¹The model in this chapter differs substantially from the model presented in Chapter 4. For a full description of the model discussed here, please see the methods section of this chapter.

under control. The CEI framework [8] describes this overall structure consisting of four modules: communication, plan, belief, and risk perception. The model proposed here instantiates the CEI framework by implementing these four modules. The methods section provides a full specification of all modules of our model (plan, communication, belief, and risk), yet we give a summary here to help interpret the results. In our model (Figure 7.2-D), drivers plan a deterministic trajectory (i.e., a set of waypoints) by optimising comfort and speed over a time horizon. They communicate this plan implicitly through vehicle kinematics (current position, velocity, and acceleration). Drivers' velocity perception is assumed to be noisy. The drivers' belief about the actions of the other vehicle is represented as a set of probability distributions for the other vehicle's positions at specific times in the future. The recent behaviour of the other vehicle influences the variability in the belief (i.e., inconsistent behaviour increases the variance). The perceived risk is calculated by evaluating the probability that the positions of the ego and the other vehicle (plan and belief) overlap (i.e., the probability of a collision).

We assume every driver has two dynamic risk thresholds: an upper threshold and a lower threshold. The plan is updated either when a) the upper threshold is exceeded (to prevent a collision) or when b) the perceived risk stays below the lower threshold for a certain amount of time τ (to revert to "normal" behaviour when the conflict is resolved). Both thresholds are dynamically adjusted by an incentive function reflecting traffic rules and customs. The rationale behind this is that two drivers perceive the same amount of risk, but traffic rules and customs provide a higher incentive for one of them to act. For example, the following vehicle is usually responsible for preventing collisions in a car-following scenario.

7.2.1. Model simulations

The model uses 10 parameters designed to reflect the scenario, which were equal for all simulations². We used a grid search to find individual risk threshold parameters to describe the nine pairs (18 drivers) from the experiment [11] (i.e., one upper and lower threshold per driver is used across conditions). The incentive functions are the same for all drivers and were fitted to all experimental data using linear regression. We simulated the same number of trials as in the experiment: 9 participant pairs, 11 conditions, and 10 repetitions of each condition per pair (990 total trials). The model simulations run faster than real-time with an average run time of 2.6 s (Intel Xeon E5 quad-core) for an average real-time duration of 14.2 s. An example of a simulated trial for participant pair 3 can be found in Figure 7.2-B. In the remainder of this section, we will evaluate the model behaviour on the three behavioural levels presented in Figure 7.1-C, both for individual and joint behaviour. Some trials ended in a collision; these are excluded from the results because they represent edge cases. Collisions happened infrequently and in all conditions for human drivers (28/990) and model simulations (29/990). The online supplementary materials contain more details on how the collisions were distributed over conditions [40].

Characteristics and magnitude of control inputs

Empirical evidence showed that human drivers use intermittent piece-wise constant acceleration control to solve merging conflicts in the simplified merging scenario [11]. This type of control results in piece-wise linear velocity patterns (roughly triangular in the plots) that indicate clear decision moments when drivers change

²Please see the methods section for more information.

their control input to help resolve the merging conflict. The model's control behaviour is qualitatively similar to that of the human drivers (Figure 7.3-A); it replicates the characteristic patterns in the velocity plots.

Pair 3 is highlighted in Figure 7.3-A to facilitate comparison. In this pair, the left human driver mostly decelerated at the tunnel exit to prevent a collision, although in some cases, they accelerated. The model replicates this behaviour. In one trial (each), the left model and left human maintained their initial velocity for about one second before accelerating. Both the timing of the re-planning (i.e., the location of the peak of the velocity profile) and the absolute maximum deviation from the initial velocity were consistent between the left driver and the left model. Contrary to the left driver, the right driver in pair three barely acted to mitigate the risk. Only in one case (each) the right model and the human driver decelerated at the tunnel exit to prevent a collision. This happened in the same trial where the left driver delayed their initial response and then accelerated for both the model and the human driver pair.

The magnitude of control inputs applied by each driver was consistent between the model and human behaviour for most pairs (Figure 7.3-B). Some drivers consistently provided very little input (e.g., the left driver in pair 2), while others used higher input levels (e.g., the right driver in pair 7). The model accurately reflected this quantitative difference through the personalised risk thresholds. However, in some specific conditions for specific pairs, the model produced different average behaviour caused by outliers (e.g., pair 1, left driver, condition 0_8). In this specific example, the simulated left driver came to a complete stop in 2 out of 10 trials. In both cases, the simulated left driver initially accelerated but quickly changed strategies and started to brake. At this point, the right driver had already responded to the initial acceleration and had also started braking. Thus, both vehicles were braking, and the only safe solution was for the left driver to come to a complete standstill. This sequence of events can be understood as a miscommunication caused by a strategy switch. These miscommunications also happened with human drivers (e.g., in Figure 7.3-A, the right human driver sometimes decelerates briefly before accelerating and going first). However, with human drivers, these trials ended in a collision or with a less extreme maximum deviation; complete standstills do not occur in the human data. This difference could be due to the velocity perception noise or how communication is translated to a belief in the model. However, since these are edge cases that happen infrequently, a more detailed investigation is needed to understand these miscommunications fully.

Across all participants, both human drivers and the model used lower-magnitude inputs with increasing absolute projected headways (Figure 7.3-C). Absolute relative velocities, however, had opposite effects on human and model behaviour. This can partly be explained by the fact that only for human behaviour there is a significant interaction effect. We believe the absence of this effect in the model stems from the independent effects velocity and position have on the belief (i.e., there is no interaction effect in the belief construction). The origins and inner workings of this phenomenon in human behaviour are unknown. In general, there is a strong correlation between the model's and human drivers' inputs (Figure 7.3-D) across all conditions and participants. The model exhibited slightly higher maximum absolute deviations from the initial velocity than human drivers (difference 0.31 m/s , 95 % CI [$0.23 \text{ m/s} - 0.40 \text{ m/s}$]). We believe this difference is due to the earlier explained outliers.

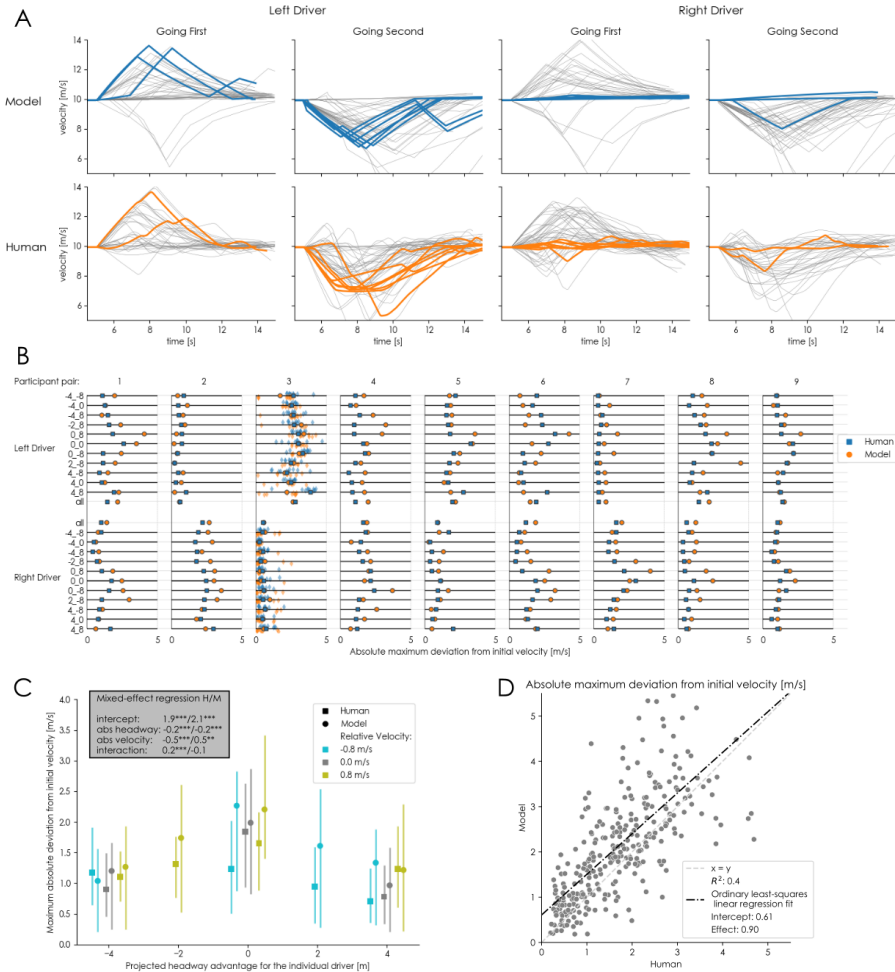


Figure 7.3: An overview of a comparison of the **control inputs** performed by 9 pairs of human drivers, and by the model fit to capture these pair's merging behaviour. Trials that ended in a collision were excluded here. **A**: all velocity traces for the model (top row) and human drivers (bottom row) in condition 0_0 (i.e., no initial velocity or position difference between the vehicles), for each driver (left/right) in the pair, divided into the high-level outcome of the trial (i.e., which driver merged first). To facilitate the comparison for a single pair, the velocities of pair 3 are highlighted. **B**: mean absolute maximum deviation from the initial velocity for the left and right driver in all participant pairs for all experimental conditions. We use the absolute deviation to compare conditions because a single condition can contain both outcomes within one pair (left going first and right going first, e.g., see panel A). Thus a single driver could accelerate in some trials and decelerate in others within one condition. The minimum and maximum deviations will be studied separately in section 7.2.1 as individual contributions to the decision. For pair 3, all underlying data are visualised (10 trials per condition). One outlier (9.5 m/s, model trial, condition 4_0) is not shown. **C**: mean absolute maximum deviation from the initial velocity, aggregated over all drivers for different relative velocities and projected headways. The error bars represent interquartile ranges. The headway and velocity values in panel C are shown from the perspective of the individual driver. This means that a projected headway of 4 m represents an advantage from the driver's perspective, while -4 m represents a disadvantage. The inset indicates coefficients of mixed-effects linear regression models predicting the mean absolute maximum deviation from initial velocity as a function of projected headway and relative velocity (* : $p \leq 0.05$, ** : $p \leq 0.01$, *** : $p \leq 0.001$). Full results of the statistical analyses are available in the supplementary materials. **D**: the relationship between human and model behaviour for all participant pairs and all conditions (i.e., all points from panel B).

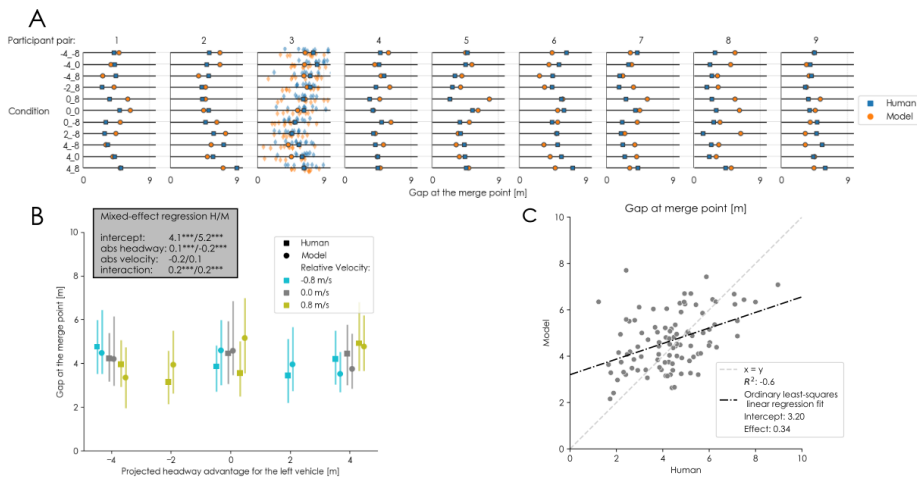


Figure 7.4: An overview of the joint **gap-keeping behaviour** of the model and human drivers. **A**: mean gap (safety margin) at the merge point for all participant pairs and all conditions, excluding collisions. For pair 3, all data points (i.e., all trials) are shown. **B**: mean gaps per condition (aggregated over all drivers); the error bars represent interquartile ranges. The inset indicates coefficients of mixed-effects linear regression models predicting the mean gap as a function of projected headway and relative velocity (* : $p \leq 0.05$, ** : $p \leq 0.01$, *** : $p \leq 0.001$); full results of the statistical analysis are available in the supplementary materials. **C**: the relationship between mean gaps of human drivers and those produced by the model for all participant pairs and all conditions (i.e., all points from panel A).

Safety margin in terms of the size of the gap between the vehicles

Safety margins can be evaluated individually or at a joint level (Figure 7.1); drivers individually contribute to achieving a certain realised safety margin (i.e., gap) on the joint level. These individual contributions can be observed from the absolute control input behaviour shown in Figure 7.3-B. We use the absolute deviation from the initial velocity to compare conditions because in some conditions drivers accelerate in some trials and decelerate in other trials.

In participant pair 3, the left driver mostly contributed to the safe solution; the right driver did not greatly deviate from their initial velocity (Figure 7.3-B). The model replicates these unequal contributions for this and other pairs (e.g., pairs 2 and 7, and to a lesser extent in pair 5). For the other participant pairs, keeping the safety margin is more of a joint effort, a phenomenon described by the model as well. The model reflects these differences in individual contribution through the baseline risk threshold levels (Table 7.2). Drivers with a lower tolerance for risk (i.e., upper risk threshold) will act to mitigate their perceived risk, while drivers with a higher tolerance will remain passive. Drivers in pairs with equal contributions (e.g., pairs 4, 6, and 9) have similar upper risk thresholds, while drivers in pairs with unequal contributions (e.g., pairs 2, 3, and 7) have larger individual differences (Table 7.2). In some human pairs, the drivers contribute equally to the safety margin, but their relative contributions differ for different conditions (Figure 7.3-B). For example, in pair 6, both drivers always act. However, in the conditions with a projected headway advantage for the left driver (i.e. a positive number), the right driver tends to do more, while for the negative projected headways, the left driver acts. This aligns with the assumption that the following driver has more incentive to solve a conflict. The model captures this behaviour (Figure 7.3-B) through the incentive functions. However, in pair 6, the differences between conditions are smaller in

the model than in the human behaviour. The limited extent to which the model shows this phenomenon can be attributed to the assumption of identical incentive functions for all pairs. Because only some human pairs show this type of behaviour, the averaged incentive functions reduce the extent to which this phenomenon is visible in the model. A method to estimate individual incentive functions could improve this in future versions of the model.

In terms of joint gap-keeping behaviour, the human data showed no substantial qualitative differences between pairs (Figure 7.4-A); the effects of kinematic conditions on the gap are significant but small (Figure 7.4-B). There is much variability within pairs, where gaps range between 0 – 9 m within one condition (pair 3, Figure 7.4-A) and within conditions, with interquartile ranges of 2 – 3 m (Figure 7.4-B). This large variability results in a limited correlation between the model and human behaviour, where the model overestimates some smaller gaps and underestimates some larger gaps (Figure 7.4-C). However, overall the model keeps gaps that are comparable in size to human behaviour for all participant pairs and all conditions (Figure 7.4-B). The mean gap for the model was only slightly larger than that of human drivers (4.8 m vs 4.5 m; difference 0.3 m, 95 % CI [0.09 m–0.51 m]).

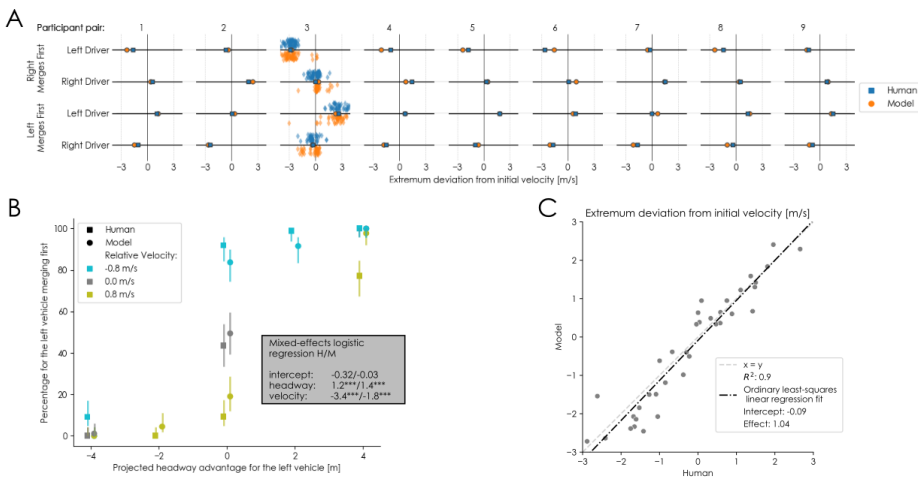


Figure 7.5: An overview of the **decision-making behaviour** of the model and human drivers. **A**: mean individual contributions of drivers to the high-level outcome of a trial as the maximum or minimum deviation from the initial velocity (i.e., amount of acceleration or deceleration); collisions are excluded. All trials are shown for a representative participant pair (pair 3). **B**: joint interaction outcome (i.e., who merged first) for all pairs in every condition. The error bars represent the 95% confidence intervals. The inset indicates coefficients of mixed-effects linear regression models predicting the mean gap as a function of projected headway and relative velocity (* : $p \leq 0.05$, ** : $p \leq 0.01$, *** : $p \leq 0.001$); full results of the statistical analysis are available in the supplementary materials. **C**: the relationship between human and model behaviour for all participant pairs and all conditions (i.e., all points from panel A).

Decisions

Finally, there is the high-level outcome of a merging interaction, which on a joint level can be summarised by the answer to the question: "Who merged first?". In human merging interactions, the probability that a driver merges first increases with their projected headway advantage and decreases with their relative velocity advantage [11]. The model replicates these effects (Figure 7.5-B), although the velocity effect is smaller for the model than for human drivers. The difference in

this effect size is especially evident in conditions with pure velocity differences (i.e., a 0 m projected headway) (Figure 7.5-B). In these conditions, the slower vehicle (i.e., the one that merges first most often) approaches the merge point ahead of the other vehicle (so that they arrive simultaneously). Potential explanations for the discrepancy in effect size between the humans and the model could be that humans systematically underestimate velocity differences or that following vehicles prefer braking over accelerating to prevent collisions. The noise on velocity and the cost of accelerating or decelerating are both assumed to be symmetrical in the model. The similar effect sizes indicate that our proposed combination of kinematics-based probabilistic beliefs and the concept of risk-based control in individual drivers is a strong potential explanation of the underlying principles that govern the high-level outcome in merging interactions.

The intercept of the logistic regression (Figure 7.5-B) quantifies the asymmetry between outcomes (and participants). The model's intercept is closer to zero than the human drivers', although this effect is insignificant (Figure 7.5-B). This can be explained by the fact that the model uses the same symmetrical functions for the belief and incentive across all drivers; therefore, the high-level outcome can be expected to be symmetrical as well (i.e., have an intercept of 0.0).

The individual contributions of the drivers to the high-level outcomes can be seen as the question for an individual driver to "go or yield"; i.e., the decision to accelerate or to brake (Figure 7.5-A). As with the individual contributions to the safety margins, some drivers consistently contribute very little to the high-level outcome (e.g., the left driver in pair 2). The model reflects this phenomenon for multiple drivers using their individual thresholds (left in pairs 2 and 7, right in pairs 3, 5, and 8). However, some drivers only contribute to the high-level outcome when they go second (e.g., the right driver in pair 1), which is reflected by the model through the incentive function. Finally, some drivers always contribute to the outcome (mostly in interactions with drivers that do nothing, e.g., the left driver in pair 3). The model reflects all these three qualitative phenomena (Figure 7.5-A). Quantitatively, there is a strong correlation between the humans' and models' decision behaviour (Figure 7.5-C).

Besides the individual decisions that lead to a joint high-level outcome, the model also describes how long the drivers take to reach a decision on a safe outcome. This duration can be measured with the Conflict Resolution Time (CRT) [38]. We found that the model captured the previously observed relationship between initial kinematics and CRT (see supplementary materials for details [40]).

7.3. Discussion

We presented a model based on the Communication-Enabled Interaction (CEI) framework [8] that accurately describes driver behaviour in a simplified interactive merging scenario (Figure 7.2-A). Our model captures the driver behaviour on three levels (Figure 7.1-C): the control input behaviour of individual drivers, the safety margins kept by pairs of drivers and how individual contributions establish these, and the high-level decisions of individual drivers (i.e., to merge or yield) and the pair (i.e., who goes first?). Because the model quantitatively and qualitatively captures individual and joint driver behaviour on all three levels we consider it likely that the underlying mechanisms of the model (communication-based belief and risk-based re-planning) correspond to the mechanisms underlying human interactive driving behaviour.

These underlying mechanisms bear a resemblance to mechanisms previously used in models of traffic interactions. The communication-based belief is related to the concept of Theory of Mind (ToM) [41], used in other models of traffic interactions (e.g. [42], [43]). The main difference is that ToM assumes humans to have an internal representation of the motivations of others, while our model uses a more basic belief of future kinematics; i.e., our model does not care if another driver prioritises speed over safety, it observes a higher velocity and updates a kinematic belief.

Risk-based re-planning is a mechanism previously used by Kolekar et al. to model isolated driver behaviour in 7 real-world driving scenarios [44]. Our model extends the concept of re-planning when the perceived risk exceeds a threshold to an interactive scenario. However, our definition of risk (perceived probability of a collision) is much more simplified than the Driver Risk Field (DRF) used by Kolekar et al. The DRF considers the risk posed by different events (e.g. going off-road or colliding with a tree) and includes steering behaviour. Our simplified scenario does not require such a sophisticated definition. However, including the DRF in a new model based on the CEI framework could enable the modelling of real-world merging scenarios.

Finally, the prevalent approach to modelling multiple drivers in interactions is game theory (e.g., [43], [45]–[48]). Our results have shown that including multiple drivers in a single model (rather than modelling a single driver that responds to their environment) is important because the same joint behaviour can stem from multiple individual contributions, and both drivers continuously update their behaviour based on the other driver's actions. However, for traffic interactions between vehicles and pedestrians, it has been shown that using game theory to optimise a short-term payoff value is not enough to explain the complex phenomena observed in the real world; instead, a range of more complex mechanisms such as a ToM and implicit communication were needed [42].

Our model and the CEI framework can have important further implications on multiple fronts. The model could be used to further improve our understanding of interactive driving behaviour for the development of automated driving technologies, and our model could potentially be generalised to other scenarios with traffic interactions.

The underlying mechanisms of the model enabled it to replicate human driving behaviour on multiple levels. Therefore, the model could help researchers to understand better how these mechanisms function in human behaviour [34]. For example, implicit communication (through vehicle movements) and how it influences a driver's belief is observed often but only partly understood [9], [39]. The same holds for the perception of risk, which has been investigated for isolated drivers [44], [49], but our model could help extend this to interactions. Our model could facilitate research in these directions, leading to an increase in fundamental knowledge.

This knowledge could facilitate more practical applications, such as the design of movements that convey a clear message about the intent of an automated vehicle [4]. These movements could consider other drivers' expectations regarding high-level outcomes and communicative actions. Matching the automated vehicle behaviour with expectations might increase behavioural acceptance, although this should be investigated further. Second, the model might be used to inform the real-time interaction planning of automated vehicles. In particular, it could inform the AV about the potential future actions of other drivers (as predic-

tion models are generally being used [5], [6], [50]). Finally, our model could also be valuable in the development phase of automated behaviour by being part of an interactive and dynamic environment for benchmark testing where models are used to evaluate the behaviour of autonomous vehicles (e.g., [29], [51]).

Finally, the scenario we modelled in this work bears a resemblance to other interactive scenarios. In essence, the scenario in Figure 7.2-A entails a continuous and dynamic interaction where participants search for a mutually beneficial solution. Although there are minor individual advantages to be gained regarding speed and comfort, the most important goal is mutual: to prevent a collision. These aspects are comparable to other interactive scenarios such as traffic interactions between vehicles and pedestrians (e.g., [42], [46]), pedestrian interactions in a crowd [52], or physical human-robot interactive tasks [53] (e.g., [47], [54]). The shortcomings of other models that led to the development of the CEI framework [8] – the assumption of rationality, absence of communication, and difficulties in extending game theory beyond high-level decisions – also apply to these related scenarios. Thus, exploring CEI-based models, such as the one presented here, for other interactive scenarios can be an interesting topic for future work.

Our work has three important limitations: the simplified scenario, the manually chosen model parameters, and the complexity of fitting an intermittent behaviour model. To start with the first limitation, we have used a simplified merging scenario in this work to gain insight into the complex dynamics of driver interactions [38]. This scenario enabled us to uncover the characteristics of human behaviour regarding accelerations but does not include two important aspects of merging: steering and traffic rules (such as the right of way). We previously found that the characteristics of human behaviour in our simplified simulator correspond to those found in real-world driving [11], [55], [56]. Therefore, we are confident that our model captures an important aspect of human behaviour in real traffic: intermittent piecewise constant control. Since the other underlying principles of the model (i.e., risk perception [44], communication [9], [39], and a belief about intent [42]) have also been observed in real traffic and used in other successful models we believe that our model can be generalised to realistic merging scenarios. Nonetheless, extending the model to realistic scenarios (especially realistic beliefs that cover multiple tactical responses to the same situation [57]) remains a topic for future work.

Second, the model uses 10 manually chosen parameters to match the scenario (these were heuristically determined, based on literature, or tuned to fit the data; see the Methods section for details). Among these parameters are the length of the planning horizon for the drivers and the saturation time τ that governs how long it takes drivers to re-plan when the conflict is resolved. Although we found the model robust to changes in these parameters, how they generalise to other scenarios (e.g., with different dimensions of the track) is unknown. It is possible that all scenarios need a specific set of parameters and that to generalise the model to work in multiple scenarios simultaneously, these parameters need to be dynamically adjusted.

Finally, due to the intermittent nature of our model, it is complex to fit it to human data. Because of the mechanism where a risk threshold triggers a planning update, individual trials only provide limited information about the threshold value that would describe a driver best. If the driver acts, the threshold is exceeded, and if they don't act, the threshold is not exceeded. The amount of action is not related to the upper risk threshold. Therefore, we used a grid search method to ob-

tain individual values for risk thresholds. However, this method is computationally inefficient, imprecise, and hard to use with more complex scenarios with multiple control inputs (i.e., when including steering). Because intermittent control is a key aspect of the CEI framework and our model, this is a fundamental limitation to the potential to generalise the model. More work is needed to develop a more robust method to fit intermittent models such as ours.

In conclusion, our model hypothesised that a communication-enabled kinematic belief combined with risk-based intermittent actions underlie human interactive behaviour in merging. In contrast to the currently prevalent game-theoretic models of traffic interactions, our model does not rely on the assumption of rationality and explicitly includes implicit communication between drivers. Despite its simplicity, our model could accurately describe the joint behaviour of human drivers and their individual contributions in merging interactions on three levels: control inputs, safety margins, and decisions. We believe our model could be a useful tool to increase the fundamental understanding of the effects a vehicle's kinematics actions have on the beliefs of other drivers. Therefore, we hope our model represents a step towards understanding driving interactions and developing interaction-aware automated driving.

7.4. Methods

In this work, we evaluate a Communication-Enabled Interaction (CEI) model in a simulated environment and compare it to human behaviour data that was previously collected in a simulator experiment. Here we reiterate the details of the experiment, present the design of the four modules of the CEI model (plan, communication, belief, and risk perception), and discuss the parameters we used for the model and the fitting procedure. Finally, we present the details of the software and data we used in this work which is all available online from public repositories.

7.4.1. Experiment and simulation environment

The data on human driver behaviour we used in this work was previously gathered in an experiment in a coupled top-down-view driving simulator [11], [38]. Eighteen volunteers (6 female, 12 male, mean age: 25, std: 2.6) participated in the study and were divided into 9 fixed pairs (i.e., each participant interacted with the same counterpart in all trials). This experiment was approved by TU Delft's Human Research Ethics Committee and all participants gave their informed consent before participating.

The participants controlled the acceleration of their vehicle using the gas and brake pedal of a steering-wheel game controller (Logitech Driving Force GT). The headings of the vehicles were fixed (i.e., equal to the heading of the road). Participants each sat behind a computer screen that showed a top-down view of the simulation. They were seated in the same room behind a curtain to prevent them from seeing each other. The drivers wore noise-cancelling headsets (Sony WH-1000XM3) with ambient music to prevent them from communicating in any other way than through vehicle kinematics.

In the experiment, the drivers followed a track consisting of three sections of equal length (50 m each, total track length 150 m): a tunnel, an approach and a car-following section. The vehicle dimensions were 4.5 m x 1.8 m. In the tunnel, the drivers could observe both vehicles and their initial velocity. The initial velocities were either equal (10 m/s) for both vehicles or one of the vehicles had a 0.8 m/s advantage (9.6 m/s - 10.4 m/s). If the vehicles maintained their velocity, they

would collide at the merge point with varying headways (i.e., distance from front bumper to front bumper). We called this the projected headway and varied it between 0, 2, and 4 *m* for each vehicle. Conditions were labelled according to the projected headway and relative velocity (e.g., -4_8) where positive numbers denote an advantage for the left vehicle. Drivers were told: "Maintain your initial velocity yet prevent a collision. No vehicle has the right of way. Remain seated, use one foot on the gas or brake pedal, keep both hands on the steering wheel, and do not communicate by making sounds or noise. Remember that this is a scientific experiment, not a game or a race."

The participants approached the merge point from the left or the right side (randomised before each trial). However, only their own view was varied; in the experimenter's view, and in all results discussed here, the same driver in a pair is referred to as the left or right driver. To facilitate velocity perception, the steering wheels provided vibration feedback when vehicles deviated from their initial velocity. This feedback increased with the magnitude of the deviation. In case of a collision, the simulation was paused for 20 seconds. This time penalty lasted longer than the duration of a single trial (~ 16 seconds), thereby providing an incentive to avoid collisions.

Both in the experiment and the model simulations, we simulated the vehicles as point mass objects, their dimensions were only used for collision detection. The vehicles were subject to a negative acceleration due to resistance of $a_r = 0.5 + 0.005v^2$, where v is the vehicle's velocity. Vehicle velocities were always positive.

7.4.2. Model

Our proposed model is based on the CEI modelling framework [8]. According to this framework, joint driver behaviour can be understood as a combination of four modules: plan, communication, belief, and risk perception.

Plan The drivers are assumed to have a deterministic plan for the near future. The model plans to maintain a constant acceleration input over its planning horizon; the acceleration value is obtained by minimising a cost function over this horizon (Equation 7.1). This cost function c penalises deviating from the desired velocity v_d (which in the experiment is equal to the initial velocity) and large values of the acceleration input a_t

$$c = \sum_{t=0}^T (v_t - v_d)^2 + a_t^2, \quad (7.1)$$

where v_t denotes the velocity at time t , and T is the time horizon. Importantly, the cost function does not include a term for collision avoidance; instead, the CEI framework assumes that drivers manage safety by keeping the risk below the risk threshold, which is imposed as an optimisation constraint on the planning module of the model [8].

With the optimal constant acceleration, a trajectory (i.e., a set of waypoints over time) is constructed over the time horizon T . This trajectory is later used to evaluate the perceived risk. The planned constant acceleration is applied with added noise to execute the plan. This noise represents the discrepancy between a planned acceleration and the gas pedal input (i.e., inaccuracies in the driver's neuromuscular system and internal model). The noise is added after the optimisation in the planning phase and remains constant until the next re-plan. Noise is drawn from a scaled normal distribution: $\mathcal{N}(\mu_n = 0, \sigma_n^2 = \frac{1}{40^2})$.

If the optimisation fails because no solution can be found within the constraints, the model falls back to either full braking or full acceleration. If the ego vehicle is behind the other vehicle, heading for a collision, and no solution to the planning problem can be found, the model applies full braking. Similarly, if the ego vehicle is ahead of the other vehicle but cannot find a feasible plan, it applies full acceleration. In both cases, a new optimisation is triggered at the next time step until a valid solution can be found again.

Communication In our model, drivers communicate through vehicle kinematics. Explicit communication (e.g., with indicator lights) is not included in the model or the experiment for simplicity. Drivers observe the other vehicle's position, velocity, and acceleration at every time step. Position and acceleration observations are assumed to be perfect. Velocity perception is assumed to be noisy to account for the fact that drivers sometimes accelerate and sometimes decelerate at the tunnel exit in the same condition (Figure 7.3-A). This behaviour can be explained by the fact that drivers under- or overestimate the other vehicle's velocity in the tunnel.

The noise in the velocity perception is inspired by evidence accumulation, a concept used in driver decision-making studies before [42], [58]. Specifically, we assume that drivers update their perceived velocity of the other vehicle v^p at every time step with an observation affected by noise

$$v_t^p = v_{t-1}^p + dv^p \quad (7.2)$$

$$dv^p = \alpha(v_t - v_{t-1}^p) + \beta dW. \quad (7.3)$$

In Equations 7.2 and 7.3, subscript t denotes time, the superscript p denotes the perception of the other vehicle's velocity, dv^p is the perception update, v is the other vehicles true velocity, α denotes the update rate, β the noise level, and W is a stochastic Wiener process (thus dW is a sample from a normal distribution $\mathcal{N}(\mu = 0, \sigma^2 = dt)$).

Belief The observed communication is used to create a belief about the future positions of the other vehicle. This belief consists of belief points at a specific belief frequency f_b over the same time horizon T as the one used in the plan. Every belief point is a probability distribution over positions. Each of these belief points are represented by the sum of two normal distributions:

$$b_t = \frac{1}{2} \mathcal{N}(\mu_t, \sigma_t^2) + \frac{1}{2} \mathcal{N}(\mu_t, \phi \sigma_t^2), \quad (7.4)$$

where b is the belief point representing the probability distribution over positions for the other vehicle at time t and ϕ is a scaling factor. The first part of this equation represents the positions of the other vehicle that are kinematically feasible within the bounds of comfortable acceleration. The second part is kinematically infeasible within comfortable bounds and can be interpreted as a belief that something unexpected will happen (e.g., an emergency braking). The motivation behind including two components to the belief distribution was that a single normal distribution either assigns similar probabilities to the kinematically likely and unlikely positions (when σ_t is high) or only considers high-risk scenarios (when σ_t is low). Our belief model addresses this issue by emphasising the kinematically likely outcomes, but at the same time including a safety margin in case of errors or unlikely events

(e.g., emergency braking, long reaction times, or a perception error of the other driver).

The parameters μ_t and σ_t are based on a normally distributed expected acceleration ($\mathcal{N}(\mu_a, \sigma_a^2)$) which is constructed based on driver's memory about recent acceleration observations. Drivers keep a memory of recent acceleration observations of the other vehicle M_a

$$M_a = [a_{-T_m}, \dots, a_t] \quad (7.5)$$

$$\mu_a = \bar{M}_a \quad (7.6)$$

$$\sigma_a^2 = \left(\frac{1}{3}a_c\right)^2 + \mathbf{var}(M_a) \quad (7.7)$$

Here, the mean expected acceleration μ_a is calculated as the average of the remembered values over the past T_m seconds (Equation 7.5). The standard deviation σ_a of the expected acceleration is based on the maximum comfortable acceleration a_c (assuming that 99.7% of observed accelerations fall within $\pm a_c$) and an added variance $\mathbf{var}(M_a)$. The latter part increases the expected variance in future accelerations if inconsistent behaviour has recently been observed.

The parameters for the belief point distributions are then constructed using point mass kinematics with this normally distributed acceleration

$$\mu_t = \frac{1}{2}(t - t_0)^2\mu_a + v_0(t - t_0) + p_0, \quad (7.8)$$

$$\sigma_t^2 = \frac{1}{2}(t - t_0)^2\sigma_a^2, \quad (7.9)$$

where p_0 denotes the observed position of the other vehicle, v_0 is the perceived (noisy) velocity of the other vehicle, and t denotes the time of this belief point and t_0 the current time.

Risk perception Risk perception combines the planned trajectory and the belief about the future positions of the other vehicle to calculate the probability of a collision. For every belief point, the model determines the bounds of collisions [8]; these are the extremum positions of the other vehicle that will result in a collision. The believed probability that the other vehicles will be between these bounds is the perceived probability of a collision. This probability is assumed to be the perceived risk.

The perceived risk is evaluated against two dynamic risk thresholds, the upper (ρ_u) and the lower (ρ_l) threshold. Both thresholds consist of a driver's individual base value (θ), which is adjusted with an incentive function:

$$\rho_u^d = \theta_u^d + \lambda_{u,1}\Delta p + \lambda_{u,2}\Delta v + \lambda_{u,3}\Delta p\Delta v \quad (7.10)$$

$$\rho_l^d = \theta_l^d + \lambda_{l,1}\Delta p + \lambda_{l,2}\Delta v + \lambda_{l,3}\Delta p\Delta v. \quad (7.11)$$

In these equations, superscript d denotes a specific driver, Δp and Δv are the relative position and velocity from this driver's perspective, and λ are the incentive parameters which are assumed to be constant over the population.

A re-plan is triggered if the upper risk threshold ρ_u is exceeded. This re-plan aims to find an acceleration that brings the perceived risk below $0.8\rho_l$. If the perceived risk stays below the lower threshold ρ_l for longer than the saturation time τ , the conflict is assumed to be resolved, and another re-plan is triggered to revert to "normal" behaviour. In this case, the risk is constrained to be below $0.6\rho_u$. Finally, if the desired velocity is reached while the vehicle is accelerating or decelerating, another re-plan is triggered to allow the vehicle to maintain the preferred velocity.

7.4.3. Model parameter fitting

Our model and simulation use 10 parameters with values that were manually designed, their values are shown in Table 7.1a. The timing parameters for the simulation (dt), planning (T), and belief (T_m, f_b) were chosen such that they are suitable for the scenario yet enable reasonable computation times. The noise parameters σ_n and β , and saturation time τ were manually tuned to reflect the human data. The belief scaling factor ϕ was designed to obtain sufficient resolution in the risk signals. The parameters α (which denotes how long drivers need to observe a vehicle to estimate its velocity) and a_c (the maximal comfortable acceleration) were based on empirical literature: α [59], [60], a_c [61].

Table 7.1: Model and simulation parameters with constant values for all participants

(a) Manually designed parameters

T	6.0 s	dt	0.05 s
T_m	4.0 s	f_b	4 Hz
σ_n	$\frac{1}{40}$	β	0.6
τ	1.6 s	ϕ	3.0
α	0.5	a_c	$1.0 \frac{m}{s^2}$

(b) The fitted population-level parameters for the incentive functions

Parameter	Value	Std. Err.	z	p
$\lambda_{u,1}$	0.003	0.001	3.27	0.001
$\lambda_{u,2}$	0.018	0.006	2.97	0.003
$\lambda_{u,3}$	-0.006	0.001	-6.38	0.000
$\lambda_{l,1}$	0.004	0.001	3.55	0.000
$\lambda_{l,2}$	0.016	0.008	1.95	0.051
$\lambda_{l,3}$	-0.003	0.001	-2.03	0.042

The risk thresholds and incentive function parameters were fitted to the data using a grid search. We created a 25×25 grid for every kinematic condition using upper thresholds in the range [0.3, 0.9] and lower thresholds in the range [0.01, 0.4]. For these grids, we disabled the incentive functions and only used the base thresholds (θ in Equations 7.10 and 7.11). We ran one simulated trial per set of thresholds per condition with all the noise in the model disabled, resulting in 11 grids of 625 trials. In every trial, we simulated the behaviour of a single CEI driver against a vehicle travelling at constant velocity to obtain the immediate response of the (CEI) driver at the tunnel exit before any interaction takes place. For every such trial, we recorded the modelled driver's deviation from the initial velocity after 1.0 second. Then, for every trial of every human participant, we searched the grid for the risk threshold values that best described that driver's velocity deviation after 1.0 second. This resulted in individual base threshold values (ρ_u and ρ_l) for every trial, in total 110 sets of thresholds for every participant. With 18 participants, this gave us 1980 upper and lower threshold values for 11 different kinematic conditions.

To determine participant-level thresholds based on these trial-level thresholds, we used two linear mixed-effect models, one for the lower and one for the upper threshold. Both models used the form: $\rho \sim \Delta p * \Delta v$, where ρ denotes the threshold and Δp and Δv are the relative position and relative velocity respectively. Random intercepts were included per participant. The resulting fixed-effect coefficients (Table 7.1b) were used as incentive parameters (λ in Equations 7.10 and 7.11), and the random-effect coefficients (Table 7.2) were used as participant-level base values θ_l and θ_u .

7.4.4. Software and data

We implemented a custom simulation environment in Python to run the experiment and simulate the model. The optimisation in the planning part of the model was implemented with Casadi [62]. The statistical analyses were performed in python using the statsmodels package [63].

Table 7.2: Personal base values for the upper and lower thresholds for each driver

Pair	Driver	θ_l	θ_u
1	left	0.165	0.495
	right	0.260	0.562
2	left	0.245	0.635
	right	0.058	0.493
3	left	0.058	0.488
	right	0.245	0.631
4	left	0.183	0.537
	right	0.201	0.524
5	left	0.113	0.498
	right	0.269	0.585
6	left	0.246	0.550
	right	0.161	0.546
7	left	0.320	0.736
	right	0.201	0.522
8	left	0.165	0.525
	right	0.246	0.586
9	left	0.178	0.519
	right	0.227	0.543

Our code is publicly available on Github [64]. The data from the experiment [65] and model simulations [66] are available on the 4TU data repository. Interactive plots of the human data are available online [67].

Acknowledgement

We thank Lorenzo for his help with the Casadi implementation and Tom for his valuable feedback on the manuscript.

Bibliography

- [1] C. D. Harper, C. T. Hendrickson, S. Mangones, and C. Samaras, "Estimating potential increases in travel with autonomous vehicles for the non-driving, elderly and people with travel-restrictive medical conditions", *Transportation Research Part C: Emerging Technologies*, vol. 72, pp. 1–9, Nov. 2016, ISSN: 0968090X. DOI: [10.1016/j.trc.2016.09.003](https://doi.org/10.1016/j.trc.2016.09.003). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2016.09.003> [20https://linkinghub.elsevier.com/retrieve/pii/S0968090X16301590](https://linkinghub.elsevier.com/retrieve/pii/S0968090X16301590).
- [2] L. M. Clements and K. M. Kockelman, "Economic Effects of Automated Vehicles", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2606, no. 1, pp. 106–114, Jan. 2017, ISSN: 0361-1981. DOI: [10.3141/2606-14](https://doi.org/10.3141/2606-14). [Online]. Available: <http://journals.sagepub.com/doi/10.3141/2606-14>.
- [3] S. Pettigrew, "Why public health should embrace the autonomous car", *Australian and New Zealand Journal of Public Health*, vol. 41, no. 1, pp. 5–7, Feb. 2017, ISSN: 13260200. DOI: [10.1111/1753-6405.12588](https://doi.org/10.1111/1753-6405.12588). [Online]. Available: <http://doi.wiley.com/10.1111/1753-6405.12588>.
- [4] B. Brown, M. Broth, and E. Vinkhuyzen, "The Halting problem: Video analysis of self-driving cars in traffic", *Conference on Human Factors in Computing Systems - Proceedings*, 2023. DOI: [10.1145/3544548.3581045](https://doi.org/10.1145/3544548.3581045).
- [5] D. Sadigh, N. Landolfi, S. S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state", *Autonomous Robots*, vol. 42, no. 7, pp. 1405–1426, Oct. 2018, ISSN: 0929-5593. DOI: [10.1007/s10514-018-9746-1](https://doi.org/10.1007/s10514-018-9746-1). [Online]. Available: <https://doi.org/10.1007/s10514-018-9746-1> [20http://link.springer.com/10.1007/s10514-018-9746-1](http://link.springer.com/10.1007/s10514-018-9746-1).
- [6] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [7] O. Siebinga, A. Zgonnikov, and D. Abbink, "A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles", *ACM Transactions on Human-Robot Interaction*, vol. 11, no. 4, pp. 1–21, Dec. 2022, ISSN: 2573-9522. DOI: [10.1145/3538705](https://doi.org/10.1145/3538705). [Online]. Available: <https://dl.acm.org/doi/10.1145/3538705>.
- [8] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Modelling communication-enabled traffic interactions", *Royal Society Open Science*, vol. 10, no. 5, May 2023, ISSN: 2054-5703. DOI: [10.1098/rsos.230537](https://doi.org/10.1098/rsos.230537). [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsos.230537>.
- [9] B. Brown, E. Laurier, and E. Vinkhuyzen, "Designing Motion: Lessons for Self-driving and Robotic Motion from Human Traffic Interaction", *Proceedings of the ACM on Human-Computer Interaction*, vol. 7, no. GROUP, pp. 1–21, 2023. DOI: [10.1145/3567555](https://doi.org/10.1145/3567555).
- [10] H. Kita, "A merging-giveway interaction model of cars in a merging section: A game theoretic analysis", *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3-4, pp. 305–312, 1999, ISSN: 09658564. DOI: [10.1016/S0965-8564\(98\)00039-1](https://doi.org/10.1016/S0965-8564(98)00039-1).
- [11] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Human Merging Behavior in a Coupled Driving Simulator: How Do We Resolve Conflicts?", *IEEE Open Journal of Intelligent Transportation Systems*, vol. 5, no. October 2023, pp. 103–114, 2024, ISSN: 2687-7813. DOI: [10.1109/OJITS.2024.3349635](https://doi.org/10.1109/OJITS.2024.3349635). arXiv: [2308.04842](https://arxiv.org/abs/2308.04842). [Online]. Available: <http://arxiv.org/abs/2308.04842> [20https://ieeexplore.ieee.org/document/10380755/](https://ieeexplore.ieee.org/document/10380755/).
- [12] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, vol. 62, no. 2, pp. 1805–1824, 2000, ISSN: 1063651X. DOI: [10.1103/PhysRevE.62.1805](https://doi.org/10.1103/PhysRevE.62.1805). arXiv: [0002177](https://arxiv.org/abs/0002177) [cond-mat].

- [13] N. Kauffmann, F. Winkler, F. Naujoks, and M. Vollrath, "“What Makes a Cooperative Driver?” Identifying parameters of implicit and explicit forms of communication in a lane change scenario", *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 58, pp. 1031–1042, 2018, ISSN: 13698478. DOI: [10.1016/j.trf.2018.07.019](https://doi.org/10.1016/j.trf.2018.07.019). [Online]. Available: <https://doi.org/10.1016/j.trf.2018.07.019>.
- [14] Q. Zhang, R. Langari, H. E. Tseng, D. Filev, S. Szabowski, and S. Coskun, "A Game Theoretic Model Predictive Controller with Aggressiveness Estimation for Mandatory Lane Change", *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 1, pp. 75–89, 2020, ISSN: 23798858. DOI: [10.1109/TIV.2019.2955367](https://doi.org/10.1109/TIV.2019.2955367).
- [15] A. Kondyli and L. Elefteriadou, "Modeling driver behavior at freeway-ramp merges", *Transportation Research Record*, no. 2249, pp. 29–37, 2011, ISSN: 03611981. DOI: [10.3141/2249-05](https://doi.org/10.3141/2249-05).
- [16] J. A. Laval and L. Leclercq, "Microscopic modeling of the relaxation phenomenon using a macroscopic lane-changing model", *Transportation Research Part B: Methodological*, vol. 42, no. 6, pp. 511–522, 2008, ISSN: 01912615. DOI: [10.1016/j.trb.2007.10.004](https://doi.org/10.1016/j.trb.2007.10.004).
- [17] C. F. Choudhury, "Modeling Driving Decision with Latent Plans (PhD Dissertation)", no. 2002, 2007.
- [18] R. M. Michaels and J. Fazio, "Driver behavior model of merging", *Transportation Research Record*, no. 1213, pp. 4–10, 1989, ISSN: 03611981.
- [19] Kazi Iftekhar Ahmed, "Modeling Drivers' Acceleration and Lane Changing Behavior", Ph.D. dissertation, 1999.
- [20] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model MOBIL for car-following models", *Transportation Research Record*, no. 1999, pp. 86–94, 2007, ISSN: 03611981. DOI: [10.3141/1999-10](https://doi.org/10.3141/1999-10).
- [21] Q. Yang and H. N. Koutsopoulos, "A microscopic traffic simulator for evaluation of dynamic traffic management systems", *Transportation Research Part C: Emerging Technologies*, vol. 4, no. 3 PART C, pp. 113–129, 1996, ISSN: 0968090X. DOI: [10.1016/S0968-090X\(96\)00006-X](https://doi.org/10.1016/S0968-090X(96)00006-X).
- [22] P. Hidas, "Modelling lane changing and merging in microscopic traffic simulation", *Transportation Research Part C: Emerging Technologies*, vol. 10, no. 5-6, pp. 351–371, 2002, ISSN: 0968090X. DOI: [10.1016/S0968-090X\(02\)00026-8](https://doi.org/10.1016/S0968-090X(02)00026-8).
- [23] X. Wan, P. J. Jin, F. Yang, J. Zhang, and B. Ran, "Modeling Vehicle Interactions during Merge in Congested Weaving Section of Freeway Ramp", *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2421, no. 1, pp. 82–92, 2014, ISSN: 0361-1981. DOI: [10.3141/2421-10](https://doi.org/10.3141/2421-10).
- [24] W. Daamen, M. Loot, and S. P. Hoogendoorn, "Empirical analysis of merging behavior at freeway on-ramp", *Transportation Research Record*, no. 2188, pp. 108–118, 2010, ISSN: 03611981. DOI: [10.3141/2188-12](https://doi.org/10.3141/2188-12).
- [25] F. Marczak, W. Daamen, and C. Buisson, "Merging behaviour: Empirical comparison between two sites and new theory development", *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 530–546, 2013, ISSN: 0968090X. DOI: [10.1016/j.trc.2013.07.007](https://doi.org/10.1016/j.trc.2013.07.007). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2013.07.007>.
- [26] C. Dong, J. M. Dolan, and B. Litkouhi, "Smooth Behavioral Estimation for Ramp Merging Control in Autonomous Driving", *IEEE Intelligent Vehicles Symposium, Proceedings*, vol. 2018-June, no. Iv, pp. 1692–1697, 2018. DOI: [10.1109/IVS.2018.8500576](https://doi.org/10.1109/IVS.2018.8500576).
- [27] H. Liu, W. Xin, Z. Adam, and J. Ban, "A game theoretical approach for modelling merging and yielding behaviour at freeway on-ramp section", *Transportation and Traffic Theory*, no. January, pp. 1–15, 2007. [Online]. Available: [http://www.ce.umn.edu/\\$%5Csim\\$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf](http://www.ce.umn.edu/$%5Csim$liu/publication/2007%5C_ISTTT17%5C_Liu%5C_Xin%5C_final.pdf).
- [28] S. Coskun, Q. Zhang, and R. Langari, "Receding Horizon Markov Game Autonomous Driving Strategy", in *2019 American Control Conference (ACC)*, vol. 2019-July, IEEE, Jul. 2019, pp. 1367–1374, ISBN: 978-1-5386-7926-5. DOI: [10.23919/ACC.2019.8815251](https://doi.org/10.23919/ACC.2019.8815251). [Online]. Available: <https://ieeexplore.ieee.org/document/8815251/>.

- [29] N. Li, D. W. Oyler, M. Zhang, Y. Yildiz, I. Kolmanovsky, and A. R. Girard, "Game theoretic modeling of driver and vehicle interactions for verification and validation of autonomous vehicle control systems", *IEEE Transactions on Control Systems Technology*, vol. 26, no. 5, pp. 1782–1797, 2018, ISSN: 10636536. DOI: [10.1109/TCST.2017.2723574](https://doi.org/10.1109/TCST.2017.2723574). arXiv: [1608.08589](https://arxiv.org/abs/1608.08589).
- [30] A. Ji and D. Levinson, "A review of game theory models of lane changing", *Transportmetrica A: Transport Science*, vol. 9935, no. May, pp. 1–19, 2020, ISSN: 2324-9935. DOI: [10.1080/23249935.2020.1770368](https://doi.org/10.1080/23249935.2020.1770368). [Online]. Available: <https://doi.org/10.1080/23249935.2020.1770368>.
- [31] R. Krajewski, J. Bock, L. Kloeker, and L. 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)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 2118–2125, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569552](https://doi.org/10.1109/ITSC.2018.8569552). arXiv: [1810.05642](https://arxiv.org/abs/1810.05642). [Online]. Available: <https://ieeexplore.ieee.org/document/8569552/>.
- [32] J. A. Michon, "A Critical View of Driver Behavior Models: What Do We Know, What Should We Do?", in *Human Behavior and Traffic Safety*, Boston, MA: Springer US, 1985, pp. 485–524, ISBN: 0306422255. DOI: [10.1007/978-1-4613-2173-6_19](https://doi.org/10.1007/978-1-4613-2173-6_19). [Online]. Available: http://link.springer.com/10.1007/978-1-4613-2173-6_19.
- [33] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion", *ACM/IEEE International Conference on Human-Robot Interaction*, vol. 1, pp. 301–308, 2013, ISSN: 21672148. DOI: [10.1109/HRI.2013.6483603](https://doi.org/10.1109/HRI.2013.6483603).
- [34] O. Guest and A. E. Martin, "How Computational Modeling Can Force Theory Building in Psychological Science", *Perspectives on Psychological Science*, vol. 16, no. 4, pp. 789–802, 2021, ISSN: 17456924. DOI: [10.1177/1745691620970585](https://doi.org/10.1177/1745691620970585).
- [35] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-Feasible Trajectory Forecasting With Heterogeneous Data", 2020. arXiv: [2001.03093](https://arxiv.org/abs/2001.03093). [Online]. Available: <http://arxiv.org/abs/2001.03093>.
- [36] B. Brito, H. Zhu, W. Pan, and J. Alonso-Mora, "Social-VRNN: One-Shot Multi-modal Trajectory Prediction for Interacting Pedestrians", *Proceedings of Machine Learning Research*, vol. 155, no. CoRL, pp. 862–872, 2020, ISSN: 26403498. arXiv: [2010.09056](https://arxiv.org/abs/2010.09056).
- [37] A. Mészáros, J. F. Schumann, J. Alonso-Mora, A. Zgonnikov, and J. Kober, "TrajFlow: Learning the Distribution over Trajectories", 2023. arXiv: [2304.05166](https://arxiv.org/abs/2304.05166). [Online]. Available: <http://arxiv.org/abs/2304.05166>.
- [38] O. Siebinga, A. Zgonnikov, and D. Abbink, "Interactive merging behavior in a coupled driving simulator: Experimental framework and case study", *Human Factors in Transportation*, vol. 60, pp. 516–525, 2022. DOI: [10.54941/ahfe1002485](https://doi.org/10.54941/ahfe1002485).
- [39] Y. M. Lee, R. Madigan, O. Giles, *et al.*, "Road users rarely use explicit communication when interacting in today's traffic: implications for automated vehicles", *Cognition, Technology & Work*, vol. 23, no. 2, pp. 367–380, May 2021, ISSN: 1435-5558. DOI: [10.1007/s10111-020-00635-y](https://doi.org/10.1007/s10111-020-00635-y). [Online]. Available: <https://doi.org/10.1007/s10111-020-00635-y>. <https://link.springer.com/10.1007/s10111-020-00635-y>.
- [40] O. Siebinga, A. Zgonnikov, and D. Abbink, *Supplementary material to the paper "A model of dyadic merging interactions explains human drivers' behavior from input signals to decisions"*. [Online]. Available: https://www.olgersiebinga.nl/files/supplementary_materials_cei_model.pdf#7D.
- [41] D. Premack and G. Woodruff, "Premack and Woodruff: Chimpanzee theory of mind", *Behavioral and Brain Sciences*, vol. 4, pp. 515–526, 1978.
- [42] G. Markkula, Y.-s. Lin, A. R. Srinivasan, *et al.*, "Explaining human interactions on the road by large-scale integration of computational psychological theory", *PNAS Nexus*, vol. 2, no. 6, S. Gavrillets, Ed., pp. 1–13, May 2023, ISSN: 2752-6542. DOI: [10.1093/pnasnexus/pgad163](https://doi.org/10.1093/pnasnexus/pgad163). [Online]. Available: <https://doi.org/10.1093/pnasnexus/pgad163> <https://academic.oup.com/pnasnexus/article/doi/10.1093/pnasnexus/pgad163/7202259>.

- [43] R. Tian, M. Tomizuka, and L. Sun, "Learning Human Rewards by Inferring Their Latent Intelligence Levels in Multi-Agent Games: A Theory-of-Mind Approach with Application to Driving Data", in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Sep. 2021, pp. 4560–4567, ISBN: 978-1-6654-1714-3. DOI: [10.1109/IROS51168.2021.9636653](https://doi.org/10.1109/IROS51168.2021.9636653). arXiv: [2103.04289](https://arxiv.org/abs/2103.04289). [Online]. Available: <http://arxiv.org/abs/2103.04289%20https://ieeexplore.ieee.org/document/9636653/>.
- [44] S. Kolekar, J. de Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model", *Nature Communications*, vol. 11, no. 1, p. 4850, Dec. 2020, ISSN: 2041-1723. DOI: [10.1038/s41467-020-18353-4](https://doi.org/10.1038/s41467-020-18353-4). [Online]. Available: <http://dx.doi.org/10.1038/s41467-020-18353-4%20https://www.nature.com/articles/s41467-020-18353-4>.
- [45] V. Gabler, T. Stahl, G. Huber, O. Oguz, and D. Wollherr, "A game-theoretic approach for adaptive action selection in close proximity human-robot-collaboration", *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 2897–2903, 2017, ISSN: 10504729. DOI: [10.1109/ICRA.2017.7989336](https://doi.org/10.1109/ICRA.2017.7989336).
- [46] F. Camara, N. Bellofto, S. Cosar, *et al.*, "Pedestrian Models for Autonomous Driving Part II: High-Level Models of Human Behaviour", *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–20, 2020, ISSN: 1524-9050.
- [47] S. Nikolaidis, S. Nath, A. D. Procaccia, and S. Srinivasa, "Game-Theoretic Modeling of Human Adaptation in Human-Robot Collaboration", *ACM/IEEE International Conference on Human-Robot Interaction*, vol. Part F1271, pp. 323–331, 2017, ISSN: 21672148. DOI: [10.1145/2909824.3020253](https://doi.org/10.1145/2909824.3020253). arXiv: [1701.07790](https://arxiv.org/abs/1701.07790).
- [48] A. Turnwald, W. Olszowy, D. Wollherr, and M. Buss, "Interactive navigation of humans from a game theoretic perspective", *IEEE International Conference on Intelligent Robots and Systems*, no. Iros, pp. 703–708, 2014, ISSN: 21530866. DOI: [10.1109/IROS.2014.6942635](https://doi.org/10.1109/IROS.2014.6942635).
- [49] S. Kolekar, J. de Winter, and D. Abbink, "Which parts of the road guide obstacle avoidance? Quantifying the driver's risk field", *Applied Ergonomics*, vol. 89, no. July, p. 103196, Nov. 2020, ISSN: 00036870. DOI: [10.1016/j.apergo.2020.103196](https://doi.org/10.1016/j.apergo.2020.103196). [Online]. Available: <https://doi.org/10.1016/j.apergo.2020.103196%20https://linkinghub.elsevier.com/retrieve/pii/S0003687018307373>.
- [50] E. Ward, N. Evestedt, D. Axehill, and J. Folkesson, "Probabilistic Model for Interaction Aware Planning in Merge Scenarios", *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 2, pp. 1–1, 2017, ISSN: 2379-8904. DOI: [10.1109/TIV.2017.2730588](https://doi.org/10.1109/TIV.2017.2730588). [Online]. Available: <http://ieeexplore.ieee.org/document/7987753/>.
- [51] R. Queiroz, D. Sharma, R. Caldas, *et al.*, "A Driver-Vehicle Model for ADS Scenario-based Testing", pp. 1–16, May 2022. arXiv: [2205.02911](https://arxiv.org/abs/2205.02911). [Online]. Available: <http://arxiv.org/abs/2205.02911>.
- [52] F. Martínez-Gil, M. Lozano, I. García-Fernández, and F. Fernández, "Modeling, evaluation, and scale on artificial pedestrians: A literature review", *ACM Computing Surveys*, vol. 50, no. 5, 2017, ISSN: 15577341. DOI: [10.1145/3117808](https://doi.org/10.1145/3117808).
- [53] A. De Santis, B. Siciliano, A. De Luca, and A. Bicchi, "An atlas of physical human-robot interaction", *Mechanism and Machine Theory*, vol. 43, no. 3, pp. 253–270, 2008, ISSN: 0094114X. DOI: [10.1016/j.mechmachtheory.2007.03.003](https://doi.org/10.1016/j.mechmachtheory.2007.03.003).
- [54] B. Sadrfaridpour and Y. Wang, "Collaborative Assembly in Hybrid Manufacturing Cells: An Integrated Framework for Human-Robot Interaction", *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 3, pp. 1178–1192, 2018, ISSN: 15455955. DOI: [10.1109/TASE.2017.2748386](https://doi.org/10.1109/TASE.2017.2748386).
- [55] G. Markkula, "Modeling driver control behavior in both routine and near-accident driving", *Proceedings of the Human Factors and Ergonomics Society*, vol. 2014-Janua, pp. 879–883, 2014, ISSN: 10711813. DOI: [10.1177/1541931214581185](https://doi.org/10.1177/1541931214581185).
- [56] G. Markkula, E. Boer, R. Romano, and N. Merat, "Sustained sensorimotor control as intermittent decisions about prediction errors: computational framework and application to ground vehicle steering", *Biological Cybernetics*, vol. 112, no. 3, pp. 181–207, 2018, ISSN: 14320770. DOI: [10.1007/s00422-017-0743-9](https://doi.org/10.1007/s00422-017-0743-9). arXiv: [1703.03030](https://arxiv.org/abs/1703.03030). [Online]. Available: <https://doi.org/10.1007/s00422-017-0743-9>.

- [57] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Uncovering Variability in Human Driving Behavior Through Automatic Extraction of Similar Traffic Scenes from Large Naturalistic Datasets", in *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, IEEE, Oct. 2023, pp. 4790–4796. DOI: [10.1109/SMC53992.2023.10393913](https://doi.org/10.1109/SMC53992.2023.10393913). eprint: [2206.13386](https://arxiv.org/abs/2206.13386). [Online]. Available: <https://ieeexplore.ieee.org/document/10393913/>.
- [58] A. Zgonnikov, D. Abbink, and G. Markkula, "Should I Stay or Should I Go? Cognitive Modeling of Left-Turn Gap Acceptance Decisions in Human Drivers", *Human Factors*, 2022, ISSN: 15478181. DOI: [10.1177/00187208221144561](https://doi.org/10.1177/00187208221144561).
- [59] A. J. Fath, M. Lind, and G. P. Bingham, "Perception of time to contact of slow- and fast-moving objects using monocular and binocular motion information", *Attention, Perception, and Psychophysics*, vol. 80, no. 6, pp. 1584–1590, 2018, ISSN: 1943393X. DOI: [10.3758/s13414-018-1517-8](https://doi.org/10.3758/s13414-018-1517-8).
- [60] B. Jörges and L. R. Harris, "Object speed perception during lateral visual self-motion", *Attention, Perception, and Psychophysics*, vol. 84, no. 1, pp. 25–46, 2022, ISSN: 1943393X. DOI: [10.3758/s13414-021-02372-4](https://doi.org/10.3758/s13414-021-02372-4).
- [61] L. L. Hoberock, "A survey of longitudinal acceleration comfort studies in ground transportation vehicles", *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, vol. 99, no. 2, pp. 76–84, 1977, ISSN: 15289028. DOI: [10.1115/1.3427093](https://doi.org/10.1115/1.3427093).
- [62] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control", *Mathematical Programming Computation*, 2018.
- [63] S. Seabold and J. Perktold, "Statsmodels: Econometric and statistical modeling with python", in *9th Python in Science Conference*, 2010.
- [64] O. Siebinga, *simple-merging-experiment [Software]*, 2022. [Online]. Available: <https://github.com/tud-hri/simple-merging-experiment>.
- [65] O. Siebinga, A. Zgonnikov, and D. A. (Abbink, "Data underlying the publication: Interactive merging behavior in a coupled driving simulator: Experimental framework and case study", 2022. DOI: [10.4121/19550377.v1](https://doi.org/10.4121/19550377.v1). [Online]. Available: https://data.4tu.nl/articles/dataset/Data_underlying_the_publication_Interactive_merging_behavior_in_a_coupled_driving_simulator_Experimental_framework_and_case_study/19550377.
- [66] O. Siebinga, A. Zgonnikov, and D. Abbink, "Data underlying the publication: A model of dyadic merging interactions explains human drivers' behaviour from input signals to decisions", 2023. DOI: [10.4121/4126c919-1d0c-4ba9-80fa-3960f49e8cd7](https://doi.org/10.4121/4126c919-1d0c-4ba9-80fa-3960f49e8cd7). [Online]. Available: <https://doi.org/10.4121/4126c919-1d0c-4ba9-80fa-3960f49e8cd7>.
- [67] O. Siebinga, A. Zgonnikov, and D. Abbink, *Supplementary materials to the paper "Merging in a Coupled Driving Simulator: How do drivers resolve conflicts?"* [Online]. Available: <https://tud-hri.github.io/simple-merging-experiment>.

8

Conclusion

Three main contributions form the foundation of this thesis (Figure 8.1): 1) research on using available naturalistic highway data for driver model validation, 2) a novel model framework for interactive driving, and 3) a controlled experiment with a simplified merging scenario that yielded new empirical evidence on joint driver behaviour which aided to validate and improve that model. Combined, these contributions supported the aim of this thesis: to increase the fundamental understanding of merging and lane-changing interactions and capture this knowledge in a joint driver model.

Regarding the first main contribution, Chapters 2 and 3 focused on using available naturalistic highway data for driver model validation. Chapter 2 proposed a framework to validate driver models underlying interaction-aware controllers. A case study using the framework showed that a state-of-the-art inverse-reinforcement-learning-based model, which was used in interaction-aware control, failed to accurately describe operational and tactical driver behaviour in a large number of real-world lane changes. This chapter highlighted the importance of separately evaluating driver behaviour models' tactical and operational aspects. Chapter 3 revealed that these two levels of behaviour also play an important role in the variability in drivers' responses to similar situations. This chapter proposed a method to obtain similar situations from naturalistic datasets automatically. In a real-world situation where drivers had to respond to a slower-moving vehicle in their lane, there was variability in tactical responses (e.g., some slowed down, others changed lanes) and operational responses (e.g., in the amount of acceleration). These findings motivate the need for a new type of driver model which can accurately describe behaviour on multiple levels.

Regarding the second main contribution, Chapter 4 provided a more extensive investigation into what type of driver model would be suitable to capture interactive driver behaviour. In this chapter, we identified the most important characteristics of driver behaviour in traffic interactions and discussed to what extent existing modelling approaches can capture these characteristics. The chapter argued that no existing model type could simultaneously capture the joint dynamics of multiple drivers, include communicative actions, and account for the idea that drivers are not rational decision-makers. These limitations of existing approaches strengthened the motivation for developing a new type of driver model specifically suitable for interactions. To meet those requirements, we proposed the Communication-Enabled-Interaction (CEI) framework to describe driver behaviour in merging interactions. Chapter 4 outlined the CEI-modelling framework and presented an example of a model implementation that showed plausible human-like behaviour in 4 simplified merging simulations. In a 5th car-following scenario, human-like gap-keeping behaviour emerged from the probabilistic belief and risk perception in the model.

Chapters 5 and 6 were part of the third main contribution and investigated using controlled experiments to study driver behaviour in interactions. Previous controlled experiments focused on the behaviour of a single driver, therefore a suitable simulator scenario and analysis tools for experiments with multiple drivers were not available. Chapter 5 laid the necessary foundation for interaction experiments by proposing a scenario and three new analysis tools. In this scenario, two drivers resolve a simplified, top-down-view merging conflict in a coupled driving simulator. Chapter 6 utilised the experiment's results for an empirical analysis of driver behaviour. This analysis revealed that the scenario's kinematics, and not individual differences between drivers, mostly determined the interaction's high-level




Figure 8.1: The three pillars in which the work in this thesis was divided.

8

outcome (who merges first) and that individual drivers use intermittent piecewise-constant control to solve the merging conflict. These empirical findings served as a basis for improving the model proposed in Chapter 4.

Finally, Chapter 7 brought the three main contributions together in a new CEI model and validated this model on the empirical findings from the experiment. The new version of the model accurately reproduces quantitative and qualitative aspects of driver behaviour on multiple levels. The model captures the high-level joint outcome of the interactions (who merges first) and how this outcome is established from individual decisions (e.g. to yield or not). The model describes the gap between the vehicles, jointly kept by the drivers, and their individual contributions to this gap. Finally, the intermittent piece-wise-constant control that drivers use for their vehicle inputs is also captured by the model.

8.1. Overarching conclusions

All chapters contain specific conclusions relevant to their content; these will not be repeated here. However, based on the combined results of all chapters, I draw three overarching conclusions:

1. Different drivers respond with different tactical and operational behaviours to similar interactive situations; therefore, driver models should capture operational and tactical variability, which should be assessed independently.

In Chapter 2, we argued, based on theory and literature, that driver models should be compared with human behaviour on both tactical and operational levels separately. The main argument was the particular importance of regarding multiple levels of behaviour in models used for interaction-aware automated driving. Traditional driver models (e.g. the Intelligent Driver Model [1]), typically regard a single tactical behaviour (e.g., car following). This is sufficient for studying traffic flow or the transportation system as a whole because multiple models of different tactical behaviours can be combined to describe the population. However, interaction-

aware automated driving has to cope with individual interactions with unique drivers. These drivers might respond in different ways to the automated vehicle. A driver might first follow the AV and then decide to change lanes. Therefore, we argued that for a model to be useful to inform automated driving decisions, it should be able to express multiple tactical behaviours as a unified model.

The empirical evidence that *drivers respond with different tactical and operational behaviours to similar situations* was provided in Chapter 3 (Figure 3.5). The results of a case study, in which we studied the responses of 250 drivers to similar interactive car-following scenarios, showed variability on both the tactical and operational levels. In the same scenario, some drivers change lanes while others keep following the vehicle in front of them. Earlier studies into variability focused on operational variability within one tactical behaviour (e.g., [2]–[4]) or on tactical variability in the population instead of in responses to a particular scenario (e.g., [5]).

Going back to Chapter 2, the case study there showed that distinguishing between operational and tactical behaviour is important when evaluating a driver model that aims to capture multiple tactical behaviours (Figure 2.5). An inverse-reinforcement-learning model showed different tactical behaviour than human drivers in the same situations. On an operational level, the model always showed behaviour with the same characteristics, independent of its tactical behaviour. This operational behaviour corresponded to human behaviour during lane changes but differed from that during car following. This meant that no matter if the model changed lanes or followed another car, it always used the same acceleration profile that human drivers use to change lanes. To make this observation, these operational behaviours should not be aggregated when evaluating the model but should be assessed separately. Thus, *operational variability should be assessed independently of tactical variability*.

This conclusion has two implications in light of recent literature. First, many machine-learning-based approaches to driver modelling or trajectory prediction use single metrics such as Final Displacement Error (FDE) or Root Mean Squared Error (RMSE) to evaluate their performance [6]. These approaches are agnostic to levels of behaviour and learn their behaviour automatically from large datasets containing many behaviours; therefore, they often describe multiple tactical behaviours. Using a single metric to evaluate their behaviour will thus aggregate the operational behaviour over multiple tactical behaviours. While this issue has been described and investigated earlier [6], [7], the findings in this thesis show that it is a practical concern in the development of interaction-aware AVs and further underline its importance. Some machine-learning-based approaches for trajectory prediction even explicitly aim to predict multiple tactical behaviours (e.g. [8]–[10]). In machine-learning literature, this is referred to as "multi-modal" prediction, where tactical behaviours are regarded as (behavioural) modes. How to evaluate these modes and their predicted distribution and how to compare different prediction approaches is still an open problem [6]. The proposed approach of Chapter 2 could provide a solution for validating the machine-learned trajectory prediction approaches (beyond models used in IAC) because it shows how knowledge from driver behaviour literature can be combined with open datasets to validate models that capture multiple tactical behaviours.

Finally, many interaction-aware automated driving approaches regard human behaviour as "noisy" (e.g., [11]) or "uncertain" (e.g., [12]). However, based on this first overarching conclusion, I argue that human behaviour is *variable*, but

not *uncertain* or *noisy*. The main difference is that variability can be understood, quantified, and explained, while noise and uncertainty are inherently random. Chapters 2 and 3 have shown that we can extract and quantify operational and tactical variability in human behaviour from naturalistic datasets. This could help us gain a better understanding of the variability in driver behaviour, which in turn could help develop better interaction-aware automated driving.

To conclude, this first conclusion provided three main insights. First of all, Chapter 3 provided more evidence to support what we argued for in Chapter 2: it is important for interaction-aware AVs to regard multiple tactical behaviours in a model of interactions. Chapter 2 showed that a model used in an IAC does not generalise well to real-world scenarios; therefore, new models describing multiple tactical behaviours should be developed for this use case. Finally, the empirical evidence in this thesis shows that it is important to validate these driver models on both levels and regard these levels separately during the validation.

2. An important requirement for a joint driver model is the ability to capture that drivers do not continuously (rationally) optimise their acceleration inputs; instead, they use intermittent piece-wise constant control – as was empirically observed in a simplified merging scenario.

8

Chapters 5 and 6 are dedicated to using a controlled experiment in a driving simulator to study the interactive behaviours of drivers. Chapter 5 introduced the necessary novel scenario and analysis tools. The empirical analysis in Chapter 6 revealed that *humans do not continuously rationally (in a game theoretic sense) optimize their control inputs during interactions* (Figure 6.10). Instead, they apply a constant control input until something triggers a decision. At this decision moment, the continuous input suddenly changes. We refer to this control behaviour as: "*intermittent piecewise-constant control*, where intermittent refers to the observed decision moments, and piecewise-constant to the constant acceleration levels in between." Both the absence of rationality in human behaviour and the intermittent control have been found previously in other (driving) tasks.

Previous experiments have shown that humans do not continuously and rationally optimize their decisions when playing simple economic games [13] or when choosing their velocity during driving [14]. Instead, the concept of bounded rationality, as introduced by Simon [15], seems to provide a better explanation for the behaviour of isolated drivers [14], [16] and car-following drivers [17]. This thesis is the first to show that these findings generalise to interactive driving behaviour. Furthermore, it is the first to present a model of traffic interaction based on satisficing (a form of bounded rationality [18]).

Intermittent control behaviour in humans is a long-standing topic of research. The results are many observations, theories, and models of general intermittent (motor) control (e.g. [19], [20]) and intermittent control in isolated driving (e.g., [21], [22]). For example, intermittent control behaviour has previously been observed in the acceleration inputs of truck drivers in naturalistic data [23]. However, as with bounded rationality, this thesis generalises these findings to traffic interactions during merging and presents a model of such interactions that integrates these characteristics.

This second conclusion has implications for research and applications in two aspects. First, driver models that aim to capture the operational control behaviour of interacting drivers must consider intermittent piecewise-constant control. This is needed to describe specific subtle communicative actions such as the miscommunication example in Chapter 7 (Section 7.2.1), which was caused by a strategy

switch (i.e., intermittent behaviour). Models that only target to describe higher levels of behaviour could ignore this conclusion and assume rational continuous control. However, as argued in Chapter 7, a unified model that can accurately describe multiple levels of behaviour is more likely to have actually captured the underlying mechanisms of driver behaviour. I will revisit the distinction between operational and tactical behaviour in rationality in the Further Outlook section: Section 8.5.

The second implication relates to control approaches for interaction-aware automated driving. For interaction-aware control, it is particularly important to consider the identified characteristics of human behaviour. Understanding the finesses of human operational (intermittent) control is important when trying to understand and interpret implicit communication (through vehicle motion). We need to understand nominal behaviour before we can distinguish implicit communication in that behaviour. Furthermore, the intermittent nature of human operational control could also play a role in the acceptance of autonomous behaviour, i.e., behaviour that deviates too much from human behaviour might not be understood and thus not be acceptable. Further investigations in this direction could be interesting opportunities for future research.

Finally, the assumption that interacting humans behave rationally is commonly made in interaction-aware control literature. However, violations of this assumption (e.g., Chapter 6) could have implications for the safety of interaction-aware autonomous vehicles. Other drivers might respond differently than the AV expected, and passengers might perceive the AV behaviour as unsafe if it drives too competitively. If, and to what extent, these issues are problems to consider should be further investigated. I will revisit how we could potentially leverage the CEI-model for such investigations in Section 8.4.

3. With communication-enabled, risk-based intermittent control, the proposed CEI-model can describe abstract merging interactions between two drivers, including their decisions (who goes first), safety margins over time, and underlying individual contributions and control inputs (brake/accelerate).

Chapter 4 outlined the Communication-Enabled Interaction (CEI) model framework based on theory and literature (Figure 4.1). This chapter demonstrated the potential of a CEI model to produce plausible, human-like interactions. In Chapter 7, a CEI model implementation was presented based on the empirical findings of Chapter 6. Chapter 7 showed that the CEI model can not only produce plausible behaviours but can also accurately reproduce the high-level outcome of merging trials, the safety margins drivers maintain, the input characteristics of individual drivers, and the individual differences between participant pairs.

The CEI framework is a unique and novel approach to implementing interaction models. Three key features specifically distinguish the CEI model from other driver and interaction models. First, it explicitly incorporates communication between drivers, an important aspect of traffic interactions [24]. This is a unique feature for a model of driving interactions. Second, it models risk perception and assumes drivers act based on their perceived risk. This assumption was previously used in a model that could explain driver behaviour in seven non-interactive scenarios [25]. Finally, the CEI model uses intermittent piecewise constant control, consistent with the findings in Chapter 6.

The final CEI model presented in Chapter 7 covers several phenomena observed in human behaviour in the driving simulator experiment (Chapter 6). Of these phenomena, five can be directly linked to features of the model's implementation.

These model features could provide insight into the underlying principles of how human drivers interact.

First, the individual differences between drivers (i.e., when they act and how much they do to prevent a collision) are reflected in the model by adjusting a single parameter: the risk threshold. The model assumes that all drivers perceive the same amount of risk but that some drivers have a higher tolerance for perceived risk. Because changing the risk threshold is enough to describe individual differences accurately, this assumption might explain the differences between drivers in risk-based actions. A study on risk perception by Kolekar et al. also showed that a unified definition of perceived risk could explain the behaviour of multiple drivers [26] (they used uniform parameters for a risk model although they assumed that risk perception would differ per individual). However, untangling risk perception (i.e., how risky a situation is in one's perception) and risk thresholds (i.e., at which level of risk one acts) in experiments can be challenging since the amount of action (e.g., steering angle) is often used as the measured signal [26], [27]. A certain action can therefore happen at a low perceived risk with a low risk threshold or a high perceived risk with a high threshold. Untangling risk perception and the resulting risk-based action, as well as the differences between drivers in this respect, is needed to investigate if the model's assumption is valid. This could be an opportunity for future research.

Second, the probabilistic belief of the modelled drivers results in two phenomena being replicated by the model. First, the assumption that the other vehicle will travel at constant velocity while considering that it might accelerate or decelerate is the leading mechanism behind the high-level outcome of the simulations. That is, it explains the relationship between the kinematics of a scenario and which vehicle merges first. Second, the human-like, velocity-dependent, gap-keeping behaviour in Chapter 4 emerged from this belief representation.

The third model feature directly linked to observed behaviour is the planning of a constant acceleration value (or gas pedal input). This model feature resulted in the characteristic intermittent piecewise-constant control. However, this feature is not unique to our model; it has been used before in models of non-interacting drivers [22], [23].

The final phenomenon is only partially explained by the model: it is the fact that some drivers in an interaction (e.g., the following driver in car following) act more often to lower the risk. This is connected to the first phenomenon in the sense that the model assumes dynamic risk thresholds which cause the driver to act once they are in a certain position: we named this the incentive function. The model uses the same incentive function for all drivers. However, for some individuals, the model could not accurately reflect the differences in behaviour between conditions. More work is needed to investigate if a personal incentive function could explain this.

The CEI model could have several applications related to automated driving technologies. It could be used online, in an automated vehicle, offline during the development of new technologies, or in scientific research to understand driver actions and communication better. A more extensive discussion on the possible applications of the model is presented in section 8.4.

8.2. General limitations

Every chapter reported and discussed the limitations specific to the work it described. However, there are two overarching limitations of the work in this thesis

–primarily regarding its scope– that deserve an extended discussion here. The first limitation concerns driver behaviour in the simulation environment and how it could be generalised to the real world. The second is related to the scope of the interactions that are studied.

The simulation used throughout this thesis consists of a simplified merging scenario. The major simplification in this scenario is the fact that only the vehicles' velocities can be controlled. There is no steering involved. As explained in Chapter 5, this simplification was needed to enable the analysis of the actions of an interacting pair of human drivers, which is a novel contribution of this thesis. Furthermore, the human participants in the experiment viewed the simulation in a top-down view. These are two simplifications to the scenario, and the task drivers were facing. Therefore, they could have affected the human behaviour discussed in Chapters 5 and 6.

The experimental scenario simplified the control task of the drivers by reducing the controlled degrees of freedom from two to one. The visual representation of the scenario was also simplified by reducing the two areas of interest in a first-person view (i.e., the road in front of the vehicle and the rear-view mirror) to a single top-down view containing all necessary information. These simplified controls and observations could have decreased the operational variability in the observed behaviour compared to real-world driving. However, the characteristics of the found behaviour likely resemble real-world traffic interactions because of the similarity of the real-world task, the task instructions (i.e., "this is a scientific experiment, not a race or a game"), the requirement that participants had to have a valid driver's license, and the gas pedals that were used as input devices.

The simplification of the control inputs had another effect. By reducing the degrees of freedom of the inputs, the driver's options regarding tactical decisions were also reduced. The only tactical decision a driver had to make was whether they go first. Other tactical behaviours that can be used in real-world scenarios, such as making a courtesy lane change or overtaking a vehicle before merging, were impossible in the simulation. Although the model successfully reproduced the tactical behaviour in this simplified scenario, it is an open question of how it will handle the larger tactical variability in real-world driving.

The second limitation is related to the type of traffic participants and the interaction in the simulation in general. This thesis only considered a highway merging interaction between two car drivers. This is only a small part of the traffic interactions present in the real world. These interactions can occur between many different types of traffic participants (e.g., pedestrians, (motor)cyclists, truck drivers, etc.) in many different situations (e.g., on the highway, an urban intersection, a roundabout, etc.). How the findings in this thesis generalise to other traffic participants and other situations has to be investigated.

In the specific case of a one-on-one highway merging interaction, the work only considered a subset of the kinematic possibilities. In the simulations in this thesis, the vehicles were always on a collision course at the start of the simulation (except for Scenario A in Chapter 4). While in real-life interactions, this is not always the case [7]. The underlying assumption is that driver behaviour in the more critical situations (i.e., with an impending collision) will generalise to the less critical situations. However, it is important to note that these situations were not so critical that only an emergency action could solve them. Whether this assumption of generalisation holds should be investigated.

8.3. Communication

The concept of communication in traffic and its role throughout this thesis also deserves some further discussion. The title of this thesis stems from the framework presented in Chapter 4 and indicates how big this role is. The proposed framework explicitly incorporates the possibility of communication between drivers in traffic interactions, which is a fundamental difference from the prevalent modelling approach using game theory where the absence of communication is one of the fundamental assumptions. When John Nash presented his famous equilibrium in 1951 [28], the fourth sentence of his paper stated:

"Our theory, in contradistinction [to von Neumann and Morgenstern [29]], is based on the absence of coalitions in that it is assumed that each participant acts independently, without collaboration or communication with any of the others."

Despite this fundamental assumption of the Nash equilibrium, the continuous observation of the state of the world in some game-theoretic models bears resemblance to the CEI-based models in Chapters 4 and 7. The main differences were already discussed in the discussion section of Chapter 4: "The communication in our framework allows drivers to construct and update a belief about the other vehicle's plan without the need for any prior information about the other driver" and "in game theory, observations are not remembered. They only serve to determine the current state, which is enough to reason about the other players' actions". So, the difference lies not only in the fundamental assumption but it propagates to the practical implementation of models.

However, it must be noted that the models presented in the case study in Chapter 4 and in Chapter 7 are only basic implementations of the CEI framework when it comes to communication: they do not leverage the full potential (and freedom) provided by the framework. These models only consider implicit communication and do not provide the possibility of actions that are performed purely to be communicative. Both models allow for communication through motion, without assuming that all motion is (meant as) communication. This makes it impossible to precisely identify which part of the (modelled) human behaviour is actually meant as communication. Therefore, I must conclude that the CEI framework holds much unused potential for future studies targeting formalized models of communication-enabled traffic interactions (as will be discussed in the following section).

8.4. Practical applications

The scientific work presented in this thesis could have practical applications. The main motivation behind the work was to improve interaction-aware autonomous driving, but other fields could also benefit from the results. The main practical potential lies with the proposed CEI model, but also the experimental framework and its analysis methods and the developed software have potential for further practical applications. All three aspects will be discussed in this section.

8.4.1. Applications of the model

There are many possible extensions and potential applications of the model that could be worth investigating. This section will outline some possible practical and scientific applications of the model. An overview of the potential applications of the model is shown in Figure 8.2. This section will follow the structure with three application areas mentioned in the figure (Autonomous Driving (AD), development, and scientific applications) to discuss the potential use cases.

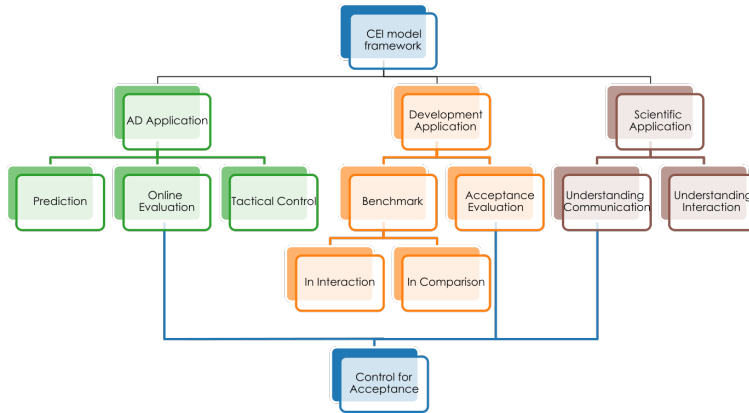


Figure 8.2: An overview of the potential applications of the model. This figure shows three potential application areas: online application in Autonomous Driving (AD), offline application during the development of controllers, and scientific applications. Each has multiple potential use cases for the model. Combined, some of these applications could lead to *control for acceptance*.

One important remark in this discussion of future applications is that the CEI model is a model of human behaviour. It was designed to capture and predict human behaviour, including all aspects of human behaviour that are undesirable in AV behaviour (such as collisions in the simulated environment). Therefore, the model should not be used for direct control in automated driving. Instead, it could be leveraged to predict and understand how humans interact in traffic to create better AV behaviour.

AD application

The first application area of interest is the online Autonomous Driving (AD) application. This encompasses all applications of the model where the model is used onboard an intelligent vehicle. Based on existing literature, I formulated three potential use cases.

Prediction The first potential use case is to make predictions about other drivers to inform an existing AD controller. This would be comparable to the interaction-aware controllers discussed earlier (e.g., [11], [30]–[32]). To achieve this, the model must be split into two separate drivers. One of the two will be swapped with the vehicle's controller. The predicted response of the other vehicle could be obtained by inserting the controller's future plan into the model and simulating it for a short time. The controller can then use this prediction to make better decisions. There are two main drawbacks to this approach. First, because the controller's future plan does not depend on the other participant's observed (future) communication, this would mean making the one-way interaction (Chapter 4), defying the whole reason for the development of the model. Second, most existing interaction-aware controllers use game theory to efficiently construct a plan based on the predictions. However, since the model abandoned the most important assumptions of game theory (rationality and no communication), finding a feasible plan based on the prediction will be mathematically and computationally challenging.

The main advantage of using the proposed model over the existing approaches is that it could provide better predictions of human behaviour during interactions. As demonstrated in Chapter 2, this could mean a substantial improvement compared to existing approaches.

Online evaluation Another potential online application of the model would be to evaluate an existing AV controller's future plan. Instead of directly using the model to obtain predictions of the other drivers' future actions, we could use the model to predict their perceived risk. This different prediction would allow an AV controller to reason about how their behaviour is perceived. This way, an AV controller could actively try to minimise the perceived risk by passengers and other road users.

Another option would be to adapt the AV behaviour (online) to actively minimise the Conflict Resolution Time (CRT). To the best of our knowledge, using a model to control for such high-level interaction metrics has not been proposed before for interactions between vehicles. More research is needed to determine which metrics would be useful to predict. How to alter the AV's plan to optimise for these metrics is also an open question.

8

Tactical control The final proposal for an online application of the model is based on a paper by Fisac et al. published in 2019 [33]. In this work, they propose to use a hierarchical controller for autonomous vehicles. A low-level, high-frequency controller is used for operational control. This controller works much like traditional AV controllers. It maximises safety while providing passengers with a fast and comfortable ride. On top of this operational controller, Fisac proposes to use a tactical control loop. This is a more high-level control loop that runs at a lower frequency.

The job of this tactical control loop is to make higher-level driving decisions in a human-like way, which is hypothesised to increase acceptance. a CEI model could take the job of this tactical control loop. For example, in a merging scenario, the tactical control loop can decide to yield to another vehicle or decide into which gap to merge. This information is passed on to the operational control loop, which tries to stick to the plan of the tactical loop yet prioritises safety. This would allow for safe AV control while adhering to human-like interactive behaviour whenever possible.

Development application

Besides using the model onboard intelligent vehicles, there is potential in using the model during the development of new automated driving technologies. This includes use in laboratories and driving simulators during the development of new AV controllers. These applications show similarities with onboard use, but the main difference is that the simulation outcomes here are used to simulate human behaviour, not to predict it.

Benchmark Benchmark testing for autonomous vehicles is a topic for which multiple driver models were specifically proposed (e.g., [34], [35]). The general idea is to create a simulation with "human" traffic participants whose behaviour is generated by models to test newly developed AV controllers. The main advantage of this method is that it is possible to test new controllers in a repeatable, time and cost-efficient way without the need for extensive human-in-the-loop experiments.

The CEI model could be used for benchmark testing in two ways. First, in the traditional way for generating human behaviour. The main advantage over existing models could be that our model was specifically developed for interactive scenarios and thus might produce more realistic interactive behaviour. Second, the model could be used to compare the controller's behaviour to generated human behaviour. This would allow developers to evaluate to what extent the controller's behaviour is human-like in interactive traffic. Such an evaluation could give valuable insights into how to improve the controller.

Acceptance evaluation The evaluation of acceptance metrics, such as perceived risk and CRT, was already discussed for online applications. This same idea could also be applied during development. Instead of using the outcome to improve the controller's plan directly, it could be used to improve the design of the controller. For example, suppose two AV controllers are designed specifically to target merging scenarios. Using the CEI model, developers could easily create hundreds or thousands of hypothetical merging scenarios. The model can produce metrics such as CRT or perceived risk in interaction with the new controllers. These metrics can then be used to evaluate the differences between the controllers. This would allow engineers to select and improve the best-performing controller.

Scientific application

Finally, there are specific opportunities for model applications in scientific research. First and foremost, the model can help understand the communication between traffic participants. Communication is an important part of both the model and human driving. The model can help us understand how humans translate their plans into communicative actions and how they observe such actions to form a belief. This can be done by letting humans interact with different versions of the model and by trying to replicate human behaviour in simulations.

A good understanding of drivers' communicative actions could greatly help design better AV control algorithms. Other traffic participants will better understand an AV controller that can communicate its plan to others. This could mean it will be accepted more easily. Furthermore, an AV controller that can understand the communication sent by other traffic participants can act accordingly, which could also help to be accepted. However, previous models mostly regard the behaviour of a single driver without considering the possibility of them communicating through their actions. The main advantage of using the CEI model over combining two existing models is that the CEI framework explicitly includes converting a plan to a communicative action and converting this communication to a belief. This allows the study of these mechanisms. Current approaches to understanding communication in traffic mostly rely on the use of naturalistic data from real traffic (e.g., [24]), making it difficult to obtain (causal) relationships between drivers' beliefs, plans, and communication, and impossible to gain insight into a driver's belief.

Besides understanding communication, the model could help us understand human traffic interactions in general. There are many potential research directions to be pursued. Potential research questions could be: What is the influence of different acceleration profiles? (E.g., does the asymptotic velocity profile often used in reward functions have implications?) How do risk thresholds relate to other drivers' perceptions of the interaction? Can we measure perceived acceptance with a metric other than interviews/surveys?

Control for Acceptance

By leveraging the proposed applications for online and offline evaluation and understanding communication, a higher-level research goal could be to develop *control for acceptance*. Control for acceptance means developing AV controllers that actively optimise their plan to be accepted by their passengers and other road users. This could be done alongside the current optimisation for safety and comfort. By actively increasing acceptance, AV behaviour could be better understood and, therefore safer.

The three proposed applications are needed to develop control for acceptance for different reasons. Online evaluation will teach us how to optimise the plan of a controller for acceptance metrics, such as perceived risk. Offline evaluation will provide the opportunity to experiment with many controllers, metrics and strategies to increase acceptance. Finally, understanding communication will allow the design of controllers that can efficiently communicate with other drivers, which is most likely a key aspect of being accepted. By aiming for control for acceptance, we hope the CEI framework can eventually contribute to safe and more acceptable AV behaviour.

8

8.4.2. Experimental framework and CRT

The experimental framework presented in Chapter 5 and used in Chapter 6 proved to be a useful way to study traffic interactions. It enabled controlled experiments in a coupled simulator, providing insight into the underlying mechanisms of merging interactions. The simplification to a single input dimension allowed the usage of multiple analysis tools (see Chapter 5). More insight into human interactions could be obtained with this experiment and these tools. For example, the generalisability of the findings in this thesis could be investigated by extending the experiment to other interactive scenarios, such as intersections or roundabouts.

One of the new analysis tools is the Conflict Resolution Time (CRT). This metric provides insight into the duration of a conflict and, therefore, into the difficulty of solving the conflict. Currently, the CRT is only defined in the simplified scenario, but if it can be extended to real-world scenarios with longitudinal and lateral control, there are three potential use cases.

The CRT metric could prove a useful metric in the design of autonomous behaviour. It can serve as a performance metric in benchmark tests. For example, it could be used to compare multiple controllers to see which version solves conflicts the fastest. But also to compare AV behaviour to human behaviour, especially since we found a correlation between CRT and kinematic conditions in Chapter 6. AVs might be able to replicate this behaviour or improve it. Besides benchmarks, there might be a possibility of actively controlling an automated vehicle based on the CRT (i.e., the duration of the conflict). However, how other drivers and AV passengers perceive such optimized behaviour is unknown and should be investigated.

There is also potential for using CRT in traffic control if the relationship between vehicle kinematics and CRT is known for real-world traffic interactions. Signalized intersections and highway on-ramps, or vehicle-to-infrastructure communication schemes could be used to control vehicle kinematics to lower the CRT actively. This lowers the duration of the conflicts and could therefore potentially increase safety, decrease control effort, and increase traffic flow.

8.4.3. Software and data

Large parts of this thesis rely on software simulations or the processing of naturalistic data. Specific software has been developed for both use cases. To increase the potential for re-use of the work, all software developed for the work in this thesis has been published online ¹.

The most extensive software package has also been published as a paper which is included in Appendix A. This package is called TraViA, which stands for Traffic Visualisation and Annotation tool. It has been used in Chapters 2 and 3. TraViA was developed to visualise and annotate open naturalistic datasets. It supports the HighD [36], ExiD [37], PNeuma [38], and NGSim [39] datasets. Extension packages for TraViA, which provide implementations of the IRL validation framework of Chapter 2, and the Hausdorff method of Chapter 3, are available.

Besides TraViA, the environment to conduct the simplified merging experiment has been published online, along with scripts for the analysis tools of Chapter 5. All scripts to reproduce the statistical analysis and plots of Chapter 6 are available online. Finally, the implementations of both versions of the CEI model (Chapters 4 and 7) and its simulation environment can be found online.

All data gathered in the driving simulator experiment and the model simulations are available on the 4TU data repository. Please see the respective chapters for links to all data and software repositories.

8.5. Further outlook

Besides the thesis' conclusions, immediate possibilities for future work, and the potential applications and impact of the CEI model, three topics deserve further discussion. First, the distinction between levels of driver behaviour in rationality and control; second, the commonly-used underlying assumptions of modelling variability in driver behaviour; and third, the usage of naturalistic data to understand individual interactions. These topics have been touched upon in this thesis but were not the main subjects of discussion. However, these three topics pose difficult open problems that need to be solved and understood to reach a level of driving automation (and maybe robotic automation in general) where technology is truly capable of handling non-verbal contact-less interactions with humans. Therefore, they deserve a (brief) discussion on the final pages of this thesis.

Levels of behaviour The role of tactical and operational behaviour in driver model validation has been discussed extensively in previous chapters. However, other aspects of driver modelling might benefit from an analysis from the perspective of levels of behaviour as well. Two –in particular– have been touched upon but have not been discussed or explored before: the role of levels of behaviour in human rationality and in the variability within predictions.

The scientific discussion of whether humans behave rationally –and if so, to what extent– has been going on for decades and is still not settled. The debate stems from an early assumption in economics and game theory that humans are rational. These early studies mostly considered high-level choices: for example, if a human would buy product A or B or if a player would play strategy X or Y. This statement already reveals that these original studies mostly considered strategic behaviour. This is reflected in early game theoretic driving interaction models (e.g., Kita's model from 1999 [40]). These early models showed promising results

¹github.com/tud-hri

in explaining human one-shot decision-making in traffic, arguing in favour of the rationality assumption regarding tactical driving behaviour.

However, when the focus is shifted to operational behaviour, as in this thesis, humans do not seem to optimize their behaviour continuously (see Chapter 6 or [14], [16]). Based on these studies, it is tempting to conclude that humans do not behave rationally, period. However, the diversity in levels of behaviour between studies might (partly) explain why the rationality debate is still not settled. A hypothesis could be that humans behave more like rational utility maximisers when making high-level tactical decisions (where the implications of a decision are often easier to comprehend), but that this rationality declines when operational behaviour is considered. More fundamental research into this topic could increase our understanding of driving behaviour and human behaviour in general. This could allow for models that make different assumptions for the different levels of behaviour.

Another aspect that might benefit from existing knowledge on levels of behaviour is the predictions that AVs make of human driver behaviour. As discussed before, traditional driver models mostly describe the operational behaviour of drivers for a specific tactical behaviour. Interaction-aware AVs need to consider multiple tactical responses and, therefore, mostly rely on predictions made by machine-learned models, such as the inverse-reinforcement-learning-based model in Chapter 2. These machine-learned models are usually agnostic to the levels of behaviour. The inverse-reinforcement-learning-based model in Chapter 2 is a good example. This model learns a single reward function that has to explain all human behaviour. Other examples are deep-learning-based models that predict multiple future trajectories (e.g., [8], [10]) without distinguishing between tactical and operational variations in these trajectories. Incorporating knowledge of the levels of behaviour could potentially improve the predictions of these machine-learned driver models for two reasons: first, human operational behaviour differs for different tactical behaviours (Chapter 2), and second, there is a lot of literature on operational behaviour within specific tactical behaviours (Chapter 2).

Assumptions of behavioural variability Another potential benefit of incorporating levels of behaviour in machine-learned driver models would be to add structure to the underlying assumptions of variability. Many deep-learning-based approaches assume that the variability in behaviour can be modelled with parametric distributions (often a Gaussian distribution) (e.g., [8], [10]). This assumption is agnostic to the levels of behaviour. However, the results of the case study in Chapter 3 (specifically Figure 3.5) could be interpreted as a sign that this assumption will not hold for all tactical behaviours. Especially for lane-changing vehicles, the trajectory waypoints at the same points in time (measured from the selected scene) do not seem to be normally distributed. This is, of course, only an example from a single case study on a naturalistic dataset, which poses no definitive proof that the assumption of Gaussian distributions is invalid. It should merely be interpreted as an encouragement for more research into the underlying distributions of behavioural variability for different tactical behaviours.

Naturalistic datasets The same chapter (Chapter 3) also illustrated the difficulties of working with naturalistic data. We needed to develop a sophisticated method to compare human responses to similar traffic scenes and get insight into behavioural variability in real traffic data. However, this chapter also showed that

naturalistic data can be more than just training data for machine-learned models. It can be a rich study ground to understand driver behaviour in the real world better. While there have been extensive empirical studies into the traffic system as a whole [5], [41] (mostly focused on traffic flow), the current literature lacks tools to study individual interactions between drivers in naturalistic data. In this thesis, we made a start with developing such tools for simulator studies (Chapter 5) and naturalistic data (Chapter 3). But there is much more work to be done on this front before we can fully leverage the large amounts of naturalistic data recently published [36], [38], [39].

8.6. Final remarks

To conclude, although much more work remains to be done on the model proposed in this thesis, I hope it will prove to be a substantial step towards interaction-aware automated driving and interaction-aware robotics in general. Creating automation that can interact with humans in a natural and safe manner is one of the major modern-day challenges in robotics. The ability to interact will make or break the potential for robots to make an impact in the real world, not only for automated driving but also for robots in public spaces and the workplace. Even though this thesis only considered a very small subset of human-robot interactions, that of interactions on the road, it outlined the importance of understanding and modelling human behaviour to improve our technology. We can still learn a lot from –and about– “the human controller”.

Bibliography

- [1] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, vol. 62, no. 2, pp. 1805–1824, 2000, ISSN: 1063651X. DOI: [10.1103/PhysRevE.62.1805](https://doi.org/10.1103/PhysRevE.62.1805). arXiv: 0002177 [cond-mat].
- [2] A. Economou, I. Beratis, E. Papadimitriou, G. Yannis, and S. G. Papageorgiou, "Intraindividual variability in driving simulator parameters of healthy drivers of different ages", *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 78, pp. 91–102, 2021, ISSN: 13698478. DOI: [10.1016/j.trf.2021.02.002](https://doi.org/10.1016/j.trf.2021.02.002). [Online]. Available: <https://doi.org/10.1016/j.trf.2021.02.002>.
- [3] J. M. Cooper, N. Medeiros-Ward, and D. L. Strayer, "The impact of eye movements and cognitive workload on lateral position variability in driving", *Human Factors*, vol. 55, no. 5, pp. 1001–1014, 2013, ISSN: 00187208. DOI: [10.1177/0018720813480177](https://doi.org/10.1177/0018720813480177).
- [4] S. Kolekar, W. Mugge, and D. Abbink, "Modeling intradriver steering variability based on sensorimotor control theories", *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 3, pp. 291–303, 2018, ISSN: 21682291. DOI: [10.1109/THMS.2018.2812620](https://doi.org/10.1109/THMS.2018.2812620).
- [5] F. Marczak, W. Daamen, and C. Buisson, "Merging behaviour: Empirical comparison between two sites and new theory development", *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 530–546, 2013, ISSN: 0968090X. DOI: [10.1016/j.trc.2013.07.007](https://doi.org/10.1016/j.trc.2013.07.007). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2013.07.007>.
- [6] S. Mozaffari, O. Y. Al-Jarrah, M. Dianati, P. Jennings, and A. Mouzakitis, "Deep Learning-Based Vehicle Behavior Prediction for Autonomous Driving Applications: A Review", *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 1, pp. 33–47, 2022, ISSN: 15580016. DOI: [10.1109/TITS.2020.3012034](https://doi.org/10.1109/TITS.2020.3012034). arXiv: 1912.11676.
- [7] A. R. Srinivasan, M. Hasan, Y.-S. Lin, et al., *Comparing merging behaviors observed in naturalistic data with behaviors generated by a machine learned model*, 2021. arXiv: [2104.10496](https://arxiv.org/abs/2104.10496) [cs.LG].
- [8] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-Feasible Trajectory Forecasting With Heterogeneous Data", 2020. arXiv: [2001.03093](https://arxiv.org/abs/2001.03093). [Online]. Available: [http://arxiv.org/abs/2001.03093](https://arxiv.org/abs/2001.03093).
- [9] B. Brito, H. Zhu, W. Pan, and J. Alonso-Mora, "Social-VRNN: One-Shot Multi-modal Trajectory Prediction for Interacting Pedestrians", *Proceedings of Machine Learning Research*, vol. 155, no. CoRL, pp. 862–872, 2020, ISSN: 26403498. arXiv: [2010.09056](https://arxiv.org/abs/2010.09056).
- [10] A. Mészáros, J. F. Schumann, J. Alonso-Mora, A. Zgonnikov, and J. Kober, "TrajFlow: Learning the Distribution over Trajectories", 2023. arXiv: [2304.05166](https://arxiv.org/abs/2304.05166). [Online]. Available: <http://arxiv.org/abs/2304.05166>.
- [11] D. Sadigh, N. Landolfi, S. S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state", *Autonomous Robots*, vol. 42, no. 7, pp. 1405–1426, Oct. 2018, ISSN: 0929-5593. DOI: [10.1007/s10514-018-9746-1](https://doi.org/10.1007/s10514-018-9746-1). [Online]. Available: <https://doi.org/10.1007/s10514-018-9746-1> <http://link.springer.com/10.1007/s10514-018-9746-1>.
- [12] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and Decision-Making for Autonomous Vehicles", *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 187–210, 2018, ISSN: 2573-5144. DOI: [10.1146/annurev-control-060117-105157](https://doi.org/10.1146/annurev-control-060117-105157).
- [13] C. F. Camerer, *Behavioral game theory: Experiments in strategic interaction*. 2003, ISBN: 0691090394. DOI: [10.1016/j.socec.2003.10.009](https://doi.org/10.1016/j.socec.2003.10.009).
- [14] M. Schmidt-Daffy, "Prospect balancing theory: Bounded rationality of drivers' speed choice", *Accident Analysis and Prevention*, vol. 63, pp. 49–64, 2014, ISSN: 00014575. DOI: [10.1016/j.aap.2013.10.028](https://doi.org/10.1016/j.aap.2013.10.028). [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2013.10.028>.
- [15] H. A. Simon, "A Behavioral Model of Rational Choice", *The Quarterly Journal of Economics*, vol. 69, no. 1, p. 99, 1955, ISSN: 00335533. DOI: [10.2307/1884852](https://doi.org/10.2307/1884852).

- [16] M. A. Goodrich, W. C. Stirling, and E. R. Boer, "Satisficing revisited", *Minds and Machines*, vol. 10, no. 1, pp. 79–110, 2000, ISSN: 09246495. DOI: [10.1023/A:1008325423033](https://doi.org/10.1023/A:1008325423033).
- [17] E. R. Boer, "Car following from the driver's perspective", *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 2, no. 4, pp. 201–206, 1999, ISSN: 13698478. DOI: [10.1016/S1369-8478\(00\)00007-3](https://doi.org/10.1016/S1369-8478(00)00007-3).
- [18] H. A. Simon, "Rational choice and the structure of the environment.", *Psychological Review*, vol. 63, no. 2, pp. 129–138, 1956, ISSN: 1939-1471. DOI: [10.1037/h0042769](https://doi.org/10.1037/h0042769). [Online]. Available: <http://doi.apa.org/getdoi.cfm?doi=10.1037/xge0000013%20http://doi.apa.org/getdoi.cfm?doi=10.1037/h0042769>.
- [19] P. Gawthrop, I. Loram, M. Lakie, and H. Gollee, "Intermittent control: A computational theory of human control", *Biological Cybernetics*, vol. 104, no. 1-2, pp. 31–51, 2011, ISSN: 03401200. DOI: [10.1007/s00422-010-0416-4](https://doi.org/10.1007/s00422-010-0416-4).
- [20] I. D. Loram, H. Gollee, M. Lakie, and P. J. Gawthrop, "Human control of an inverted pendulum: Is continuous control necessary? Is intermittent control effective? Is intermittent control physiological?", *Journal of Physiology*, vol. 589, no. 2, pp. 307–324, 2011, ISSN: 00223751. DOI: [10.1113/jphysiol.2010.194712](https://doi.org/10.1113/jphysiol.2010.194712).
- [21] G. Markkula, E. Boer, R. Romano, and N. Merat, "Sustained sensorimotor control as intermittent decisions about prediction errors: computational framework and application to ground vehicle steering", *Biological Cybernetics*, vol. 112, no. 3, pp. 181–207, 2018, ISSN: 14320770. DOI: [10.1007/s00422-017-0743-9](https://doi.org/10.1007/s00422-017-0743-9). arXiv: [1703.03030](https://arxiv.org/abs/1703.03030). [Online]. Available: <https://doi.org/10.1007/s00422-017-0743-9>.
- [22] U. Durrani, "A New Car-Following Model with Incorporation of Markkula's Framework of Sensorimotor Control in Sustained Motion Tasks", *University of Windsor*, pp. 1–207, 2022. [Online]. Available: <https://www-proquest-com.ledproxy2.uwindsor.ca/docview/2664805698?pq-origsite=primo>.
- [23] G. Markkula, "Modeling driver control behavior in both routine and near-accident driving", *Proceedings of the Human Factors and Ergonomics Society*, vol. 2014-Janua, pp. 879–883, 2014, ISSN: 10711813. DOI: [10.1177/1541931214581185](https://doi.org/10.1177/1541931214581185).
- [24] Y. M. Lee, R. Madigan, O. Giles, *et al.*, "Road users rarely use explicit communication when interacting in today's traffic: implications for automated vehicles", *Cognition, Technology & Work*, vol. 23, no. 2, pp. 367–380, May 2021, ISSN: 1435-5558. DOI: [10.1007/s10111-020-00635-y](https://doi.org/10.1007/s10111-020-00635-y). [Online]. Available: <https://doi.org/10.1007/s10111-020-00635-y>.
- [25] S. Kolekar, J. de Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model", *Nature Communications*, vol. 11, no. 1, p. 4850, Dec. 2020, ISSN: 2041-1723. DOI: [10.1038/s41467-020-18353-4](https://doi.org/10.1038/s41467-020-18353-4). [Online]. Available: <http://dx.doi.org/10.1038/s41467-020-18353-4> <https://www.nature.com/articles/s41467-020-18353-4>.
- [26] S. Kolekar, B. Petermeijer, E. Boer, J. de Winter, and D. Abbink, "A risk field-based metric correlates with driver's perceived risk in manual and automated driving: A test-track study", *Transportation Research Part C: Emerging Technologies*, vol. 133, no. October, p. 103428, 2021, ISSN: 0968090X. DOI: [10.1016/j.trc.2021.103428](https://doi.org/10.1016/j.trc.2021.103428). [Online]. Available: <https://doi.org/10.1016/j.trc.2021.103428>.
- [27] S. Kolekar, J. de Winter, and D. Abbink, "Which parts of the road guide obstacle avoidance? Quantifying the driver's risk field", *Applied Ergonomics*, vol. 89, no. July, p. 103196, Nov. 2020, ISSN: 00036870. DOI: [10.1016/j.apergo.2020.103196](https://doi.org/10.1016/j.apergo.2020.103196). [Online]. Available: <https://doi.org/10.1016/j.apergo.2020.103196> <https://linkinghub.elsevier.com/retrieve/pii/S0003687018307373>.
- [28] J. Nash, "Non-Cooperative Games", *The Annals of Mathematics*, vol. 54, no. 2, p. 286, Sep. 1951, ISSN: 0003486X. DOI: [10.2307/1969529](https://doi.org/10.2307/1969529). [Online]. Available: <https://www.jstor.org/stable/1969529?origin=crossref>.
- [29] J. von Neumann and O. Morgenstern, *Theory of games and economic behavior*. 1944.

- [30] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [31] N. Evestedt, E. Ward, J. Folkesson, and D. Axehill, "Interaction aware trajectory planning for merge scenarios in congested traffic situations", *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, pp. 465–472, 2016. DOI: [10.1109/ITSC.2016.7795596](https://doi.org/10.1109/ITSC.2016.7795596).
- [32] E. Ward, N. Evestedt, D. Axehill, and J. Folkesson, "Probabilistic Model for Interaction Aware Planning in Merge Scenarios", *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 2, pp. 1–1, 2017, ISSN: 2379-8904. DOI: [10.1109/TIV.2017.2730588](https://doi.org/10.1109/TIV.2017.2730588). [Online]. Available: <http://ieeexplore.ieee.org/document/7987753/>.
- [33] J. F. Fisac, E. Bronstein, E. Stefansson, D. Sadigh, S. S. Sastry, and A. D. Dragan, "Hierarchical game-theoretic planning for autonomous vehicles", *IEEE International Conference on Robotics and Automation*, vol. 2019-May, pp. 9590–9596, 2019, ISSN: 10504729. DOI: [10.1109/ICRA.2019.8794007](https://doi.org/10.1109/ICRA.2019.8794007). arXiv: [1810.05766](https://arxiv.org/abs/1810.05766).
- [34] R. Queiroz, D. Sharma, R. Caldas, et al., "A Driver-Vehicle Model for ADS Scenario-based Testing", pp. 1–16, May 2022. arXiv: [2205.02911](https://arxiv.org/abs/2205.02911). [Online]. Available: <http://arxiv.org/abs/2205.02911>.
- [35] N. Li, D. W. Oyler, M. Zhang, Y. Yildiz, I. Kolmanovsky, and A. R. Girard, "Game theoretic modeling of driver and vehicle interactions for verification and validation of autonomous vehicle control systems", *IEEE Transactions on Control Systems Technology*, vol. 26, no. 5, pp. 1782–1797, 2018, ISSN: 10636536. DOI: [10.1109/TCST.2017.2723574](https://doi.org/10.1109/TCST.2017.2723574). arXiv: [1608.08589](https://arxiv.org/abs/1608.08589).
- [36] R. Krajewski, J. Bock, L. Kloeker, and L. 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)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 2118–2125, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569552](https://doi.org/10.1109/ITSC.2018.8569552). arXiv: [1810.05642](https://arxiv.org/abs/1810.05642). [Online]. Available: <https://ieeexplore.ieee.org/document/8569552/>.
- [37] T. Moers, L. Vater, R. Krajewski, J. Bock, A. Zlocki, and L. Eckstein, "The exid dataset: A real-world trajectory dataset of highly interactive highway scenarios in germany", in *2022 IEEE Intelligent Vehicles Symposium (IV)*, 2022, pp. 958–964. DOI: [10.1109/IV51971.2022.9827305](https://doi.org/10.1109/IV51971.2022.9827305).
- [38] E. Barmponakis and N. Geroliminis, "On the new era of urban traffic monitoring with massive drone data: The pNEUMA large-scale field experiment", *Transportation Research Part C: Emerging Technologies*, vol. 111, no. November 2019, pp. 50–71, 2020, ISSN: 0968090X. DOI: [10.1016/j.trc.2019.11.023](https://doi.org/10.1016/j.trc.2019.11.023). [Online]. Available: <https://doi.org/10.1016/j.trc.2019.11.023>.
- [39] U.S. Department of Transportation Federal Highway Administration, *Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data. [Dataset]*, 2016. [Online]. Available: <https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajectory/8ect-6jqj>.
- [40] H. Kita, "A merging-giveway interaction model of cars in a merging section: A game theoretic analysis", *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3-4, pp. 305–312, 1999, ISSN: 09658564. DOI: [10.1016/s0965-8564\(98\)00039-1](https://doi.org/10.1016/s0965-8564(98)00039-1).
- [41] W. Daamen, M. Loot, and S. P. Hoogendoorn, "Empirical analysis of merging behavior at freeway on-ramp", *Transportation Research Record*, no. 2188, pp. 108–118, 2010, ISSN: 03611981. DOI: [10.3141/2188-12](https://doi.org/10.3141/2188-12).


Appendices

The background of the page is a vibrant blue sky filled with soft, white, wispy clouds. In the bottom-left corner, there is a cluster of several orange circles of varying sizes, some of which are partially cut off by the edge of the page. A white diagonal line is also visible in the bottom-left corner, extending from the edge towards the center.



A

TraViA: a Traffic data
Visualization and Annotation
tool in Python



In recent years, multiple datasets containing traffic recorded in the real world and containing human-driven trajectories have been made available to researchers. Among these datasets are the HighD, pNEUMA, and NGSIM datasets. TraViA, an open-source Traffic data Visualization and Annotation tool was created to provide a single environment for working with data from these three datasets. Combining the data in a single visualization tool enables researchers to easily study data from all sources. TraViA was designed in such a way that it can easily be extended to visualize data from other datasets and that specific needs for research projects are easily implemented.

A.1. Statement of need

The combination of drones, cameras, and image recognition techniques might sound like a recipe for a spy movie. But actually, this combination allows for the collection of rich traffic datasets. The recipe is straightforward: hover a drone above a location with traffic, record a video, and use image recognition to generate bounding boxes for all vehicles. The result is a dataset containing human-driven trajectories at the location of interest that can be used for many scientific purposes, e.g., to study traffic flow, model human behavior, or design autonomous vehicle controllers.

Because the required ingredients are easily accessed all over the world, multiple such datasets have been published in recent years. In Germany, the highD project [1] recorded all traffic at 6 different high-way locations; in Athens, Greece, all traffic in the city's business district was recorded using 10 drones for 5 days in the pNEUMA project [2]; and American highway traffic was recorded using fixed base cameras in the NGSIM project [3]. Combined, these datasets span different countries, types of vehicles, and environments, a combination valuable for researchers with different backgrounds. Example usages of these datasets are validating human behavior models (e.g., [4] and [5]) or testing autonomous vehicle controllers (e.g., [6]).

Currently, it is difficult to leverage the powerful combination of multiple datasets because all the datasets come in different formats, and it is often difficult to get a good and real-time visualization of the data. Some visualization tools exist (one is provided with the highD data [1] and another example for NGSIM data can be found in [7]) but they are specifically made for only one of these datasets and are very basic in the sense that they provide little control over simulation time and no insight in raw values per vehicle per frame. In addition to difficulties with visualization, finding, and annotating situations of interest in these massive datasets is a time-consuming task and keeping track of the annotations for the different datasets requires some bookkeeping skills.

TraViA was developed to provide a solution for these problems. TraViA can be used to visualize and annotate data from highD, pNEUMA, and NGSIM and uses generic vehicle objects to store the state of vehicles at a specific time. This makes it possible to validate and test models or controllers on multiple datasets in parallel, without having to cope with the different dataset formats.

A.2. Software functionality

TraViA is written in Python 3 and has a graphical user interface developed in PyQt5. A screenshot of TraViA is provided in Figure A.1. This screenshot shows the capabilities of TraViA in a single image. The main features of TraViA are:

- Advanced information display based on raw data for every vehicle in every dataset by leveraging generic vehicle objects
- Dynamic visualization of the traffic scene with possibilities to zoom, pan, and rotate for an optimal view
- Exporting the visualization to a video or single image
- An interactive timeline that shows dataset annotations, which are saved as Python objects for easy manipulation

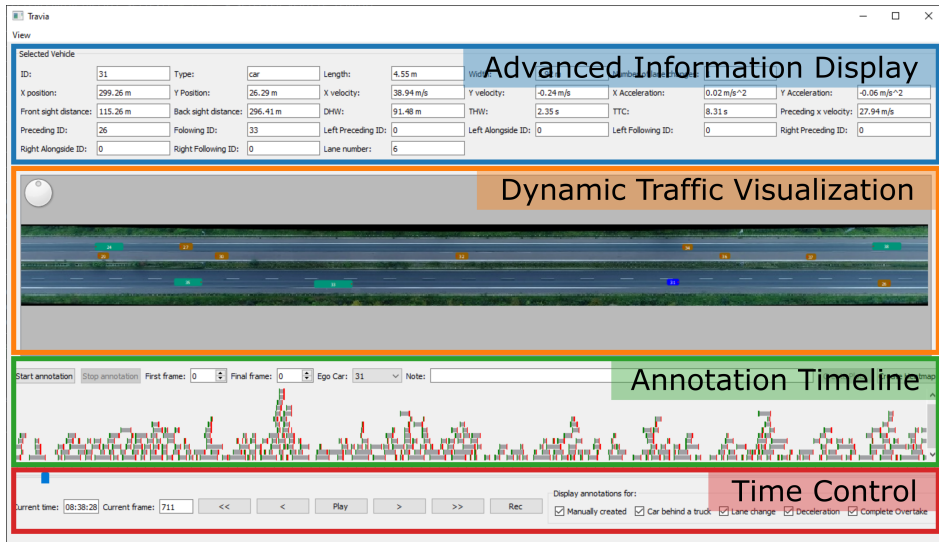


Figure A.1: A screenshot of the TraViA software visualizing a frame of the highD dataset. The main features of TraViA are highlighted in this image.

TraViA was designed for use as a stand-alone program. It uses abstract classes as a basis for all dataset-specific objects to enable easy implementation of new datasets (for a class diagram and more information on how to do this, please see the readme file in the repository). It was specifically created to serve as a tool for generic visualization and annotation such that it can be used by researchers from different fields. To show the capabilities of TraViA and to provide a starting point for other researchers who want to use TraViA for their work, three example implementations of tools for specific purposes are included with TraViA. The first example is the functionality to automatically detect and annotate particular scenarios (e.g., lane changes), the second is the functionality to plot specific vehicle signals over the course of an annotation, and the third is a function to plot a heatmap overlay for use in autonomous vehicle reward function development. All of these example tools are only implemented for use with the highD dataset.

A.3. Usage of TraViA in science and education

Currently, TraViA is being used by the author for model validation of an inverse-reinforcement-learning-based driver model. A publication on this validation is currently being prepared for submission. Besides that, TraViA is used for educational purposes, allowing students at TU Delft to explore big naturalistic datasets by providing them with an accessible, GUI-based starting point.

Acknowledgements

I thank Joris Giltay for his help with testing the instructions in the readme file.

Bibliography

- [1] R. Krajewski, J. Bock, L. Kloeker, and L. 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)*, vol. 2018-Novem, IEEE, Nov. 2018, pp. 2118–2125, ISBN: 978-1-7281-0321-1. DOI: [10.1109/ITSC.2018.8569552](https://doi.org/10.1109/ITSC.2018.8569552). arXiv: [1810.05642](https://arxiv.org/abs/1810.05642). [Online]. Available: <https://ieeexplore.ieee.org/document/8569552/>.
- [2] E. Barmounakis and N. Geroliminis, "On the new era of urban traffic monitoring with massive drone data: The pNEUMA large-scale field experiment", *Transportation Research Part C: Emerging Technologies*, vol. 111, no. November 2019, pp. 50–71, 2020, ISSN: 0968090X. DOI: [10.1016/j.trc.2019.11.023](https://doi.org/10.1016/j.trc.2019.11.023). [Online]. Available: <https://doi.org/10.1016/j.trc.2019.11.023>.
- [3] U.S. Department of Transportation Federal Highway Administration, *Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data. [Dataset]*, 2016. [Online]. Available: <https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajectory/8ect-6jqj>.
- [4] A. Talebpour, H. S. Mahmassani, and S. H. Hamdar, "Modeling lane-changing behavior in a connected environment: A game theory approach", *Transportation Research Part C: Emerging Technologies*, vol. 59, pp. 216–232, 2015, ISSN: 0968090X. DOI: [10.1016/j.trc.2015.07.007](https://doi.org/10.1016/j.trc.2015.07.007). [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2015.07.007>.
- [5] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, vol. 62, no. 2, pp. 1805–1824, 2000, ISSN: 1063651X. DOI: [10.1103/PhysRevE.62.1805](https://doi.org/10.1103/PhysRevE.62.1805). arXiv: [0002177 \[cond-mat\]](https://arxiv.org/abs/0002177).
- [6] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles", *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, Dec. 2019, ISSN: 0027-8424. DOI: [10.1073/pnas.1820676116](https://doi.org/10.1073/pnas.1820676116). [Online]. Available: <http://www.pnas.org/lookup/doi/10.1073/pnas.1820676116>.
- [7] C. Sazara, *NGSIM-trajectory-animation*, <https://github.com/cemsaz/NGSIM-trajectory-animation>, 2017. (visited on 03/31/2021).

Acknowledgements

Zoals de verteller van Spongebob zou zeggen: "Vier jaar later" (en dat kon je niet zonder dat accent lezen), zo voelt het nu een beetje. Vier jaar geleden besloot ik om terug te komen naar de TU Delft om aan een promotieonderzoek te beginnen. Toen leek vier jaar om dat onderzoek af te ronden een hele lange tijd. Een week later werd de eerste covid-lockdown plotseling van kracht en leek het even alsof het echt een eeuwigheid zou gaan duren. Maar uiteindelijk zijn de afgelopen vier jaar voorbij gevlogen. Sinds ik David liet weten dat ik het ging doen heb ik geen moment spijt gehad. Een leuk onderzoek, leuke onderwijstaken, en leuke mensen om me heen; zoals mijn moeder zou zeggen: "wat wil een mens nog meer".

David, ik begin natuurlijk bij jou. Iets meer dan zeven jaar geleden was ik op zoek naar een begeleider voor mijn afstudeeropdracht. Ik dacht die man kan wel leuk college geven, en hij heeft vast wel zin om met mij aan "iets van exoskeletten en perceptie" te werken; dus dat mailtje was snel gestuurd. Sinds dat eerste gesprek loop ik nog steeds na elke meeting bomvol nieuwe ideeën en motivatie je kantoor uit. Je kan motiveren en kritisch zijn tegelijk, nieuwe inzichten bieden terwijl je het werk wat eerder gedaan is niet onderuit haalt, en bovenal mensen meenemen in je aanstekelijk enthousiasme. Dat maakt je een fantastische begeleider, mentor, en coach. Ik heb enorm veel van je geleerd de afgelopen jaren en ik ben heel blij dat we in de toekomst samen werk gaan maken van de toekomst van werk.

Arkady, as much as I want to help motivate you to learn Dutch, I think switching to English is better for now. You joined our project after my first year to fill the void in the team and take up the role of co-promotor, and I'm very happy you did. Your door was always open, and you brought new perspectives and fresh insights to the table. I've learned a lot from our collaboration in the last three years, and although I'm shifting my research focus away from vehicles a bit, I hope our collaboration will continue. Thanks for always being around when I needed help and for being a great sparring partner!

But just a promotor and a co-promotor don't make the team complete. In the last four years, I've received much support from our colleagues at Nissan: **Yamamura-san, Bouet-san, Takagi-san, Takada-san, Kondoh-san, Sakuma-san, and Daniel-san** and more recently also **Takamatsu-san** and **Sato-san**. Thank you all for the many interesting discussions and conversations, both online and at our workshops in Delft. I really appreciated your support and ideas, as well as your patience, which was sometimes much needed with our scientific approach. Finally, thanks to "team aerospace": **Max, René, and Rowenna** (who were working on a sister project) for the conversations and exchanges of ideas.

Buiten het projectteam waren er veel andere fantastische collega's. **Tom**, bedankt voor de lunches, random ideetjes, sparringsessies, en voor het aanhoren van mijn frustraties af en toe. Ik hoop dat je (als jij straks ook klaar bent) af en toe nog terug komt naar de TU om een pasta te halen bij Pulse. **Nicky**, bedankt voor koffietjes en leuke gesprekken. Het was erg leuk om samen voor de eerste keer naar een congres te gaan, al werden niet alle verwachtingen waargemaakt misschien. **Wilbert**, thanks for the great time we had in NY and for all the nice borrels. **Vishal**, thanks for the chats on all sorts of random topics. **Lorenzo**, bedankt voor je hulp om de demo in de simulator op tijd af te krijgen, zonder jou was dat nooit gelukt! Tenslotte de mensen van het oude Haptics Lab: **Niek, Timo, Sarvesh** en **Joris**. Bij jullie kwam ik vier jaar geleden binnen in een hele leuke groep, en ondanks corona en de lockdowns voelde ik me er heel snel thuis, bedankt!

Bijna 15 jaar geleden begon mijn avontuur in Delft met een klein clubje, en we weten het nog altijd met elkaar uit te houden. **Bart, Hidde, Kristiaan, Roeland, Sjoerd, Sol, en Thijmen**, thanks voor alle weekendjes, wintersporten, zomervakanties, borrels, activiteiten, vergaderingen, en boompjes. We hebben de afgelopen jaren enorm veel samen gedaan en meegemaakt, en ik weet zeker dat er nog een hoop gaat volgen! Wat ze ons voorspelden was een beetje een cliché maar het bleek wel waar; ik heb 15 jaar geleden in dat Rijswijkse parkje, doodop en stinkend in mijn vieze shirt, toch wel echt vrienden voor het leven gemaakt.

En daar kwamen er snel nog een paar bij. **Tjeerd en Jelle**, bedankt dat jullie mijn paranimfen zijn! Toen ik twee mensen moest kiezen die me tijdens mijn verdediging moesten ondersteunen was die keuze snel gemaakt; gewoon de twee die me al jaren door dik en dun steunen. Thanks voor de steun en alle dingen die ik van jullie beiden geleerd heb. Tjeerd, wij zijn zo'n beetje op elk podium in Nederland (en daarbuiten) samen geweest denk ik, van een champignonschuur in Wageningen tot het centraal station van Antwerpen, en van het Paard tot Tivoli. Altijd onderweg met een hoop grappen en goede gesprekken. Thanks voor alles! Jelle, aan de bar of aan de keukentafel, als wij eenmaal zitten te praten zijn we nooit snel klaar; het gaat altijd echt ergens over en het verveelt nooit. Al die gesprekken zijn een extra reden dat je een goede paranimf bent, ik heb je denk ik zoveel verteld dat jij zo'n beetje het meeste weet van iedereen over wat er allemaal in dit boekje staat.

Maar de groep "vrienden van Henk" is natuurlijk veel groter. Ik kan jullie niet allemaal bij naam noemen helaas, want dat kost teveel ruimte¹, maar bedankt voor het zijn van deze fanatische familie! Ik hoop dat er nog vele weekendjes, (wintersport) vakanties, en natuurlijk evenementen mogen volgen. Ik blijf het fantastisch vinden om met jullie bezig te zijn, op het podium of ernaast en ik waardeer de goeie vriendschappen die uit onze gezamenlijke hobby zijn ontstaan!

Dan van mijn surrogaat familie naar mijn echte familie, te beginnen met "de meiden": **Rixt, Hinke, en Marte** (sorry Mart, die volgorde kom je nooit vanaf). Jullie hebben alle drie iets bijgedragen aan dit boekje, misschien zonder het te weten, maar toch erg belangrijk. Rixt, bedankt voor je hulp met de vormgeving! Je had al eerder laten zien dat je een proefschrift echt mooi kan maken, en dat hebt je nu zeker herhaald. Hinke, bedankt dat je me hebt aangespoord om toch voor een PhD te gaan. Jouw goede voorbeeld heeft me erg goed gedaan. En Marte, bedankt dat je me steeds weer laat zien dat lol maken en genieten van alles wat je doet prima samen gaat met presteren. Maar alle drie vooral bedankt voor alle leuke dingen die we samen doen, ik hoop dat er nog vele broer-zus activiteiten mogen volgen! Rikkie Out!

Boris, Atse, en Hester, jullie hebben je de afgelopen jaren je weg moeten vinden bij dit zootje ongeregeld, maar volgens mij is dat fantastisch gelukt. Ik kijk alvast uit naar de volgende feestjes van het jaar, en ben blij dat jullie er bij zijn gekomen! Tenslotte, lieve **pap en mam**, bedankt voor alle hulp en steun die ik altijd van jullie heb gekregen. Jullie hebben me altijd aangespoord om te doen waar ik gelukkig en blij van word en me geholpen om dat mogelijk te maken. Daar ben ik echt

¹ Oké, voetnoten hebben een kleiner lettertype, ik doe het lekker toch: Tjeerd, Renata, Joost, Lucas, Philip, Wouter, Erwin, Ubels, Sander, Liset, Laurent, Wouter, Carli, Nout, Martin, Ivo, Bram, Daniël, Willem, Bart, Quirine, Iza, Bas, Jelle, Nathan, Tijmen, Charmaine, Job, Martin, Rini, Hidde, Mathilde, Jelle, Berry, Eva, Niels, Emilie, Jasper, Maaïke, Rens, Roald, Loïs, Johan, Daan, Robbert, Frank, Tjerk, Julian, Kim, Maartje, Thijn, Nicolette, Marloes, Thijs, Sam, Wikash, Susanne, Loes, Matthijs, Charlotte, Bjarne, Dennis, Arya, Anna, Eline, Max, Fimme, Lieke, Benthe, en Martijn.

heel blij mee en heel dankbaar voor. Het doen van mijn PhD is hier slechts één voorbeeld van. Bedankt voor jullie oneindige steun, energie en tijd om me te motiveren.

Toen ik aan mijn PhD begon dacht ik nog "ik ga nooit zo'n afgezaagd dankwoord schrijven". Maar ik ben bang dat dat het daar toch behoorlijk bij in de buurt is gekomen. Iedereen die dat nu ook denkt, wil ik vooral uitdagen om vier jaar aan een onderzoeksproject te besteden, een boekje van bijna 200 pagina's op te pennen, en het dan te presteren om de laatste pagina's leeg te laten en niet even terug te denken aan alle mensen die belangrijk voor je zijn.

Curriculum Vitæ

Olger Siebinga

22–12–1990 Born in Heerenveen, the Netherlands.

Education

2020–2024 PhD Cognitive Robotics
Delft University of Technology, Delft, the Netherlands
Thesis title:
*Communication-Enabled Interactions in Highway Traffic-
A joint driver model for merging*

2014–2018 M.Sc. Mechanical Engineering
Delft University of Technology, Delft, the Netherlands
Thesis title:
Shared Control in Exoskeletons - a proof of concept

2009–2014 B.Sc. Mechanical Engineering
Delft University of Technology, Delft, the Netherlands

2003–2009 Voorbereidend Wetenschappelijk Onderwijs (VWO)
OSG Sevenwolden, Heerenveen, the Netherlands

Work

2018–2020 Software Engineer
Fleet Cleaner B.V., Delft, The Netherlands

2016 Researcher
Institute for Human and Machine Cognition (ihmc),
Pensacola, Florida, USA

Experiences

2016 Intern
Institute for Human and Machine Cognition (ihmc),
Pensacola, Florida, USA

Awards

2022 Early Career Best Paper Award 2nd Rank
Human Factors and Ergonomics Society Europe Chapter

List of Publications

Journal Papers

6. **Olger Siebinga**, Arkady Zgonnikov, David Abbink, *Human Merging Behavior in a Coupled Driving Simulator: How Do We Resolve Conflicts?* *IEEE Open J. Intell. Transp. Syst.*, **5**, (2024).
5. **Olger Siebinga**, Arkady Zgonnikov, David Abbink, *Modelling communication-enabled traffic interactions*, *Royal Society open science* **10**, 5 (2023).
4. Niek Beckers¹, **Olger Siebinga**¹, Joris Giltay, André van der Kraan, *JOAN: a framework for human-automated vehicle interaction experiments in a virtual reality driving simulator*, *Journal of Open Source Software* **8**, 82 (2023).
3. **Olger Siebinga**, Arkady Zgonnikov, David Abbink, *A human factors approach to validating driver models for interaction-aware automated vehicles*, *ACM Transactions on Human-Robot Interaction* **11**, 4 (2022).
2. **O. Siebinga**, *TraViA: a Traffic data Visualization and Annotation tool in Python*, *Journal of Open Source Software* **6**, 65 (2022).
1. Robert Griffin, Tyson Cobb, Travis Craig, Mark Daniel, Nick van Dijk, Jeremy Gines, Koen Kramer, Shriya Shah, **Olger Siebinga**, Jesper Smith, Peter Neuhaus, *Stepping forward with exoskeletons: Team IHMC's design and approach in the 2016 cybathlon*, *IEEE Robotics & Automation Magazine* **24**, 4 (2017).

Peer-Reviewed Conference Papers

2. **Olger Siebinga**, Arkady Zgonnikov, David Abbink *Uncovering variability in human driving behavior through automatic extraction of similar traffic scenes from large naturalistic datasets*, *IEEE Conference on Systems, Man, and Cybernetics* (2023).
1. **Olger Siebinga**, Arkady Zgonnikov, David Abbink *Interactive Merging Behavior in a Coupled Driving Simulator: Experimental Framework and Case Study*, *13th International Conference on Applied Human Factors and Ergonomics* (2022).

Pre-print Papers

1. **Olger Siebinga**, Arkady Zgonnikov, David Abbink *A merging interaction model explains human drivers' behaviour from input signals to decisions*, *arXiv*, Dec. 15, (2023).

Datasets

1. **Olger Siebinga**, Arkady Zgonnikov, David Abbink, *Data underlying the publication: Interactive merging behavior in a coupled driving simulator: Experimental framework and case study*, *4TU.ResearchData* (2022).

¹Equal contribution

