Modelling Individual Driver Trajectories to Personalise Haptic Shared Steering Control in Curves

Barendswaard, S.

DOI
10.4233/uuid:7292e35d-d45a-4ad1-9663-ae2b5c5a9f16

Publication date
2021

Document Version
Final published version

Citation (APA)
Barendswaard, S. (2021). *Modelling Individual Driver Trajectories to Personalise Haptic Shared Steering Control in Curves*. [Dissertation (TU Delft), Delft University of Technology].
https://doi.org/10.4233/uuid:7292e35d-d45a-4ad1-9663-ae2b5c5a9f16

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

MODELLING INDIVIDUAL DRIVER TRAJECTORIES TO PERSONALISE HAPTIC SHARED STEERING CONTROL IN CURVES

Sarah Barendswaard

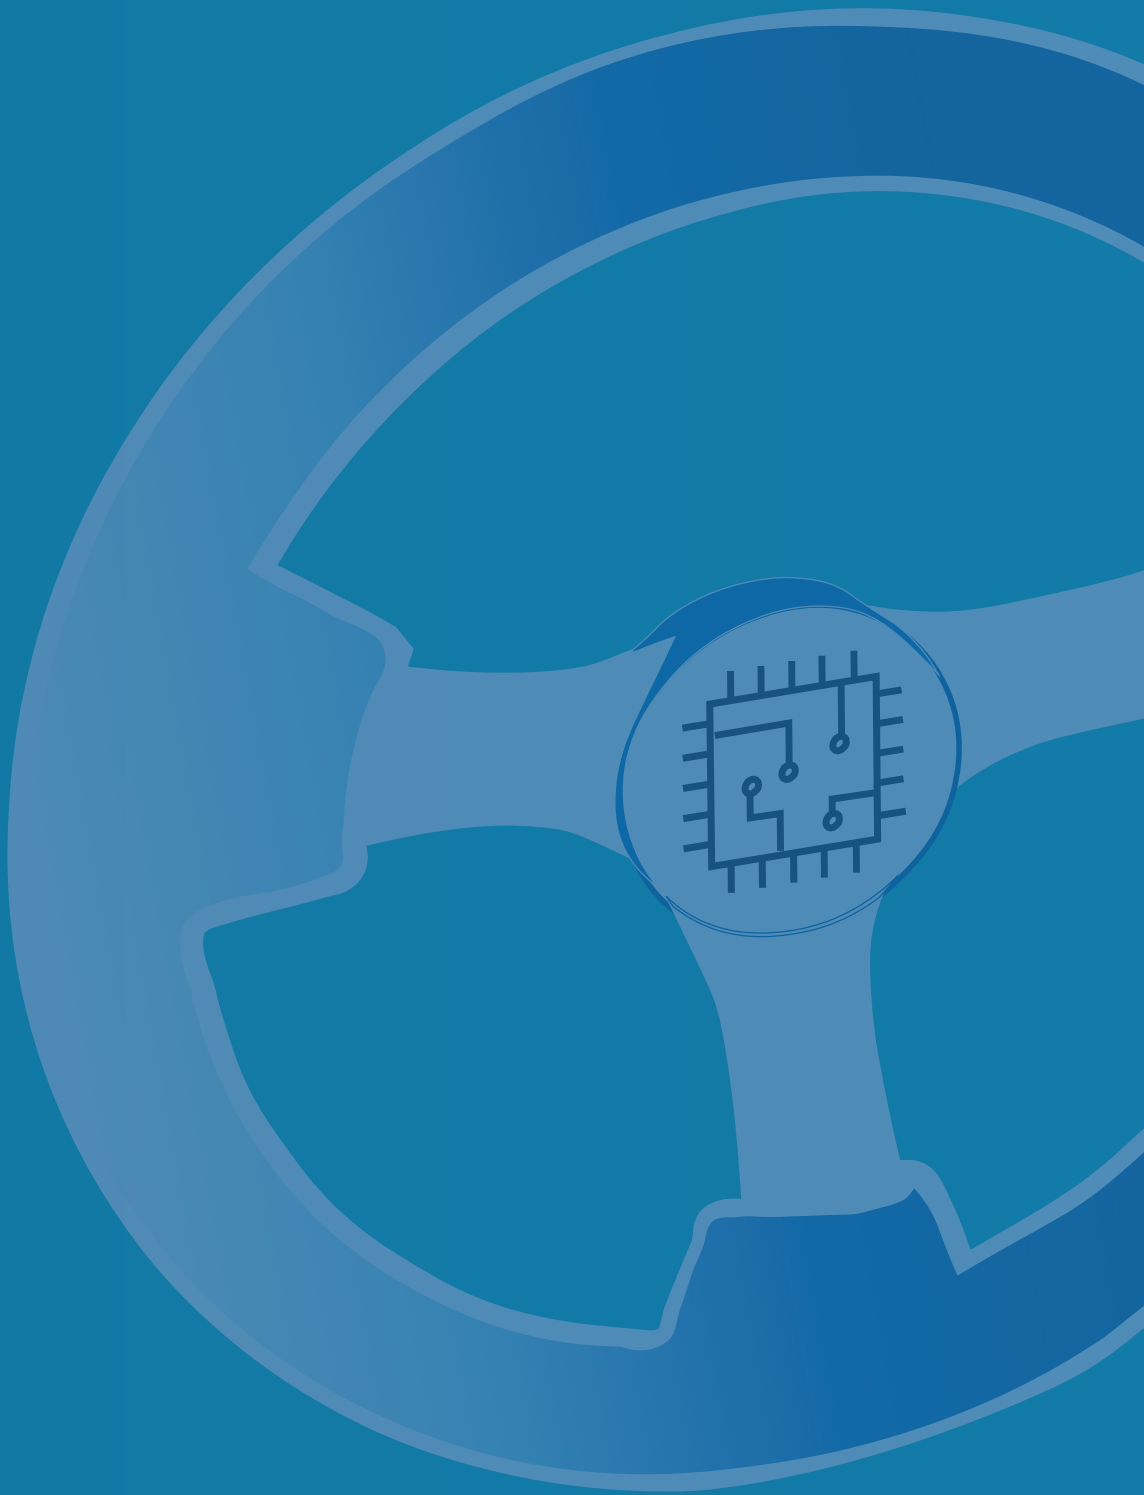

Modelling Individual Driver Trajectories to Personalise Haptic Shared Steering Control in Curves

MODELLING INDIVIDUAL DRIVER TRAJECTORIES TO PERSONALISE HAPTIC SHARED STEERING CONTROL IN CURVES

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 16 Maart 2021 om 15:00 uur

door

Sarah BARENDSWAARD

Ingenieur in de Luchtvaart en Ruimtevaart,
Technische Universiteit Delft, Nederland,
geboren te Sittard, Nederland.

Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie bestaat uit::

Rector Magnificus	voorzitter
Prof. dr. ir. D.A. Abbink	Technische Universiteit Delft
Dr. ir. D.M. Pool	Technische Universiteit Delft
Dr. ir. E.R. Boer	Technische Universiteit Delft
Onafhankelijke leden:	
Prof. dr. ir. B. van Arem	Technische Universiteit Delft
Prof. dr. Ing. F. Flemisch	RWTH Aachen, Duitsland
Prof. dr. D. de Waard	Groningen Universiteit, Groningen
Dr. D. Cole	Cambridge University, Engeland
Prof. dr. ir. F.C.T. van der Helm	Technische Universiteit Delft, reservelid

Dr. ir. M.M. van Paassen heeft in belangrijke mate aan de totstandkoming van het proefschrift bijgedragen.

The work presented in this thesis was made possible by the Dutch Technology Foundation STW (VIDI project 14127), which is part of the Netherlands Organization for Scientific Research (NWO).

Keywords: Driver Modelling, Personalisation, Haptic Shared Control, Advanced Driver Assistance Systems, Driver Identification, Driver Classification

Printed by: Ipskamp

Front & Back: Designed by Eiman Abubakr and Sarah Barendswaard

ISBN 978-94-6421-255-6

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

To my beloved mother, Eiman

CONTENTS

SUMMARY

Road safety is still a challenging issue. In 2020, 1.35 million people have died as a result of traffic accidents, where the number one cause of death for young adults between the age of 5 and 29 is car accidents. In an attempt to improve road safety, the automotive industry has developed numerous types of Advanced Driver Assistance Systems (ADAS). These systems are in general effective in improving safety. However, these systems will only be used if and only if drivers perceive the assistance as intuitive and cooperative. It is recently found that 61% of drivers sometimes switch off the assistance, 23% feel that current assistance are annoying and bothersome, whereas only 21% find them helpful. A safe system that is not used has no safety benefits.

A promising way to improve driver acceptance and to increase safety is to employ haptic shared control (HSC), which is an effective way of keeping drivers in the active control loop. Support in the form of HSC benefits situation awareness and ensures effective monitoring of the environment and automation. However, torque conflict resulting from opposing intentions of driver and automation is reported to be a bottleneck for drivers' acceptance of HSC. Particularly, such conflicts are found to be most debilitating in curves. With each driver having an individual driving style, with different preferences and skill levels, the current standard 'one-size-fits-all' assistance approach to HSC, and driver support in general, is not satisfactory for every individual.

An effective approach to increase acceptance in ADAS, and a reliable way to align the automation to the driver's preferences, is through personalisation. Here, personalisation is generally defined as *'making something suitable for the needs and preferences of a particular person'*. For HSC, personalisation can be effectively realised by adapting the system's adopted trajectory to that of the driver. Therefore, the personalisation of HSC requires a driver modelling approach that predicts an individual driver's behaviour.

Before this thesis, the personalisation of HSC was attempted by adjusting the gains of a corrective feedback HSC, as though it were a driver steering model itself. What was missing was 1) a HSC that allows for personalisation, i.e., a framework where a personalisable reference trajectory is independent of the haptic controller and, 2) a computational driver model or a data-driven driver classification approach that is able to describe individual drivers.

When this thesis was started, a theoretical HSC concept, the 'Four-Design-Choice-Architecture' (FDCA) was introduced within our group. This promising concept was, however, not realised or implemented yet. As for modelling individual drivers, it was not known what type of driver steering and trajectory model(s) are suitable to generate personalised trajectories, if any, due to the lack of a standardised way to compare and evaluate the output performance of driver behaviour models with different structures and complexities. It was not known exactly how to achieve successful personalisation in curves, nor was the needed level of personalisation understood, i.e., adapting to the

intricacies of each individual or adapting to a more general style. Moreover, whether personalisation in itself improves the acceptance of HSC systems, was still to be verified.

These challenges are addressed in the four parts of this thesis:

1. *Driver model assessment*: The development of an assessment method and application on prominent control-theoretic driver models in the literature.

2. *Driver trajectory classification*: Understanding and categorising the types of individual driver trajectories present in the driving population.

3. *Driver prepositioning*: Understanding and modelling driver prepositioning behaviour, a behaviour found to be an essential, yet mostly overlooked aspect of curve-driving behaviour.

4. *Application to Haptic Shared Control*: Apply and evaluate personalised haptic shared control.

With the rich body of driver models available in the literature, it is not only problematic to choose which model to use for personalisation, but it is also unknown whether any can accurately reproduce individual driver steering and trajectory behaviour. A new driver model assessment methodology developed in this thesis quantifies the ability of any control-theoretic driver model to reproduce trajectories of individual driver behaviour (*descriptiveness*), as well as be parsimonious to allow for identification (*realistic identifiability*). This assessment procedure is applied to three prominent control-theoretic driver models for comparison: the Mars model, the Van der El model and the Van Paassen model. According to the developed driver model assessment procedure, the Van Paassen model is most suitable for identifying individual driver steering and trajectory behaviour. Nevertheless, it is classified as 'under-parameterised', i.e., the model is parsimonious but is unable to reproduce individual drivers sufficiently. None of the tested models can capture the drivers' anticipating and preparing for curve entry (prepositioning). This shortcoming severely degrades the descriptiveness of these models, restricting these driver models to being able to describe only 17% of drivers at most. Application of the developed model assessment technique not only highlighted what model is most suitable but also pointed out some serious shortcomings even in the most prominent control-theoretic driver model in the literature.

How differently do drivers take curves? Can individual trajectories be categorised, and used as templates for personalisation? The lack of a mathematical algorithm to categorise such driver trajectories in literature was addressed in this thesis by developing novel rule-based trajectory classifiers for curve driving. Based on knowledge from empirical data of a dedicated driving simulator experiment, rule-based trajectory classifiers were developed. The resulting classes categorise different trajectory types, ranging from *severe curve-cutting* to *severe counter curve-cutting*, and are found to classify all drivers in two separate experiments successfully. From the classification results, performed on right and left turns separately, it is found that most drivers fall into different classes in right and left turns. This means that curve driving in left and right curves should be treated independently. Moreover, it is found that drivers change

their trajectory style when the curve geometry changes, i.e., more drivers curve cut when the curve becomes sharper. The average behaviours of each class were also used for the model assessment methodology to gain insight into what type of behaviour a given model can reproduce. Moreover, this trajectory classifier is also used in this thesis to implement class-average personalisation.

All drivers show a consistent way in which they laterally position themselves before entering a curve, defined in this thesis as prepositioning. However, most driver steering models do not account for the occurrence of prepositioning. This thesis improves the understanding of prepositioning through geometric quantification and investigating the effect of velocity and curvature. In the dedicated prepositioning experiment, 88% of the individual driver trajectories exhibit significant prepositioning behaviour. Moreover, lateral prepositioning displacement increases with increasing driving velocity and increasing curvature. The gained knowledge of prepositioning behaviour is used to develop a sigmoidal prepositioning path model. This path model describes all prepositioning behaviour from a dedicated prepositioning experiment accurately. When combined with the Van Paassen model, which is under-parameterised, the combined model is able to describe all different types of individual trajectory categories accurately. Hereby, a key novelty of this thesis is the first individual control-theoretic driver model. This model is capable of capturing a wide range of different driving styles including the prepositioning that earlier models were not considering.

On the application side, what was missing is a personalisable haptic shared controller. This thesis is the first to realise and implement the theoretical concept of the Four-Design-Choice-Architecture for HSC proposed in our group. This architecture for HSC is unique in that it separates the reference trajectory from the haptic controller, which facilitates independent personalisation of the controller's reference (trajectory). To test the newly implemented Four Design Choice (FDC) controller, an experiment was designed where the FDC HSC was compared with a conventional feedback HSC. Two different realisations of the FDC were tested: one with a one-size-fits-all controller reference and one with a personalised reference. These references were generated using the Mars driver model (i.e., no prepositioning). It is found that both FDC implementations reduce the occurrence of conflicts by a factor 2.3 and reduces driver torque by a factor 3.2 compared to the conventional controller. The significant improvement of the FDC HSC compared to conventional HSC stems from the novel feedforward torque in the FDC HSC. This type of torque only guides rather than corrects, where correction from centreline is what most (conventional) HSC apply. This guidance allows the driver to deviate from the reference trajectory without penalty. Hereby, the feedforward guidance is a new feature in the FDC HSC that changes the way HSC is perceived in general. Subjectively, the personalised FDC implementation is found to improve both satisfaction and usefulness compared to the conventional HSC. This level of increased acceptance was highly encouraging. We expected that by also including prepositioning into the personalised guidance, we would be able to reach even higher levels of acceptance without sacrificing performance.

A final experiment was conducted to evaluate personalisation with the developed FDC HSC, whilst including prepositioning. The experiment was carried out using two different driver groups. One consisting curve-cutters and the other offset drivers. Each

driver group evaluated two different levels of personalisation (i.e., full personalisation and class-average personalisation), against the centreline trajectory (the ADAS industry's standard), and the class-average of the other driver group. The results show that full personalisation is the most capable of aligning torques between automation and driver, and thereby significantly minimises conflicts. However, no evidence is found that this torque alignment is perceived as beneficial by drivers. Subjectively, drivers rate fully personalised, class average personalised and centreline with equal acceptance. This is not to say that trajectory type does not matter, offset drivers particularly give low ratings to a curve-cutting guidance. With this, there is no evidence found that personalisation through mimicking the drivers behaviour or style can leverage extra acceptance, however personalisation through adapting to a driver's preferences is essential. The drivers unexpected high acceptance ratings to centreline is explained by 1) the centreline trajectory generally falling within drivers' preferences and, 2) the novel feedforward torque in the FDC HSC gives the drivers' the allowance to deviate from the controller's reference without penalty. Hereby, further personalising through mimicking the driver with the FDC HSC is not recommended. Therefore, as long as the reference trajectory falls within the preferences of the driver and the shared controller is not (only) correcting, drivers typically will accept the HSC as a useful and satisfactory support system.

This thesis has achieved it's highest level goal, which is to improve the acceptance of the haptic shared control driver support. This thesis provides an improved understanding and new insights into 1) how the novel FDC HSC has solved much of the acceptance issue put forward, and 2) an understanding of *how* to personalise with the FDC HSC. In terms of modelling tools and methods, this thesis has contributed with: 1) a model assessment procedure that can highlight the strengths and weaknesses of any control theoretic model, 2) a trajectory classifier, which can categorise different types of drivers, 3) a prepositioning path model, which, when combined with the Van Paassen control-theoretic driver model results in the first individual control-theoretic driver model, i.e., a model that can capture all main styles of individual driver behaviour and 4) the first personalisable HSC, where the developed modelling methods are applied to evaluate personalised haptic shared control. The findings and insights from this thesis have contributed to design guidelines and, can accelerate future research. Some examples include 1) using the individualised driver steering model, personalisation of ADAS can now be done in real-time, 2) using the developed trajectory classifier, explicit personalisation can be achieved, i.e., the driver can select the type of trajectory guidance he may want, and, 3) the driver trajectory modelling methods developed in this thesis can be used for the personalisation of path-planning in fully autonomous-vehicles.

SAMENVATTING

Het vergroten van de verkeersveiligheid is nog steeds een grote maatschappelijke en technologische uitdaging. Alleen al in 2020 zijn er 1.35 miljoen mensen omgekomen als gevolg van verkeersongevallen. Auto-ongelukken zijn de belangrijkste doodsoorzaak voor kinderen en jongvolwassenen tussen de 5 en 29 jaar. Om de verkeersveiligheid te verbeteren heeft de auto-industrie talloze Advanced Driver Assistance Systems (ADAS) ontwikkeld. Dit soort systemen zijn veelal effectief in het verbeteren van de veiligheid, maar worden alleen door bestuurders gebruikt wanneer zij de assistentie als intuïtief en coöperatief ervaren. Recentelijk is gebleken dat 61% van de automobilisten de assistentie wel eens uitschakelt, 23% de huidig beschikbare assistentie als vervelend en hinderlijk ondervindt, en dat slechts 21% van de bestuurders aangeeft ADAS nuttig te vinden. Een veilig systeem dat niet wordt gebruikt zorgt niet voor meer veiligheid.

Een veelbelovende manier om de acceptatie van ADAS door bestuurders te verbeteren, en daarmee de veiligheid te vergroten, is het gebruik van Haptic Shared Control (HSC). Ondersteuning met HSC is een effectieve manier van assistentie geven die bestuurders actief betrokken houdt bij het besturen van de auto en daarmee het situatiebewustzijn en de alertheid in het letten op de omgeving en de automatisering verbetert. Echter, tegenstellingen in de krachten die door de bestuurder en haptische assistentie op het stuurwiel worden uitgeoefend ("conflicten"), als gevolg van tegengestelde intenties van bestuurder en automatisering, beperken op dit moment de acceptatie van HSC door bestuurders. Dergelijke conflicten zijn in de praktijk vaak het meest storend bij sturen van bochten. Aangezien elke bestuurder een individuele rijstijl heeft, als gevolg van voorkeur en vaardigheidsniveau, is de huidige standaard van "one-size-fits-all"-assistentie in HSC en andere ADAS over het algemeen niet bevredigend voor elke bestuurder. Een effectieve aanpak om de acceptatie in ADAS te vergroten, en een betrouwbare manier om de assistentie af te stemmen op de voorkeuren van de individuele bestuurder, is door middel van *personalisatie*. Personalisatie wordt doorgaans gedefinieerd als het *'aanpassen aan de behoeften en voorkeuren van een specifieke persoon'*. De personalisatie van HSC kan worden gerealiseerd door de baanpositie die het HSC systeem probeert te volgen aan te passen aan hoe de bestuurder zelf rijdt. Dus, voor de personalisatie van HSC zijn technieken nodig die het gedrag van individuele bestuurders kunnen modelleren en voorspellen.

Voor het uitvoeren van dit promotieonderzoek werd doorgaans geprobeerd HSC te personaliseren door de parameters van corrigerende feedback HSC systemen individueel aan te passen. Wat dus ontbrak was 1) een HSC systeem wat in de opzet personalisatie expliciet mogelijk maakt door het gebruik van een personaliseerbare referentiebaanpositie die onafhankelijk is van de haptische controller, en, 2) een betrouwbare aanpak, gebaseerd op wiskundige bestuurdersmodellen of data-gedreven

bestuurdersclassificatie, voor het beschrijven van het stuurgedrag van individuele bestuurders.

Bij aanvang van het promotieonderzoek dat wordt beschreven in dit proefschrift is, binnen onze onderzoeksgroep, een nieuwe en verbeterde opzet voor HSC systemen, de 'Four-Design-Choice-Architecture' (FDCA), op theoretische basis voorgesteld. Dit veelbelovende concept was echter nog niet geïmplementeerd of gebruikt voor HSC. Voor het modelleren van het stuurgedrag van individuele bestuurders was eigenlijk niet bekend wat voor modellen geschikt zouden zijn voor het implementeren van gepersonaliseerde referentiebaanposities binnen HSC, en of die al beschikbaar waren, door het ontbreken van een gestandaardiseerde manier om de beschrijvende capaciteit van verschillende bestuurdersmodellen expliciet te vergelijken en te evalueren. Daarnaast was nog onbekend hoe succesvolle personalisatie voor het assisteren van het sturen door een bocht kon worden bereikt, alsook het daarvoor benodigde niveau van personalisatie: is het nodig om volledig te personaliseren, dus inclusief de specifieke bijzonderheden van elk individu, of is het ook voldoende te variëren tussen meer algemene stijlen van sturen? Op de vraag of personalisatie als techniek überhaupt de acceptatie van HSC-systemen verbetert was ook eigenlijk nog geen duidelijk antwoord.

Deze uitdagingen worden stapsgewijs in de vier delen van dit proefschrift aangepakt:

1. *Driver Model Assessment*: Het ontwikkelen van een nieuwe beoordelingsmethode voor de beschrijvende capaciteit van bestuurdersmodellen en de toepassing daarvan op een aantal welbekende modellen uit de literatuur.

2. *Driver Trajectory Classification*: Het categoriseren en data-gedreven classificeren van de verschillende vormen waarin de baanpositie verandert bij echte autobestuurders tijdens het rijden door bochten.

3. *Driver Prepositioning*: Het beter begrijpen en wiskundig modelleren van voorpositioneringsgedrag ("prepositioning") door bestuurders, een essentieel maar doorgaans vergeten onderdeel van hoe bestuurders een bocht door sturen.

4. *Application to Haptic Shared Control*: Het toepassen en evalueren van deze technieken voor het personaliseren van haptic shared control.

Ondanks het feit dat er een groot aantal verschillende bestuurdersmodellen is voorgesteld in de wetenschappelijke literatuur was er nog geen algemeen geaccepteerde aanpak om het meest geschikte model te vinden voor een bepaalde toepassing, zoals personalisatie. Daarnaast was eigenlijk ook onbekend of er wel een model bestond dat het individuele stuurgedrag van bestuurders nauwkeurig genoeg kon reproduceren voor deze toepassing. Om die reden is in dit proefschrift een nieuwe evaluatiemethode ontwikkeld die het vermogen van bestuurdersmodellen om de baanposities van verschillende individuele bestuurders te reproduceren (*descriptiveness*) kwantificeert, en daarnaast beoordeeld of de modellen voldoende spaarzaam geparameteriseerd zijn om betrouwbare schattingen van hun parameters toe te laten (*realistic identifiability*). Daarnaast is de beoordelingsprocedure toegepast voor het vergelijken van drie bekende bestuurdersmodellen: het Mars-model, het Van der El-model en

het Van Paassen-model. De ontwikkelde beoordelingsprocedure laat duidelijk zien dat van deze drie het Van Paassen-model het meest geschikt is voor het identificeren van individueel stuurgedrag. Toch is ook dit model slechts in staat om maar een klein deel van de rijstijlen die worden gebruikt door individuele bestuurders voldoende te modelleren en wordt het dus geclassificeerd als 'ondergeparametriseerd'. Dit komt omdat geen van de geteste modellen het anticiperende gedrag dat veel bestuurders laten zien ter voorbereiding op het ingaan van een bocht ('prepositionering') voldoende kan beschrijven. Deze tekortkoming beperkt de beschrijvende capaciteit ('descriptiveness') van deze modellen dusdanig dat ze het gedrag van hooguit 17% van de gebruikte database van bestuurders afdoende kunnen modelleren. De ontwikkelde modelbeoordelingstechniek heeft dus niet alleen duidelijk gemaakt welk model het meest geschikt is voor personalisatie van HSC, maar heeft ook aangetoond dat eigenlijk alle bekende regeltechnische bestuurdersmodellen uit de literatuur dezelfde ernstige tekortkoming hebben.

Hoe groot zijn de verschillen tussen bestuurders bij het nemen van een bocht? Kunnen de baanposities zoals gevolgd door individuele bestuurders worden gecategoriseerd en gebruikt als sjablonen voor personalisatie? Het ontbreken van een geautomatiseerd algoritme voor het categoriseren van baanpositie-variaties in de wetenschappelijke literatuur is de reden dat in dit proefschrift een nieuw gestructureerd en data-gedreven classificatie-algoritme voor deze toepassing is ontwikkeld. Dit classificatie-algoritme is ontwikkeld op basis van de empirische meetdata uit twee, in dit proefschrift beschreven, rijsimulator-experimenten en kan op basis van het gereden traject op de weg onderscheiden tussen verschillende rijstijlen, variërend van 'beginnend aan de buitenkant van de baan, eindigend aan de binnenzijde' (*severe curve-cutting*) tot 'beginnen aan de binnenkant van de bocht, eindigend aan de buitenzijde' (*severe counter curve-cutting*). De classificatieresultaten, waarin bochten naar rechts en naar links apart zijn meegenomen, laten zien dat de meeste bestuurders in bochten naar rechts en naar links doorgaans zeer verschillende manieren van baanpositionering toepassen. Dit resultaat toont aan dat het rijden in een bocht naar links en naar rechts altijd gescheiden moet worden behandeld. Bovendien is met deze analyse aangetoond dat bestuurders hun rijstijl aanpassen aan de geometrie van de bocht. Zo gaan bestuurders bijvoorbeeld steeds meer de bocht afsnijden wanneer bochten scherper worden. Op basis van deze classificatie van bestuurdersgedrag is ook een realisatie van het 'gemiddelde' gedrag voor elke klasse geselecteerd om te gebruiken binnen de voorgestelde beoordelingstechniek voor bestuurdersmodellen, om daarmee beter inzicht te krijgen in welke soorten gedrag wel of niet door een bepaald model kunnen worden beschreven. Dit resultaat van deze nieuwe classificatie-aanpak is in dit proefschrift ook verder gebruikt om HSC personalisatie op basis van een klasse-gemiddelde te implementeren.

In alle experimentele meetgegevens die verzameld zijn voor dit proefschrift laten bestuurders een sterke consistentie zien in hoe ze hun auto lateraal op de weg positioneren alvorens een bocht in te gaan. In dit proefschrift wordt dit gedrag 'voorpositionering' ('*prepositioning*') genoemd. In de meeste wiskundige bestuurdersmodellen wordt echter het optreden van dit voorpositioneringsgedrag niet expliciet meegenomen. Dit proefschrift verbetert ons begrip van voorpositionering met een geometri-

sche model en een kwantitatieve analyse van hoe snelheid en kromming van de weg dit gedrag beïnvloeden. In een speciaal voor dit onderzoek uitgevoerd experiment vertoont 88% van de bestuurders duidelijk voorpositioneringsgedrag en neemt de laterale verplaatsing bij het voorpositioneren toe met toenemende rijsnelheid en kromming van de weg. Verder zijn deze gegevens gebruikt om een sigmoïde-vormig geometrisch model van voorpositioneringsgedrag af te leiden. Dit model kan al het voorpositioneringsgedrag zoals dat is geobserveerd in de verzamelde experiment data nauwkeurig beschrijven. Als toevoeging aan het (ondergeparameteriseerde) Van Paassen-model zorgt het meenemen van voorpositioneringsgedrag ervoor dat alle verschillende gemeten rijstijlen nu nauwkeurig kunnen worden beschreven. Dus een essentiële bijdrage van dit proefschrift is het eerste volledig individualiseerbare en complete bestuurdersmodel dat in staat is om een breed scala aan verschillende rijstijlen te modelleren. Voor de praktische toepassing van deze toegevoegde kennis ontbrak het aan een, in de basis, personaliseerbare HSC architectuur. In dit proefschrift is het theoretische concept van de Four-Design-Choice-Architecture (FDCA, 'Vier-Ontwerp-Keuze-Architectuur') voor HSC, zoals die in onze onderzoeksgroep is voorgesteld, voor het eerst gerealiseerd en geïmplementeerd. Een uniek aspect van deze HSC architectuur is dat de referentiebaanpositie expliciet gescheiden wordt van het haptische besturingsalgoritme, wat een volledig onafhankelijke personalisatie van de referentiebaan mogelijk maakt. Om de implementatie van de Four-Design-Choice (FDC) HSC te testen is een experiment uitgevoerd waarbij de FDC HSC is vergeleken met een conventionele HSC architectuur. In dit experiment zijn twee verschillende realisaties van de FDC HSC getest: één met een 'one-size-fits-all' baanreferentie en een tweede met een gepersonaliseerde referentie. Voor dit experiment werden de referentiebaanposities gegenereerd met het Mars-bestuurdersmodel (i.e., zonder voorpositionering). De resultaten laten zien dat beide FDC-implementaties conflicten tussen bestuurder en HSC met een factor 2.3 verminderen, en leiden tot 3.2 keer lagere stuurkrachten door de bestuurder, vergeleken met de conventionele HSC. Deze aanzienlijke verbetering door het gebruik van de FDC HSC komt vooral doordat in de FDC HSC de 'begeleidende' feedforward-krachten volledig gescheiden zijn van de 'corrigerende' feedback-krachten, terwijl conventionele HSC voornamelijk proberen te corrigeren tot het volgen van de middellijn van de weg. De actieve feedforward begeleiding stelt de bestuurder in staat om vrijelijk, zonder afstraffende stuurkrachten, af te wijken van de referentiebaanpositie, wat de gebruikservaring van de nieuwe FDC HSC sterk verandert ten opzichte van conventionele HSC. Subjectieve beoordelingsdata laten zien dat de gepersonaliseerde FDC-implementatie zowel de 'tevredenheid' als de 'nuttigheid' van de assistentie sterk verbeteren. Ookal was deze verbetering van de acceptatie al zeer bemoedigend, de verwachting is dat deze wanneer ook voorpositionering wordt toegevoegd aan de gepersonaliseerde begeleiding nog verder zal verbeteren.

Het laatste experiment in dit proefschrift is uitgevoerd om de effecten van personalisatie bij het gebruik van de FDC HSC te evalueren, waarbij nu ook de voorpositionering door bestuurders is meegenomen. Het experiment bekeek en vergeleek de twee meest voorkomende rijstijlen in twee groepen van proefpersonen: één groep met bestuurders die bochten afsnijden (de 'curve-cutters') en een andere met mensen die op een vaste baanpositie ten opzichte van de middellijn blijven rijden (de 'offset-

drivers'). Elke bestuurdersgroep testte vier verschillende HSC instellingen voor de referentiebaan: twee verschillende personalisatieniveaus (d.w.z., volledige personalisatie en personalisatie op basis van een klasse-gemiddelde), de middellijn van de weg (de huidige ADAS-industriestandaard), en het klasse-gemiddelde van de andere bestuurderssgroep. De resultaten tonen aan dat volledige personalisatie het meest geschikt is om de stuurkrachten tussen de automatisering en de bestuurder op elkaar af te stemmen en zo conflicten tot een minimum te beperken. Deze verbeterde afstemming van de stuurkrachten wordt alleen niet duidelijk door bestuurders als een structurele verbetering ervaren, want subjectief beoordelen bestuurders de volledig gepersonaliseerde, klasse-gemiddeld gepersonaliseerde, en de middenlijn referentiebaanposities op hetzelfde acceptatieniveau. Dit wil echter niet zeggen dat de referentiebaanpositie geen essentiële factor in acceptatie is, want de 'offset-drivers' waarderen een 'curve-cutting'-assistentie duidelijk niet. De onverwacht hoge acceptatiegraad van het gebruik van de middenlijn als baanreferentie wordt verklaard door het feit dat 1) de middenlijn over het algemeen binnen het bereik van de voorkeuren van de bestuurders valt en 2) de ontkoppelde feedforward-kracht in de FDC HSC de bestuurders de mogelijkheid geeft om, zonder dat dat resulteert in sterke corrigerende haptische krachten, af te wijken van de referentie van de HSC. Hierdoor lijkt een sterke personalisatie door het nabootsen van hoe een bestuurder zelf rijdt geen duidelijke extra voordelen te hebben bij het gebruik van de FDC HSC. Zolang de referentiebaanpositie binnen de acceptatiegrenzen van de bestuurder valt en de HSC niet (alleen) corrigeert, zullen bestuurders de haptische assistentie doorgaans accepteren als een nuttig en bevredigend ondersteuningssysteem.

Door direct bij te dragen aan het verbeteren van de acceptatie van de ondersteuning van bestuurders door HSC is het hoofddoel van dit proefschrift bereikt. Dit proefschrift heeft ons begrip van dit probleem vergroot en nieuwe inzichten gegeven in 1) hoe de nieuwe FDC HSC een groot deel van het acceptatieprobleem van HSC systemen kan oplossen, en 2) *hoe* personalisatie met zo effectief mogelijk kan worden bereikt voor de FDC HSC. Daarnaast heeft het onderzoek een aantal nieuwe en breed toepasbare methoden en algoritmes voor het modelleren van het gedrag van bestuurders opgeleverd: 1) een modelbeoordelingsprocedure die de sterke en zwakke punten van regeltechnische bestuurdersmodellen kan kwantificeren, 2) een classificatie-algoritme voor het categoriseren van de verschillende manieren waarop bestuurders een bocht door sturen, 3) een voorpositioneringsmodel dat in combinatie met het Van-Paassen bestuurdersmodel het eerste volledig individualiseerbare bestuurdersmodel oplevert, d.w.z., een model dat alle belangrijke stijlen van individueel rijgedrag kan beschrijven, en 4) de eerste personaliseerbare HSC, waarbij de ontwikkelde modelleringsmethoden worden toegepast om gepersonaliseerde HSC te evalueren. De bevindingen en inzichten uit dit proefschrift hebben verder bijgedragen aan het opstellen van richtlijnen en kunnen toekomstig onderzoek versnellen. Enkele belangrijke voorbeelden zijn: 1) met behulp van geïndividualiseerde stuurmodellen van de bestuurder kan de personalisatie van ADAS nu in real-time worden aangepast, 2) met behulp van het ontwikkelde classificatiealgoritme voor gereden baanposities kan een expliciete vorm van personalisatie worden geïmplementeerd, d.w.z., door bestuurders het type baanreferentie dat hun voorkeur heeft zelf te laten kiezen, en 3) de modelleringsmethoden die in dit

proefschrift zijn ontwikkeld kunnen ook worden gebruikt voor de personalisatie van pad-planning in volledig autonome voertuigen.

1

INTRODUCTION

1.1. DRIVER ACCEPTANCE OF DRIVING AUTOMATION

In the era of the fourth industrial revolution, Autonomous Driving (AD) plays a vital role [113] [153]. The current vision within the automotive industry is that fully autonomous vehicles are the future [42]. This is due to the potential benefits of minimising accidents, enhancing driver comfort, decreasing environmental impact, etc [73]. However, for a successful realisation of fully autonomous driving, driver acceptance is imperative [190] [100].

The biggest roadblocks for the successful realisation and adoption of fully Autonomous Driving is the psychology of the driver, not the needed technology [156]. Important psychological factors that affect driver acceptance of AD include perceived ease of use, perceived usefulness, perceived risk and trust [81]. A recent survey carried out by SAE finds that 73% of respondents preferred to at least share the control with their highly intelligent vehicle, where it is found that *'the biggest barrier to adoption of automated vehicles is ... acceptance.'* [3]. Thus drivers are not comfortable with handing over full control and trusting an autonomous system. The most prominent reasons stem from safety concerns due to potential hacking attacks and system malfunction [88].

Various studies on the adoption of AD predict that in the late 2020s AD will be available for the public [104], whereas widespread adoption and acceptance of AD will only manifest three decades later, around 2050 [106]. In the meanwhile, the underlying AD technologies will be used in Advanced Driver Assistance Systems (ADAS), which contribute to 'paving' the way towards fully autonomous driving.

ADAS are electronic driving automation systems that assist the driver in preventing accidents and increase driver comfort and convenience. Most ADAS execute their assistance through either employing binary alarms, traded control, others use continuous guidance through (haptic) shared control [36]. ADAS systems include Lane Keep Assistance (LKA) [70], Lane Departure Warning (LDW) [83], collision avoidance systems [111] and the autopilot systems [34]. These type of systems have been theo-

retically tested through massive computer and driver simulations, showing objective improvements in safety and driving performance [7] [142].

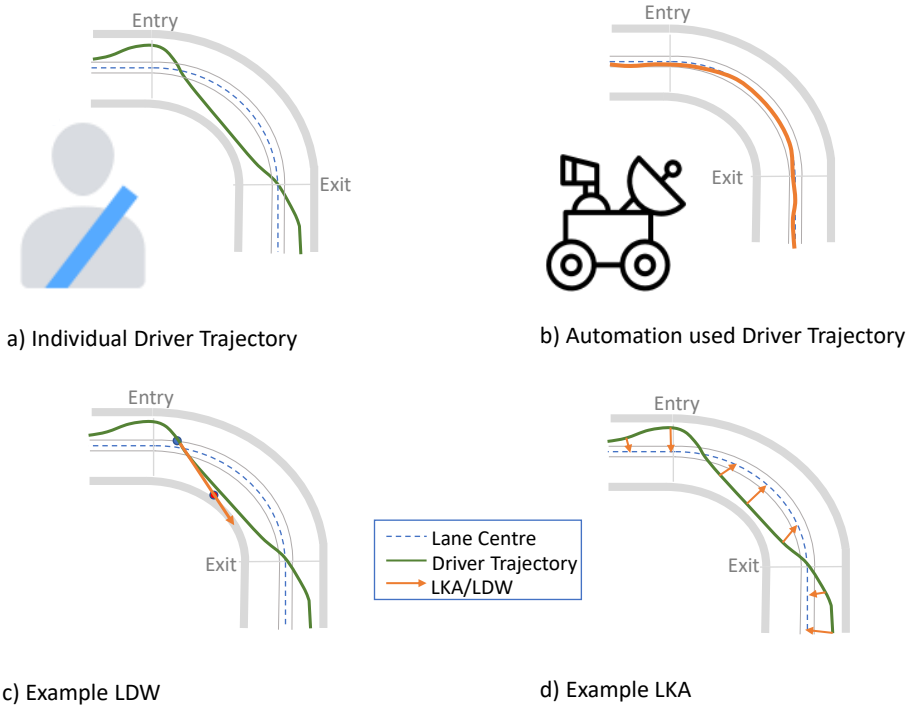

Figure 1.1: The first illustration a) shows the trajectory of an actual driver, b) shows the path of a trajectory-planning-algorithm which many LKA and LDW build upon [55] [82], c) shows how inaccurate prediction of a driver's preferred trajectory can cause false LDW, the arrow represents the prediction algorithm, whereas the red dot represents the false alarm. In d) a continuous LKA system is considered, where the driver from a) will be unnecessarily guided towards the lane center.

Many steps are taken where the industry simply expects drivers to accommodate and accept any added automation. However, drivers only use ADAS when they perceive the automation as intuitive and cooperative [72]. With each driver having an individual driving style, with different preferences and skill levels, the current standard 'one-size-fits-all' assistance approach to driver support may not be satisfactory for every individual. ADAS designs, which genuinely challenge the traditional role of drivers, are currently not explicitly optimised for driver acceptance. In fact, in a survey conducted by the Tech Experience Index study [51], 23% of drivers felt that current lane-keeping and centering systems are annoying and bothersome, whereas 61% sometimes switch off the assistance, compared to only 21% that find them helpful. Hence, the acceptance of ADAS remains a significant issue today [3]. A lack of understanding of an individual's trajectory can cause alarms and support advisories that are unnecessary and untimely. Fig. 1.1 illustrates such a mismatch in trajectory between ADAS and driver, where an example for curve negotiation is given with an individual driver from a) and how a current automatic controller may be designed to drive in b).

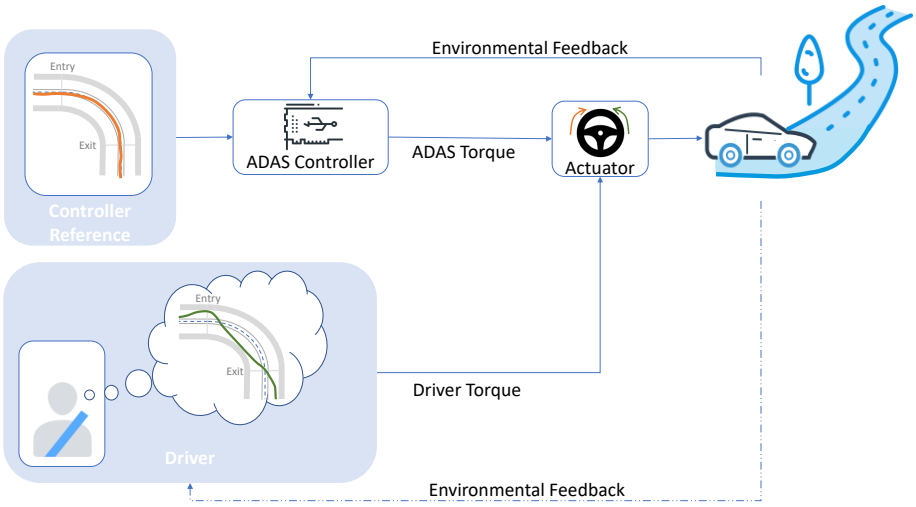

Figure 1.2: This figure illustrates the a general ADAS control diagram for LKA systems. Here, both the ADAS controller and the driver receive feedback from the environment, however the reference of the driver and that of the controller are different, resulting in a mismatch of control approaches

If the trajectory that an ADAS uses as its reference driving path does not match that of the driver, the ADAS will assist when the driver does not feel it is needed. Such mismatches in driver trajectory would cause unnecessary guidance and suboptimal support given by the LDW and LKA systems, as shown in c) and d). This mismatch is can be felt at the control level, as illustrated in the control diagram in Fig. 1.2. Different trajectories can result in different steering strategies. Hence, using an inaccurate algorithm to predict an individual driver's trajectory or action may result in alarms that are perceived as false, either through binary beeps or continuous guidance [137]. This assistance presents itself as a source of annoyance, where the driver may experience the *'cry wolf phenomenon'* [33] (for binary warnings), rejecting the system altogether [135]. Therefore, *'many drivers turn off this nannying technology and don't want it on future vehicles'* [51].

Ignoring human capabilities, limitations, and strengths by *'engineering the automation and expecting the human to accommodate to it can be a recipe for disaster.'* [158]. Until fully autonomous vehicles are the *de-facto* standard, designing assisting driving automation must be done in a human-centred fashion to ensure driver acceptance.

1.2. HUMAN-CENTERED DESIGN

In 1996 the concept of human-centred design was introduced to help mitigate the accidents that occurred in the commercial aviation domain [27]. With the flight management system and the autopilot, the cockpit provided a high level of automation, which resulted in the human acting as a supervisor, and thus a 'back-up' system. This gave way to the well-documented, classic problems of loss of situation awareness, misuse,

disuse, skill degradation etc. These human-automation problems seem to be echoing back into today's advancements in automotive automation. In fact, these problems are now even more serious because the available time the driver has to possibly regain control in a critical situation is much smaller than that of a pilot. Therefore, it is vital to learn from and follow these guidelines.

The human-centred design guidelines give general requirements for *what* needs to be present in automation that is ultimately steered by humans [27]. Yet, Charles Billings, the author, made it clear that specifications on *how to* execute these guidelines is not possible since each system has a particular context. In this section, the human-centred design guidelines are outlined to understand the types of requirements that can help improve acceptance in today's ADAS:

Human Centered Design Guidelines [27]

1) Humans must be in command.
2) Humans must be actively involved.
3) Humans are adequately informed.
4) Humans are able to monitor the automation.
5) The automation must be predictable by the human.
6) The automation must monitor and understand the human.

The first guideline, *the human must be in command*, may be puzzling to some, as some commands given by humans happen to be questionable, resulting in accidents. However, when switching on an assistance system, the ultimate responsibility still lies with the human driver. Even with the highest (SAE) level of driving automation that is currently legally allowable, Level 3 [53], where the human is expected to uphold vigilance whilst monitoring the driving automation, this *'fall-back ready user is expected to take over without undue delay'* [4]. For this reason, being in command is the reality of the driver today.

To effectively be in command, the human needs to adhere to the second guideline: to be *actively involved*. Thus, the concept of the human being a 'back-up' system through unrealistic vigilant monitoring is an approach incompatible with this guideline. Unfortunately, long hours of vigilant monitoring of 'autopilot' driver assistance systems is required by car manufacturers like Tesla, who blame the inattentive driver in case of a fatal accident caused by a system inaccuracy [1]. In fact, the lack of active involvement is known to lead to a loss of situation awareness, which makes it humanly impossible to intervene without undue delay [123]. To effectively stay in the control loop, drivers should engage in tasks that include *'perceptual, cognitive, and psychomotor components so that the drivers must perceive or detect, think about, and respond actively to some stimulus'* [27]. Without such an engaging stimulus, regaining control may put the driver in an even more difficult situation than without any automation [91].

To be actively involved, the driver needs to adhere to the third guideline: to be *adequately informed*. Without information about automated features of the assistance, involvement is delayed, as the actions become unpredictable to the human. This lack of information could easily result in mistrust of the system. *To be adequately informed*,

the driver needs to adhere to the fourth guideline: to be able to *monitor the automation*. This holds not only about the current actions the system undertakes, but also future activities, i.e., to predict and understand the system's actions. Moreover, being actively involved through monitoring is vital for building a correct mental model of the dynamics of the system. This places systems such as Haptic Shared Control, where the system's action can be directly felt [44] at an advantage. Hereby, the first four guidelines can easily be applied when realising the driver assistance concept of *haptic shared control*, which is further elaborated in Section 1.3.

Finally, the last two guidelines dictate that the automation must be predictable by the human and, that the automation must effectively monitor and understand the human. The fifth guideline can partly be satisfied when being successfully and actively involved with the automation and the human building a correct mental model of the system. Still, to be intuitive, the automation must respond in a way that is agreeable to the driver. Hence, the automation must understand the driver as well, which is the sixth guideline. Due to the human not being error-free, the automation needs to how the driver drives normally (whilst being attentive) and correct the driver in times of fluctuating attention or external perturbations. These last two guidelines call for each entity understanding the intent of the other intelligent entity, creating effective *team-player* dynamics between driver and highly automated vehicle. Therefore, to effectively design driving automation, the automation must *understand the driver and act accordingly*. Each driver has their respective unique characteristics, which essentially calls for the need of *personalised* driving automation, as further discussed in Section 1.4.

1.3. HAPTIC SHARED CONTROL

A promising solution to realise guidelines 1-4 of the human-centred design principles, defined in Section 1.2, for the framework of steering ADAS is Haptic Shared Control (HSC). Compared to the conventional *traded control* approaches we find in various ADAS today (i.e., either the driver or the ADAS is in control at a specific moment in time), HSC is found to be more accepted by drivers [131]. It is found that the drivers' trust in the system correlates with their ability to regain control (guidelines 2-4), an advantage that is inherently present with HSC.

Haptic Shared Control consists of two combined elements: *Haptics* and *Shared Control*. It is important first to define these terms to understand what is meant by Haptic Shared Control. The definition of Haptics in the context of shared control applications, as is used in this thesis, is:

Haptics

The use of electronically or mechanically generated movement that a user experiences through the sense of proprioceptive feedback as part of an interface.

Shared Control in the context of human-robot interaction is defined as given in [5].

1

Shared Control

In shared control, human(s) and robot(s) are interacting congruently in a perception-action cycle to perform a dynamic task that either the human or the robot could execute individually under ideal circumstances.

Haptic shared control is where the human and the car automation share the control of the vehicle through torques on the shared control interface. Haptic shared control is generally implemented such that the contribution of both controllers is needed to achieve the highest level of safety and comfort most consistently. This not only brings the driver back into the active control loop, it can also help and guide the driver to drive better and safer, effectively diminishing the chance of accidents [5].

For driving, the control interface is either the steering wheel or gas pedal. When the human can feel the additional torques on the control interface in a *coupled* haptic shared control system, this also provides a platform to interact and communicate intentions with the driving automation [47]. Forcing the driver to stay in the active control loop does not only maintain situation awareness, but may even improve drivers' manual skills through learning from the guidance torques [167].

As a design guideline for HSC systems, the concept of *'the horse metaphor'* is often used [57] [112]. This is where the relationship between driver and haptic shared controller can be seen, metaphorically, to rider and horse. The horse, being an intelligent being that always wants to stay safe, takes commands from the rider. Fortunately, when the rider is slightly distracted when, e.g., reading a map, the horse does not stop and keeps ongoing. Likewise, when the rider gives a command that results in an unsafe situation, e.g., to gallop into a tree, the horse will evade this command, as it can see the tree (and other obstacles on the road) and act upon it (change heading and speed). All these, whilst intentions of both intelligent beings are interactively communicated through push and pull torques on the reins. Although it would be naive to think we can create an automation that is as interactive and intelligent as a horse, the horse-rider phenomenon is a universally understandable metaphor. An ideal to strive towards for HSC [6].

There have already been numerous initiatives to apply haptic shared control. In general, there are two design approaches: 1) using virtual fixtures and 2) optimal trajectory guidance. The first approach essentially sets a grid of go and no-go areas, i.e., the torques are activated when the event of passing a threshold occurs, e.g., placing repellent forces around obstacles or road boundaries, or the flight envelope in aviation [172]. The second approach, which is the area of interest of this thesis, *continuously* guides the driver on a specified trajectory. Continuous haptic shared control has been developed for multiple manoeuvres within steering control tasks, where performance enhancements are found in manoeuvres such as lane-keeping [32], lane changes [169], evasive manoeuvres [136] and, above all, curve negotiation [129] [31].

Although haptic shared control has proven beneficial in terms of improved performance (safer driving) and reduced effort in curves [129], it is also proven that current systems suffer from user acceptance issues [137]. With the existing controllers, user experiences report guidance that is perceived as too strong [45], on trajectories

that feel 'unnatural' [31], especially in curves [128] [129]. These non-human-like, non-personalised controllers result in disagreement between haptic shared controller and driver, evident from opposing torques on the control interface; a phenomenon called a haptic (torque) conflict [6]. This means that the driver fights the shared controller, which causes annoyance, mistrust and rejection of the system altogether. Similar to the mismatch presented in LDW and LKA ADAS, as illustrated in Fig. 1.1, the haptic shared controller must also understand the preferences of the human driver.

To increase the acceptance of haptic shared control, this thesis proposes to add guideline 6 of the human-centred design framework, which is achieved through effective personalisation. As an example of personalisable haptic shared control, Fig. 1.3 gives an illustration of the structure of the recently proposed 'four design choice' HSC architecture [177]. As the name suggests, the control structure comprises four design choices. The haptic shared controller is essentially triggered by the first design choice: the Human Compatible Reference (HCR). Upon changing the driver, the HCR must adapt accordingly. The other design choices, the Strength of Haptic Feedback (SoHF), Level of Haptic Support (LoHS) and Level of Haptic Authority (LoHA), deal with haptic feedforward and feedback strengths, and authority, are elaborated in detail in Chapter 7.

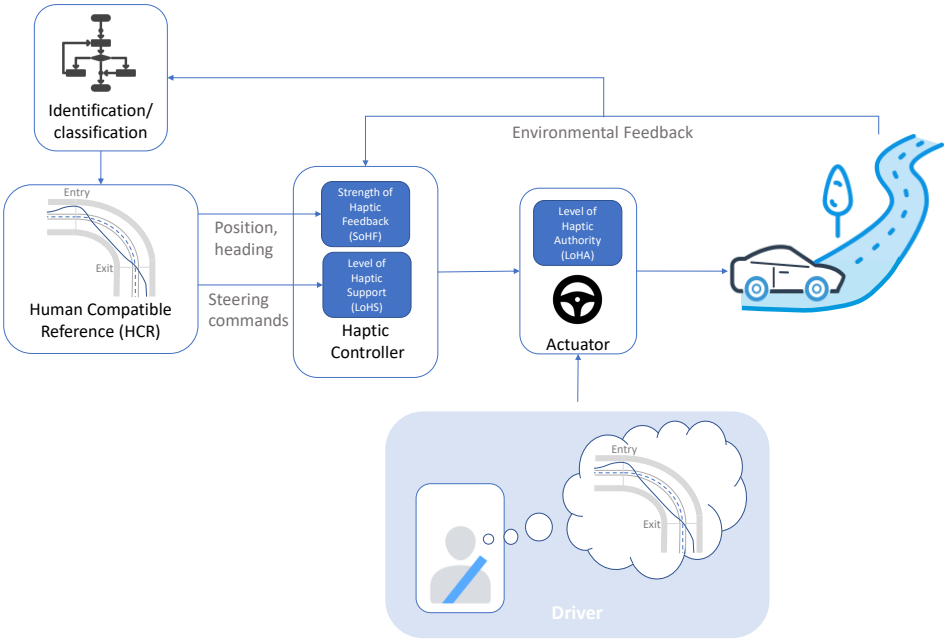

Figure 1.3: The figure illustrates the control structure for the Four-Design-Choice haptic shared controller, derived from the Four-Design-Choice-Architecture [177]. These comprise of a Human Compatible Reference (HCR), Level of Haptic Support (LoHS), Strength of Haptic Feedback (SoHF) and Level of Haptic Authority. More details about how the HCR is generated in the identification/classification block is given in Fig. 1.4.

Attempts to personalise haptic shared control have so far not shown dramatic improvements [31], which can be traced down to the disability of various underlying reference trajectories to re-construct the individual driver's behaviour [155]. A haptic shared controller, synonymous to a horse, can only adapt to the driver when it possesses the intellect to distinguish between different drivers, where this intellect is analogous to a driver model able to perceive, monitor and reproduce individual driver behaviour or the driver's style. With such adaptations, the interaction between driver and shared controller becomes one that increases not only acceptance, but also reliance and trust.

What if the currently inflexible 'horses' would be granted the ability to become aware of their master's needs and preferences across varying situations, and adapt themselves to best serve and satisfy them? What if drivers and cars could become trustworthy companions that can rely on each other, like Alexander the Great and his famous horse Bucephalus.

Prof. David Abbink

1.4. PERSONALISATION

One of the key concepts mentioned in Section 1.2 is that of the automation understanding the human. Therefore, to ensure optimal user experience, *personalisation* of systems is essential. The general definition of personalisation is [72]:

Personalisation

To make something suitable for the needs and preferences of a particular person.

Specifically for ADAS personalisation, the goal is *'to improve the driving experience and the performance of the assisted drivers by adapting the assistance system to their preferences and needs'* [72]. Unfortunately, most technology-centred approaches do not take into account the driver's preference and assume the human always to accept an 'ideal' path, mostly the lane centre [98]. This is evident from trajectory planning algorithms used in various driver assistance systems such as lane-keeping assistance, departure warning, and autopilot [55], as shown in Fig. 1.1. As a step forward, general human-like behaviour is taken into account in several studies, as researchers are starting to believe that it is not only essential to develop intelligent vehicles that are safe but path planning and prediction that is also human-like [69] [96] [86]. Incorporating such human-like models for ADAS is found to improve performance and acceptance of the overall driving. In fact, industrial projects like Nissan's HumanDrive [2] will take such concepts to the commercial market.

As humans drive differently and each has different preferences, merely basing a support system on a one-size-fits-all human-like algorithm seems superficial and insufficient. We see that there is an infinite number of 'human-like' ways to drive, taking an average may support a substantial slice of the driving population; however, not everyone. Driver classification studies find that there are not only numerous ways to

classify different drivers [107] [179], some studies also find over 50 clusters of driving styles [76]. Driver modelling studies also find a continuous range of identified parameters, indicating the presence of a continuum of individual driving behaviour. These findings all support the notion of personal driving style and characteristics.

Adapting ADAS to an individual driver may leverage more significant benefits than only considering a general 'human-like' algorithm. Personalisation improves the driving experience with ADAS in general [101], enhancing driver comfort and safety [97]. The need for personalisation is further highlighted with driver safety being strongly correlated to the individual's style [143], where ADAS adaptations to a driver's style improve performance and safety [94].

There are two (high-level) ways to achieve personalisation, either explicitly or implicitly [54]. Explicit personalisation facilitates the driver to choose a form of guidance to his/her liking, which leaves the driver in control. However, drivers may also need to do extra cognitive work to understand which mode they may prefer. Implicit personalisation is where the ADAS adapts itself to the driver over time (assuming the car is used by a single driver, which can be traced through driver identification or facial recognition etc.). This offers a more fine-tuned and automated form of personalisation. On the other hand, the driver may not always fully understand the adaptations made.

For both implicit and explicit personalisation, driver modelling is cardinal. Driver models are incorporated to predict and quantify the driver's trajectories, intentions, and state–all to adapt or intervene through the support system. Personalisation was pioneered for longitudinal control through adaptive cruise control [72]. More recently, lateral (steering) control ADAS such as lane keep assist and lane change support systems are studied for personalisation [72]. A general control loop for *steering* ADAS is presented in Fig. 1.2, whereas a personalisable HSC is illustrated in Fig. 1.3. Here the generated Human Compatible Reference shown as an individual trajectory on a curved road is generated by the identification/classification block, which employs an algorithm that detects the (individual) characteristics of the driver from the available environmental data and can capture/replicate it. The generated reference is then used as input for the HSC controller, who uses this information to cater to the needs of the individual driver.

Personalisation of ADAS through predicting driver behaviour has shown to improve driving performance and driver acceptance, even when compared to a human-like model with fixed parameters [188]. These efforts include personalisation through identifying driver state through driver facial classification, shown to improve the performance of a collision avoidance ADAS by 50% [64]. Additionally, using a driver model to estimate driver state and predict their future trajectory, the false alarms for collision avoidance and lane departure were reduced by 50% [160]. Including statistics of the driver's previous habits through classification decreases the number of false alarm for lane departure warnings by 66% [195]. Satisfaction increases by 30% when including driver behaviour through modelling into their lane departure warning algorithm [189]. Driver control activity of an LKA haptic shared controller is reduced by up to 40% when personalising the parameters of an underlying driver model through reinforcement learning [184]. Therefore, *driver behaviour modelling* plays a central role in trajectory personalisation, elaborated in the next section.

1

1.5. DRIVER BEHAVIOUR MODELLING

Human driving behaviour contains intricate characteristics of nonlinearity, ambiguity, and randomness, thereby making the modelling of such behaviour not straightforward. Modelling human drivers, in contrast to their time-invariant counterpart, the vehicle, is not an easy task which, until today, is relatively poorly understood. Consequently, the expensive and time-consuming prototype evaluation with human test drivers remains an integral part of developing any ADAS [132]. As a result, understanding human driver behaviour through modelling and identification is a topic under intensive research [180] [103]. Driving a car requires a rich combination of cognitive and motoric processes pertaining to make useful observations about the surrounding environment and of the traffic situation. Accordingly, all-round driver models should include higher-level cognitive processing and lower-level operational control [105]. In fact, there are three categorised levels of driving [125]: *strategic*, where the route is planned, and goals are set, *tactical*, where specific manoeuvres are selected which have short term objectives, and lastly, the *operational* level that realises such manoeuvres. Although all these levels are important, consistent with a bottom-up approach, this thesis targets on modelling the low-level operational behaviour. Nevertheless, individual differences are a result of adopting different strategies (at the strategic level), which impacts how a manoeuvre is planned (tactical level) and are thereby differently manifested at the operational level.

Four main types of driver steering models are described in the literature:

1. *Data-driven modelling* [8], which scrutinizes large amounts of (training) data to make a mathematical algorithm that makes predictions or decisions about driver state [64], driver style [181], driver intent [171] and also driver steering behaviour [11] [24].

2. *Analytic modelling*, where driver algorithms are analytically derived from geometric relationships [25] [157], or cognitive implications [147].

3. *Cost-function-based models* which assume that drivers formulate a function, or criteria and constraints for acceptable behaviour [78], where weights are optimised to attain certain objectives [86] [87], such as minimising the angle to the tangent point [30] and stabilising position around a preview point [108].

4. *Control-theoretic modelling*, where the focus is placed on the human's control dynamics, identification studies have shown that humans use feedforward (preview) control to anticipate the oncoming road and feedback control to stabilise the car [163].

The driver modelling used in the thesis focuses only on both data-driven modelling and control-theoretic modelling. The choice of these specific modelling types is made due to the increased diversity in modelling that these modelling frameworks provide, i.e., individual differences and the diversity and specifics of individual trajectories maybe not be captured when assuming a driver to *optimise* certain costs, or restrict the driver to *analytical* derivations of behaviour. In the human controller, there is a bit of inconsistency that may not be captured with a fixed cost function or with

theoretical derivations of how drivers are assumed to behave. In fact, in this thesis, we show that drivers under normal driving conditions do not appear to consistently follow behaviour that would emerge from either optimal control models.

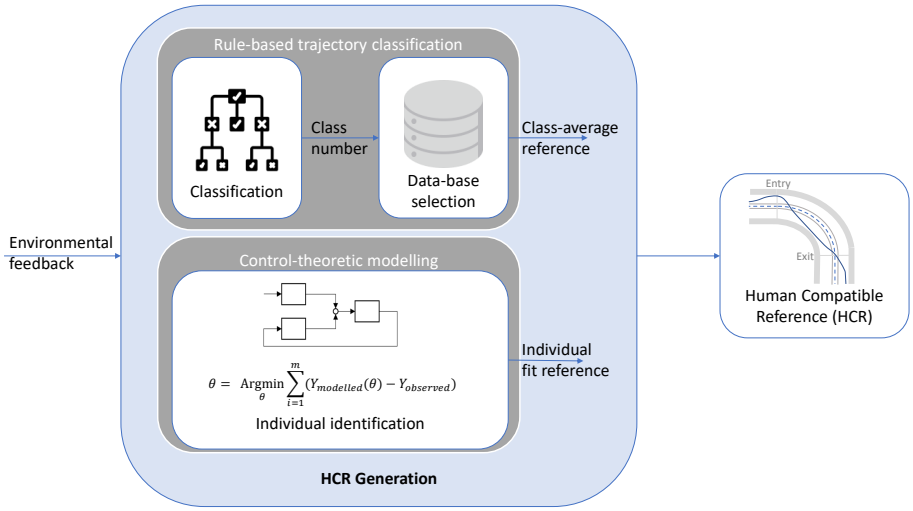

Figure 1.4: This figure illustrates two selected types of driver modelling which can potentially describe personalised trajectories. The modelling methods can be used to generate the Human Compatible Reference for the haptic shared controller outlined in Fig. 1.3. The two modelling approaches are rule-based trajectory classification and control-theoretic modelling. Both these approaches use environmental feedback (possibly CAN signals) to generate the HCR. The classification technique first classifies a driver's trajectory and then uses the resulting class to select the 'average' driver behaviour in that class from a database. Note that the classifier itself does not generate the HCR, for that, a database of (past) driving behaviour is needed. The control-theoretic model identification method uses the environmental feedback data to optimise model parameters which ensure the model's best fit to the individual's data. Hereby, using only environmental feedback and a (considerable) processor, the HCR can be generated.

Data-driven modelling [8], are based upon vast amounts of data, which may come from driving simulator data [11], the vehicle's data output Controller Area Network (CAN) [61], or more recently, smart-phone data [134]. This approach embodies several mathematical techniques, including *basic statistical classification*, where patterns in the data are used to categorise and anticipate driver manoeuvres [97] [13] and individual styles [191]. Another mathematical technique is *data-generative models*, where patterns between input features such as sensed road curvature, and driver output signals such as steering angles are analysed through machine learning methods such as neural networks [11] or Gaussian Mixture Models [182]. From naturalistic driving studies, there is evidence that there are different categories of drivers from trajectory analysis [162]. However, a mathematical algorithm to formalise these categories is missing. Therefore, this thesis develops a type of basic statistical classifier, in particular, to classify different types of driver steering behaviour through a novel *rule-based trajectory classifier*. How this classifier is used to generate a trajectory for the HCR illustrated in the FDC control structure of Fig. 1.3, is illustrated in Fig. 1.4. The classifier itself is not

used to generate a trajectory, but, previously classified driver data (from a database) is used to generate such a reference.

Control-theoretic driver steering models characterise the driver using a combination of transfer functions. Each block represents fundamental characteristics, like an anticipatory gain, time delay, and neuromuscular dynamics. These individual blocks receive anticipatory (road following) feedforward and (various) compensatory (disturbance rejection) feedback signals, transforming them into signals that together make up the drivers steering behaviour [48]. Such control-theoretic models have been developed through both analytic derivations [146] [177] [150], and objective human steering control identification [119] [175]. These models can provide an accurate prediction of the individuals' control output, i.e., steering angles and trajectory. As illustrated in Fig. 1.4 control-theoretic models can be used to generate the HCR. Moreover, when coupled with an (unscented) Kalman filter, real-time updates of meaningful parameters such as time delay and neuromuscular stiffness can be traced [12]. These models, adapted in real-time, can directly be used for personalisation. Albeit, an appropriate selection method to determine which of these models would best allow for personalised guidance is still missing.

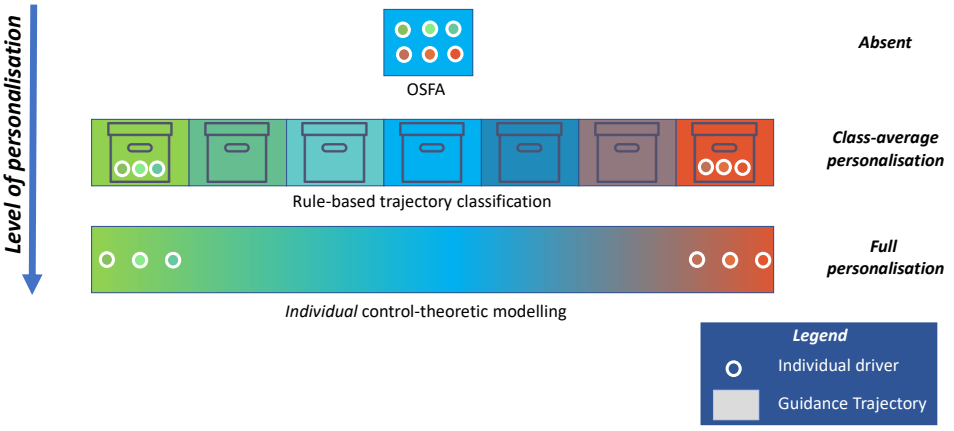

Figure 1.5: This figure illustrates the differences in the level of personalisation with different driver steering and trajectory modelling techniques. Control-theoretic trajectory modelling has the *potential* of providing full personalisation, *if and only if* the model is able to capture these different drivers, through *individual* control-theoretic modelling. Data-driven, rule-based trajectory classification, essentially groups drivers together, providing a class-average type of personalisation. The One Size Fits All (OSFA) condition illustrates the grouping of all drivers into one category.

This thesis focuses on both *control-theoretic driver steering and trajectory modelling* and *rule-based trajectory classification* to faciliate personalisation in HSC. Although both methods output *driver steering and trajectory behaviour*, their differences should be clearly understood. Fig. 1.5 illustrates this difference effectively through metaphorically inferring individual drivers as dots, where their style is a specific colour. The corresponding box that the individual driver falls into is the respective driver (trajectory) behaviour that the specific model outputs, where the colour corresponds to

the modelled trajectory style. For illustration purposes, only three drivers on the opposite ends of the 'style spectrum' are shown.

Descriptive control-theoretic modelling has the *potential*, with a model that can reproduce different individual driver behaviour, of providing full personalisation. Here, the understanding of the human by the automation would be unique to each individual, just as the continuous spectrum of colours are. This type of personalisation can only be implicit. Moreover, compared to explicit personalisation, its is more computationally laborious due to its fine-tuning. It also means that possible adaptations are made without the driver being informed (due to the implicit form of personalisation).

Rule-based trajectory classification can provide personalisation at a 'group' level, or give class-average personalisation. This is done by first applying the classifier to a dataset, where the average in a particular class outcome can be used for the group's guidance. Compared to full personalisation, see Fig. 1.5, three individual drivers on opposite ends of the spectrum are each placed in the same class, where the class-average delivers the average behaviour in that respective class. This type of personalisation can be implemented both implicitly and explicitly. While using the class-average as an explicit form of personalisation may not only be computationally efficient, it gives the driver the freedom to choose the type of trajectory and is a bounded solution, i.e., it will not adapt to potentially unsafe behaviour. Finally, an *average driver* type guidance, provides One-Size-Fits-All support, where all the drivers in the spectrum are placed together into *one* 'box'.

1.6. RESEARCH GAP

To summarise, driver acceptance of Advanced Driver Assistance Systems is a roadblock to enhance driver safety. A promising way to improve driver acceptance and to increase safety is to employ haptic shared control, which is an effective way to keep drivers in the active control loop, maintaining the benefits of situation awareness, and ensuring effective monitoring of the environment and automation. However, torque conflict (opposing intentions between driver and automation manifested in the control level) is reported to be a bottleneck for a driver's acceptance to haptic shared control. Particularly, these conflicts are found to be highest in curves.

A promising approach found to increase acceptance in other ADAS and, a reliable way to align the 'intentions' of driver and automation is through personalisation. This personalisation can be effectively realised through adapting the system's adopted trajectory. Exactly how to achieve successful personalisation in curves is still not effectively understood, nor is the required level of personalisation known. It is as of yet still unknown what type of driver steering and trajectory model(s) are suitable to generate personalised trajectories, if any, due to the lack of a standardised way to compare and evaluate the output performance of driver behaviour models with different structures and complexities. Moreover, whether using individualised driver steering and trajectory models improves acceptance of HSC systems by reducing conflicts between user and shared-controller, is only hypothesized and not verified yet.

1

1.7. RESEARCH GOAL

The challenges of the research gap to increase driver acceptance as mentioned in the previous section forms the backbone of the research goal of this thesis. Therefore the goal statement is:

Thesis goal

Develop model-based approaches to personalise trajectories for steering through curves to enhance acceptance in haptic shared control systems.

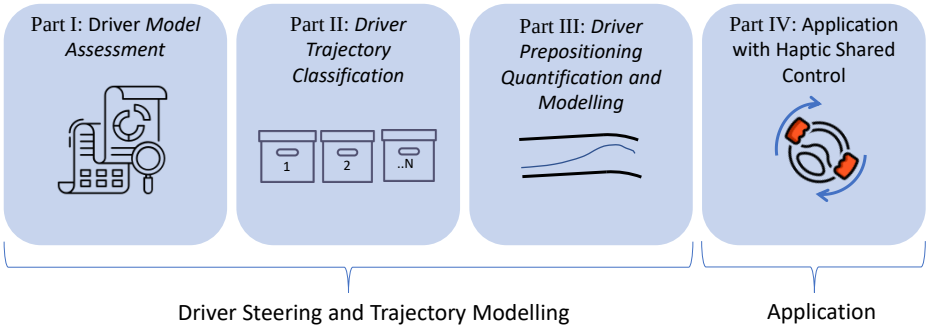

Figure 1.6: This figure illustrates the four parts carried out in this thesis to achieve the goal. These are also indicated in Fig. 1.7.

1.8. APPROACH

The realisation of the goal of this thesis is carried out in four parts. These are illustrated in Fig. 1.6.

I *Driver Model Assessment* This part aims to develop a selection technique to objectively determine what model(s) are most suitable to describe and identify individual drivers.

II *Driver Trajectory Classification* This part aims to develop a classifier that can effectively categorise the different types of observable driver trajectory styles in the driving population.

III *Driver Prepositioning* This part aims to develop a control-theoretic driver steering and trajectory model technique that can describe individual trajectories, which, as shown in this thesis, can only be achieved by incorporating prepositioning (driver behaviour before curve entry). This individualised driver model can facilitate a high level of personalisation.

IV *Haptic Shared Control Application* This part applies and evaluates the developed personalised driver steering and trajectory model(s) in a haptic shared controller, to verify the effectiveness of personalisation.

The approach taken in these four parts are elaborated in the following sections.

1.8.1. PART I: DRIVER MODEL ASSESSMENT

With the rich body of driver models available in literature [25] [132] [114] [180] [103] [95], it is not only problematic to choose what model to use, but it is also unknown whether any can accurately reproduce individual driver steering and trajectory behaviour. In literature, a systematic method to evaluate and compare the performance of control-theoretic driver models is absent. Therefore an effective methodology for rating and comparing the key characteristics of different driver models is essential. In this thesis, the development of the assessment method and assessing different models is carried out over two chapters. First, a baseline methodology is layed out for simple and similar models. Second, this model assessment methodology is updated to allow for comparison of different types of (more advanced) control-theoretic models.

In the first chapter of this part the baseline methodology is formalised. The baseline methodology introduces a set of general criteria comprising of *descriptiveness*, i.e., the ability of a model to replicate individual drivers, and *identifiability*, i.e., the extent to which these models can produce unique and meaningful parameters. To illustrate the proposed baseline assessment method, it is applied for a detailed comparison of two (simple) driver models.

In the second chapter of this part, the baseline assessment procedure is updated in two ways: 1) it is made driver trajectory style dependent, from the driver trajectory classes defined in Part II, and 2) it is standardised, i.e., the obtained assessment results can be used to compare different models that exhibit different types and number of parameters. The improved and updated data-driven model assessment method is applied to three control-theoretic driver steering and trajectory models: the Mars model [155], Van der El Model [175] and Van Paassen Model [177]. The Mars model is chosen as this model was applied to haptic shared control and has also been used for real-time driver disturbance detection [12]. The Van der El model is a recently developed control-theoretic driver steering model, derived using objective system identification of driver steering behaviour. The Van Paassen model is a recent novel driver steering and trajectory model, that is explicitly developed for haptic shared control. The standardised assessment method is applied to these three models, each exhibiting different numbers of parameters and structural elements, to show which model is most suited for personalisation.

1.8.2. PART II: DRIVER TRAJECTORY CLASSIFICATION

The first criterion outlined in Part I is descriptiveness: the ability of a model to capture curve driving steering and trajectory behaviour of individual drivers. This part formalises the specific types of trajectories that a model should be able to capture. To determine these trajectory types, Part II executes a driving simulator experiment with 45 subjects driving on three representative roads, at three different driving velocities. From the collected data, two rule-based trajectory classifier are developed, one to classify empirical data and the other to additionally account for trajectories produced by conventional control-theoretic models. Conventional control-theoretic models assume drivers to start the curve at centerline, however, there is no empirical evidence

for this. Therefore, an additional classifier is formulated to explicitly account for trajectories that are generated by these models. The classifiers are found to not only provide a mathematical algorithm that can distinguish between the different types of trajectory patterns in the driving population, but also formalise the different trajectory types that driver models should be able to replicate, i.e., descriptiveness. Additionally, it can also be used by itself as a 'driver trajectory model' for personalisation. The empirical classifier is used as a means to achieve 'class-average' personalisation (as outlined in Fig. 1.5) in Part IV.

1.8.3. PART III: DRIVER PREPOSITIONING

All drivers show a consistent way in which they laterally position themselves before entering a curve, defined in this thesis as prepositioning. Unfortunately, the nature of control-theoretic driver steering models does not allow for any anticipation or lateral positioning for the curve before curve onset. This is because most models are 'activated' by road curvature, or a change in heading, with at most 1 s of preview. To improve the current understanding of this behaviour, this part is carried out in two chapters. First, the composition of prepositioning is geometrically quantified and how it is affected by changing road radius and velocity is investigated. Second, a mathematical model is proposed to capture drivers' prepositioning behaviour. Fitting metrics are used to verify the developed model for all types of drivers. Moreover, the prepositioning model is merged with the best performing control-theoretic model from Part I. Finally, the model assessment criterion: descriptiveness, is applied to show the added value of including prepositioning behaviour to a driver model.

1.8.4. PART IV: APPLICATION WITH HAPTIC SHARED CONTROL

Application with haptic shared control is carried out over two chapters. First, a personalisable haptic shared controller is realised and implemented. Second, personalised haptic shared control, and the levels of personalisation are evaluated.

In the first application chapter, a realisation of the novel 'Four-Design-Choice-Architecture' [177] is newly implemented as the 'Four-Design-Choice' (FDC) haptic shared controller. This haptic shared controller was chosen because it is the first to separate controller reference from haptic controller part, which facilitates independent personalisation of the controller's reference (trajectory). To test out the efficacy of the control structure and the controller's potential to personalise, a human-in-the-loop experiment is carried out. Three haptic shared control implementations are tested 1) the conventional HSC [31] 2) the FDC HSC with a driver modelled reference for the average driver, a One-Size-Fits-All implementation and, 3) the FDC HSC with a driver modelled reference for individual behaviour, a (limited) personalised implementation where prepositioning is not considered. The driver model used to identify driver behaviour and thereby generate the controller reference is the Mars model, which is a prominent control-theoretic driver model used for haptic guidance [155], and assessed in Part I.

In the second chapter of this part, personalisation is evaluated. The experiment carried out in this chapter tests two different groups of drivers (classified using the trajectory classifier of Part II). These are the most occurring driver class *offset drivers*,

and drivers that try to optimise their steering demand and lateral acceleration, *curve-cutters*. Four different types of independent FDC controller references are examined: 1) a fully personalised reference, 2) a class-average personalised reference, 3) the centreline (considered as a baseline because it is the industries standard) and, 4) the other groups class-average reference. Additionally, these controller references can be served to the driver in different ways, determined by the HSC feedforward gain. Two different feedforward gains are tested because the first chapter of this part finds that this gain has an impact on driver acceptance. Note that for the fully personalised reference, recorded data were used; however, the individualised control-theoretic model from Part III could also be used. The results of this part verify not only whether personalisation increases acceptance, but also what level of personalisation may be required.

1.8.5. SCOPE

All the driver steering and trajectory models in this thesis assume fixed speed, thereby focusing on steering applications such as curve-negotiation. This assumption is mainly made for fundamental understanding. When adding another axis of control, motoric and cognitive interferences may take place [20], cluttering the analysis of steering behaviour. With the focus on curve driving, this thesis takes an in-depth look into driver *trajectory* behaviour. Moreover, the driver models tested in this thesis that generate these trajectories, are control-theoretic, linear time-invariant models, although real drivers may show time-variant artefacts such as fatigue. To focus on a fundamental understanding of how drivers use visual information to negotiate a curve, in this thesis, all experiments are performed in a fixed-based driving simulator. Even though, in real-life, drivers perceive and respond to vestibular cues when negotiating a (sharp) curve. Nevertheless, this vestibular feedback is only used in the curve, prepositioning is triggered purely based on visual information.

This thesis thus aims to improve driver acceptance by experimentally demonstrating the benefits personalisation of haptic shared control through both driver trajectory behaviour classification and state-of-the-art control-theoretic driver modelling, determined offline. This thesis gains a deep understanding of the fundamentals of individual driver modelling and the carried out experiments are used to develop new design guidelines which can be used for real-time implementations of personalisation.

1.9. THESIS OUTLINE

The outline of this thesis, sectioned into four parts, is presented in Fig. 1.7.

Part I consists of Chapters 2 and 3, where the methodology for evaluating and comparing control-theoretic driver models is developed and applied to five different models in total. In Chapter 2, the assessment method is developed, allowing there to be a formalised rating for driver steering models. In Chapter 3, the assessment method undergoes two further developments. First, it becomes mathematically standardised, to allow for comparisons between models of a different number of parameters. Secondly, the method becomes data-driven, i.e., dependent on the classes defined in Chapter 4 as representatives for the types of drivers found in the driving population. In practice, Chapter 4 is carried out before Chapter 3.

Part II consists of Chapter 4, which focuses on driver trajectory classification. Two rule-based classification methods are introduced, and the effects of changing driver speed and road curvature on trajectory class outcomes is determined.

Part III consists of Chapter 5 & 6. These address driver steering behaviour before curve entry, defined as prepositioning. Geometric quantifications of this behaviour in terms of trajectory are made in Chapter 5, based on collected experimental data. At the same time, these geometric quantifications are used in Chapter 6 to develop a mathematical model that can capture the different prepositioning behaviours present in the population. In Chapter 6 the developed prepositioning model is combined with the best performing control-theoretic model from Chapter 3, to show its benefits for different classes of curve driving steering and trajectory behaviour.

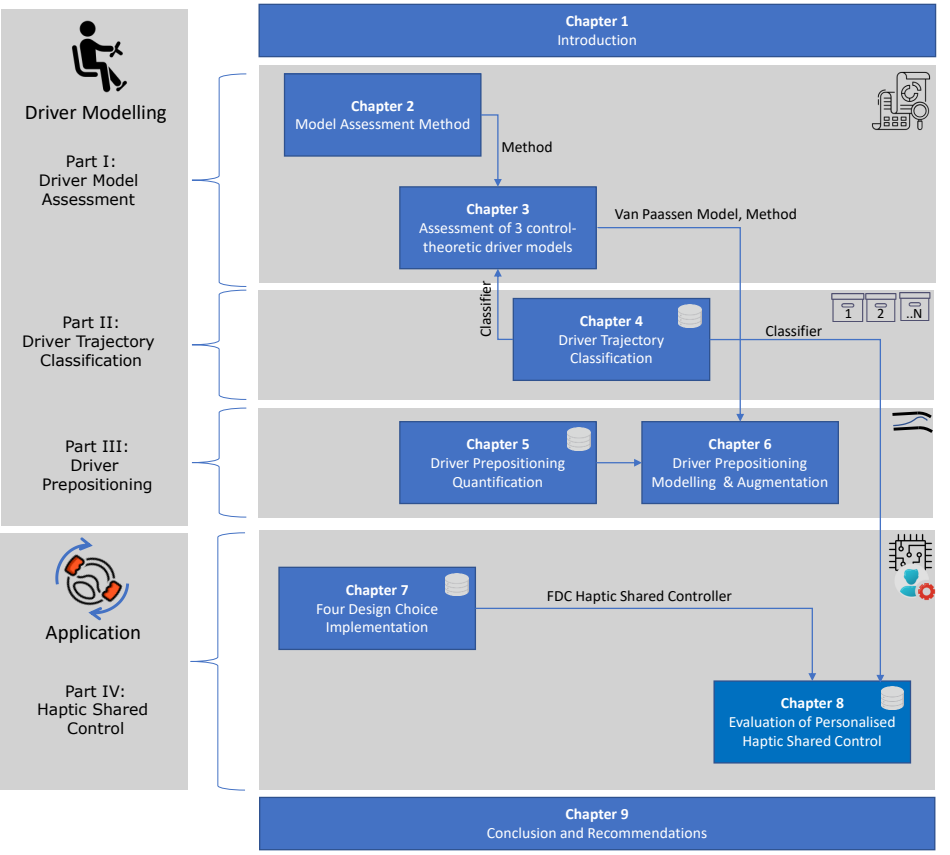

Figure 1.7: This figure illustrates the outline of the thesis in relation to the research parts, see Fig. 1.6. The cylinder shape indicates that a dataset has been collected in a particular Chapter through a human-in-the-loop experiment.

Part IV consists of Chapter 7 & 8. Here the effects of personalisation are tested in an application with haptic shared control. In Chapter 7, limited personalisation is implemented using a state-of-the-art control-theoretic driver model. Note that in practice,

Chapter 7 is carried out before Chapter 3. In Chapter 8, class-average and full personalisation are directly compared, facilitated by the developed rule-based trajectory classifier from Chapter 4 and the developed FDC HSC from Chapter 7.

Finally, Chapter 9 wraps-up this thesis where conclusions, design guidelines and recommendation are given.

I

Driver Model Assessment

2

Model Assessment Method: Simple Driver Models

Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful.

George E.P. Box

Part I of this thesis, which comprises two chapters, deals with driver model assessment. This first chapter sets the foundations by proposing a model assessment method, whereas the second chapter compares and assesses three different and prominent control-theoretic driver steering models. This chapter introduces the criteria of the basic assessment method: 1) descriptiveness, the model's ability to capture different types of steering behaviour, 2) identifiability, the ability of the model for unique mapping between a steering behaviour and a parameter combination, and 3) realism, the parameter span resulting in realistic steering behaviour. The utility of the introduced assessment method is shown by analysing and comparing two driver models: both assuming proportional control on a predicted lateral position, comprising a gain and lookahead time.

The contents of this chapter have been published as:

Barendswaard, S., Pool, D.M., Boer, E.R., Abbink, D.A., 'A Method to Assess Individualized Driver Models: Descriptiveness, Identifiability and Realism', *Transportation Research Part F: Psychology and Behaviour*,Volume 61, Pages 16-29.

2.1. INTRODUCTION

For some driver assistance systems (DAS) mathematical models of human driver control behaviour are indispensable [6, 124]. Specifically, DAS that collaborate with the human on an operational level also needs to be tuned to the human's timely skill-based behaviour [79]. For a good tuning, we need a proper grip on human behaviour in real-time, which can be achieved through identifying computational driver models (that consider environmental and human changes) online [12].

In literature, the focus of driver modelling is directed towards a fundamental understanding of how humans drive. Research has been done on how human sensors pick-up information from the environment through; visual, vestibular and somatosensory receptors to form a deeper understanding by developing relevant models [132]. Many models focus on the visual receptors of the human, selecting environmental cues from the complex three-dimensional visual scene, with both a perception of road path geometry and optic-flow [62]. For example, most driver steering models are based on the hypothesis of parallel high- and low-frequency compensation [48, 75, 174], often coupled to dedicated "far" and "near" preview/tangent points [116, 145, 155], respectively. Driver steering models currently implemented in DAS are often, for practical reasons, simple – e.g., two-parameter (single preview point) – driver models [129, 144].

What is still missing in the framework of driver (steering) modelling, is a structured approach to assess driver models' practical capabilities. Specifically, their abilities in capturing different realistic driver steering behaviour and the corresponding ability of unique mapping between that behaviour and an identified parameter set are still missing. The lack of a systematic assessment method is a bottleneck to determining which elements of a driver model, be it structure (i.e., inputs and outputs, model equations) or the number of parameters, is beneficial and necessary for a model. Driver models' practical capabilities in capturing different types of drivers can be enhanced by using more parameters. However, then the mapping between a modelled driver output and parameter set becomes less unique, resulting in a tradeoff [10, 49]. A lack of a formal quantification of model characteristics has prohibited the fair comparison of existing driver models and thereby, the determination of an optimal driver model, for a given application. More specifically, we lack a comprehensive analysis of model performance in terms of 1) *descriptiveness*, the capacity for capturing realistic human behaviour and style variations, 2) *identifiability*, the success of unique parameter retrieval and 3) *realism*, avoiding unrealistic driver model steering behaviour.

As an attempt to bridge the missing link in the driver (steering) modelling literature, this chapter introduces the needed criteria for a systematic quantitative assessment method for the sake of model performance analysis. The introduced assessment method is demonstrated via comparing a *linear prediction driver model*, used for individualisation of haptic shared control by means of driver identification [31], with a *nonlinear prediction driver model* used as part of an underlying controller in [128, 129]. The models' *descriptiveness* is determined to evaluate whether the models are able to capture different trajectories in the lateral position domain. Real human driver data for comparison is also illustrated, taken from recent simulator experiments [31].

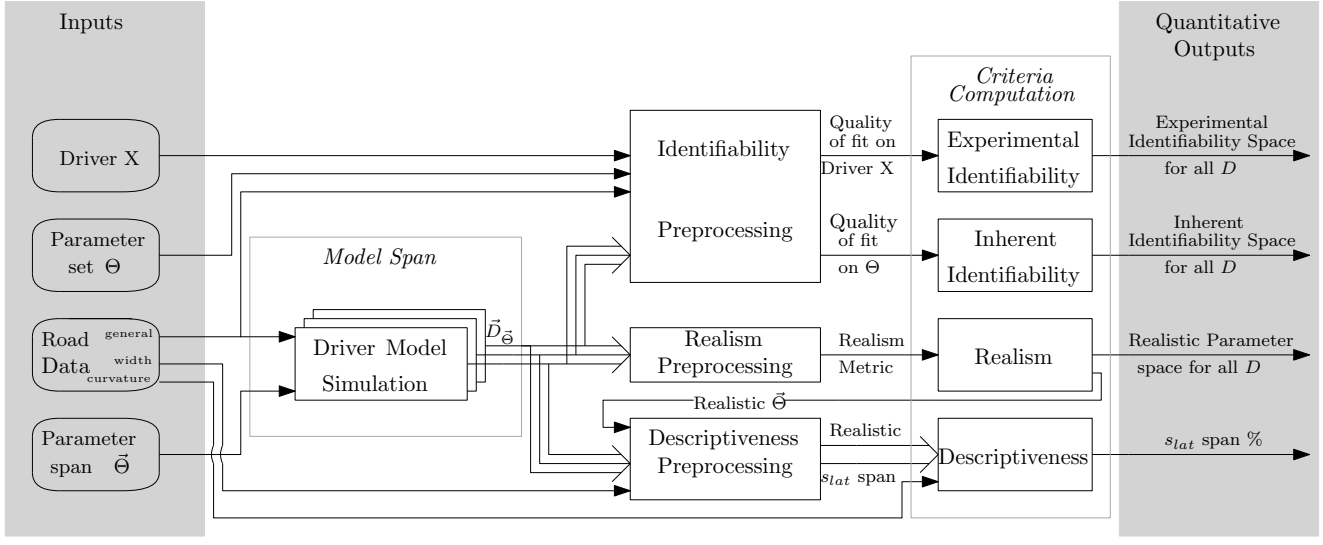

Figure 2.1: This figure illustrates a flow-diagram that presents how the three criteria (Descriptiveness, Identifiability, Realism) are *generally* computed and evaluated. The inputs comprise a Driver X, a reference parameter set Θ, road data and the considered parameter span $\vec{\Theta}$. The Driver X, i.e., empirical driver behaviour of an human driver, is used for experimental identifiability. The parameter set Θ is used for inherent identifiability. The road data and considered parameter span are used to compute the driver model response at every entry of the parameter span $\vec{\Theta}$ which results in the model span $\vec{D}_{\vec{\Theta}}$. The model span is used in the identifiability preprocessing block to evaluate how close Driver X or the model response to Θ is to each of the entries in the model span. For the realism criteria, each entry in the model span is evaluated in terms of whether it is realistic (where being realistic is defined in Section 2.2.1) and descriptiveness determines the trajectory area that the model span covers in the lateral position domain.

The *identifiability* will be investigated in twofold, through an inherent model- sensitivity type analysis and by examining the sensitivity of parameter retrieval based on experimental data. These analyses are done in both the lateral position and steering angle domain. The models' capacity for capturing realistic human steering deflections is verified by quantifying human-like behaviour (*realism*) in terms of avoiding many steering wheel oscillations.

This chapter is structured as follows: Section 2.2 gives a description of the three assessment criteria and elaborates on how they are quantified. Section 2.3 describes the driver models from [31] and [128, 129] that are assessed in this chapter. Section 2.4 presents the results of the assessment, along with a summarised objective comparison of the models. Finally, Sections 2.5 and 2.6 are the discussion and conclusion.

2.2. DESCRIPTION OF THE ASSESSMENT METHOD

This model assessment method provides a complete model performance analysis and provides a platform for a fair comparison between multiple driver models. The method is based on three criteria: Realism, Identifiability and Descriptiveness. These criteria are computed based on four inputs and a driver model, as illustrated in Fig. 2.1. The inputs are the road data (consisting of road geometry, road curvature and road width), the parameter span $\vec{\Theta}$ and the identifiability inputs (parameter set Θ and Driver X). The road geometry provides the environmental cues for the driver model simulation, where the exact signals extracted from the road geometry depend on the driver model. The parameter span $\vec{\Theta}$, for all driver model parameters within Θ, is a subspace of $(-\infty,\infty)$. It is essential to choose this span wisely to include for extremes of the driver steering model. The parameter set Θ is input for the inherent identifiability, whereas Driver X is input for the experimental identifiability.

For all three criteria, generating the models *Model Span* is necessary. Given the road geometry, the model span simulates the driver outputs of the driver model for each parameter combination defined in the parameter span $\vec{\Theta}$. This vector of driver model outputs $\vec{D}_{\vec{\Theta}}$ (where steering angle and lateral position, amongst others, are driver model outputs $\{\delta_s, s_{lat}\} \in D$) is then used as input for all three criteria preprocessing blocks; Identifiability preprocessing, Realism preprocessing and Descriptiveness preprocessing. Each of which is elaborated in the criteria sections. The output of the criteria computation gives a quantitative output which can be used to classify a model and also a means to compare models with each other.

These three criteria are elaborated, giving both their general definitions' and specific implementations' in the following.

2.2.1. REALISM

Realism is generally defined as: *the parameter span of the driver model generating realistic driver (steering) behaviour.*

The realistic parameter span results from examining the *Realism Metric* as given in Fig.2.1, specified in this chapter by analysing the modelled steering wheel deflections δ_s, through the steering reversal rate SRR. It can be seen in Fig. 2.2, that the signal of interest extracted from $\vec{D}_{\vec{\Theta}}$ is the steering angle $\vec{\delta}_{s_{\vec{\Theta}}}$. The matrix signal $\vec{\delta}_{s_{\vec{\Theta}}}$ is examined

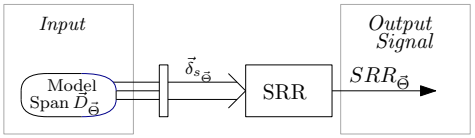

Figure 2.2: This figure illustrates a flow-diagram that presents how the **Realism Preprocessing Block** is *specifically* implemented. The input is the model span, from this the steering angles (δ_s) are selected

for its steering reversal rate SRR, which tests for its number of oscillations. The steering reversals are defined as the number of times that the steering wheel was reversed by a magnitude greater than 2 deg [118] around a local minima and maxima of steering wheel angle δ_s. The SRR block produces another matrix signal $\vec{SRR}_{\vec{\Theta}}$ which has the same dimension as $\vec{\delta}_{s_{\vec{\Theta}}}$. The signal $\vec{SRR}_{\vec{\Theta}}$ becomes an input to the Realism criteria which decides upon a realistic $\vec{\Theta}$ (less than 5 oscillations per curve) for which the steering oscillations are acceptable, providing the parameter area of $\vec{\Theta}$ as quantitative output. A large realistic parameter area means more flexibility for choosing a parameter set, which is generally good. Another output from this criteria is the parameter span of realistic steering reversals *realistic* $\vec{\Theta}$, input to the descriptiveness block.

2.2.2. DESCRIPTIVENESS

Descriptiveness is generally defined as: *the extent to which a driver (steering) model is capable of capturing different driver (steering) behaviour.*

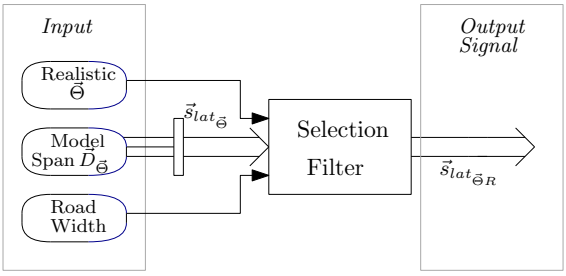

Figure 2.3: This figure illustrates a flow-diagram presenting how the **Descriptiveness Preprocessing Block** is *specifically* implemented. The model span, realistic parameter span and road width are taken as inputs. The realistic parameter span and road width are used in the selection process to select trajectories (s_{lat}) that are produced by the realistic parameter span and fall within the road boundaries.

In this chapter, we define the extent of capturing different driver steering behaviour in the lateral position domain, as a distinction between different types of driving steering behaviour for constant velocity simulations is predominantly found in *driver trajectories* (the s_{lat} response). In Fig. 2.3, the signal of interest is the lateral deviation $\vec{s}_{lat_{\vec{\Theta}}}$, which is extracted from the driver model span $\vec{D}_{\vec{\Theta}}$. The complete modelled lateral deviation span is filtered by only taking through the signals which are contained within the road width, and that is not resulting from a *unrealistic* parameter set. The filtered matrix signal $\vec{s}_{lat_{\vec{\Theta}R}}$ has a smaller dimension than $\vec{s}_{lat_{\vec{\Theta}}}$ due to leaving out un-

realistic signals. The signal $\vec{s}_{lat_{\tilde{\Theta}R}}$ (the *realistic* s_{lat} *span* in Fig.2.1), is used as input to the descriptiveness criteria where the realistic lateral position span is quantified as a percentage area of the total *descriptiveness area*. The larger this percentage, the more descriptive the model becomes.

We define a *descriptiveness area* of s_{lat} (lateral position) in which all the extremes of *curve cutting* and *over turning* that lie within the limits of the road-width are contained. All types of driver behaviour, labelled as swinging, correcting, cutting, and drifting [162], should lie within the bounded descriptiveness area. The descriptiveness area is given as a grey block illustrated in Fig. 2.4. The descriptiveness area starts 1 second before the curvature begins and ends 1 second after the curvature ends because, for curve driving, preview/prediction time is usually around 1s [174] [31].

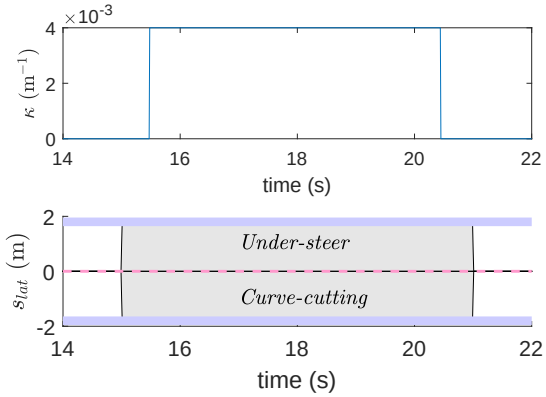

Figure 2.4: Top Figure: curvature profile with a curve radius of 250m, for a velocity of 25m/s; a curve length of 150m. Bottom Figure: representation of ideal model descriptiveness area given in grey for the corresponding curvature profile and a 3.6m road width, the dashed line at s_{lat}=0 represents the lane center.

2.2.3. IDENTIFIABILITY

Identifiability is generally defined as: *the ability for unique mapping between a given driver (steering) behaviour and a parameter combination.*

Identifiability looks into the sensitivity of a given driver steering behaviour to variations in model parameters (that would produce almost the same behaviour). This chapter realises this by defining two different categories of driver steering behaviour, that which comes from the driver model itself, and that which comes from experimental data of a human driver. Resulting in two different types of identifiability; inherent and experimental driver identifiability. Inherent identifiability evaluates *the extent to which a simulated model time series can be uniquely and accurately identified using the same model.* In contrast, experimental driver identifiability evaluates *the extent to which an experimental driver time series can be uniquely and accurately identified.* Mathematically, given a driver model where the parameter space Θ is finite, we declare the driver model output D to be identifiable if the mapping $\Theta \rightarrow D_{\Theta}$ is unique. This implies that:

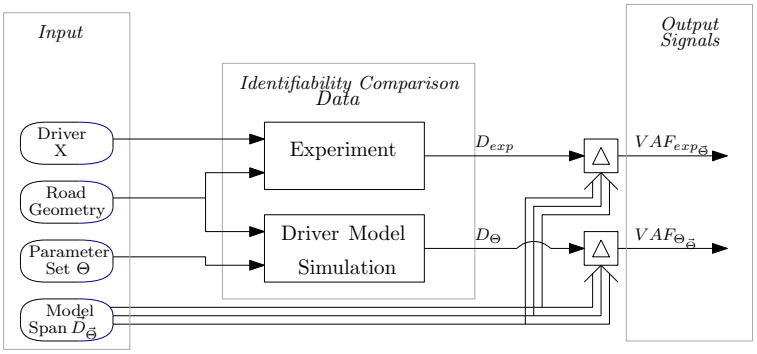

Figure 2.5: This figure illustrated a flow-diagram that presents how the **Identifiability Preprocessing Block** is *specifically* implemented. There are two types of identifiability considered, that of experimental identifiability and inherent identifiability. For the experimental identifiability an empirical drivers output is used (D_{exp}) and compared to the model span to output the experimental identifiability VAF matrix ($VAF_{exp_{\vec{\Theta}}}$). For inherent identifiability, a single modelling output (D_Θ) is used for comparison with the model span to output the inherent identifiability VAF matrix ($VAF_{\Theta_{\vec{\Theta}}}$).

$$D_{\Theta_1} = D_{\Theta_2} \rightarrow \Theta_1 = \Theta_2 \quad for\ all\ \Theta_1, \Theta_2 \in \Theta \tag{2.1}$$

As given in Fig.2.5, this unique mapping is tested for both inherent and experimental identifiability. For inherent identifiability, the model's output given a parameter setting Θ is D_Θ. Whereas for experimental identifiability a real driver X is to do an experiment to obtain their driver steering behaviour D_{exp}. Simulated time domain response $D_{\hat{\Theta}}$ (a single realisation of the *model span* $\vec{D}_{\vec{\Theta}}$) is compared through the VAF metric to the reference time domain series D_Θ for inherent identifiability and D_{exp} for experimental driver identifiability. Equation 2.2 illustrates the VAF equation for inherent identifiability, whereas for experimental driver identifiability D_Θ is substituted with D_{exp}.

$$\text{VAF} = \left(1 - \frac{\sum_{k=1}^{N} |D_\Theta[k] - D_{\hat{\Theta}}[k]|^2}{\sum_{k=1}^{N} D_\Theta^2[k]}\right) \times 100\% \tag{2.2}$$

The VAF represents the normalised sum of errors between two signals in the time domain subtracted from unity, where N is the number of samples. The higher the VAF, the better the driver model output $D_{\hat{\Theta}}$ from the model span is able to fit the baseline time series D_Θ (for inherent identifiability) or D_{exp} (for experimental driver identifiability). A VAF of 100% means that measured and modelled signals are identical, and the *quality of fit* (as given in Fig.2.1) is high. In principle, other metrics that measure quality of fit can be used; however, VAF is preferred as it is widely used in identification literature [19, 174]. The Variance Accounted For (VAF) between the signals D_Θ, D_{exp} and the $\vec{D}_{\vec{\Theta}}$ is used to obtain $VAF_{\Theta_{\vec{\Theta}}}$ and $VAF_{exp_{\vec{\Theta}}}$ respectively. These are used as input signals for the identifiability criteria.

A quantification of identifiability is defined as the identifiability space, i.e., the parameter space where the VAF is between 95-100%. When this space is large, the mapping between driver behaviour and parameter set becomes less sensitive, less 'unique'

and less identifiable (i.e. the steering behaviour resulting from Θ is almost the same as 2Θ). Therefore there is an inversely proportional relationship between identifiability and identifiability space.

In this chapter, two driver outputs D are of interest: lateral position s_{lat} and the steering deflections δ_s, as these two model outputs are mostly used for human controller identification [12, 19, 31].

2.3. MODELS FOR ASSESSMENT

This section carefully describes two simple driver steering models, which are assessed in this chapter using the criteria defined in Section 2.2. Both models have the same underlying concept of how drivers control a car: by performing proportional control (with control gain K_y) on a predicted lateral position (a function of look-ahead time t_{LH}). Although both models, are two-parameter models, having a gain K_y and look-ahead time t_{LH}, they are very different on a detailed level, making them suitable for such model performance comparison and scrutiny. The two models tested in this chapter are referred to as the 'linear' and 'nonlinear' prediction model. Both models are lateral control models, assuming constant velocity. To facilitate a fair comparison with the experimental driving data in [31], the vehicle dynamics used for both models in this chapter is identical to [31], a modified bicycle model, presented in Appendix A.1.

2.3.1. LINEAR PREDICTION MODEL

The linear prediction model, see Fig. 2.6, was used for individualization of haptic shared control through driver identification in [31]. It is one of the lowest-order and most simple steering models of curve driving behaviour available.

Although this model controls on predicted lateral position, the structure in Fig. 2.6 shows that the model is a compensatory PI controller, taking current heading as a compensatory input cue, $(\psi_r(t) - \psi_g(t))$. The vehicle's lateral velocity with respect to the road center $v_{lat_r}(t)$ can be estimated using a small angle approximation and a constant longitudinal car velocity v_{long} assumption, as shown in Eq. (2.3). This causal linear approximation of lateral position based on a current heading cue uses a similar fundament as the approximation of lateral velocity made in the vehicle/road dynamics of [145].

$$v_{lat_r}(t) = v_{long}(\psi_r(t) - \psi_g(t)) \tag{2.3}$$

Here $\psi_r(t)$ is the heading angle of the road and $\psi_g(t)$ is the heading angle of the car. The approximated lateral position can be continuously integrated to obtain an approximation of lateral displacement from the reference lane $\hat{s}_{lat}$, as given in Eq. (2.4).

$$\hat{s}_{lat}(t) = \int_0^t v_{lat_r}(t)dt \tag{2.4}$$

A prediction of the future lateral position is given as a *linear* extrapolation given the current lateral velocity $v_{lat_r}(t)$ multiplied by a look-ahead time t_{LH}, summed to the current approximation of lateral position $\hat{s}_{lat}(t)$, as given in Eq. (2.5).

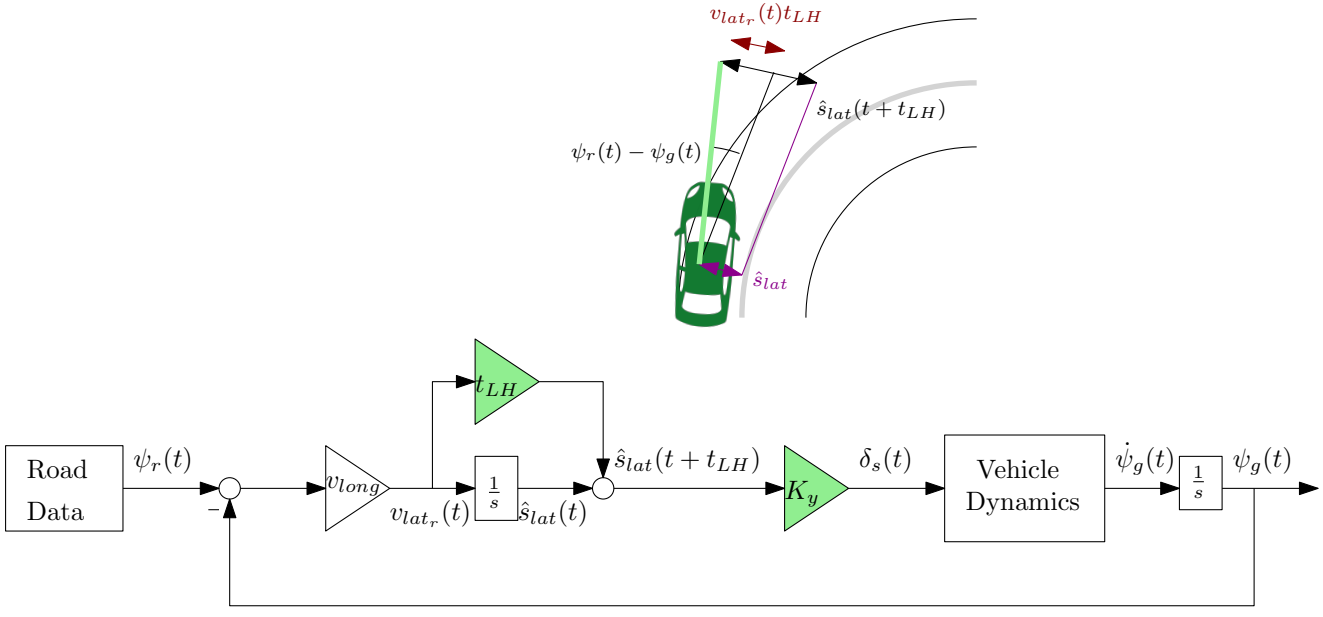

Figure 2.6: This figure illustrates the linear-prediction driver-model architecture used for identification in [31]. The model parameters are indicated in green: the lookahead time (t_{LH}) and the compensatory position error gain (K_y). The road heading $\psi_r(t)$ at current time is taken as reference input, then, the heading error ($\psi_r(t)-\psi_g(t)$) is scaled with longitudinal velocity (v_{long}) to obtain the lateral velocity ($v_{lat_r}(t)$), by making a small angle approximation. The lateral velocity is integrated to obtain an estimate of the current lateral position ($\hat{s}_{lat}(t)$). Additionally, the lateral velocity is multiplied by lookahead time (t_{LH}) and summed to the current estimate of lateral position to obtain a linear estimate of lateral position ahead ($\hat{s}_{lat}(t+t_{LH})$). The lateral position ahead is then scaled with gain K_y.

$$\hat{s}_{lat}(t + t_{LH}) = \hat{s}_{lat}(t) + v_{lat_r}(t)\, t_{LH} \tag{2.5}$$

Note this causal prediction of lateral position is then not compared to the future nearest road point, in contrast to [186] and [144]. The driver is modelled to perform a proportional control on a linear causal prediction of the future lateral position as given in Eq. (2.6).

$$\delta_s(t) = K_y \hat{s}_{lat}(t + t_{LH}) \tag{2.6}$$

This type of linear causal prediction as a function of heading error results in large inaccuracies of lateral position prediction during a curve, that could already be improved when taking heading rate into account [145]. Moreover, with the lack of a feedback loop on the road ahead, this model does not include any anticipation or preview.

2.3.2. NONLINEAR PREDICTION MODEL

Eq. (2.6), presenting the human performing proportional control K_y on predicted lateral position, a function of look-ahead time t_{LH}, is also applicable to the nonlinear prediction model as given in Fig 2.7. However, in contrast to the linear prediction model, the nonlinear model predicts by taking vehicle yaw rate $\dot{\psi}_g(t)$ and body lateral velocity $v_{lat}(t)$ into account. Estimating $\hat{s}_{lat}(t + t_{LH})$ is then done by closing the future lateral position loop. Consequently, the nonlinearity stems from both the nonlinear mapping between look-ahead time and the time-advanced vehicle states and the geometrical calculations needed to determine the future position and lateral position (w.r.t the road).

This driver steering model was implemented and tested as a part of the haptic shared controller in several previous studies [128, 129]. Contrary to a driver modelling study, the focus of these studies was on the practical effect of the magnitude of force and/or stiffness feedback of a haptic shared controller. Even though one of the underlying variables used for control is the predicted future lateral position, the gain and look-ahead time used to determine predicted lateral position were not varied, nor was their effect studied. Nevertheless, this type of lateral position prediction for haptic shared control has also been used in [32].

From a conceptual stance, this driver steering model has an interesting way of combining elements of the Successive Organisation of Perception [89], where instead of applying compensatory stabilizing control to current vehicle states and feedforward to the future road ahead, we apply a compensatory loop to a driver's estimation of future vehicle state on the previewed road ahead. This is is similar to one of the loops in the multi-loop driver steering model in [185].

Where this model differs from the time-advanced lateral position loop in [185] is how the time-advanced lateral position is calculated. In the algorithm 1 we show how the vehicle state is extrapolated from its vehicle state ($X(t)$, $Y(t)$, $\psi_g(t)$, $x_{vd}(t)$), with a constant steering wheel input δ_s to the predicted position vector ($\hat{X}(t + t_{LH})$,$\hat{Y}(t + t_{LH})$). Essentially manifesting a simulation within the simulation presented in Fig. 2.7. Algorithm 1 uses two functions, VehicleDyn() and Body2Global() which are the same as the blocks *Vehicle Dynamics* and *Body to Global Transformation* in Fig.2.7, both elaborated in Appendix A.1 and A.2 respectively.

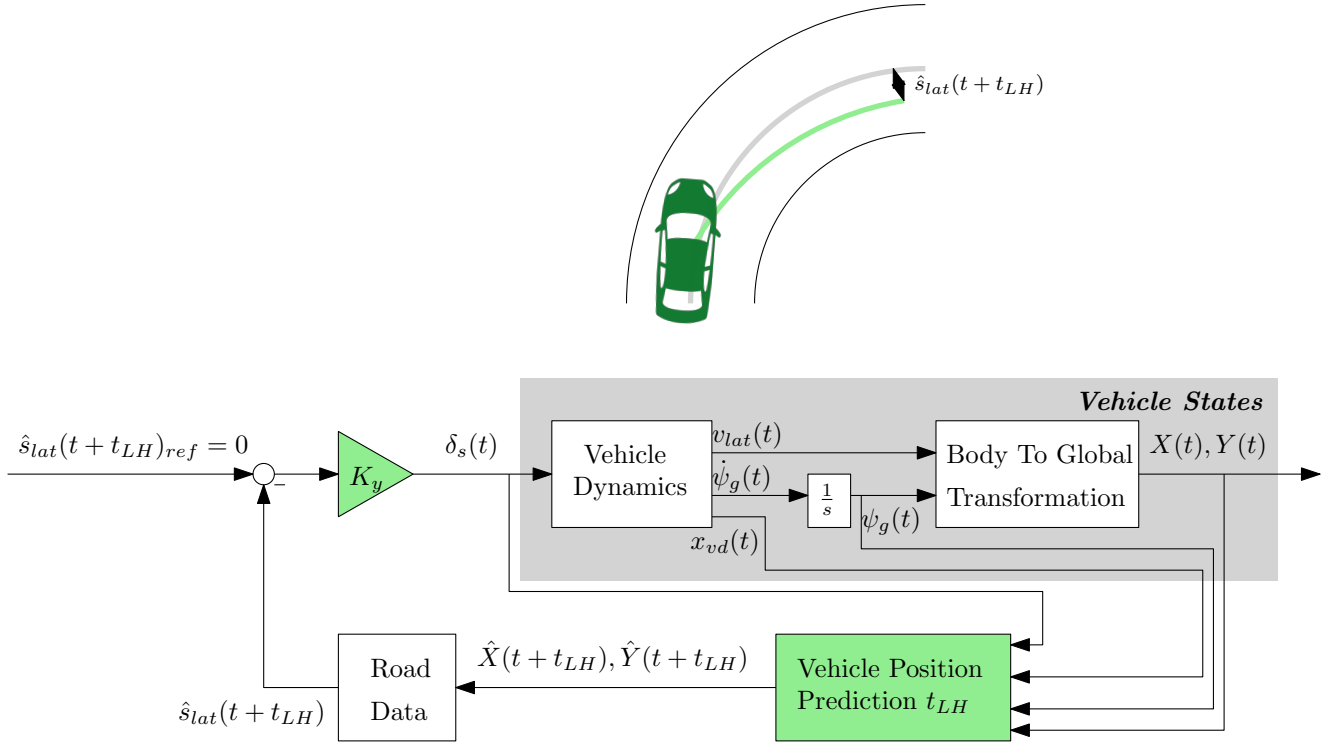

Figure 2.7: This figure illustrates the nonlinear prediction driver-model architecture. The model parameters are highlighted in green, these are the prediction time t_{LH} and compensatory gain K_y. Given the current steering angle ($\delta_s(t)$), vehicle states ($x_{vd}(r)$), heading angle ($\psi_g(t)$) and position ($X(t), Y(t)$), a nonlinear prediction (including yawrate) is made iterating the vehicle dynamics t_{LH} s ahead. The predicted position ahead ($\hat{X}(t+t_{LH}), \hat{Y}(t+t_{LH})$) is mapped onto the physical previewed road ahead to obtain the estimate for future lateral position ($\hat{s}_{lat}(t+t_{LH})$). This position is then compensated for by gain K_y.

Algorithm 1: Vehicle Position Prediction t_{LH}

Input : $X(t)$, $Y(t)$, $\psi_g(t)$, $x_{vd}(t)$, $\delta_s(t)$
Output : $\hat{X}(t+t_{LH})$, $\hat{Y}(t+t_{LH})$
Constant Parameters: v_{long}, t_{LH} , dt

for i = t:dt:t+t_{LH}-dt **do**
  $[\ v_{lat}(i)\ \dot{\psi}_g(i)\ x_{vd}(i+1)\ ]^T$ = VehicleDyn($x_{vd}(i)$,δ_s)
  $\psi_g(i+1) = \psi_g(i) + \dot{\psi}_g(i)$dt
  $[\ v_y(i)\ v_x(i)\ ]^T$ = Body2Global(v_{long}, $v_{lat}(i)$, $\psi_g(i)$)
  $X(i+1) = X(i) + v_x(i)$dt
  $Y(i+1) = Y(i) + v_y(i)$dt
end

With the linear vehicle dynamics (given in the discrete state-space VehicleDynamics function), the mapping between steering input $\delta_s(t)$ and iterated vehicle state outputs $v_{lat}(t+t_{LH})$, $\dot{\psi}_g(t+t_{LH})$ is linear. However, the mapping between steering input δ_s and vehicle position $\hat{X}(t+t_{LH})$,$\hat{Y}(t+t_{LH})$ is not, due to the Euler transformation matrices.

Merely iterating this state space an t_{LH}/dt number of times does not make the calculation nonlinear. However, the relationship between look-ahead time t_{LH} and the vehicle output states $v_{lat}(t)$ and $\dot{\psi}_g(t)$ is. With the current stable vehicle dynamics, feeding in a constant steering wheel input and iterating the vehicle's state, results in a step input. For a stable LTI state space, all states will reach a steady-state response, implying a nonlinear relationship between state outputs and iteration time. Moreover, the relationship between t_{LH} and vehicle position $\hat{X}(t+t_{LH})$,$\hat{Y}(t+t_{LH})$ is also nonlinear due to the transformation matrices. Finally, the predicted position $\hat{X}(t+t_{LH})$,$\hat{Y}(t+t_{LH})$ is used to find the smallest Euclidean distance to the road, which is then $\hat{s}_{lat}(t+t_{LH})$.

2.4. APPLICATION OF ASSESSMENT METHOD

This section presents the outcome of assessing both the linear and the nonlinear prediction models with the developed criteria. First, the descriptiveness of the models is evaluated, including a comparison to real human - driving data. Then, the *inherent* identifiability of the models for D_Θ is evaluated along with the *experimental* driver identifiability of curve cutting subject 5, representing D_{exp} from [31]. Finally, the realism criteria is assessed.

2.4.1. DESCRIPTIVENESS

Figs. 2.8 & 2.9 show example results of the driver model descriptiveness test, where the descriptiveness of the models is shown through the shaded model capabilities area given in Fig. 2.8 for the linear prediction model and Fig. 2.9 for the nonlinear prediction model. This area encompasses the full set of solutions in the s_{lat} domain within

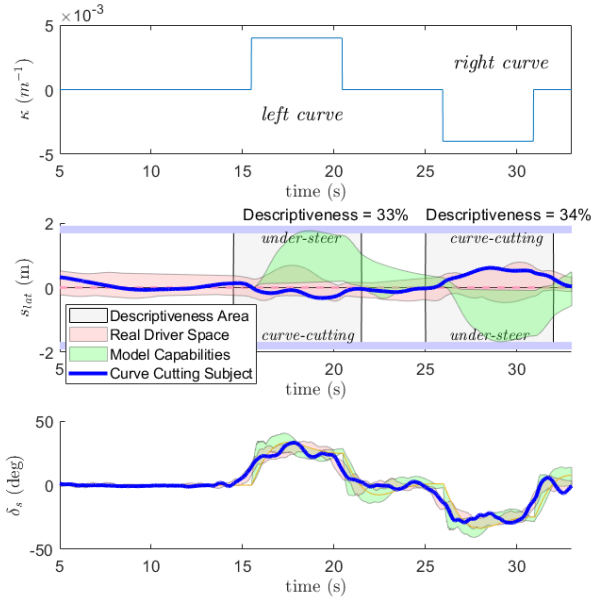

Figure 2.8: An illustration of Descriptiveness for the linear prediction model. Top figure: curvature profile, middle figure: Model Descriptiveness of Curve-Negotiation in s_{lat} domain with real data from a curve cutting subject from [31], bottom figure: corresponding δ_s

the limitation of realistic steering reversals. The corresponding steering deflections δ_s is also given.

For the linear prediction model in the s_{lat} domain, it is clear that only under-steer driver steering behaviour can be reproduced, without even being able to track the exact centre of the curve. Whereas for the nonlinear prediction model, we see that the model capabilities cover both under-steer and curve-cutting. Due to the predicted lateral position *feedback loop,* we see that this model is also able to steer into the curve before the curve starts. However, this model is not able to steer into the curve on the outer part of the lane.

Data recorded from a real subject (Subject 5 from [31]) obtained in a driving simulator experiment where the subject negotiated the exact same curvature profile is given in blue, representing curve-cutting behaviour. The span of the curve negotiation behaviour of all 12 participating subjects presented in the *real driver space* is shaded in light pink. It is remarkable to see that for the linear prediction model, only a fraction of the real driver space (14%) falls within the model capabilities in the s_{lat} domain. In contrast, for the nonlinear prediction model, the overlap is much more significant (89%). In the steering angle domain δ_s, the linear prediction model has more overlap with 68% falling within the model capabilities. Yet, the nonlinear prediction model has even more overlap with 93% falling within the model capabilities. Variability in steering behaviour is more apparent and distinguishable in lateral position than in steering

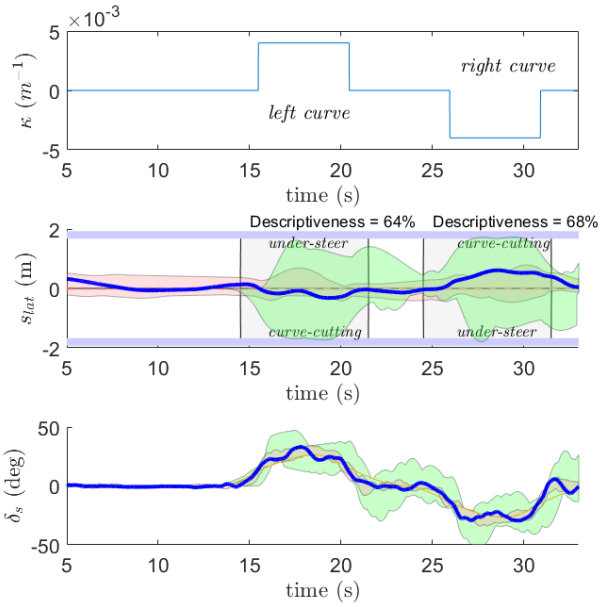

Figure 2.9: An illustration of Descriptiveness for the nonlinear prediction model Top figure: curvature profile, middle figure: Model Descriptiveness of Curve-Negotiation in s_{lat} domain with real data from a curve cutting subject from [31], bottom figure: corresponding δ_s

wheel angle. The suitability of a model to capture different driver steering behaviour in curve negotiation should therefore not just be evaluated on steering wheel angle, but also lateral position.

2.4.2. IDENTIFIABILITY

The ability of model parameters to be accurately identified from data can be quantified in two ways. Firstly through inherent identifiability, where the quality of fit between the models steering behaviour given a particular Θ and the model span (based on parameter combinations defined in the parameter span) is determined. The sensitivity of the driver model's steering behaviour (to changes in parameters) is visible with parameter sets that also result in a good fit around the ground truth given as the underlying Θ. With experimental identifiability, an optimal parameter combination is found by finding the best quality of fit for Driver X within the model span, without such ground truth.

The inherent identifiability of a model is evaluated by calculating the VAF for a wide range of parameter combinations resulting in driver model outputs' $\vec{D}_{\hat{\Theta}}$, against a modeled/simulated data set D_{Θ}. Whereas the identifiability of an experimental data run is calculated by evaluating the VAF for all parameter combinations $\vec{D}_{\hat{\Theta}}$ against an experimental dataset D_{exp}, which in this case represents curve-cutting. The range of

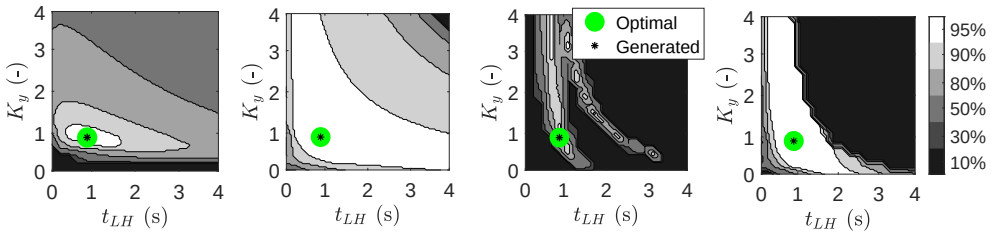

(a) In the s_{lat} domain, simulated data has an identifiability space of 0.5

(b) In the δ_s domain, simulated data has an identifiability space of 7.5

Linear Prediction Model ((a) & (b))

(c) In the s_{lat} domain, simulated data has an identifiability space of 0.01

(d) In the δ_s domain, simulated data has an identifiability space of 3.8

Nonlinear Prediction Model ((c) & (d))

Figure 2.10: Inherent identifiability and model sensitivity, in terms of Variance Accounted For (VAF) based on simulated data where K_y=1, t_{LH} = 1s. The green dot shows the parameter set for the optimal VAF, and the black asterisk shows the parameter set used to generate D_Θ

values evaluated to obtain the matrix $\vec{D}_{\hat{\Theta}}$, for parameters $\{ K_y, t_{LH} \}$ are $K_y \in (0,4]$ (from 0 till 4, excluding 0 and including 4) and $t_{LH} \in (0,4]$. For the two-parameter models used, the large VAF matrices are illustrated visually using a *heat-map* as done in Fig 2.10.

INHERENT IDENTIFIABILITY

For the inherent identifiability analysis, the simulated data set D_Θ is for parameter values K_y =1, t_{LH} = 1s. As it is impossible to illustrate the inherent identifiability for all initial sets of parameters we chose this single parameter set as this was the most frequently identified parameter combination in [31] and the default settings used for [128]. Figs. 2.10a,b and Figs. 2.10c,d illustrate the VAF data in the s_{lat} and δ_s domains for the linear and nonlinear prediction models, respectively. It can be seen that the parameter set corresponding to the optimal VAF is identical to the generated data set for $K_y = 1$, $t_{LH} = 1$s.

Fig. 2.10a presents the inherent identifiability space of the lateral position of the linear prediction model with a total integrated white VAF area of 0.5s, with an asymmetrical distribution, spanning $t_{LH} \in [0.4, 1.9]$s and $K_y \in [0.7, 1.3]$. For the s_{lat} domain this is a wide range.

Fig. 2.10b illustrates the inherent identifiability of the steering wheel deflections for the linear prediction model. The identifiability space has a very large range almost spanning half of the solution space presented, spanning values of $K_y \in [0.1,4]$ and $t_{LH} \in [0.1,4]$s, indicating poor identifiability.

Fig. 2.10c presents the inherent identifiability space of the lateral position of the nonlinear prediction model with an area of 0.01s, spanning three very small areas around the coordinates (t_{LH}, K_y) at (1,1) ,(1.1, 3.4) and (2.4, 0.9). Here we see an example of good identifiability.

Fig. 2.10d illustrates the inherent identifiability of the steering wheel deflections for the nonlinear prediction model. The identifiability space has a distribution spanning half of the solution space presented, spanning values of $K_y \in [0.1,4]$ and $t_{LH} \in [0.2,2.5]$s, indicating a poor identifiability when only considering δ_s. However, if we

combine the δ_s identifiability space by the s_{lat} identifiability space (dot product), we find that only one small area around (t_{LH}, K_y) = (1,1) will come out as a solution.

When comparing the inherent δ_s identifiability spaces of the linear model with the nonlinear model, it is interesting the see that the nonlinear prediction model shows the same insensitivity to gain $K_y \in [\ 0.1,4\ ]$. However, the sensitivity of the lookahead time is larger than the linear model, reducing this range to $t_{LH} \in [\ 0.2,2.5\ ]$s. This indicates that the lookahead time distinguishes between different δ_s steering behaviour better in the nonlinear model, which, can be related to the fact that there is a feedback loop on predicted lateral error for the nonlinear prediction model.

EXPERIMENTAL IDENTIFIABILITY

Figs. 2.11a, 2.11b, 2.11c and 2.11d illustrate the experimental identifiability of the models given a real curve-cutting data set D_{exp} of Subject 5 from [31]. The optimal parameter set that maximizes the VAF is given in green. A solution within the 95-100% VAF limit is outlined in red, which is scrutinized from another (realism) perspective in Section 2.4.3. This example illuminates the challenges of accurate identification given models of different descriptiveness.

Fig. 2.11a illustrates the experimental identifiability map of the linear prediction model for Subject 5 in the s_{lat} domain. In combination with Fig. 2.8, the linear prediction model is clearly unable to replicate the experimentally observed curve-cutting in the s_{lat} domain. The highest VAF found in the parameter solution space is around 8%, which is low.

Fig. 2.11b illustrates the experimental identifiability map of the linear prediction model for Subject 5 in the δ_s domain. The identifiability space is significantly reduced compared to the inherent analysis of Fig. 2.10b. In Fig. 2.8 it is seen that steering angle is not perfectly replicated; however, we still find a relatively large identifiability space of 1.1s. In combination with the lack of identifiability space in the s_{lat} domain, this indicates that based on identification in the δ_s domain, it is not possible to discriminate between different driver steering behaviour.

Fig. 2.11c shows the experimental identifiability map of the nonlinear prediction model for Subject 5 in the s_{lat} domain. Although the trajectory of Subject 5 lies within the model capabilities, see Fig. 2.9, the maximum VAF that can be achieved is a low 58%, instead of 95-100%, indicating that the model can not fully replicate this driver's trajectory.

Fig. 2.11d illustrates the experimental identifiability map of the nonlinear prediction model for Subject 5 in the δ_s domain. The identifiability space is somewhat reduced from Fig. 2.10d. In Fig. 2.9, we see that steering inputs of subject 5 mostly falls within the model's capabilities, except in entering and leaving the curve. Given that the range of identified values of t_{LH} is between 0.5 and 1.7s and K_y between 0.2 and 4. This shows that in the δ_s domain, the identifiability of t_{LH} is much better than K_y. Meaning that, the parameter t_{LH} differentiates (more) between the different drive lines encompassed in Fig. 2.9, as the further you look, the more you curve-cut.

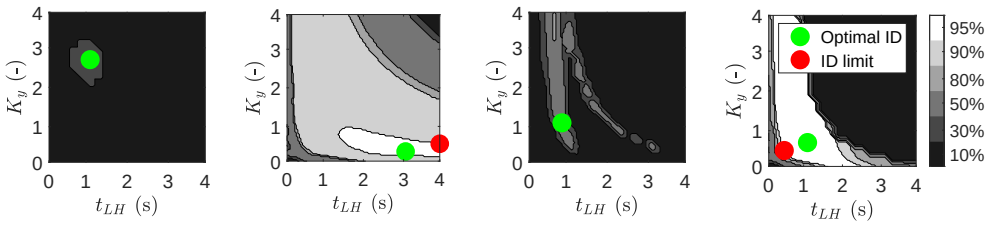

(a) In the s_{lat} domain, *subject 5* has no identifiability space

(b) In the δ_s domain, *subject 5* has an 1.1s identifiability space

(c) In the s_{lat} domain, *subject 5* has no identifiability space

(d) In the δ_s domain, *subject 5* has an 2.5s identifiability space

Linear Prediction Model ((a) & (b))

Nonlinear Prediction Model ((c) & (d))

Figure 2.11: Experimental identifiability, in terms of Variance Accounted For (VAF) based on real curve-cutting subject data. The green dot shows the parameter set for the optimal VAF

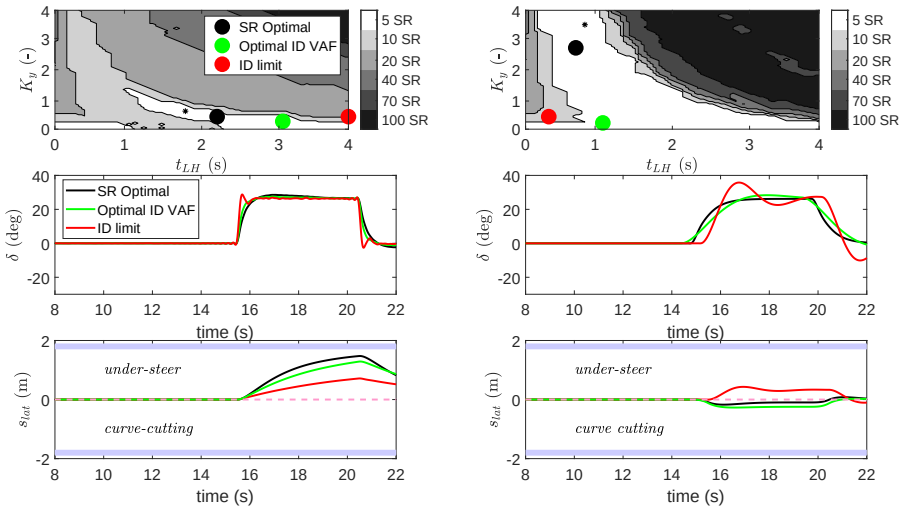

(a) Linear Model

(b) Nonlinear Prediction Model

Figure 2.12: Illustration of Realism Heatmap of steering reversal rates (top panels) for the linear model (left panels) and the non-linear model (right panels), including time traces of three different sets of parameters. SR optimal, the parameter set corresponding to the optimal SRs, Optimal ID the parameter set corresponding to the optimal VAF of experimental identifiability in the δ_s domain and ID limit a parameter set on the limits of the identifiability space of experimental identifiability in the δ_s domain. Middle panels: a comparison of steering wheel deflections relating to points on the heatmap and a human δ_s, bottom panels: illustration of the corresponding s_{lat}.

2.4.3. REALISM

The realism criteria focuses on defining a feasible parameter space that results in realistic driver steering behaviour. In this chapter this is quantified as an acceptable number steering reversals. For each parameter combination in the parameter space $\{K_y, t_{LH}\}$ for $K_y \in (0,4]$ and $t_{LH} \in (0,4]$, the steering reversals are evaluated and presented in a SR-heatmap as illustrated in Fig. 2.12a. Three points are shown in the map,

the *SR optimal* (minimal amount of steering reversals), the *optimal ID VAF*, a point corresponding to the optimal VAF for subject 5 illustrated in Fig. 2.11b, and the *ID limits* is also a point that falls within the bounds of the identifiability parameter space for subject 5 in the domain of δ_s. Fig. 2.12 middle illustrates the δ_s corresponding to the three points on the SR-heatmap and Fig. 2.12 bottom presents the corresponding s_{lat}.

Fig 2.12 shows that both the optimal SR and the optimal ID parameter combination for Subject 5 produce acceptable δ_s, in both the linear and nonlinear prediction model. However, this is not the case for δ_s corresponding to a parameter combination on the identifiability limits (ID limit). It is still remarkable to see that the limiting solution that fits within the identifiability parameter space (95-100% VAF), even producing a smooth trajectory in s_{lat}, may have unacceptable δ_s oscillations. Both the linear and the nonlinear prediction model produce a solution which may not be suitable for some driver assistance system applications; nevertheless, the linear prediction model produces more oscillations. This is due to the SR gradient on the considered boundary of the acceptable steering oscillations space; the nonlinear prediction model has a more smooth gradient on the chosen boundary location than the linear prediction model. Moreover, the nonlinear prediction model can provide a larger space of acceptable steering oscillations. Where the linear prediction model has an area of 2.5, the nonlinear prediction model has an area of 6.8.

2.4.4. QUANTITATIVE MODEL PERFORMANCE COMPARISON

This section objectively compares the performance of the two models using the criteria elaborated in the previous sections. The model assessment results are summarized in Table 2.1.

Table 2.1 facilitates a pairwise objective summarized comparison between the two models, which shows that the nonlinear prediction model outperforms the linear prediction model in many aspects. The nonlinear prediction model has a larger descriptiveness percentage, meaning it is capable of capturing a wider range of driver steering behaviour. Moreover, it has better inherent identifiability (smaller identifiability area) in both the δ_s and s_{lat} domains, and a larger realism space (larger SR area). However, at first glance, for the experimental curve-cutting identifiability, it is not clear whether the linear prediction model is better or worse than the nonlinear prediction model. Incorporating knowledge from the descriptiveness area, we see that the reduced parameter solution space for δ_s for the linear prediction model is in fact because the model is not able to capture this type of trajectory. This is reflected by the fact that the maximum VAF for the nonlinear prediction model is 58% in Fig. 2.11c whereas for the linear prediction model it is only 8% (in Fig. 2.11a). Moreover, suppose we superimpose both the s_{lat} and the δ_s identifiability spaces. In that case, the optimal solution in s_{lat} for the nonlinear model falls within the identifiability solution space of δ_s, illustrating an overlapping solution. Whereas for the linear prediction model, the solutions are not coherent, proving that one can not reproduce this type of trajectory. Nevertheless, the fact that we do not see a 95-100% replicability from the nonlinear prediction model in the s_{lat} domain means that this type of individualized trajectory, with its distinctive pattern, is still not fully captured by this driver steering model.

Table 2.1: Quantitative comparison between the linear and nonlinear prediction models

Criteria	Linear Prediction Model	Nonlinear Prediction Model
Descriptiveness		
s_{lat} Span	33 %	66 %
Identifiability		
Inherent: $K_y, t_{LH} = (1,1)$		
δ_s	7.5	3.8
s_{lat}	0.5	0.01
Experimental: Curve-cutting		
δ_s	1.1	2.5
s_{lat}	N.A.	N.A.
Solution Overlap?	No	Yes
Realism		
SR area	2.5	6.8

2.5. DISCUSSION

In this chapter, three driver model assessment criteria are defined: descriptiveness, identifiability and realism. These criteria determine the degree to which a given model structure is capable of capturing individual differences in driving steering behaviour. Moreover, these criteria are applied to find a model suitable for individualising intelligent DAS. Descriptiveness is defined as the ability of the model to capture a variety of different trajectories in the s_{lat} domain, quantified as a percentage (a 100 % indicates that the model can replicate all different trajectories). Identifiability has two different branches; inherent identifiability for a specific parameter set and experimental identifiability for a specific driver. Inherent identifiability indicates how sensitive a simulated time series (from a particular parameter set) is to a range of different parameters. Whereas experimental identifiability indicates how accurately the steering behaviour of an empirical driver can be identified given a particular model. An identifiability space is defined as the model parameter space satisfying a VAF > 95%. In this chapter, experimental identifiability of a curve-cutter (Subject 5) from [31], is evaluated. Finally, the realism criterion determines the space of solutions with acceptable oscillatory δ_s behaviour. As a demonstration of the proposed assessment procedure, these three criteria are applied to two models, each having two parameters, and each applies the concept that drivers perform proportional control on a predicted lateral position. These two models comprise 1) a linear (two-parameter) prediction model from [31], which was used for individualising haptic shared control, and 2) a nonlinear

(two-parameter) prediction model, that was used as underlying driver steering model for the haptic shared control studies [128, 129].

Within the framework of cybernetic driver steering models, there is much thought put into the structure of a model by for example, inferring the driver's mental model [155]. However, many are unable to capture a range of driver steering behaviour. This is reflected by kinematic models such as [139] only allowing for positive lateral deviation (curve-cutting), or in [157] where under-steer can not be described. This may not be problematic for a multitude of applications, however, when applied for individualisation of DAS such as in [31], the ability of the model to be *descriptive* becomes of paramount importance. The two-parameter model of [31] has a descriptiveness of 33% only allowing for under-steer. Despite its limited capabilities to capture different trajectories, it is used for identification and individualisation. The identifiability of the linear model (of [31]) to empirical curve cutting behaviour, had *no identifiability space* in the s_{lat} domain and a maximum VAF of 8%. This means that one can not describe the curve-cutting trajectory using the linear model, which, may pose a dispute about the conclusions drawn in [31], i.e., whether individualisation was truly realised in the implemented haptic shared controller. With low descriptiveness, the identifiability results become questionable and trajectory dependent. Moreover, using a model with low descriptiveness to individualise DAS is quite controversial.

During the identification of cybernetic models, a cost function is generally optimised. The results of this chapter show that the type(s) of driver output(s) selected for optimisation is important. In some studies a number of driver outputs are optimized simultaneously [138], others use only s_{lat} [139] or δ_s [31] [133], in combination with steering torque [146]. It is found in this study that although optimising for only δ_s may give a high VAF in the δ_s domain; it may be far from a good replication in the s_{lat} domain, where different trajectories are illuminated. This is because small variations in steering, which do not make a large difference in overall δ_s VAF can change the trajectory significantly. Therefore a high VAF for δ_s can be misleading. Still, the human control-action output is used to objectively identify human dynamics in many frequency domain analysis of manual pilot-control [19].

Inherent identifiability of a model for all modelling outputs is essential to obtain as a prerequisite before identification. Otherwise, the found results can be misleading, i.e., large identifiability spaces indicate non-unique solutions and large insensitivity to parameter variations. This becomes especially important when testing the effect of environmental factors or distraction on driver identified parameters [12]. Without knowledge of identifiability space or parameter sensitivity, the cause of parameter variations is ambiguous. These variations could either result from the independent variable tested (experimentally) or the inherent identification insensitivity. The inherent identifiability space in δ_s is vast for the linear model [31], which indicates non-unique solutions and bad identifiability. Whereas for s_{lat} it is much smaller, indicating a more promising identification domain for discriminating between different behaviours.

Although the s_{lat} domain facilitates better identifiability, one should not ignore the δ_s domain. Without constraining the parameter solution space to that of realistic SR, good trajectory fits will result but with possibly unrealistic δ_s. For applications

such as haptic shared control where the human haptically feels the provided δ_s, these deflections need to be comfortable. Comfort can be provided when the provided δ_s is coherent with the humans inherent neuromuscular filtering properties [109]—thereby suggesting to add constraints of realism in the identification process. This is because a parameter difference of 0.5s for t_{LH} may not make much difference in VAF; however, it can result in oscillations and user-rejection of the system.

This chapter defines three metrics in a general sense; however, there are limitations to these. Descriptiveness is quantified as a percentage of modelled area covered relative to the descriptiveness area (defined as the area spanning the road width during a curve). Although most trajectories should fall within the bounds of the descriptiveness area, the discontinuous edges of the descriptiveness area (a block) are unnatural. Moreover, the descriptiveness area does not indicate the shape of the trajectory, which is also crucial in defining a trajectory. This shortcoming in the definition of the descriptiveness criteria is evident with the nonlinear model. The trajectory of Subject 5 from [31] can not be successfully described using the nonlinear model, at a VAF of 58%. Yet, the trajectory of Subject 5 falls within the 'model capabilities' outlined in the descriptiveness test. Therefore defining a model's capabilities based on modelling output area is insufficient. An alternative would be to define driving styles that a model should capture. One could base these driving styles on inertial behaviour [35]; defensive, neutral and offensive or, trajectory typology categorisation [162]. Hereby, the descriptiveness metric can improve to include for such refinements.

This chapter gives a general definition of each criteria, and also provides a means to realise this general definition. This does not mean that the way in which these criteria are realised is the only way. We evaluate two driver model outputs δ_s and s_{lat} in our analysis. However, in variable velocity simulations, other driver outputs such as lateral acceleration and time-to-line crossing could also be used for defining different driver behaviours (descriptiveness) and identification, with an additional realistic lateral acceleration metric for realism.

We consider two models that are based on the same concept, but are mathematically realised very differently. The different criteria have all suggested that the nonlinear model outperforms the linear prediction model. However, the different criteria can be weighed to a specific application. For a study where the only objective is personalisation, i.e., the focus is on describing different driver behaviour, identifiability becomes less critical. Thereby, the 'best' model may not be the model that you *need* for your *application*.

Overall, these assessment criteria indicate a quantitative difference in performance between the two models. However, what it doesn't specify is the cause of this difference. On a superficial level, both the linear and nonlinear prediction models do the same: both assume that humans perform proportional control on a predicted future lateral position. Taking a closer look, the difference stems from both the way the predicted future lateral position is calculated and closing the loop on the future lateral position prediction. These two elements seem to be essential for driver modelling performance and quality improvement.

2.6. CONCLUSION

This chapter proposes a three-part method by which one can assess the practical capabilities of a given driver model. The three criteria outlined are descriptiveness, identifiability and realism. Descriptiveness quantifies to what extent the model can capture different driver steering behaviour. Identifiability evaluates the extent to which the mapping between a given driver steering behaviour and a parameter set is unique. Realism determines the parameter space that results in realistic driver steering behaviour.

Two simple two-parameter driver steering models are compared to demonstrate the effectiveness of this assessment method. Both models assume the human performs proportional control on a predicted lateral position; however, they differ in the implementation of the prediction. Therefore, they are named the linear-prediction and the nonlinear- prediction models. The assessment methodology was able to discriminate between these two models' performance for all three criteria. The nonlinear model outperformed the linear model in all three quantified criteria, it exhibits 1) twice as large descriptiveness, 2) twice as good inherent identifiability in the steering angles domain, 3) a 50 times better inherent identifiability in the lateral position domain and 4) 2.7 times larger realism space.

Overall, the presented results show that the proposed assessment method can quantify key differences between driver steering models. By providing explicit insight into different factors that quantify a model's success in replicating real driver behaviour, the method scrutinizes models by giving them an absolute quality 'grade' and thereby presents a platform for a fair comparison between different models.

3

Assessment of Three Control-Theoretic Driver Steering Models

Chapter 2 has introduced a driver model assessment method that is only applicable to compare similar driver models (i.e. with equal number and type of parameters). This chapter upgrades the assessment methodology presented in Chapter 2 to compare and assess different control-theoretic driver models with a different number of parameters. The upgrade results in two standardised metrics based on empirical data, which are: 1) descriptiveness, defined as the capacity of a model to describe individual trajectories and, 2) realistic identifiability, defined as the uniqueness of mapping between identified parameters and realistic trajectories - as constrained by staying in the lane, and sufficiently smooth steering angles. This upgraded assessment method is applied to three prominent control-theoretic models and the results are discussed.

The contents of this chapter are to be published as:

Barendswaard, S., Pool, D.M., Boer, E.R., Abbink, D.A., 'Comparison of Control-Theoretic Driver Models for Personalised Trajectory Generation using an Assessment Methodology', *IEEE Transactions on Human Machine Systems.*

3.1. INTRODUCTION

Personalisation of Advanced Driver Assistance Systems (ADAS), plays a vital role in increasing the acceptance of an ADAS, allowing it to support drivers with safe and comfortable driving [72]. For such personalised systems, the main challenge is to accurately match and predict a driver's steering behaviour, trajectories and states [114]. The aim is to support personal styles within the realm of safe driving behaviour. Applications of personalised ADAS with underlying driver steering models include: 1) Lane Departure Warnings (LDW), where the predicted driver state is used by ADAS to apply preventive measures [65, 160, 194] and, 2) Lane Keeping Assistance (LKA) where a drivers steering and trajectory preference is understood by predicting driver's desired steering through (individualised) models [145, 181]. Outside the traditional ADAS applications, personalised trajectories can also be used with haptic shared control guidance [152] and fully autonomous path prediction [71, 149], to support safe driving behaviour and increase acceptance. With an increasing interest in personalised ADAS systems, more effective estimation and prediction accuracies of driver behaviour are needed [8].

In one of the most basic steering manoeuvres, that of curve-driving, personalisation can be effectively realised through the automation identifying an individual driver's steering behaviour in real-time. Individuals are known to exhibit a range of different styles [41] and trajectory classes in curves [18], which indicates the presence of diversity in driver steering and trajectory behaviour. There are two properties a driver model must have to identify and describe the diverse driver trajectories available in the driving population. Firstly the model should be able to capture different types of drivers and, secondly the identified parameters for every kind of driving trajectory behaviour should be unique [17].

Driver steering models available in literature can be categorised into four different groups: 1) *data driven modelling*, [11, 24, 64, 171, 181], 2) *analytic modelling*, [25, 147, 157], 3) *cost-function based modelling* [30, 78, 86, 108] and, 4) *control-theoretic modelling* [116, 150, 173]. Although all four groups of driver models are able to describe individual driver trajectories and steering behaviour, only the control theoretic driver models can explicitly *identify* driver steering and trajectory behaviour with *meaningful parameters*. A model that exhibits meaningful parameters facilitates a platform in which we can understand human driver behaviour. Control-theoretic models place the focus on the human's control dynamics, where the driver is represented by meaningful blocks of transfer functions, each block representing elemental characteristics, like an anticipatory gain, time delay, and neuromuscular dynamics, which are parameters that have physical meaning. These individual blocks receive anticipatory (road following) feedforward and (various) compensatory (disturbance rejection) feedback signals [163], transforming them into signals that together generate the drivers steering behaviour [48].

Within the domain of control-theoretic driver steering models, it is common practice to demonstrate the performance of a driver steering model in terms of its capacity for matching measured steering angles. This is usually carried out on a (small) set of human driving data, without the aim of describing *individual* driver behaviour, either in the time domain [145] or the frequency domain [119, 163, 173]. However, it is easier

to match steering angles than trajectories, and fitting on steering angles approach does not verify a model's ability to capture *individual* driver steering behaviour. For that, trajectory behaviour needs to be described [17]. A complicating factor in judging the appropriateness of a control-theoretic model lies in the well-known trade-off between the ability to capture all types of behaviour, and the danger of over-parametrisation. For example, a model with a large number of parameters may be very descriptive, i.e., able to capture a variety of different individual driver trajectories, however, may not be very identifiable, i.e., exhibit unique parameter retrieval [49]. To quantify such a trade-off, Chapter 2 introduced a systematic analysis methodology that evaluates the performance of driver steering models. However, studies to *compare* how well different control-theoretic driver steering models can capture individualised trajectories and steering behaviour is, to the best of our knowledge, not yet dealt with in literature.

Therefore, this study introduces a comprehensive assessment methodology (as a refinement of Chapter 2) to *compare* prominent driver models. The refined assessment methodology comprises two metrics: descriptiveness and realistic identifiability. Descriptiveness determines the capacity of a model to capture different individual driver steering and trajectory behaviour. Realistic identifiability determines the standard degree of uniqueness of mapping between estimated parameter and driver steering and trajectory behaviour.

This comprehensive assessment methodology is applied to compare the following models:

1. **The Mars model** [114, 145], chosen for its popular use to estimate driver steering in haptic shared control applications [74, 145, 152, 183].

2. **The Van der El model** [173], chosen as this model is the state-of-the-art of objective (frequency domain) identification of driver perception of the road and execution of the driving task.

3. **The Van Paassen model** [177], a model developed within our group based on lessons learnt from previous shortcomings of personalisable haptic shared controllers [31].

This chapter is structured as follows: Section 3.2 describes the models for assessment and additional adjustments made for the comparison, the comprehensive methodology is presented in Section 3.3, assessment results in Section 3.4, discussion in Section 3.5 and last but not least conclusions in Section 3.6.

3.2. MODELS FOR ASSESSMENT AND ADJUSTMENTS

This section describes the adaptations made to facilitate a fair comparison between models, the models and the limitations of each model. Three models are assessed, namely: the Mars model [155], the Van der El model [173] and the Van Paassen model [177].

For a fair comparison, the simulated neuromuscular dynamics and vehicle/road dynamics are made identical. The respective vehicle/road dynamics consists of a bicycle model and linear computation of heading error ψ_L and trajectory (defined in

this chapter as lateral position relative to road centerline, without velocity information) s_{lat}. These vehicle dynamics use the road curvature as (extra) input, as given in the state space of Eq. (3.1). This system characterisation is based on the *bicycle/road* dynamics used in [145] [116] [74].

$$\begin{bmatrix} \dot{\beta} \\ \dot{r} \\ \dot{\psi}_L \\ \dot{s}_{lat} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 \\ 0 & 1 & 0 & 0 \\ V & l_s & V & 0 \end{bmatrix} \begin{bmatrix} \beta \\ r \\ \psi_L \\ s_{lat} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ -V \\ -l_s V \end{bmatrix} \kappa + \begin{bmatrix} a_{15} \\ a_{25} \\ 0 \\ 0 \end{bmatrix} \delta_s \tag{3.1}$$

The model of Eq. (3.1) consists of a bicycle model (see the a_x coefficients, elaborated in Appendix B.1), with the slide slip angle, β and the yaw rate, r as states. Additionally, the states heading angle error ψ_L and lateral deviation from centerline s_{lat} (trajectory) are linearly approximated as a function of road curvature input κ, the slide slip angle β and yaw rate r. The derivatives of these states: side slip rate $\dot{\beta}$, yaw acceleration $\dot{r}$, heading error rate $\dot{\psi}_L$ and lateral velocity $\dot{s}_{lat}$ are computed as part of the state space.

The neuromuscular dynamics H_{nm}, present in each of the models, is represented as a filter, as given in Eq. (3.2), to prevent steering inputs with too high frequencies. See Appendix B.1 for its value.

$$H_{nm} = \frac{1}{T_n s + 1} \tag{3.2}$$

To gain deeper insights into the reasons for possible assessment outcomes, it is important to gain a fundamental understanding of the assessed models. Therefore, the description and elaboration of each model is carried out in four parts: 1) concept, explained with the help of the concept diagrams in Figs. 3.1, 3.2 & 3.3, 2) explanation, elaborated based on the control diagrams in Figs. 3.1,3.2 & 3.3, 3) limitations and 4) degrees of freedom, i.e. number of free parameters.

3.2.1. MARS MODEL

CONCEPT

The Mars model as described in [155], is based on the hypothesis that drivers use near and far points ahead on the road for steering control, as illustrated in Fig. 3.1. This is realised in the model by considering two inputs, the near angle θ_{near} to a road waypoint close to the car and the far angle θ_{far} to the tangent point on the curve. Hereby, the near angle drives a stabilising compensatory control loop and the far angle enables anticipatory control in curves.

EXPLANATION

The control structure in Fig. 3.1, shows that the feedforward input signal is not directly θ_{far}, but the current road curvature κ. In fact θ_{far} can be approximated through Eq. (3.3) [155].

$$\theta_{far}(t) = D_{far}\kappa(t) \tag{3.3}$$

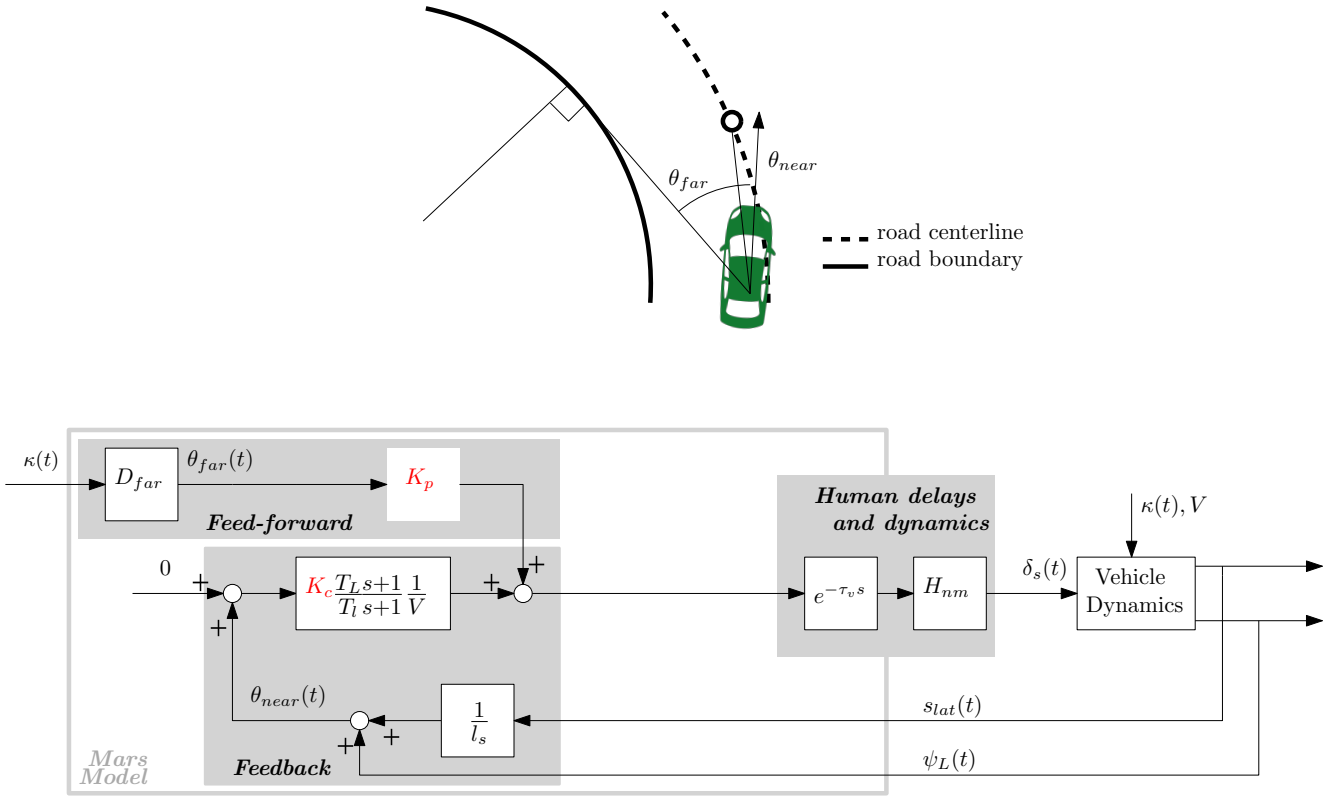

Figure 3.1: The top figure illustrates a schematic of the concept of the Mars model. This concept uses stabilising feedback on the near angle θ_{near} (i.e., angle to the near point, which is 5m ahead of the car on the road centreline, indicated with a white circle) and anticipatory feedforward on the far angle θ_{far} (i.e., angle to the tangent point on the curve, where this tangent makes a 90^o angle to the radius vector of the curve). The bottom figure illustrates the control diagram of the Mars model. The parameters used for identification are indicated in red.

The assumptions for this equation to hold are 1) the curvature of the road must be constant and, 2) the lateral distance from the centerline is constant. The geometric derivation for Eq. (3.3) is included in Appendix B.2 for reference. The $\theta_{far}(t)$ signal is scaled with a prediction gain K_p, to obtain the complete anticipatory response.

For the feedback loop, the heading error ψ_L and relative lateral position w.r.t. to road centerline s_{lat} are used to compute θ_{near}, given in Eq. (3.4). Here we see the summation of two angles, the heading error and the angle the car makes to the near point. This near point is assumed to be a waypoint on the road centerline a distance l_s ahead of the car. The assumption for this computation to be valid is that we are on a straight road, as elaborated in the geometric derivation in Appendix B.2.

$$\theta_{near}(t) = \frac{s_{lat}(t)}{l_s} + \psi_L(t) \tag{3.4}$$

The compensatory loop is then also scaled with gain K_c and consequently filtered, with a lead-lag filter, to reduce any noise present in the stabilisation loop. Both feedback and feedforward contributions are summed, and sent through the human-like time delay element $e^{-\tau_v s}$ and neuromuscular dynamics H_{nm} to obtain the steering wheel input δ_s.

LIMITATIONS

Although this model includes both feedback and feedforward input signals, with a structure very similar to that of McRuer et al. [119], it is lacking in terms of not having any explicit preview. The concept of the model claims to be driven by an anticipatory tangent point. However, this anticipation is based on linear extrapolation of the *current* road curvature, rather than the true road ahead. Moreover, the feedforward loop does not include any smoothing filter, which is problematic when encountering non-clothoidal roads.

DEGREES OF FREEDOM

The degrees of freedom for this model are two: the 'preview' gain K_p and compensatory gain K_c. The lead-lag values T_L and T_l are not considered as degrees of freedom as the original authors [155] have restricted their design values to serve a certain purpose: filtering out the noisy feedback data. The visual time delay τ_v is fixed to the set value (0.04s) by the original authors.

3.2.2. VAN DER EL MODEL

CONCEPT

The Van der El model [175] is a driver model that was obtained from objective frequency domain identification, following the tradition of McRuer, who used this method to identify the fundamentals of human control in tracking tasks [120]. The model is deduced from frequency domain data collected through a curve driving experiment with sinusoidal roads.

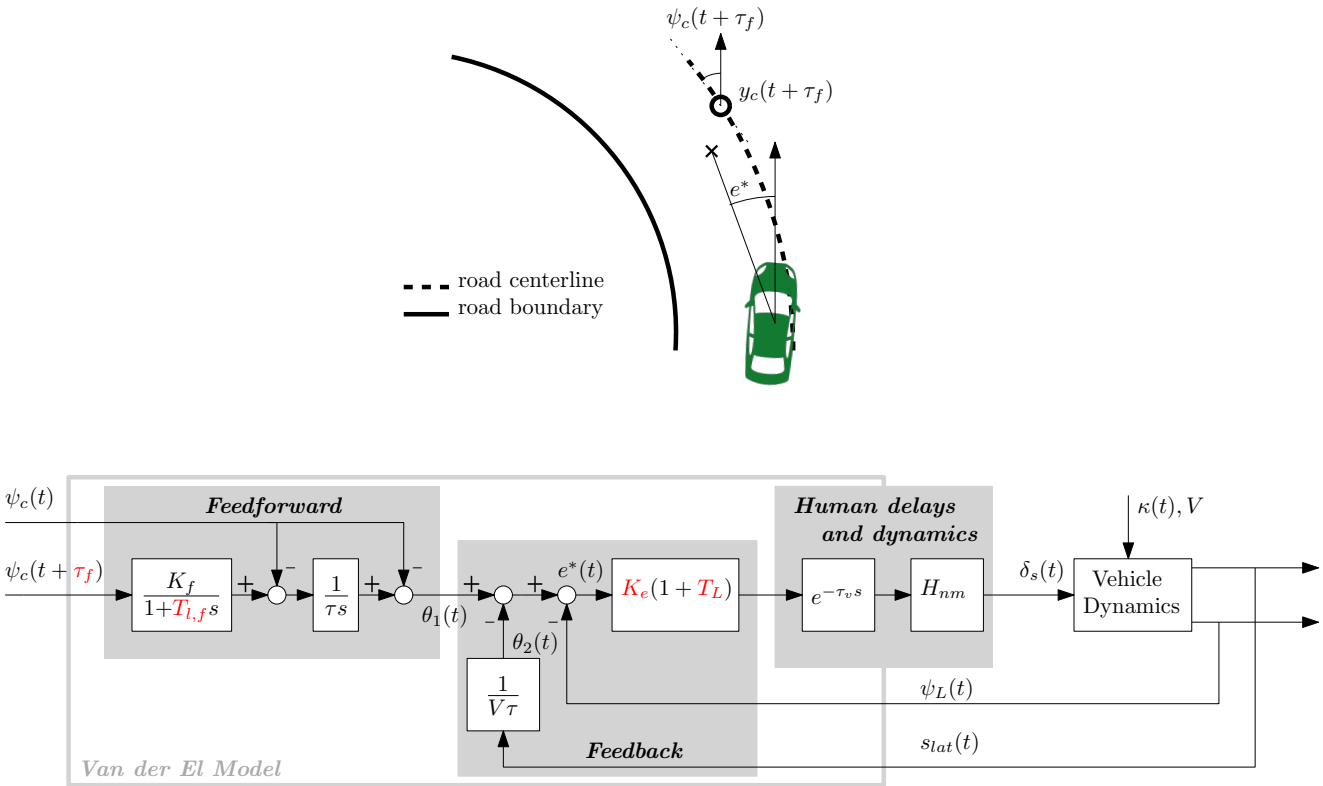

Figure 3.2: The top figure illustrates a schematic of the concept of the Van der El model. The Van der El model applies feedback on the angle (e^*) the cars axis makes to a previewed, filtered road waypoint ahead $y_c(t+\tau_f)$, this point is indicated as a "X" on the schematic of the concept. The bottom figure illustrates the control diagram of the *transformed* Van der El model. For details on the transformation, see Appendix B.3. The parameters used for identification are indicated in red.

Results of this objective identification procedure have shown that in the *abstract* experiment carried out by [175], drivers react to three inputs: 1) absolute car heading feedback 2) absolute car position feedback and 3) previewed road centerline position feedforward $y_c(t+\tau_f)$. The findings have resulted in a control-theoretic model that accurately explains the steering behaviour of human drivers in the *frequency domain*.

The identified model also has a geometric interpretation. The human is thought to anticipate a future waypoint $y_c(t+\tau_f)$ that is filtered (illustrated as 'x' in Fig. 3.2) of the previewed road centerline position. Through the additional heading and position feedback, the cascaded transfer functions result in a specific element being controlled in a compensatory fashion: the angle the car's principal axis makes to the filtered future waypoint ahead 'x' (e^*). This means the control concept is compensatory control on the angle the principal axis of the car makes to a filtered anticipated waypoint ahead.

EXPLANATION

The input signals to the model are objectively deduced from an abstract experiment where the road is a sum of sinusoids. Therefore, the input signals of the original model are not compatible with a practical driving situation. A key contribution of this chapter is transforming the Van der El model to one that applies to practical driving situations. The algebraic conversion to more realistic inputs as used here and shown in Fig. 3.2, is given in Appendix B.3.

The adapted control structure, shown in Fig. 3.2, makes the summation of three angles:

θ_1 the heading angle difference between the nearest road point heading and the road heading at τ_f.

θ_2 the angle between s_{lat} and the future filtered waypoint 'x'.

ψ_L the heading error w.r.t the road heading.

The first angle θ_1 represents the target, whereas the summation of θ_2 and ψ_L are compared to the target to obtain $e^*(t)$. The concept of the geometric interpretation of these angles are illustrated in Fig. B.3 in Appendix B.3.

For the computation of the first angle, we see the feedforward loop takes the road heading at a time τ_f ahead $\psi_c(t+\tau_f)$ and filters it through the transfer function $\frac{K_f}{1+T_{l,f}s}$. With a time constant of $T_{l,f}$ this reduces the effective lookahead to a point which is less than τ_f, as given in Eq. (3.5).

$$\tau = \tau_f - T_{l,f} \tag{3.5}$$

The filtered road heading at time τ ahead is then compared to current road heading, obtaining the heading 'increment'. This is then sent through two blocks which is merged into one in Fig. 3.2, as shown in Eq. (3.6).

$$\frac{1}{\tau s} = \frac{V}{s}\frac{1}{V\tau} \tag{3.6}$$

Firstly, the future heading error of the road is multiplied by velocity and integrated, and this results in position error the road makes to the future waypoint on the road. Secondly, the position error to future waypoint is divided by $V\tau$, the longitudinal distance to the waypoint at time τ ahead. With a small angle approximation, this results in the angle from the current road position to the future road waypoint. Subsequently, to obtain the angle increment needed, the current heading is subtracted from this amount, resulting in θ_1.

θ_2 is then added to θ_1. The car's current lateral position error s_{lat} is divided by $V\tau$, which computes the relative angle made to this waypoint, due to the lateral position offset to road centerline (θ_2). Finally, the third angle ψ_L the heading error is added in the last feedback loop. The resulting value e^* is the angle from the car's axis to the future waypoint. This is fed through an equalisation block, that exhibits a compensatory gain and lead time constant, where the functionality of this block is to minimise e^*.

LIMITATIONS

This model does not exhibit an independent feedforward path that can be separately tuned, which can potentially diminish its descriptive capabilities.

DEGREES OF FREEDOM

In the Van der El model there are four degrees of freedom, highlighted in red as given in Fig. 3.2. The lookahead time τ_f, the filtering time constant $T_{l,f}$, the equalisation gain K_e and the lead time constant T_L. The parameter K_f is kept constant, as it is found to always exhibit a value of 1 for curve driving [175]. The value of τ is a function of τ_f and $T_{l,f}$, as given in Eq. (3.5). The visual time delay τ_v is taken as constant value of 0.35s, as identified in [175], this is inherent in human response mechanisms.

3.2.3. VAN PAASSEN MODEL

CONCEPT

The Van Paassen model was built based upon knowledge and understanding of previous driver modelling shortcomings [31], it was developed for the sake of facilitating the personalisation of haptic shared control. This model [177] makes use of pure position feedback on a predicted (curve cutting) position ahead ($y_{cc}(t)$) and an independent feedforward path that provides the necessary steering to make the curve, through explicitly coupling the curvature feedforward input with the car's velocity.

Hereby, the concept is that the steering rate target to make the curve is set by the feedforward loop, whereas trajectory fine-tuning is left for the position feedback loop. What is noteworthy, is that the feedback loop predicts the position of the car ahead, and sets a future curve cutting reference y_{cc} for the feedback loop to follow. This anticipatory feedback concept is visualised in Fig. 3.3. This model is the first to explicitly declare a non-zero position feedback reference, which makes sense as drivers are found not to follow the centreline when driving [18].

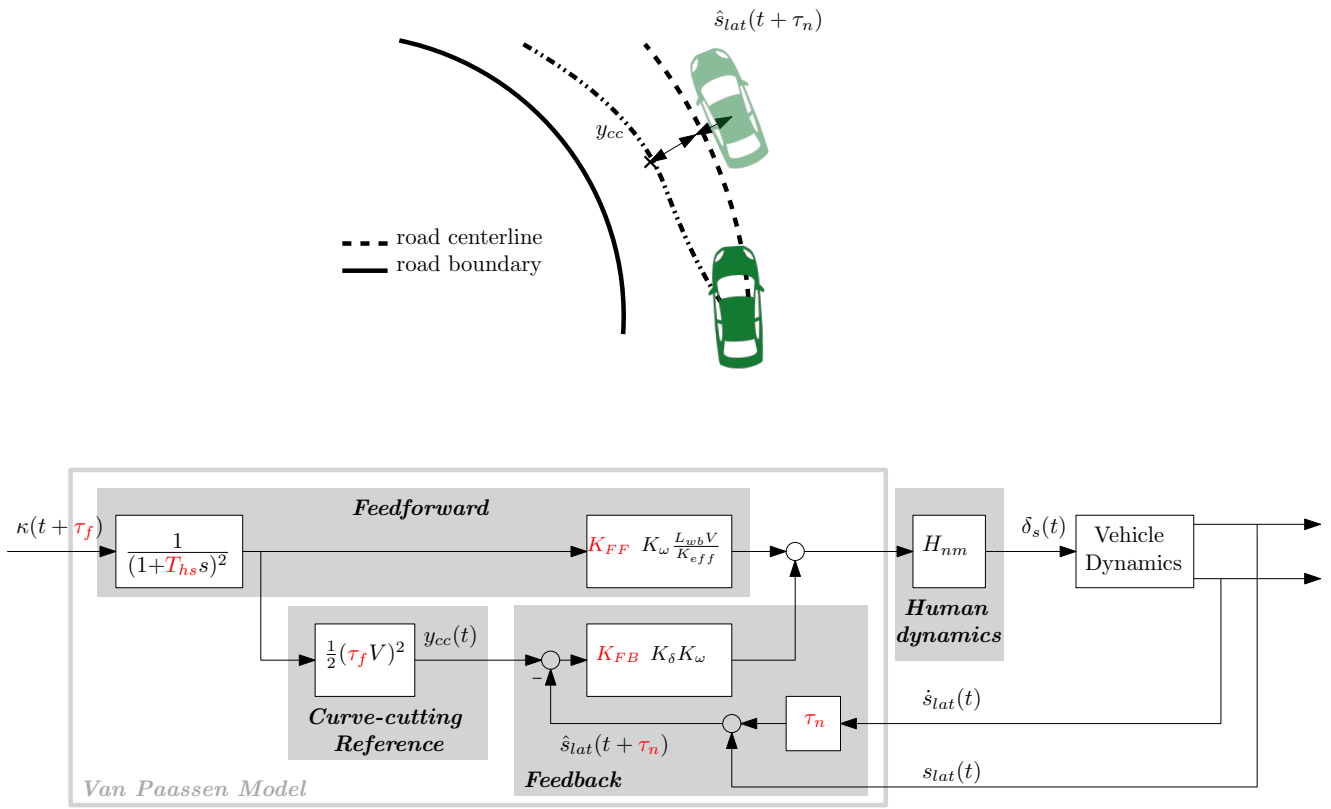

Figure 3.3: The top figure illustrates a schematic of the concept of the Van Paassen model. The Van Paassen driver model uses feedforward to set the steering rate needed to make the curve and uses feedback on the predicted lateral position ahead $\hat{s}_{lat}(t+\tau_n)$ with a curve-cutting distance y_{cc} as reference. Note that the concept illustration of the Van Paassen model excludes the explicit feedforward contribution, as given in the control diagram. The bottom figure illustrates the control diagram of the Van Paassen model. The parameters used for identification are indicated in red.

EXPLANATION

In Fig. 3.3 the feedforwarded signal $\kappa(t+\tau_f)$ is filtered with a second-order filter with a break frequency at T_{hs}, this is especially important when the road is non-clothoidal, to prevent any transition responses from the dynamical transfer functions. The filtered curvature signal is then used for both the feedback loop and the feedforward path.

In the feedback loop, the filtered curvature is used to compute the curve cutting distance as given in Eq. (3.7), the geometric derivation that makes this relationship possible is given in Appendix B.4.

$$y_{cc}(t) = \frac{1}{2}(\tau_f V)^2 \kappa(t+\tau_f) \tag{3.7}$$

The curve cutting distance is taken as a reference in the predicted position tracking feedback loop. The predicted position is computed by linear approximation at the prediction look ahead time τ_n (ideally $\tau_n = \tau_f - T_{hs}$). K_{FB} then compensates this error.

In the feedforward loop, the filtered curvature is scaled by driven velocity to obtain a steering angle that is proportional to the required yaw rate in a curve, given in Eq. (3.8).

$$r = \frac{V}{R} = V\kappa \tag{3.8}$$

This means that when the feedforward gain K_{FF} is too high, the steering input δ_s is appropriate for a sharper curve.

Note that Van Paassen does not exhibit a visual time delay. This could be due to it being redundant with preview time, i.e. the effective preview time is the visual time delay subtracted from the preview time. Hence to skip this step, the considered preview time τ_f is the *effective* preview time.

LIMITATIONS

The concept of setting a reference other than centerline for position tracking is unique to this model. Unfortunately, the computation of y_{cc} as given in Eq. (3.7), results in a fixed value that scales with curvature magnitude. Moreover, the term is squared, which means that it will always be positive, implying that the reference of the position tracking will always lead to curve cutting. This structure may prove inflexible when fitting for hypothetical counter curve cutters (i.e. drivers that exhibit a negative y_{cc}).

DEGREES OF FREEDOM

The degrees of freedom in the Van Paassen model are five: the lookahead time τ_f, the filter time constant T_{hs}, the prediction time τ_n, the feedback gain K_{FB} and the feedforward gain K_{FF}. The other parameters shown in the model such as the wheelbase gain K_w, wheelbase distance L_{wb}, the effective steering ratio K_{eff} and K_δ are all constants and properties of the vehicle driven. These properties of the vehicle dynamics are in the Van Paassen driver model part, as he explicitly assumes that drivers invert specific properties of the car.

3.3. ASSESSMENT METHODOLOGY

The assessment methodology introduced in this Section provides a platform for fair comparison between different control-theoretic driver models. The fundament of this assessment methodology is based on Chapter 2. For the aim of detailed comparison, the methodology of Chapter 2 is modified through upgrading the metric *Descriptiveness* and introducing the standardised (and combined) metric *Realistic Identifiability*. The detailed explanation of the additional improvements of the methodology is elaborated in Appendix B.5.

The assessment methodology computes two standardised properties (*Descriptiveness* and *Realistic Identifiability*) of the assessed model for a *defined manoeuvre*. The *Descriptiveness* is defined as the capacity of the model to describe different types of driver (trajectory) behaviour. The *Realistic Identifiability* defines the extent to which the resulting identified model parameters are unique to a given driver behaviour, relative to the parameter space which results in desired driving behaviour. The more unique, the smaller the realistic identifiability value and thus the more meaningful the considered parameters are.

Section 3.3.1 explains what the implications of the assessed model are with different Descriptiveness and Realistic Identifiability outcomes. Section 3.3.2 gives a stepwise guideline of the assessment procedure. Section 3.3.3, explains the realism boundary conditions for desirable modelling outputs, Section 3.3.4 explains the experimental data generation and Section 3.3.5 explains the processing blocks.

3.3.1. IMPLICATIONS OF THE ASSESSMENT METHODOLOGY

The implications of the results of the assessment methodology are illustrated by a four-quadrant diagram in Fig. 3.4. The four-quadrant diagram provides the interpretation of the combined results of descriptiveness and realistic identifiability. In the case of a low descriptiveness and high realistic identifiability (Q4), the model is over-parametrised, i.e. exhibits parameters that do not have a unique function/purpose (meaningless) and is unable to describe a sufficient portion of the tested driving population. This model needs fundamental revision. In the case of a low descriptiveness and low realistic identifiability (Q3), the model exhibits meaningful parameters but is unable to describe a sufficient portion of the driving population. Thereby, this model needs to be structurally extended. In the case of a high descriptiveness and high realistic identifiability (Q2), the model is over-parametrised. This implies parameter redundancy in the model, i.e. the parameters have lost meaning. However, the model is complete (it doesn't need an extension) as it can describe a sufficient portion of the driving population. In the case of a high descriptiveness and low realistic identifiability, the model is both parsimonious and able to sufficiently describe the driver population tested, this is considered an ideal model (Q1).

An ideal model would be able to describe all the categorised types of driving trajectories, thereby achieve a score of 100% *Descriptiveness*. Whereas, an ideal *Realistic Identifiability* is achieved when the percentage ratio between parameter space that describes a specific type of driving behaviour and the total parameter space producing desirable driving behaviour, is lower than 1%.

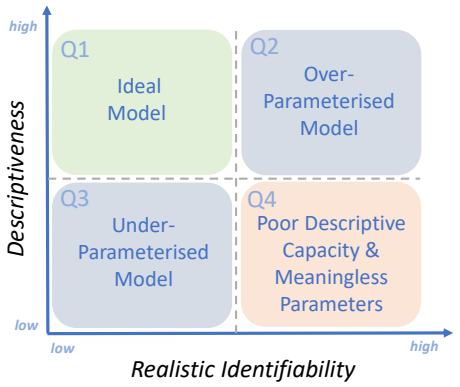

Figure 3.4: This figure illustrates the assessment methodology's (quantitative) implication quadrants. The first quadrant (Q1) implies the model is ideal, i.e., a model with a high descriptiveness and low realistic identifiability. The second quadrant (Q2) implies the model is over-parametrised, i.e., a model that has high descriptiveness and high realistic identifiability. The third quadrant (Q3) implies the model is under-parametrised, i.e., it has low descriptiveness and low realistic identifiability. The fourth quadrant (Q4) implies the model exhibits poor descriptive capacity and meaningless parameters, i.e., it has low descriptiveness and high realistic identifiability.

3.3.2. ASSESSMENT PROCEDURE GUIDELINE

This assessment procedure is developed for control-theoretic steering models which describe curve driving behaviour. However, it can also be applied to (identifiable) models that can describe any driving manoeuvre. The assessment procedure guideline consists of eight steps, which are given (as **S1** ... **S8**) in Fig. 3.5. The guideline steps on how to carry out the assessment procedure are the following:

S1 To start with, the *Scenario Settings* defines the road scenario to be driven and modelled. In the case of (fixed speed) curve driving, this is the road width W_r, the road curvature profile κ and the fixed vehicle velocity V.

S2 In this step: *Experimental Data Generation* (elaborated in Section 3.3.4), the driving simulation is first created using the scenario settings (W_r, κ, V). The experiment is executed by testing a given number of human drivers. The *Experimental Data Generation* produces the trajectories of all the individual drivers $(\vec{s}_{lat,e})$. Note that the vector notation in $(\vec{s}_{lat,e})$ is used when dealing with a matrix of data, as opposed to $s_{lat,e}$ which is a single trajectory.

S3 In this step: the *Realism Boundary Conditions* (RBC) are set (elaborated in Section 3.3.3). One of these conditions is based on the scenario settings, specifically the road width W_r.

S4 In this step: *Classification and Identification* is the first processing block (elaborated in Section 3.3.5), where the experimental trajectories $\vec{s}_{lat,e}$ are first categorised by a trajectory classifier. Then a representative driver from each class is selected. The *Assessed Model* is then used in the identification of the class representative drivers (one representative driver per class).

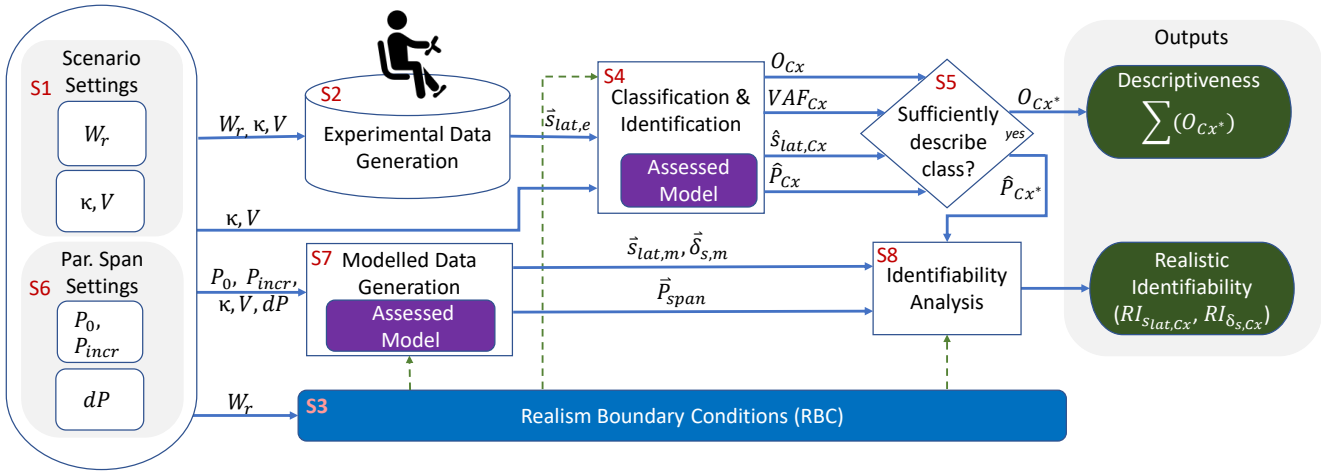

Figure 3.5: This figure illustrates a flow diagram to outline the assessment procedure. There are three *Scenario Settings*, the road width W_r, the road curvature profile κ and (fixed) velocity V. There are three *Parameter Span Settings* which are decided based on knowledge of the model: the initial realistic parameter set P_0, parameter exploration increment P_{incr} and parameter span resolution dP. There is one database, that of experimental driving simulator data explained in Section 3.3.4 and one decision node *Sufficiently describe class?* defined in Section 3.3.5. There are three main processing blocks: *Classification and Identification* (elaborated in both Fig. 3.6 and Section 3.3.5), *Modelled Data Generation* (elaborated in both Fig. 3.7 and Section 3.3.5) and *Identifiability Analysis* (elaborated in both Fig. 3.9 and Section 3.3.5). These processing blocks incorporate the Realism Boundary Conditions.

The inputs ($\kappa, V,$) are used to compute the *Assessed Model* response in the identification procedure, where the RBC is used to constrain the identification optimisation. The *Classification and Identification* block results in four outputs: 1) the occurrence per trajectory class O_{Cx}, 2) the model fit evaluation metric per class VAF_{Cx} (Variance Accounted For, defined in Eq. (3.11)), 3) the modelled trajectory per class $\hat{s}_{lat,Cx}$ and, 4) the parameter estimates per class $\hat{P}_{Cx}$.

S5 In this step, a decision is made: which classes the assessed model can *sufficiently describe*. For curve driving trajectories, the sufficiency criteria is outlined in Section 3.3.5. The resulting *Descriptiveness* is the summation of the class occurrence O_{Cx^*} (consequent from the experimental data in **S3**) which the model can sufficiently describe, resulting in the first metric of this methodology: *Descriptiveness*.

S6 This step determines the *Parameter Span Settings* for **S7**. These are the initial parameter set P_0, a parameter search resolution P_{incr} and a parameter span resolution dP. These values are determined based on knowledge and experience with the assessed model.

S7 This step: *Modelled Data Generation*, the second processing block (elaborated in Section 3.3.5), produces all the modelled trajectories and steering angles considered ($\vec{s}_{lat,m}, \vec{\delta}_{s,m}$) for the *Identifiability Analysis*. First this step explores the maximum possible value of each parameter that would still satisfy the defined RBC, this is done in steps of P_{incr}. Second, the resulting parameter span $\vec{P}_{span}$, with parameter resolution dP, is used to compute the model span $\vec{s}_{lat,m}, \vec{\delta}_{s,m}$ which is the model response for each parameter combination in $\vec{P}_{span}$.

S8 This step: *Identifiability Analysis*, the third processing block (elaborated in Section 3.3.5), determines how sensitive the model response is to changes in model parameter values. This analysis is only performed for the classes the model can *sufficiently describe*. Essentially it computes the comparison of the model output from the parameter estimate $\hat{P}_{Cx^*}$ and the model span $\vec{s}_{lat,m}, \vec{\delta}_{s,m}$. This is performed independently on both the steering angle $\vec{\delta}_{s,m}$ and trajectory $\vec{s}_{lat,m}$ domains. The output from this block is the metric *Realistic Identifiability*, which is the second output of this methodology.

3.3.3. REALISM BOUNDARY CONDITIONS

Realism boundary conditions are defined as the necessary conditions for desired driver modelling outputs. These are:

1 The trajectory must be within the road boundaries, i.e. $|s_{lat}| < W_r$.

2 The response of steering must be smooth, i.e. no oscillation in δ_s.

The first boundary condition directly depends on the road width. The second boundary condition determines the smoothness of the steering response, determined by the number steering oscillations, which is computed by the number of steering reversals [118]. The steering reversals is conventionally computed using human steering angles to indicate the drivers workload. To filter out human noise, the steering angles signal is first filtered and then a reversal is recognised when reaching a steering magnitude greater than 2 deg around a local minima and maxima. In contrast to the conven-

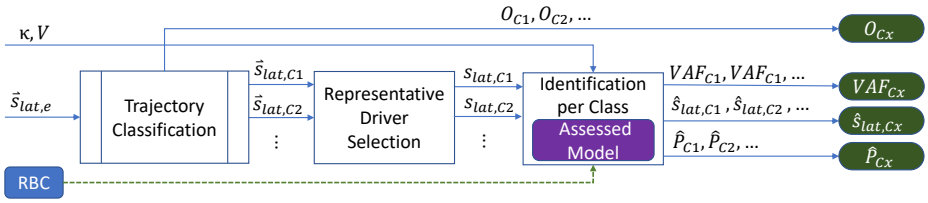

Figure 3.6: This figure illustrates the Classification and Identification processing block (elaborated in Section 3.3.5), which, consists of three sub-processing blocks. The experimental trajectories $\vec{s}_{lat,e}$ are categorised in *Trajectory Classification*. Each class of trajectories $\vec{s}_{lat,Cx}$ undergoes *Representative Driver Selection*, where the resulting selected trajectories $s_{lat,Cx}$ are used for identification in the *Identification per Class* block.

tional method, this methodology deals with the output of driver models which do not exhibit (human) noise. Therefore, the filtering of steering angle is removed because this may obscure high-frequency oscillations resulting from some model settings. Additionally, to detect (and prevent) such small unwanted high frequency oscillations, instead of 2 deg to define an oscillation, a factor 10 less is used: 0.2 deg.

For a strictly smooth steering response, using the above method to compute steering reversals, 2 steering reversals are needed per curve. The proposed methodology allows for four extra steering reversals, setting the limit to 6 steering reversals per curve. This amount of 6 steering reversals per curve is found to be the maximum for human data from Chapter 4.

3.3.4. EXPERIMENTAL DATA GENERATION

The experiment is performed in a driving simulator, testing individual human drivers. The trajectories from the experiment ($\vec{s}_{lat,e}$) are needed for driver trajectory style classification (as outlined in Chapter 4), performed in the *Classification and Identification* processing block. Hence, the generalisability of the class outcomes, increases with an increasing number of subjects (individual drivers) that participate in the experiment. The experimental data is needed to define the types of driving behaviour that the driver model should be able to capture (in *Classification and Identification*), which, is used to evaluate descriptiveness.

3.3.5. PROCESSING BLOCKS

The processing blocks which are dealt with in the following sections are *Classification and Identification*, *Modelled Data Generation* and *Identifiability Analysis*, shown in Fig. 3.5.

CLASSIFICATION AND IDENTIFICATION

Classification and Identification is carried out through three consequent sub-processes as outlined in Fig. 3.6. These sub-processes are the *Trajectory Classification*, *Representative Driver Selection* and *Identification per Class*. The first process is *Trajectory Classification* where experimental driver trajectories are placed into categories using the 11-class classifier defined in Chapter 4. The classifier sorts the individual trajectories

from the experiment ($\vec{s}_{lat,e}$) into groups of trajectories ($\vec{s}_{lat,C1}, \vec{s}_{lat,C2}, ...$), where each of these groups present a class. The groups are ordered based on three decision nodes: 1) where is the driver at curve entry? (outer part of curve, centerline band or inner part of curve), 2) if centerline, to which side do you drift in the curve? (outer, centerline, inner), 3) how many in curve centreline band transitions are made? (0,1,2). This rule-based classifier results in 11 classes which categorises all the trajectories from the driving simulator study.

The class occurrence per class $O_{C1}, O_{C1}, ...$ is a calculated from the percentage of drivers falling into each class. From these classes of trajectories the *Representative Driver Selection* sub-processing block computes the average trajectory per class and then selects the driver that exhibits the closest fit (via standard least-squares computation) to the class average trajectory. The representative trajectories per class ($s_{lat,C1}$, $s_{lat,C2}$,..) are input for the final sub-processing block: *Identification per Class*.

Identification per Class is performed only on trajectories, because this driver modelling signal is the best for discriminating between individual drivers and driving styles, as found in Chapter 2. Fitting on only trajectory can result in oscillating steering (Chapter 2), to avoid this, the identification optimisation is constrained to stay within the *Realism Boundary Conditions (RBC)*. The equation for identification is outlined for (hypothetically) fitting on class x in Eq. (3.10).

$$\hat{P}_{Cx} = \arg\min_{P} \sum_{i=1}^{n} (\hat{s}_{lat}(\kappa_i, V, P) - s_{lat,Cx}(i))^2 \tag{3.9}$$

$$\text{subject to} [\hat{s}_{lat} \quad \hat{\delta}_s] = Model(\kappa, V, P) \quad satisfies \quad RBC \tag{3.10}$$

Hereby, the output of the identification block results in a parameter estimate $\hat{P}_{Cx}$ and corresponding trajectory $\hat{s}_{lat,Cx}$ per class: $\hat{P}_{C1}, \hat{P}_{C2}$... and $\hat{s}_{lat,C1}, \hat{s}_{lat,C2}, ...$ The accuracy of fit is evaluated using the Variance Accounted For (VAF) metric, a common system identification measure [20, 175], described in Eq. (3.11).

$$\text{VAF}_{Cx} = \left(1 - \frac{\sum_{k=1}^{N} |s_{lat,Cx}[k] - \hat{s}_{lat,Cx}[k]|^2}{\sum_{k=1}^{N} s_{lat,Cx}^2[k]}\right) \times 100\% \tag{3.11}$$

For the VAF_{Cx} computation in Eq. (3.11), the identified trajectory responses $\hat{s}_{lat,Cx}$ resulting from the model, are compared to the real trajectory that represents the class average $s_{lat,Cx}$. The VAF represents the normalized sum of errors subtracted from unity, where N is the number of samples. The higher the VAF, the better the model is able to capture the trajectory behaviour. A VAF of 100% means that the model captures 100% of the signal, which is the goal for an *ideal* model.

The four outputs of *Classification and Identification* are the class occurrence per class O_{Cx}, the evaluated VAF for the modelled trajectories per class VAF_{Cx}, the estimated trajectories per class $\hat{s}_{lat,Cx}$ and the parameter estimates $\hat{P}_{Cx}$ per class.

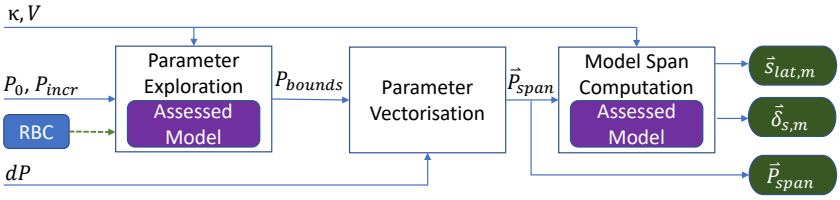

Figure 3.7: This figure illustrates the Modelled Data Generation processing block (elaborated in Section 3.3.5), which, consists of three sub-processing blocks. The bounds of the parameters still satisfying the RBC are explored in *Parameter Exploration*. The discrete set of parameters for the Parameter span P_{span} are vectorised in *Parameter Vectorisation* as defined in Eq. (3.12). The modelling outputs $(\vec{s}_{lat,m}, \vec{\delta}_{s,m})$ corresponding to each parameter combination in P_{span} is determined in *Model Span Computation*.

3

SUFFICIENTLY DESCRIBING CLASS BEHAVIOUR

The decision node in Fig. 3.5 decides whether the assessed model can sufficiently describe the representative drivers from the classified trajectory data $(\vec{s}_{lat,C1}, \vec{s}_{lat,C2}, ...)$. This is determined by satisfying both of the following two criteria:

- $VAF_{Cx} \geq 80\%$.
- Modelled response should classify into the same class.

First, the model fit should be good within the curve, a value of 80% is chosen to be sufficient for trajectories in this chapter. Second, the modelled response should be of the same class as the data it aims to describe.

MODELLED DATA GENERATION

This block is feeded by the inputs κ, V and incorporates the Road Boundary Conditions. *Modelled Data Generation* is carried out through three consequent sub-processes as outlined in Fig. 3.7. These sub-processes are the *Parameter Exploration*, *Parameter Vectorisation* and *Model Span Computation*.

Parameter Exploration is the first sub-process, where the aim of this process is to find the maximum range of each model parameter that still produces 'realistic' model behaviour. For this, both the road properties κ, V and the Realism Boundary Conditions are needed.

The search algorithm used is specifically designed for this method and is described in Algorithm 2. We start at a low parameter value set starting point P_0, which should be realistic. This is determined off-line based on domain knowledge. While the model is within the realistic space, there are two different search types to consider, the principal parameter update, denoted by P where all parameters are updated all at once with the incremental value P_{incr} and, the temporary parameter update $P_t[k]$, which has the same step size value as P_{incr}. The second type of search explores the incremental increase of a single parameter, whilst keeping all other parameters constant. This search continues until reaching out of the realistic space. All the parameters explored are stored in P_s, which is used to find the Parameter bounds P_{bounds}.

Algorithm 2: Parameter Exploration

Require: $P_0,\ P_{incr},\ \kappa,\ V$
  $P = P_0$
  while $[\hat{s}_{lat}\ \ \hat{\delta}_s] = Model(\kappa, V, P)\ \ satisfies\ \ RBC$ **do**
    for $k = 1 : N$ **do**
      $P_t = P$
      while $[\hat{s}_{lat}\ \ \hat{\delta}_s] = Model(\kappa, V, P)\ \ satisfies\ \ RBC$ **do**
        $P_t[k] = P_t[k] + P_{incr}[k]$
        $P_s = [P_s\ \ P_t]$
      end while
    end for
    $P = P + P_{incr}$
  end while
  $P_{bounds} = max(P_s)$

A special halting process (not explicitly shown in the exploration algorithm) is introduced in a temporary parameter update $P_t[k]$ exploration for lead/lag time constants (for e.g. the Mars model [116]). As is known for lead/lag filters, they can serve as two different purposes: 1) determine the break frequency of a given filter, 2) act as effective gain. Once the time constant is outside of the bounds of the excitation frequencies of the road, the lead/lag contribution reduces to an effective gain. Thereby, the effectiveness in changing the response by increasing the time constant parameter may asymptotically stagnate, as illustrated in Fig 3.8. Therefore, to avoid very large time constant parameter bounds, an additional halting criteria is implemented for this type of parameter: when the VAF between consecutive responses of incremented time constant exploration trials is larger than 98%.

Parameter Vectorisation is a simple step, but important for the accuracy of Realistic Identifiability. The formula for vectorisation is given in Eq. (3.12).

$$\vec{P}_{span} = [0, dP, 2dP, 3dP, \ldots, P_{bounds}] \tag{3.12}$$

As dP increases, the accuracy of sensitivity estimation decreases. When dP decreases, the computational load increases. Therefore a trade-off between accuracy and number of responses to compute must be considered. The relationship between number of responses to be computed, number of parameter-values chosen per parameter (N and M for a model with 2 degrees of freedom (DoF)) are that of: number of $[s_{lat}, \delta_s]$ responses = NM. Hereby the computational time scales linearly with the number of $[s_{lat}, \delta_s]$ responses.

It is important to note that this algorithm computes the span P_{span} based on the maximum possible value per parameter P_{bounds} abiding the the RBC, however, this does not mean that all parameter combinations in P_{span} abide to RBC aswell.

Model Span Computation is the last sub-process of *Modelled Data Generation*. This subprocess takes the parameter span $\vec{P}_{span}$ as input and evaluates the model's response to each parameter combination in a grid format. Since, the model outputs

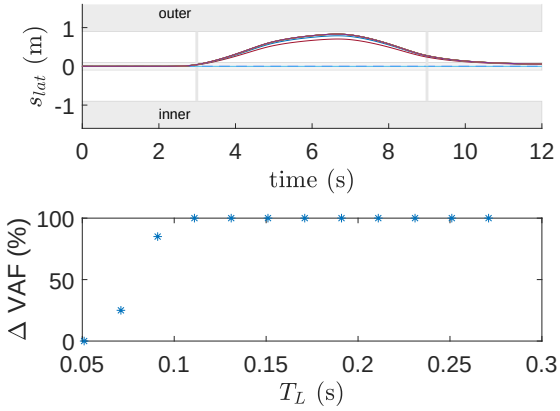

Figure 3.8: Illustration of the halting process for a lag time constant. The trajectory s_{lat} illustrated for a range of different lead time constant T_L values. After increasing T_L above 0.11s the response does not change.

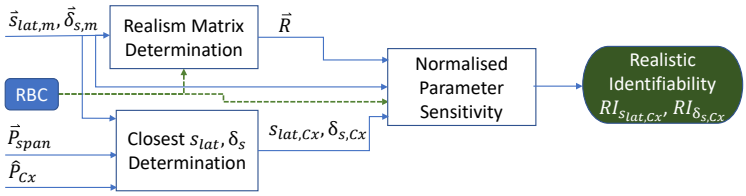

Figure 3.9: This figure illustrates the Identifiability Analysis processing block (elaborated in Section 3.3.5), which, consists of three sub-processes. The *Realism Matrix Determination* evaluates whether each entry in the model span $\vec{s}_{lat,m}, \vec{\delta}_{s,m}$ satisfies RBC. The *Closest* s_{lat}, δ_s *Determination* finds the model entry that best matches $\hat{P}_{Cx}$. The *Normalised Parameter Sensitivity* evaluates the sensitivity space of P_{Cx} within the RBC, normalised to the models' realistic space $\vec{R}$.

not only trajectories s_{lat}, but also steering angles δ_s, this results in the model's span along two domains $\vec{\delta}_s$ and $\vec{s}_{lat}$.

IDENTIFIABILITY ANALYSIS

This block is fed by the inputs the parameter span and selected parameter estimates ($\hat{P}_{Cx^x}, \vec{P}_{span}$) and model span ($\vec{s}_{lat,m}, \vec{\delta}_{s,m}$) and, incorporates the Road Boundary Conditions. *Identifiability Analysis* is carried out through three sub-processes, two of which are executed in parallel, as outlined in Fig. 3.9. These sub-processes are the *Closest* s_{lat}, δ_s *Determination*, *Realism Matrix Determination* and *Normalised Parameter Sensitivity*.

The *Closest* s_{lat}, δ_s *Determination* sub-process takes in the selected parameter estimate $\hat{P}_{Cx^*}$ and finds the index that corresponds to the closest parameter in $\vec{P}_{span}$. This index is then placed in the model span $\vec{s}_{lat,m}, \vec{\delta}_{s,m}$ to find the the corresponding $s_{lat,Cx}$, $\delta_{s,Cx}$. Note that for this step to be valid, the parameter increment dP in the

Parameter Span $\vec{P}_{span}$ should not be too large. When this is the case, the identified P_{Cx^*} and selected response from the model span may be quite different.

The *Realism Matrix Determination* evaluates each grid entry in the model span ($\vec{s}_{lat,m}, \vec{\delta}_{s,m}$) and determines whether a respective parameter combination results in a response that satisfies the RBC. This results in a binary matrix $\vec{R}$ containing ones and zeros, that defines the realistic space of parameter combinations.

Algorithm 3: Normalised Parameter Sensitivity

Require: $\vec{R}$, $s_{lat,Cx}, \delta_{s,Cx}$, $\vec{s}_{lat,m}, \vec{\delta}_{s,m}$
  for $i = 1 : N$ **do**
    for $j = 1 : M$ **do**
      if $R(i,j) == 1$ **then**
        $\text{VAF}_{s_{lat}}(i,j,:) = \text{VAF}(s_{lat,m}(i,j,:), s_{lat,Cx})$
        $\text{VAF}_{\delta_s}(i,j,:) = \text{VAF}(\delta_{s,m}(i,j,:), \delta_{s,Cx})$
      end if
    end for
  end for
  $RI_{s_{lat,Cx}} = \frac{size(find(VAF_{s_{lat}} \geq 95))}{size(find(R==1))}$
  $RI_{\delta_{s,Cx}} = \frac{size(find(VAF_{\delta_s} \geq 95))}{size(find(R==1))}$

The *Normalised Parameter Sensitivity* compares the selected class response ($s_{lat,Cx}$, $\delta_{s,Cx}$) to all the responses in the model span ($\vec{s}_{lat,m}, \vec{\delta}_{s,m}$) through the validation metric VAF. In Algorithm 3 this is demonstrated on a model that has two DoF, with N and M number of parameters in each dimension. The identifiability volume is determined by the parameter space surrounding $\hat{P}_{Cx^*}$ that produces a VAF $\geq$ 95% and satisfies the RBC. Note that a VAF of 95% was chosen heuristically. This DoF-dimensional volume is then normalised by the total realistic volume in $\vec{R}$ to produce the final standardised metric of this methodology, *Realistic Identifiability*. This standardisation is a key change w.r.t. Chapter 2 and makes it possible to compare models of different numbers and types of parameters. However, a model that has more degrees of freedom, may be susceptible to a smaller realistic identifiability volume, since normalising a unit on a 2-D circle gives a larger percentage than a unit in a sphere. In terms of acceptable realistic identifiability values, a model with less than 1% realistic identifiability is ideal, this means that 1% within the realistic space of parameters can sufficiently describe the specified driver behaviour.

Table 3.1: The Table lists the values chosen for P_0, P_{incr} used for the *Parameter Exploration* parameter processing block. The resulting P_{bounds} and chosen dP required to compute P_{span} is also reported.

	Mars		**Van der El**				**Van Paassen**				
	K_p	K_c	K_e	τ_f	T_L	$T_{l,f}$	K_{FF}	K_{FB}	T_{hs}	τ_f	τ_n
P_0	3	15	2	0.3	0.2	0.2	0.1	0.15	0.1	0.1	0.1
P_{incr}	1	3	1	0.1	0.1	0.01	0.05	0.05	0.1	0.05	0.05
P_{bounds}	22.3	77	64	3.4	6.2	0.9	1.2	5.3	27.3	2.8	1.8
dP	0.75	2	1.1	0.2	0.33	0.1	0.08	0.4	2	0.3	0.2

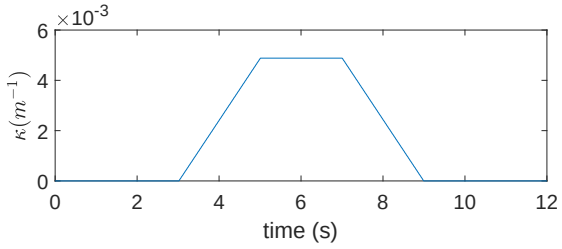

Figure 3.10: This figure illustrates the clothoidal road curvature from curve 3 in Chapter 4. This is curvature profile that stems from the road used for the empirical driver data, and is used for driver modelling input.

3.4. ASSESSMENT RESULTS

Three prominent models, Mars [116], Van der El [173] and Van Paassen [177], were assessed using the method outlined in Section 3.3. This section comprises three parts. The first part describes the applied input parameters for assessment. The second part presents the descriptiveness results and the third part presents the realistic identifiability results. The last two parts give the assessment per model and a comparison of the models.

3.4.1. APPLIED INPUT PARAMETERS FOR ASSESSMENT

This section deals with the specific input parameter which are used in the assessment of the three models. The road width W_r used is 3.6m. The velocity used was fixed at 80km/h. The clothoidal road used is illustrated in Fig. 3.10, taken from Chapter 4. In a single run, there are 5 right and 5 left curves, where the taken trajectories are averaged across the 5 runs. Therefore there are two trajectories, one per curve direction, considered per driver. The dataset considers 45 drivers. Details of P_0, P_{incr} and dP for each model parameter are given in Table 3.1.

3.4.2. DESCRIPTIVENESS

The descriptiveness outlines to what extent a model can capture different types of driver behaviour. The types of driver behaviour tested in this Chapter are defined by an extended trajectory classifier (11-class classifier of Chapter 4). Fig. 3.11 shows the representative driver of each class, as well as the identification fit results per class. Ta-

ble 3.2 summarises the identification fit results per class. The degree of importance to fit a certain class is indicated by the number of drivers classified in a particular class, the class occurrence O_{Cx}, as shown in Fig. 3.11 and Table. 3.2.

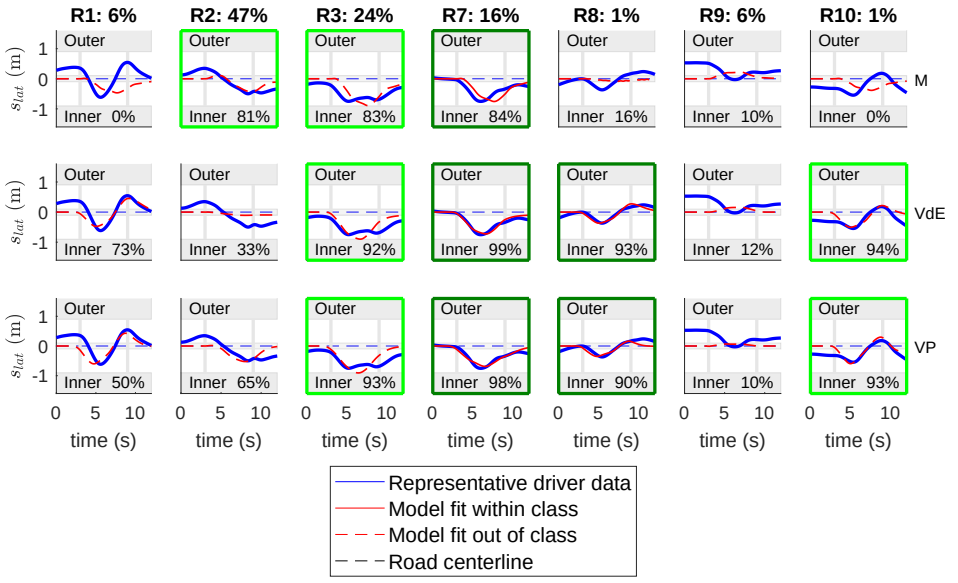

Figure 3.11: This figure shows the result of model fits in the lateral position domain (s_{lat}) for different classes of right curve driving (indicated by R), with the corresponding class occurrence (O_{Cx}) percentage indicated by Rx: O_{Cx}%. The class occurrence is based on data from a fixed based driving simulator experiment from Chapter 4. The first row presents the model fits of Mars (M), the second row presents the model fits of Van der El (VdE) and the third row presents the model fits of Van Paassen (VP). The horizontal grey bars represent the effective road boundary starting at 0.9 and -0.9 m. The two vertical grey lines represent the beginning and end of a curve. The representative driver of each class is illustrated by blue line. The model fits, which result in other classes than the representative driver, are shown with a dashed red line. The model fits resulting in the same class are represented by a (non-dashed) red line. When the fit is above 80 %, the plot has a dark green box frame. When the fit is above 80% and not within the same class, the plot has a light green frame.

Table 3.3 gives an overview of the descriptiveness value and weighted average VAF fit value.

ASSESSMENT

The *Mars* model has a descriptiveness of 16% as shown in Tables 3.2, 3.3 and Fig. 3.11. This model can only capture one class of driver behaviour, class 7, as shown in Table 3.2 and Fig. 3.11. Nevertheless, the weighted average VAF, that is the average VAF weighted by class occurrence (see Table 3.2), is still high for the Mars model at 72%. This is because a high VAF fit value 81% is achieved for class 2, which comprises 47% of drivers in our dataset.

The *Van der El* model has a descriptiveness of 17% shown in Table 3.3, and captures two classes of driver behaviour, class 7 and 8, as illustrated in Table 3.2. However, the

Table 3.2: The identification fit per class for the three models. The columns represent a class number, the resulting percentage occurrence O_{Cx} of each respective class. The identification fits given in VAF per class. When bold the identified response falls in the same class as the observed data used to identify the model parameters.

Class	O_{Cx} (%)	**Mars** (%)	**Van der El** (%)	**Van Paassen** (%)
1	6	0	73	50
2	47	81	33	65
3	24	83	92	93
7	16	**84**	**99**	**98**
8	1	0	**93**	**90**
9	6	10	12	10
10	1	0	94	93

Table 3.3: Overview of descriptiveness score ($\sum O_{Cx^*}$) and weighted average (by the class occurrence O_{Cx}) VAF fit score from Table 3.2, for the models.

Metric	**Mars**	**Van der El**	**Van Paassen**
Descriptiveness (%)	16	17	17
Weighted average VAF (%)	72	64	74

average fit value is 64%; which is mostly because the model shows a bad descriptive capacity to class 2, which comprises 47% of drivers in our dataset.

The *Van Paassen* model has a total descriptiveness of 17%, see Table 3.3, and also captures two classes of driver behaviour, class 7 and 8 at high accuracy. Moreover, the average fit value is 74%, see Table 3.3.

COMPARISON

Table 3.3 shows that the Van Paassen model and Van der El models are both able to capture two classes of curve driving, in comparison to the Mars model that can only capture one class. With the information on class occurrence given in Table 3.2, this means 17% descriptiveness for both Van Paassen and Van der El, whereas 16% for Mars. This indicates that the Van Paassen and Van der El have the same value. However, when taking into account the weighted average fit value, we see that Van Paassen is the most flexible and descriptive model.

All control-theoretic driver models tested in this Chapter can only capture class 7 and 8 and stay within the same class. This is because, before curve entry, these models are not capable of replicating any prepositioning behaviour, a phenomenon that drivers commonly undertake, as deduced from Chapter 5. Hence, only the classes that start at centerline upon curve entry (potentially also classes 4,5 and 6 in another dataset or curve direction) are those which the models can replicate, this accounts for 17% of the drivers in the dataset from Chapter 4. This low descriptiveness percentage is a consequence of these control theoretic models ignoring steering behaviour before curve entry.

For the other classes, we can see that sometimes, the in-curve fit is good, with the Mars model reaching 81% for class 2 and Van Paassen reaching 93% in class 10. This

Table 3.4: The Realistic Identifiability values for the classes 7 & 8: the classes that the model can sufficiently capture. Given for both the steering wheel δ_s and trajectory s_{lat} domains. The corresponding parameters P_{Cx} can be found in Appendix A.

Models	**Realistic Identifiability (%)**			
	Class 7		**Class 8**	
	δ_s	s_{lat}	δ_s	s_{lat}
Mars	88	9	N.A.	N.A.
Van der El	97	3	100	1.7
Van Paassen	6	0.06	5.6	0.05

indicates how flexible the models are in capturing different types of drivers. Nevertheless, some classes are more important to fit than others. For example, class 2 is the most prevalent at 47% in Chapter 4, and is a curve cutting class. This means that it is an essential class of driver behaviour to model. Interestingly, although the Mars model does not fit most classes as well as Van der El and Van Paassen, it fits class 2 better than both of them.

3.4.3. REALISTIC IDENTIFIABILITY

Realistic Identifiability indicates to what extent a specific parameter combination P_{Cx^*} has unique mapping to its modelled driver steering behaviour $[\hat{s}_{lat} \;\; \hat{\delta}_s]$. In this assessment procedure, realistic identifiability is only evaluated for the classes that the model can sufficiently describe. Hence the parameters evaluated are P_{C7}, P_{C8}, for class 7 and 8 respectively. The realistic identifiability results for the N-dimensional parameter spaces of the models' are summarised in Table 3.4. A 2-dimensional identifiability heatmap slice of the N-dimensional parameter space is presented (for illustration purposes) in Fig. 3.12, 3.14 and 3.15 for class 7. The presented identified P_{Cx} is rounded up to the closest resolution point in the defined P_{span}, as defined in the assessment methodology see Fig. 3.9.

ASSESSMENT

The *Mars model* can only sufficiently capture class 7, as given in Table 3.4. Fig 3.12 illustrates this realistic identifiability for the identified parameter set of class 7, P_{C7}. Here a heatmap of VAF fit results are shown for different parameter combinations (defined by the parameter axes) relative to the modelled output (either s_{lat} or δ_s) for P_{Cx}. The blue boundary illustrates the Realistic Boundary Conditions (RBC). These boundary conditions are a consequence of two requirements: 1) staying within the road boundaries W_r and maintaining the number of steering reversals below 6. These criteria are separately illustrated in Fig. 3.13, where the righter part of the boundary is constrained by the steering reversals requirement whereas the left part is constrained by staying within road boundaries. The realistic identifiability percentage is the 95% area (white space) divided by the area within the RBC bounds.

For class 7, 9% of all realistic trajectories, fall within 95% VAF of the model trajectory response (s_{lat}) for P_{C7}. Whereas, 88% of the realistic model steering angle responses fall within 95% VAF of the model steering response (δ_s) to P_{C7}. This means

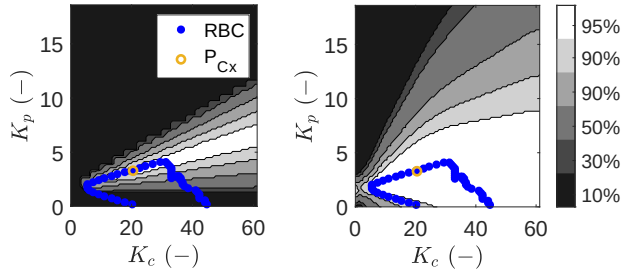

(a) s_{lat} (b) δ_s

Figure 3.12: These plots illustrate a 2-dimensional representation of Realistic Identifiability for class 7 of the Mars model, showing the feedforward K_p and feedback K_c gains. Plot a) presents the trajectory domain s_{lat} and plot b) presents the steering angle δ_s domain.

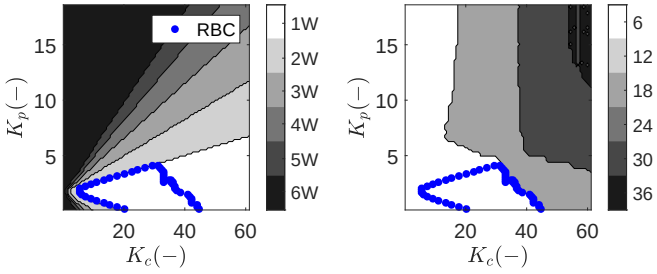

(a) within W_r? (b) Steering Reversals

Figure 3.13: These plots illustrate a 2-dimensional representation of Realistic Parameter space of the Mars model, showing the feedforward K_p and feedback K_c gains. Plot a) presents the first RBC constraint, whether s_{lat} is within the road boundaries W_r, plot b) presents the second RBC constraint, the number of steering reversals.

that the model's δ_s response is not sensitive to changes in parameters, within the realistic space of parameter combinations, whereas the model's s_{lat} response is relatively better. On observing Fig. 3.12a), it can be seen that there is a ratio between the feedforward gain K_p and the feedback gain K_c that produces similar trajectories, this can be observed through the slanted, nearly linear white space on the s_{lat} heatmap. For Fig. 3.12b), the sensitivity space is larger, it overlaps a lot with the realistic space $\vec{R}$ shown with the RBC, confirming the realistic identifiability percentage of 88% given in Table 3.4.

For the *Van der El model*, we observe a higher discrepancy between the identifiability of δ_s and s_{lat} modelling outputs, for both class 7 and 8, as given in Table 3.4. On *average* (across class 7 and 8), the δ_s realistic identifiability is 98.5%, whereas s_{lat} realistic identifiability is 2.35%. Fig. 3.14 illustrates the class 7 realistic identifiability, for only two parameters (degrees of freedom): feedback gain K_e v.s. preview time τ_f, which is a 2 dimensional slice of a 4-dimensional parameter space. These parameters were selected as they can best compare to the anticipatory gain K_p and feedback gain K_c from the Mars Model. The two requirements that define the RBC is illustrated in Ap-

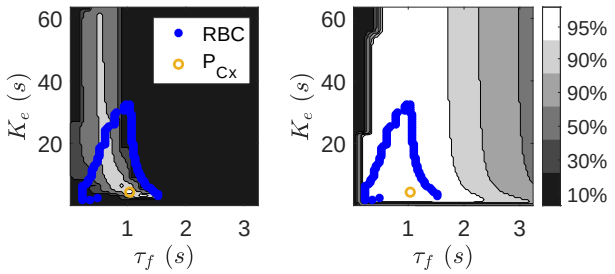

(a) s_{lat} (b) δ_s

Figure 3.14: These plots illustrate a 2-dimensional representation of Realistic Identifiability for class 7 of the Van der El model, showing the error gain K_e and the look-ahead time t_f.Plot a) presents the trajectory domain s_{lat}, plot b) presents the steering angle δ_s domain.

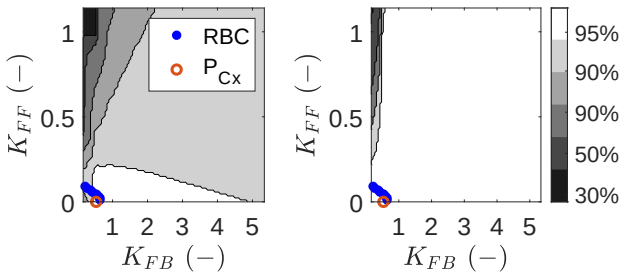

(a) s_{lat} (b) δ_s

Figure 3.15: These plots illustrate a 2-dimensional representation of Realistic Identifiability for class 7 of the Van Paassen model, showing the feedforward gain K_{FF} and the feedback gain K_{FB}. Plot a) presents the trajectory domain s_{lat}, plot b) presents the steering angle δ_s domain.

pendix B.7. Fig. 3.14 illustrates a small s_{lat} identifiability space within the RBC. For a 90% VAF, it seems that parameter K_e can exhibit a large range of values without much change in the response, however, the s_{lat} response is significantly more sensitive to changes in preview time τ_f. For both the s_{lat} and δ_s domains, a smaller range for τ_f than for K_e is seen, this means that changing τ_f can directly result in a wider range of trajectories.

For the *Van Paassen model*, the realistic identifiability in both the s_{lat} and δ_s domains are small. On *average* (across class 7 and 8), δ_s realistic identifiability is 5.8% whereas s_{lat} realistic identifiability is 0.055% as given in Table 3.4. However, when we decrease the Realistic Identifiability threshold from 95% to 90%, the sensitivity volume of the Van Paassen model increases to 92% and 94% for δ_s in class 7 and 8. In fact, the ratio between the realistic identifiability of s_{lat} to δ_s is about a factor 100, for both class 7 and 8. This still means that identifying individual driver behaviour in the s_{lat} domain is more effective. Fig. 3.15, illustrates the class 7 realistic identifiability, for only two parameters (degrees of freedom): feedforward gain K_{FF} v.s. feedback gain K_{FB}, which is a 2 dimensional slice in a 5-dimensional parameter space. These parameters were selected as they can best compare to the anticipatory gain K_p and feedback gain K_c from the Mars Model. The two requirements that define the RBC is illustrated in Appendix B.7. From Fig. 3.15, the δ_s parameter sensitivity is very low, however, that of s_{lat} is better. It is interesting to see that parameter K_{FB} has less sensitivity than K_{FF}. Where the relative indifference of K_{FB} is similar to K_e in the Van der El model, interestingly, both of these gains are providing feedback on an anticipated point ahead. So, increasing the feedback gain does not change the output response much.

COMPARISON

The comparison is given in Table 3.4. The Mars model exhibits the largest realistic identifiability value for s_{lat} at 9%. The Van der El model exhibits better realistic identifiability values at an average of 2.35 %, whereas Van Paassen has the best realistic identifiability value, at an average of 0.055%. Hereby, the Van Paassen realistic identifiability exhibits the most unique mapping between identified model parameters and driver behaviour, which means that it is the most promising to meaningfully describe a specific driver behaviour. It is interesting to observe that the realistic identifiability is not very different for class 7 and 8 for Van Paassen and Van der El. This, however, could have to do with the classes being similar.

3.5. DISCUSSION

The aim of this chapter was to introduce a comprehensive assessment methodology (as a refinement of Chapter 2) and, to use this methodology to compare the quality of three prominent control-theoretic models. The refined assessment methodology comprises two metrics: descriptiveness and realistic identifiability. Descriptiveness determines the capacity of a model to capture different individual driver steering and trajectory behaviour. Realistic identifiability determines the standard degree of uniqueness of mapping between estimated parameter and driver steering and trajectory behaviour. The assessed models are 1) **the Mars model** [114, 145], chosen for its popular use to estimate driver steering in haptic shared control applications

[74, 145, 152, 183], 2) **the Van der El model** [173], chosen as this model is the state-of-the-art of objective (frequency domain) identification of driver perception of the road and execution of the driving task, and, 3) **the Van Paassen model** [177], a model developed within our group based on lessons learnt from previous shortcomings of personalisable haptic shared controllers [31].

Both this chapter and Chapter 2 finds that the realistic identifiability is substantially higher for s_{lat} than it is for δ_s for all models tested. This is because there is a certain amount of steering input δ_s needed to make a curve. In fact, this amount is proportional to the yaw rate required of the vehicle, which is a direct function of curvature ($r = \kappa V$). That is, when increasing the yaw rate, this results in a steering magnitude that is needed for a sharper curve. Small deviations from this amount result in different trajectories; thereby, the difference in steering between trajectories is much smaller than the magnitude of steering required to make a curve. Hence, obtaining a good fit on steering angles may only indicate that you are negotiating the same road, rather than capturing the style of an individual driver. It is, therefore, necessary to include s_{lat} whilst identifying individual driver steering and trajectory behaviour.

The control-theoretic models tested in this assessment are 'triggered' by a feedforward signal. This signal is current road curvature $\kappa(t)$ for the Mars model, previewed road heading $\psi(t+\tau_f)$ for the Van der El model and previewed curvature $\kappa(t+\tau_f)$ for the Van Paassen model. With preview time being around 1s [119, 163, 173], at best, the control-theoretic driver model can exhibit steering behaviour around 1s before curve entry. However, 83% of the drivers tested in this chapter exhibit clear steering behaviour before curve entry, a phenomenon called prepositioning, that on average starts 7s before curve entry, as deduced later in Chapter 5. With ignoring this curve initiation phase (prepositioning), the descriptive capacity of these control-theoretic models is bounded to driver trajectory behaviour that starts around curve entry (class 7 and 8). This imposes a maximum descriptiveness threshold at 17%, which is quite low. This means that at best, control-theoretic models are classified as 'under-parametrised'. Therefore extending current control-theoretic models to account for prepositioning is a *must* to sufficiently describe all the considered trajectory classes of the driving population.

The dimensionality of a model (number of free parameters) may place its realistic identifiability at an advantage. One can understand this advantage through the following example: normalising a unit block w.r.t to a 2-dimensional circle gives a higher fraction value than normalising w.r.t a 3-dimensional sphere. In this Chapter, Van Paassen has a factor 100 better realistic identifiability compared to Van der El, in both the δ_s and s_{lat} domains. The Van Paassen model has five parameters, resulting in 5-dimensional parameter space, Van der El has four parameters, thereby the Van Paassen model has a larger space to normalise in. However, even with four parameters assumed, the Van Paassen model outperforms the Van der El model, see table B.7 in the Appendix.

Interestingly, the Van Paassen model seems to perform best in both descriptiveness and realistic identifiability, even with reduced dimensionality. This could be traced down to the flexible structure of the model; having independent feedforward and feedback loops. Since a certain amount of steering is needed to make a curve, the feed-

forward is used to obtain the required steering (to stay on the road). The feedback loop is used to fine-tune to trajectory position by exhibiting an explicit non-centerline reference for future (τ_f) position tracking. Finally, the Van Paassen model does not exhibit a time-delay. Although time-delays are undeniably a human element of control behaviour, it may not be a necessary element to model human driving behaviour. In fact, when removed, it improves the stability margins of the model, making more parameter combinations realistic. A time delay could even be redundant when using preview time, as it can directly compensate for it.

Identifiability in the realm of driver modelling is the ease of replicating driver behaviour with meaningful parameters. The meaningfulness of the parameters identified can be reflected by looking to the Realistic Identifiability. When the realistic identifiability percentage (on the s_{lat} domain) is low, this means the mapping between behaviour and parameters is unique, which implies meaningfulness. However, the uniqueness of mapping does not necessarily indicate the ease of identification. In fact, with more parameters to optimise for during identification, the dimensionality of the (nonconvex) optimisation increases, resulting in more local minima to avoid which means that it can become harder to obtain the true parameters. Therefore it is important to investigate the relationship between identifiability and realistic identifiability.

In previous studies based on the statistical modelling [10, 49], it is found that increasing the number of polynomial parameters improves the fit of the model to the data, however, decreases the identifiability. There is a parallel to be drawn with the findings of this study; however, there is a sheer difference as well. Statistical models can only add or remove polynomial terms, whereas, in this study, we are also considering different model structure, which adds another dimension to the issue. In fact, in Chapter 2, two models of the same number of parameters were tested; however, due to structural differences, there was an apparent discrepancy in both descriptiveness and identifiability. Moreover, in this study, the Van der El model has twice as many parameters that the Mars model, however, exhibits a similar realistic identifiability percentage, and even exhibits a lower weighted average class fit value. Whereas the Van Paassen model has one more parameter than Van der El, but achieves an improvement in realistic identifiability and improves by 10% for weighted average class fits. Hence, a clear-cut tradeoff in descriptiveness and realistic identifiability v.s. number of parameters is not found, due to a critical dimension: the control structure. This means that with control-theoretic models, an increase or decrease in number of parameters can not be scaled with an increase in descriptiveness or decrease in realistic identifiability, without considering the control structure.

The proposed comprehensive methodology can explicitly indicate the types of behaviour a given model can capture, and to what extent. In the previously proposed assessment methodology in Chapter 2, only the area the modelled trajectories can span along a curve is considered for descriptiveness. However, from driving simulator data, we find that not only may the whole area of the curve be unnecessary to describe, but also that a mere area obscures the important pattern or shape of the trajectory which has to be described. The *realistic identifiability* is a novel standardised metric that is a combination of the previously proposed metrics *identifiability* and *realism*. If we were to use the 'old' metrics, these would produce incomparable values. These

conventional metrics compute a volume that is not independent of parameter type or numbers of degree of freedom in the model. Thereby the previously proposed metric can only be used to compare models that have the same *type* and *number* of free parameters. Whereas in this comprehensive methodology, the identifiability space is normalised by the realism space, resulting in a dimensionless fraction which can be used to compare models regardless of parameter type and number of free parameters.

A model that is both descriptive and exhibits meaningful parameters is potentially instrumental for the concept of meaningful human control in driving automation. One of the two conditions for meaningful human control in automated driving systems is the traceability of system (and driver) behaviour to its source (i.e. through parameters), for the purpose of accountability [148]. Meaningful parameters are essential for not only the understandability of the model but also for physical interpretation and mapping of driver steering and trajectory behaviour. Hereby, when (real-time) identified driver parameters of a descriptive model are meaningful, one can trace accountability.

In literature, there are two ways found to validate through *assessing* driver steering model performance. First through fits on steering angle, and second, through fitting the trajectory of an ISO lane change. Validation through fits on steering angle is found in both the body of control-theoretic driver modelling [119, 145, 163, 173], and with cost-function based driver models[40, 84]. However, fitting on steering angles is found to exhibit poor discriminability between different driver behaviours (identifiability) as found in Chapter 2, which, questions the significance of validating on steering angle alone. Validation by showing the fit of a couple of drivers in specifically the ISO lane change manoeuvre is performed for other (non-control theoretic) driver models [26, 52, 56]. However, the ISO lane change manoeuvre, traditionally driven by an expert driver to test the handling performance of a vehicle, is hardly a representation of normal driving. Therefore, applying the descriptiveness metric can improve the quality of the current validation methods because 1) it is performed on trajectories which is proven to provide discriminability between individual drivers and 2) provides a platform to validate on both representative drivers and manoeuvres, compared to a handling-quality test ISO lane change.

The assessment results presented in this chapter are based on a specific fixed-based driving simulator dataset from Chapter 4. This dataset exhibits a particular profile of curvature of a clothoid road that is 6s long with 10s straight sections in between curves. The road profile taken at a constant speed of 80 km/h is representative of a highway road in the Netherlands. This representative (right-curve) road may give a characteristic indication of the quality performance of a driver model; however, can not represent the assessment results when tested on other (realistic) roads. Therefore, for a practical assessment result, it is recommended to apply this assessment on data from a real-world driver test-track.

3.6. CONCLUSIONS

This chapter assesses and compares the quality of performance of three different control-theoretic driver models for personalised trajectories. These are the Mars model, Van der El model and Van Paassen model. The assessment grades the models using two

metrics: *descriptiveness* and *realistic identifiability*. The descriptiveness metric signifies what per cent of drivers from the driving population (defined in the used dataset) the model can sufficiently capture. For an ideal model, the descriptiveness is 100%. This computation is based on the model's identification on a representative driver from each driving trajectory class, where the trajectory classes categorise different types of drivers within the specific driving population. The realistic identifiability evaluates the degree of uniqueness of mapping between identified parameter set and resulting driver behaviour. Promising realistic identifiability is achieved when the percentage ratio between parameter space that describes a specific class behaviour and the total realistic parameter space, is lower than 1%.

The descriptiveness results are the Mars model 16%, Van der El 17% and Van Paassen 17%. Whereas the corresponding weighted (by class occurrence) average fit values are: the Mars model 72%, van der El 64% and Van Paassen 74%. Out of the seven trajectory classes tested, the Mars model is only able to describe one class (class 7, i.e. normal curve cutting starting at road centerline) sufficiently. In contrast, the Van der El model and Van Paassen models can describe two classes (class 7 & 8 i.e. severe curve cutting and normal curve cutting starting at road centreline) sufficiently. For all models the descriptiveness is low and can not be further improved without accounting for driver steering behaviour before curve entry, i.e. prepositioning. Steering angle realistic identifiability is poor for all models. This indicates that the steering angle is not appropriate for individualised driving trajectory behaviour identification. The average realistic identifiability in the trajectory domain is Mars model 9%, Van der El 2.35 % and Van Paassen 0.055%. The assessment results imply that the Van Paassen model exhibits the best descriptiveness and also the best realistic identifiability. With a high realistic identifiability and low descriptiveness, the best model, Van Paassen is 'under-parametrised'. It is therefore recommended to extend the Van Paassen model to account for prepositioning. This extension is carried out in Chapter 6.

II

Driver Trajectory Classification

4

DRIVER TRAJECTORY CLASSIFICATION

We are moving slowly into an era where data is the starting point, not the end.

Pearl Zhu

When taking a curve, drivers follow their own unique trajectory. Most driver style classifiers in literature are based on inertial inputs (i.e. gas pedal deflection, throttle, braking, acceleration), denoting whether a given driver is aggressive or calm. However, this does not give a direct indication of a drivers trajectory style, i.e. whether a driver typically cuts curves or stays at a constant offset in the lane even during curve negotiation. To fill this void, this chapter, which makes part II of the thesis, introduces two novel rule-based classifiers that categorise seven and eleven different trajectory styles. The classifiers are applied to data from a fixed-base driving simulator study in which 45 subjects drove on three roads, at three different velocities and with three corresponding radii.

The contents of this chapter have been published as:

Barendswaard, S., Pool, D.M., Boer, E.R., Abbink, D.A., 'A Classification Method for Driver Trajectories during Curve-Negotiation', *IEEE Systems Man Cybernetics, Bari, Italy, 2019.*

4.1. INTRODUCTION

Driver-style classification is emerging as a critical factor in driver categorisation to adjust warnings and lane guidance [166], road safety in the context of usage based insurance [143], and even driver-modelling [17]. Driver assessment and profiling is important for power management [46] where calm drivers consume less fuel than aggressive drivers in the same scenarios [28]. Unsurprisingly, there is also a relationship between driving style and road safety [143], with 23% of deaths in traffic being related to 'aggressive' driving styles [181]. This has been a motivation for the development of a number of different ADAS, of which *personalised* implementations explicitly adapt their algorithm to a given driver style [94].

A number of different ways to classify human driver styles have been proposed ranging from classifying groups of people [13] to differentiating between individual people [107] and different genders [179]. Studies on supervised classification techniques have based the definition of their classes on the friction circle [35], a subjective Driver-Style-Questionnaire (DSQ) [179], or a rule-based decision tree on throttle aggressiveness [102]. Whereas unsupervised techniques use feature extraction techniques such as Principle Component Analysis [41] in combination with hierarchical clustering [159].

Most data-mining efforts to classify human driver style have focused on *inertial* behaviour (i.e. gas pedal deflection, throttle, braking, acceleration) [179] [130] [41] [107] [35] [102] [159] [13]. In this approach, generally thresholds are defined to indicate whether drivers show a particular level of agressiveness [13] [35], however, a driver's level of inertial behaviour can change within a single manouvre. Studies find that a single driver exhibits multiple styles for a single manouvre (segmented into multiple time windows), sometimes showing three times more clusters than drivers [76] or that there are 5 segments where a single driver shows different inertial behaviour within a given curve [24]. Moreover, these inertial classifications do not directly distinguish between different trajectories nor can they categorise these. For example defining a curve-cutting style, which is an output that can directly be used with trajectory driven ADAS. Alternatively: trajectory classification is inherently consistent for a single manoeuvre, and can categorise different trajectories.

Trajectory classification has been attempted in [165], by using identified parameters of a driver model and the steering angle as the features for classification. The results have shown poor discriminative properties due to the use of steering wheel angle as an input feature, where steering wheel deflection is proven to be a bad metric to discriminate between drivers' trajectory as found in Chapter 2. Conversely, a naturalistic driver study describes different drivers by indicating a driver trajectory typology in curves [162]. However, this typology is without any numerical quantification.

To bridge this gap, this chapter introduces two newly developed numerical trajectory classifiers for curve driving. They are based on lateral position on a road, before and during a curve, as features for defining rule-based *curve-trajectory classifiers* of driver behaviour. The rules depend on where the car is positioned (w.r.t. the lane centre) before curve entry and how many transitions the driver's trajectory makes across the centerline band during curve negotiation. To show the effectiveness of the proposed trajectory-based classifications, the classifiers are applied to a large dataset col-

lected from 45 subjects, that drove over three different curves in a dedicated driving simulator experiment. The curves were designed such that the experienced lateral acceleration was the same in each curve when following the center line (i.e. sharper curves were traversed at a slower fixed speed to yield the same center line lateral acceleration).

Two classifiers are developed: a 7-class and an 11-class trajectory classifier. The 7-class classifier is stand-alone published work. The 11-class classifier is presented as part of an extension of the 7-class classifier.

This chapter is structured as follows: the details of the human-in-the-loop driving experiment performed for data collection that comprises the dataset used to define both classifiers is outlined in Section 4.2. The classifiers are defined in Section 4.3. The results of applying the classifier on the data is presented in Section 4.4. The discussion and conclusion are in Sections 4.5 and 4.6.

4.2. DATASET

The dataset used to test the proposed classification comes from 45 subjects steering through three different curves, where speed is axiomatically controlled, in a fixed-based driver simulator experiment. The details of the road design and experimental procedures are given in this section.

4.2.1. ROAD DESIGN

A single-lane curve-driving study tested three different road profiles which were tested through three separate drives. The radii and respective velocities tested are listed in Table. 4.1. The three different car velocities are chosen as limit velocities in Dutch traffic rules [140]. The corresponding radii are chosen such that all roads have a maximum centerline lateral acceleration of 2.41 ms^{-2}, which is the maximum lateral acceleration for road design [151].

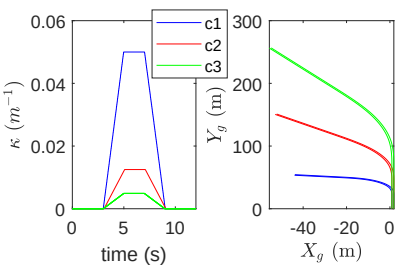

Figure 4.1: The first figure illustrates the curvature profile κ of the tested curves c1, c2 and c3. The second figure illustrates the corresponding birds-eye-view of these curves as a function of global X_g and Y_g coordinates.

The transitions between straight and constant curve sections are designed as clothoids, as illustrated in Fig. 4.1. 10-second straight section intervals were inserted in between curves, and within each curve, the clothoid sections at curve entry and exit lasted 2

Table 4.1: The designed curve radii, car velocities and peak lateral acceleration of the three curves tested.

Condition	Radius	Velocity (km/h)	a_{lat} (ms^{-2})
c1	R= 20	25	2.41
c2	R= 80	50	2.41
c3	R= 204	80	2.41

seconds. The maximum curvature section lasted 2 seconds, making the total time in the curve 6 seconds. Each curve was repeated 10 times, i.e. 5 right curves and 5 left curves, in alternating order.

4.2.2. CONTROL TASK

Subjects performed a curve negotiation (lateral control) driving task in a fixed-base simulator (the Human-Machine Interface Laboratory in TU Delft) at a fixed speed. A heavy sedan of 1.8 m wide was used to simulate the vehicle on a single lane road visually. A vehicle dynamics identical to previous investigations [152], approximated by a bicycle model, was controlled in a simulation environment with apparatus identical to previous experiments [152].

4.2.3. EXPERIMENTAL-SETUP AND PROCEDURE

Before participating in the experiment, participants signed a consent form. The conditions outlined in Table 4.1 were presented in a randomised order to each subject. For each curve, the participants were given a familiarisation run of 160 s before collecting data for each condition.

4.2.4. SUBJECTS AND INSTRUCTIONS

The experiment was performed by 45 subjects between the age of 18 and 31 years (average of 22 years and a standard deviation of 3.1 years). The range of driver experience was between 0 (only experience with driving for the drivers' license exam) and 10 years, with an average of 3.3 years and a standard deviation of 2.7 years. All participants were instructed to drive as they normally would and to hold their hands on the steering wheel at a "ten-to-two" position.

4.3. TRAJECTORY CLASSIFIERS

In this section, both the 7-class classifier and the 11 class classifier are described. Both proposed rule-based classifiers use the lateral position on the road s_{lat} at different locations relative to the curve start on a known curvature profile κ to define the classes. The introduced rules are knowledge-based, i.e., they are based on the author's observation (human cluster detector) of the types of trajectories drivers take when driving a curve in the driving simulator. The possible lateral positions are defined relative to the road centerline at two different instances: at curve entry and during a curve. The details on how trajectories are distinguished differently are explained in the following.

4.3.1. 7-CLASS CLASSIFIER

Fig. 4.2 illustrates the trajectory outcomes that define the 7-class classifier. These trajectory shapes are the result of the rules presented in the decision tree illustrated in Fig. 4.3. The root node looks at curve entry (where the clothoid curve transition starts), where you can either be on the inner or the outer part of the curve. The first decision node tackles the number of transitions along the centerline *band* that can be made during the curve. The second decision node determines whether you have stayed within the centerline band, where this is only of concern if you have made 0 transitions along this band.

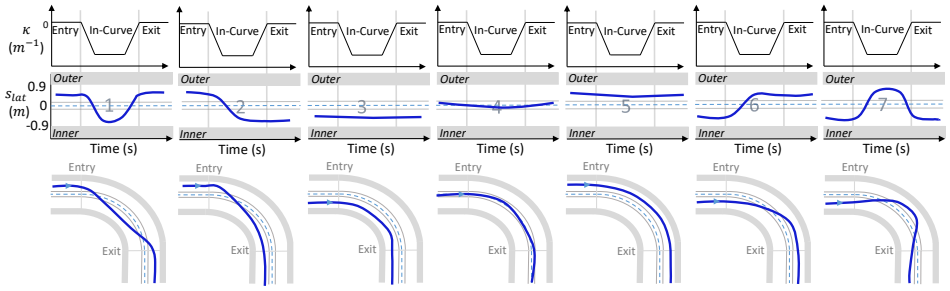

Figure 4.2: This figure illustrates the trajectory classes of the proposed 7-class rule-based classes as distinguished by the classifier of Figure 4.3. The diagram illustrates the trajectory classes for right curves, but this is symmetric for left curves. The first row illustrates the curvature profile, indicating curve entry point and the in-curve phase. The second row illustrates the trajectories in the lateral position domain (s_{lat}). The third row illustrates a corresponding birds-eye-view schematic of the trajectories presented in the s_{lat} domain.

The lateral road position at which drivers are during curve entry is crucial as it could indicate whether they intend to cut the curve [192]. How many transitions are taken along the centerline band defines how the curvature of the trajectory differs from that of the road centerline [30]. The choice of having transitions across a band of ± 0.1 m rather than the centerline line is made here such that only drivers who significantly cross this region are distinguished. The value of ± 0.1 m is chosen as it was found to be the average standard deviation of *straight lane driving* in previous studies [152], i.e. it is considered a 2 σ uncertainty band. This, however, needs to be established for real-world driving.

4.3.2. 11-CLASS CLASSIFIER

MOTIVATION

The 7-class classifier is developed based on empirical findings of driver behaviour. As a result of this, the root node categorises two types of drivers, those who start at the inner or outer part of the curve at curve entry. This binary decision does not consider a trajectory that starts precisely at centreline, which is the case for most conventional control-theoretic driver models. Therefore, the extension of the 11-class classifier is explicitly made for classifying the trajectories resulting from conventional control-theoretic models, as is done in Chapter 3. Additionally, there are a fraction of drivers

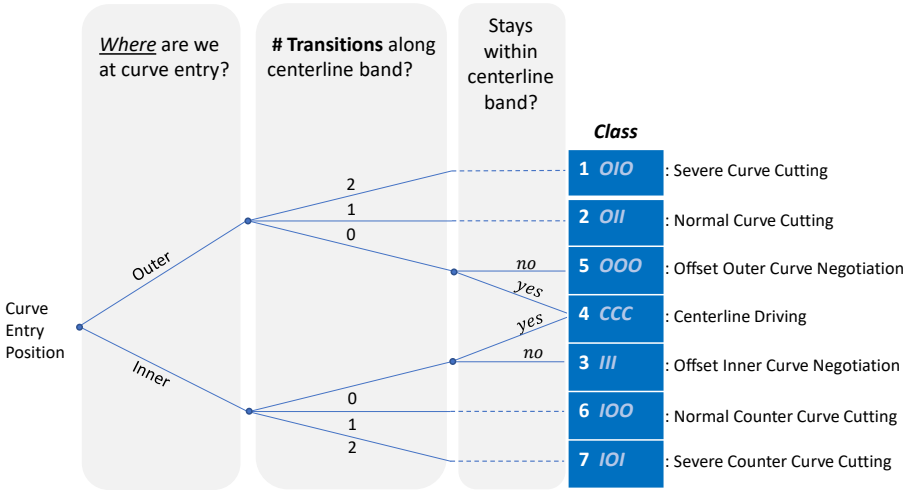

Figure 4.3: This figure illustrates the decision tree of the 7-class rule-based classifier, which is deduced from knowledge of driver trajectories. The centerline band is taken at ±0.1 m from the centerline. The value of ± 0.1 m is chosen as it was found to be the average standard deviation of straight lane driving in Chapter 7. The class numbers are complemented with a code XXX, where X ∈ O, I, C, representing outer, inner, centerline.

who similarly do not preposition and thereby also start at centreline. To precisely classify these drivers, the 11-class classifier is needed.

CLASSIFIER

Fig. 4.4 illustrates the trajectory outcomes that define the 11-class classifier. These trajectory shapes are the result of the rules presented in the decision tree illustrated in Fig. 4.5. The root node looks at curve entry (where the clothoid curve transition starts), where you can either be on the outer, centerline-band or inner part of the curve. The first decision node is exclusive to the drivers that start at centerline. It distinguishes whether the trajectory deviates out of the centerline band once in the curve, and if so, on which side of the curve. If the trajectory stays in the centerline band, the trajectory is directly classified as class 6 (centerline driving, CCC). The next decision node determines how many *complete* transitions across the centerline band are made. Compared to the 7-class classifier, classes 4 (COI) to 8 (CIO) now separately distinguishes trajectories that enter the curve at centerline band.

4.4. RESULTS

The outcomes of both classifiers on a dataset of 45 drivers on 3 different curves are described, the effect of curve direction is analysed and driver consistency is indicated for the 7-class classifier in Section 4.4.1 and for the 11-class classifier in Section 4.4.2.

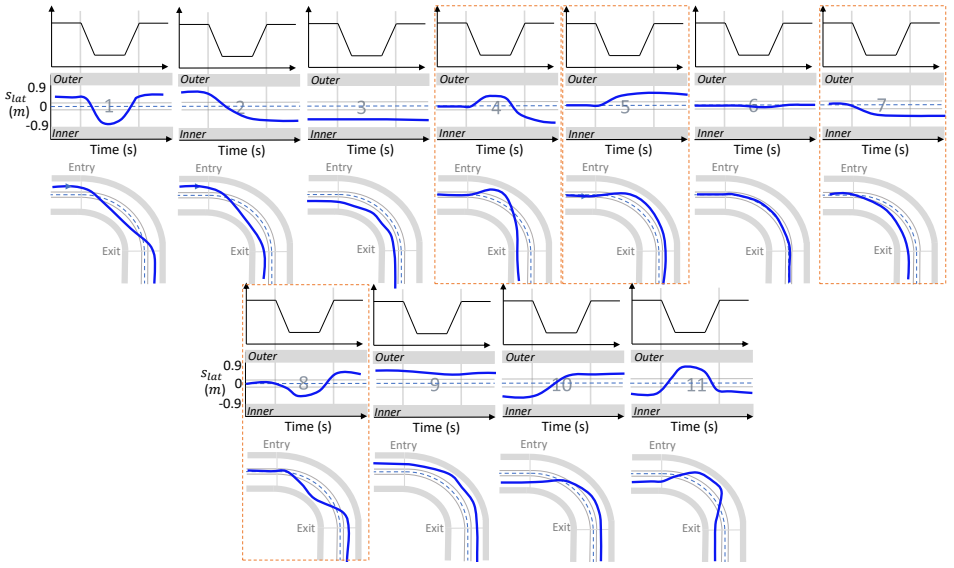

Figure 4.4: This figure illustrates the trajectory classes of the proposed 11 rule-based classes as distinguished by the classifier of Figure 4.5. The new classes which start at centerline are indicated with a red dashed box around them. The diagram illustrates the trajectory classes for right curves, but this is symmetric for left curves. The first row illustrates the curvature profile, indicating curve entry point and the in-curve phase. The second row illustrates the trajectories in the lateral position domain (s_{lat}). The third row illustrates a corresponding birds-eye-view schematic of the trajectories presented in the s_{lat} domain.

4.4.1. 7-CLASS CLASSIFIER

CLASS OUTCOMES

Fig. 4.6 illustrates the class outcomes of the 7-class classifier. This figure shows the lateral position s_{lat}, starting 3 seconds before curve entry and ending 3 seconds after curve exit, for all seven classes. A vertical grey line indicates curve entry and exit, i.e., indicating entry phase to curve phase and curve phase to exit phase. The centerline is indicated at position $s_{lat} = 0$ and the effective road boundaries are illustrated by the shaded area, also indicating the inner and outer side of the curve.

This graph shows that within the dataset collected, a wide variety of driver trajectory styles exists even in a relatively homogeneous test group. Nevertheless, not all classes are found to occur. For example, no instances of the centerline driving (4, CCC) class are found for both right and left curves. Moreover, some classes are more prevalent than others, as is shown with the percentage occurrence in Fig. 4.7. For left curves, class 5 (offset outer curve negotiation, OOO) occurrence is 60% for curve 1 (c1), 48% for curve 2 (c2) and 44% for curve 3 (c3). On the other hand, class 7 (severe counter curve cutting, IOI) occurs only in 7% of the data for c3. For right curves class 2 (normal curve cutting, OII) and 3 (offset inner curve negotiation, III) are the most frequently occurring classes, contributing up to 96% together for c3. Some classes seem to correlate with curvature; for example, class 1 (severe curve cutting, OIO) occurs more frequently for tighter curves (higher curvature). In contrast, class 2 (normal

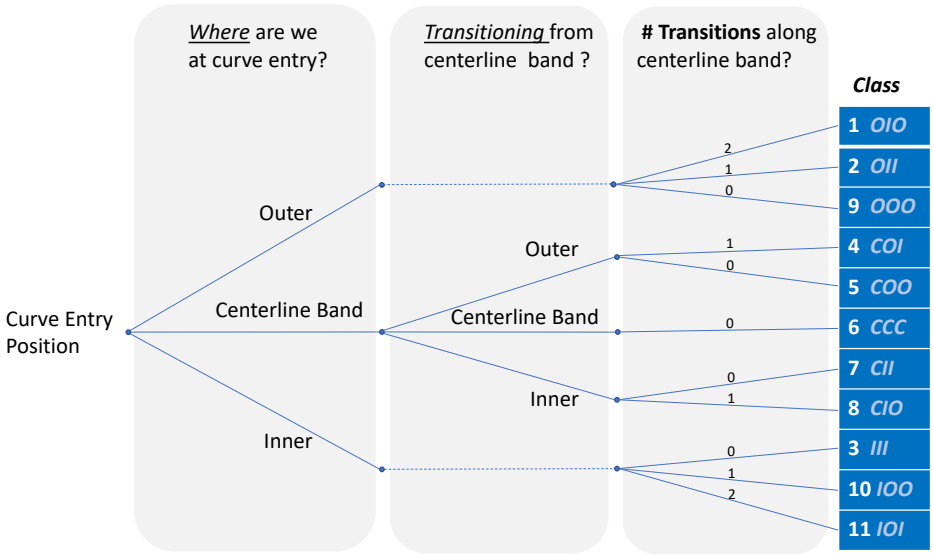

Figure 4.5: This figure illustrates the decision tree of the 11-class rule-based classifier, which is deduced from knowledge of empirical driver trajectories and control-theoretic driver model trajectories. The centerline band is taken at ±0.1m from centerline. The class numbers are complemented with a code XXX, where X ∈ O, I, C, representing outer, inner, centerline.

curve cutting, OII) occurs more frequently for smaller curvature. Seeming that some class 2 (normal curve cutting, OII) drivers switch to class 1 (severe curve cutting, OIO) for increasing curvature.

EFFECT OF CURVE DIRECTION

One can see a clear difference in trajectory class occurrence between right and left curves in Fig 4.7. Specifically, for right curves, class 2 (normal curve cutting, OII) and class 3 (offset inner curve negotiation, III) are most prevalent (86% occurrence together on average). Whereas for left curves, class 5 (offset outer curve negotiation, OOO) and 2 (normal curve cutting, OII) are most prevalent (69% occurrence together on average). Seeming that class 3 (III) and 5 (OOO) switch places for right and left curves. On observing these two classes, one can see that both trajectories have a negative s_{lat} compared to centerline as seen in Fig. 4.6 (which is on the right side of the road). This prevalence of a driver driving on the right side of the road stems from the driver's straight road bias which was found to be -0.1 m on average [152]. Drivers' lateral position is biased to the right side of the road, which makes the inner part of the right curve and the outer part of the left curve. What does this mean? The trajectories on left curves are often classified as outer curve negotiation (class 5, OOO), whereas for right curves as offset inner curve negotiation (class 3, III).

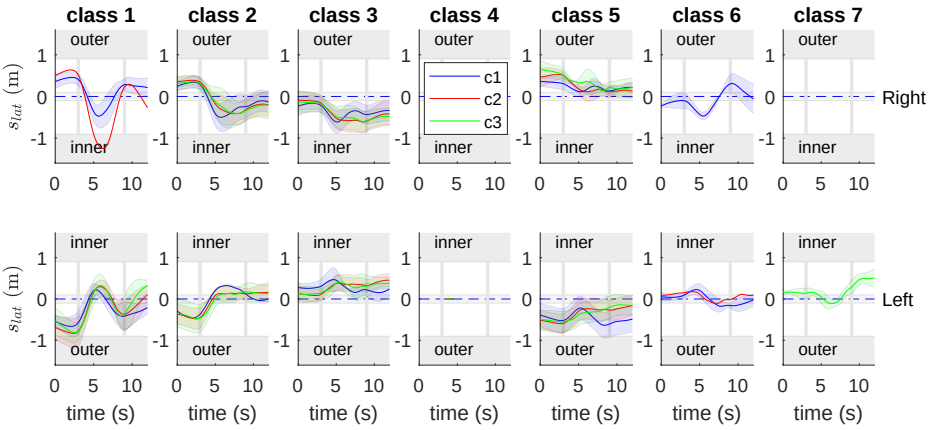

Figure 4.6: This figure illustrates the average and standard deviation of the outcomes of the classification for each class in the lateral position (s_{lat}) domain, respectively. The figure illustrates 3 s before curve entry, a 6 seconds curve and 3 s after curve exit. The curve entry and exit points are indicated with vertical grey lines. The outcome of the classifications is shown for both right and left curves respectively.

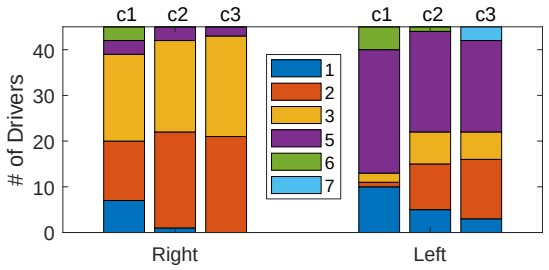

Figure 4.7: This figure illustrates a stacked bar chart which indicates the occurrence of each class (i.e., number of drivers in each class) for the 7-class classifier across each curve, i.e., curve 1 (c1), curve 2 (c2) and curve 3 (c3).

DRIVER CONSISTENCY ACROSS CURVES

The consistency of drivers indicates how generalisable the classification results are when applied to different curves. Fig. 4.8 illustrates pie charts that give the number of drivers who are *consistent* in their trajectory class across 3 curves, *semi-consistent* (consistent across 2 curves) and *non-consistent* (varying classes across all three curves). One can see that drivers are more consistent in their trajectory class for right curves, with 62% being consistent, 31% being semi-consistent and 7% inconsistent. For left curves 44% are consistent, 40% semi-consistent and 16% inconsistent. Where the most inconsistent curve negotiation styles are found for the sharpest curve (c1), as can be seen in the class occurrence distribution in Fig. 4.7. Inconsistent drivers could be a

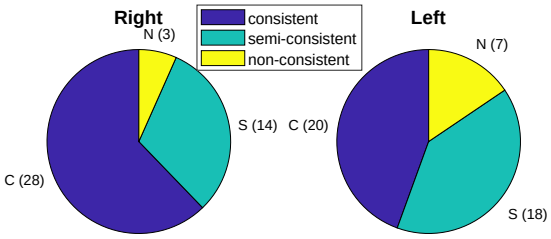

Figure 4.8: This figure illustrates a pie chart indicating the number of drivers who are consistent, semi consistent or inconsistent in their driver trajectory style. *Consistent*: all three curves are driven with the same trajectory class, *semi-consistent*: two curves are driven with the same trajectory class, and *non-consistent*: all three curves are driven in different trajectory classes.

result of a driver struggling to be consistent over the 10 driven curves per curve. Nevertheless, in general, the observed difference from c1 could be a result of having a curvature much larger than c2 and c3, as shown in Fig. 4.1. Where c2 has a curvature 2.5 times larger than c3, c1 has a curvature that is 10 times larger than c3. With c1 being the most demanding curve, demanding the largest steering inputs, a difference in skill due to the elevated demand may be a factor that influences this inconsistency. Moreover, a more demanding curve may instigate the driver to reduce their workload by cutting the curve. Therefore, the classification results can be more reliably generalisable between curves with similar curvature, or steering demands. Hence in a real-world application such as a trajectory-driven ADAS, reclassification of driver trajectories on curves with similar curvature may not be necessary.

4.4.2. 11-CLASS CLASSIFIER

CLASS OUTCOME

Fig. 4.9 illustrates the classification outcome using the 11-class classifier. This graph shows that within the dataset collected, a wide variety in driver trajectory styles exists with the 11-class classifier. Nevertheless, not all classes are found to occur, no instances of the centerline driving (6, CCC) class are found for both right and left curves. Moreover, some classes are more prevalent than others, as is shown with the percentage occurrence in Fig. 4.10.

For left curves, class 9 (OOO), which is equivalent to offset outer curve negotiation, is the most commonly observed with an occurrence of 53 % for c1, 48% for c2 and 44% for c1. This class is very similar, but not equal to class 5 (OOO) in the 7-class classifier. The new class is not equal to class 5 in the 7-class classifier because the 7-class classifier includes trajectories which start at centreline band. The new classes

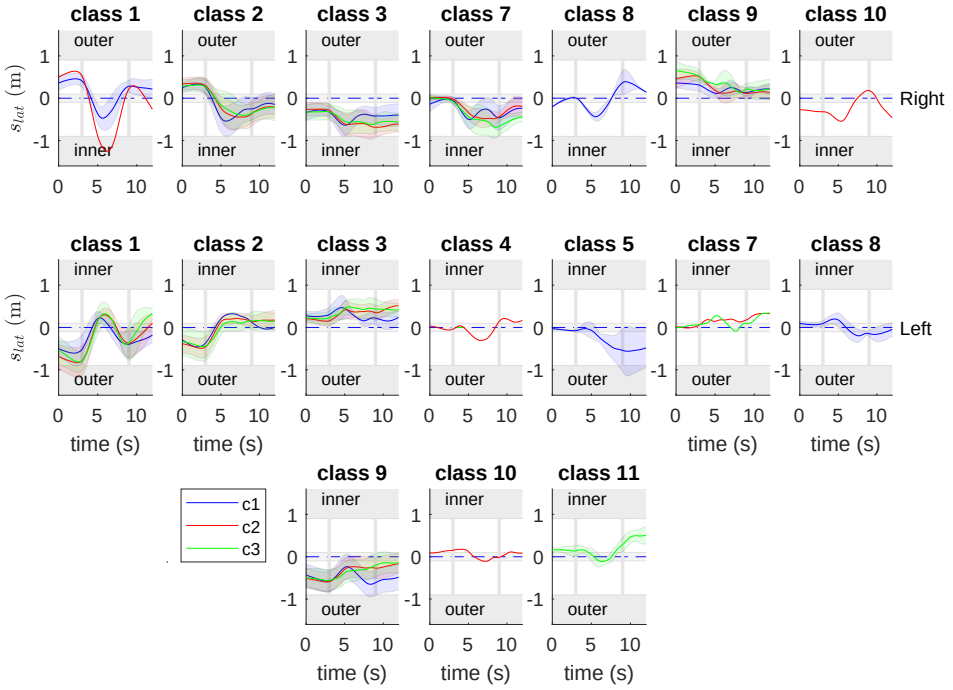

Figure 4.9: This figure illustrates the average and standard deviation of the outcomes of the classification for each class of the 11-class classifier in the lateral position (s_{lat}) domain, respectively. The classes which are empty, such as class 4, 5, 6 and 11 for right curves, are not shown. The figure illustrates 3 s before curve entry, a 6 seconds curve and 3 s after curve exit. The curve entry and exit points are indicated with vertical grey lines. The outcome of the classifications is shown for each curve respectively.

that start at centreline (classes 4 to 8) occur 15.6% for c1, 4% for c2 and 2% for c1. The curve cutting classes (classes 1, OIO and 2, OII) occur 26% for c1, 33% for c2 and 37% for c3.

For the right curves, class 2 (normal curve cutting, OII) is the most predominant. It is also very similar, but not equal, to class 2 (OII) in the 7-class classifier. The new class is not equal to class 2 in the 7-class classifier because the 7-class classifier includes trajectories which start at centreline band. The occurrence for class 2 (OII) is 29% for c1, 46% for c2 and 48% for c3. Combined with class 1 (severe curve-cutting, OIO), the occurrence of curve-cutting is 44% for c1, 49% for c2 and 48% for c3. In the 7-class classifier, the most predominant class for right curves was class 3 (III), however, in the 11-class classifier, class 3 of the 7-class classifier is shared between class 3 (III) and class 7 (CII) of the 11-class classifier. These two classes are similar. However, class 7 distinguishes the trajectories that start at lane centre. The occurrence for class 3 (III) is 26% for c1, 29% for c2 and 32% for c3. The occurrence for class 7 (CII) is 15% for c1, 15% for c2 and 15% for c3. The cumulative occurrence for the classes that start at the centerline (classes 4 to 8) is 20% for c1, 15% for c2 and 15% for c3.

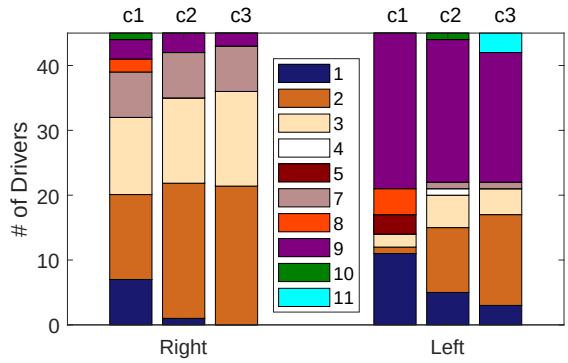

Figure 4.10: This figure illustrates a stacked bar chart which indicates the occurrence of each class (i.e., number of drivers in each class) for the 11-class classifier across each curve, i.e., curve 1 (c1), curve 2 (c2) and curve 3 (c3).

In general, some class occurrences also seems to relate with curvature type for the 11-class classifier. Class 1 (severe curve cutting, OIO) occurs more frequently for larger curvature, whereas class 2 (OII) occurs more frequently for smaller curvature. Class 11 (IOI), on the other hand (equivalent to class 7 in the 7-class classifier) only occurs for the smallest curvature curve (c3) in left curves.

DRIVER CONSISTENCY ACROSS CURVES

Fig. 4.11 illustrates pie charts that give the number of drivers who are *consistent* in their trajectory class across 3 curves, *semi-consistent* (consistent across 2 curves) and *non-consistent* (not consistent across all three curves).

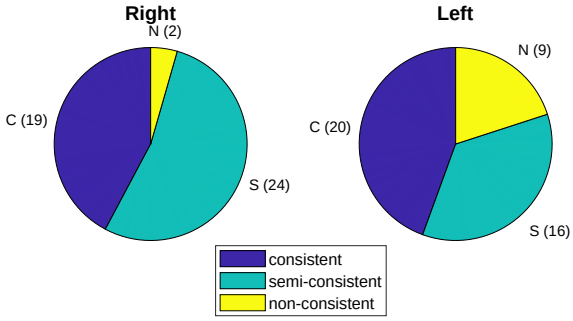

Figure 4.11: This figure illustrates a Pie chart indicating the number of drivers who are consistent, semi consistent or inconsistent in their driver trajectory style for the 11-class classifier. *Consistent*: all three curves are driven with the same trajectory class, *semi-consistent*: two curves are driven with the same trajectory class, and *non-consistent*: all three curves are driven in different trajectory classes.

Compared to the 7-class classifier, there is a drop in consistency for right curves. Previously this was 62%, now 42%. This drop is expected due to making trajectory

distinction more refined. The drivers who lowered the consistency scores, moved to the semi-consistent category, increasing this from 31% to 53%. Interestingly, the non-consistent group decreased from 7% to 4%. This can be traced down to a driver who starts in the centerline-band, where previously, a hard-line was drawn if you started your trajectory on the positive or negative axis. Whereas with the 11-class classifier, as long as you start within centreline band, whether positive or negative, it is considered the same trajectory.

For left curves, the scores for being consistent have not changed substantially. The semi-consistent category has decreased from 40% to 35%, increasing the non-consistent group from 15% to 20%. These results show that still, in line with the 7-class classifier, drivers are most consistent in their trajectory in right curves.

4.5. DISCUSSION

The aim of this chapter was to develop a trajectory classifier that is able to categorise driver trajectories in curves mathematically. Two rule-based classifiers are developed: a 7-class classifier to categorise empirical trajectories and an 11-class classifier to additionally categorise trajectories generated by conventional control-theoretic models (i.e., the trajectory starts at centreline at curve onset) [116] [174] [11]. We find that both developed classifiers can successfully categorise all the considered empirical trajectories resulting from the dedicated driving simulator experiment, with the additional 'centreline-start' classes in the 11-class classifier amounting to 7% of driver on left curves and 17% of drivers on right curves.

The fact that class 4 (CCC) in the 7-class classifier and class 6 (CCC) in the 11 class classifier (centerline driving) is not found to occur, is interesting given that the centreline trajectory is the industry's standard for autonomous path planning and driver assistance guidance [98] [55]. Moreover, if we were to classify trajectories coming from a conventional driver model, with a zero curve entry bias, it would not classify using the 7 class classifier because of the defined root node. The knowledge based on empirical-drivers did not result in a root node including centerline in the 7 class classifier; rather, the lateral position should either be positive or negative. This shows a mismatch between empirical and modelled driver trajectories. Also, the philosophy of control-theoretic driver models always aiming to reduce lateral deviations from the centerline seems to have no ground in driving data, suggesting that designing the control reference to be a 'driver trajectory' makes more sense [177], for efficient modelling.

In terms of trajectory shape, it can be argued that classes 2 (normal curve cutting, OII) and 3 (offset inner curve negotiation, III) are not very different from each other, especially in right curves. What seems to distinguish between them is a consistent lateral displacement. In fact, a majority of trajectories (except class 7 or 11 (IOI)) exhibit a similar trajectory shape, i.e., after curve entry point, drivers drive towards the inner part of the curve, even if they don't reach it. Class 1 (OIO) contributes to the largest 'travel' towards the inner part of the curve, which is also reflected in the corresponding average lateral acceleration. For the 7-class classifier in Fig. 4.12, class 1 (severe curve cutting, OIO) and 7 (severe counter curve cutting, IOI) are clearly distinguished (medians 0.5 ms^{-2} apart). Whereas classes 2 (normal curve cutting, OII), 3 (offset inner curve negotiation, III), 5 (offset outer curve negotiation, OOO) and 6 (normal counter

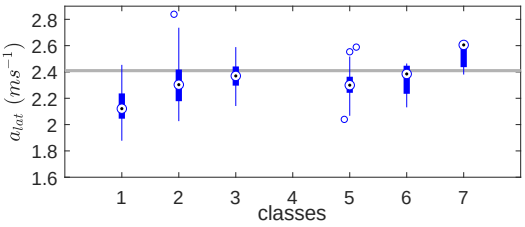

Figure 4.12: This figure illustrates the differences in average a_{lat} in the curve, given for all curves driven for the 7 class classifier. A line at 2.41 ms^{-2} is outlined to represent a_{lat} when the centerline of the curve would be followed.

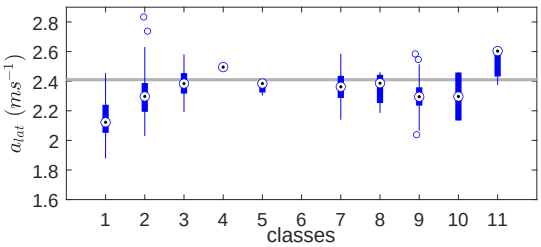

Figure 4.13: This figure illustrates the differences in the average a_{lat} in the curve, given for all curves driven for the 11 class classifier. A line at 2.41 ms^{-2} is outlined to represent a_{lat} when the centerline of the curve would be followed.

curve cutting, IOO) have medians with only 0.08 ms^{-2} apart. The lateral acceleration distribution across 11-class classifier is illustrated in Fig. 4.13. Here class 1 (OIO) and 11 (IOI) is distinguished with medians that are also 0.5 ms^{-2} apart, whereas the other classes fall within 0.1 ms^{-2} from each other. This can be explained by the fact that a_{lat} does not vary with any consistent bias on a trajectory because the change in driven radius is small. Moreover, the difference between class 1 (OIO) and 7 (IOI), or 1 (OIO) and 11 (IOI) stems from drivers taking larger-than-centerline radii trajectories for class 1 (OIO) and smaller-than-centerline radii trajectories for class 7/11. However, when curves become longer, the maximum and minimum radii achievable tend towards that of the centerline value, meaning that the longer the curve, the difference in lateral acceleration achievable between class 1 (OIO) and 7 (IOI) will tend to zero [30] (also explained in Chapter 5). This means that trajectory-based classification provides better discriminative abilities for this dataset, whereas a classifier based on lateral acceleration could obscure such refined differences in driver trajectory, as acceleration is not only 'blind' to consistent bias but also becomes increasingly indifferent between trajectories in longer curves.

Drivers are more consistent in their style across curve 2 (c2) and curve 3 (c3) than they are compared to curve 1 (c1). This difference stems from the greater steering

demand required to make c1, where more drivers curve cut, as an attempt to minimise this steering demand. Moreover, c1 is the curve with the shortest curve length at a velocity of 20 km/h. With a short length, the maximum achievable radius increases, i.e. the added benefit of curve-cutting in terms of decreasing steering demand and improving TLC increases [30]. Hereby, trajectory style is not a fixed property of the driver; it can be influenced by curve type. More specifically, demanding curves like c1 instigate more drivers to curve-cut.

With two classifiers introduced in this chapter, one may expect a trade-off between them. However, there is no one 'perfect' classifier. The 7-class classifier is useful when classifying empirical data, whereas the 11-class classifier is explicitly needed when including driver modelled trajectories in the classification. However, the 11 class classifier comes at the cost of degraded consistency, which is expected due to the refinement of classes, i.e. the more detailed the definitions of the class, the less variability allowed within a class.

4.6. CONCLUSION

This chapter introduces two novel rule-based classifiers that categorise 7 and 11 different trajectory styles. The classifiers are applied to a dataset of 45 drivers negotiating three different curves, each with different curvature and velocity.

The classification results show that normal curve cutting and biased inner curve negotiation are the most prevalent classes for right curves, making up for 86% of class occurrence for the 7-class classifier and 77% for the 11-class classifier on average. Offset outer curve negotiation and normal curve cutting are the most prevalent classes for left curves, making up 69% of class occurrence for the 7-class classifier and 65% for the 11-class classifier on average. The 11-class classifier was introduced to distinguish the drivers that start at centreline, these make up 7% of left curves on average and 17% of right curves on average.

For both classifiers, drivers are not always consistent in their style (on average 52% in right curves and 38% in left curves), where most inconsistency stems from a curve that requires significantly greater steering demand and has the shortest curve length. Here we see that drivers adapt their style to curve type.

Out of the two introduced classifiers, one does not outperform the other. Rather, one can use the 7-class classifier for empirical data classification, whereas the 11-class classifier is necessary when classifying trajectories generated by conventional control-theoretic driver model. Additionally, both introduced trajectory classifiers can directly be used in any personalised ADAS (as is done in Chapter 8), primarily trajectory-driven ADAS, thereby providing an alternative to conventional driver modelling.

III

DRIVER PREPOSITIONING

5

DRIVER PREPOSITIONING QUANTIFICATION

Part III deals with bridging the gap between straight road and curve-driving through describing and modelling prepositioning. This part consists of Chapter 5 & 6. This chapter will only deal with geometrically quantifying prepositioning behaviour and investigating the effect of changes in velocity and curve radius of driver prepositioning. The findings of this chapter are used in Chapter 6.

The contents of this chapter have been published as:

Barendswaard, S. and van Breugel, L. and Schelfaut, B. and Sluijter, J. and Zuiker,L. and Pool, D.M. and Boer, E.R. and Abbink, D.A., 'Effect of Velocity and Curve Radius on Driver Steering Behaviour before Curve Entry', *IEEE Systems Man Cybernetics conference, Bari, Italy, 2019.*

5.1. INTRODUCTION

In naturalistic driving studies curve driving behaviour is generally split into four distinct phases [38, 117] as illustrated in Fig. 5.1 1) *estimation*, where drivers estimate the curvature and adjust their speed for curve entry, 2) *anticipation*, where drivers anticipate the curvature and (pre)position accordingly, 3) *entry and curve negotiation*, where drivers adjust their steering and lateral acceleration to stay within road boundaries, and finally, 4) *exit*, accelerating to the desired speed again. Two of these phases, the *estimation* and *anticipation* phases, occur before curve entry, wherein the *anticipation* phase, drivers show the first steering inputs and prepositioning behaviour. However, our current understanding of curve driving behaviour, reflected in driver steering models, does not give enough consideration to the behaviour before curve entry. Some driver steering models ignore the occurrence of prepositioning by assuming that drivers enter the curve at lane centre [116] [11], other steering models assume they start to turn into the curve a few moments before curve entry at a certain preview time [174][177], or always preposition *optimally* when the tangent point is detected [30]. In fact, preview time is always much shorter than the time before a curve when we see steering changes, i.e., prepositioning phase.

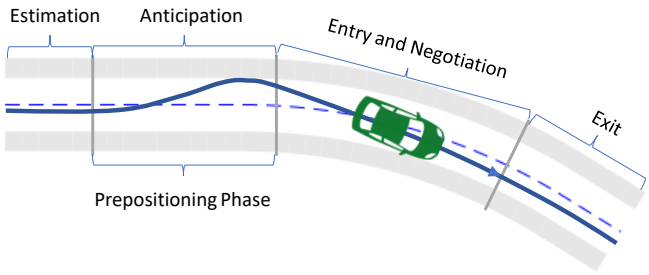

Figure 5.1: Illustration of the four phases of curve negotiation [117] [38], where this chapter focuses on the anticipation or the *prepositioning* phase

The prepositioning phase is found to play a crucial role in defining the initial conditions for the in-curve trajectory [162]. This phase determines whether the driver has the intention of minimising their trajectory's curvature, yaw rate and lateral acceleration to cut the curve for increased comfort [192]. This is done, ideally, by following the 'optimal race-line', which is harder to drive for longer curves [30]. This, in turn, will also affect the driver's speed choice [170]. Moreover, it is recently found that drivers may make a decision to accept or reject trajectory-driven haptic-shared-control ADAS based on the interactions during this phase (as found in Chapter 7, conducted before Chapter 5). Hence an overall better understanding of driver behaviour before curve entry, and the factors which can affect it, is much needed and still lacking.

To improve on the current understanding of driver prepositioning behaviour, this chapter investigates the effect of changing vehicle velocity and road curvature on prepositioning behaviour. The prepositioning phenomenon is investigated through empirical data from a fixed-base driving simulator experiment, where the velocities $V = 50$

and 80 km/h, and radii $R = 204$ and 350 m, were tested. Geometrical changes in prepositioning were analysed. Moreover, a possible instigator to preposition, the Time-to-Line-Crossing (TLC), which is a road safety-margin perception metric, is investigated along with the extent to which drivers adhered to optimal prepositioning trajectories.

This chapter is structured as follows, in Section 5.2 the theoretical kinematic relationships between prepositioning and lateral acceleration and TLC are given. The independent and dependent variables, road design and hypotheses for the simulator experiment are described in Section 5.3. Section 5.4 presents the experimental findings, interpretations and discussion. Finally the conclusion is presented in Section 5.5.

5.2. BACKGROUND

Prepositioning is a spontaneous phenomenon that is found to occur in naturalistic driving studies [38, 117]. However, only *optimal* prepositioning behaviour is mathematically modelled and quantified in literature. In fact, prepositioning is an emergent behaviour from all optimal control models, which race-line drivers undertake whilst pursuing an optimal trajectory [30]. In the light of trajectory optimal curve negotiation in lateral acceleration for a given speed, this section provides the theoretical background to understand the benefit of and need for prepositioning.

For optimal curve negotiation, one of the performance criteria is to minimise lateral acceleration. Lateral acceleration a_{lat} is given in Eq. (5.1), where κ is curvature and V is velocity.

$$a_{lat} = \kappa V^2 \tag{5.1}$$

Minimising lateral acceleration in the curve can be accomplished by reducing speed as well as pursuing a trajectory with a reduced curvature. Here, a particular trajectory, called the optimal race-line, would be achieved without compromising velocity, but rather, minimising curvature. A minimum curvature is achieved by driving a trajectory that exhibits a radius larger than that of the road. The largest possible radius R_i that can be taken through a constant curvature curve is illustrated in Fig. 5.2 with a dashed line.

Using trigonometric relations the following equation for the maximum radius R_i achievable in a curve is given in Eq. (5.2), taken from [30].

$$R_i = \frac{R_2 - R_1 \cos(\phi/2)}{1 - \cos(\phi/2)} \tag{5.2}$$

Here, R_1 and R_2 are the respective inner and outer radii of the curve boundaries, whereas ϕ is the angular curve length. This formulation also indicates a distance λ that defines the needed distance to start 'turning into the curve'. This distance is geometrically given in Eq. (5.3).

$$\lambda = \sin(\phi/2)(R_i - R_1) \tag{5.3}$$

This distance λ increases with increasing ϕ and increasing $R_i - R_1$. Eq. (5.2) also gives us information about conditions to maximise the considered R_i. When the angu-

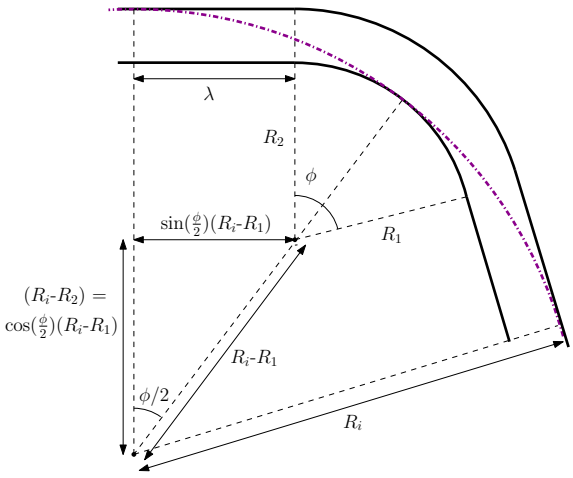

Figure 5.2: This figure illustrates a curvature minimising trajectory [30] in dotted purple. The angular curve length spans 2ϕ, the radius of the inner boundary of the curve is R_1, that of the outer boundary is R_2 and the maximum radius of the optimal lateral acceleration a_{lat} trajectory is R_i. The distance needed to optimally turn into the curve is λ.

lar curve length tends towards π (180 deg), R_i reaches R_2 meaning there is no benefit that arises from prepositioning anymore. Hereby, the larger the curve length, the less the potential reduction of the experienced a_{lat}, achieved from prepositioning.

5.2.1. TIME-TO-LINE-CROSSING AND LATERAL ACCELERATION

The need for prepositioning can be motivated from a lateral acceleration perspective, i.e. Eq. (5.1). A second consideration, as lateral acceleration is not a signal that can always be perceived by drivers (e.g. in a fixed based driving simulator), is that prepositioning helps maximising the Time-to-Line-Crossing (TLC). The TLC is perceived visually and indicates the safety margins to the road [110] [121]. It has been found in an early study that drivers may try to maintain an upper bound to TLC in curves, where for sharper curves drivers may preposition or slow down more [63].

In the following, the relationship between a_{lat} and TLC is derived. Fig. 5.3 illustrates a vehicle negotiating a curve with radius R at a lateral displacement δ from the inner road boundary on a road width of w. If the driver were to take their hands off the steering wheel, the vehicle would exit the road boundary ahead of the vehicle at a distance D. The time to reach this road boundary is the constant heading 'straight' TLC, considered in this chapter.

Using Fig. 5.3, the following geometric relations can be established. Using Pythagoras' theorem Eq. (5.4) is formulated.

$$(R+\delta)^2 + D^2 = (R+w)^2 \tag{5.4}$$

Assuming that δ and w are both small compared to R so that both δ^2 and w^2 can be neglected, the following equation is achieved:

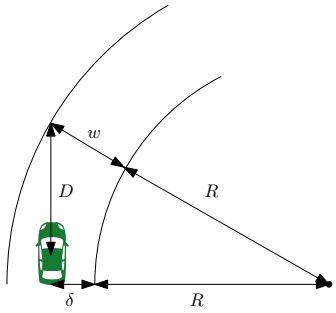

Figure 5.3: This figure illustrates a geometric diagram which is used to derive trigonometric relationships related to TLC. The car is driving on a single lane curve with radius R and width w, at a distance δ to the road boundary. If the driver were to let go of the steering wheel, the car would exit the lane at a distance D ahead of the car.

$$D = \sqrt{2R(w-\delta)} \tag{5.5}$$

From Fig. 5.3 it follows that the TLC is given by Eq. (5.6).

$$\text{TLC} = \frac{D}{V} \tag{5.6}$$

$$V = \frac{D}{\text{TLC}} \tag{5.7}$$

Substituting Eq. (5.5) into Eq. (5.1) the relationship between a_{lat} and TLC is established:

$$a_{lat} = \kappa V^2 = \frac{2(w-\delta)}{\text{TLC}^2} \tag{5.8}$$

As is clear from Eq. (5.8), lateral acceleration is inversely proportional to the TLC^2 in a curve. Hence, minimising lateral acceleration can be achieved by maximising TLC. Therefore, prepositioning (and continuing to follow the optimal curve trajectory) increases TLC. Moreover, the amount of TLC that can be maximised depends on curve length, just as the amount of a_{lat} that can be minimised depends on curve length (as deduced from Eq. (5.2)). Specifically, one can gain more TLC with a shorter curve length.

5.3. EXPERIMENT

5.3.1. CONTROL TASK

In the experiment, subjects were asked to perform a curve negotiation (lateral control) driving task at a fixed speed (i.e., steering only) in a fixed-base driving simulator. A heavy sedan of 1.8 m wide was used to visually simulate the vehicle on a single lane road of 3.6 m width. Simplified vehicle dynamics identical to those used in previous investigations [152] approximated by a bicycle model, was used to simulate the lateral-directional car dynamics.

Table 5.1: The centerline Time-to-Line-Crossing (TLC), full curve length and the time spent driven in the curve at the respective constant velocity.

Condition	TLC (s)	curve length (m)	curve time (s)
R= 204 m, V=50 km/h	1.39	122	8.76
R= 204 m, V=80 km/h	0.87	122	5.47
R= 350 m, V=50 km/h	1.82	178	12.81
R= 350 m, V=80 km/h	1.14	178	8.01

5.3.2. INDEPENDENT VARIABLES

This study has two factors; car speed and road curvature, each having two levels comprising $V \in \{50, 80\}$ km/h and curvature with radius $R \in \{204, 350\}$ m. The choice of the speeds is based on the maximum velocity on different roads in the Netherlands: 50 km/h in urban roads and 80 km/h in some motorways. The corresponding radii are chosen based on the Dutch empirical road design formula in Eq. (5.9) [140].

$$R \geq \frac{7V^2}{210 - V + 9p} \tag{5.9}$$

This equation sets radii limitations to maintain driver comfort empirically, where V stands for velocity in km/h and p stands for the road tilt in a percentage between 0 and 100%. A road tilt of 10% is taken, as is common in a highway T-section [140]. The radius of 204 m is the minimum radius for a velocity of 80 km/h whilst that of 350 m is the minimum radius for a velocity of 100 km/h. A velocity of 100 km/h is taken to design the road radius, as this velocity is a consecutive limit velocity in Dutch traffic rules and the corresponding radius is acceptable for both velocities of 50 and 80 km/h. This is done to avoid a push for prepositioning to reduce the curvature, so that most likely, the captured data is consistent with natural well-learned behaviour.

The designed curves with corresponding fixed velocities, naturally have different centreline Time-to-Line-Crossing (TLC), curve length and curve time, as given in Table 5.1. Using Eq. (5.5) and Eq. (5.6), the constant heading, straight TLC [29] values are computed assuming constant curvature (the maximum curvature of the road) and that the driver drives on the centreline. With a smaller radius (larger curvature) and larger velocity, the centreline TLC decreases, meaning a less 'safe' driving situation.

5.3.3. ROAD DESIGN

Five road design aspects are of interest for this experiment: the road width, the straight section length, the curve length, the clothoid length, and the order of presentation of right and left curves. The road width was chosen as 3.6 m in line with Dutch road specifications and previous studies [68, 152].

The design of the *straight section length* determines whether prepositioning is studied independently for each curve. To minimise the chance of interference with 'post-positioning' after curve exit, the straight section length was fixed to 25 seconds.

The design of the *curve length* is optimised with respect to the extended tangent point (ETP). That is, the minimum curve length such that at curve entry the curve exit is not visible [68], i.e., the curve length is always twice the distance to the tangent point, when prepositioning maximally. This is necessary for a fair comparison between both curvatures. The angle ψ to the tangent point, for a curve with radius R and road width w, from the outer curve boundary at curve entry can be computed with Eq. (5.10). The angle to the extended tangent point is twice this value, where the corresponding curve length l_{ETP} is calculated according to Eq. (5.11).

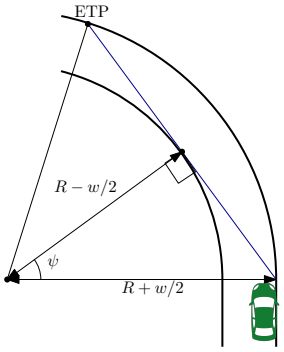

Figure 5.4: This figure illustrates the maximum achievable distance to the extended tangent point (ETP). This occurs when a driver prepositions maximally. The angular length to the tangent point is ψ, whereas the angular length to the extended tangent point is 2ψ.

$$\psi = \arccos\left(\frac{R - w/2}{R + w/2}\right) \tag{5.10}$$

$$l_{ETP} = 2\psi R \tag{5.11}$$

The *clothoidal length* of the curve l_c is designed according to Eq. (5.12), which is an empirical relation taken from Dutch road design principles [140]:

$$l_c = \frac{1}{9}R \tag{5.12}$$

Finally, the order of presentation of left and right curves was randomised, as it should be unpredictable to avoid precognitive behaviour where drivers preposition based on expectation.

5.3.4. APPARATUS

The driving task was performed in the fixed-based driving simulator at the Human-Machine-Interaction Laboratory at TU Delft. The visual scenery was presented using three LCD projectors covering a horizontal field-of-view of 180 deg and a vertical field-of-view of 40 deg, with an update rate of 50 Hz. A MOOG FCS Ecol8000S Actuator running at 2500 Hz was used for generating self-aligning haptic torques on the steering wheel.

5.3.5. EXPERIMENTAL-SETUP AND PROCEDURE

Before participating in the experiment, participants indicated whether they had a driver's license and how experienced they were and provided written informed consent. The conditions were presented in randomised order which was balanced across participants. Before each condition, the participants were given a familiarisation run of 160 s.

5.3.6. SUBJECTS AND INSTRUCTIONS

The experiment was performed by twenty-four subjects between the age of 20 and 63 years (average of 32.8 years and standard deviation of 16 years). There was an average of 28 years of driving experience with a standard deviation of 12.8 years. Within this population, there are 4 novice drivers (i.e., did not have a drivers license yet) and 4 highly experienced drivers (i.e., had between 39 to 42 years of driving experience). All subjects were instructed to drive as they normally would and to hold their hands on the steering wheel at a" ten-to-two" position.

5.3.7. DEPENDENT VARIABLES

There are two sets of dependent variables. The first defines geometric properties that were used to quantify participants' prepositioning behaviour, whilst the second set of metrics is used to verify whether prepositioning improves TLC (i.e. whether the prepositioning was performed to achieve an improvement in TLC or equivalently, lateral acceleration).

QUANTIFYING PREPOSITIONING BEHAVIOUR

For the quantification and characterisation of the measured differences in prepositioning behaviour, several geometric metrics derived from the driven trajectory are extracted, as illustrated in Figure 5.5. This figure shows a curvature profile κ and the corresponding lateral position within the road boundaries s_{lat}, both as a function of time. As the road width is 3.6 m and the car width is 1.8 m, the effective road width is 1.8 m, as indicated in shaded grey in Fig. 5.5. The optimal race line deduced according to [30], is illustrated in the s_{lat} domain as well as in a 2D mapping. The following six metrics were derived from the s_{lat} trajectory before curve entry to characterise prepositioning behaviour:

- y_{max}: the maximum prepositioning displacement relative to road centerline. Note that the presented $y_{max}*$ in the results section is y_{max} corrected for curve direction, i.e., always showing positive values in the direction of the outer curve boundary.

- y_e: the curve entry position relative to road centerline. Again, y_e* is y_e corrected for curve direction.

- τ_{in}: the time before curve entry that the driver decides to turn *into* the curve, i.e., the time at which y_{max} occurs.

- y_b: The trajectory bias on the straight section, here calculated as the average between 20 and 10 s before curve entry. This is based on a previous study, where drivers start to decelerate before curve entry, 4.6 s on average, with a 1.5 s σ [68].

- Δy_{max}: the maximum prepositioning displacement relative to the straight section bias y_b.

- Δy_e: the curve entry position relative to the straight section bias y_b.

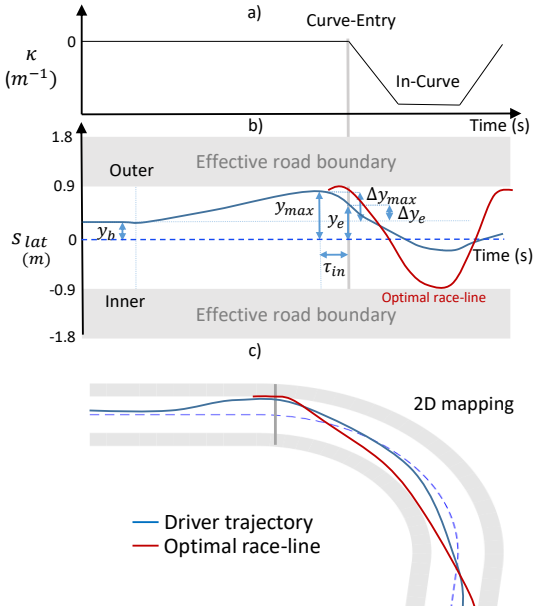

Figure 5.5: This figure illustrates the geometric overview of the metrics used to quantify prepositioning behaviour, with an equivalent 2D schematic. Plot a) illustrates the curvature profile, indicating the prepositioning phase before curve entry and the in-curve phase. Plot b) illustrates a possible prepositioning trajectory in the lateral position domain s_{lat}. Geometric quantifications of the prepositioning trajectory are indicated. These include: straight road bias y_b, maximum prepositioning displacement y_{max}, curve entry position y_e, turn in time τ_{in}, relative maximum prepositioning displacement Δy_{max} and relative curve entry position Δy_e. Additionally, the optimal race-line as deduced from [30] is illustrated in red. Finally, plot c) illustrates a 2D birds-eye-view schematic of the driver trajectory and optimal race-line.

For the evaluation of the statistical significance of differences in these metrics, two two-way repeated-measures Analysis Of Variance (ANOVA) was performed, i.e., one for each curve direction, with the velocity and road radius as factors. This is done for two reasons: 1) drivers drive differently on the right and left curves, as found in Chapter 4 and, 2) the effect of curve direction on prepositioning behaviour is not the focus of this study. Assumptions for this statistical test to be valid include 1) *normality*, which was tested using the Shapiro-Wilk test, and, 2) *sphericity*, which was corrected for using Greenhouse-Geisser correction in case of a violation.

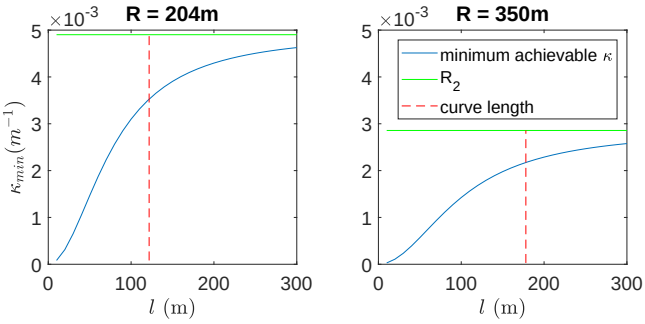

Figure 5.6: This figure illustrates the minimum achievable κ (using Eq. (5.2)) through following the optimal raceline. This is plotted against curve length for the different considered curves. The curve length considered for the designed roads is indicated by a red vertical line. R_2 is the radius of the outer boundary of the road, as outlined in Fig. 5.2. The minimum achievable κ is plotted on the left for R = 204m and on the right for R = 350m.

METRICS TO VERIFY WHETHER PREPOSITIONING IMPROVES TLC

In addition to comparing prepositioning trajectories themselves, we also focus on TLC and lateral position s_{lat} during the curve. Specifically, we compare the trajectory generated by participants to the theoretically calculated optimal race-line. For this, the following metrics are considered:

- The average lateral position s_{lat} w.r.t. road centreline across 5 runs (20s before curve entry and including the curve part).
- The 'straight' Time-to-Line-Crossing (TLC), i.e., the constant heading, constant velocity TLC [29]. Calculated through an iterative predictive calculation on the future physical road ahead with an *effective* road width of 1.8 m.
- Optimal-race-line Variance Accounted For (VAF). The VAF between measured driver trajectories and the optimal race-line in the lateral position domain (as described in [30]) was used to quantify the similarity of the two trajectories. The VAF is a fitting metric [17], where 100% indicates a perfect match.
- TLC ratio signifies $\frac{\text{TLC}_{driver}}{\text{TLC}_{centerline}}$, which is the ratio between the average TLC of the driver (across the constant curvature section) and that of the centerline. These ratios are computed for the in-curve section. A value larger than 1 signifies increased TLC relative to the centreline values, also given in Table 5.1.

5.3.8. HYPOTHESES

Four hypotheses were formulated for this experiment:

H.I Maximum prepositioning displacement y_{max} and Δy_{max} increase with increasing road curvature and with increasing velocity.

H.II τ_{in} (related to the distance λ) decreases with increasing road curvature and with increasing velocity.

Table 5.2: The optimal τ_{in} values derived from the required λ distance (in Eq. (5.3)) needed for an optimal race-line.

Optimal τ_{in} (s)	V = 50 km/h	V = 80 km/h
R= 204 m	1.72	1.07
R= 350 m	2.03	1.26

H.III All subjects display prepositioning behaviour, which is shown by exhibiting an $\Delta y_{max} > 0.05$ m.

H.IV By prepositioning drivers can obtain a more optimal TLC in curves.

H.I is based on Fig. 5.6, computed based on Eq. (5.2) which gives the potential decrease in curvature when drivers preposition before a curve as a function of curve length. For R = 204 m prepositioning can potentially minimise driven curvature by 0.0014 m^{-1}, whereas for R = 350 m by 0.0006 m^{-1}. Since the curve with larger curvature (R= 204) facilitates a larger improvement in visual safety margins (TLC) when prepositioning, it is hypothesised that prepositioning increases with curvature. Moreover, when speed increases, the TLC further decreases, so it is hypothesised that the perceived need by the drivers to increase their TLC by prepositioning may increase as well.

H.II is based on theoretical derivations. Substituting the considered roads in Section 5.3.3 into Eq. (5.3) and Eq. (5.2), the needed τ_{in} to achieve the optimal race-line is given in Table 5.2. Here it is clear that τ_{in} decreases with increasing curvature and velocity.

H.III is based on the previous finding that drivers start to decelerate before curve entry 4.6 s on average [68]. Thereby preparation for the curve by prepositioning is complementary and expected from everyone.

H.IV is based on Fig. 5.2, i.e., following the optimal race-line can improve TLC in curves. However, not all drivers follow such an optimal trajectory, as found in Chapter 4. Therefore, it is hypothesised that by only prepositioning, one may also harvest a fraction of these TLC benefits.

5.4. RESULTS

5.4.1. EFFECT OF VELOCITY AND CURVE RADIUS ON PREPOSITIONING BEHAVIOUR

Figure 5.7 shows boxplots of $y_{max}*$, Δy_{max}, y_e* and Δy_e. The schematic on the far right illustrates the physical realisation of how the average y_{max} position (average across the number repetitions of each participant and all participants through each specific curve) is affected by curve radius and velocity in left curves. The illustration of the y_{max} position is done at $s_{long}(\tau_{in})$ (equivalent to the drivers 'turn in distance' λ), which is different for each velocity. It can be seen that on average drivers increase their y_{max} with velocity and curvature.

Similarly, $y_{max}{}^*$ for right curves increases with increasing curvature ($\frac{1}{R}$) and velocity, as seen in Fig. 5.7. An average displacement increase (for both right and left

curves) of 0.08 m is found for increasing road curvature, and an average displacement increase of 0.09 m is found for increasing velocity. For y_{max} significant effects (without any normality violations) for velocity (right: ($F(1,23) = 6.4$, $p < 0.05$), left: ($F(1,23) = 18.3$, $p < 0.01$)) and radius (right: ($F(1,23) = 7.3$, $p < 0.05$), left: ($F(1,23) = 13.2$, $p < 0.01$)) were found.

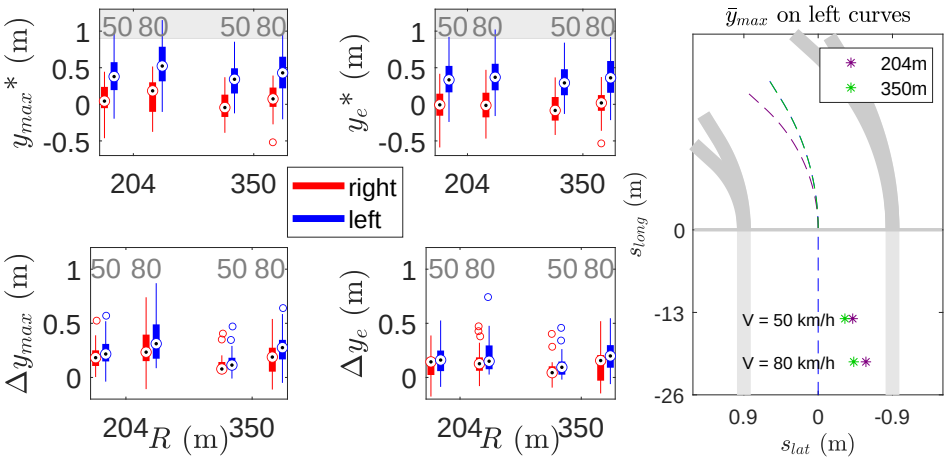

Figure 5.7: The figure illustrates the metrics used to describe prepositioning behaviour: $y_{max}*$, y_e*, Δy_{max} and Δy_e. The grey bar in the $y_{max}*$ and y_e* plots signify the effective road boundary of the *outer* part of the curve. The last figure on the right illustrates the average maximum displacement $\bar{y}_{max}$ at the average turn in position (s_{lat}, s_{long}) at the average turn-in time τ_{in} = 0.92 s, for a left curve

The relative prepositioning displacement Δy_{max} (relative to y_b, where y_b is unaffected by road radius or velocity, average of -0.15 m) is also affected by curvature and radius for both right and left curves. Here we see an average displacement increase of 0.08 m for increasing road curvature, and an average displacement increase of 0.1 m for increasing velocity. For Δy_{max} significant effects (without any normality violations) for velocity (($F(1,23) = 6.7$, $p < 0.05$)-*right curves*,($F(1,23) = 22.4$, $p < 0.01$)-*left curves*) and radius (($F(1,23) = 23.3$, $p < 0.01$)-*right curves*,($F(1,23) = 27.8$, $p < 0.01$)-*left curves*) were found.

The absolute and relative entry position y_e and Δy_e seems only to be significantly affected by velocity in left curves. For left curve, an average y_e increase of 0.05 m and an average Δy_e increase of 0.04 m are found with increasing in velocity. Statistical analysis shows significant effects for y_e (without any normality violations) for velocity in left curves (($F(1,23) = 6.7$, $p < 0.05$) and for Δy_e (without any normality violations) for velocity in left curves ($F(1,23) = 7.07$, $p < 0.05$).

The drivers' turn-in time τ_{in}, as shown in Figure 5.8, is not significantly affected by road radius and velocity. We do see, however, a difference between right and left curves. Across both curve directions, τ_{in} occurs at 1.04 s before curve entry on average, with left curves exhibiting a τ_{in} that is 0.25 s lower than found for right curves on average.

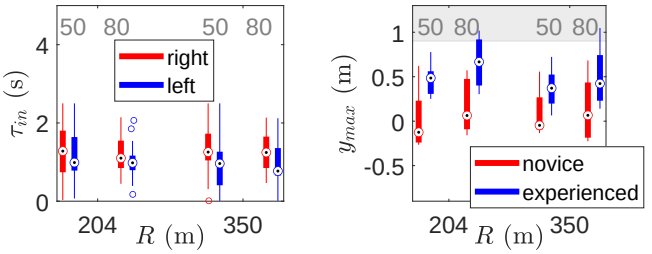

Figure 5.8: The left figure illustrates boxplots which present τ_{in} for both right and left curves. The right figure illustrates the maximum prepositioning y_{max} exclusively for left curves, distinguishing experienced from novice drivers, i.e., experienced have more than 40 years of driving experience and novice do not have a drivers license.

The average straight section bias y_b is -0.15 m, where for left curves it is -0.18 m and for right curves it is -0.13 m. This effect is significant when considering a 3-way ANOVA with curve direction as the third factor ($F(1,23)$ =8.9, p <0.01). This could indicate that drivers start to drift in the direction of prepositioning before 10 s (where the straight section bias is computed). This is not surprising, since the curve direction can be distinguished from the horizon line (in the driving simulator) about 16 s before curve entry.

These results show that drivers consistently and significantly exhibit larger y_{max} (preposition more) with increasing velocity and curvature. Hereby, the drivers' chosen prepositioning seems to increase when decreasing the default safety margin (decrease the centerline driving TLC, i.e., an increase in the theoretical curve demand) of the simulated road environment as given in Table 5.1, thereby H.I is accepted. Moreover, the significant effect for left curves is consistently more substantial, with drivers exhibiting both larger absolute and relative y_{max}, a significant increase in y_e for velocity and even turning into the curve sooner. In fact, the significant increase in y_e with increasing velocity, which is only found for left curves, could be a consequence of drivers exhibiting a larger Δy_{max} of 0.08 m on left curves. Drivers exhibit a preferred y_b sign (negative sign), which also means that the room for prepositioning is larger for left curves than it is for right. Moreover, with an average τ_{in} of 0.92 s (for left curves) before curve entry, the more you swing wide, the harder it is to maintain the same y_e position.

Theoretically, the distance λ for following the optimal race-line increases with reducing road curvature. This translates to an optimal turn in time $\tau_{in,opt}$ that decreases with increasing road curvature and velocity. However, the parameter τ_{in} is not significantly affected by curvature or velocity. This rejects H.II. This also implies that not all drivers follow the optimal race-line, in line with the findings of Chapter 4.

5.4.2. DOES EVERYONE PREPOSITION?

For the collected experiment data, it is found that 88% of the runs tested shows a significant prepositioning contribution (defined here as having a $\Delta y_{max} > 0.05$ m), as shown in Fig. 5.7. This validates H.III. Moreover, it seems that prepositioning behaviour could

be related to experience in driving, especially without any available vestibular feedback as given in Fig 5.8. In our experiment, the inexperienced driver tended to drive near the centerline. In contrast, the experienced drivers could be exercising a habit that is known for its TLC maximising and a_{lat} reducing benefits. This data suggests that prepositioning behaviour is in fact, a 'learned' behaviour, but a better experiment (with more participants in each driver group) is needed to make a conclusive statement.

For those who do preposition, it is found that drivers start to preposition at 7.3 s before curve entry on average. This value is markedly different from the 'preview' time most driver models implement [174] [177], suggesting that this is a separate category of anticipation that occurs for active steering inputs. Nevertheless, the average turn-in time τ_{in}, independent of curve radius or velocity, is found to be 1.04 s on average, which is similar to the human control preview time found in preview identification studies [174].

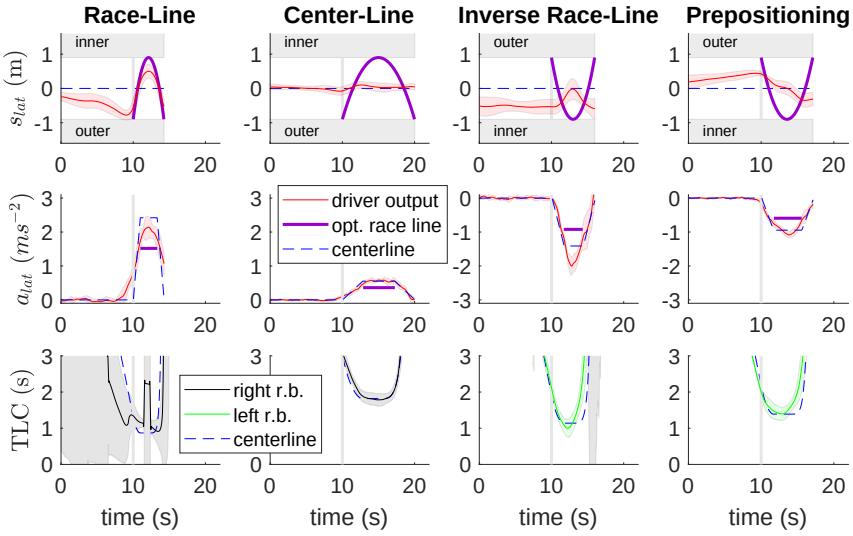

Figure 5.9: This figure illustrates the average and standard deviation of 10 runs of four drivers: a Race-line driver, a center-line driver, an inverse race-line driver and a driver that prepositions only. The first row illustrates the lateral position s_{lat}, second row a_{lat} and third row TLC. The centerline output is illustrated as a baseline reference in all domains: s_{lat}, a_{lat} and TLC. The optimal race line is illustrated for comparison. The minimal achievable a_{lat} through following the optimal raceline is computed using Eq. (5.2) and Eq. (5.1).

5.4.3. DOES PREPOSITIONING BENEFIT TLC?

Drivers preposition more with increasing velocity and curve radius. At the same time, the curves that have a larger curvature ($\frac{1}{R}$) and velocity have decreased centerline TLCs, as shown in Table 5.1. TLC can be related to the driver's safety margin [110] [121], where a higher TLC implies a better safety margin. Therefore, it seems that drivers

could be reacting to those decreased centerline driving TLC by prepositioning more. As outlined in Section 5.2, race-line drivers adopt this prepositioning strategy to optimise their lateral acceleration by taking a radius larger than the road. This results in a reduced curvature of the driven trajectory ($\kappa = 1/R$) and reduced required steering due to reduced yaw rate ($r = \kappa V$). This results in an efficient trajectory that 'cuts the curve', which significantly increases TLC in curves, both due to the 'cutting' position on the inner part of the curve and the acquired heading. This trajectory thus also better approximates the optimal race-line [30].

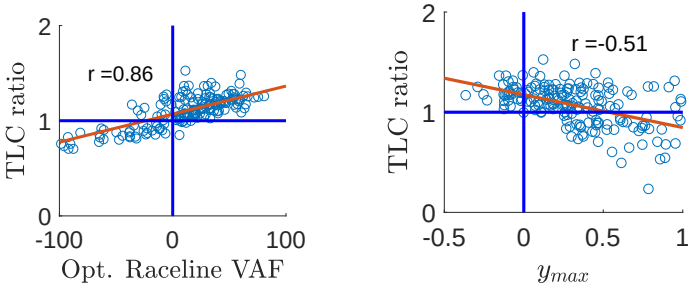

(a) TLC ratio v.s. optimal race-line VAF (fit) (b) TLC ratio v.s. y_{max}

Figure 5.10: This figure illustrates scatter plots and the best linear fit to these data points. Moreover, the correlation between the presented metrics are computed.

Figure 5.9 illustrates the lateral position s_{lat}, lateral acceleration a_{lat} and TLC of 4 different drivers: a *Race-Line* driver that closely follows the optimal race-line through the curve, a *Center-Line* driver maintaining close proximity to the centerline, a *Inverse Race-Line* driver who follows a strategy that is opposite to the optimal race-line through the curve and a *Prepositioning* driver who *only* prepositions. Fig. 5.9 clearly shows that following the optimal race-line decreases a_{lat} and increases TLC, as expected. Whereas taking an opposite strategy (inverse race-line) increases a_{lat} and decreases TLC further. However, merely prepositioning does *not* result in the same benefits as following the optimal race-line.

Fig. 5.10 illustrates two correlation plots, relating the y_{max} position and TLC ratio and the extent to which drivers follow the optimal race-line (optimal race-line VAF) and the TLC ratio. Although both cases show a statistically significant correlation, that for the optimal race-line VAF is much more robust, which means it is more reliable. The correlation of 0.86 indicates a strong relationship between following the extent to which one follows the optimal race-line and an improvement in TLC. The correlation of -0.51 suggests that the less drivers preposition, the higher their TLC, which shows that prepositioning alone does not help with increasing your TLC. In fact, some drivers with a y_{max} close to 1 obtain a substantially small TLC ratio (possibly by following an outer-curve negotiation style (OOO)). This result rejects H.IV. Hereby, following the optimal race-line, which includes prepositioning, results in larger TLC values in curves, however, prepositioning alone may not.

If the curve length were designed to be longer, however, such correlations may become weaker. The relative reduction of curvature of the optimal race-line is known to

become smaller with increasing curve length as implicated by Eq. (5.2), resulting in a reduced heading difference which further reduces the increase in TLC of the optimal race-line. This also means that the benefit of prepositioning decreases with increasing curve length, as theoretically expected. Conversely, it can also be argued that drivers preposition in relation to curve length. The conditions with a radius of 204 m have not only a larger curvature, but also a smaller curve length compared to a radius of 350, as given in Table 5.1. Hereby the significant increase in prepositioning could be a result of both a smaller curve radius (sharper curve) and a shorter curve length, which is an unforeseen possible confound. Therefore it is recommended to perform an experiment that can independently vary curve length and curve radius.

5.5. CONCLUSION

This chapter concludes that the maximum prepositioning displacement is found to increase with increasing velocity and decrease with increasing road radius. The turn in time, i.e., the time before curve entry where the driver exhibits their maximum prepositioning position is not affected by curvature or velocity (contrary to theory for *optimal* prepositioning), with an average of 1.04 s. Prepositioning behaviour is evident for 88% of the driven curves by the participants, where drivers exhibit a prepositioning displacement relative to road bias of more than 0.05 m.

It seems that drivers adapt their prepositioning to the perceived centerline Time-to-Line-Crossing. However, the driver can increase their TLC *only* by following the optimal race-line, which, inherently exhibits prepositioning. Prepositioning alone, without tending to follow to optimal race-line is not found to improve TLC. Overall, this chapter shows the importance of incorporating these findings in driver modelling (Chapter 6), which will bridge the gap between straight road and in-curve driving behaviour, bolstering the descriptive capacity of these models.

6

DRIVER PREPOSITIONING MODELLING AND INTEGRATION

Model building is the art of selecting those aspects of a process that are relevant to the question being asked.

J.H. Holland

In the previous chapter prepositioning was geometrically quantified. This chapter deals with the modelling of prepositioning. A stand-alone prepositioning-path generation model is introduced, which can be coupled with any control-theoretic driver steering model with an explicit position feedback loop. The proposed prepositioning-path model is combined with the Van Paassen control-theoretic driver steering model, which is assessed in Chapter 3. The combined model is evaluated using the descriptiveness assessment metric.

The contents of this chapter are to be published as:

Barendswaard, S., Pool, D.M., Van Paassen, M.M., Boer, E.R., Abbink, D.A., 'A pre-curve prepositioning model to enhance the descriptiveness of specific control-theoretic driver steering models.', *IEEE Transactions on Human Machine Systems.*

6.1. INTRODUCTION

The understanding of human driving behaviour in curve negotiation is of increasing importance in the design of Advanced Driver Assistance Systems (ADAS) and autonomous cars [69] [98] [86]. The acceptance of some types of ADAS systems (Lane Keeping Assitance, Lane Departure Warnings) is highly dependent on the automation being able to understand and thereby capture an individuals' driver behaviour and respond accordingly [152]. Moreover, the acceptance of fully autonomous cars also depends on their ability to adjust motion and path planning algorithms with realistic human driving behaviour, because only human-like manoeuvres feel 'natural' and can build trust between driver and automation [154].

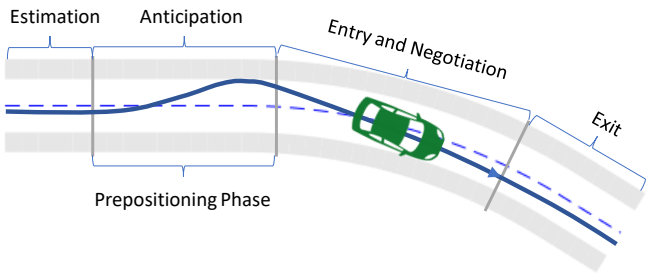

Figure 6.1: Illustration of the four phases of curve negotiation [117] [38], where this chapter focuses on the curve anticipation or the *prepositioning* phase

Naturalistic driving studies show that human drivers' curve driving behaviour is split into four distinct phases [117][38], as illustrated in Fig. 6.1: 1) *estimation*, where drivers estimate the curvature and adjust their speed for curve entry, 2) *anticipation*, where drivers anticipate the curvature and (pre)position accordingly, 3) *entry and curve negotiation*, where drivers adjust their steering and lateral acceleration to stay within road boundaries, and finally, 4) *exit*, accelerating to the desired speed again. Two of these phases, *estimation* and *anticipation*, occur before curve entry, wherein the *anticipation* phase, drivers demonstrate prepositioning behaviour. A majority of drivers exhibit prepositioning behaviour, namely 88 % of the runs tested in the experiment from Chapter 5. Drivers may practice prepositioning as an attempt to increase their Time-to-Line-Crossing or decrease their steering demand in the curve. Contrary to the preview time identified in driver models, typically between 0.6-1.2s [174][177], prepositioning (for a straight section of 25s) is found to be initiated 7s on average before curve entry in Chapter 5, though there is a large variation between subjects.

Most driver steering models ignore the occurrence of prepositioning by assuming that drivers enter the curve at lane centre [116] [11]. Other steering models assume they start to turn into the curve a few moments before curve entry at a certain preview time [174][177] (around 1s), or always preposition *optimally* through resorting to optimal control with cost function minimisation [30]. Prepositioning is an emergent behaviour in optimal preview control models. Due to being controlled with a cost-function, the degree of prepositioning behaviour is linked with the in-curve be-

haviour. However, as shown in Chapter 5, naturalistic driver prepositioning behaviour does not show such a link. For example, normal curve cutters and offset outer curve negotiators (driver class trajectory behaviour outlined in Chapter 4) preposition, but do not follow the optimal race-line in the curve. Therefore, this chapter places its focus on control-theoretic driver models because they 1) facilitate increased flexibility in modelling prepositioning and in-curve behaviour and 2) include parameters that have physical meaning such as time delay and neuromuscular gain.

State-of-art control-theoretic modelling suffers from their capacity to describe driver behaviour accurately. Chapter 3 has concluded that such models are not able to describe more than 17% of the driving population (for right curves), attributable to the inability to describe prepositioning. The reason for this shortcoming is twofold: 1) current control-theoretic models are validated in a limited domain (steering angles) [116] [11] and, 2) control-theoretic driver models are largely developed for abstract and unrealistic roads consisting of a continuum of curvature (sum of sinusoids), which is a limited domain not allowing for any explicit prepositioning that is observed in a real world curve separated by straight sections. [174].

To bridge the modelling gap, this chapter introduces a novel prepositioning-path model to be used as an extension with control-theoretic driver steering models [16]. The Van Paassen control-theoretic driver model was chosen for two reasons: 1) it exhibits such exclusive position feedback and, 2) resulted in the best model performance from the assessment in Chapter 3. This chapter augments the Van Paassen model with a prepositioning-path reference input. The effectiveness of this augmentation is tested by applying the descriptiveness assessment procedure from Chapter 3, to both the augmented and baseline Van Paassen model.

This chapter is structured as follows: Section 6.2 presents the proposed prepositioning-path model and its capacity to capture different types of prepositioning behaviour. The Van Paassen model is elaborated in Section 6.3. The augmentations made to the Van Paassen model to account for prepositioning is given in Section 6.4. The descriptiveness assessment of the augmented model relative to the un-augmented model is given in Section 6.5, discussion in Section 6.6. Finally, the conclusions are presented in Section 6.7.

6.2. PREPOSITIONING-PATH MODEL

This section introduces the prepositioning path model by first explaining the relevant geometric parameters in Section 6.2.1, which are used for the Model Formulation in Section 6.2.2. The parameters of the formulated model are explicitly explained in Section 6.2.3. The results of verifying the introduced path model is given in Section 6.2.4.

6.2.1. GEOMETRIC PARAMETERS

A prepositioning trajectory expressed in terms of lateral position on the road s_{lat} is geometrically defined by seven measurements, using the metrics illustrated in Fig. 6.2, which are introduced in Chapter 5.

These geometric quantities can be used to compare the different types of prepositioning behaviour in the driving population, or how road characteristics and velocity

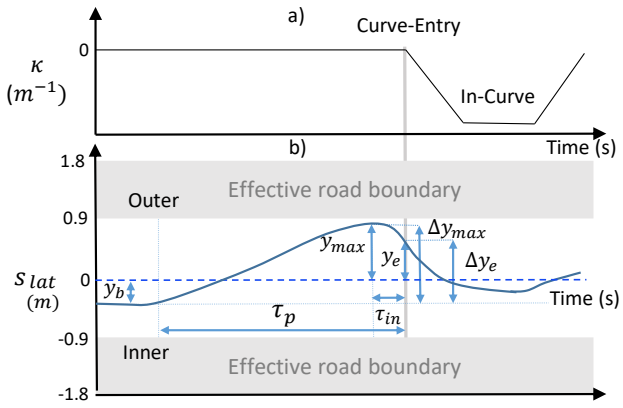

Figure 6.2: A geometric quantification of prepositioning behaviour. These comprise the straight section bias (y_b), the maximum prepositioning distance (y_{max}), the curve entry position (y_e), the maximum prepositioning distance relative to the straight section bias (Δy_{max}), the curve entry position relative to the straight section bias (Δy_e), turn in time (τ_{in}) and prepositioning time (τ_p).

can affect prepositioning, as was done in Chapter 5. In fact, for modelling prepositioning, the following three geometric parameters are of importance.

- y_b: straight section average bias
- τ_p: prepositioning time
- Δy_{max}: total prepositioning displacement relative to y_b

With these three parameters, the prepositioning response between y_b and y_{max} is defined, the rest of the geometric measurements follow.

The straight section average bias y_b is the biased position on the road that the driver already exhibits before prepositioning initiation (all drivers exhibit a lateral bias on straight roads which depends on the lane of travel, especially in multi-lane motorways [92]). When the driver spots the onset of a curve, on average at 7s before curve entry (as found in Chapter 5), the driver starts to drift away from the initial road bias y_b to a displacement which is equivalent to y_{max}, positioning a net distance of Δy_{max} at τ_{in} s before curve entry (around 1s).

6.2.2. MODEL FORMULATION

The geometric description of prepositioning can be modelled using a bias and a sigmoid function. However, the control-theoretic model becomes 'active' at curve entry, see section 6.3. Therefore, to minimise the s_{lat} modelling interference, the contribution of the prepositioning path model should be smoothly damped down to zero immediately after curve entry, which requires a second sigmoidal function.

As given in Eq. (6.1) and illustrated in Fig. 6.3, the proposed model starts at the constant bias level y_b. The first sigmoidal function is activated around the prepositioning time τ_p, which, is related but not equal to the sigmoid time shift τ_1. Instead

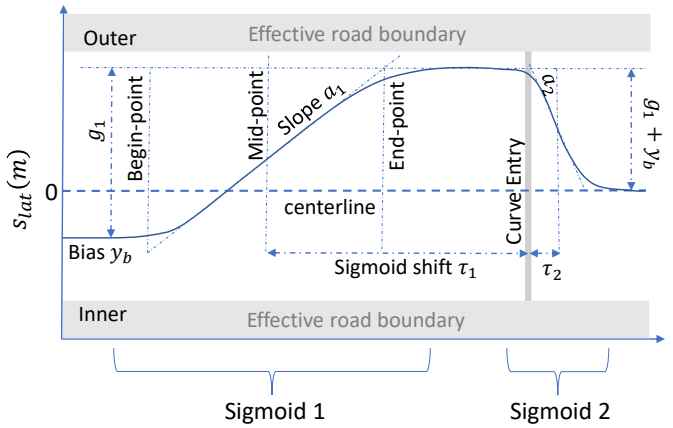

Figure 6.3: Illustration of the prepositioning path model, with the sigmoidal parameters graphically presented [37]. The model starts at the bias level y_b. The sigmoid amplitude is g_1 for the Sigmoid 1 and $g1 + y_b$ for Sigmoid 2. The sigmoid slopes a_1 and a_2 and the sigmoid time shifts τ_1 and τ_2, marking the midpoint of the sigmoid from the curve entry point.

6

τ_1 represents the middle of the sigmoidal slope, given by a_1. The amplitude at this first sigmoid is g_1, equivalent to Δy_{max}. Once the curve onset is detected the second (sharper sloped sigmoid) is activated around τ_2, which facilitates 'turning into the curve' at time τ_{in} (around 1s before curve entry as found in Chapter 5) and reduces the prepositioning response back to zero (with an amplitude of $y_b + \Delta y_{max}$, or $y_b + g_1$). The parameters τ_p and τ_{in} are not explicitly included in the model formulation, however, they emerge from the choices made for τ_1, τ_2, a_1 and a_2.

$$y_{prep}(t) = y_b + \frac{g_1}{1 + e^{-a_1(t-\tau_1)}} - \frac{g_1 + y_b}{1 + e^{-a_2(t-\tau_2)}} \tag{6.1}$$

The model of Eq. (6.1) only outputs a path, relative to the road centreline. This path is independent of velocity. The path model also does not produce the crucial driver output of steering angles. This path is meant to be used as a position reference in control-theoretic driver steering models.

6.2.3. MODEL PARAMETERS

The prepositioning part is essentially captured by the amplitude g_1, the slope a_1 and the time τ_1 as given in Fig. 6.3. The bias y_b can be extracted from the apparent bias in the data. This is especially clear from a dataset that has a long straight section before curve entry, as is the case in Chapter 5. The second sigmoid has an amplitude that is not independently tunable ($g_1 + y_b$), whereas the slope a_2 and time τ_2 are. The function of the prepositioning path model ends at curve entry. Therefore, the parameters a_2 and τ_2 must ensure that the prepositioning contribution reduces to zero as fast as possible without incurring any modelling artefacts. To reduce the path as fast as possible, the parameter a_2 should be large and τ_2 should be zero. However, very high a_2

and low τ_2 results in a near step-like response which causes steering oscillations as a modelling artefact. To avoid such artefacts, the sigmoid time τ_2 is fixed to a value of 0.5 s and the parameter a_2 is restricted not to exceed 2 m/s. For a y_{max} at 0.9 and y_b at 0, this second sigmoid reduces to zero at a velocity of 0.6 ms^{-1}, acceleration of 0.57 ms^{-2} and jerk of 1.31 ms^{-3}. Nevertheless, since the control-theoretic model is 'active' in this region, the decrease rate may not be perceived as strongly. Additionally, the centripetal acceleration required to make the curve is significantly larger at 2.41 ms^{-2}. Also, the steering demand is much lower for prepositioning than it is to make the curve, i.e. the instantaneous curvature needed to preposition is around 1/10th of that for a curve with 0.005 m^{-1}.

6.2.4. MODEL VERIFICATION

The prepositioning-path model is verified on the dataset of Chapter 5, which includes 24 participants. This dataset is chosen because the experiment performed to generate the dataset is designed with a straight section of 25 s to detect prepositioning behaviour accurately. For identification with the prepositioning path model, only 20 s before curve entry is taken, to exclude any possible post-curve repositioning behaviour. The identification of the prepositioning path model is performed using Eq. (6.2). Here the equation finds the argument P (a parameter vector $[g_1, \tau_1, a_1, a_2, y_b]$) that minimises the expression which defines the absolute difference between the modelled prepositioning path $\hat{s}_{lat}(P)$ and the true path s_{lat}.

$$\hat{P} = \arg\min_{P} \sum_{i=1}^{n} (\hat{s}_{lat}(P) - s_{lat}(i))^2 \tag{6.2}$$

IDENTIFICATION RESULTS

Relating to the trajectory classes defined in Chapter 4, there are three different prepositioning categories, i.e., where y_e (from Fig. 6.2) is either at the *Outer*, *Middle* (centerline band) or *Inner* part of the curve. The explicit mapping between prepositioning category and trajectory class as defined in Chapter 4 is illustrated in Fig. 6.4.

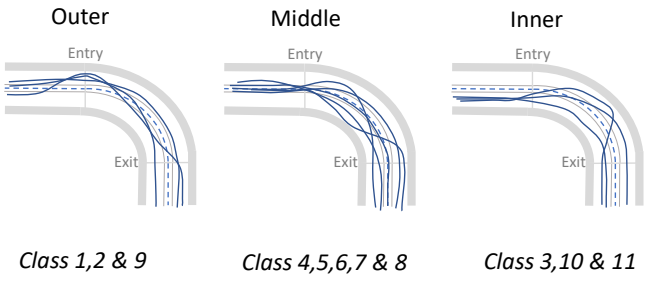

Figure 6.4: Prepositioning category and how these relate to the classes of the 11-class classifier outlined in Chapter 4, shown from a birds eye-view perspective. Due to the exaggerated scale of the road width, the curves look unnatural. Class 1 & 2 denote curve cutting, 3 biased inner-curve drivers, 4-8 centerline onset drivers, 9 biased outer-curve drivers and, 10 & 11 counter curve cutting.

Therefore, the prepositioning paths are grouped into three categories per curve direction, as illustrated in Fig. 6.5. Here, a representative driver, i.e. a driver that best represents the average of the category, is selected, and the fit of the respective driver's trajectory is illustrated. The VAF (a signal validation metric, the Variance Accounted For, where a 100% means that the model can describe 100% of the behaviour, i.e. a perfect trajectory match) of the model fit on the representative driver is also listed on the upper right road boundary. On the lower right road boundary, the percentage occurrence *O* of each prepositioning category from the considered dataset is listed. The corresponding identified parameters of the representative drivers are given in Table. 6.1. All the model fits for each trajectory are illustrated in Fig. C.1 in Appendix C.1, where a sample of these is illustrated in Fig. 6.6. The corresponding identified parameters are shown in Fig. 6.7.

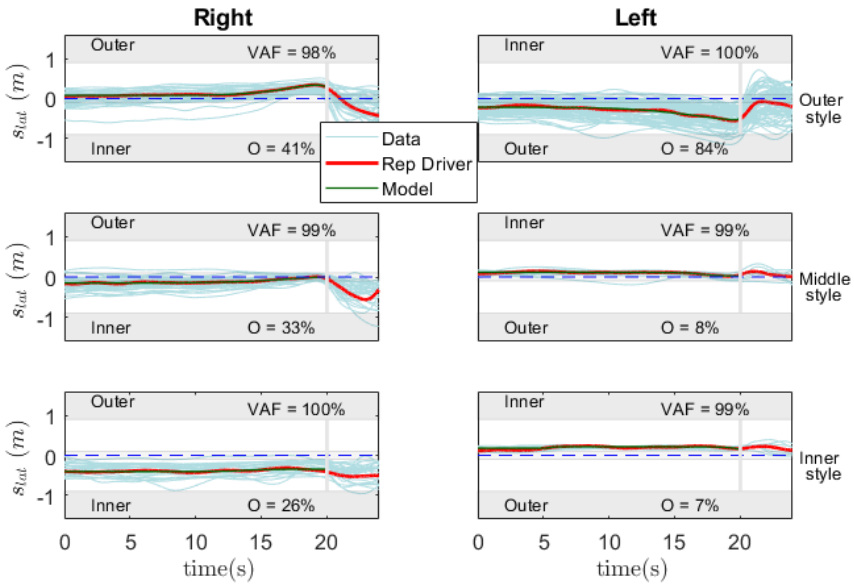

Figure 6.5: This figure illustrates the fitting (in the lateral position domain s_{lat}) of the proposed prepositioning model on a representative driver (given in red) in each of the different prepositioning categories. Each row represents a different prepositioning category: Outer, Middle and Inner, whereas each column represents a curve direction (left, right). The prepositioning path model fit is illustrated with a dark green line. The measurement data is presented in light blue. The corresponding model fit value of the representative driver is indicated by the VAF value. The percentage of drivers falling into a given prepositioning category is indicated by the Occurrence (O).

Without prepositioning, a regular control-theoretic model would stay on the centerline in the straight sections. Without accounting for prepositioning, the computed VAF for the straight sections is zero, even for the Middle prepositioning cases. Hereby, this path model shows great improvement. However, the prepositioning path model is unable to describe oscillatory behaviour on straight sections, which is apparent with

Table 6.1: Identified parameters of the representative drivers of each prepositioning category: Outer, Middle and Inner, for each curve direction Right (R) and Left (L).

P. type	y_b (m)		g_1 (m)		τ_1 (s)		a_1 (m/s)		a_2 (m/s)	
	R	L	R	L	R	L	R	L	R	L
Outer	0.08	-0.21	0.25	-0.44	5.6	3.3	0.33	0.38	2	2
Middle	-0.14	0.11	0.20	-0.07	2.1	3.0	0.42	1.6	2	0.28
Inner	-0.42	0.17	0.07	-0.03	0.03	0.01	0.14	1.1	1.8	1.4

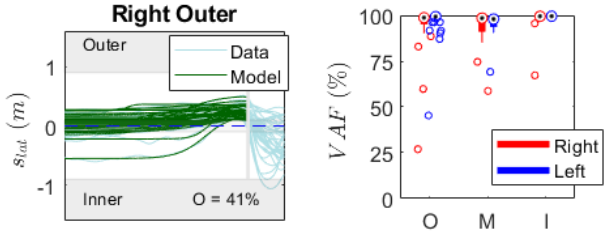

Figure 6.6: The first figure illustrates a sample of all the resulting prepositioning model fits to the driver data illustrated in Fig. 6.5. The second figure illustrates the corresponding VAF fit values for each prepositioning category, *O* outer, *M* middle and *I* inner.

some drivers. This is illustrated in Fig. 6.6, where the modelled data does not always capture the variation in the straight Section, explicitly apparent with the two responses starting at the inner side of the curve. The incapacity of the prepositioning path model to capture oscillations is also evident with the low VAF outliers in Fig. 6.6. This model is not designed to capture human stochastic variability; instead, it should capture an average trend that drivers may want to follow or may want to be guided by in trajectory-driven ADAS. Such pre-curve oscillatory behaviour most likely has its origin in other non-curve related factors such as distraction, attention fluctuation, confusion or simply a small slip in steering.

For both right and left curves, the most commonly observed prepositioning category is that of *Outer* curve prepositioning. As illustrated in Fig. 6.5 for right curves *Outer* has 41% occurence, *Middle* has 33% occurrence and *Inner* has 26% occurrence. Whereas for left curves, *Outer* has 84% occurrence, *Middle* has 8% occurrence and *Inner* has 7% occurrence. The high *Outer* curve occurrence for both right and left curves could be related to most drivers exhibiting a significant Δy_{max}. The distinctively higher left *Outer* occurrence of 84% could be related to the straight road bias having an average of -0.1m, i.e. -0.1m already falls in the outer part of the curve.

For the representative driver selected from the dataset, the Δy_{max} is larger for left curves than for right curves. This is indeed reflected in the parameters identified, given in Table 6.1. The value for g_1 for right curves is 0.25 m whereas for left curves it is -0.44 m. This representative parameter set is in line with the aggregate findings illustrated in Fig. 6.7. The median g_1 for left curves is at -0.31 m whereas for the right curves it is 0.20 m. This means that on average drivers make greater lateral displacements to preposition for a left curve, than for a right curve. This is likely because drivers are

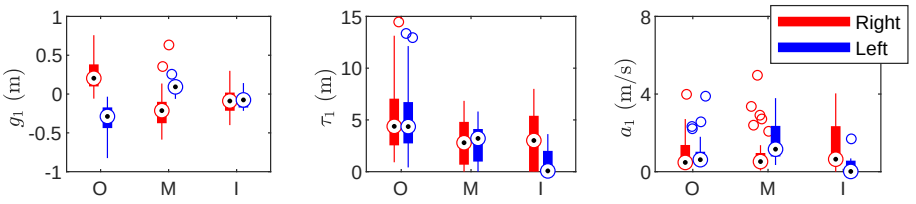

Figure 6.7: These plots illustrate the range of identified parameters resulting from fits on the dataset of Chapter 5, for the proposed prepositioning path model y_{prep}. These parameters are the prepositioning amplitude g_1, the sigmoid time shift τ_1 and the sigmoid slope a_1. They are presented separately for prepositioning responses that enter the curve at the *O*uter, *M*iddle and *I*nner part of the curve

on the left side of the car, which, causes a bias offset to the left. Moreover, they can perceive the distance to the left road edge more accurately, which means they can cut a left curve more easily.

For both right and left curves, the identified g_1 value for Outer prepositioning is larger than for both Middle and Inner prepositioning. This is evident from the identified parameters of both the representative drivers in Table 6.1 and Fig. 6.5. The value of g_1 being distinctively larger for Outer prepositioning cases means that most drivers preposition to strategically optimise their steering demands in the curve. Moreover, the magnitude of the sigmoidal time shift constant τ_1 seems to decrease in parallel with Outer exhibiting the largest prepositioning times. Here we see that when the prepositioning displacement is large (g_1) the prepositioning time is usually larger as well. This is because enough time is required to steer toward y_{max}; otherwise, the prepositioning manoeuvre may induce an uncomfortable lateral acceleration. Finally, the slope parameter a_1 does not seem to abide by a trend between the different prepositioning categories.

6.3. VAN PAASSEN DRIVER STEERING MODEL

The Van Paassen model [177] makes the use of feedback on predicted curve cutting position ($y_{cc}(t)$, independent of y_{prep}) and an independent feedforward path that sets the steering rate requirements in the curve at preview time τ_f. With this, the concept is that the steering requirement to make the curve results from the feedforward loop, whereas trajectory fine-tuning is left for the position feedback loop.

For *'trajectory fine-tuning'*, the feedback loop predicts the position of the car ahead, i.e. $\hat{s}_{lat}(t+\tau_n)$. It sets a future curve cutting reference $y_{cc}(t)$, which is an exact trajectory position for the feedback loop to follow, see Fig. 6.8. This $y_{cc}(t)$ position is only non-zero when in the curve (non-zero curvature input). This explicit position reference also enables other position references that are not explicitly linked to the current or previewed curvature input.

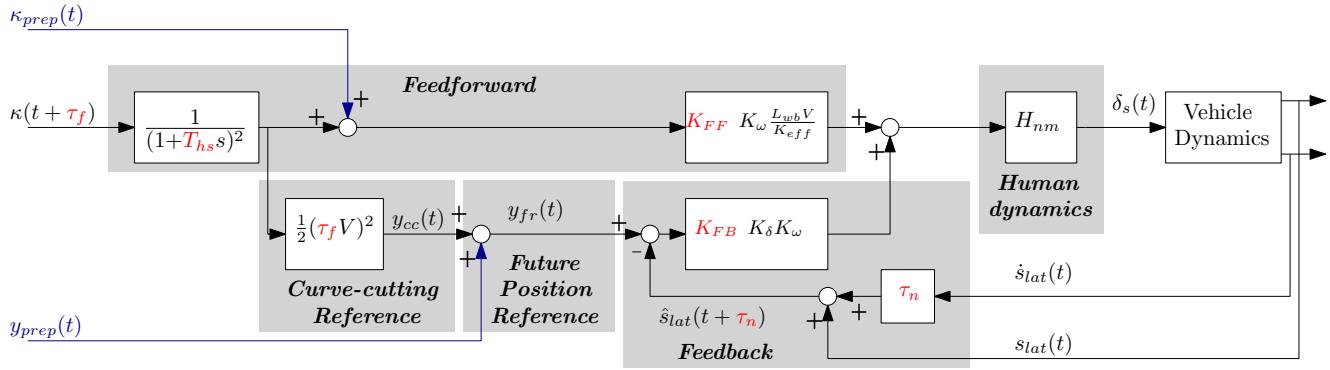

Figure 6.8: Control structure for the adapted Van Paassen model introduced in this chapter, where the adaptations are indicated in blue. The original Van Paassen model is given in Fig. 3.3. Here the previewed curvature is taken as feedforward input, along with the additional prepositioning curvature κ_{prep}, whereas the feedback is on lateral position and lateral velocity. Moreover, the prepositioning position reference y_{prep}, is added to the curve-cutting position reference $y_{cc}(t)$ to obtain the future position reference $y_{fr}(t)$.

6.3.1. OUTLINE

The feedforward input signal (the previewed curvature $\kappa(t+\tau_f)$) is filtered with a second-order filter with a break frequency at T_{hs}, see Fig. 6.8. This is especially important when the road is non-clothoidal, to prevent any transition responses from the dynamical transfer functions. The filtered curvature signal is then used for both the feedback loop and the feedforward path.

In the feedback loop, the filtered curvature is used to compute the curve cutting offset as given in Eq. (6.3), further geometric explanations are given in Chapter 3 and Appendix B.4.

$$y_{cc}(t) = \frac{1}{2}(\tau_f V)^2 \kappa(t+\tau_f) \tag{6.3}$$

The curve cutting distance is taken as a reference in the position tracking feedback loop. The predicted position (the fed-back signal $\hat{s}_{lat}(t+\tau_n)$) is computed by linear approximation at the prediction look ahead time τ_n (ideally $\tau_n = \tau_f - T_{hs}$ [173]). This error is then compensated with feedback gain K_{FB}.

PARAMETERS FOR IDENTIFICATION

The considered parameters in the Van Paassen model for identification are identical to those considered in Chapter 3. These are five: the lookahead time τ_f, the filter time constant T_{hs}, the prediction time τ_n, the feedback gain K_{FB} and the feedforward gain K_{FF}. The other parameters shown in the model such as the wheelbase gain K_w, wheelbase distance L_{wb}, the effective steering ratio K_{eff} and K_δ are all constants and properties of the vehicle driven. The values for these are taken from [177].

6.4. AUGMENTATIONS MADE TO THE VAN PAASSEN MODEL

This chapter augments the Van Paassen model by adding a prepositioning position reference y_{prep} to the curve cutting reference, which augments the signal from only being a curve cutting reference $y_{cc}(t)$ to a future position reference $y_{fr}(t)$, as illustrated in Fig. 6.8. Additionally, the corresponding prepositioning curvature κ_{prep} is fed into the feedforward loop to provide a matching feedforward contribution.

In the feedforward loop, the filtered curvature is summed with the corresponding prepositioning curvature $\kappa_{prep}(t)$. This prepositioning curvature is computed based on the identified prepositioning position, as given in Eq. (6.4) and Eq. (6.5). For the model simulation, the analytical expressions are first explicitly evaluated, using the parameterised prepositioning model of Eq. (6.1).

$$a_{prep}(t) = \frac{d^2 y_{prep}(t)}{dt^2} \tag{6.4}$$

$$\kappa_{prep}(t) = \frac{a_{prep}(t)}{V^2} \tag{6.5}$$

The prepositioning path consists of two sigmoids, where, after curve entry, the contribution of $a_{prep}(t)$ should not be considered for prepositioning curvature. This is because, after curve entry, the second sigmoid's sole purpose is to damp out the position reference to zero as fast and smoothly as possible, to allow for the control-theoretic

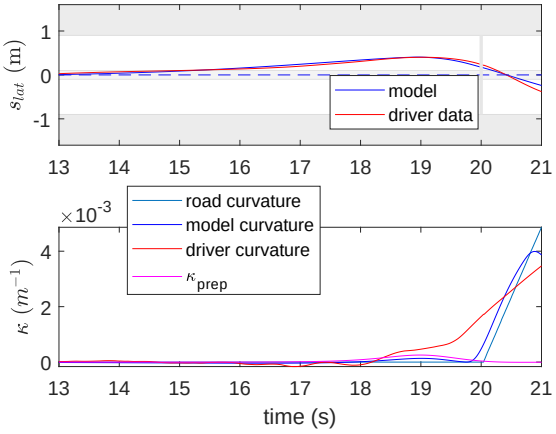

Figure 6.9: The first figure is in the lateral position domain (s_{lat}) and illustrates driver data in red and the corresponding prepositioning model fit in blue. The second figure, in the (instantaneous) curvature domain, shows the road curvature in light blue, the modelled curvature in blue, the driver curvature in red and the prepositioning curvature (κ_{prep}) in pink.

model to manage the steering response in the curve independently. Therefore, the computed lateral acceleration resulting from the second sigmoid after curve entry defined in Fig.6.3, after curve entry, is a modelling artefact and should not be considered for κ_{prep}. To avoid this, a fade-out function can be used to attenuate the κ_{prep} to zero in the curve. To illustrate that this combination does not result in a peak in driven curvature when implemented with the model, Fig. 6.9 shows the prepositioning response of a given driver and the corresponding instantaneous curvature of the driver and the model. The driver, illustrated in red in Fig. 6.9 exhibits a maximum jerk of 0.94 m/s^3, whereas the modelled prepositioning curvature, illustrated in pink, exhibits a maximum jerk of 0.05 m/s^3. Note that comfortable levels of (guidance) jerk should lie below 0.6 m/s^3 [14].

The combined curvature feedforward is scaled by the driven velocity to obtain a steering angle that is proportional to the required yaw rate in a curve, as given in Eq. (6.6).

$$r = \frac{V}{R} = V\kappa \tag{6.6}$$

This means that when the feedforward gain K_{FF} is too high the steering wheel deflections is tuned to a sharper curve.

In the augmented Van Paassen model, the prepositioning path is implemented in both the feedforward and feedback loop, which may seem redundant. To illustrate the added benefit of including the prepositioning path in both the feedforward loop through κ_{prep} and the feedback loop through y_{prep}, model responses are given in Fig. 6.10. The pure feedback response (red) shows a lag in the tracking of y_{prep}. Thereby, due to the build-up of error, overshoots the desired Δy_{max}. The pure feedforward

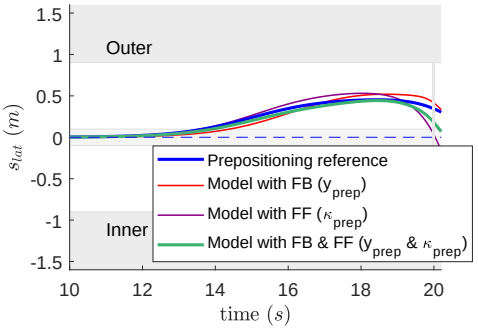

Figure 6.10: Illustration of a prepositioning reference y_{prep} in the lateral position domain (s_{lat}) is given in blue. The augmented Van Paassen model tracking response with only feedback (on y_{prep} position feedback) is given in red, only feedforward (κ_{prep}) is given in purple and both feedback and feedforward (κ_{prep} + y_{prep}) is given in green. The parameters used to model $y_{prep}(t)$ are g_1 = 0.5 m, τ_1 = 5 s, a_1 = 1 m/s, τ_2 = 0.5 s and a_2 = 2 m/s.

model (purple) does not display any lag and initiates prepositioning at the right time; however, due to lack of position feedback also overshoots Δy_{max}. When both feedback and feedforward contributions are combined in the model, the best response is achieved. The VAF of the response with only feedback is 96%, with only feedforward is 95% and with both feedforward and feedback is 98.5 %. In addition, from a human factors perspective, overshoots may be highly negatively perceived by drivers because these overshoots may be out of the comfort zone of drivers, moreover, with haptic shared controllers, drivers may need to give extra torques to overcome these overshoots.

For the simulation of this driver model, both the vehicle/road dynamics and neuromuscular dynamics are needed to close the control loop. The details of these dynamics are given in Chapter 3.

6.5. DESCRIPTIVENESS ANALYSIS

This section considers the 'descriptiveness' of the Van Paassen model with and without the prepositioning model. The descriptiveness is a property of a model that is evaluated by the assessment procedure outlined in Chapter 3. Descriptiveness shows the capability of a given model to capture different types of drivers, categorised into 11 classes of trajectories introduced in Chapter 4. A representative driver from each trajectory class is taken, and the assessed model is fit to this driver. When the model fit to the data for a particular class is higher than 80% VAF and falls within the same class, the model is considered to sufficiently describe the population in that trajectory class. The descriptiveness is then the sum of the class occurrence (percentages of drivers within a class) of the trajectory classes that the model can sufficiently describe (i.e., the total percentage of drivers that can be modelled accurately).

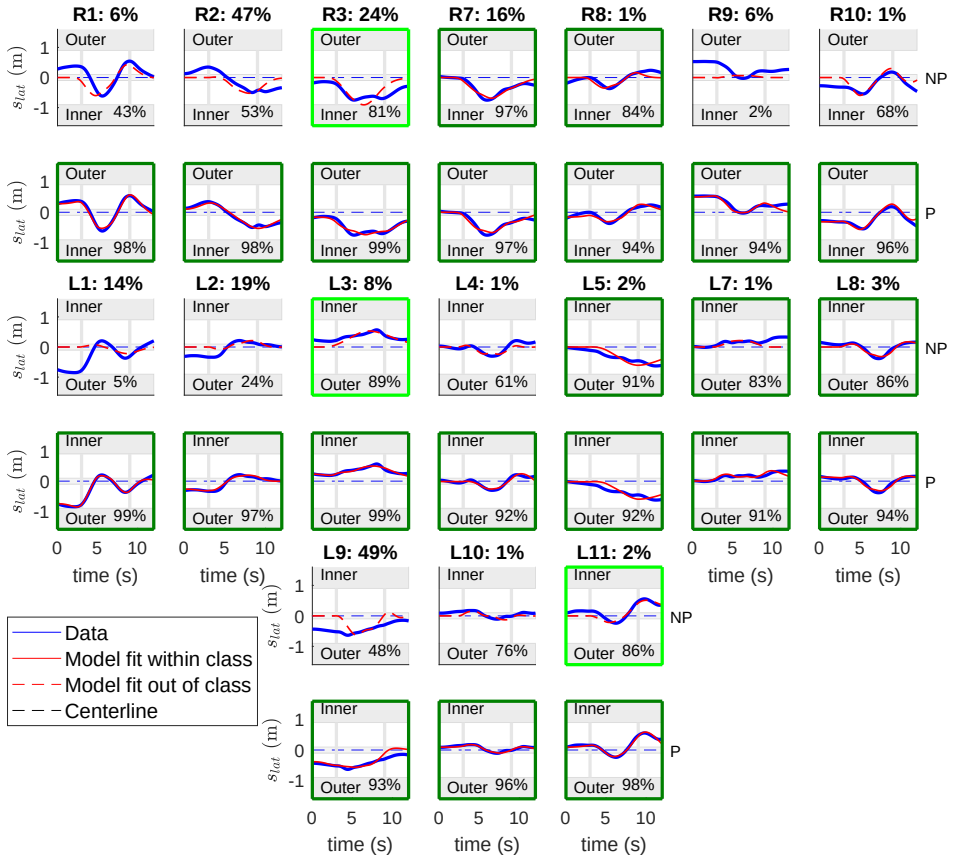

Figure 6.11: This figure illustrates the model fit results for right and left curves for the Van Paassen model without (NP) and with prepositioning-path model (P). The grey bars around 0.9 to 1.1m and -0.9 to -1.1m represent the effective road boundary. The vertical grey lines represent the beginning and end of a curve. The representative driver of each class is illustrated by the 'Data' (blue line). The model fits which result in other classes, are shown in a dashed red line. The model fits resulting in the same class are represented by a (non-dashed) red line. When the fit is above 80 %, the plot has a dark green box frame. When the VAF is above 80 % and not within the same class, the plot has a light green frame.

6.5.1. VAN PAASSEN MODEL WITH AND WITHOUT PREPOSITIONING PATH MODEL

In this section, the impact of including the proposed prepositioning path model in the Van Paassen model is explicitly shown. Fig. 6.11 illustrates the fit of both the Van Paassen model without prepositioning (NP), in the odd rows, and with prepositioning (P), in the even rows. The graph illustrates the fit of the model in red and the representative driver of each class in blue. The curve direction and class are indicated in the title along with the class occurrence. The VAF value of the fit is given in the bottom right road boundary. The classes that exhibit significant prepositioning offsets such as

class 1, 2 and 9 for both right and left curves, demonstrate a large improvement in VAF value when including prepositioning. Classes 4, 5, 7 and 8 are classes that do not really preposition (middle), as illustrated in Fig. 6.4. Thereby, the difference in fit with and without prepositioning is indeed minimal. Classes 3, 10 and 11 deal with trajectories that enter the curve at the inner side of the curve. Since the prepositioning magnitude is generally not large with these classes, the difference in fit is also not as large. Table C.1 in the Appendix shows the identified parameters for each response shown in Fig. 6.11. These parameters also generally show a larger difference between P and NP for classes with large prepositioning.

The descriptiveness of a model is the sum of the class occurrences of the classes the model can sufficiently capture. This is evaluated differently for left and right curves. For right curves, the descriptiveness is 17% without prepositioning, as only R7 and R8 are accurately modelled. With prepositioning, descriptiveness for right curves is 100% as the representative driver data of all classes can be accurately described. For left curves, the descriptiveness is 7%, as only L5, L7 and L8 are accurately modelled without prepositioning, and also 100% with prepositioning.

6.6. DISCUSSION

The aim of this chapter was to develop a prepositioning path model that can capture a variety of different prepositioning path behaviours and, combine this path model with the Van Paassen control-theoretic model. Overall, including prepositioning in the Van Paassen model has made a considerable improvement in the types of driver trajectories the considered control-theoretic driver model can capture in a curve. This increase highlights the importance of considering prepositioning in driver models. This importance is further stressed from a human-factors perspective. When the prepositioning phase is in line with the driver's intentions, this good 'first impression' may make the driver more likely to accept the in-curve guidance advisory, even when the in-curve driving behaviour is not the same as theirs, as found in Chapter 7. Therefore, it is recommended to use the proposed augmented driver model, or the prepositioning path model in combination with another suitable driver model, for future ADAS applications.

The prepositioning path model was combined with the Van Paassen model in this chapter, however, can also be combined with other control-theoretic driver steering models. Adding a position and curvature reference was straight forward in the Van Paassen model due to an already available explicit position feedback reference y_{cc} and feedforward on curvature. For other driver models, those who work with tracking waypoint angles [173] or tangent points [116], may need more restructuring to account for prepositioning. Hereby, the prepositioning path model is not restricted for application with the Van Paassen Model.

For future ADAS and autonomous driving (AD) applications, accounting for prepositioning is essential. However, prepositioning and thereby prepositioning parameters are not only dependent on individual style (Chapter 4), but also on upcoming curve details (Chapter 5). Hereby, the next step is to understand how parameters such as g_1 and τ_1 are affected by different types of curves, i.e. varying curve length, curve width, curve radius.

When a driver prepositions in a way that would allow for minimising their lateral acceleration or maximising their time-to-line crossing, this does not mean that this objective actually comes into fruition in the curve (as concluded in Chapter 5). In many cases drivers preposition to set themselves up for greater comfort or safety in the upcoming curve, nevertheless once in the curve, it may not be demanding enough to continue to follow the optimal race line. In other words, if a driver prepositions in a way that would allow for minimising lateral acceleration, it is not guaranteed that his in-curve trajectory will continue to satisfy this objective. This is evident from having three prepositioning categories and 11 classes, as illustrated in Fig. 6.4. There is a link between prepositioning category and trajectory type. However, there is no singular mapping. Such a disconnect places the current prepositioning modelling approach at an advantage compared to the only other means to model prepositioning behaviour: optimal preview control [30]. Therefore, the proposed method of modelling prepositioning in a piece-wise fashion provides greater flexibility than optimising for a certain objective. This flexibility is useful for ADAS and AD to match driver expectations (Chapter 4 & 5).

Drivers preposition predominantly towards the outer part of the curve. This is evident by the most commonly observed prepositioning category being the 'Outer' type and the sigmoidal gain g_1 being predominantly positive for right and negative for left curves. Consequently, such drivers proceed to travel towards the inner part of the curve. Most drivers (except class 11 (IOI) drivers) 'travel' towards the inner part of the curve relative to their curve entry point, even when they don't make a transition across the centreline band. This finding is illustrated by the representative drivers illustrated in Fig. 6.11. This finding also supports the notion of the curve cutting reference in the Van Paassen model $y_{cc}(t)$ always being positive as outlined in Eq. (6.3).

The prepositioning path model assumes that the bias y_b of the driver is constant. However, making this assumption may not be a representation of reality, due to driver straight section fluctuations and variance (as for example illustrated in Fig. 6.6) which may differ from driver to driver. Nevertheless, from a modelling perspective (especially for driver guidance in ADAS), this assumption would even be desired. Guiding a driver with such variance and possible oscillations may not only decrease safety but also reduce acceptance as such guidance is not expected nor predictable.

The prepositioning responses that fall in the 'Middle' and 'Inner' styles could be modelled accurately using a pure gain y_b. Therefore, identifying physical parameters that assume a significant g_1 may result in over-fitting. In fact, to model a constant bias prepositioning response, τ_1 should be close to zero, but can also be very large, in combination with a_1 being very small. Nevertheless, for all other responses exhibiting a significant g_1 (>0.05 m), the proposed model does not over-fit.

The purpose of the second sigmoid in Eq. (6.1) is to 'damp' out the prepositioning response in the proposed prepositioning path model to zero at curve entry and to do so as smoothly as possible without inducing any modelling artefacts. The considered dataset set used to model prepositioning has a 6 s long curve that exhibits 2 s clothoidal transitions. For this dataset, the considered τ_2 was fixed to 0.5 s, which means that the sigmoid's time constant is set at a quarter of the clothoidal curve transition, i.e., the prepositioning response is damped out well before the constant curva-

ture part. However, when considering different realistic sharp curves, i.e., those with a transition of 1s, the considered τ_2 should decrease to avoid any recurring modelling artefacts. Hence, parameters a_2 and τ_2 should depend on the curvature profile whilst not exceeding comfortable jerk levels [14]. Therefore, it is recommended to investigate exactly how these relate to each other explicitly.

6.7. CONCLUSIONS

State-of-the-art control-theoretic driver models are unable to describe more than a fifth of the driving population in and near curves accurately. This shortcoming is primarily traced down to the inability to describe prepositioning behaviour, i.e., driver behaviour before curve onset. This chapter introduces a new prepositioning path model for drivers' lateral position changes prior to curve entry with parameters that have geometrical significance. This model is shown to describe various prepositioning behaviours successfully and is verified on a dataset explicitly designed to study prepositioning.

The prepositioning path model is used to augment (extend) the Van Paassen control-theoretic driver model by adding an extra position reference. The descriptiveness of the Van Paassen model with and without the prepositioning augmentation is assessed. The Van Paassen model without prepositioning can accurately describe 17% of drivers in right curves and 7% of drivers in left curves, less than a fifth of the drivers considered. The Van Paassen model with prepositioning can sufficiently describe 100% of drivers in both right and left curves. The findings of this chapter introduce a new modelling concept, that can be used to model individual curve drivers at fixed speed accurately. Moreover, the individualised model can directly be used to personalise various ADAS in curves accurately, where these ADAS would need to detect curve entry 10s prior (to determine the prepositioning path model) and the curvature profile around preview time (τ_f) for the control-theoretic model.

IV

Application with Haptic Shared Control

7

FOUR DESIGN CHOICE HAPTIC SHARED CONTROLLER TO INCREASE DRIVER ACCEPTANCE

Part IV consists of two Chapters which will deal with applying and evaluating some of the personalisation methods considered in the previous parts of this thesis with haptic shared control. This chapter introduces the first implementation of the novel 'Four-Design-Choice' (FDC) haptic shared controller. One of the design choices is an independent controller reference, which can be separately personalised. A human-in-the-loop experiment is carried out to evaluate three implementations of haptic shared control: 1) the FDC with a personalised reference, 2) the FDC with 'One-Size-Fits-All' reference and, 3) a conventional haptic shared controller. Since this chapter was carried out before the model assessment findings of Chapter 3, the Mars model is used as driver model to identify the personalised and 'One-Size-Fits-All' controller reference trajectories. The haptic shared control realisations are evaluated in terms of conflict occurence, driver torque and subjective acceptance.

The contents of this chapter have been published as:

Scholtens, W. and Barendswaard, S. and Pool, D.M. and Van Paassen, M.M. and Abbink, D.A., 'A New Haptic Shared Controller Reducing Steering Conflicts', *IEEE Systems Man Cybernetics conference, Miyazaki, Japan, 2018.*

7.1. INTRODUCTION

Sharing control through haptics is an alternative approach that mitigates the negative side effects that arise from traded control in highly automated vehicles [59][137], such as loss of awareness [122] and skill degradation [45]. Haptic Shared Control (HSC) supports drivers in the steering task by providing assisting torques on an actuated steering wheel. The driver and automation jointly exert torques thereby negotiating the final steering wheel angle [129]. Several studies with different HSC implementations have shown that compared to manual driving, driving with HSC increases safety margins [67] and decreases control activity (i.e., steering wheel reversals) [137].

However, these improvements come at a price of increased driver torque magnitude [31], which is attributable to a misalignment of intent between driver and HSC, generally defined as a conflict [80]. There are different ways this misalignment can take place, resulting in various types of conflicts, where the critical element of all conflicts is when the driver experiences resistive torques [80]. Therefore, in this study, conflict occurrence is quantified by instances when the signs of torque from the driver and HSC are opposite, indicating opposing intentions and resistance. The control actions by the driver and automation must ideally complement each other, and not counteract unnecessarily. For the most HSC design philosophies, these conflicts (opposing torques) are a source of annoyance [135][137], reducing the acceptance and thereby benefits of HSC.

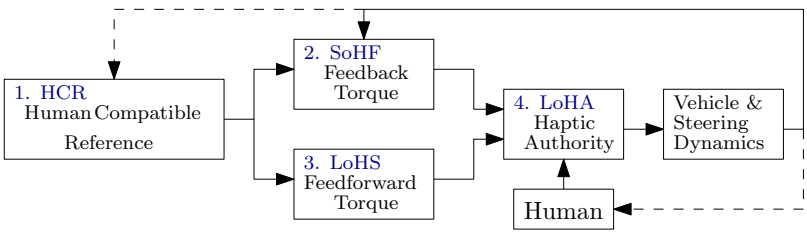

Figure 7.1: The Four-Design-Choice-Architecture design philosophy from [177].

The Four-Design-Choice-Architecture (FDCA) design philosophy for haptic shared control was proposed to reduce driver torques and conflicts and increase acceptance through its inherent control structure [177]. The design philosophy is depicted in Fig. 7.1. The haptic controller architecture exhibits an independent and separated Human-Compatible-Reference (HCR) and haptic controller part. The HCR is the controller's reference and can be generated by a human-like driver steering model or through recorded data. The essence of the haptic controller part that follows the independent HCR is twofold: a pure feedback component to provide compensatory torques reducing errors with the HCR (SoHF), and a pure feedforward component to provide pursuit or anticipatory torques for curve negotiation (LoHS). This decoupling of the error correction (SoHF) from the support and guidance to follow the controller's reference (LoHS), is expected to improve acceptance compared to only feedback [177]. These two loops are summed together contributing to the total torque. How much influence the controller has on the final steering wheel angle is determined through the Level

of Haptic Authority (LoHA), in essence, this authority is set by adjusting the stiffness setting of the steering wheel around an optimal guidance angle.

The considered FDCA offers a more flexible way to mimic how drivers control than conventional approaches, through its independent HCR. It is therefore important to note that generating an appropriate HCR is crucial. No matter how good the haptic controller, when the HCR is not suitable to a driver, the driver-automation conflicts may not be resolved.

A few methods have been explored to reduce conflicts in HSC including aligning the reference with the human's trajectory (personalisation) [31] [43] or adapting the Level of Haptic Authority to grip strength (i.e., if the human grips more, the system grips less) [161]. However, investigating the effect of adjusting the inherent controller construct on driver acceptance has not been studied yet. This study is the first to test two key aspects: 1) the effectiveness of the main design philosophy of the new haptic controller construct: the Four-Design-Choice-Architecture [177] by implementing a realisation of the architecture: the Four-Design-Choice haptic shared controller (FDC-HSC) and, 2) personalising the HCR through state-of-art control-theoretic driver steering modelling [116]. This implementation tests two types of HCRs: 1) a One-Size-Fits-All (OSFA) average driver reference and, 2) a Personalised driver reference. Both HCRs are generated through identification on real driver data using the Mars driver model [116].

This study evaluates the conflict reducing capabilities of the implemented FDC-HSC in a fixed-speed curve negotiation task. Moreover, the added benefit of personalising the HCR is compared to a 'One-Size-Fits-All' HCR. A two-phase simulator experiment is performed where the first phase is used to collect the empirical driver data needed to generate both the One-Size-Fits-All and personalised HCRs. In the second phase, four experimental conditions are tested, comprising:

1. The FDC-HSC with a OSFA HCR,
2. The FDC-HSC with a personalised HCR,
3. The conventional Meshed-HSC of [129] (taken as baseline)
4. Unaided manual driving (also taken as baseline).

The results are evaluated in terms of driver-HSC conflicts, subjective ratings and driver torque.

This chapter is structured as follows: Section 7.2 describes the implemented Four-Design-Choice haptic shared controller, along with the conventional Meshed haptic shared controller. The driver model used to generate the 'One-Size-Fits-All', and Personalised HCRs is elaborated in Section 7.3. After which, the experimental design, the setup of the simulator experiment, experimental conditions, dependent variables and hypothesis are elaborated in Section 7.4. The experimental results summarising the outcomes of conflicts, workload and subjective measures and a discussion are presented together in Section 7.5. Finally, the main conclusions are presented in Section 7.6.

7.2. HAPTIC SHARED CONTROL DESIGNS

This section describes two different types of shared control implementations that are each derived from a different shared control concept. An implementation of the novel Four-Design-Choice-Architecture (FDCA) shared control architecture from [177] is described in Section 7.2.1 and the conventional Meshed (M) shared controller of [129] is described in Section 7.2.2.

7.2.1. FOUR-DESIGN-CHOICE (FDC) HSC

The implemented FDC-HSC, shown in Fig. 7.2, is an implementation of the Four-Design-Choice-Architecture (FDCA) [177]. In the following each design choice is explained, although, only three design choices were altered and implemented in this realisation.

HUMAN COMPATIBLE REFERENCE (HCR)

As shown in Fig. 7.2, the FDC-HSC uses the HCR as its reference. The outputs of the *Modelled Driver Trajectory* from Section 7.3, are the coordinates of the HCR: the trajectory $(\vec{X}_R, \vec{Y}_R)$, the heading $\vec{\Psi}_R$ and the steering angles $\vec{\delta}_R$. The 'reference selector' block finds the index that minimizes the distance to the car's current position $(X_{car}(t), Y_{car}(t))$, then, the corresponding $(X_R(t), Y_R(t))$, $\Psi_R(t)$ are taken as input for the SoHF and $\delta_R(t)$ for LoHS. This HCR (i.e., trajectory $(\vec{X}_R, \vec{Y}_R)$, heading $\vec{\Psi}_R$ and steering angles $\vec{\delta}_R$) can be generated by a driver steering model as is done in this chapter by using the model outlined in Section 7.3, or through recorded driver data.

STRENGTH OF HAPTIC FEEDBACK (SOHF)

The Strength of Haptic Feedback block takes the HCR heading $\psi_R(t)$ and position $(X_R(t),\ Y_R(t))$ as inputs and compares these to the car heading $\psi_{car}(t)$ and position $(X_{car}(t),\ Y_{car}(t))$. The operator Δ in Fig. 7.2: calculates the Euclidean distance between $(X_R(t), Y_R(t))$ and $(X_{car}(t),\ Y_{car}(t))$. The resulting feedback errors Δs_{lat} and $\Delta\Psi$, are weighted through gains K_s and K_Ψ, summed and multiplied by gain K_{SoHF} to obtain the total feedback torque T_{SoHF}. The SoHF essentially acts as a 'spring' that pulls the steering angle to the target. The values for the gains K_s, K_Ψ and K_{SoHF} used in this chapter are given in Table 7.1 and were heuristically tuned. The feedback torque was tuned not to feel 'too strong' and at the same time avoid oscillatory behaviour, even when given an impulse driver torque.

LEVEL OF HAPTIC SUPPORT (LOHS)

The Level of Haptic Support feedforwards the HCR's $\delta_R(t)$ in an open-loop manner. To obtain the feedforward torque $T_{LoHS}(t)$, the $\delta_R(t)$ is multiplied with a gain K_{LoHS}. Note that: although $\delta_R(t)$ is a part of the HCR, one can also generate it by taking road curvature processed by inverse vehicle dynamics to obtain the required steering angles for *centreline driving*. This feedforward torque provides a distinctively different type of guidance as even if there is no deviation from the HCR trajectory you are still supported. Note the distinction between these two haptic support components, while $T_{SoHF}(t)$ is only non-zero when there is a deviation from the reference, $T_{LoHS}(t)$ is independent of the car lateral or heading error, and only depends on where the driver is along the HCR.

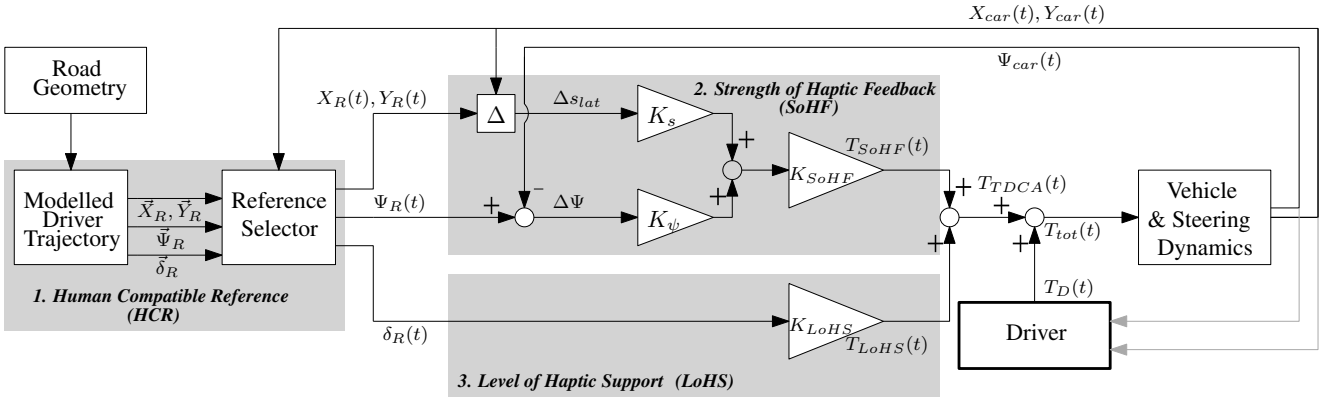

Figure 7.2: The implemented Four-Design-Choice (FDC) Haptic Shared Controller, a novel realisation of the 'Four-Design-Choice-Architecture' of [177]. The Human Compatible Reference (HCR) is the FDC haptic controllers reference that comprises of a trajectory ($\vec{X}_R, \vec{Y}_R$), heading $\vec{\Psi}_R$ and steering angles $\vec{\delta}_R$. This reference can be generated by a driver model (as is done in this chapter), driver trajectory classifier, or be a play-back of recorded data. The Strength of Haptic Feedback (SoHF) regulates the feedback strength of the controller based on the lateral error to the trajectory reference Δs_{lat} and the heading error $\Delta\Psi$. Note the Δ operator signifies the Euclidean distance between, in this case, ($X_R(t), Y_R(t)$) and ($X_{car}(t), Y_{car}(t)$). The Level of Haptic Support (LoHS) guides the driver with anticipatory feedforward torques, determined by the controller reference's steering angles $\delta_R(t)$.

Table 7.1: Haptic Shared Control parameter values used for the implemented FDC-HSC and M-HSC

Gain	FDC-HSC	**Gain**	M-HSC
K_s	0.05 [N]	D	0.08 [N]
K_Ψ	0.03 [Nm/deg]	P	0.9 [Nm/deg]
K_{SoHF}	1.5 [-]	K_f	2.0 [-]
K_{LoHS}	0.45 [Nm/rad]	t_{LH}	0.7 [s]

The gain K_{LoHS} (0.45 Nm/rad) given in Table 7.1 was heuristically tuned. Between the reference steering angles $\delta_R(t)$ and the actual steering angle realised on the wheel $\delta_s(t)$, there is the K_{LoHS} gain that transforms $\delta_R(t)$ into $T_{LoHS}(t)$ and the steering wheel dynamics then transforms $T_{LoHS}(t)$ into $\delta_s(t)$. In terms of mapping between $\delta_R(t)$ and $\delta_s(t)$ w.r.t magnitude, a factor K_{LoHS}/K_w is established, where K_w is the stiffness of the steering wheel. Therefore when $K_{LoHS} = K_w$, there is a 1:1 *magnitude mapping* between $\delta_R(t)$ and $\delta_s(t)$. As found (later) in Chapter 8, the stiffness of the steering wheel is 0.49 Nm/rad, so with the K_{LoHS} gain used in this chapter, the magnitude mapping corresponds to a 0.45/0.49 = 92 % share of control, i.e. the controller performs 92% of the steering demand set by δ_R.

LEVEL OF HAPTIC AUTHORITY (LOHA)

The LoHA determines the authority the automation has over the final control output. This is done by adjusting the stiffness of the steering wheel around a particular optimal steering angle. A high authority setting results in a very rigid system, where the contribution of the driver is diminished, i.e. due to a higher steering wheel stiffness, the driver needs to exert substantially more torque to change the steering wheel angle. Above a certain threshold, the system can become so rigid that it becomes impossible to move the steering wheel. This can be particularly useful in critical situations, as investigated in [196]. In this investigation the LoHA was not implemented, as reflected in Fig. 7.2, and thereby not varied. One could say that it was set to a value of 1 because the standard steering wheel stiffness was being used with no adaptation to a particular optimal steering angle.

7.2.2. MESHED (M) HSC

The Meshed Haptic Shared Controller (M-HSC) as shown in Fig. 7.3, is named as such because the strength of the generated haptic feedback is *inherently lumped* with the driver model structure that generates the trajectories. In other words, one cannot change the haptic feedback or reference (trajectory) independently from one another. Typically, the input for a meshed controller is the road centerline reference, obtained from selecting a road reference position and heading that is t_{LH} s ahead of the closest road point to $X_{car}(t), Y_{car}(t)$. The resulting look-ahead reference position $X_R(t+t_{LH}), Y_R(t+t_{LH})$ and road heading $\Psi_R(t+t_{LH})$ are then compared to the predicted vehicle states $(\hat{X}_{car}(t+t_{LH}), \hat{Y}_{car}(t+t_{LH}))$ and $\hat{\Psi}_{car}(t+t_{LH})$, respectively.

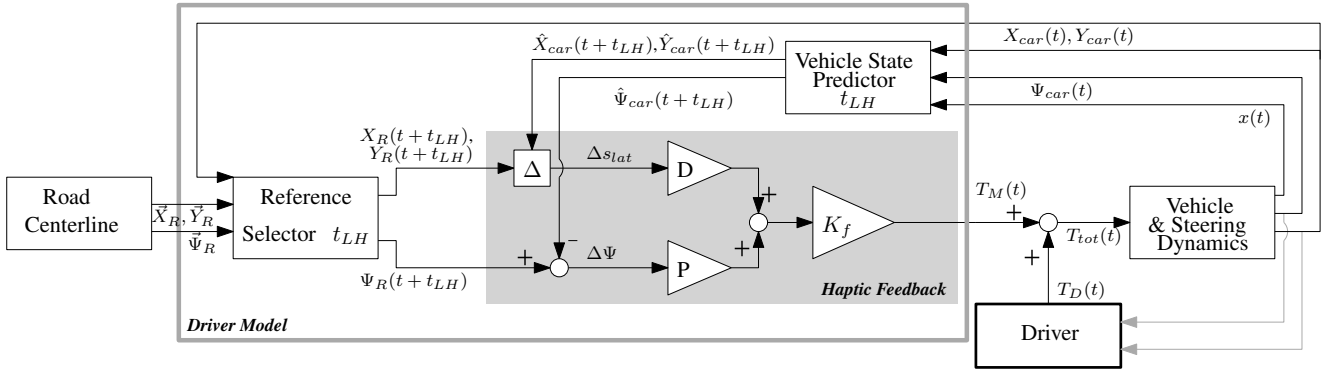

Figure 7.3: The implemented Meshed (M) Haptic Shared Controller previously applied in [129]. The control structure takes the road centerline ahead $X_R(t+t_{LH}), Y_R(t+t_{LH})$ and road heading $\Psi_R(t+t_{LH})$ at lookahead time t_{LH}, as reference for the pure feedback controller. The trajectory and heading references are compared to the predicted vehicle position $(\hat{X}_{car}(t+t_{LH}), \hat{Y}_{car}(t+t_{LH}))$ and heading $\hat{\Psi}_{car}(t+t_{LH})$, at lookahead time t_{LH}. The position and heading feedback are independently tuned with gains D and P, respectively.

The prediction is made for a constant steering wheel input with initial conditions: the bicycle model states $x(t)$, car heading $\Psi_{car}(t)$ and position $(X_{car}(t), Y_{car}(t))$, iterating the bicycle dynamics t_{LH} s ahead. Two independent feedback loops that correct for the predicted errors Δs_{lat} and $\Delta\Psi$ are weighted separately by gains D and P before being multiplied by the feedback gain K_f, to result in the total meshed torque $T_M(t)$. The values for the gains used in this controller are given in Table 7.1 and taken from [129].

Framing this controller within the FDCA framework, it is like having an HCR that is the road centerline t_{LH} s ahead, with a quickened (controlling on lookahead time t_{LH} rather than current time t) SoHF block, without any LoHS and with the same LoHA as implemented in FDC-HSC.

7.3. CONTROL-THEORETIC DRIVER MODEL

In this chapter, the HCRs were realised by offline model identification, using a modification of the Mars driver model [116], which was assessed in Chapter 3. This model was used instead of a more 'descriptive' model due to the chronology of this thesis; i.e., this chapter was completed before Chapter 3 was started.

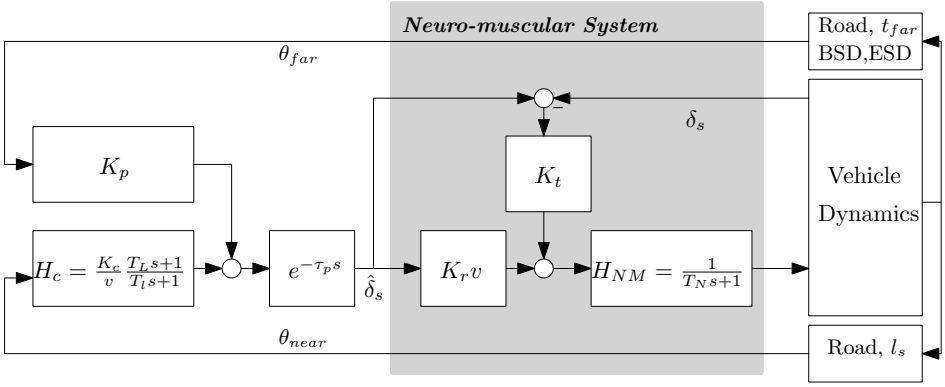

Figure 7.4: This driver-model is used to generate the references and is based on the work from [116]. The near and far angles are the input in the driver-model and Kc and Kp determine their influence respectively.

A modified driver model of [116] illustrated in Fig. 7.4, combines feedback and feedforward loops, assuming that driver's act upon a near angle θ_{near} that is compensated for and far angle θ_{far} that is anticipated. The components in the model include the visual compensation (K_c, T_l, T_L) and anticipation (K_p), processing delay (τ_p), angle to torque coefficient (K_r), gain of the stretch reflex (K_t) and neuromuscular time constant (T_N). Out of these parameters, only K_c and K_p are considered degrees of freedom because these parameters influence the resulting trajectory most.

As opposed to [116], which is assessed in Chapter 3, small-angle approximations are not used in the evaluation of the model in this chapter, resulting in a body-to-global-transformation (Euler rotation matrices are used instead of linear approximations of s_{lat}). This means that an actual calculation on a physical road is executed

every time the world is transformed around the body of the car. Compared to small-angle approximations and linearisation, this is more in line with how an autonomous vehicle senses the world at every time step. Finally, an alternative non-linear method of calculating θ_{near} and θ_{far} is used.

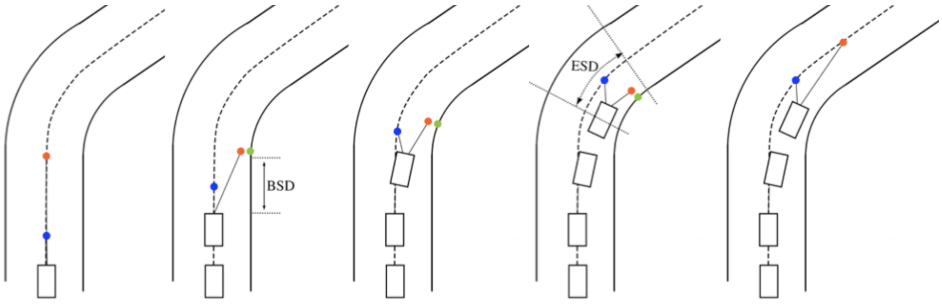

Figure 7.5: Visual representation of the near point (blue dot) and far point (orange dot) over the length of the curve, the tangent point is indicated with the green dot. The steering initiation distances before the begin and end of the curve are indicated with respectively BSD and ESD. In the most right plot t_{far} determines how far in the future the far point is located.

Here θ_{near} is determined as the physical angle (no approximations) the car makes to the near point: a point that is l_s m in front of the car on the road centerline. The θ_{far} angle is computed using the road data and three additional parameters: BSD, ESD and t_{far}. The working principle of these additional parameters is visualised in Fig. 7.5. The lookahead-time t_{far} acts as a searching horizon, looking to find a tangent point, and, when found, the far point becomes a point closely offset from the tangent point, a reachable target point [30]. When not found, the far point is a point t_{LH} ahead on road centre. The far angle θ_{far} is then the angle the car makes to the far point (equivalent to the target point). The target point can be geometrically detected 5s before curve entry; however, this is not when we would want the car to start steering into the curve. Hereby, the Begin Steer Distance (BSD) and End Steer Distance (ESD) are both threshold values that control initiation and halting of steering behaviour before curve entry and after curve exit.

7.4. EXPERIMENT

7.4.1. CONTROL TASK

Subjects performed a curve negotiation (lateral control) driving task in a fixed-base simulator. They were asked to drive with a fixed speed of 24 m/s on a road of width 3.6 m. Subjects negotiated five left and five right clothoidal curves with 240 m straight sections in between curves. The curve radius was 300 m with a curve length of 108 m including two clothoidal transitions of 18 m at beginning and end, with a total road length of 3.7 km.

A heavy single-track sedan of 1.8 m wide was used to visually simulate the vehicle, with a vehicle dynamics identical to previous investigations [129] approximated by a bicycle model.

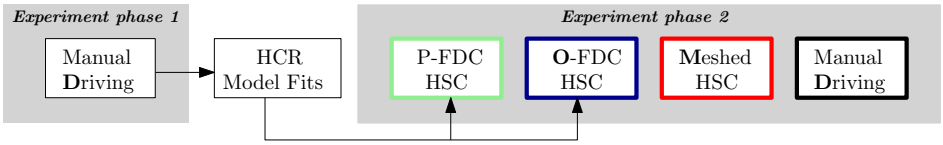

Figure 7.6: Two-phase experiment design. The first phase comprises manual driving, where the individual drivers' trajectory behaviour was recorded. This data was used to generate the pHCR and oHCR. In the second phase four different experimental conditions were tested: **P** the personalised HCR with FDC-HSC, **O** the 'One-Size-Fits-All' HCR with FDC-HSC, **M** the Meshed HSC and **D** the Driver only (without HSC).

7.4.2. APPARATUS

The driving task was performed in the fixed-based driving simulator at the Human-Machine-Interaction Laboratory at Delft University of Technology. The scenery was presented using three LCD projectors covering a horizontal field-of-view of 180 degrees and a vertical field-of-view of 40 degrees, with an update rate of 50Hz and an image generation delay of 10 ms.

A MOOG FCS Ecol8000S Actuator running at 2500 Hz was used for generating haptic torques on the steering wheel. Its stiffness was set to 4.2 Nm/rad over the complete deflection range: steering wheel inertia was set to 0.3 Nm/rad, and the damping coefficient was 2 Nms/rad.

7.4.3. EXPERIMENTAL CONDITIONS

This study investigates one independent variable, namely: the type of haptic shared controller, comprising four levels:

D No shared control, i.e., unaided *D*river steering behaviour.

M The conventional *M*eshed HSC

O The newly implemented FDC HSC with a '*O*ne-Size-Fits-All' HCR (*oHCR*)

P The newly implemented FDC HSC with a *P*ersonalised HCR (*pHCR*)

7.4.4. EXPERIMENTAL-SETUP AND PROCEDURE

As shown in Fig. 7.6, the experiment had two phases. The first phase was used to measure two controller references: 1) the average driver behaviour across all participants (One-Size-Fits-All) and, 2) the average driver behaviour per participant across 5 curves (Personalised). These references were used to generate oHCR and pHCR, explained in Section 7.4.6.

In the second experiment phase, the experimental conditions elaborated in Section 7.4.3 are presented to the driver in randomised order, following a Latin square design. During both experimentation phases, the participants were given a familiarisation run of 160 s before collecting data for each condition. After each run (comprising 180 s) during the second phase, the subjects were asked to fill out subjective ratings using a Van Der Laan acceptance questionnaire [176] to assess the satisfaction and the usefulness of the guidance subjectively.

Table 7.2: Driver model parameter ranges for HCR fitting along with the right and left curve fit values.

Parameter	Range	Right Curve Fit	Left Curve Fit
K_p [-]	2:0.125:3.25	2.5	2.375
K_c [ms^{-1}]	0:1:5	4	0
BSD [m]	2.5:2.5:25	12.5	5
ESD [m]	40:5:55	55	40
t_{far} [sec]	2:0.5:3.5	3.5	3.5

7.4.5. SUBJECTS AND INSTRUCTIONS

The experiment was performed by sixteen subjects between the age of 23 and 28 years (average of 26.5 years), all owned a valid driver's license. They were instructed to drive as they normally would and to hold their hands on the steering wheel at a "ten-to-two" position.

7.4.6. FITTING THE HCR

Two types of HCRs are considered in this chapter: the 'One-Size-Fits-All' reference (oHCR) and the personalised reference (pHCR). Both HCRs are generated by first estimating the model parameters that best describe the average driver and individual drivers. Then those model parameters were used to generate the model response. The model used to generate the HCR is the Mars model defined in Section 7.3. Two types of reference data (in the lateral position domain s_{lat} and steering angle domain δ_s) from phase 1 of the experiment are used for identification: 1) the average reference of the sixteen drivers for oHCR and, 2) the average personal reference for pHCR (where only data from a single driver is used). Note that the left and right curves were treated separately.

These trajectories are used to find the parameter vector P (K_c, K_p, ESD, BSD and τ_{far}) that results in the best fit. The driver modelling outputs are generated by using the road defined in Section 7.4.1 as input with a parameter range outlined in Table 7.2. The best fit of the *lateral position* s_{lat} trajectories, was optimised for using a cost function outlined in Eq. (7.1). This optimisation was constrained within the range defined in Table 7.2 to keep the identification results within realistic boundaries.

$$\hat{P} = \arg\min_{P} \sum_{i=1}^{n} (\hat{s}_{lat}(P) - s_{lat}(i))^2 \tag{7.1}$$

The parameters of the driver steering model considered as constants are taken from [116]: T_l = 1s, T_L = 3s, τ_p = 0.03s, K_r = 0.3, K_t= 0.5, T_N = 0.1s and l_s = 5m.

Fig 7.7 illustrates the oHCR fit results. The RMSE between the average trajectories of all drivers and the oHCR reference in the lateral position domain is 0.04 m for right curves and 0.17 m for left curves during the curve part. In terms of the Variance Accounted For (VAF) (a model fit validation metric, where 100% means the model can describe 100% of the data), the quality of fit for right curves is 95% and for left curves it is 0%. This low percentage stems from the VAF giving offsets like the one observed for the left fit, a high penalty.

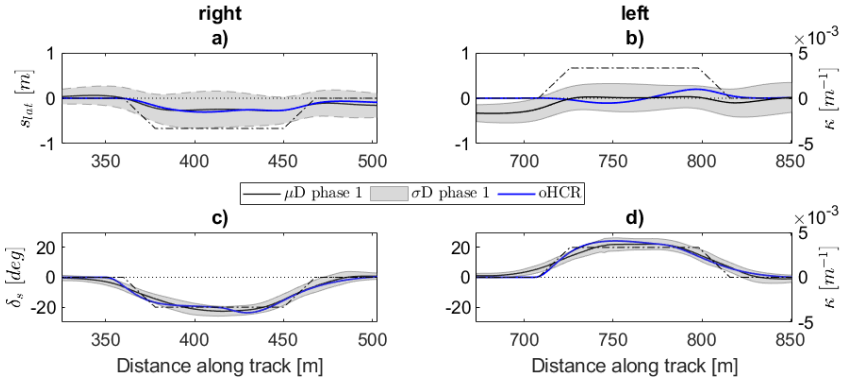

Figure 7.7: This figure illustrates the average driver behaviour in black and the oHCR in blue. The measured lateral position s_{lat} is shown in the first row and steering angle δ_s time traces in the second row. The presented driver data is averaged over all participants during phase 1, with the curvature profile κ on the right axis.

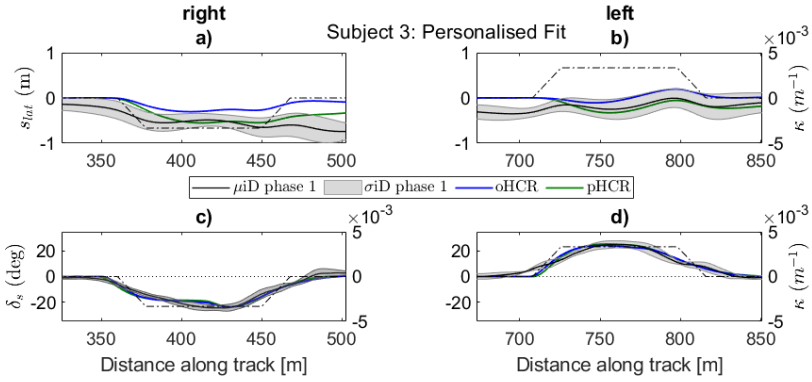

Figure 7.8: This figure illustrates the average behaviour of Subject 3 in black, the oHCR in blue and the pHCR of Subject 3 in green. The measured lateral position s_{lat} is illustrated in the first row and steering angle δ_s time traces in the second row. The presented data of Subject 3 is averaged over 5 runs during phase 1, with the curvature profile κ on the right axis.

Fig. 7.8 illustrates the pHCR fit results for Subject 3 only. The RMSE between this driver and the oHCR is 0.37 m for right curves and 0.22 m for left curves, whereas the RMSE for the pHCR is 0.19 m for right curves and 0.17 m for left curves. In terms of the VAF for the oHCR, 49% is achieved for right curves and 0% for left curves. For the pHCR, 85% is achieved for right curves and 36% for left curves, which is still not good according to the descriptiveness criteria defined in Chapter 3. Nevertheless, the pHCR has a better fit to the individual subject data than the oHCR. So, this also means that the personalised HCRs can be considered 'limited' in terms of the level of personalisation, outlined in the Introduction Fig. 1.5.

The improvement in fit results for pHCR is generally reflected in Fig. 7.9, which illustrates the VAF values for the oHCR and pHCR, per subject and curve direction. The right curves generally obtain better fits than left curves, even with pHCR. This is predominantly because the considered model does not account for any prepositioning behaviour. As was found in Chapter 5, the prepositioning on left curves is significantly more pronounced than for right curves, with an average prepositioning displacement Δy_{max} increase of 0.08 m. This may stem from drivers exhibiting a natural bias on the road. In phase 1, the average road bias is -0.1 m, and the average Δy_{max} is 0.1 m towards the outer part of the curve. This resulted in only minor offsets for right curve entries and a -0.2 m offset for left curve entry. Due to the driver model's inability to account for this, an in-balance in model fit performance is evident.

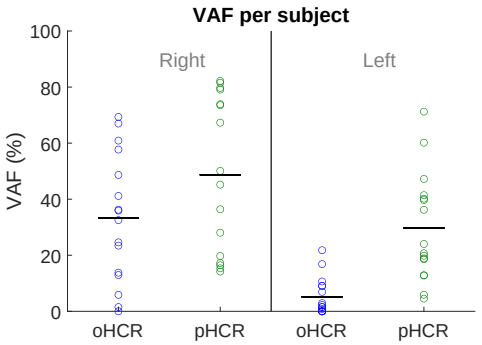

Figure 7.9: The VAF of the manual driving behaviour of each subject compared to oHCR and thier personalised pHCR, per curve direction

7

7.4.7. DEPENDENT VARIABLES

The dependent variables are used to evaluate both objective and subjective acceptance of the driver to the different HSC implementations outlined in Section 7.4.3. The considered dependent variables are:

OBJECTIVE ACCEPTANCE MEASURES

Objective acceptance is quantified using two metrics: driver torque, and conflict occurrence, described in the following:

1. The physical workload is quantified by the mean driver torque applied (T_D) in the curve.
2. The conflict measure in this study is quantified by *occurrence of conflicts*, determined when the driver torque and the HSC torque are opposite in direction. For a value of 0, the HSC torque is aligned with the driver.

$$O_c = \begin{cases} 0 & \text{if } sign(T_D) = sign(T_{HSC}) \\ 1 & \text{if } sign(T_D) \neq sign(T_{HSC}) \end{cases} \tag{7.2}$$

The occurrence of conflict O_c is evaluated as a percentage of time the HSC conflicts with the driver in the curve phase, i.e. T_{HSC} has an opposing direction and thereby sign to T_D.

SUBJECTIVE ACCEPTANCE MEASURE

To assess the subjective acceptance of the system, the questionnaire from Van der Laan [176] is used, which uses two dimensions: a usefulness scale and a satisfaction scale. The questionnaire consists of nine items; each is graded with a score between -2 to +2:

1. Useful - Useless
2. Pleasant - Unpleasant
3. Bad - Good
4. Nice - Annoying
5. Effective - Superfluous
6. Irritating - Likeable
7. Assisting - Worthless
8. Undesirable - Desirable
9. Raising Alertness - Sleep-inducing

The usefulness score is determined through the average of items 1, 3, 5 and 7. The satisfactory score is determined by the average of items 2, 4, 6 and 8.

7.4.8. STATISTICS

To evaluate the statistical significance of the dependent variables defined in Section 7.4.7, a two-way repeated-measures ANOVA was performed taking the controller type and curve direction as factors. Post-hoc pairwise comparisons were also conducted.

7.4.9. HYPOTHESES

Five hypotheses were formulated for this experiment:

H.I A decrease in the occurrence of conflicts $O_{conflict}$ with the oHCR FDC-HSC compared to M-HSC.

H.II A decrease in the occurrence of conflicts $O_{conflict}$ with the pHCR FDC-HSC compared to oHCR FDC-HSC.

H.III A decrease in driver torque T_D with both the oHCR and pHCR FDC-HSC compared to both the M-HSC and manual driving.

H.IV Higher subjective usefulness and satisfaction scores for the oHCR FDC-HSC compared to M-HSC.

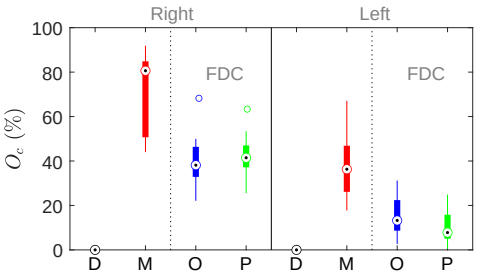

Figure 7.10: Mean absolute occurrence of conflicts during curve negotiation. D - unaided *D*river, M - *M*eshed-HSC, O - FDC-HSC with *o*HCR and P - FDC-HSC with *p*HCR.

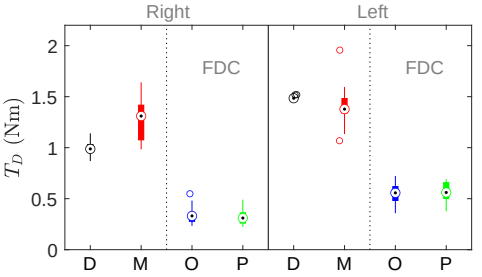

Figure 7.11: Mean driver torque during curve negotiation. D - unaided *D*river, M - *M*eshed-HSC, O - FDC-HSC with *o*HCR and P - FDC-HSC with *p*HCR.

H.V Higher subjective usefulness and satisfaction scores for the pHCR FDC-HSC compared to oHCR FDC-HSC.

H.I and H.IV are hypothesised because it is expected that the inherent control construct of the FDC shared-controller reduces conflicts and thereby improves acceptance ratings. H.II and H.V are hypothesised because it is expected that personalising the HCR will increase acceptance and thereby reduces conflicts compared to a One-Size-Fits-All reference. H.III is hypothesised because it is expected that the feedforward torque contribution of the FDC controller would be accepted and consequently decrease the driver's input.

7

7.5. RESULTS AND DISCUSSION

Conflicts between human and HSC, quantified by the occurrence of conflicts O_c shown in Fig. 7.10, show a significant effect for HSC type ($F(3,10) = 212$, $p<0.01$) and curve direction ($F(1,10) = 84$, $p<0.01$). The pairwise comparisons show a significant decrease for both curve directions in O_c for O compared to M, with an average of 27% ($p<0.01$). There is also a significant decrease (in both curve directions) found in O_c, for P compared to M, with an average of 28% ($p<0.01$). However, no significant difference between O and P is found, see Fig. 7.10. Hereby, we find a significant decrease in conflicts for the FDC-HSC (both P and O) compared to Meshed-HSC for both right and left curves ($p<0.01$) with an average reduction factor of 2.3. This accepts H.I and rejects H.II.

The driver torque T_D shown in Fig. 7.11, shows a significant effect for HSC type ($F(3,10) = 425$, $p<0.01$) and curve direction ($F(1,10) = 96$, $p<0.01$). The pairwise comparisons show a significant decrease in T_D for both curve directions when comparing O and M, with an average of 0.9 Nm ($p<0.01$). There is also a significant decrease (with both curve directions) found in T_D for P compared to M, with an average of 0.91 Nm ($p<0.01$). Again, no significant difference between O and P is found. Therefore, driver torque with the FDC-HSC (both P and O) significantly decreases with respect to both M-HSC and manual driving ($p<0.01$), with average reduction factors of 3.2 and 2.8, respectively. As a result of this, H.III is accepted. To steer through a given curve a certain

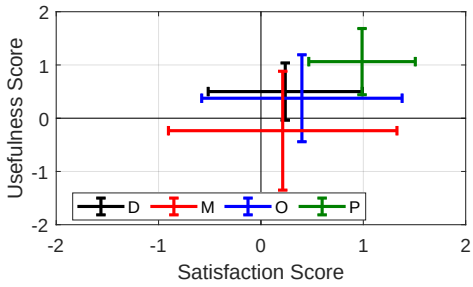

Figure 7.12: Mean and STD of the Usefulness and Satisfaction score from 'Van der Laan' questionnaire. D - unaided *D*river, M - *M*eshed-HSC, O - FDC-HSC with *o*HCR and P - FDC-HSC with *p*HCR.

amount of torque is needed, the fact that driver torques were much smaller during FDC-HSC reflects not only that the FDC is contributing with feedforward torque, but also that the driver accepts this torque. If the driver did not accept these guidance torques, he would be fighting the system, which means not only a rise in conflict but also larger driver torques. Moreover, there is a significant increase in driver torque with the Meshed-HSC compared to manual driving during the right curves ($p<0.01$), in line with previous findings [129].

Figure 7.12 shows the subjective Usefulness and Satisfaction scores obtained from the Van der Laan questionnaire. For Usefulness and Satisfaction, there are significant effects for HSC type: ($F(3,10) = 5.35$, $p<0.05$) and ($F(3,10) = 5.68$, $p<0.01$). From the pairwise comparisons, this significance stems from a difference between M and P only. There is no significant difference between O and P. For Usefulness there is a significant increase of on average 1.3 points with P compared to M. For Satisfaction, there is a significant increase of on average 0.8 points with P compared to M. These results indicate that H.IV and H.V can not be accepted. With a better driver model that can fully describe driver behaviour as outlined in Chapter 3, a difference between O and P may become significant. Nevertheless, there is an insignificant average increase that shows a trend between M and O and also between O and P.

As illustrated in Figs. 7.10, 7.11 and 7.14 the variability (standard deviation) of the occurrence of conflicts is significantly larger for the M than it is for both O and P ($p<0.01$). A concrete reason for this is not found yet. The significant increase in the standard deviation of driver torque ($p<0.01$) for M compared to both P and O is better understood when related to *Motor Unit Recruitment principles*. In this principle, when the driver exerts more force, more driver muscle neurons (motor neurons) are activated, which results in more muscle fibres being activated, where the variability of the strength of each of these recruited muscle fibres add up. This is clear when we understand that the amount of driver torque increases with a stronger dynamic haptic feedback [115] and that the M has larger feedback gains than both O and P, as given in Table 7.1.

From a control-theoretic point of view, a good controller follows its reference accurately. Likewise, the FDC-HSC should guide the driver to drive along its reference: the modelled average driver (oHCR) or the personalised reference (pHCR). However,

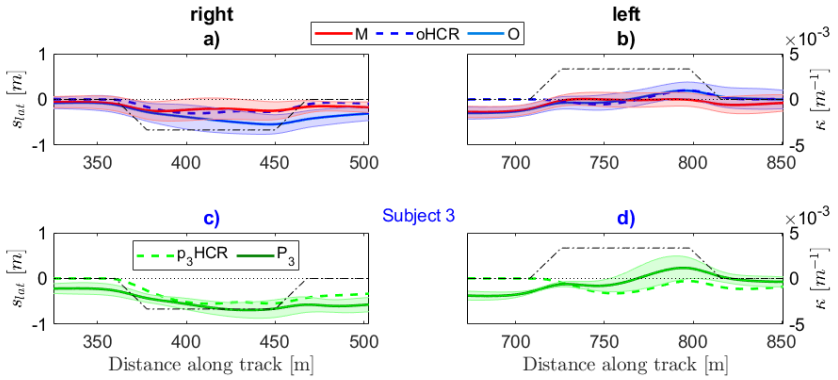

Figure 7.13: Illustration of lateral position for both left and right curves, showing the oHCR and the outcomes the the driver with M-HSC and O-HSC a)-b). An example of Subject 3 is illustrated for the Personalised case, showing this subjects' pHCR and the outcome of Subject 3 with P-HSC.

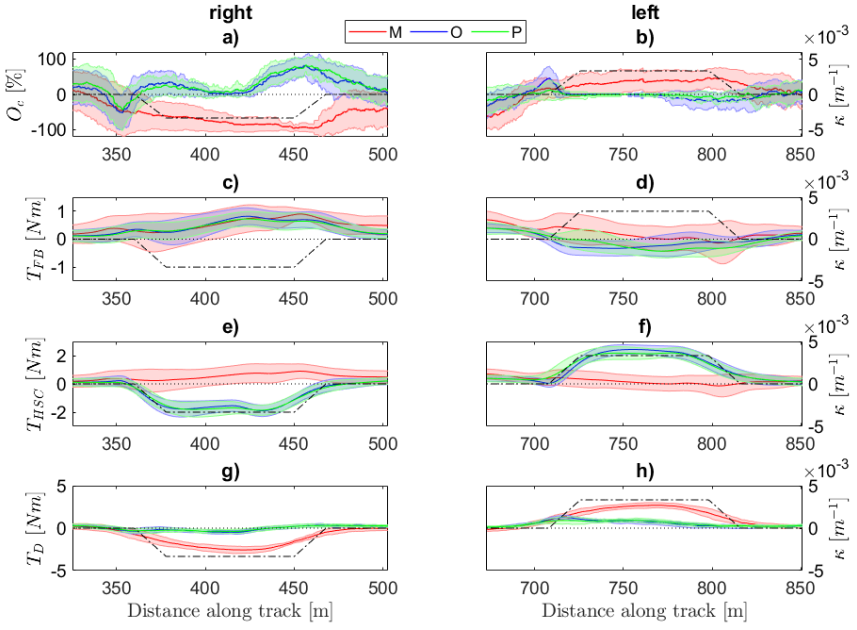

Figure 7.14: The occurence of conflicts averaged over all drivers, the haptic feedback torque T_{FB}, total HSC torque T_{HSC} and applied human driver torque T_D are shown for left and right curves, for the M-HSC, O-HSC and P-HSC respectively a)-h).

since the driver can reject the guidance by giving opposing torques, we see that following the HCR only happens when the driver accepts the guidance. From Fig. 7.13, it can be seen that for the left curves, the average of all drivers closely follows the oHCR. However, this is not found to be the case for the right curves. Consequently, we see that for the right curves there are significantly more conflicts than left curves, which is true for both O and P. This results in a moderately strong correlation between the extent to which the FDC-HSC follows its HCR in the curve and conflicts in torque (r = 0.55).

The FDC-HSC combines feedback with feedforward (both with respect to the HCR), as shown in Fig. 7.2, whereas the Meshed controller consists of a haptic feedback loop with respect to the road centre t_{LH} s ahead, as shown in Fig. 7.3. From Fig. 7.14c)-d), the magnitude of the feedback torque T_{FB} is seen to be very similar (especially for right curves) between all HSC implementations. This implies that from a torque perspective, the improvement in acceptance between M and both FDC implementations likely stems from the feedforward torque. From Fig. 7.14e)-h), it is clear that the added feedforward torque results in a largely diminished driver torque. The driver's perception of these two different types of guidance in a dynamic closed loop is very different, however, analogous to a teacher who continuously corrects each small deviation you make (feedback), to a teacher who "shows you the way" (feedforward). This is similar to the metaphor for haptic shared control: *"the instructor-student during aviation lessons"* [58] [77]. The instructor can assist the student using two different types of guidance: 1) active, by exerting torques on the system to execute manoeuvres which helps the student learn by example (feedforward), or 2) passive, by monitoring the student and intervening when he/she errs (feedback). This confirms that the *Level of Haptic Support* (i.e., the feedforward component) may be crucial for the acceptance of the system.

One would expect that right curves have fewer conflicts than left curves because the HCR is a (much) better fit for right curves for both oHCR and pHCR as shown in Fig. 7.9. However, the opposite is true: conflicts for left curves are found to be much lower. This is evident from Fig. 7.13a)-b), which shows that the average trajectory driven with O for left curves is a closer match with its reference (oHCR), implying agreement between driver and HSC. Conversely, for right curves, the HCR seems less accepted, resulting in a large adaptation of trajectory as shown in Fig. 7.13a). This unintuitive result proves that the occurrence of conflicts is not singularly dependent on the accuracy of replicating the *exact* human driver behaviour for the HCR. It is possible that the driver was more accepting of the (inaccurate) left curve oHCR and pHCR triggered guidance because the left curve entry offset means that the HCR is in the direction of curve-cutting. This allows for the driver's initial intent to curve-cut at curve-entry, which builds trust, to the extent that the drivers later follow a trajectory that does not represent their (average) driver style. This allowance and agreement with the HSC is also evident from a significant decrease in absolute feedback torque compared to right curves, shown in Fig. 7.14c)-d). Again, this is not unique for the FDC-HSC, as the M-HSC, also shows less occurrence of conflicts for left curves. Hereby, when the first torque experience matches expectations, drivers are more likely to accept and follow. If a conflict is experienced right at the start of the curve, then that sets the driver up

for constant conflict and thus low acceptance. Hence, as long as the (initial) direction of the driver intended trajectory is supported, the human driver will experience fewer conflicts.

Unexpectedly, no significant differences were found between P and O in any of the considered metrics. Nevertheless, the subjective ratings demonstrate the expected trend, with satisfaction scores increasing 0.6 scores on average with ($p = 0.3$) and usefulness scores improving 0.7 scores on average with ($p = 0.27$). The insignificant effect could be related to the minor improvement in fit between oHCR and pHCR. For oHCR, the average VAF is 19.2% whereas for pHCR it is 39.5%. This poor fit is a result of using the Mars model, whose descriptiveness (capacity to describe different driver behaviour) is only 16% (Chapter 3). The main bottleneck for this model's poor descriptiveness is the inability to describe prepositioning behaviour (driver behaviour before curve entry). Hereby, this model is unable to describe a large portion of the drivers considered, evident with only two pHCR fits achieving a VAF> 80%. Although the P condition is labelled 'personalised', the driver model's lacking ability to describe driver behaviour questions whether one can consider this condition as personalised. Therefore, to obtain a fully personalised HCR, it is recommended to use the Van Paassen model with a prepositioning extension (outlined in Chapter 6), which has a descriptiveness of 100%.

7.6. CONCLUSION

The goal of this study was to evaluate an implementation of the Four-Design-Choice-Architecture (FDCA) haptic shared controller: the FDC-HSC, in terms of occurrence of opposition torques (conflicts) between human driver and controller. A human-in-the-loop simulator experiment was conducted comparing three HSC implementations: 1) the FDC with a 'One-Size-Fits-All' controller reference (O), 2) the FDC with a Personalised controller reference (P) and, 3) a previously developed Meshed haptic shared controller (M-HSC).

The personalised FDC-HSC (P) resulted in significantly higher usefulness and satisfaction than the M-HSC implementation with average improvements of 1.3 and 0.8 scores on a -2 to 2 scale, respectively. For all metrics, there is no significant improvement between O and P. In general, the FDC-HSC (both O and P) significantly reduces occurrence of conflicts by an average factor of 2.3 (43%), and significantly reduces driver torque by an average factor of factor 3.2 (31%) compared to the M-HSC. Detailed analysis of steering wheel torque data shows that this can be attributed to two inherent design elements in the FDC-HSC: the independent feedforward (LoHS) contribution and an HCR that supports (future) curve entry intentions (i.e., the acceptance depends on whether the initial HSC torques at curve entry are perceived as resistance or as guidance). This proves that the Four-Design-Choice-Architecture is effective and has (larger) potential for acceptance with different design choice settings and a 'fully' personalised HCR.

8

EVALUATION OF PERSONALISED HAPTIC SHARED CONTROL

Good design, when it's done well, becomes invisible. It's only when it's done poorly that we notice it

Jared Spool

Current haptic shared control (HSC) implementations for curve driving suffer from user acceptance issues: there exists disagreement between haptic shared controller and driver, reflected by opposing torques on the steering wheel, a phenomenon called a haptic (torque) conflict. A promising solution to enhance driver acceptance is minimizing misalignment between controller and driver by personalising the controller's reference trajectory to that of the driver. Chapter 7 introduced the first attempt to personalise the trajectories of the FDC shared controller, however, due to the limitations of both the driver model used to generate the personalised controller references and the limited controllers ability to accurately track its reference, a satisfactory implementation of personalisation was lacking. This chapter aims to evaluate the potential benefit of fully personalised guidance to class average and centreline. Additionally, the personalisation guidances are evaluated at different levels of haptic support.

The contents of this chapter are to be published as:

Barendswaard, S., Ghys, E., Van Paassen, M.M., Boer, E.R., Abbink, D.A., Pool, D.M. 'Personalised Haptic Shared Control Approaches to Improve Driver Acceptance in Curve Driving', *IEEE Transactions on Human Machine Systems*, In preparation.

8.1. INTRODUCTION

Conventional human-automation interaction is based on trading control between human and automation, resulting in intermittent supervisory control. Continuous sharing control through haptics is an alternative approach that mitigates many of the classic 'ironies of automation' [15] that arise from traded control in highly automated vehicles [59][137], such as loss of situation awareness [122], complacency, misuse [135] and skill degradation [45]. Haptic Shared Control (HSC) is a type of Advanced Driver Assistance System that supports drivers in the steering task by providing assisting torques on an actuated steering wheel. The driver and automation share control of the vehicle by jointly exerting torques on the steering wheel or gas pedal [129]. Several studies with different HSC implementations have shown that compared to manual driving, driving with HSC increases safety margins [67] and decreases control effort [137].

Although HSC has proven beneficial in terms of improved performance (safer driving) and reduced effort in curves [129], it is also proven that current systems suffer from user acceptance issues [137]. With the existing controllers, user experiences report guidance that is perceived as too strong [45], or trajectories that feel 'un-natural' [31]. These non-human-like, non-personalised controllers result in disagreement between haptic shared controller and driver, evident from opposing torques on the control interface, a phenomenon called a haptic (torque) conflict [6]. This means that the driver fights the shared controller, which causes annoyance, mistrust and rejection of the system altogether.

Currently, most HSC systems guide drivers towards the centerline or a One-size-Fits-All path [146]. However, humans drive differently and exhibit different preferences. Driver classification studies find that there are not only numerous ways to classify different drivers [107] [179], some studies also identify over 50 clusters of driving styles [76]. To increase driver acceptance of HSC, it is essential for the automation to understand the individual driver [27] and adapt the system accordingly, i.e. personalise the system.

The personalisation of HSC for steering support has not been attempted in literature yet. However, other advanced driver assistance systems have. In general, personalisation improves the driving experience with ADAS [101], enhancing driver comfort and safety [97]. A driver's safety is found to be strongly correlated to the individual's style [143], where ADAS adaptations to a driver's style improve performance and safety [94]. Acceptance of autopilot systems is improved by predicting individual driver behaviour through driver modelling [188]. Satisfaction increases by 30% when including individualised driver behaviour models in a lane departure warning algorithm [189]. Therefore, to increase the acceptance of the driver, personalisation of HSC is a promising option.

There are, however, multiple ways in which the personalisation can be achieved with haptic shared control for steering through curves. Personalising the haptic shared controller's reference can be achieved on basis of 1) the unique individual, through driver identification [189] ('full' personalisation) or 2) on the more general 'style', through driver classification [94] ('class-average' personalisation). The high-level differences between these two different types or 'levels of personalisation' are outlined in Fig. 8.1. Moreover, specifically for the considered 'Four Design Choice' HSC introduced in Chapter 7, the personalised controller reference can be 'served to' or 'imposed on' the driver in different ways. This is determined by the 'Level of Haptic Support', which is the feedforward gain of the haptic controller. With different personalisation approaches available, it is still unknown which approach achieves the best driver acceptance, or whether acceptance changes at all. Hereby, this study investigates the effect of the *'level of personalisation'* and *'level of haptic support'* on driver acceptance.

For the sake of effective statistical analysis, two driver groups, each falling within a specific trajectory style, are considered: a) the offset-driver group and, b) curve-cutters group. The 7-class rule-based trajectory classifier from Chapter 4 defines these groups. A two-phase simu-

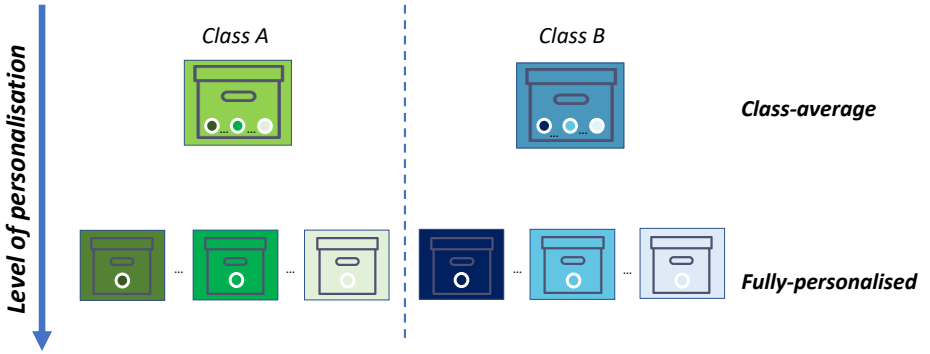

Figure 8.1: Illustration of personalisation levels. The class-average considers the average of a class as personalised reference, Class A and Class B are examples or arbitrary classes. Whereas fully personalised considers the unique individual's behaviour as reference.

lator experiment is carried out, where the first phase is needed to collect empirical driver data to select the desired driver groups and to generate the personalised controller references. During the second phase, eight different conditions are tested comprising four controller references and two levels of haptic support. The four tested controller references per driver group are: 1) the 'fully' personalised reference, 2) a 'class-average' personalised reference, 3) the centerline reference (which is the industry standard [98] [55], and baseline condition) and, 4) the class-average of the other driver group. It is hypothesised that the difference between 'full' and 'class-average' personalisation will be small in terms of acceptance, but that both would be a clear improvement compared to the centreline reference. The two levels of haptic support considered are 1) 0.5, where the controller provides 50% of the steering task (a straight forward implementation of shared control) and, 2) 0.92, where the controller provides 92% of the required steering inputs (the equivalent share used in [152] [178] [196]). Here 100% translates to a hands-free mode. The differences between these eight haptic shared control realisations are measured in terms of subjective ratings, driver torque and conflicts.

This chapter is structured as follows: Section 8.2 explains an upgrade made to the 'Four Design Choice' HSC to facilitate appropriate personalisation. Section 8.3 provides the details of the experimental design, experimental conditions and the metrics used to evaluate acceptance. In Section 8.4 the results of the experiment are presented, followed by the discussion and conclusion in Sections 8.5 and 8.6.

8.2. UPGRADED FOUR DESIGN CHOICE HSC IMPLEMENTATION

For appropriate personalisation, a HSC system that is capable to accurately track its (personalised) controller reference is essential. For this study the 'Four Design Choice' (FDC) controller

is chosen due to its greater acceptance rate compared to conventional controllers (as found in Chapter 7).

More importantly, this specific controller has shown the potential for accurate personalisation, due to the separation between the controllers reference (which includes the trajectory) and the haptic controller in the HSC architecture [177]. However, the previous implementations of the FDC HSC in Chapter 7 and in [196] [178] were, in hindsight, not able to accurately follow the set controller reference. Therefore, this study has upgraded the FDC from Chapter 7 to ensure accurate reference tracking, which is essential to test personalisation accurately.

The FDC HSC as introduced in Chapter 7 consists of four design choices as illustrated in Fig. 8.2. These are: 1) the Human Compatible Reference (HCR), which is the haptic shared controller's reference consisting of a trajectory ($\vec{X}_R(t), \vec{Y}_R(t)$), heading ($\vec{\Psi}_R(t)$) and steering angles ($\vec{\delta}_R(t)$) of a (modelled/selected) driver. 2) The Strength of Haptic Feedback (SoHF) determines the strength by which the control system corrects the driver after the vehicle deviates from the reference path. 3) The Level of Haptic Support (LoHS) is where the HCR's steering angles are feedforwarded to provide trajectory-following guidance. Essentially, it regulates how much the control system steers the vehicle in the reference direction without being affected by the car's position and orientation. This loop guides the driver with anticipatory torques to drive on the designated reference. 4) The Level of Haptic Authority (LoHA), although not implemented in this thesis, regulates the balance between the torque input of the human and the torque input of the controller by adjusting the stiffness of the steering wheel around the optimal control angle (explored in [196]).

The upgrade made to the FDC HSC for the current experiment comprises the addition of the inverse of the steering wheel dynamics $H_s^{-1}(s)$ in the LoHS feedforward loop, the third design choice, as highlighted in purple in Fig. 8.2. This compensates for the steering wheel dynamics $H_s(s)$. The reference HCR steering $\delta_R(t)$ is (pre) filtered by the inverse steering wheel dynamics, which produces the necessary reference torque $T_R(t)$ that would produce the desired reference HCR steering (for $K_{LoHS} = 1$, ideally $\delta_s(t) = \delta_R(t)$). The reference torque $T_R(t)$ is scaled by K_{LoHS} to obtain the resulting feedforward torque $T_{LoHS}(t)$.

This inverse not only facilitates accurate *controller reference* tracking, but also transforms the meaning of the parameter K_{LoHS}. Now, the parameter K_{LoHS} indicates the fraction of torque that the feedforward path uses to drive the reference. When $K_{LoHS} = 1$, then 100% of the driving task is performed by the HSC.

8.2.1. STEERING WHEEL DYNAMICS

The transfer function $H_s(s)$ of the physical steering wheel dynamics in the experimental set-up was identified in the frequency domain by assuming LTI mass-spring-damper dynamics. This was achieved through exciting the steering wheel with sinusoidal force perturbations. Ten sinusoids were tested ranging from 0.1 to 10 rad/s. The driven speed of the simulation was kept constant at 80km/h. The mass-spring-damper system considered is given in Eq. (8.1).

$$H_s(s) = \frac{1}{J_w s^2 + B_w s + K_w} \tag{8.1}$$

The identification results show that for the steering wheel used in this experiment $J_w = 0.026$ kgm^2, $B_w = 0.11$ Nms/rad and $K_w = 0.49$ Nm/rad.

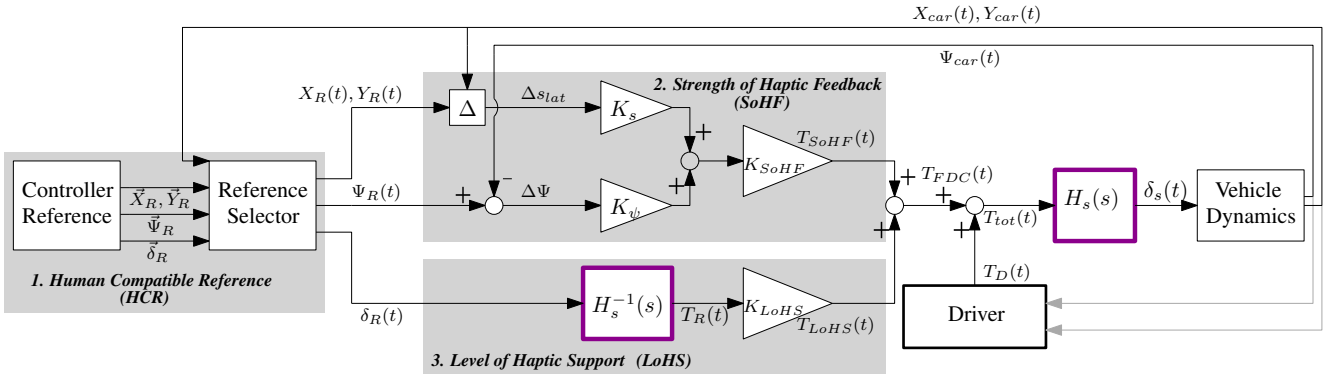

Figure 8.2: A control diagram of the upgraded Four-Design-Choice (FDC) Haptic Shared Controller from Chapter 7. The Human Compatible Reference (HCR) is the FDC haptic controllers reference that comprises of a trajectory $(\vec{X}_R, \vec{Y}_R)$, heading $\vec{\Psi}_R$ and steering angles $\vec{\delta}_R$. This reference can be generated by a driver model or be a play-back of recorded data. The Strength of Haptic Feedback (SoHF) regulates the feedback strength of controller based on the lateral error to the trajectory reference Δs_{lat} and the heading error $\Delta\Psi$. Note the Δ operator signifies the Euclidean distance between, in this case, $(X_R(t), Y_R(t))$ and $(X_{car}(t), Y_{car}(t))$. The Level of Haptic Support (LoHS) guides the driver with anticipatory feedforward torques, determined by the controller reference's steering angles $\delta_R(t)$. The (inverted) steering wheel dynamics is highlighted in purple and compensates for the steering wheel dynamics of the physical steering wheel.

8.2.2. VERIFICATION

The upgraded haptic shared controller was verified using the driving simulator apparatus defined in Section 8.3.7. Two realisations of the FDC were tested for the data from all 45 subjects driving on curve 3 from Chapter 4 (which is the same road used in this study, as described in Section 8.3.1). Subject 14, a curve cutter, is taken as an example and is plotted with the HCR as reference for comparison: with the inverse steering dynamics $H_s^{-1}(s)$ and without, as illustrated in Fig. 8.3a. The figure illustrates the tracking performance of the HCR for a left and right curve, where the curve entry and exit points are indicated by vertical grey lines. The VAF fit value across all 45 subjects, with and without $H_s^{-1}(s)$ is given in Fig. 8.3b, here a VAF of 100% means the vehicle's trajectory matches the HCR perfectly. The controller tracking performance for the upgraded FDC clearly outperforms the FDC without $H_s^{-1}(s)$. This improvement is reflected by the VAF improving from 34% to 99% when including $H_s^{-1}(s)$. This difference stems from two sources: 1) a heuristically tuned gain for K_{LoHS} equal to 0.45 Nm/rad, which should be 0.49 Nm/rad to achieve 100% of steering contribution according the identified steering wheel stiffness in Section 8.2.1 and, 2) a lag that is clearly present in the steering response between the red and blue response in Fig. 8.3a. The lag in the steering angles is small, however, the effect on the driven trajectory is larger. Therefore, the introduced upgraded form of the FDC provides a clear improvement compared to the conventional FDC implementation.

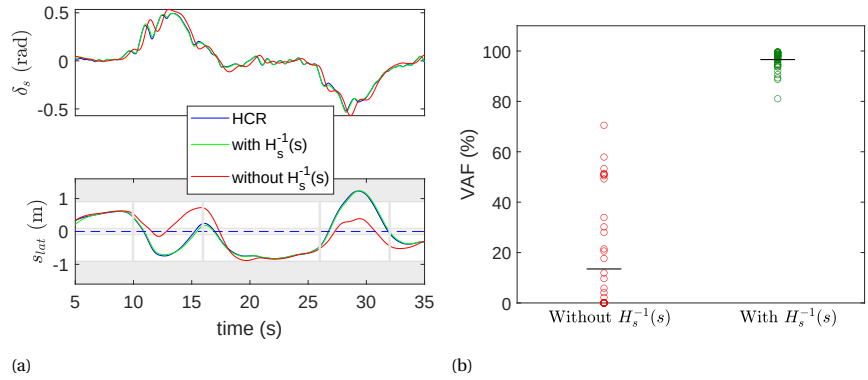

Figure 8.3: Verification of the addition of $H_s^{-1}(s)$. For these results, the controller parameter values are K_{LoHS} = 1 for the response with inverse dynamics and for the response without the values were taken from Chapter 7, K_{LoHS} = 0.45. All other parameter values (K_s, K_{psi}, K_{SoHF}) are taken from Chapter 7. On the left, a comparison is illustrated of the tracking performance of the FDC HSC with and without the inverse steering wheel dynamics $H_s^{-1}(s)$. The HCR is Subject 14 from the dataset of Chapter 4, with the same road design, i.e., the minimum radius is 204 m and velocity is 80 km/h. On the right, a comparison of the HCR tracking performance in terms of VAF across the tracking performance of 45 drivers from the dataset in Chapter 4.

8.3. EXPERIMENT

8.3.1. ROAD DESIGN

For the current experiment the driven road is identical to curve 3 in Chapter 4. The curves are designed to be clothoidal having 2 seconds for curve entry and exit transitions and 2 seconds at the maximum curvature, making the total in-curve time 6 s. The maximum curvature was 0.0049 m^{-1} (radius of 204 m), which with a fixed speed of 80 km/h means that a maximum lateral acceleration of 2.41 ms^{-2} is achieved, which is the maximum allowable lateral acceleration in

a curve to avoid tire slip or nonlinear vehicle dynamics operation [151]. The straight sections inbetween curves are 10 s each, with each curve repeated 5 times, i.e., 5 right curves and 5 left curves, in alternating order. The road is designed to be a single lane with a width of 3.6 m wide. The width of the car is 1.8 m, making the 'effective' road width 1.8 m.

8.3.2. DRIVER TRAJECTORY CLASS GROUPS

The experiment is designed to focus on two trajectory-class groups with very different curve-negotiation styles. The classifier used is the 7-class classifier, which, is used due to its enhanced consistency and its intended empirical application, as discussed in Chapter 4. The trajectory class groups are defined by the combination of the drivers' class in the right and left curves. The convention used in this chapter to refer to a class combination is RxLy, where x refers to the considered class for right curves, and y refers to the considered class for left curves. The two considered trajectory classes are: the offset-drivers (R3L5) and the curve-cutters (R2L1 & R2L2). Sixteen participants were considered for each driver group, i.e., 16 offset-drivers and 16 curve-cutters. The motivation to select these groups is elaborated in the following:

Offset-drivers → The R3L5 class combination is the most commonly occurring driver trajectory class combination in the dataset of Chapter 4, comprising 36% of tested drivers. This class of drivers performs biased inner curve negotiation on right curves and biased outer curve negotiation on left curves.

Curve-cutters → The R2L1 class combination represents the class combination of drivers that cut curves the most in the dataset of Chapter 4. This comprises normal curve-cutting for right curves and severe curve-cutting for left curves. Due to the scarcity of drivers that fall into the R2L1 class combination (see Fig. 8.4), R2L2 drivers are included in the driver group to conduct a statistically powerful experiment. The R2L2 group is seen as a relevant addition because both R2L1 and R2L2 are curve-cutting drivers. These drivers efficiently negotiate a curve by taking a trajectory that exhibits a larger radius than that of the curve. Thereby they optimise the amount they steer, which is reflected by a decrease in the driven curvature, lateral acceleration [30] and an increase in Time-to-Line-Crossing, as outlined in Chapter 5.

8.3.3. EXPERIMENT CONDITIONS

This experiment tests two factors: 1) the HCR type, comprising four levels and 2) the K_{LoHS}, comprising two levels. This results in 4 x 2 = 8 conditions in total. The considered HCRs are elaborated in the following:

CL The **C**enter**L**ine HCR, equivalent to R4L4. This HCR represents the *baseline* condition because it represents the industry standard.

OD The class-average of **O**ffset-**D**rivers, equivalent to the R3L5 HCR. This HCR is considered as 'class-average' personalisation for the offset-drivers trajectory-class group. This trajectory class combination is highlighted in blue for right and left curves in Fig. 8.5.

CC The class-average of **C**urve-**C**utters, equivalent to the R2L1 HCR. This HCR is considered as 'class-average' personalisation for the curve-cutters (R2L1 and R2L2 drivers). This trajectory class combination is highlighted in yellow for right and left curves in Fig. 8.5.

P This 'fully' **P**ersonalised HCR is defined separately for each driver and represents the level 'full' personalisation.

The second factor in the experiment is the value of K_{LoHS} which determines the share of control torques that the HSC provides the driver, i.e., with a value of 1 the driver can resort to a 'hands free' mode. Hence K_{LoHS} determines the fraction of control share in the driving task.

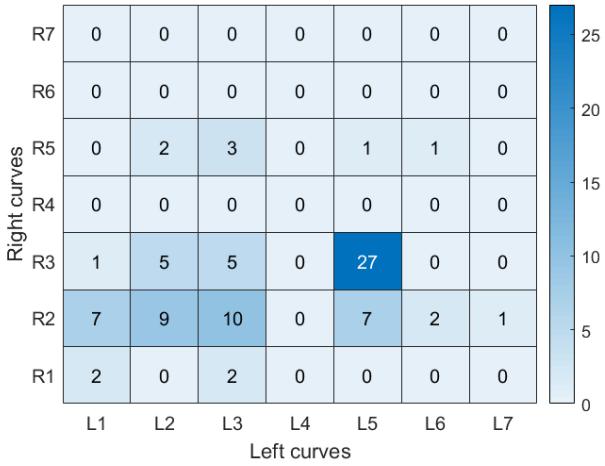

Figure 8.4: The class combination outcomes (classification of each participant for the right and left curves combined) of the manual driving experiment in the first session (n = 85). The numbers in the grid indicate the number of drivers that are classified in each class combination.

The precise share required for haptic shared control is not (yet) precisely defined as given in [5]. Hence, in this experiment two values for K_{LoHS} are tested:

- K_{LoHS} = 0.5: An implementation of haptic shared control with equal share of torque contributions from the driver and automation.
- K_{LoHS} = 0.92: The corresponding support share for the conventional FDC which translates to 92% share of control. This particular support share with the conventional FDC implementation was previously used in [152] [178] [196], and has proven to improve driver acceptance.

8

8.3.4. DETERMINATION OF HCRS

The HCRs considered in this chapter are determined in two ways: through driver modelling and through averaging (classified) recorded data. To generate the CL reference, a non-linear two-parameter driver model from [17] was used. Recorded data was not used due to the lack of driver occurrence in this class from the dataset in Chapter 4.

For the 'class-average' HCRs, i.e. OD and CC, a representative driver that best represents the average data of drivers (averaged over 12 repetitions) that fall into these trajectory-class combinations in the dataset (of the 7-class classifier outcome in Chapter 4) is selected. For the fully personalised HCRs the driven trajectories per individual participant, are averaged over the 12 curve repetitions. For all recorded averaging, the section covers averaging over 12 seconds: 3 seconds before curve entry, 6 seconds during the curve, 3 seconds after the curve. Then the HCR spanning the full road scenario (comprising of six left and right curves) is reconstructed by repeating the averaged right and left curve references along the full road scenario. The straight sections in-between the 12 seconds averaged references are smoothly interpolated. The (reconstructed) averaged reference comprises the trajectory positions ($\vec{X}_R$, $\vec{Y}_R$), heading ($\vec{\Psi}_R(t)$) and steering angles ($\vec{\delta}_R$).

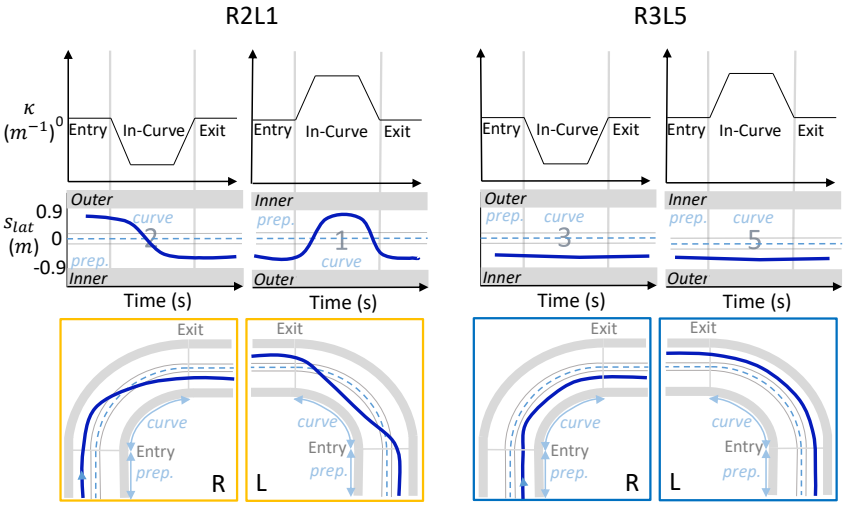

Figure 8.5: Outline of the selected trajectory classes combinations for the 'class-average' HCR. The classes are based on the trajectory classifier of [21]. The selected class combination for offset-driving (R3L5) is highlighted in blue, whereas for curve-cutting (R2L1) it is highlighted in yellow.

PERSONALISED HCRS

For the two selected driver class groups (16 offset-drivers and 16 curve-cutters), the 16 personalised and single class-average trajectories are illustrated for comparison in Figure 8.6. The class-averaged trajectories are shown in black. For some drivers, the class-average trajectory is similar to the average of their respective recorded trajectories, but still, drivers within the same class can show quite a bit of variability in driven trajectories.

8

8.3.5. EXPERIMENTAL SETUP & PROCEDURE

The experiment was carried out in two sessions. In between sessions, the data collected from the first session was processed, to select the desired trajectory class combination groups outlined in Section 8.3.2. The procedure is elaborated in the following:

FIRST SESSION

In this session, participants drove a specified road scenario defined in Section 8.3.1 without haptic support. The first session consisted of one training run comprising six left and right curves for familiarisation, and two more identical runs for data collection. This means that data was gathered for twelve left and right curves per participant. In total 85 participants performed this first session, of which only 32 were also invited for the second session.

PROCESSING FIRST SESSION DATA

The trajectories of all drivers from the first session were averaged over the 12 right and left curves separately, per driver. Consequently, each of these trajectories per driver were run through the 7-class classifier of Chapter 4. Each driver's class was determined by the most frequently occurring trajectory class (per curve direction). Only the drivers with a desired trajectory class combination (16 offset-drivers and 16 curve-cutters) were re-invited for the second session and for further processing. The recorded trajectories, heading and steering angles of the selected

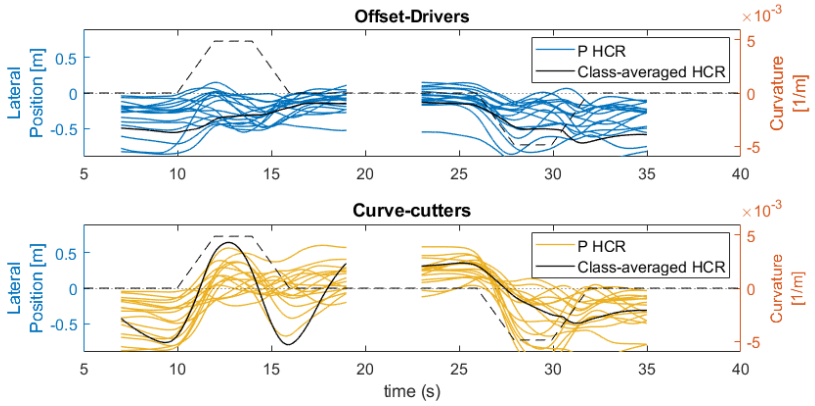

Figure 8.6: Time traces of the fully personalised HCR's and class-averaged HCR's for left and right curves, for both considered driver-groups. These are the curve-cutters falling in both R2L1 and R2L2, and offset-drivers falling in R3L5 (driving with an offset to the centreline).

drivers were averaged over each curve direction separately, to obtain the 'fine-tuned' personalised HCRs.

SECOND SESSION

In this session, the re-invited driver trajectory class groups (16 offset-drivers and 16 curve-cutters) drove the same road scenario as in the first session with different haptic guidance settings. Specifically, the four HCRs and two K_{LoHS} were altered, as defined in the Conditions Section 8.3.3.

8.3.6. PARTICIPANTS AND INSTRUCTIONS

Eighty-five subjects (13 women and 72 men) participated in the first session. The driver's ages ranged between 19 to 60 years (median = 24 years, inter-quartile-range (IQR) = 3 years) and all were in possession of a driving license for at least a year (median = 6.2 years, IQR = 3.7 years). Before the start of the experiment in the first session, the drivers signed a consent form and were informed about the goal of the research, and that based on their manual trajectory style, they might be asked to return for the second session. After classification and selection, 32 drivers returned for the second session. In the second session, participants tested all conditions in randomised order, determined by a 8x8 Latin square matrix. Each participant performed two training runs: one run of manual driving and one driving with the 'baseline' HSC condition (HCR = CL, & K_{LoHS} = 0.92). All subjects participated on a voluntary basis and no financial compensation was provided. Participants were asked to keep their hands at a "ten-to-two" position on the steering wheel and to drive as they normally would. After the participants completed each condition run, they were asked to fill in the Van der Laan [176] and CARS [93] questionnaires. The experiment was approved by the Human Research Ethics Committee (HREC) of the Delft University of Technology (application number: 991).

8.3.7. APPARATUS

The human-in-the-loop experiment was performed in the fixed-base driving simulator at the Human-Machine Interaction Laboratory (HMILab) at Delft University of Technology. The visual scenery was displayed using three LCD projectors, each covering a horizontal field-of-view of

180 deg and a vertical field-of view of 40 deg, at an update rate of 50 Hz and an image generation delay of 10 ms. A MOOG FCS Ecol800S actuator, operating at a rate of 2500 Hz, was employed for the generation of haptic torques on the steering wheel.

8.3.8. DEPENDENT VARIABLES & METRICS

The dependent variables used to analyse the drivers' acceptance over the different experiment conditions in the second session, are divided in two categories: objective metrics and subjective. The objective metrics consist of driver torque, feedback torque, conflict time and conflict torque while the subjective variables are the results of the CARS and Van der Laan questionnaires.

OBJECTIVE VARIABLES

- *Driver torque:* The driver torque is the torque exerted by the driver on the steering wheel while driving with guidance. It is an indication of the driver's effort and willingness to comply to the guidance.
- *Feedback torque:* The feedback torque is a contributor to the total FDC HSC torque, which is the summation of both feedback and feedforward torques. The feedback contribution is generally perceived as penalising by the driver, as it directly attempts to correct any deviation from the reference, as deduced in Chapter 7. A high feedback torque is an indication of disagreement, because the feedback torque is computed (just as indicated in Fig. 8.2) as a summation of the deviation in trajectory and heading compared to the HCR.
- *Conflict time:* For each time step in the experiment, it is determined whether or not the driver and FDC haptic support are in conflict. Instances of conflict arise when the driver torque T_D is opposite to the total controller torque T_{HSC} with an magnitude larger than 10% of T_{HSC} (just noticeable difference denoted by the Weber fraction [90]). The qualitative definition of conflict is augmented from Chapter 7 and is:

$$O_{conflict} = \begin{cases} 1, & \text{if } T_D \cdot T_{HSC} < 0 \\ & \textit{and} \quad T_D > 0.1 \cdot |T_{HSC}| \\ 0, & \text{otherwise} \end{cases}$$

  The conflict time gives an indication of objective acceptance. The conflict time is calculated from $O_{conflict}$ of each time step and is expressed as a percentage of time that $O_{conflict} = 1$.
- *Conflict torque:* The conflict torque equals the difference between the driver torque T_D and HSC torque T_{HSC} when a conflict is present. Conflict torque is complementary to conflict time.

These objective metrics are analysed separately per curve direction and curve phase. Curve direction (left and right) is treated separately because the designed HCRs are also different for each curve direction (with the exception to CL). Two phases of curve negotiation are considered: the prepositioning phase (analysed across 7s before curve entry [21]) and the in-curve phase. It is important to separate these because the prepositioning phase is previously postulated to affect in-curve acceptance, in Chapter 7. For both these phases the average of the above mentioned objective metrics is taken over 5 runs.

SUBJECTIVE VARIABLES

- *Controller Acceptance Rating Scale (CARS):* The CARS scale [93] is a subjective rating system, with scores ranging between 1 and 10. A score of 1 indicates that the system is untrustworthy and improvement is necessary to guarantee safe operation, whereas a 10 shows that deficiencies are rare and the driver does not need to compensate for errors.

- *Van der Laan:* The Van der Laan questionnaire [176] subjectively determines the "usefulness" and "satisfaction" of a particular haptic support setting. This is determined through nine features, four to describe the satisfaction level and five for the usefulness, which are graded by each participant with a score between −2 and 2 (the full scale is given in Chapter 7).

8.3.9. STATISTICAL ANALYSIS

For the objective measures, a between-subject two-way repeated-measures ANOVA is performed. The two factors tested are HCR type and K_{LoHS} feedforward gain, as defined in Section 8.3.3, whereas between-subject is included to account for the two driver groups tested, outlined in Section 8.3.2. Instead of a three-way ANOVA to account for curve direction, the proposed two-way ANOVA was repeated for each curve direction separately. This is because the designed HCRs are also different in each curve direction. The normality of the data was checked using the Shapiro-Wilk test, whereas sphericity is checked and if needed, corrected for using the Greenhouse-Geisser correction. To perform pair-wise comparisons, a post hoc test was done comparing the main effects with a Bonferroni adjustment. The significance of the subjective measures was tested using the non-parametric Friedman test, based on the final rating of acceptance, usefulness and satisfaction per condition. The pairwise comparisons were also tested for the non-parametric results with Bonferroni corrections.

8.3.10. HYPOTHESES

The following three hypothesis were tested with the experiment:

H.I No significant difference in acceptance exists between support based on a class-averaged HCR or based on a fully-personalised HCR.

H.II Support based on a centerline (CL) HCR will cause the highest conflict torques and will have the worst subjective ratings, leading to a low acceptance rate.

H.III A higher feedforward gain K_{LoHS} decreases conflict time and torque whilst receiving increased subjective acceptance.

H.I is based on previous research, which found that during an automated over-take manoeuvre drivers like their own overtake style played back to them as much as a similar style [66]. H.II is based on previous research that found that centerline driving is not a style exhibited by any driver [18]. Moreover, drivers seem to prefer similar driving styles to their own [66]. H.III is based on findings from Chapter 7, where it was concluded that the feedforward component in the FDC significantly contributes to conflict reduction, even though more torque is provided by the HSC system.

8.4. RESULTS

8.4.1. OBJECTIVE MEASURES

The objective measures are presented in two parts, where each part focuses on the effect of one of the two examined factors: 1) HCR type and 2) K_{LoHS} (FF) level. With respect to the presented hypothesis, H.I and H.II both relate to HCR type, whereas H.III relates to the effect of K_{LoHS}.

The objective measures presented are driver torque and conflict time. These objective measures are analysed independently per curve direction, i.e. left and right curves, and per curve

phase, i.e. prepositioning and in-curve phases. For the sake of brevity feedback torque and conflict torque are presented in Fig. D.5 and Fig. D.6 in the Appendix.

As an example, the time series of all objective measures, with HSC torque included, are illustrated in Fig. D.1 for a curve-cutter and Fig. D.2 for an offset-driver, both in Appendix D.1. For each driver the average values of twelve left and right curves are shown, where a positive curvature, κ, represents a left curve.

EFFECT OF HCR TYPE

This section focuses on the effects of HCR type in the two tested driver groups. To illustrate the observed trends without excess data clutter, only the K_{LoHS} = 0.92 is illustrated in Figures 8.7 and 8.8. The data for K_{LoHS} = 0.5 can be found in Fig. D.3 and Fig. D.4 and follows similar trends as with K_{LoHS} (FF) = 0.92.

The **driver torque** T_D during *prepositioning* for *right* curves, illustrated in Fig. 8.7 a) shows a different trend for each driver group. The highest median torque of 0.4 Nm is achieved with curve-cutting guidance for offset-drivers, whereas the highest median torque of 0.36 Nm is achieved with offset-driving guidance for curve-cutters. For both driver groups the fully personalised condition results in the lowest driver torque at an average median torque of 0.13 Nm. Statistics shows that T_D is significantly affected by HCR ($F(1.9,58.9) = 10.9$, $p<0.01$). A significant interaction between HCR and driver group is found in this phase ($F(1.9,48.9) = 12.4$, $p<0.01$). For curve-cutters, pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 0.13 Nm ($p<0.05$) and a significant decrease between OD and P with an average of 0.23 Nm ($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and CC with and average of 0.32 Nm ($p<0.01$), a significant increase between OD and CC with an average of 0.28 Nm ($p<0.05$) and a significant decrease between CC and P with an average of 0.3 Nm ($p<0.01$). For both driver groups there is no significant difference between class-average and P even though class-average has a higher median torque (supporting H.I). For both driver groups CL does not result in the highest driver torque (opposing H.II).

The driver torque during the *prepositioning* phase in *left* curves, illustrated in Fig. 8.8 a) shows no particular change in trend between driver groups. Both groups exhibit the lowest driver torque for the fully personalised condition at an average median of 0.13 Nm and similar median torques for all the other conditions at an average median of 0.3 Nm. Statistics shows that T_D is significantly affected by HCR ($F(1.4,41.4) = 5.3$, $p<0.05$). Contrary to right curves, for left curves there is no significant interaction effect between HCR and driver group in the prepositioning phase. Pairwise comparisons show that there is a significant decrease between CL and P with an average of 0.18 Nm ($p<0.01$), a significant decrease between OD and P with an average of 0.15 Nm ($p<0.01$) and a significant decrease between CC and P with an average of 0.16 Nm ($p<0.01$). Here, the lowest torque is found for P, with the other HCRs having no significant differences from eachother. This is expected because both an CC and OD preposition toward the outer part of the left curve. Both driver groups experience significant differences between class average and P (opposing H.I) and CL does not instigate the largest driver torque (opposing H.II).

The driver torque T_D during the *curve* in *right* curves, illustrated in Fig. 8.7 b) shows no particular change in trend between driver groups. Both driver groups exhibit the largest driver torque for the centerline guidance at an average median of 0.7 Nm, whereas the smallest driver torque results from both offset driving (at average median of 0.32 Nm) and fully personalised guidance (at average median of 0.35 Nm). Statistics shows that T_D is significantly affected by HCR type ($F(2.1,61.7) = 81.9$, $p<0.01$). There is no significant interaction between driver group and HCR. Pairwise comparisons show that this effect stems from a significant decrease between CL and OD with an average of 0.46 Nm ($p<0.01$), a significant decrease between CL and CC with

an average of 0.29 Nm ($p<0.01$), a significant decrease between CL and P with an average of 0.33Nm ($p<0.01$), a significant increase between OD and CC with an average of 0.16 Nm and a significant increase between OD and P with an average increase of 0.12 Nm. For offset-drivers there is a small but significant difference between class-average and P, where the driver torque is even higher for P (opposing H.I). For curve-cutters there is no significant difference between class average and P (supporting H.I). The highest median driver torque (thereby highest driver effort) is found for CL compared to the other HCRs (supporting H.II).

The driver torque during the *curve* in *left* curves is illustrated in Fig. 8.8 b). For both driver groups the highest driver torque is found for offset driving guidance at an average median of 0.6 Nm. For offset drivers the lowest driver torque is found for centreline guidance at a median of 0.28 Nm, followed by curve-cutting guidance at 0.32 Nm, an unexpected trend. For curve-cutters the lowest median torque is found for the fully personalised guidance at 0.3 Nm followed closely by curve-cutting guidance at 0.34 Nm. Statistics shows that T_D is significantly affected by HCR ($F(2.4,70.8) = 80.2$, $p<0.01$). There is a significant interaction between driver group and HCR ($F(2.3,70.2) = 12.9$, $p<0.01$). For curve-cutters, pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 0.26 Nm ($p<0.01$), a significant decrease between CL and CC with an average of 0.14 Nm ($p<0.01$), a significant decrease between CL and P with an average of 0.11 Nm ($p<0.01$), a significant decrease between OD and CC with an average of 0.39 Nm ($p<0.01$) and a significant decrease between OD and P with an average of 0.37 Nm ($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and OD with an average of 0.29 Nm ($p<0.01$), a significant increase between CL and P with an average of 0.16 Nm ($p<0.01$), a significant decrease between OD and CC with an average of 0.32 Nm ($p<0.01$) and a significant increase between CC and P with an average of 0.2 Nm ($p<0.01$). For both driver groups there is no significant effect between class-average and P (supporting H.I), nor does CL instigate the largest driver torque (opposing H.II).

The **conflict time** O_c during *prepositioning* in *right* curves, illustrated in Fig. 8.7 c), shows is a clear difference in trends between driver groups. For curve-cutters the largest conflict time stems from offset-driving guidance at median of 82% whereas for offset-drivers the largest conflict time stems from curve-cutting guidance at median of 67%. For both driver groups the lowest conflict times stem from the fully personalised guidance at median of 0%. Statistics shows that O_c is significantly affected by HCR ($F(1,30) = 9.2$, $p<0.01$). There is a significant interaction between HCR and driver group ($F(1,30) = 7.8$, $p<0.01$). For curve cutters, pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 30% ($p<0.01$) and a significant decrease between OD and P with an average of 57% ($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and CC with an average of 52% ($p<0.01$) and a significant decrease between CC and P with an average of 42% ($p<0.01$). Both driver groups achieve the lowest conflict time with P with no significant difference to class-average (supporting H.I). Here, the curve-cutters achieve significantly the highest conflict time with OD and the offset-drivers achieve significantly the highest conflict time with CC (opposing H.II).

The conflict time O_c during the *prepositioning* phase in *left* curves, illustrated in Fig. 8.8 c), shows no difference in trend between driver groups. For both driver groups the highest median conflict time is found for offset driving at around 78%, followed by centreline at around 60%, whereas the lowest conflict time is evident for the fully personalised condition at average median of 5%. Statistics show that O_c is significantly affected by HCR ($F(1.6,49.1) = 5.4$, $p<0.05$). There is no significant interaction between driver group and HCR. Pairwise comparisons show that there is a significant decrease between OD and CC with an average of 12 % ($p<0.01$), a significant decrease between OD and P with an average of 38% ($p<0.01$) and a significant decrease

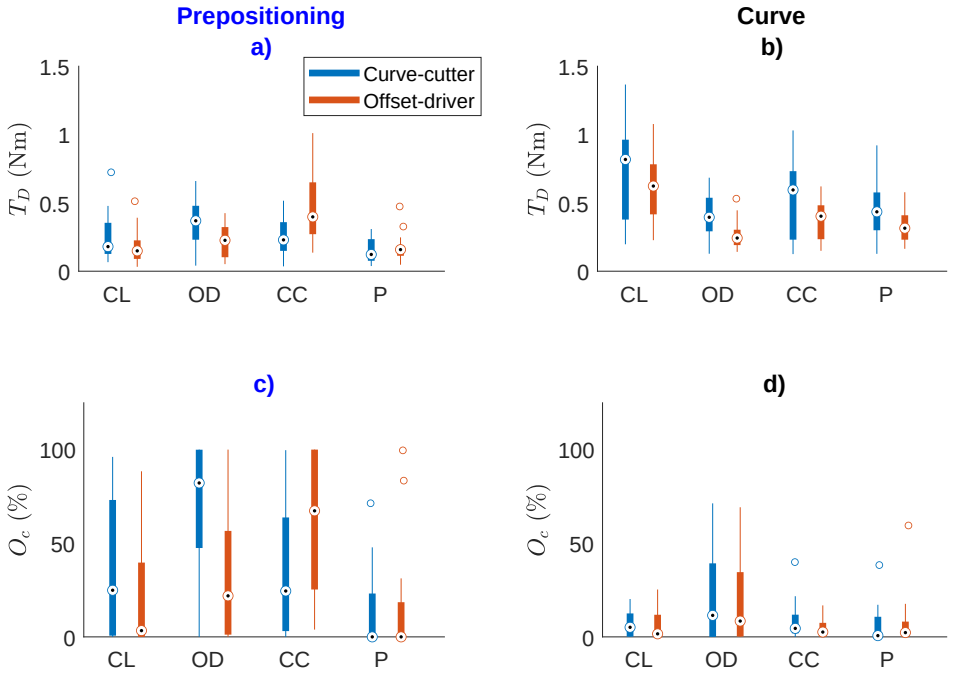

Figure 8.7: This figure illustrates the conflict occurrence O_c and driver torque T_D trends across HCR type for **right curves**, where the HCRs on the x-axis are Center-Line (CL), Offset-Driving (OD), Curve-Cutting (CC) and fully Personalised (P). The driver torque T_D is given in a) & b) and the conflict time O_c is given in c) & d). The measures in the prepositioning phase are presented in a) & c) and in the curve phase in b) & d). These results are exclusively for *Right* curves, with FF = 0.92. The overall results including FF = 0.5 included are presented in Fig. D.3.

8

between CC and P with an average of 20% ($p<0.05$). Conflict time is lowest with P with significant increase for only the class-average HCRs for both driver groups (opposing H.I and H.II).

The conflict time O_c during the *curve* in *right* curves, illustrated in Fig. 8.7 d), shows no difference in trends between driver groups. The offset-driving guidance results in the largest median torque at 12 %, with all other conditions achieving similar (minimal) conflict time. Statistics shows that O_c is significantly affected by HCR ($F(1.6,48.5) = 4.29$, $p<0.01$). There is a non-significant interaction between driver group and HCR. Pairwise comparisons show that this stems from a significant decrease between OD and P with an average of 9.1% ($p<0.05$). For curve-cutters this supports H.I, whereas for offset drivers it opposes H.I. The CL does not result in largest conflict time (opposing H.II).

The conflict time O_c during the *curve* in *left* curves, illustrated in Fig. 8.8 d) shows different trends between driver groups. For curve-cutters the fully personalised guidance achieves lowest conflict time at a median of 0% with other conditions exhibiting a similar (larger) median conflict around 10%. Offset drivers exhibit the largest conflict time with curve-cutting at median 8%, other conditions show similar minimal median conflict times around 2%. Statistics show that O_c is significantly affected by HCR ($F(2.4,70.2) = 7.5$, $p<0.01$). There is however no significant interaction between driver group and HCR. Pairwise comparisons show that there is a significant decrease between OD and P with on average 5.4 % ($p<0.01$) and a significant decrease between

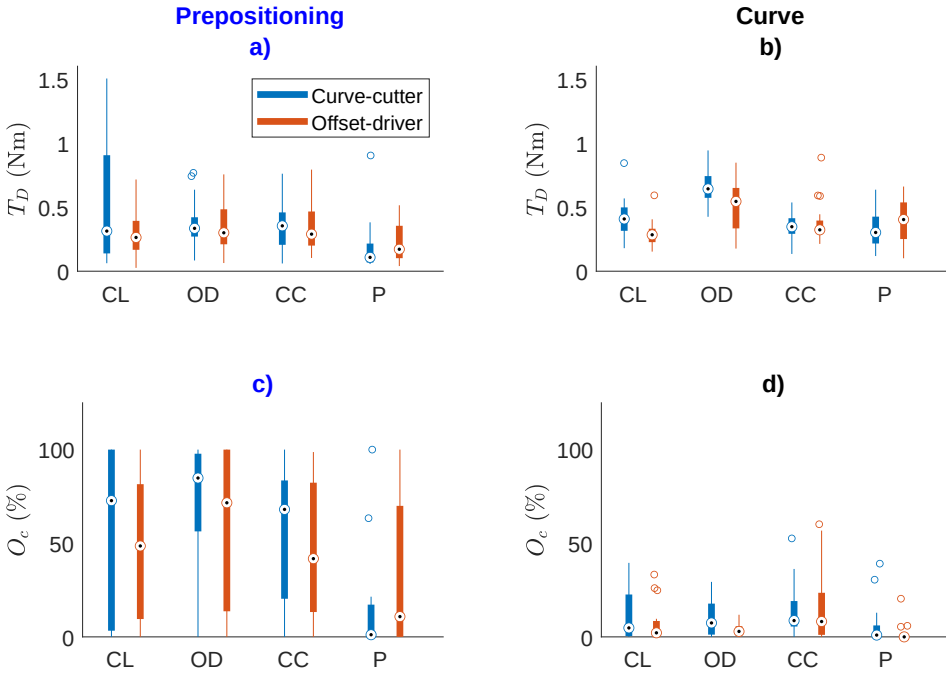

Figure 8.8: This figure illustrates the O_c and driver torque T_D trend across HCR type for **left curves** , where the HCRs on the x-axis are Center-Line (CL), Offset-Driving (OD), Curve-Cutting (CC) and fully Personalised (P). The driver torque T_D is given in a) & b) and the conflict time O_c is given in c) & d). The measures in the prepositioning phase are presented in a) & c) and in the curve phase in b) & d). These results are exclusively for *Right* curves, with FF = 0.92. The overall results including FF = 0.5 included are presented in Fig. D.4.

CC and P with on average 8.4 % ($p<0.01$). Conflict time is lowest with P with significant increase for only the class-average HCRs for both driver groups (opposing H.I and H.II).

Since the reference trajectories for right and left curves are different, clear differences in acceptance across the considered driver groups are present with curve direction. For right curves significant interaction between HCR type and driver group is only present in the prepositioning phase, whereas for left curves a significant interaction is only found for the in-curve phase. This is relatable to the similarities in the class-average trajectories during the prepositioning and in-curve phases, as can be observed in Fig. 8.5. More analysis is presented in the Discussion.

EFFECT OF LEVEL OF HAPTIC SUPPORT (FF)

This section focuses on the effects of variations in the Level of Haptic Support, determined by the K_{LoHS} gain of the FDC HSC, also known as the FeedForward (FF) gain, in the two tested driver groups. To illustrate the observed trends without excess data clutter, all the four driven HCRs at particular K_{LoHS} gain are presented together, with the grand averages indicated by a thick line, as illustrated in Fig. 8.9 and Fig. 8.10.

The **driver torque** T_D during *prepositioning* is shown in Fig. 8.9 a) for right curves and Fig. 8.10 a) for left curves. In both curve directions T_D is not significantly affected by the Level of Haptic Support. A significant change in driver torque is not present during the prepositioning phase, due to the contribution of feedforward torque being small.

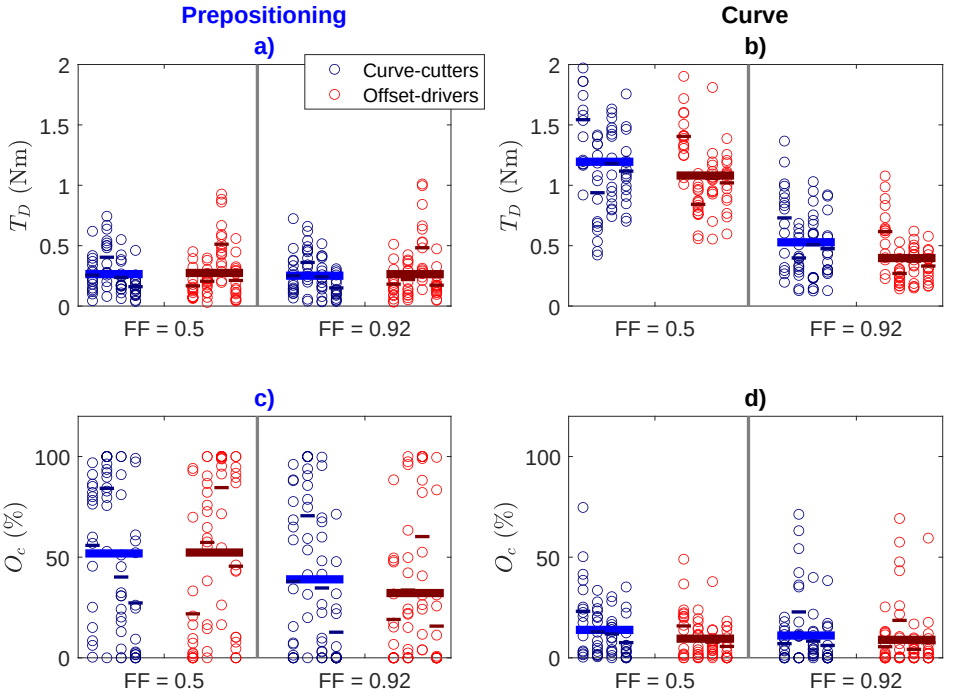

Figure 8.9: For **right curves**, the driver torque T_D is given in a) & b) and conflict time O_c is given in c) & d). The measures in the prepositioning phase are presented in a) & c) and in the curve phase in b) & d). Four columns of data are presented per driver group per FF level, each of these columns represent one of the four tested HCRs in the following order: Center-Line (CL), Offset-Driving (OD), Curve-Cutting (CC) and Personalised (P).

The driver torque during the *curve* phase, illustrated in Fig. 8.9 b) for right curves and Fig. 8.10 b) for left curves. In both curve directions T_D, shows a decrease which is also significant with increasing K_{LoHS} (FF), right – gain (F(1,30) = 1122.4, p<0.01) and left – (F(1,30) = 707.8, p<0.01). For right curves, an average decrease of 0.66 Nm is found for curve cutters and an average decrease of 0.68 Nm is found for offset-drivers (supporting H.III). For left curves, an average decrease of 0.62Nm is found for curve cutters and an average decrease of 0.61Nm is found for offset-drivers, also supporting H.III.

The **conflict time** O_c during *prepositioning* is illustrated in Fig. 8.9 c) for right curves and Fig. 8.10 c) for left curves. In both curve directions O_c shows a decrease which is significant with increasing K_{LoHS} (FF) gain, right – (F(1.6,48.5) = 4.29, p<0.01) and left – (F(1,30) = 27.8, p<0.01). For right curves, an average decrease of 12.8% is found for curve-cutters and 20.1% for offset-drivers, when increasing the K_{LoHS} gain (supporting H.III). For left curves an average decrease of 16% is found for curve cutters and 18% for offset drivers, when increasing the K_{LoHS} gain, also supporting H.III.

The conflict time O_c during the *curve*, illustrated in Fig. 8.9 d) for right curves and Fig. 8.10 d) for left curves. In both curve directions O_c is not significantly affected by the Level of Haptic Support. Unlike the effects found for HCR type, the driver torque and conflict occurrence across right and left curves with variations in K_{LoHS}, are equivalent.

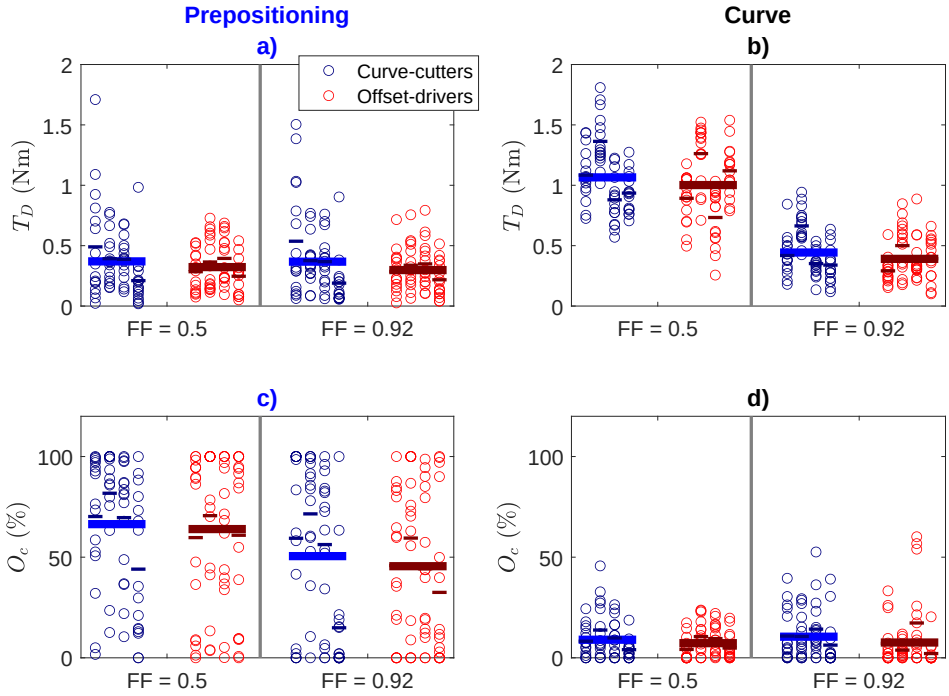

Figure 8.10: For **left curves**, the driver torque T_D is given in a) & b) and conflict time O_c is given in c) & d). The measures in the prepositioning phase are presented in a) & c) and in the curve phase in b) & d). Four columns of data are presented per driver group per FF level, each of these columns represent one of the four tested HCRs in the following order: Center-Line (CL), Offset-Driving (OD), Curve-Cutting (CC) and Personalised (P).

8.4.2. SUBJECTIVE MEASURES

This section first presents the CARS ratings and then the usefulness and satisfaction scores obtained from the Van der Laan questionnaire. The measures are presented in the order of the hypothesis, i.e. first the effects of HCR type, thereafter the effect of K_{LoHS} (FF) gain.

CARS SCORE

The CARS scores provided by all thirty-two drivers is illustrated in a cumulative distribution in Figure 8.11. The corresponding median CARS score is indicated with a vertical dashed line. The CARS rating is categorised as follows: 1 means the system is not perceived as safe, 2-4: support system is barely controllable to maintain safety, 5-7: support system is unsatisfactory, 8-10: support system is satisfactory.

HCR type While there is no significant effect for HCR type for curve-cutters, for offset-drivers there is a significant effect for HCR type ($\chi^2(3)=13.4$, p = 0.004, W = 0.15). Pairwise comparisons show that this effect resulted from the significant difference between P and CC for K_{LoHS}= 0.92 ($p<0.05$). The offset-drivers give high acceptance scores with a median of 9 for CL, OD and P HCR. The significant exception lies with CC, where a strong feedforward gain is not favored. Fig. 8.11 shows no significant differences between P and the class-average for both

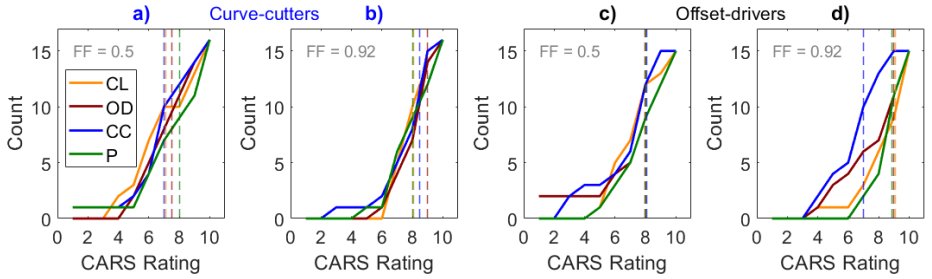

Figure 8.11: The CARS rating of the curve-cutters are given across a) to b). The CARS rating of the offset-drivers are given across c) to d). The different levels of haptic support (FF) are plotted separately. The corresponding medians for each HCR are indicated by dashed vertical lines.

driver groups (supporting H.I). There is also no difference between P, class-average or CL for both driver groups (opposing H.II).

K_{LoHS} For curve-cutters, a feedforward gain of $K_{LoHS} = 0.92$ receives significantly higher subjective acceptance scores than $K_{LoHS} = 0.5$ ($\chi^2(1)$=7.681, p = 0.006, W = 0.12), with a median increase of 1 point. For offset-drivers, such a significant improvement is not present due to the low score achieved with CC for $K_{LoHS} = 0.92$. When the CC HCR is omitted, a feedforward gain of $K_{LoHS} = 0.92$ receives significantly higher subjective acceptance scores with the three other HCRs ($\chi^2(1)$=6.429, p = 0.011, W = 0.143), with a median increase of again 1 point. Hereby, a larger feedforward gain results in significantly better acceptance ratings for P, class-average and CL (supporting H.III).

SATISFACTION AND USEFULNESS

The subjective satisfaction and usefulness ratings obtained from the Van der Laan questionnaire are summarised in Fig. 8.12. Here the results of both the usefulness and satisfaction scores are presented together in one plot, illustrating both the average and standard deviations with error bars.

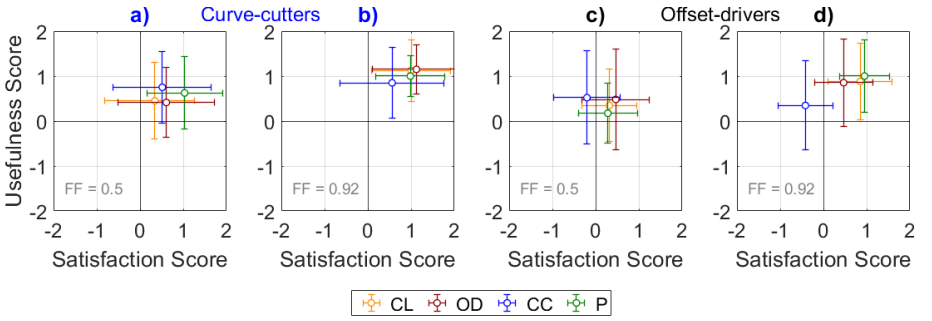

Figure 8.12: Satisfaction-Usefulness plots resulting from the Van Der Laan questionnaire for both the curve-cutters in a) & b) and offset-drivers in c) & d). Here the average ratings are given by the dots with 2 standard deviations as error bars.

HCR type For the curve-cutters, there is no significant effect for satisfaction and usefulness with HCR type, in line with the CARS scores. For the offset-drivers a significant effect for satisfaction with HCR type for only the K_{LoHS} = 0.92 case (χ^2(3)=8.293, p = 0.04, W = 0.184). Pairwise comparisons show a significant difference between only P and CC (p<0.05), with an average decrease of 1.4 scores. There is also a significant effect for usefulness with HCR for K_{LoHS} = 0.92 (χ^2(3)=11.17, p = 0.011, W = 0.248). Pairwise comparisons show a significant effect between only P and CC (p<0.05), with an average decrease of 0.7 points. Furthermore, no significant difference between P and class-average HCRs is found (supporting H.I). Finally, there is also no difference between P, class-average or CL for both driver groups (opposing H.II).

K_{LoHS} For curve-cutters there is a significant effect for satisfaction with K_{LoHS} level (χ^2(1) = 4.9, p = 0.027, W = 0.077) with an average increase of 0.3 points. There is also a significant effect for usefulness with K_{LoHS} level (χ^2(1)=17.47, p = 0.000, W = 0.273), with an average increase of 0.47 points. The offset drivers also rate the usefulness of K_{LoHS} = 0.92 conditions significantly higher (χ^2(1)=4.923, p = 0.027, W = 0.082), at an average increase of 0.4 points. However, they clearly do not accept a CC HCR at K_{LoHS}=0.92, nevertheless, this HCR becomes indistinguishable from other HCRs and thereby acceptable at K_{LoHS}=0.5. The perceived usefulness for both driver groups significantly increases with increasing feedforward gain (supporting H.III), whereas the perceived satisfaction significantly increases with an increased feedforward gain for the curve-cutters only.

8.5. DISCUSSION

The aim of this chapter is to gain an understanding of how to personalise haptic shared control. This understanding is gained through exploring different levels of personalisation and different levels of haptic support. An experiment is carried out on two driver groups, curve-cutters and offset-drivers, classified according to the 7-class classifier in Chapter 4. These two driver groups were presented with four different controller references: fully personalised, class average personalised, centreline and the class-average of the other driver group. Additionally, these references were presented at a high and low setting of haptic support. This experiment was carefully designed to gain insight into 1) the level of personalisation (fully personalised, v.s. class average personalised) and, 2) the level of haptic support required for the best driver acceptance of haptic shared control. The following three hypotheses were made: 1) no difference in subjective acceptance measures or objective conflict is expected between full personalisation and class average personalisation, 2) the centreline guidance results in lower subjective acceptance and more extensive duration in conflict than both forms of personalisation and, 3) a higher feedforward gain is expected to increase acceptance and decrease conflict.

THE ACCEPTANCE OF THE HYPOTHESES

There is no evidence for any differences between class-average and fully personalised HCRs if you ask participants (*accepting H.I*). Nevertheless, significant differences between personalisation levels are found in the considered objective measures. Where, without doubt, the fully personalised HCR results in significantly the least conflict reflected by 1) the lowest conflict duration, at mean conflict times below 10% of the total driving time (in both the prepositioning and in-curve phases) and, 2) the lowest conflict torques, at mean conflict torque below

0.2 Nm in both the prepositioning and in-curve phases, for both driver groups. This indicates the best controller-driver agreement is found for fully personalised HCR in terms of aligning torques. *Contrary to H.I*, the class-average personalised guidance is sometimes found to be significantly different from fully personalised guidance. The instances when significant differences are found is summarised in Table 8.1. This objective conflict is found for both curve-cutters in left curves and offset-drivers for right curves. For curve-cutters a significant in-curve peak in conflict torque is found for the CC HCR in left curves (illustrated in Fig. D.6), whereas for offset-drivers a significant peak in conflict time and torque is found for the OD HCR in right curves (illustrated in Fig. 8.7 & Fig. D.5). Careful analysis finds that this unexpected objective conflict torque and time is dependent on two factors: 1) whether the driver has a similar trajectory to the HCR used by the HSC and, 2) the adaptiveness of the driver to the provided HCR, i.e. the driver follows the HCR. Fig. 8.13 illustrates two drivers with similar driving styles and their adaptiveness to the CC HCR guidance. The non-adaptive driver (driver 2) receives high conflict time and torque (43% and 0.9 Nm), whereas the adaptive driver (driver 1) receives low conflict time and torque (2% and 0.1 Nm). Therefore, the willingness of the driver to adapt is an essential factor that affects torque alignment and thereby torque conflict.

Table 8.1: The significance of pairwise comparisons between fully personalised and class-average HCRs for both driver groups. The extra metrics feedback torque T_{SoHF} and conflict torque T_c are given in the appendix. When σ = -, p > 0.05, when σ = *, p <0.05, when σ = **, p <0.01.

Metric	Curve-cutters				Offset-drivers			
	Right		Left		Right		Left	
	Prep	Crv	Prep	Crv	Prep	Crv	Prep	Crv
T_D	-	-	**	-	-	**	**	-
T_{SoHF}	-	-	**	*	-	-	*	-
O_c	-	-	*	**	-	*	**	**
T_c	-	-	**	*	-	-	**	-

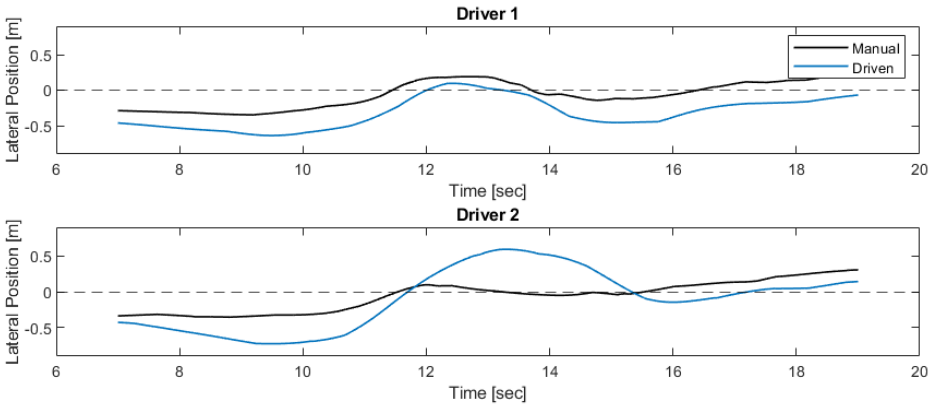

Figure 8.13: Comparison of two drivers trajectory behaviour and their resulting, driven trajectories when subjected to CC HCR in left curves. The resulting adaptation is an indication of driver agreement

The industry standard of centerline guidance for intelligent vehicles [98] [55] is not found to cause the highest conflict torques, nor does it exhibit the worst subjective ratings, meaning that hypothesis *H.II is rejected*. In this study, and Chapter 4, there are no drivers whose trajectory style is centerline driving, which means it is not a 'human-like' trajectory. In fact, the highest driver torques and feedback torques are found for CL in right curves. Nevertheless, both driver groups do not distinguish it from personalised HCRs on the subjective scale, crediting this particular HCR 'human-compatible'. This poses the question of whether acceptance is affected by what drivers expect from an intelligent vehicle, i.e. whether the guidance is predictable, rather than only a similarity in trajectory style and aligning torques. In fact, the fifth guideline of human-centred design is that *'the automation must be predictable by the human'* [27]. Hence, when a guidance reference is predictable, this can build trust, which can leverage acceptance.

Increasing the level of haptic support K_{LoHS} shows a significant change in perceived acceptance. A level of haptic support that has the HSC contribute 92% of the driving task receives significantly higher acceptance than the 50% support level, with 1) acceptance ratings significantly improving by 1 score (on a 1-10 scale), 2) usefulness significantly improving by a factor 1.9 (on a -2 to 2 scale), 3) in-curve driver torque significantly decreasing by 0.64 Nm, 4) prepositioning conflict time significantly decreasing by 17% and 5) prepositioning conflict torque significantly reducing by 0.1 Nm. This finding means that *H.III is accepted*. A system that has a higher level of haptic support, i.e., a more substantial feedforward torque, supports the driver with more guidance, making the ratio between guidance (feedforward) and correction (feedback) torques larger. Also, a larger feedforward gain is closer to full automation: with a K_{LoHS} =1, the driver can 'drive' hands-free, as the support will provide all the torques necessary to drive the curve, which drivers may find more useful. These findings are in line with the conclusions of Chapter 7, where it was found that the FDC controller increases driver acceptance compared to a pure feedback haptic controller due to the feedforward torques. Although both are torques, they are dynamically perceived very differently by drivers. Analogous to coaching, feedback torque is like a coach who constantly corrects and penalises you *after* you make a mistake. That is, feedback provides corrections that require most of the initiation to be done by the driver, where drivers may end up giving anticipatory torques over correction torques. On the other hand, feedforward torque is like a coach who guides you with anticipatory advice along the desired path. Feedforward gives earlier guidance than feedback, especially when entering a curve. When the timing and the magnitude of the guidance are accurate, acceptance will be high. Hereby, it is proven that the feedforward torque plays a crucial role in the acceptance of haptic shared control. Therefore, rather than using K_{LoHS} = 0.92, which was heuristically tuned in Chapter 7, it is recommended to further investigate the ideal value of the feedforward gain K_{LoHS} for optimal acceptance.

EFFECT OF TRAJECTORY TYPE ON SUBJECTIVE ACCEPTANCE

No evidence was found that full personalisation is necessary for the FDC haptic shared controller if you ask participants. The subjective measures show that drivers from both tested groups do not perceive any significant change in acceptance, satisfaction or usefulness when the guiding controller reference changes between the fully personalised, class-average personalisation and centerline guidance. Even though subjective questionnaires may not always convey accurate preferences due to possible time-variant artefacts during the experiment such as fatigue (especially when comparing four different HCRs), these (non-significant) effects and trends between trajectory types are very consistent across all three subjective questionnaires. This relative indifference to the level of personalisation and centerline could be caused by the inherent nature of the FDC controller exhibiting feedforward guiding torques, which allows drivers to

deviate from the set HCR without much-felt penalty. This is in contrast to pure feedback haptic systems [31] [129] where any deviation is penalised. This allowance of the FDC controller could then result in no subjectively perceived added benefit to personalisation through aligning torques.

Although no evidence was found to support the use of the considered personalised HCRs in this experiment from a subjective stance, this is not to say that HCR type or that personalisation in the general sense is not important. The goal of personalisation in ADAS is *'to improve the driving experience and the performance of the assisted drivers by adapting the assistance system to their preferences and needs'* [72]. Hereby personalisation is not limited to mimicking the driver's behaviour (as is considered in this chapter), but rather adapting to their preferences. In fact, mimicking a driver's style in automated vehicles is not (always) found to leverage the greatest acceptance, or fall within a driver's preferences [149] [71] [22] [193]. Nevertheless, the results of this experiment indeed do show differing *preferences* between driver groups. Offset-drivers particularly show a significant decrease in acceptance, satisfaction and usefulness when presented with curve-cutting guidance. Curve-cutting guidance is not what offset-drivers expect from a haptic controller and is out of the 'comfort zone' of these drivers. In fact, a system that explicitly does not guide an offset driver with curve cutting guidance is adapting to their preferences and is thereby personalised in the general sense [72]. Conversely, curve-cutters readily accept an offset driving HCR, without any significant difference to fully personalised, class-average or centerline guidance. Hereby, one could relate the acceptance of HCR to the predictability of the trajectory, i.e., for curve-cutters, both CL and OD could be more predictable than CC is to an offset-driver. Therefore, offset-drivers, in contrast to curve-cutters are more sensitive to changes in guidance trajectory and thereby, may benefit more from systems that explicitly adapt to their preferences and are thereby personalised (in a general sense).

EFFECT OF DRIVER GROUP

Compared to curve-cutters, the offset-drivers show significantly reduced conflict times, with an average decrease of 4% and significantly less feedback torque, with an average decrease of 0.2 Nm, across all conditions. This could be a result of their 'willingness to adapt'. Figs. D.8 and D.9 in Appendix D.5 show the six left and right trajectories driven for the eight different test conditions, for a curve-cutter and offset-driver, respectively. The curve-cutter persistently maintains their trajectory class with guidance, whereas the offset-driver adapts to the presented HCR. With these adaptations, the conflict received is less. Hereby, driver type can affect the magnitude of conflict in general. Driver type seems to correlate with adaptation and therefore, further investigating the adaptiveness of the other the driver types shown in Fig. 8.4, is recommended.

Since the curve-cutter group is comprised of two subgroups of driver types (R2L1 & R2L2, i.e., in left curves one group is classified as severe curve-cutting and the other normal curve-cutting), it can be speculated that this is one of the reasons for the increased significant differences between fully personalised and class-average personalisation in left curves specifically, given in Table 8.1. In fact, the trend in conflict time and conflict torque does not significantly change with sub-driver group. Fig. D.7 in Appendix D.4 illustrates that the high conflict stems from both sub-groups of drivers (R2L2 & R2L1). Rather than a binary difference in class, this could be related to the fact that the considered L1 is not a class average of the tested drivers, since it is the average of only 3 drivers from the dataset in Chapter 4.

ACCEPTANCE OF MORE 'EFFICIENT' HCRS

Driver's do not readily accept 'more efficient' HCRs; in fact, the class-average CC HCR exhibits significantly lower acceptance, satisfaction and usefulness ratings for offset-drivers. The L1 trajectory (the left trajectory of the CC HCR) is a realisation of an 'optimal race-line' [30], which

inherently maximises the Time to Line Crossing (thereby safety), minimises the in-curve lateral acceleration and reduces the driver's steering effort (as deduced in Chapter 5). Although these are objective benefits that most drivers would find desirable, when subjected to this 'unconventional' guidance, the offset-drivers are not accepting of it. Nevertheless, this could be a result of experimenting in a fixed-base driving simulator, where these drivers cannot perceive the benefit of following a curve-cutting trajectory that minimises lateral acceleration.

TORQUE ALIGNMENT AND DRIVER GROUP

Controller-driver agreement in terms of aligning torques is generally found during instances when the driver's and the HCR's trajectories are aligned, even without personalised guidance. This is clearly observed when presented with the other group's class-average guidance. Curve cutters try to minimise their lateral acceleration by starting at the outer part of the curve in the prepositioning phase and gravitate towards the inner part of the curve in the in-curve phase, this behaviour is seen both for right and left curves. Offset-drivers try to follow a straight path along the centreline (but not on the centreline), they drive along the inner part of the curve for right curves and along the outer part of the curve for left curves. More specifically, the offset-drivers preposition on the inner side of a right curve, whereas curve-cutters preposition on the outer side of a right curve. Whilst in right curves, both trajectories gravitate towards the inner curve. Both offset-drivers and curve-cutters preposition towards the outer part of the left curve, whereas in the left curve offset-drivers stay on the outer curve whereas curve-cutters gravitate towards the inner part of the curve. This is illustrated graphically in Fig. 8.14. For this reason, differences between the considered driver groups are reflected in the prepositioning phase for right curves and in the curve phase for left curves. Whereas during prepositioning in left curves and being in the curve during right curves, the controller-driver intent is the same, which is reflected by significantly shorter conflict times and smaller conflict torques.

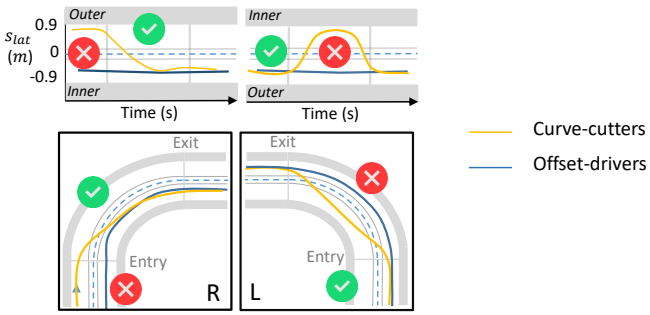

Figure 8.14: Similarity and differences between the two driver types

SUBJECTIVE ACCEPTANCE CAN NOT BE EXPLAINED BY OBJECTIVE CONFLICT OCCURRENCE

The objective conflict measures considered in this chapter do not show a correlation with the subjective ratings given by the drivers for the considered FDC controller. For example, the OD HCR is given high subjective ratings by both groups of drivers. However, the OD HCR results in the highest in-curve conflict torque for both driver groups, in right curves. Surprisingly, the OD HCR does result in a low amount of (penalising) feedback torque. Whereas the HCR with the

highest penalising feedback torque, CL receives minimal conflict. This implies that the computed conflict for OD may result from the feedforward torque (it is, however, not possible to distinguish precisely whether the driver has a conflict with a feedback or feedforward torque as they are always provided together). Moreover, we find that with a higher feedforward torque, which is subjectively favoured, the in-curve conflict torque significantly increases in left curves with 0.15 Nm on average. However, a 'conflict' with feedforward torque is not penalising nor unpleasant. Some drivers may like to 'lean on the system' and feel the haptic torques, in this way the opposing torque may feel like the system 'protects' the driver [168]. However, leaning behaviour is directly recorded as an objective conflict. Conversely, an offset driver that clearly 'fights' curve-cutting guidance may struggle with the guidance and receive the same magnitude of conflict time. Nevertheless, the defining line between 'leaning' and 'struggling' may be related to *conflict frequency*. Therefore, to bridge the gap between the objective and subjective acceptance measures, it is essential to investigate different methods to compute conflict in a way that matches what the driver perceives as conflict, rather than only a computation of opposing torques.

8.6. CONCLUSIONS

This chapter investigated different personalisation approaches to improve acceptance of haptic shared control. It was assumed that when guidance is strongly linked to drivers' driving style, acceptance will be higher. The investigation is based on empirical data from two driver groups with distinct driving styles: the most commonly observed driver type referred to as offset-drivers and curve-cutters. These drivers are presented with a Haptic Shared Control system that is driven by four different Human-Compatible-References (HCRs) and two different feedforward gains (K_{LoHS}). The four tested HCRs are: 1) a fully personalised reference, 2) a class-average personalised reference, 3) the centerline reference and, 4) the other driver group's class-average reference. These HCRs were applied with two feedforward gains: K_{LoHS} = 0.5 and K_{LoHS} = 0.92. The acceptance is evaluated by both subjective questionnaires and objective measures of conflict time, conflict torque, feedback torque and driver torque.

The main takeaways from the tested experimental conditions in this chapter are:

8

1. Full personalisation significantly improves torque alignment between driver and controller which is reflected by 1) significantly the lowest conflict times, at mean conflict times below 10% and, 2) significantly the lowest conflict torques, at mean conflict torque below 0.2 Nm in both the prepositioning and in-curve phases, for both driver groups. Hereby, a class-average HCR leverages significantly larger conflict compared to fully personalised guidance. This is unexpected and opposes H.I

2. Drivers do not perceive this torque alignment from the fully personalised guidance as an added benefit. Subjectively, there is no significant change in acceptance, satisfaction or usefulness when the HCR changes from fully personalised to class-average to centerline, for both tested driver groups. This means that all drivers accept a HSC that drives like them, regardless of being fully personalised or class average personalised, which is expected and supports H.I. The centreline driving guidance (industry's standard) is also equally well accepted subjectively, even though objectively, conflict torques are significantly higher. This is unexpected and opposes H.II; nevertheless, this acceptance seems to stem from the predictability of the reference.

3. We find that some (computed) conflict torques are not experienced as 'conflicting'. We expect that conflict in feedforward torques (exclusive to the FDC HSC) are accepted while conflict in the form of feedback torques are not accepted.

4. The level of haptic shared control gain K_{LoHS} plays a crucial role in driver acceptance, with acceptance ratings significantly improving by 1 score (on a scale of 1-10), usefulness significantly improving by a factor 1.9, in-curve driver torque significantly decreasing by 0.64 Nm and prepositioning conflict time significantly decreasing by 17%. This result supports H.III.

The design guidelines that can be extracted from the findings of this paper are fourfold: 1) for most applications personalisation (through mimicking the driver's trajectory) is not necessary with the FDC controller; instead, the centreline reference achieves equal subjective acceptance (due to the predictability of this trajectory) and can be used for computationally efficient HSC implementation, 2) for applications where torque alignment is explicitly necessary a fully personalised reference should be used, 3) the preference of the majority of drivers, i.e. offset-drivers, should be considered by not presenting them with curve-cutting guidance and, 4) a feedforward gain of 0.92 should always be used; however, it is recommended to investigate other gain values that could further improve driver acceptance.

9

CONCLUSIONS AND RECOMMENDATIONS

This final chapter provides a comprehensive overview of the achievements, the conclusions and the recommendations of this thesis. The conclusions and recommendations are presented in four sections. Firstly, a brief overview of the achievements and the key conclusions of each part. Secondly, the overall conclusions of the thesis. Thirdly, design guidelines for haptic shared control in curves. Fourthly, possible applications of the developed tools. Finally, the recommendations section gives a clear description of essential follow-up research in the field of personalisation in driving automation.

RECAPITULATION OF RESEARCH GOAL

Most current advancements in driving automation follow a technology-centred approach, placing their attention on engineering solutions where human cognitive and motoric functions are ultimately replaced with assisting automation. These systems theoretically improve safety and comfort; however, in practice, the human has difficulty in adopting, accepting and trusting the assisting driving automation. This thesis aims to *improve driver acceptance* in haptic shared control, a specific type of Advanced Driver Assistance System that provides lane-keeping assistance during lateral control tasks. In this type of assistance, conflicts, which are correcting torques that indicate opposing intentions, are often reported during curve driving, thereby decreasing the acceptance.

This acceptance can be improved by considering the presence of the human driver through applying human-centred design principles, i.e., integrating the human's preferences and limitations into the design of HSC to better align the system with the human. With the support providing a continuous share of the control, which are felt through torques on the steering wheel, the need for the controller's reference trajectory to be compatible with the individual driver becomes even more relevant. In this thesis, we assume that this compatibility between driver and automation is best achieved through the notion of *personalisation*, where, the controller's reference trajectory would mimic the trajectory the individual would take when driving manually. Therefore, the goal of this thesis is to:

Thesis goal

Develop model-based approaches to personalise trajectories for steering through curves to enhance acceptance in haptic shared control systems.

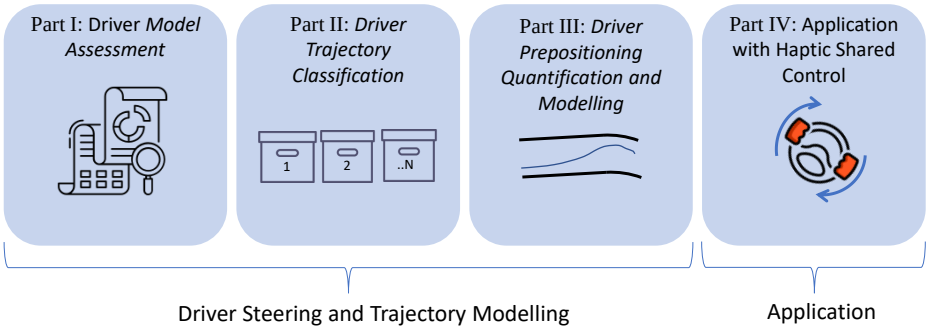

Figure 9.1: This figure shows the four parts of the research carried out in this thesis.

To realise an improvement in acceptance through personalisation, the following four research parts are considered, see Fig. 9.1:

I *Driver Model Assessment* This part aims to develop a selection technique to objectively determine what model(s) are most suitable to describe and identify individual drivers.

II *Driver Trajectory Classification* This part aims to develop a classifier that can categorise the different types of observable driver trajectory styles in the driving population.

III *Driver Prepositioning Quantification and Modelling* This part aims to develop a control-theoretic driver steering and trajectory model technique that can describe individual tra-

jectories, which can only be achieved by incorporating prepositioning. This individualised driver model can facilitate a high level of personalisation.

IV *Haptic Shared Control Application* This part applies and evaluates the developed personalised driver steering and trajectory model(s) in a haptic shared controller, that allows for personalisation.

Driver steering and trajectory modelling are dealt within the first three parts, which, focus on improving and developing methods for modelling of personalised steering and trajectories using accurate control-theoretic modelling and rule-based trajectory classification. Part IV, *Application*, deals with applying the outcomes of Part I-III into a final study to compare the impact of personalised driving assistance on driver behaviour and acceptance. These two components are shown in Fig. 9.1.

The following sections include an overview of the achievements in the research parts (Chapters 2-8), summarising the key conclusions. In addition, the overall thesis conclusions, design guidelines, developed tools and recommendations for further research are discussed.

9.1. OVERVIEW OF RESEARCH PARTS

9.1.1. PART I: DRIVER MODEL ASSESSMENT

CHAPTER 2: DRIVER MODEL ASSESSMENT METHOD

This chapter has introduced a novel assessment methodology for driver models that produce steering profiles and trajectories. Three criteria are developed based on understanding the requirements for identifying individual driver behaviour in an efficient and meaningful way. The three developed criteria are *descriptiveness*, *identifiability* and *realism*. Descriptiveness quantifies to what extent the model can capture different driver trajectories in a curve, where a larger descriptiveness value is good. Identifiability evaluates the extent to which the mapping between a given driver steering behaviour and a parameter set is unique. A smaller identifiability space is means a more unique solution which is desirable. Finally, realism determines the parameter space that results in realistic driver steering behaviour (i.e., remains within road boundaries and limited steering wheel reversals), where, a large realistic parameter space implies a model with many feasible solutions, which is desirable.

The effectiveness of this assessment method is demonstrated on two simple (i.e., comprising only a single feedback loop) driver steering and trajectory models. One model uses linear prediction (with information of current time only) for future lateral error estimation, whereas the other model uses non-linear prediction (with explicit preview information). The assessment methodology was able to discriminate between these two models' performance for all three criteria. The non-linear model outperforms the linear model in all three criteria, i.e., it shows: 1) twice as large descriptiveness, 2) twice as good model identifiability for steering angle, and 3) a 50 times better ability for unique mapping between parameter set and modelled lateral trajectory position.

Apart from these differences between driver models, the assessment method revealed an essential result: fitting on steering angles will always give an accurate fit value, given that the resulting trajectory remains within the boundaries of the road. As a result, for capturing all individual differences in steering behaviour, fitting on steering angles alone is not sufficient. Including trajectories in identification will make the distinction between individual drivers valid since trajectories provide a refined way of discriminating successfully between drivers.

CHAPTER 3: ASSESSMENT OF THREE CONTROL-THEORETIC DRIVER MODELS

The main contribution of this chapter is an upgrade of the model assessment introduced in *Chapter 2*, making it both standardised and trajectory-class dependent. In *Chapter 2*, three cri-

teria were developed, whereas in *Chapter 3* these are compacted to two criteria via standardisation. Here, the calculated identifiability is standardised with a model's realism space, introducing the metric 'Realistic Identifiability'. Moreover, a meticulous process chain is designed to obtain these metrics in a data-driven way, i.e., dependent on each trajectory class (as defined in *Chapter 4*). Hereby, *descriptiveness* is defined as the degree to which the model can capture (reproduce) driver behaviour within a specific trajectory class. In the updated methodology, the *realistic identifiability* is evaluated only for the driver classes that the model can adequately capture.

The upgraded methodology is used to assess three different driver steering models: the Mars model, the Van der El model and the Van Paassen Model. These models are different than those assessed in Chapter 2 as they exhibit both feedback and feedforward components, compared to the pure feedback models considered in Chapter 2. The descriptiveness results are the Mars model (16%), Van der El model (17%) and the Van Paassen model (17%). These low descriptiveness values are caused by the models' inability to describe driver behaviour before curve entry: prepositioning. The corresponding weighted average fit values (i.e., each model has a fit per class and each class has a different percentage of occurrence) are the Mars model (72%), the Van der El model (64%) and the Van Paassen model (74%). The average realistic identifiability in the trajectory domain are the Mars model (9%), the Van der El model (2.35%) and the Van Paassen model (0.055%). The assessment results imply that the Van Paassen model exhibits the best descriptiveness and also the best realistic identifiability. With high realistic identifiability and low descriptiveness, Van Paassen however, is 'under-parametrised'. It is recommended to extend the Van Paassen model to account for prepositioning as this is identified as the biggest bottleneck for all considered models. Here an extension does not imply merely an addition of parameters, but a structural extension of the model.

Part I: Key conclusions

- The developed driver model assessment methodology quantifies the ability of any control-theoretic driver model to produce trajectories and steering angles that can capture a wide range of individual driver behaviours (*descriptiveness*), as well as be parsimonious to allow for identification (*identifiability*).
- According to the developed driver model assessment procedure, the Van Paassen model, is most suitable for identifying individual driver steering and trajectory behaviour.
- The main aspect that all tested models can not capture is the drivers' anticipating and preparing for curve entry (prepositioning), which severely degrades the descriptiveness of these models. A driver model extension is needed to account for this.

9.1.2. PART II: DRIVER TRAJECTORY CLASSIFICATION

CHAPTER 4: DRIVER TRAJECTORY CLASSIFICATION

How differently do drivers take curves? Can individual trajectories be categorised, and used as templates for personalisation? The lack of a mathematical algorithm to categorise such driver trajectories in literature was the driving force in this research to develop novel rule-based trajectory classifiers for curve driving. Two novel rule-based classifiers that categorise 7 and 11 different trajectory styles are introduced. The classifiers distinguish drivers based on two main decision nodes: 1) where drivers start at curve onset, and 2) the number of crossings along the lane centre during the curve. The two classifiers categorise different trajectory types, ranging

from *severe curve-cutting* to *severe counter curve-cutting*. The classifiers are applied to a dataset of 45 drivers negotiating three different curves in a fixed-base driving simulator, each with different curvature and fixed velocity.

The classification results show that for the experimental conditions studied normal curve cutting and offset inner curve negotiation are the most prevalent classes for right curves, making up for 86% of class occurrence for the 7-class classifier and 77% for the 11-class classifier, on average. Offset outer curve negotiation and normal curve cutting are the most prevalent classes for left curves, making up 69% of class occurrence for the 7-class classifier and 65% for the 11-class classifier on average. The 11-class classifier was introduced to distinguish the drivers that start at centreline, these make up 7% of left curves on average and 17% of right curves on average.

Ultimately, these classifiers formalise the types of driver trajectory behaviour in curves available in the driving population and give insight into types of driver trajectories that are important to predict, based on prevalence.

Out of the two introduced classifiers, one does not outperform the other. Rather, the 7-class classifier is intended to be used for empirical data classification, whereas the 11-class classifier is necessary when classifying trajectories generated by conventional control-theoretic driver models. Control-theoretic models all assume that drivers start at lane centre at curve entry, to account for this, the 11 class classifier is introduced.

Part II: Key conclusions

- The developed algorithms for trajectory classification can be used as an alternative to control-theoretic driver modelling, and be directly applied in personalised ADAS.
- Centerline trajectory is not a human-like trajectory, as the data from the driving simulator study shows that no drivers fall into this class.
- When personalising ADAS, left and right curves should be treated independently, as drivers drive different trajectories on the left and right curves because they exhibit a non-zero straight section bias.

9.1.3. PART III: DRIVER PREPOSITIONING

CHAPTER 5: DRIVER PREPOSITIONING QUANTIFICATION

This chapter has dealt with the impact of curve radius and driving velocity on driver steering and trajectory behaviour before curve entry, i.e., *prepositioning*. The findings of this chapter, based on data from a dedicated driving simulator experiment, are 1) the maximum prepositioning displacement consistently increases with increasing (fixed) velocity and increases with increasing road curvature, 2) prepositioning behaviour is evident for 88% of the driven curves and, 3) drivers adapt their prepositioning to the perceived centreline Time-to-Line-Crossing (TLC), where the driver increases their TLC by following the optimal race-line, which, inherently exhibits prepositioning. This means that many drivers preposition, however, they do not always end up harvesting the full potential benefit of prepositioning.

Incorporating the findings of this chapter in future driver modelling (e.g., *Chapter 6*) will bridge the gap between straight road and in-curve driving behaviour, thereby bolstering the descriptive capacity of these models.

CHAPTER 6: DRIVER PREPOSITIONING MODELLING AND INTEGRATION

This chapter has introduced a new prepositioning path generation model for drivers' lateral position changes before curve entry with parameters that have geometrical significance.

The prepositioning path model is used to augment the Van Paassen control-theoretic driver model (the best model from the model assessment in *Chapter 3*) by adding a prepositioning component in the form of a prepositioning reference trajectory and reference prepositioning curvature prior to curve entry. The descriptiveness of the Van Paassen model with and without the prepositioning augmentation is assessed. The Van Paassen model without prepositioning can accurately describe 17% of drivers in right curves and 7% of drivers in left curves, less than a fifth of the drivers considered. The Van Paassen model with prepositioning can sufficiently describe 100% of drivers in both right and left curves.

Part III: Key conclusions

- Prepositioning displacement significantly increases with velocity and curvature.
- Prepositioning is a phenomenon that is triggered by the objective of the driver to maximise the time-to-line-crossing.
- Incorporating prepositioning behaviour is vital to bridge the gap between straight and in-curve driver steering and trajectory modelling.
- Combining the developed prepositioning path model with the Van Paassen control-theoretic driver steering model results in a model that can accurately predict all driver class trajectories that were experimentally measured as part of this thesis.

9.1.4. PART IV: APPLICATION

CHAPTER 7: FOUR DESIGN CHOICE IMPLEMENTATION

The goal of this chapter is to implement and evaluate the Four-Design-Choice (FDC) haptic shared controller, the first explicitly personalisable HSC in the literature. The implemented FDC in this thesis is the (first) realisation of the 'Four-Design-Choice-Architecture' [177] for haptic shared control.

One of the four design choices is critical for personalisation: the Human-Compatible Reference (HCR), i.e., the independent controller reference comprising of trajectory, heading and steering angles. In this chapter, two variations of the HCR are evaluated with the FDC: 1) a One-Size-Fits-All (OSFA) and, 2) a personalised HCR. These HCRs are generated by the Mars driver model [116] using coefficients identified by fitting the model to observed trajectories. Note that due to the limitations of the Mars model (found in *Chapter 3*), the prepositioning phase could not be modelled. Hereby, prepositioning is not taken into account in the HSC implementations of this chapter. The two HSC implementations are compared to a conventional (quickened) pure feedback haptic shared controller (from the Delft Haptics Lab), where reference trajectory and strength of haptic assistance are inseparable, i.e., the Meshed (M) HSC [31].

Results show that the personalised haptic shared controller resulted in significantly higher subjective ratings of usefulness and satisfaction than the M-HSC, with average improvements of 1.3 and 0.8 (on a -2 to 2 scale), respectively. For all metrics, there is no significant improvement between OSFA and personalised reference trajectory in the FDC. In general, the FDC-HSC (both OSFA and personalised) significantly reduces the occurrence of conflicts by an average factor of 2.3 and significantly reduces driver torque by an average factor of factor 3.2 compared to the M-HSC. Detailed analysis of steering wheel torque data shows that this can be attributed to two

inherent design elements in the FDC-HSC. Compared to the M-HSC (which is a pure feedback HSC), the FDC-HSC has an independent feedforward contribution and an HCR that supports (future) curve entry intentions. This highlights the improved effectiveness of the Four-Design-Choice-Architecture and its (larger) potential for acceptance with different design choice settings and a fully personalised HCR.

CHAPTER 8: EVALUATION OF PERSONALISED HAPTIC SHARED CONTROL

This chapter investigated different personalisation approaches to improve acceptance of haptic shared control. The different approaches are realised through considering different levels of personalisation (which include prepositioning) and different levels of haptic support. The investigation is based on empirical data from two driver groups with distinct driving styles: the most commonly observed driver type referred to as offset drivers, as well as curve cutters. These drivers are presented with a FDC-HSC that is driven by four different Human-Compatible- References (HCR) and two different feedforward gains (K_{LoHS}). The four tested HCRs are: 1) a fully personalised reference, 2) a class-average personalised reference, 3) a centerline reference and, 4) the other driver group's class-average reference. The fully personalised references are generated by averaging recorded driver data, the class-average references are generated by classification data, the centreline reference is generated using the non-linear driver model in Chapter 2. These HCRs were applied with two FDC HSC feedforward gains: K_{LoHS} = 0.5 and K_{LoHS} = 0.92. The acceptance is evaluated by both subjective questionnaires and objective measures of conflict time, conflict torque, feedback torque and driver torque.

Full personalisation successfully improves the alignment of torques between driver and controller which is reflected by 1) significantly the lowest conflict times, at mean conflict times below 10% and, 2) significantly the lowest conflict torques, at mean conflict torque below 0.2Nm in both the prepositioning and in-curve phases, for both driver groups. Hereby, a class-average HCR does not leverage the same objective results as fully personalised. Nevertheless, drivers do not perceive this torque alignment as an added benefit. Subjectively, there is no significant change in acceptance, satisfaction or usefulness when the HCR changes from fully personalised to class-average to centerline, for both tested driver groups. This is attributable to two reasons: 1) the centreline being a predictable trajectory and, 2) the feedforward guiding torque gives the driver the allowance to deviate from the reference without penalty.

The level of haptic shared control gain K_{LoHS} plays a crucial role in driver acceptance, with acceptance ratings significantly improving by 1 score (on a scale of 1-10), usefulness significantly improving by a factor 1.9, in-curve driver torque significantly decreasing by 0.64 Nm and prepositioning conflict time significantly decreasing by 17%.

Part IV: Key conclusions

- Compared to a conventional HSC (Chapter 7), the FDC HSC significantly reduces conflict torque and driver torque, which is attributable to the feedforward guidance torque.
- Torque alignment between driver and controller is substantially improved with a fully personalised HCR in the FDC HSC.
- Drivers do not perceive any subjective difference between full personalisation, class-average personalisation and centerline guidance.
- A higher level of haptic support (stronger feedforward guidance) significantly reduces conflict time and improves subjective acceptance measures.

9.2. OVERALL CONCLUSIONS

This thesis provides novel tools and solutions to personalise haptic shared control in curves, as a means to increase user acceptance. Overall, eight main conclusions can be drawn from the different research parts of this thesis that can be of importance to relevant researchers and industry.

DRIVER MODEL ASSESSMENT IS NEEDED TO DETERMINE WHETHER A DRIVER MODEL IS SUITABLE FOR FULL PERSONALISATION

To determine whether a driver steering and trajectory model is suitable for full personalisation, detailed model assessment is indispensable, as shown in *Chapters 2 & 3*. The attempt to individualise with a driver model incapable of successfully capturing individual driver steering and trajectory behaviour is shown to have no significant additional benefits in conflict reduction (*Chapter 7*). In fact, the model assessment procedure has demonstrated that all control-theoretic models in literature are not sufficient to provide full personalisation, with the maximum amount of driver classes they can reproduce amounting to 17% (*Chapter 3*). Moreover, the assessment analysis offers the opportunity for a trade-off between the accuracy of capturing the spectrum of drivers found in the driving population and, the identifiability of these models per driver class.

IDENTIFYING INDIVIDUAL DRIVERS REALISTICALLY CAN ONLY BE DONE BY CONSIDERING *both* STEERING ANGLES AND TRAJECTORIES

When validating a model for capturing individual driver steering behaviour, this validation should never only be performed using the model's capacity for matching steering angle data. In fact, it is also possible to simply penalise oscillatory steering and solely focus on trajectory fit, as is done in *Chapter 3*. Unfortunately, many driver identification works in literature make this mistake. From *Chapter 2*, it is evident that the identifiability space of all models in the steering angle domain is at least 15 times larger than that for the trajectory. Whereas in *Chapter 3*, it is formalised that having a VAF of 90% or higher in the steering angle domain implies two things: 1) 0 lane boundary crossings and 2) do not exhibit more than six steering reversals. These findings are found for identification in the curve phase. Hence, having a proper steering angle fit only means that the resulting trajectory is within the curve boundaries without excessive oscillations. Fitting on trajectories is found to provide a more accurate prediction of trajectory behaviour. Nevertheless, the steering behaviour should still be taken into account in the identification procedure, as fitting on only trajectory can result in steering oscillations as seen in *Chapter 2*.

THE VAN PAASSEN CONTROL-THEORETIC DRIVER STEERING MODEL HAS THE BEST DESCRIPTIVENESS AND REALISTIC IDENTIFIABILITY IN LITERATURE

The Van Paassen model has the best descriptiveness and realistic identifiability out of the models tested in *Chapter 3*. Three elements discern this model from others in terms of structural components: 1) an independent feedforward loop providing the needed steering to get through the curve, 2) an anticipatory feedback reference that is a curve cutting distance, facilitating (fine-tuned) trajectory tracking on trajectories other than road centreline and 3) explicit preview on road curvature.

DRIVER PREPOSITIONING IS A FUNDAMENTAL CHARACTERISTIC OF DRIVER STEERING THAT SHOULD BE INCLUDED IN DRIVER STEERING AND TRAJECTORY BEHAVIOUR MODELS

Although most parametric driver behaviour models do not account for any prepositioning before curve entry, in this thesis, it is found that prepositioning forms an essential part of curve driving. Prepositioning behaviour is found to occur for 88% of the driving simulator data of *Chapter 5*. Moreover, the extent of prepositioning is directly linked to what style of curve negotiation you may take, shown in *Chapter 4*. It is in the prepositioning phase where drivers obtain their 'first impression' of a haptic shared controller, where they may be inclined to accept or reject the guidance in the curve (*Chapter 7*). The importance of modelling prepositioning is highlighted in *Chapter 3*, where the omission of this phenomenon in the respective control-theoretic models tested resulted in that they can only replicate 17% of the driving population sufficiently. Therefore, prepositioning is very important for individualised driver modelling and, when personalisation is aimed for in driver assistance.

THE BEST MODEL FOR FULL PERSONALISATION IS THE COMBINATION OF THE PREPOSITIONING PATH MODEL AND THE VAN PAASSEN CONTROL-THEORETIC DRIVER MODEL.

When the trajectory that a model produces after it has been identified on a trajectory (of a representative driver) in a particular class, falls within the same class and has a VAF greater than 80% (VAF is a validation metric where 100% means the model is able to describe 100% of the data), then the model is capable of describing that class sufficiently. The model is also assumed to be able to describe all the drivers in that particular class.

Without prepositioning, the Van Paassen model can only describe 17% of drivers (*Chapter 3*). After adding a prepositioning path model extension to the Van Paassen model (*Chapter 6*) 100% of the drivers are sufficiently described.

DRIVER STEERING AND TRAJECTORY BEHAVIOUR IS ASYMMETRIC FOR RIGHT AND LEFT CURVES.

Contrary to the symmetric output of current state-of-art control-theoretic driver steering models, driver trajectory class outcomes are found to be highly asymmetric between right and left curves in all data collected in this thesis. This implies inherently different behaviour on the right and left curves. Evident from *Chapter 4*, where the most popular class outcome for the right curves is class 3 (inner curve negotiation), while for left curves, it is class 5 (outer curve negotiation). Moreover, class consistency (drivers exhibiting the same trajectory class in different types of curves) levels are different, with 18% of drivers being more consistent for right than left curves. In *Chapter 5* the prepositioning displacement for left curves is significantly larger than for right curves for all conditions tested. Even the occurrence of conflict for the symmetric haptic shared controller (Meshed controller) from *Chapter 7* is asymmetric between left and right curves. These differences could stem from the driver already exhibiting a straight section bias, which inherently makes his response to equivalent curves in different directions asymmetric. This intriguing result indicates that driver behaviour modelling, and support, for left and right curves should always be treated independently.

DESPITE FULL PERSONALISATION ACHIEVING TORQUE ALIGNMENT, IT DOES NOT IMPROVE SUBJECTIVE ACCEPTANCE WITH THE FDC HSC.

Personalisation is about the automation understanding the driver's preferences and limitations and adapting to them. In this thesis personalisation is realised through adapting an independent controller reference to the behaviour of the individual driver, i.e., this means that the controller ends up 'mimicking' the driver. There are two levels of personalisation considered in *Chapter 8*, that of adapting to class-average (of the 7-class classifier outlined in *Chapter 4*) and that of full personalisation. Full personalisation is objectively found to minimise torque conflicts, where, torque alignment between controller and driver is significantly improved. However, this torque alignment is not perceived as an added benefit by drivers, as in terms of subjective preference there is no significant difference between full personalisation, class-average personalisation and even centreline guidance. Accommodating to the individual's preferences and limitations is the fundamental principle of personalisation, however, realising this accommodation by mimicking the individual driver's behaviour (as is done with personalisation in this thesis) does not lead to significant improvement. The drivers' preferences are not limited to the controller mimicking their behaviour; therefore, centreline guidance (for the two tested driver groups only) can also be considered personalised.

AS LONG AS THE GUIDANCE TRAJECTORY IS PREDICTABLE TO THE HUMAN AND THE CONTROLLER IS NOT CORRECTING, DRIVERS WILL ACCEPT THE HSC AS A USEFUL AND SATISFACTORY SYSTEM.

From *Chapter 8*, we find that centreline is a predictable trajectory and is equally acceptable to full personalisation, if you ask participants. Analysis shows that this indifference to centreline, although it is not a human-like way to drive (*Chapter 4*), could stem from the centreline being a predictable trajectory and, could also stem from the guidance provided by the feedforward torque. Being predictable is important for acceptance and is one of the human-centred design guidelines *'the automation must be predictable by the human.'* [27]. Whereas, when deviating from the feedforward guidance, unlike the correcting feedback guidance, the driver is not penalised. Hereby the feedforward torque provides a certain 'allowance' to deviate from the reference trajectory. Nevertheless, this does not mean that drivers accept any kind of trajectory in the presence of feedforward guidance, because offset drivers particularly give low ratings to a curve-cutting trajectory (*Chapter 8*). Therefore it seems that for drivers to accept a HSC guidance, the reference trajectory must be predictable and the controller should not (only) be correcting.

9

9.3. PROPOSED DESIGN GUIDELINES

The results of Part IV (Application) of this thesis has lead to three design guidelines that can be used for future implementation of haptic shared control and other driving automation. It should be stressed, that the proposed design guidelines are all based on the findings from this thesis and thus specific experiments, i.e., with fixed-speed driving, performed in a fixed-based simulator (no accelerations felt), exploring a single curve type and, negotiating curves on a single lane road. Moreover, curve negotiation may be different in the presence of road-side furniture and oncoming traffic.

FULL PERSONALISATION (ACCURATELY MIMICKING DRIVER BEHAVIOUR) IS NECESSARY ONLY FOR TORQUE ALIGNMENT

In *Chapter 8*, full personalisation was achieved with a controller that is proven to track its reference (HCR) accurately. The results in that chapter have shown that full personalisation has re-

sulted in the lowest conflict time and conflict torques. Therefore, for applications where torque alignment is explicitly necessary, full personalisation is required. Such systems are, for example, the FDC with a low setting of feedforward, or a pure haptic feedback system. Here the feedforward allowance is low or absent, which enhances the need for torque alignment.

Other examples of systems that may need full personalisation is a lane departure warning system. To avoid false alarms, the corrective prediction to depart the lane needs to be fully personalised. One does not want to be penalised for normal and sufficiently safe behaviour. On the other hand, the (continuous) guidance systems along the curve, like lane-keeping assistance, may only need to fall within the driver's preferences, which, may not be limited to mimicking the driver.

A CENTRELINE REFERENCE TRAJECTORY WITH THE FDC HSC IS AN EFFICIENT HCR IMPLEMENTATION THAT RESULTS IN HIGH ACCEPTANCE

Chapter 8 found that there was no significant effect ($p<0.05$) nor trend ($p<0.2$) between full personalisation, class average personalisation and centreline guidance, for any of the subjective ratings. This means that *for the tested conditions*, drivers do not perceive any difference with centerline guidance compared to personalised, or just do not mind. This does not mean that drivers are indifferent to the controller's reference in general, when an offset driver is given a curve-cutting reference, a clear conflict is present and subjective low ratings are given. Therefore, for an efficient implementation of the FDC, which is equally acceptable as a fully personalised FDC, it is recommended to use the centreline reference.

Nevertheless, this guideline is to be taken with caution, because the finding is valid for the specific road examined in *Chapter 8*, with the particular driver groups (curve cutters and offset drivers amounting to 52% of the driving population together). We find in *Chapter 4* that a driver's trajectory style changes when the curve length and road conditions change. Specifically, more drivers curve-cut when the curve length is shorter. Therefore, it is also likely that a driver's preferences may also change when the road type changes. For example, with a very demanding curve, the centreline may not be preferred any more. Therefore, an explicit experiment would be needed to test how a driver's preferences may change with road type.

A HIGH FEEDFORWARD GAIN IN THE FDC HSC BENEFITS ACCEPTANCE

The feedforward gain (level of haptic support) in the FDC HSC is found to play a key role in driver acceptance. One of the most significant differences between the conventional Meshed shared controller and the implemented FDC controller is the use of an independent feedforward torque (see *Chapter 7*). Analysis (in *Chapter 7 & 8*) has shown that this type of torque is perceived fundamentally different from feedback torque, resulting in less user conflict.

Adding feedforward torque reduces the penalising nature of a haptic shared controller, increasing acceptance (*Chapter 7*). In *Chapter 8* two different feedforward gains are tested: 1) 0.92 and 2) 0.5. Meaning that, whilst negotiating a curve 92% (0.92) or 50% (0.5) of the total torque is the *share* provided by the haptic shared controller. The results show that a feedforward gain of 0.92 results in significantly improved subjective acceptance, usefulness and satisfaction. These improved subjective ratings can be explained through the significant reduction of conflict occurrence, conflict torque and driver torque when applying more feedforward torque.

Furthermore, the haptic shared controller is synonymous to a coach. Where the feedback torques are penalising in nature, like a coach who corrects you, feedforward is guiding in nature, like a coach who can show you how it should be done. Moreover, there is a difference in timing between the two torques, one corrects you after you make a mistake, whereas the other offers anticipatory guidance. The feedforward torque allows the driver to deviate from the controller's

reference without penalty. Hereby, this torque is a double-edged sword, i.e., the favourable result is that the allowance it provides increases acceptance, the unfavourable consequence is that the added benefit of personalisation (by mimicking the driver) is diminished. Nevetheless, this can also be seen as a win-win situation for the higher level goal of acceptance, with a high feedforward gain, higher acceptance is achieved with an easy design. Nevertheless, as a design guideline, a high feedforward gain is advised.

9.4. POSSIBLE APPLICATIONS OF DEVELOPED TOOLS

MODEL ASSESSMENT METHODOLOGY

The model assessment methodology is a tool developed in Part I of this thesis. The standardised assessment methodology from *Chapter 3* comprises two metrics: Descriptiveness and Realistic Identifiability. These two metrics have been developed to assess whether a *control-theoretic* model is suitable for identifying individual behaviour with meaningful parameters. This assessment methodology can be used for other types of models as well. The realistic identifiability, which determines whether the model is capable identifying driver behaviour with meaningful parameters, can also be used with any other modelling framework that exhibits static parameters, such as analytical models [25] [157]. The descriptiveness, which determines the capacity of a model to describe individual driver behaviour, has a much broader application field. Descriptiveness can be used to assess the descriptive capacity of *any* model, be it one based on neural networks [11], deep learning [39] or model predictive control [85].

TRAJECTORY CLASSIFIER

The trajectory classifier is a tool developed in Part II of this thesis. There are two classifiers developed. The first one categorises empirical driver the behaviour, whereas the second one is an extension of the first, to account for trajectories that start exactly at centreline. This is done for categorising the trajectories produced by conventional control-theoretic models whose trajectories always start at the centerline.

The empirical classifier (7-class classifier) was used in *Chapter 8* to generate the class-average HCRs which were necessary to produce class-average personalisation. This type of trajectory-style personalisation using the classifier developed in this thesis is not limited to haptic shared control. In fact, it can be applied to other lane-keeping assistance and lane-departure warning systems. Moreover, it can also be used for autonomous driving. These seven different trajectory categories, or guidance styles, can simply be selected through a knob-like interface, which would then provide explicit personalisation.

INDIVIDUALISED CONTROL-THEORETIC MODEL

The individualised control-theoretic driver model is a developed tool (from Part III) of this thesis. This model is the integration of the newly developed prepositioning path model (from *Chapter 6*) with the Van Paassen driver model (assessed in *Chapter 3*). The integration of the prepositioning path model with the Van Paassen model is given in *Chapter 6*. This novel augmented Van Paassen model has a descriptiveness of 100%, which is the highest achievable descriptive grade. Moreover, the Van Paassen model exhibits the best realistic identifiability (from *Chapter 3*), achieving values lower than 1%, which means that 1% of the realistic parameter space can describe a specified driving behaviour, uniquely. Hereby, this combined model can be used to identify individual drivers in ADAS applications. Not only that, due to the good realistic identifiability, the model can be used to identify driver behaviour in real-time when using an unscented Kalman filter [25] [141]. The identified parameters can infer information about an individual's attentiveness, or fluctuations in steering behaviour, or even a change in trajectory strategy.

PERSONALISABLE FDC HAPTIC SHARED CONTROLLER

The personalisable FDC haptic shared controller is a tool developed in Part IV of this thesis. This tool is capable of following a (personalised) reference trajectory with an accuracy above 80% VAF (*Chapter 8*). This personalised trajectory tracking does not necessarily need to be used with haptics; it can also be used with input-mixing haptic shared controllers [99], i.e., systems with mechanically uncoupled steering systems [112].

Outside the field of shared control, this (personalisable) reference tracking can easily be implemented for full automation by setting the feedforward gain to a value of 1 and increasing the level of haptic authority, such that inputs from the driver do not contribute to steering. Naturally, this tracking algorithm can also be used to establish accurate lane-departure warnings and lane-keeping assistance.

9.5. RECOMMENDATIONS

The research carried out in this thesis brings forward eleven recommendations for further research.

RECOMMENDATION 1: PERSONALISATION AS AN ADAPTATION TO DRIVER PREFERENCE, RATHER THAN TORQUE ALIGNMENT

An essential part of personalisation, is understanding the preferences and limitations of individual drivers and adapting to these. In this thesis, personalisation was realised as the controller mimicking the individual trajectory, i.e., personalisation is achieved through driver controller torque alignment. However, providing personalised guidance is, in fact, broader than this. We find from *Chapter 8* that a driver's preference is not limited to his/her own fully personalised, nor his/her class-average personalised reference. The centreline guidance results in equal acceptance, even though it is not found to be a 'human-like' way to drive (*Chapter 4*). Nevertheless, different driver types are found to have different preferences. A curve cutter is accepting of offset driver guidance, whereas an offset driver does not accept curve-cutting guidance. Therefore, rather than fully personalising reference trajectories, there may be other factors that affect the preference of individual drivers. For example, *Chapter 8* shows that providing more feedforward torque results in a significant improvement in acceptance. It is expected that this is mostly due to the strong feed forward in the FDC controller. This may not hold for corrective HSCs.

Thus, to better understand where, when and how to personalise, it is recommended to investigate other factors such as feedforward gain, feedback gain which may affect acceptance. Moreover, it is also recommended to explore different driver types (other than curve-cutters and offset drivers), to investigate their differing preferences and how these should be accounted for in HSC design.

RECOMMENDATION 2: PERSONALISATION OF ADAS WITH REAL-TIME ADAPTATION USING DRIVER CLASSIFICATION

From *Chapter 8*, we find that information about an individual driver class is essential when predicting whether a driver may subjectively accept or reject certain guidance, i.e., offset drivers particularly give low acceptance ratings for curve-cutting guidance. However, from driver trajectory classification across different curves (*Chapter 4*), we find that drivers drive differently when the road environment changes. Only 63% is consistent in their class throughout all three tested right curves, and 44% is consistent in style across left curves. In real-time inertial classification studies, a single driver can exhibit multiple styles within a single curve [76] [24], whereas,

for the trajectory classification introduced in this thesis, a single driver may exhibit different trajectory styles across different curves.

In fact, some (a minority of) researchers believe that personalisation is only achieved after a cyclic process of *'Understand-Deliver-Measure'* [9]. Translated to the driving domain this would mean 1) understanding the driver (through modelling/classification), 2) delivering the personalised guidance to the driver and 3) measuring the impact and adjusting where needed. Although an iterative approach to personalisation is not done for haptic shared steering control in the automotive domain, an iterative approach or *'cyclic personalisation'* [72] has been attempted for haptic telemanipulation with significant conflict reduction compared to only (statically) personalised [43]. Therefore, rather than classifying or identifying a driver for a single curve, it is recommended to take a real-time approach to take into account the changes that the time-variant, road variant human may undertake.

An experiment to test this should consist of 2 phases, similar to *Chapter 8*, with the same road and fixed velocity profile. In the first phase, the driver's manual style is recorded and classified. During the second phase, the driver has to drive two conditions, one with a static driver reference and another with the real-time updated driver class reference, where a subsequent curve could be driven based on the classification of the previous curve (of the same curve type), this is done to improve mismatch until convergence. Hereby the classification results of the first phase are used to define the static reference trajectory and used as the initial class guidance for the first curve. Alternatively, using reinforcement learning, a trajectory can be iteratively learned by minimising conflict occurrence online (using a conflict measure that is consistent with subjective perception). Such an experiment tests the benefit of iteratively adapting to a driver's trajectory. Whereas for a practical HSC implementation machine learning techniques will be essential in mapping the preferred guidance with the changing road environment.

RECOMMENDATION 3: EXPLICIT PERSONALISATION OF ADAS

Since not all drivers are consistent with their style along the same curve, a real-time personalisation algorithm may cause the driver to mistrust the system. Once the driver has built a mental model of the working principle of the system, his mental model may be continuously falsified, due to the variable nature of a real-time HSC algorithm. Such variable behaviour is naturally expected from humans, but it is not desirable from machines. Such a real-time update may make the guidance less tangible, graspable and understandable (transparent) to the human, which is a key point in Billings human-centred design guideline *'the automation must be predictable by the human.'* [27]. Current studies support the need for the human to understand how the guidance operates [23], e.g., *'the automation is only as good as the user's capability to understand it'* [60].

Since drivers equally accept full personalisation as class average and centreline (*Chapter 8*). The driver can also choose the guidance trajectories himself through 'explicit' personalisation. Here, with a fixed series of trajectory types that the driver can select himself, the driver may quickly develop a mental model of the system that they can trust. Moreover, although this method needs a 'pre-driven' database, it is less computationally burdensome compared to real-time driver model identification. To examine the added benefit of explicit personalisation to implicit personalisation, it is recommended to perform a two-phase experiment as performed in *Chapter 7* wherein the second phase drivers are guided both implicitly and can choose their personalised guidance explicitly.

RECOMMENDATION 4: FINDING THE OPTIMAL LEVEL OF HAPTIC SUPPORT

In *Chapter 8*, it is found that a 92% torque fraction is significantly better than 50% for driver acceptance. However, it is recommended to investigate precisely which torque fraction results in the most optimal acceptance.

An experiment is suggested where six different levels of haptic support are considered. The HCR used can be centreline, given its easy implementation and high acceptance ratings. The feedforward gains must include a gain of 0 (a pure haptic feedback system) and 1 (autonomous driving with shared control). These feedforward gains, however, should be carefully tuned with the strength of haptic feedback. Further reducing the feedforward gain with the current feedback gain strength may bring the system into instability when introducing an impulse response. Therefore, investigating the desired space of 'realistic' gains is also recommended, i.e., stability analysis of the FDC HSC as a function of the feedforward and feedback gains. Moreover, there may be an interaction between level of personalisation and level of haptic support; the weaker the haptic support and the stronger the haptic feedback, the higher the required level of personalisation.

RECOMMENDATION 5: EXPLORE A METRIC THAT CAN PREDICT SUBJECTIVE ACCEPTANCE

Currently, there is a mismatch between perceived subjective acceptance and objective torque conflict computation, as for instance shown in *Chapter 8*. Here, the total haptic shared control torque, which is the summation of feedback and feedforward torques, was used in conflict computation. However, a conflict with feedforward torques is not perceived by the driver as penalising. With this feedforward, drivers could be comfortably 'leaning' on the shared controller, which is recorded as a torque conflict. Conversely, a driver that is clearly 'fighting' a given guidance may receive the same magnitude of torque conflict. Nevertheless, the defining line between 'leaning' and 'struggling' may be related to *conflict frequency*, a dimension not considered in this thesis. It is therefore recommended to further investigate a metric that can accurately predict a driver's subjective ratings. This way, we can truly understand what contributes to acceptance and deduce an objectively measurable metric that can predict whether a specific guidance will be accepted or not, without having to resort to time-consuming subjective questionnaires.

RECOMMENDATION 6: THE EFFECT OF PREPOSITIONING PHASE ON IN-CURVE ACCEPTANCE

In *Chapter 7* it was found that with the FDC HSC, left curves result in less occurrence of conflict than right curves. This is unexpected because the HCR fit (using the Mars model) to right curves was substantially better compared to left curves, with an increase of 27% in VAF on average. The result was explained by the left HCR pulling the driver towards the inner part of the curve at the curve entry phase, thereby abiding the driver's intentions.

These driver 'intentions' are motivated by a general trend (common denominator) found across the majority of drivers. From *Chapter 4* and *Chapter 6* we find that, with exception to severe counter curve-cutting (1.5% of the total curves driven), after curve entry, drivers travel towards the inner part of the curve, even if they don't make it to the inner part of the curve. To facilitate such travel towards the inner part of the curve, the controller reference (HCR) in the prepositioning phase is essential. In *Chapter 7* the controller reference for left curves starts at the centerline, moreover, the majority of drivers preposition towards the outer part of the curve for left curves (*Chapter 6*). Hereby the HCR that starts at centerline for left curves, pulls the

driver towards the inner part of the curve, which is in line with what is comfortable to drivers, thereby decreasing the experienced conflict.

However, this developed interpretation/theory from *Chapter 7* is not confirmed in *Chapter 8*. In *Chapter 8* the highest in-curve conflict for both driver groups stems from the class average inner curve negotiation reference in right curves, which is unexpected. This 'class-average' HCR, (which is actually the class average of drivers from *Chapter 4*), prepositions more to the inner part of the curve than all but one participant. Already prepositioning so far into the inner part of the curve would also facilitate 'travel' towards the inner part of the curve. Nevertheless, careful analysis shows that this increase in conflict results from conflict with the feedforward torque (*Chapter 8*).

Therefore, it is recommended to explicitly investigate the effect of the prepositioning phase on the acceptance and, whether designing the prepositioning phase of the HCR to facilitate 'travel' towards the inner part of the curve is essential to diminish in-curve conflict. I propose a between-subject study, with curve cutters and offset drivers, that tests 5 different trajectories with the FDC HSC. These trajectories preposition, and once in the curve, tend towards the centreline. Therefore, these trajectories are to be generated by the driver model developed in *Chapter 6*. All prepositioning path trajectories should have a standard prepositioning time of 7 s (average prepositioning time found in *Chapter 5*) and a 0 m bias. The first trajectory exhibits outer prepositioning to 0.8 m, the second trajectory: outer prepositioning to 0.4 m, third trajectory: 0 m prepositioning, fourth trajectory: inner prepositioning to 0.4 m, and fifth trajectory: inner prepositioning to 0.8 m.

RECOMMENDATION 7: TRAJECTORY PREFERENCES IN AD

If we consider full automation or autonomous driving (AD), the driver is a 'passive' controller. The closest form to this type of control with HSC is with a feedforward gain of 1. Therefore, the subjectively perceived acceptance with FF of 0.92 could generalise to AD. What is fundamentally different with AD is that the driver is out of the loop and is passive. Conflicts will not be perceived through opposing torques, but through comfort level, i.e., in AD, conflict may manifest as an uncomfortable vestibular cue. With passive driving, driving-styles that are safer and even less aggressive to a driver's style are more preferred [22] [193]. This is expected because as a passive driver, you do not experience sensory cancellation [132], nor can you prepare for the steering inputs given (expectation management). A similar outcome is not found from the fixed-based, fixed speed results in *Chapter 8* however, where fully personalised guidance results in the highest levels of acceptance. Nevertheless, 'driver style' in AD is mostly quantified in terms of felt accelerations and jerk in a manoeuvre [22] [193]. Within the context and limitations of this thesis, less aggressive, i.e., less lateral acceleration in the curve, can be achieved by prepositioning more or slowing down (as found in *Chapter 5*). Therefore, it is recommended to investigate trajectory preferences in AD in either a moving-base driving simulator or a real test track.

RECOMMENDATION 8: INVESTIGATE THE DRIVERS' INDIVIDUAL OBJECTIVES.

If we want to understand the driver properly, as is the sixth guideline of Billings *'the automation must monitor and understand the human'* [27], we need to know how drivers drive in many contexts. The developed classes and control-theoretic modelling techniques are descriptive of behaviour in a particular context. However, these are not generalisable to other contexts. For example, in *Chapter 4*, we find that more drivers take a severe curve cutting trajectory style when negotiating a curve that is more demanding (larger curvature). In *Chapter 5*, we find that drivers preposition more in curves with larger curvature and at larger speeds, as an attempt to improve

their TLCs. However, the developed individualised control theoretic model does not predict this. We don't understand why a particular driver drives in a certain way. Driver models that describe and successfully predict general trends found in driving based on the optimisation of specific costs and driver perceived risks to the environment are available in literature [87]. However, this prediction is not accurate at a personalised level where maybe these costs are even changing with time or, there is some non-linear contextual effects. It is recommended to further develop such models which can predict behaviour in multiple scenarios [87][86], to account for personal differences. Once this is understood, we can generalise a particular driver to different scenarios. Therefore, it is recommended to investigate the objectives of *individual* drivers and whether these may generalise to other contexts.

RECOMMENDATION 9: MOVING TO FREE-SPEED

Throughout this thesis, a fixed-speed assumption is made for analysing curve-driving. Although this may apply to a particular type of driving, i.e., with cruise control on, this assumption is generally unrealistic. The premise is made due to simplification purposes. To be able to model and understand human driver behaviour, fixed-speed scenarios are necessary to avoid the interferences that will arise when executing free speed. From the results of this thesis, we already do see that velocity has an impact. In fact, the class outcomes change with changing velocity (*Chapter 4*) and prepositioning is found to scale with velocity (*Chapter 5*).

Adding another axis of control can cause issues such as task interference, which may clutter the understanding and modelling of steering behaviour alone [20]. Now that a sufficient method of modelling driver steering behaviour is established, it is recommended to analyse the behaviour of longitudinal driver behaviour through driver identification (using the developed model in *Chapter 6*) in combination with the interference it may cause with steering behaviour. This can be carried out in a step-like approach by first investigating different fixed speeds. Then investigating speed-scheduled scenarios whilst scheduling the velocity in the bicycle dynamics model as a linear parameter varying model. Finally investigating free speed without any scheduling.

RECOMMENDATION 10: INVESTIGATE AGE GROUP RELATED BEHAVIOUR, MODELLING AND SUPPORT

Although there is evidence from literature that age affects driving style [127], this thesis has not explicitly analysed or tested different age groups for different outcomes of trajectory styles. It is generally found that older drivers exhibit the safest driving strategies [50], maintaining considerably larger car-following distances [164], but also needing longer decision times [126]. In *Chapter 5* it is found that when comparing a small group of mature drivers (having a drivers license for 40 years) and novice drivers, the mature drivers preposition significantly more, with on average 0.47 m larger prepositioning displacements. Mature drivers may gravitate towards a single (defensive) style, which is essential information that can easily be used for adaptation and personalisation of HSC to increase acceptance for drivers of all ages.

RECOMMENDATION 11: REAL-WORLD DRIVING

Data from a fixed-base simulator can result in an inaccurate representation of how human drivers drive, especially when negotiating curves of different velocity and radii. Apart from visual and proprioceptive (haptic) feedback, humans also respond to vestibular cues [187]. Not being able to sense the vestibular and vibratory feedback, may cause drivers to drive more 'aggressively,' i.e., using trajectories that exhibit inherently larger lateral accelerations and higher jerk levels. Whereas when drivers are able to sense this acceleration, they may want to diminish it by

driving a different trajectory. In the real-world, would drivers preposition more with the perception of lateral acceleration? Would the class division outcomes change? Would drivers still be asymmetric with respect to their trajectory in right and left curves? These are all essential open questions that need to be addressed for successfully personalised HSC implementations in real cars or in a high performance moving base driving simulator.

REFERENCES

[1] Tesla in fatal california crash was on autopilot. https://www.bbc.com/news/world-us-canada-43604440, 2018. BBC.

[2] Using machine learning to develop natural, human-like vehicle control. https://humandrive.co.uk/wp-content/uploads/2018/12/00702_HumanDrive-Case-Study_Updated-Version_V3.pdf, 2018.

[3] Consumers prefer to share control with self-driving vehicles, sae says. https://www.autonews.com/mobility-report/consumers-prefer-share-control-self-driving-vehicles-sae-says, 2019. Automotive News.

[4] Introducing drivepilot: An automated driving system for the highway. https://www.daimler.com/documents/innovation/other/2019-02-20-vssa-mercedes-benz-drive-pilot-a.pdf, 2019.

[5] D. A. Abbink, T. Carlson, M. Mulder, J. de Winter, F. Aminravan, T. Gibo, and E. Boer. A topology of shared control systems—finding common ground in diversity. *IEEE trans. on Human- Machine Systems*, 48(5):509–525, 2018.

[6] D. A. Abbink, M. Mulder, and E. R. Boer. Haptic shared control - smoothly shifting control authority? *Cognition, Work and Technology*, 41(5):1239–1249, 2011.

[7] H. Abdul. Autonomous emergency braking system with potential field risk assessment for frontal collision mitigation. *IEEE Conference on Systems, Process and Control (ICSPC)*, 2017.

[8] N. AbuAli and H. Abou-zeid. Driver behavior modeling: Developments and future directions. *International Journal of Vehicular Technology*, 2016.

[9] G. Adomavicius and A. Tuzhilin. Personalization technologies: A process-oriented perspective. *Commun. ACM*, 48(10):83–90, 2005.

[10] H. Akaike. A new look at the statistical model identification. *IEEE Trans. Autom. Control*, 19(6):716–723, 1974.

[11] A. Aksjonov, P. Nedoma, V. Vodovozov, E. Petlenkov, and M. Herrmann. A novel driver performance model based on machine learning. *15th IFAC Symposium on Control in Transportation Systems*, 51(9):267–272, 2018.

[12] A. Ameyoe, P. Chevrel, E. Le-Carpentier, F. Mars, and H. Illy. Identification of a linear parameter varying driver model for the detection of distraction. *IFAC- PapersOnline*, 48(26):37–42, 2015.

[13] G. S. Aoude, V. R. Desaraju, L. H. Stephens, and J. P. How. Driver behavior classification at intersections and validation on large naturalistic data set. *IEEE Transactions on Intelligent Transportation Systems*, 13(2):724–736, 2012.

[14] I. Bae, J. Moon, and J. Seo. Toward a comfortable driving experience for a self-driving shuttle bus. *Electronics*, 8(9), 2019.

[15] L. Bainbridge. Ironies of automation. *Automatica*, 19(6):775–779, 1983.

[16] M. Barbero, F. Jouault, J. Gray, and B. J. Model driven architecture-foundations and applications. *Lecture Notes in Computer Science, vol 4530. Springer, Berlin, Heidelberg*, 2007.

[17] S. Barendswaard, D. Pool, and D. Abbink. A method to assess individualized driver models: Descriptiveness, identifiability and realism. *Transportation Research Part F: Psychology and Behaviour*, 61(1):16–29, 2019.

[18] S. Barendswaard, D. Pool, E. Boer, and D. Abbink. A classification method for driver trajectories during curve-negotiation. *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, 2019.

[19] S. Barendswaard, D. M. Pool, and M. Mulder. Human Crossfeed in Dual-Axis Manual Control with Motion Feedback. In *Proc. of the 13th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Human-Machine Systems, Kyoto, Japan*, 2016.

[20] S. Barendswaard, D. M. Pool, M. M. Van Paassen, and M. Mulder. Dual-axis manual control: Performance degradation, axis asymmetry, crossfeed, and intermittency. *IEEE Transactions on Human-Machine Systems*, 49(2):113–125, 2019.

[21] S. Barendswaard, L. van Breugel, B. Schelfaut, J. Sluijter, L. Zuiker, D. Pool, E. Boer, and D. Abbink. Effect of velocity and curve radius on driver steering behaviour before curve entry. *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, 2019.

[22] C. Basu, Q. Yang, D. Hungerman, M. Singhal, and A. Dragan. Do you want your autonomous car to drive like you? *in Proc. ACM/IEEE Int. Conf. Human-Robot Interact.*, pages 417–425, 2017.

[23] M. Beggiato, M. Pereira, T. Petzoldt, and J. Krems. Learning and development of trust, acceptance and the mental model of acc. a longitudinal on-road study. In *Transportation Research Part F: Traffic Psychology and Behaviour*, volume 35, pages 75–84, 2015.

[24] A. Bender, G. Agamennoni, J. Ward, S. Worrall, and E. Nebot. An unsupervised approach for inferring driver behavior from naturalistic driving data. *IEEE transactions on intelligent transportation systems*, 16(6):3325–3336, 2015.

[25] M. Best. Real-time characterisation of driver steering behaviour. *Vehicle System Dynamics*, 57(1):64–85, 2019.

[26] L. Bi, M. Wang, C. Wang, and Y. Liu. Development of a driver lateral control model by integrating neuromuscular dynamics into the queuing network-based driver model. *IEEE Trans. Intell. Transp. Syst.*, 16(5):2479–2486, 2015.

[27] C. Billings. Human-centered aircraft automation: A concept and guidelines. *NASA Technical Memorandum 103885. Moffett Field, CA: NASA- Ames Research Center*, 1996.

[28] C. Bingham, C. Walsh and S. Carroll. Impact of driving characteristics on electric vehicle energy consumption and range. *IET Intelligent Transportation Systems*, 6(1):29–35, 2012.

[29] E. Boer. Satisficing curve negotiation: Explaining drivers' situated lateral position variability. *IFAC papers OnLine*, 49(19):183–188, 2016.

[30] E. R. Boer. Tangent point oriented curve negotiation. In *Proceedings of the 1996 IEEE Intelligent Vehicles Symposium*, pages 7–12, Sept. 1996.

[31] R. Boink, M. M. van Paassen, M. Mulder, and D. A. Abbink. Understanding and Reducing Conflicts Between Driver and Haptic Shared Control. In *Proc. of the IEEE Int. Conf. on Systems, Man and Cybernetics, San Diego (CA)*, pages 1510–1515, 2014.

[32] T. Brandt, T. Sattel, and M. Böhm. Combining haptic human-machine interaction with predictive path planning for lane-keeping and collision avoidance systems. In *Proceedings of the IEEE Intelligent Vehicles Symposium, Istanbul, Turkey*, pages 582–587, 2007.

[33] S. Breznitz. Cry wolf: The psychology of false alarms. In *Lawrence Erlbaum*, pages Hillsdale, NJ, 1984.

[34] V. Butakov and P. Ioannou. Driving autopilot with personalization feature for improved safety and comfort. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pages 387–393, 2015.

[35] G. Buyukyildiz, O. Pion, C. Hildebrandt, M. Sedlmayr, R. Henze, and F. Küçükay. Identification of the driving style for the adaptation of assistance systems. *Int. J. of Vehicle Autonomous Systems*, 13(3):244–260, 2017.

[36] C. Cabrall, A. Eriksson, F. Dreger, R. Happee, and J. De Winter. How to keep drivers engaged while supervising driving automation? a literature survey and categorization of six solution areas. *Theoretical Issues in Ergonomics Science*, 2018.

[37] M. Caglar, A. Teufel, and C. Wilke. Sicegar: R package for sigmoidal and double-sigmoidal curve fitting. *PeerJ*, 6(4251), 2018.

[38] J. Campbell, M. Lichty, J. Brown, C. Richard, J. Graving, J. Graham, M. O'Laughlin, D. Torbic, and D. Harwood. Nchrp report 600: Human factors guidelines for road systems. *Transportation Research Board of the National Academies: Washington, DC, USA*, 2012.

[39] J. Chen, Z. Wu, J. Zhang, and S. Chen. Driver identification based on hidden feature extraction by using deep learning. In *2019 IEEE 3rd Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, pages 1765–1768, 2019.

[40] D. Cole. A path-following driver–vehicle model with neuromuscular dynamics. *Veh. Syst. Dyn.*, 50(4):573–596, 2012.

[41] Z. Constantinescu, C. Marinoiu, and M. Vladoiu. Driving style analysis using data mining techniques. *International Journal Computer Communication Control*, 5(5):654–663, 2012.

[42] A. Cornet, H. Deubener, T. Möller, P. Schaufuss, and A. Tschiesner. Race 2050 – a vision for the european automotive industry. *McKinsey centre for future mobility*, 2019.

[43] A. W. De Jonge, J. G. W. Wildenbeest, H. Boessenkool, and D. A. Abbink. The effect of trial-by-trial adaptation on conflicts in haptic shared control for free-air teleoperation tasks. *IEEE Trans. Haptics*, 9(1):111–110, 2016.

[44] S. Y. de Nijs, M. Mulder, and D. A. Abbink. The value of haptic feedback in lane keeping. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 3599–3604, 2014.

[45] J. C. F. De Winter and D. Dodou. Preparing drivers for dangerous situations: A critical reflection on continuous shared control. In *IEEE Int. Conf. Syst. Man, Cybern.*, pages 1050–1056, 2011.

[46] W. Dib, A. Chasse, P. Moulin, A. Sciarretta, and G. Corde. Optimal energy management for an electric vehicle in eco-driving applications. *Control Engineering Practice*, 29:299–307, 2014.

[47] G. D'Intino, M. Olivari, and H. Bülthoff. Haptic assistance for helicopter control based on pilot intent estimation. *Journal Of Aerospace Information Systems*, -(-):509–525, 2020.

[48] E. Donges. A Two-Level Model of Driver Steering Behavior. *Human Factors*, 20(6):691–707, dec 1978.

[49] F. M. Drop, D. M. Pool, M. M. van Paassen, M. Mulder, and H. H. Bülthoff. Objective Model Selection for Identifying the Human Feedforward Response in Manual Control. *IEEE Trans. on Cybernetics*, 2016. Online preprint.

[50] D. Eby, D. Trombley, L. Molnar, and J. Shope. The assessment of older drivers' capabilities: A review of the literature. *Technical report, The University of Michigan - Transportation Research Institute*, (http://hdl.handle.net/2027.42/1245), 1998.

[51] G. Effler. Customer demand for safety technology threatened by overbearing alerts. *Tech Experience Study, J.D. Power*, 2019.

[52] A. Erséus, L. Drugge, and A. Trigell. A path tracking driver model with representation of driving skill. *Int. J. Vehicle Systems Modelling and Testing*, 6(2):81–90, 2011.

[53] T. Evas. A common eu approach to liability rules and insurance for connected and autonomous vehicles. *European Parliamentary Research Service*, 2018.

[54] H. Fan and M. Poole. What is personalization? perspectives on the design and implementation of personalization in information systems. *Journal of Organizational Computing and Electronic Commerce*, 16(3-4):179–202, 2006.

[55] M. Feniche and T. Mazri. Lane detection and tracking for intelligent vehicles: A survey. *IEEE ICCSRE*, pages 978–1–7281–0827–8/19, 2019.

[56] M. Flad, C. Trautmann, G. Diehm, and S. Hohmann. Experimental validation of a driver steering model based on switching of driver specific primitives. *IEEE International Conference on Systems, Man, and Cybernetics*, page 214–220, 2013.

[57] F. Flemisch, C. Adams, S. Conway, K. H. Goodrich, M. T. Palmer, and P. C. Schutte. The h-metaphor as a guideline for vehicle automation and interaction. *NASA, Hampton, VA, USA, Tech. Rep.*, 212672, 2003.

[58] F. Flemisch, M. Heesen, J. Kelsch, J. Schindler, C. Preusche, and J. Dittrich. Shared and cooperative movement control of intelligent technical systems: Sketch of the design space of haptic-multimodal coupling between operator, co-automation, base system and environment. In *IFAC Proceedings*, volume 43, pages 304–309, 2010.

[59] F. Flemisch, J. Kelsch, C. Loper, A. Schieben, J. Schindler, and M. Heesen. Cooperative control and active interfaces for vehicle assistance and automation. In *FISITA World Automotive Congress*, 2008.

[60] Y. Forster, S. Hergeth1, F. Naujoks1, M. Beggiato, J. Krems, and A. Keinath1. Learning and development of mental models during interactions with driving automation: A simulator study. In *Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design*, volume 35, 2019.

[61] U. Fugiglando, E. Massaro, P. Santi, S. Milardo, K. Abida, R. Stahlmann, F. Netter, and C. Ratti. Driving behavior analysis through can bus data in an uncontrolled environment. *IEEE Transactions on Intelligent Transportation Systems*, 20(2):737–748, 2019.

[62] J. J. Gibson. Perception of the visual world. *Houghton Mifflin, Boston*, 1950.

[63] H. Godthelp, P. Milgram, and G. Blaauw. The development of a time-related measure to describe driving strategy. *Human Factors*, 26(3):257–268, 1984.

[64] V. Govindarajan, K. Driggs-Campbell, and R. Bajcsy. Affective driver state monitoring for personalized, adaptive adas. *21st International Conference on Intelligent Transportation Systems*, pages 1017–1022, 2018.

[65] A. Gray, M. Ali, Y. Gao, J. K. Hedrick, and F. Borrelli. A unified approach to threat assessment and control for automotive active safety. *IEEE Trans. Int. Trans. Syst.*, 14(3):1490–1499, 2013.

[66] S. Griesche, E. Nicolay, D. Assmann, M. Dotzauer, and D. Kätner. Should my car drive as i do? what kind of driving style do drivers prefer for the design of automated driving functions? In *Braunscheiger Symposium Automatisieringssysteme, Assistenzsysteme und eingebettete Systeme für Transportmittel (AAET)*, pages 185–204. Deutches Zentrum für Luft- und Raumfahrt e.V., February 2016. ISBN 978-3-937655-37-6.

[67] P. G. Griffiths and R. B. Gillespie. Sharing control between humans and automation using haptic interface: Primary and secondary task performance benefits. *Hum. Factors J. Hum. Factors Ergon. Soc.*, 47(3):574–590, 2005.

[68] V. Gruppelaar, M. van Paassen, M. Mulder, and D. Abbink. A perceptually inspired driver model for speed control in curves. In *IEEE International Conference on Systems, Man, and Cybernetics: Miyazaki, Japan*, page 1253–1258, 2018.

[69] T. Gu, J. Dolan, and J. Lee. Human-like planning of swerve maneuvers for autonomous vehicles. *IEEE Intell. Vehicles Symp.*, pages 716–721, 2016.

[70] H. Guo, C. Shen, H. Zhang, H. Chen, and R. Jia. Simultaneous trajectory planning and tracking using an mpc method for cyber-physical systems: A case study of obstacle avoidance for an intelligent vehicle. *IEEE Transactions on Industrial Informatics*, 14(5):4273–4283, 2018.

[71] F. Hartwich, M. Beggiato, and J. Krems. Driving comfort, enjoyment and acceptance of automated driving – effects of drivers' age and driving style familiarity. *Ergonomics*, 61(8):1017–1032, 2018.

[72] M. Hasenjäger, M. Heckmann, and H. Wersing. A survey of personalization for advanced driver assistance systems. *IEEE Transactions on Intelligent Vehicles*, 5(2):335–344, 2020.

[73] Here.com. Consumer acceptance of autonomous vehicles. *3 Key Insights for the Automotive Industry*, 2017.

[74] P. Hermannstädter and B. Yang. Driver distraction assessment using driver modeling. *IEEE Proceedings on System, Man, Cybernetics*, page 1502–1508, 2013.

[75] R. A. Hess and A. Modjtahedzadeh. A control theoretic model of driver steering behavior. *IEEE Control Systems Magazine*, 10(5):3–8, Aug. 1990.

[76] B. Higgs and M. Abbas. Segmentation and clustering of car-following behavior: Recognition of driving patterns. *IEEE Transactions on Intelligent Transportation Systems*, 16(1):81–90, 2015.

[77] F. Holzmann, F. Flemisch, R. Siegwart, H. Bubb, C. Preusche, and J. Dittrich. From aviation down to vehicles—integration of a motions-envelope as safety technology. In *SAE Technical Paper*, 2006.

[78] S. Hoogendoorn, R. Hoogendoorn, M. Wang, and W. Daamen. Modeling driver, driver support, and cooperative systems with dynamic optimal control. *Transportation Research Record*, 2316(1):20–30, 2012.

[79] T. Inagaki. Smart collaboration between humans and machines based on mutual understanding. *Annual Reviews in Control*, 32(2):253–261, 2008.

[80] M. Itoh and D. Flemisch, F. anf Abbink. A hierarchical framework to analyze shared control conflicts between human and machine. *IFAC-PapersOnLine*, 49(19):96–101, 2016.

[81] P. Jing, G. Xu, Y. Chen, Y. Shi, and F. Zhan. The determinants behind the acceptance of autonomous vehicles: A systematic review. *Sustainability*, 12(5):1719, 2020.

[82] S. Kamat. Lane keeping of vehicle using model predictive control. In *5th International Conference for Convergence in Technology*, 2019.

[83] A. Kamble and S. Potadar. Lane departure warning system for advanced drivers assistance. In *Second International Conference on Intelligent Computing and Control Systems (ICICCS)*, pages 1775–1778, 2018.

[84] S. Keen and D. Cole. Bias-free identification of a linear model prediction steering controller from measured driver steering behavior. *IEEE Trans. Syst., Man, Cybern.*, 42(2):434–443, 2012.

[85] N. Kim and D. J. Cole. A model of driver steering control incorporating the driver's sensing of steering torque. *Vehicle System Dynamics*, 49(10):1575–1596, 2011.

[86] S. Kolekar, J. de Winter, and D. Abbink. A human-like steering model: Sensitive to uncertainty in the environment. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1487–1492, 2017.

[87] S. Kolekar, J. de Winter, and D. Abbink. Human-like driving behaviour emerges from a risk-based driver model. *Nature Communications*, 11(4850), 2020.

[88] M. Konig and L. Neumayr. Users' resistance towards radical innovations: The case of the self-driving car. *Transportation Research Part F: Traffic Psychology and Behaviour*, 44:42–52, 2017.

[89] E. Krendel and D. McRuer. A servomechanisms approach to skill development. *Journal of the Franklin Institute*, 269(1):24–42, 1960.

[90] A. Kumar, R. Gourishetti, and M. Manivannan. Mechanics of pseudo-haptics with computer mouse. In *2017 IEEE International Symposium on Haptic, Audio and Visual Environments and Games (HAVE)*, pages 1–6, 2017.

[91] M. Kyriakidis, J. C. F. de Winter, N. Stanton, T. Bellet, B. van Arem, K. Brookhuis, K. Martens, M. H. Bengler, J. Andersson, N. Merat, N. Reed, M. Flament, M. Hagenzieker, and R. Happee. A human factors perspective on automated driving. *Theoretical Issues in Ergonomics Science*, 20(3), 2019.

[92] G. Learmonth, G. Marker, N. McBride, P. Pellinen, and M. Harvey. Right-lateralised lane keeping in young and older british drivers. *PLoS ONE*, 13(9), 2017.

[93] K. Lee, K. Kerns, R. Bone, and M. Nickelson. Development and validation of the controller acceptance rating scale (cars): Results of empirical research. *Proceedings of the 4th US-A/Europe Air Traffic Management R&D Seminar*, 2001.

[94] S. Lef'evre, A. Carvalho, Y. Gao, H. E. Tseng, and F. Borrelli. Driver models for personalized driving assistance. *Vehicle System Dynanmics*, 53(12):1705–1720, 2015.

[95] T. Levermore, A. Ordys, and J. Deng. A review of driver modelling. In *2014 UKACC International Conference on Control (CONTROL)*, pages 296–300, 2014.

[96] A. Li, H. Jiang, Z. Li, J. Zhou, and X. Zhou. Human-like trajectory planning on curved road: Learning from human drivers. *IEEE Transactions on Intelligent Transportation Systems*, (Online Preprint), 2019.

[97] A. Li, H. Jiang, J. Zhou, and X. Zhou. Learning human-like trajectory planning on urban two-lane curved roads from experienced drivers. *IEEE Conference on Industrial Electronics and Applications*, pages 567–572, 2019.

[98] Y. Li, Y. Nan, J. He, Q. Feng, J. Zhang, and J. Fan. Study on lateral assisted control for commercial vehicles. *IEEE Conference on Industrial Electronics and Applications*, pages 567–572, 2019.

[99] Z. Li, S. Zhao, J. Duan, C. Su, C. Yang, and X. Zhao. Human cooperative wheelchair with brain–machine interaction based on shared control strategy. *IEEE/ASME Transactions on Mechatronics*, 22(1):185–195, 2017.

[100] Z. Li, S. Zhao, J. Duan, C. Su, C. Yang, and X. Zhao. Intention to use a fully automated car: Attitudes and a priori acceptability. *IEEE/ASME Transactions on Mechatronics*, 22(1):185–195, 2017.

[101] Y. Lilis, E. Zidianakis, N. Partarakis, S. Ntoa, and C. Stephanidis. A framework for personalised hmi interaction in adas systems. *Proceedings of the 5th International Conference on Vehicle Technology and Intelligent Transport Systems*, pages 586–593, 2019.

[102] C. C. Lin, S. Jeon, H. Peng, and J. M. Lee. Driving pattern recognition for control of hybrid electric trucks. *Vehicle System Dynamics*, 42(1):41–58, 2004.

[103] N. Lin, C. Zong, M. Tomizuka, P. Song, Z. Zhang, and G. Li. An overview on study of identification of driver behavior characteristics for automotive control. *Mathematical Problems in Engineering*, page Article ID 569109, 2014.

[104] T. Litman. Autonomous vehicle implementation predictions: Implications for transport planning. *Transportation Research Board (TRB) 94th Annual Meeting. Washington, D.C.*, 2020.

[105] A. Liu and D. Salvucci. Modeling and prediction of human driver behavior. *Proceedings of the 9th International Conference on Human-Computer Interaction*, 2001.

[106] M. López-Lambas and A. Alonso. The driverless bus: An analysis of public perceptions and acceptability. *Sustainability*, 11(18):4986, 2019.

[107] M. V. Ly, S. Martin, and M. M. Trividi. Driver classification and driving style recognition using inertial sensors. In *IEEE Intelligent Vehicles Symposium, Gold Coast, Australia*, page 1040–1045, 2013.

[108] C. MacAdam. Application of an optimal preview control for simulation of closed-loop automobile driving. *IEEE Trans. on Systems, Man, and Cybernetics*, 11(6):393 – 399, 1981.

[109] C. Macadam. Understanding and modeling the human driver. *Vehicle System Dynamics*, 40(1–3):101–134, 2003.

[110] S. Mammar, S. Glaser, and M. Netto. Time to line crossing for lane departure avoidance: a theoretical study and an experimental setting. *IEEE Transactions Intelligent Transportation Systems*, 7(2):226 – 241, 2006.

[111] A. A. Maranga, H. Zhu, D. Gingras, and D. Gruyer. Collision avoidance improvement using a risky area. In *2014 IEEE 12th International New Circuits and Systems Conference (NEWCAS)*, pages 297–300, 2014.

[112] M. Marcano, S. Díaz, J. Pérez, and E. Irigoyen. A review of shared control for automated vehicles: Theory and applications. *IEEE Transactions on Human-Machine Systems*, pages 1–17, 2020.

[113] B. Marr. Tech trends in practice: The 25 technologies that are driving the 4th industrial revolution. *John Wiley and Sons Inc*, isbn: 9781119646198, 2020.

[114] F. Mars and F. Chevrel. Modelling human control of steering for the design of advanced driver assistance systems. *Annual Reviews in Control*, 44:292–302, 2017.

[115] F. Mars, M. Deroo, and J. Hoc. Analysis of human-machine cooperation when driving with different degrees of haptic shared control. *IEEE Transactions on Haptics*, 7(3):324–333, 2014.

[116] F. Mars, L. Saleh, F. Chevrel, F. Claveau, and J. Lafay. Modeling the visual and motor control of steering with an eye to shared-control automation. In *Human Factors and Ergonomics Society, 55th Annual Meeting*, pages 1422–1426, 2011.

[117] F. Mauriello, A. Montella, M. Pernetti, and F. Galante. An exploratory analysis of curve trajectories on two-lane rural highways. *Sustainability*, 10(4248):1–16, Nov. 2018.

[118] J. McLean and E. Hoffmann. Steering reversals as a measure of driver performance and steering task difficulty. *Human Factors*, 17:248 —- 256, 1975.

[119] D. McRuer, R. Allen, and D. Weir. New results in driver steering control models. *Human Factors*, 19(4):381–397, 1977.

[120] D. T. McRuer and H. R. Jex. A Review of Quasi-Linear Pilot Models. *IEEE Transactions on Human Factors in Electronics*, HFE-8(3):231–249, Sept. 1967.

[121] T. Melman, D. Abbink, M. Van Paassen, E. R. Boer, and J. de Winter. What determines drivers' speed? a replication of three behavioural adaptation experiments in a single driving simulator study. *Ergonomics*, 61(7):966–987, 2018.

[122] N. Merat and A. H. Jamson. How do drivers behave in a highly automated car? In *Proc. Fifth Int. Driv. Symp. Hum. Factors Driv. Assessment*, pages 514 – 521, 2008.

[123] N. Merat and J. Lee. Preface to the special section on human factors and automation in vehicles: Designing highly automated vehicles with the driver in mind. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 54(5):681–686, 2012.

[124] D. R. Mestre, F. Mars, S. Durand, F. Vienne, and S. Espie. A Visual Aid for Curve Driving. In *Proc. of the D.S.C. , Paris*, pages 311–320, Sept. 2004.

[125] J. Michon. A critical view of driver behavior models: What do we know, what should we do? *Human Behavior and Traffic Safety*, page Plenum, 1985.

[126] D. Middleton, H. Westwood, J. Robson, and D. Kok. Assessment and decision criteria for driving competence in the elderly. *Traffic and Transport Psychology: Theory and Application*, pages 101–113, 2005.

[127] L. Molnar, D. Eby, J. Charlton, J. Langford, S. Koppel, S. Marshall, and M. Man-Son-Hing. Driving avoidance by older adults: Is it always self-regulation? *Accident Analysis and Prevention*, 57:96–103, 2001.

[128] M. Mulder and D. Abbink. Sharing control with elderly drivers: Haptic guidance during curve negotiation. *11th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Human-Machine Systems*, 43(13):310–315, 2010.

[129] M. Mulder, D. Abbink, and E. Boer. The effect of haptic guidance on curve negotiation behavior of young, experienced drivers. In *IEEE Int. Conf. on Systems, Man, and Cybernetics (SMC)*, pages 804–809, Nov. 2008.

[130] Y. Murphey, R. Milton, and L. Kiliaris. Driver's style classification using jerk analysis. *IEEE Workshop on Computational Intelligence in Vehicles and Vehicular Systems*, 2009.

[131] H. Muslim, M. Itoh, and M. Pacaux-Lemoine. Driving with shared control: How support system performance impacts safety. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 000582–000587, 2016.

[132] C. Nash, D. Cole, and R. Bigler. A review of human sensory dynamics for application to models of driver steering and speed control. *Biological Cybernetics*, 110:91–116, 2016.

[133] A. Odhams and D. Cole. Identification of the steering control behaviour of five test subjects following a randomly curving path in a driving simulator. *Int. J. Veh Auton Syst*, 12(1):44–46, 2014.

[134] Z. Ouyang, J. Niu, Y. Liu, and J. Rodrigues. Multiwave: A novel vehicle steering pattern detection method based on smartphones. In *2016 IEEE International Conference on Communications (ICC)*, pages 1–7, 2016.

[135] R. Parasuraman and V. Riley. Humans and automation: Use, misuse, disuse, abuse. *Hum. Factors J. Hum. Factors Ergon. Soc.*, 39(2):230–253, 1997.

[136] M. D. Penna, M. M. van Paassen, D. A. Abbink, M. Mulder, and M. Mulder. Reducing steering wheel stiffness is beneficial in supporting evasive maneuvers. In *2010 IEEE International Conference on Systems, Man and Cybernetics*, pages 1628–1635, 2010.

[137] S. M. Petermeijer, D. A. Abbink, M. Mulder, and J. C. F. De Winter. The effect of haptic support systems on driver performance : A literature survey. *IEEE Trans. Haptics*, 8(4):467–479, 2015.

[138] G. Prokop. Modeling human vehicle driving by model predictive online optimization. *Vehicle System Dynamics*, 35(1):19–53, 2001.

[139] A. Renski. Identification of driver model parameters. *Int. J. Of Occupational Safety And Ergonomics*, 7(1):79 — 92, Aug. 2001.

[140] Rijkswaterstaat. Handboek geometrisch weg ontwerp. *ROA, CROW KennisPlatform*, 7.4.7 Horizontaal alignement, 2012.

[141] J. Rojer, D. M. Pool, M. M. van Paassen, and M. Mulder. UKF-based Identification of Time-Varying Manual Control Behaviour. In *Proceedings of the 14th IFAC Symposium on Analysis Design and Evaluation of Human Machine Systems Tallinn, Estonia*, pages 109–114, 2019.

[142] R. Rossi, G. Gecchele, M. Gastaldi, F. Biondi, and C. Mulatti. An advanced driver assistance system for improving driver ability. design and test in virtual environment. In *2017 5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, pages 509–513, 2017.

[143] F. Sagberg, G. F. B. Piccinini, and J. Engstrom. A review of research on driving styles and road safety. *Human Factors Ergonomics Society*, 57(7):1248–1275, 2015.

[144] Y. Saito, T. Mitsumoto, and P. Raksincharoensak. Effectiveness of a Risk Predictive Shared Steering Control Based on Potential Risk Prediction of Collision with Vulnerable Road Users. In *Proc. of the 13th IFAC Symposium on Analysis, Design, and Evaluation of Human-Machine Systems, Kyoto, Japan*, pages 84–89, 2016.

[145] L. Saleh, P. Chevrel, F. Claveau, J. Lafay, and F. Mars. Shared steering control between a driver and an automation: Stability in the presence of driver behavior uncertainty. *IEEE Trans. Int. Trans. Syst.*, 14:974–983, 2013.

[146] L. Saleh, P. Chevrel, F. Mars, J. Lafay, and F. Claveau. Proceedings of the 18th ifac world congress. *Human-like cybernetic driver model for lane keeping*, pages 4368–4373, 2011.

[147] D. Salvucci. Modeling driver behavior in a cognitive architecture. *Human Factors*, 48(2):362–380, 2006.

[148] F. Santoni de Sio and J. van den Hoven. Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI*, page 15, 2018.

[149] S. Scherer. How the driver wants to be driven – modelling driving styles in highly automated driving. *Automatisiertes Fahren–Hype oder mehr? Die 7. Tagung Fahrerassistenz*, 7(4):2015–2026, 2015.

[150] S. Schnelle, J. Wang, H. Su, and R. Jagacinski. A driver steering model with personalized desired path generation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(1):111–120, 2017.

[151] W. Schofield. Engineering surveying – theory and examination problems for students. In *Butterworth-Heinemann, Oxford*, 2001.

[152] W. Scholtens, S. Barendswaard, D. Pool, M. Van Paassen, and D. Abbink. A new haptic shared controller reducing steering conflicts. *IEEE Systems, Man, Cybernetics*, In Press, Online Available, 2018.

[153] K. M. Schwab. The fourth industrial revolution. *Random House Usa Inc*, isbn: 9781524758868, 2017.

[154] W. Schwarting, J. Alonso-Mora, and D. Rus. Planning and decision-making for autonomous vehicles. *Annual Review of Control, Robotics, and Autonomous Systems*, 1:187–210, 2018.

[155] C. Sentouh, P. Chevrel, and F. Mars, F. and. Claveau. A sensorimotor driver model for steering control. *IEEE int. confere. on Systems, Man and Cybernetics*, page 2462 –2467, 2009.

[156] A. Shariff, J.-F. Bonnefon, and I. Rahwan. Psychological roadblocks to the adoption of self-driving vehicles. *Nature, Human Behaviour*, 1(10):694–696, 2017.

[157] R. Sharp, D. Casanova, and P. Symonds. A mathematical model for driver steering control, with design, tuning and performance results. *Vehicle System Dynamics*, (33):289 —- 326, 2000.

[158] T. B. Sheridan and R. Parasuraman. Human-automation interaction. *Reviews of Human Factors and Ergonomics*, 1(1):80–129, 2005.

[159] B. Shi, L. Xu, J. Hu, Y. Tang, H. Jiang, W. Meng, and H. Liu. Evaluating driving styles by normalizing driving behavior based on personalized driver modeling. *IEEE Transactions on System, Man, Cybernetics*, 45(12):1502–1508, 2015.

[160] V. Shia, Y. Gao, R. Vasudevan, K. Driggs-Campbell, T. Lin, F. Borrelli, and R. Bajcsy. Semiautonomous vehicular control using driver modeling. *IEEE Transactions on Intelligent Transportation Systems*, 15(5):2696–2709, 2014.

[161] J. Smisek, W. Mugge, J. B. J. Smeets, M. M. van Paassen, and A. Schiele1. Adapting haptic guidance authority based on user grip. *IEEE International Conference Systems, Man, and Cybernetics*, pages 1516 – 1521, 2014.

[162] P. Spacek. Curve-driving typology. *J. Of Transportation Engineering*, 131(9):669–676, Sept. 2005.

[163] J. Steen, H. J. Damveld, R. Happee, M. M. van Paassen, and M. Mulder. A Review of Visual Driver Models for System Identification Purposes. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2011*, pages 2093 – 2100, Oct. 2011.

[164] D. Strayer and F. Drews. Profiles in driver distraction: Effects of cell phone conversations on younger and older drivers. *Human Factors*, 46(4):640–649, 2004.

[165] M. Sundbom, P. Falcone, and J. Sjoberg. Online driver behavior classification using probabilistic arx models. In *16th International IEEE Annual Conference of Intelligent Transportation Systems, Hague, The Netherlands*, page 1107–1112, 2013.

[166] M. Sysoev, A. Kos, J. Guna, and M. Pogačnik. Estimation of the driving style based on the users' activity and environment influence. *Sensors (Basel)*, 17(10), 2017.

[167] S. Tada and T. Wada. Haptic shared control for backward parking and its effect on skill increase in novice drivers. *IEEE/SICE International Symposium on System Integration (SII)*, -(-), 2015.

[168] Y. Takada, E. Boer, and T. Sawaragi. Driver assist system for human–machine interaction. *Cognition Technology and Work*, 19(5):819–836, 2017.

[169] K. K. Tsoi, M. Mulder, and D. A. Abbink. Balancing safety and support: Changing lanes with a haptic lane-keeping support system. In *2010 IEEE International Conference on Systems, Man and Cybernetics*, pages 1236–1243, 2010.

[170] B. Turner, J. Woolley, and P. Cairney. An analysis of driver behaviour through rural curves: Exploratory results on driver speed. *Australasian Road Safety Conference*, 2015.

[171] L. O. Valencia-Rosado, D. Rojas-Velázquez, and G. Etcheverry. Driver intent data analysis classification. In *International Conference on Electronics, Communications and Computers (CONIELECOMP)*, pages 207–211, 2018.

[172] D. Van Baelen, J. Ellerbroek, M. M. van Paassen, and M. Mulder. Design of a Haptic Feedback System for Flight Envelope Protection. In *Proceedings of the AIAA Modeling and Simulation Technologies Conference, Kissimmee (FL)*, number AIAA 2018-0117, 2018.

[173] K. van der El, D. Pool, M. van Paassen, and M. Mulder. Identification and modeling of driver multiloop feedback and preview steering control. *IEEE International Conference on Systems, Man, and Cybernetics*, 2018.

[174] K. van der El, D. M. Pool, and M. Mulder. Measuring and modeling driver steering behavior: From compensatory tracking to curve driving. *Transportation Research Part F*, 61:337–346, 2018.

[175] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder. A Unifying Theory of Driver Perception and Steering Control on Straight and Winding Roads. *IEEE Transactions on Human-Machine Systems*, 2019.

[176] J. D. Van der Laan, A. Heino, and D. De Waard. A simple procedure for the assessment of acceptance of advanced transport telematics. *Transportation Research Part C: Emerging Technologies*, 5(1):1 – 10, 1997.

[177] M. M. Van Paassen, R. Boink, D. A. Abbink, M. Mulder, and M. Mulder. Four design choices for haptic shared control. *Adv. Aviat. Psychol.*, pages 237 – 254, 2017.

[178] W. Vreugdenhil, S. Barendswaard, D. Abbink, C. Borst, and S. Petermeijer. Complementing haptic shared control with visual feedback for obstacle avoidance. *IFAC PapersOnLine*, 52(19):371–376, 2019.

[179] A. Wahab, C. Quek, C. K. Tan, and K. Takeda. Driving profile modeling and recognition based on soft computing approach. *IEEE Transactions on Neural Networks*, 20(4):563–582, 2009.

[180] W. Wang, J. Xi, and H. Chen. Modeling and recognizing driver behavior based on driving data: A survey. *Mathematical Problems in Engineering*, pages 1–20, 2014.

[181] W. Wang, J. Xi, A. Chong, and L. Li. Driving style classification using a semisupervised support vector machine. *IEEE Transactions on Human-Machine Systems*, 47(5), 2017.

[182] W. Wang, J. Xi, and J. K. Hedrick. A learning-based personalized driver model using bounded generalized gaussian mixture models. *IEEE Transactions on Vehicular Technology*, 68(12):11679–11690, 2019.

[183] Z. Wang, T. Kaizuka, and K. Nakano. Effect of haptic guidance steering on lane following performance by taking account of driver reliance on the assistance system. *IEEE Systems, Man, Cybernetics*, In Press, Online Available, 2018.

[184] Z. Wang, Z. Yan, and K. Nakano. Comfort-oriented haptic guidance steering via deep reinforcement learning for individualized lane keeping assist. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pages 4283–4289, 2019.

[185] D. Weir and D. McRuer. Models of steering control of motor vehicles. In *Proc. 4th Annual Manual, NASA*, pages SP–192, 1968.

[186] D. H. Weir and D. T. McRuer. Dynamics of Driver Vehicle Steering Control. *Automatica*, 6(1):87–98, Jan. 1970.

[187] J. P. M. A. Wolters, K. Van Der El, H. J. Damveld, D. M. Pool, M. M. van Paassen, and M. Mulder. Effects of simulator motion on driver steering performance with various visual degradations. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 781–786, 2018.

[188] Y. Xing, C. Lv, and D. Cao. Personalized vehicle trajectory prediction based on joint time-series modeling for connected vehicles. *IEEE Transactions on Vehicular Technology*, 69(2):1341–1352, 2020.

[189] L. Xu, S. Hu, and Q. Luo. A new lane departure warning algorithm considering the driver's behavior characteristic. *Mathematical Problems in Engineering*, 2015.

[190] Z. Xu, K. Zhang, H. Min, Z. Wang, X. Zhao, and P. Liu. What drives people to accept automated vehicles? findings from a field experiment. *Transportation Research Part C: Emerging Technologies*, 95:320 – 334, 2018.

[191] D. Yi, J. Su, C. Liu, and W. Chen. Trajectory clustering aided personalized driver intention prediction for intelligent vehicles. *IEEE Transactions on Industrial Informatics*, 15(6):3693–3702, 2019.

[192] R. Y. Yuan, Z. Wei, L. Sheng, and Z. Lian. Study on vehicle track model in road curved section based on vehicle dynamic characteristics. *Mathematical Problems in Engineering*, 2012:1–17, 2012.

[193] N. M. Yusof, J. Karjanto, J. Terken, F. Delbressine, M. Hassan, and M. Rauterberg. The exploration of autonomous vehicle driving styles: Preferred longitudinal, lateral, and vertical accelerations. *in Proc. 8th Int. Conf. Automot. User Interfaces Interact. Veh. Appl.*, pages 245–252, 2016.

[194] S. Zafeiropoulos and P. Tsiotras. Design of a lane-tracking driver steering assist system and its interaction with a two-point visual driver model. *American Control Conference*, 2014.

[195] J. Zhang, J. Si1, Z. Yin1, Z. Gao, J. Moon, J. Gong, and F. Tang. Lane departure warning algorithm based on probability statistics of driving habits. *Soft Computing*, page Online Preprint, 2020.

[196] H. Zwaan, S. Petermeijer, and D. Abbink. Haptic shared steering control with an adaptive level of authority based on time-to-line crossing. *IFAC PapersOnLine*, 52(19):49–54, 2019.

A

CHAPTER 2 APPENDIX

A.1. VEHICLE DYNAMICS

The vehicle dynamics used in Chapter 2 is a linear approximation of the vehicle dynamics in the the simulator of [31], which is a modified bicycle model with non-linear tire dynamics. The linearly identified discrete state-space vehicle dynamics used in Chapter 2 is given in Eqs. A.1 and A.2. This is the elaboration of the VehicleDyn() function in Algorithm 1 presented in Chapter 2.

$$\begin{bmatrix} x_1(k+1) \\ x_2(k+1) \end{bmatrix} = \begin{bmatrix} 0.8665 & -0.0385 \\ 0.1313 & 0.9824 \end{bmatrix} \begin{bmatrix} x_1(k) \\ x_2(k) \end{bmatrix} + \begin{bmatrix} -14.2392 \\ 23.6425 \end{bmatrix} \delta_s(k) \tag{A.1}$$

$$\begin{bmatrix} v_{lat}(k) \\ \dot{\psi}_g(k) \end{bmatrix} = \begin{bmatrix} -0.2920 & -0.1938 \\ -0.0006 & 0.0283 \end{bmatrix} \begin{bmatrix} x_1(k) \\ x_2(k) \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix} \delta_s(k) \tag{A.2}$$

Note that the state vector output in Fig. 2.7 $x_{vd}(t) = [x_1(k+1)\ x_2(k+1)]^T$, which is used as input to the discrete VehicleDyn() function.

A.2. BODY TO GLOBAL TRANSFORMATION

The body to global transformation takes the body lateral velocity v_{lat} that comes from the vehicle dynamics and the constant longitudinal velocity v_{long} and transforms these to v_x and v_y components, given an angle ψ_g in the global reference frame. The angle ψ_g is taken from the cartesian positive axis, counter-clockwise as positive. Thereby, the angle needed to rotate the vehicle back to the positive x axis is $-\psi_g$. the correct rotation matrix calculation is thereby given in Eq. A.3. Afterwhich, the positions are obtained by integration.

$$\begin{bmatrix} v_y \\ v_x \end{bmatrix} = \begin{bmatrix} \cos(-\psi_g) & -\sin(-\psi_g) \\ \sin(-\psi_g) & \cos(-\psi_g) \end{bmatrix} \begin{bmatrix} v_{lat} \\ v_{long} \end{bmatrix} \tag{A.3}$$

$$X(t) = \int_0^t v_x(t)dt \tag{A.4}$$

$$Y(t) = \int_0^t v_y(t)dt \tag{A.5}$$

B

CHAPTER 3 APPENDIX

B.1. VEHICLE DYNAMICS AND NEUROMUSCULAR CONSTANTS

This appendix gives the equations used for the vehicle dynamics coefficients and the respective parameter values for the vehicle dynamics constants and neuromuscular dynamics time constant. The coefficients of the bicycle dynamics used, are given below:

$$a_{11} = -\frac{2(c_r + c_f)}{mV} \tag{B.1}$$

$$a_{12} = \frac{2(c_r l_r + c_f l_f)}{mV^2} - 1 \tag{B.2}$$

$$a_{21} = \frac{2(c_r l_r - c_f l_f)}{J} \tag{B.3}$$

$$a_{22} = \frac{2(c_r l_r^2 - c_f l_f^2)}{JV} \tag{B.4}$$

$$a_{15} = \frac{2c_f}{mVR_s} \tag{B.5}$$

$$a_{25} = \frac{2c_f l_f}{JR_s} \tag{B.6}$$

In Eq. (B.1)-Eq. (B.6) c_r and c_f are the cornering stiffness of the front and rear tires, l_f and l_r are the distance of the front and rear axles from the center of gravity. J is the moment of inertia about the yaw axis, m is the vehicle mass, V is the longitudinal velocity and R_s the steering ratio.

The values for these parameters as used in Chapter 3, are given in Table B.1.

B.2. MARS MODEL GEOMETRIC DERIVATION

This appendix explains the geometric derivations for the near and far angles in the Mars Model. The linear approximations for the near angle θ_{near} and the far angle θ_{far} are derived using geometric approximations. The approximations are based on different assumptions. These assumptions become clear when understanding the geometric relations that resulted in Eq. (3.4)

Table B.1: Constants used in both Vehicle/road dynamics and neuromuscular dynamics

Parameter	interpretation	Value
l_f	distance from gravity center to front axle	1.127 m
l_r	distance from gravity center to rear axle	1.485m
m	vehicle mass	1476 Kg
J	moment of inertia around yaw axis	1810 Kgm^2
c_f	Cornering stiffness of front tire	65000 N/rad
c_r	Cornering stiffness of rear tire	5700 N/rad
R_s	steering ratio	16
l_s	look ahead distance	5
T_N	Neuromuscular time constant	0.1

and Eq. (3.3), shown in Fig. B.1. For the near angle, θ_{near} the 'lookahead' distance l_s is constant, with a value of 5 m. Within those 5m ahead of the car, it is assumed that the curvature is zero. Based on this, θ_{near} is computed with a small angle approximation, and is the summation of the angle from the center of gravity of the car to the waypoint ahead $\frac{s_{lat}}{l_s}$ and the heading error ψ_L, as given in Eq. (B.7).

$$\theta_{near} = \tan(\frac{s_{lat}}{l_s}) + \psi_L \approx \frac{s_{lat}}{l_s} + \psi_L \tag{B.7}$$

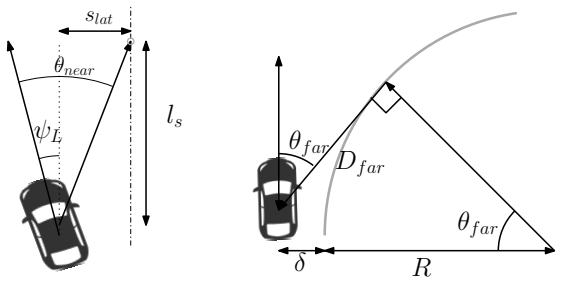

a) Near angle θ_{near} b) Far angle θ_{far}

Figure B.1: This figure illustrates a geometric interpretation of the linear near and far angle approximation in the Mars Model. This figure illustrates the assumptions made to realise the Model

For the far angle θ_{far}, depicted in Fig. B.1, the equation exhibits a small angle approximation, in other words, the estimation of the far angle is more accurate for roads with small curvature (large R), as shown in Eq. (B.8). Moreover, this geometric approximation can only be valid for constant radii curves. This assumption is taken through to the computation of the distance to tangent point D_{far}, another constant, shown in Eq. (B.9). D_{far} is essentially a function of curvature κ (or radius R) and distance from the road boundary δ. Since D_{far} is constant and R is also constant, δ has to be assumed constant as well. In this Chapter a centerline position is assumed as a value for δ, this is equal to half the road width $\frac{W_r}{2}$.

$$\begin{aligned} \theta_{far} &= tan(\frac{D_{far}}{R}) \\ &\approx D_{far}\kappa \end{aligned} \tag{B.8}$$

$$(\delta + R)^2 = R^2 + D_{far}^2$$

$$D_{far} = \sqrt{\delta^2 + 2\delta R} \tag{B.9}$$

B.3. VAN DER EL MODEL TRANSFORMATION

This appendix explains and derives the needed transformations for the Van der El model to apply it to a practical driving simulation.

The input signals of the original model, given in Fig. B.2, are not compatible with a practical driving situation. The signal y_c represents the lateral distance of the centreline of the road in global reference, i.e. from the beginning of the simulation. Likewise, the signal y is the lateral position of the vehicle in the global reference frame as well. Moreover, the absolute heading of the car ψ is also used as input; however, would produce unrealistic feedback inputs when driving on a roundabout. The lateral position is corrected by subtracting y from y_c when driving north from the start of the simulation; any deviation from centerline would produce correct lateral deviation values. However, when driving towards the east and making lateral deviations from the centerline, using the Van der El model's original inputs, the recorded deviation would be almost 0.

To make this formulation work practically, the road would have to make relatively small lateral deviations from the start point, which is achievable when considering a road that is a sum of sinusoids (as is done in [173]), rather than a realistic road. Therefore, there is a need to transform these theoretical inputs, to that of realistic driver input signals. Through algebraic manipulation we have changed the input signals to feedforward on the road heading ψ_c, feedback on the lateral position relative to centerline y_L and current heading relative to the road heading ψ_L. The algebraic derivation that leads to this manipulation is shown in Eq. (B.10), Eq. (B.11), Eq. (B.12), Eq. (B.13) and Eq. (B.14).

$$\begin{aligned}
\psi_L &= \psi_c - \psi \\
s_{lat} &= y_c - y \\
y_c &= \psi_c \frac{V}{s}
\end{aligned}$$

$$\begin{aligned}
\delta_s &= H_{y_c} y_c - H_y y - H_\psi \psi \\
&= H_{y_c} \psi_c \frac{V}{s} - H_y (y_c - s_{lat}) - H_\psi (\psi_c - \psi_L) \\
&= H_{y_c} \psi_c \frac{V}{s} - H_y (\psi_c \frac{V}{s} - s_{lat}) - H_\psi (\psi_c - \psi_L)
\end{aligned}$$

$$= (H_{y_c} \frac{V}{s} - H_y \frac{V}{s} - H_\psi)\psi_c + H_y s_{lat} + H_\psi \psi_L \tag{B.10}$$

From Eq. (B.10) and Fig. B.2, we can see that to formulate the adapted inputs, the transfer functions H_{ψ_c}, $H_{s_{lat}}$ and H_{ψ_L} need to be defined. To make such an algebraic inference, the definitions of the original functions H_ψ, H_y and H_{y_c} from [173] are the following:

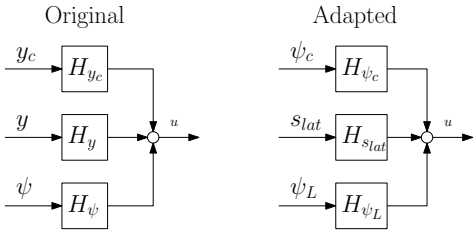

Figure B.2: The algebraic conversion from the original Van der El model which has inputs centerline position y_c (relative to starting point in the global reference frame), car position y and heading ψ. These inputs and transfer functions were deduced from an abstract experiment, therefore, to make the model more practically applicable, the transfer functions are algebraically restructured to facilitate the inputs: the road heading ψ_c, lateral position relative to centerline s_{lat} and heading error ψ_L.

$$H_\psi = K_e(1+T_L s)e^{-\tau_v s}$$
$$H_y = K_y H_\psi \quad where \quad K_y = \frac{1}{V\tau}$$

$$H_{y_c} = K_f \frac{1}{1+T_{l,f}s} e^{\tau_f s} H_y H_\psi \tag{B.11}$$

Using Eq. (B.10) and the algebraic expressions in Eq. (B.11), we can substitute the following to compute the transfer function H_{ψ_c} accompanying the input ψ_c. The substitutions result in Eq. (B.12).

$$\begin{aligned}
H_{\psi_c} &= H_{y_c}\frac{V}{s} - H_y\frac{V}{s} - H_\psi \\
&= K_f \frac{1}{1+T_{l,f}s} e^{\tau_f s} K_y K_e(1+T_L s)e^{-\tau_v s}\frac{V}{s}... \\
&\quad - K_y K_e(1+T_L s)e^{-\tau_v s}\frac{V}{s} - K_e(1+T_L s)e^{-\tau_v s} \\
&= (K_f \frac{1}{1+T_{l,f}s} e^{\tau_f s} - 1) K_y K_e(1+T_L s)e^{-\tau_v s}\frac{V}{s}... \\
&\quad - K_e(1+T_L s)e^{-\tau_v s} \\
&\quad with \quad K_y = \frac{1}{V\tau} \\
&= (K_f \frac{1}{1+T_{l,f}s} e^{\tau_f s} - 1)\frac{1}{\tau s} K_e(1+T_L s)e^{-\tau_v s}... \\
&\quad - K_e(1+T_L s)e^{-\tau_v s}
\end{aligned}$$

$$= ((K_f \frac{1}{1+T_{l,f}s} e^{\tau_f s} - 1)\frac{1}{\tau s} - 1) K_e(1+T_L s)e^{-\tau_v s} \tag{B.12}$$

Using Eq. (B.10) and the algebraic expressions in Eq. (B.11), we can substitute the following to compute the transfer function $H_{s_{lat}}$ accompanying the input s_{lat}. The substitutions result in Eq. (B.13).

$$\begin{aligned} H_{s_{lat}} &= -H_y \\ &= K_y K_e (1 + T_L s) e^{-\tau_v s} \\ \\ &= \frac{1}{\tau s} K_e (1 + T_L s) e^{-\tau_v s} \end{aligned} \tag{B.13}$$

Finally, using Eq. (B.10) and the algebraic expressions in Eq. (B.11), we can substitute the following to compute the transfer function H_{ψ_L} accompanying the input ψ_L. The substitutions result in Eq. (B.14).

$$\begin{aligned} H_{\psi_L} &= -H_\psi \\ \\ &= -K_e (1 + T_L s) e^{-\tau_v s} \end{aligned} \tag{B.14}$$

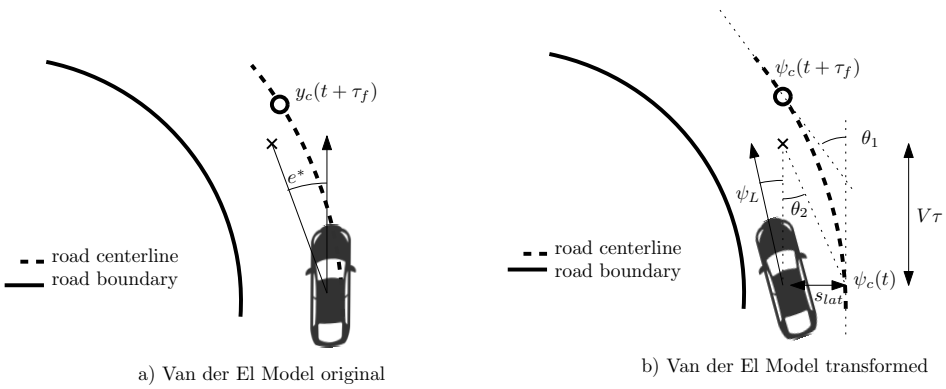

a) Van der El Model original

b) Van der El Model transformed

Figure B.3: The right figure illustrates the original Van der El model concept. The left figure illustrates the transformed Van der El model concept. Presented are the three angles summed to obtain e^* in the transformed Van der El model: θ_1, θ_2 and ψ_L. See Fig. 3.2 for the corresponding control diagram.

Fig. B.3 illustrates the concept of the three angles that make up $e*$ in the transformed Van der El model.

B.4. VAN PAASSEN CURVE CUTTING DISTANCE DERIVATION

This appendix derives the curve cutting distance y_{cc} using a geometry. The feedback reference signal y_{cc} is given in Eq. (B.15) as a function of curvature κ, velocity V and preview time τ_f. The geometric background for this formulation is given in Fig. B.4. The derivation is the following:

B

$$\sin(\theta) = \frac{\tau_f V}{R}$$
$$\cos(\theta) = \frac{R - y_{cc}}{R}$$

From these equations it can be seen that the radial angle θ is a function of preview time τ_f. Using the 2nd order Taylor expansions:

$$\sin(\theta) = \theta$$
$$\cos(\theta) = 1 - \frac{\theta^2}{2!}$$

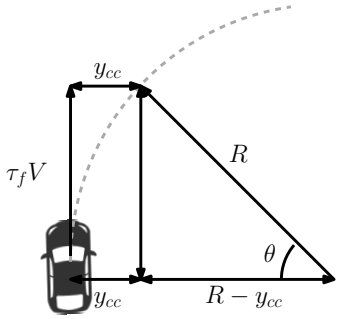

Figure B.4: The geometric interpretation the curve cutting distance y_{cc}, used as feedback reference in the Van Paassen model.

Substituting results in the following to get y_{cc}:

$$\theta = \frac{\tau_f V}{R}$$
$$1 - \frac{1}{2!}(\frac{\tau_f V}{R})^2 = \frac{R - y_{cc}}{R}$$
$$R - \frac{1}{2}(\tau_f V)^2 \frac{1}{R} = R - y_{cc}$$
$$y_{cc} = \frac{1}{2}(\tau_f V)^2 \frac{1}{R}$$

$$y_{cc} = \frac{1}{2}(\tau_f V)^2 \kappa \tag{B.15}$$

From Eq. (B.15) we see that the curve cutting distance y_{cc} is a function of the *road curvature* κ. In this equation, we assume that the vehicles trajectory radius (dotted line in Fig. B.4) equals the radius of the road. For most cases this is a good approximation, nevertheless can be highly violated when taking race-line trajectories.

B.5. REASONS FOR REFINEMENT OF ASSESSMENT METHODOLOGY

In Chapter 2, there were three criteria defined:

- *Descriptiveness*: the model's ability to capture different types of steering behaviour.
- *Identifiability*: the ability of the model for unique mapping between a steering behaviour and a parameter combination.
- *Realism*: the parameter combination space resulting in realistic steering behaviour, where realistic behaviour is defined as a model output whose trajectory stays within road boundaries and does not exhibit oscillatory steering angles.

Descriptiveness was realised as a metric through the area which the model's realistic trajectories can span along a curve. Identifiability was realised through the comparison of a model output of a given parameter combination with all the resulting model outputs from a selected parameter span P_{span}. Realism was evaluated as the realistic space of parameters within the selected parameter span P_{span}.

There is however some weaknesses identified with the way these metrics are realised: 1) for defining the descriptiveness of a model, the area which the realistic trajectories can span along a curve completely obscures information about the different patterns of trajectories that the model may need to capture. 2) The *selected* parameter span P_{span} may or may not reveal the complete identifiability or realistic space, as was shown in the results of Chapter 2. 3) The conventional identifiability space and realism space is a volume of N dimensional space (where N is the number of parameters considered as a Degree of Freedom (DoF)). This volume has units that is a multiplication of all the DoF parameter units. This is problematic when comparing models that exhibit different number of DoF parameters and different parameter units. That is, there is no way to establish a valid comparison, as not only would the units of the space be different, the dimension of the identifiability and realism volume would also be different.

To solve these weaknesses, the methodology presented is refined. This update tackles 1) the way in-which descriptiveness is computed 2) the determination of P_{span} 3) the combination of identifiability and realism to formulate realistic identifiability and, 4) the evaluation of realistic identifiability being limited to the trajectory types that the model can actually describe.

The descriptiveness of the model is now evaluated through the extent to which a model can capture different driver trajectory categories, with the trajectory classes defined in Chapter 4. Hereby the methodology clearly outlines the different types of trajectories available in the driving population and reveals the types of generic behaviour the model can capture, thereby making the methodology data-driven. The P_{span} is explored, instead of selected, thereby making sure that all realistic parameter combinations are tested. To facilitate a fair comparison between models of different number and/or type of parameters, the resulting identifiability space is made dimensionless through normalisation by the realism space, to obtain a new metric 'Realistic Identifiability'. Moreover, the Realistic Identifiability is only evaluated for the parameter combinations of real trajectory types (classes) that the model can sufficiently capture.

B.6. IDENTIFIED PARAMETERS

This appendix gives the identified parameters per class for each model. The Mars model in Table. B.2, the Van der El model in Table B.3 and, the Van Paassen model in Table B.4.

B

Table B.2: Mars Model identified parameters per class, the classes which the model can sufficiently capture are given in bold.

Class	K_p (-)	K_c (-)
1	2.7	24
2	1.8	6
3	1.2	30
7	3	21
8	1.8	23
9	3	21
10	2.7	30

Table B.3: Van der El Model identified parameters per class, the classes which the model can sufficiently capture are given in bold.

Class	T_{lf} (s)	τ_f (s)	K_e (-)	T_L (s)
1	0.032	1.2	3.4	1.1
2	0.25	0.32	7.1	0.083
3	0.09	0.30	4.93	0.35
7	0.056	1.3	5.3	0.01
8	0.095	1.07	4.5	0.78
9	0.07	1.38	5.7	0.01
10	0.05	1.64	2.74	0.05

Table B.4: Van Paassen Model identified parameters per class, the classes which the model can sufficiently capture are given in bold.

Class	K_{FF} (-)	K_{FB} (-)	τ_f (s)	T_{hs} (s)	τ_n (s)
1	0.005	0.12	1.6	0.45	1.3
2	0	1.3	0.005	0.46	0.75
3	0.0014	1.8	1.1	0.28	0.17
7	0.0040	0.25	1.8	0.78	1.2
8	0.0001	0.12	1.3	0.34	1.2
9	0.44	1.1	1.6	0.62	0.87
10	0.29	0.12	0.97	0.24	0.28

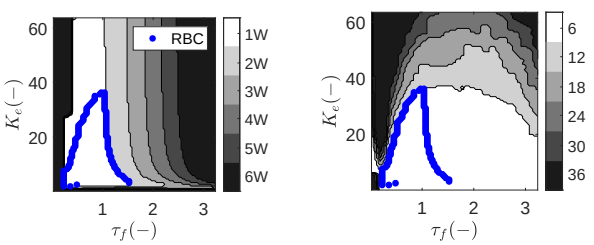

(a) within W_r? (b) Steering Reversals

Figure B.5: A 2-dimensional representation of Realistic Parameter space of the Van der El model, showing the error gain K_e and the look-ahead time t_f.(a) presents the first RBC constraint, whether s_{lat} is within the road boundaries W_r, (b) presents the second RBC constraint, the number of steering reversals.

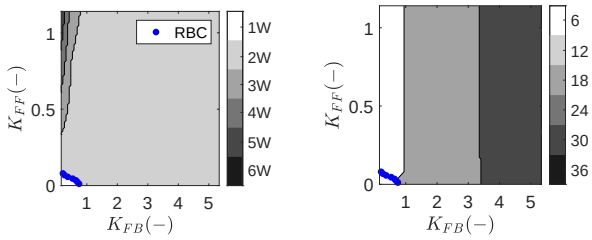

(a) within W_r? (b) Steering Reversals

Figure B.6: A 2-dimensional representation of Realistic Parameter space of the Van Paassen model, showing the feedforward gain K_{FF} and the feedback gain K_{FB}. (a) presents the first RBC constraint, whether s_{lat} is within the road boundaries W_r, (b) presents the second RBC constraint, the number of steering reversals.

Parameter Combinations	**Realistic Identifiability (%)**			
	Class 7		**Class 8**	
	δ_s	s_{lat}	δ_s	s_{lat}
$K_e\ \tau_f$	100	16	100	0.73
$K_e\ T_L$	100	12.7	100	1.3
$K_e\ T_{l,f}$	100	54	100	1.4
$\tau_f\ T_L$	98	17.1	98	6
$\tau_f\ T_{l,f}$	100	21	100	2.1
$T_L\ T_{l,f}$	100	53	100	16

Table B.5: The equivalent two-parameter set, Parameter Sensitivity values for Class 7 and 8, for the Van der El Model. They are given for the respective classes that the model can sufficiently replicate, in both the steering wheel δ_s and lateral position s_{lat} domains.

B.7. REALISM PLOTS

B.8. 2-DIMENSIONAL VAN DER EL AND VAN PAASSEN PARAMETER REALISTIC IDENTIFIABILITY VALUES

This appendix describes the 2-dimensional realistic identifiability spaces for Van der El and Van Paassen models. This is needed to facilitate comparison with the Mars model in the same dimension. These equivalent 2-dimensional realistic identifiability areas are computed with other parameters fixed; the fixed parameters take the values of class 7 and 8. The two-dimensional values for the Van der El model are given in Table B.5 and those of the Van Paassen model are given in Table B.6.

B.9. 4-DIMENSIONAL VAN PAASSEN MODEL REALISTIC IDENTIFIABILITY VALUES

This appendix gives the 4-dimensional Van Paassen Realistic Identifiability values for comparison with the Van der El model on the same dimension. One parameter is kept fixed at a time, outlined on the description row in Table B.7.

B

Parameter Combinations	Realistic Identifiability (%)			
	Class 7		Class 8	
	δ_s	s_{lat}	δ_s	s_{lat}
K_{FF} K_{FB}	100	50	100	50
K_{FF} τ_{LH}	57	28	69	23
K_{FF} τ_p	72	6.1	63	5.6
K_{FF} T_{hs}	98	0.5	99	0.4
K_{FB} τ_{LH}	75	12	66	11
K_{FB} τ_p	95	4.4	92	1.2
K_{FB} T_{hs}	100	0.5	92	0.5
τ_{LH} τ_p	81	2.3	57	3.3
τ_{LH} T_{hs}	96	1.3	97	2
τ_p T_{hs}	89	0.7	70	0.6

Table B.6: The equivalent two-parameter set, Parameter Sensitivity values for Class 7 and 8, for the Van Paassen Model. They are given for the respective classes that the model can sufficiently replicate, in both the steering wheel δ_s and lateral position s_{lat} domains.

Table B.7: 4-dimensional equivalent realistic identifiability values of Van Paassen model

Class	Fixed Parameter									
	K_{FF}		K_{FB}		T_{hs}		τ_f		τ_n	
	δ_s	s_{lat}	δ_s	s_{lat}	δ_s	s_{lat}	δ_s	s_{lat}	δ_s	s_{lat}
1	64	0.01	12	0.05	11	0.4	3	0.1	0.2	0.04
2	84	0.2	100	0.08	88	6	86	0.3	90	1.1
3	97	13	99	50	88	0.1	93	7	96	11
7	13	0.02	46	0.04	12	0.2	36	0.2	9	0.1
8	89	0.02	40	0.02	56	0.2	95	0.1	89	0.02
9	90	0.01	99	0.01	60	0.3	83	0.1	92	0.03
10	5	0.01	7	0.01	15	0.04	4	0.03	4	0.006

C

CHAPTER 6 APPENDIX

C.1. PREPOSITIONING PATH MODEL RESPONSES

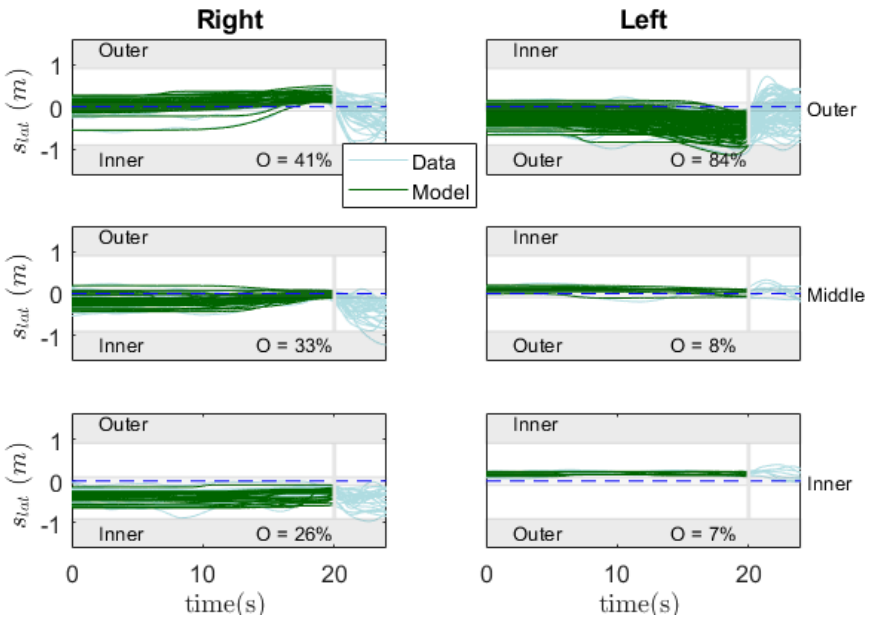

Figure C.1: This figure illustrates all data the from Chapter 5 in light blue. The corresponding model fits of all trajectories are given in dark green. The figure is sectioned by curve direction and prepositioning category: Outer, Middle and Inner. The percentage of drivers falling into a given prepositioning category is indicated by the Occurence (O).

C.2. IDENTIFIED PARAMETERS

Table C.1: The identified model parameters for the different classes, with and without prepositioning

Class	K_{FF} (-)		K_{FB} (-)		τ_f (s)		T_{hs} (s)		τ_n (s)		g_1 (m)	τ_1 (s)	a_1 (m/s)	a_2 (m/s)	y_b (m)
	NP	P	NP	P	NP	P	NP	P	NP	P					
R1	0.005	0.16	0.091	0.069	1.56	0.93	0.36	0.13	1.33	0.98	0.51	1.72	0.51	2.00	0.28
R2	0.006	0.005	1.28	1.85	0.005	0.008	0.46	1.34	0.75	0.80	0.67	0.78	1.03	1.96	0.12
R3	0.001	0.002	1.79	0.07	1.08	0.52	0.29	0.02	0.17	0.14	0.45	4.65	0.54	1.99	0.53
R7	0.004	0.012	0.25	0.08	1.81	1.62	0.78	0.65	1.15	1.53	-0.018	3.56	0.72	2.00	0.03
R8	0.008	0.003	0.12	0.04	1.27	0.45	0.35	0.02	1.17	1.84	0.29	8.01	0.93	1.75	-0.19
R9	0.44	0.21	0.11	0.27	1.59	2.42	0.63	1.41	0.87	1.07	0.05	8.52	0.08	1.96	-0.18
R10	0.29	0.31	0.12	0.10	0.97	0.69	0.24	0.17	0.28	0.01	-0.04	9.20	2.57	2.00	-0.27
L1	0.003	0.14	0.06	0.14	0.32	1.12	0.01	0.18	1.55	0.72	-0.80	4.01	0.53	1.99	-0.75
L2	0.004	0.003	0.38	0.25	0.002	0.005	0.16	0.01	0.71	0.85	-0.30	3.26	0.61	2.0	-0.31
L3	0.002	0.006	0.32	0.12	0.005	0.005	0.05	0.01	0.19	0.92	-0.52	2.82	0.71	2.0	-0.43
L4	0.07	0.25	0.44	0.06	1.14	0.53	0.18	0.20	0.29	0.29	-0.05	7.02	5.85	1.12	0.03
L5	0.08	0.14	0.008	0.006	0.21	0.47	0.02	0.15	3.51	3.1	-0.07	0.00	1.05	1.12	-0.01
L7	0.005	0.005	1.20	1.16	1.95	2.17	0.95	1.2	1.04	1.21	0.02	7.12	7.90	0.57	0.02
L8	0.001	0.001	0.05	0.06	1.25	1.12	0.58	0.53	1.57	1.56	-0.12	1.21	0.86	2.0	0.15
L9	1.52	2.18	2.23	1.85	2.19	1.68	1.75	1.86	0.63	0.27	0.15	0.37	3.41	0.99	0.24
L10	0.16	0.22	0.35	0.06	0.98	0.57	0.16	0.21	0.43	0.62	0.18	0.72	1.68	1.51	0.09
L11	0.005	0.20	0.02	0.04	0.05	0.40	0.02	0.22	2.72	1.1	0.16	2.74	1.47	2.00	0.10

D

CHAPTER 8 APPENDICES

D.1. OBJECTIVE MEASURES TIME SERIES

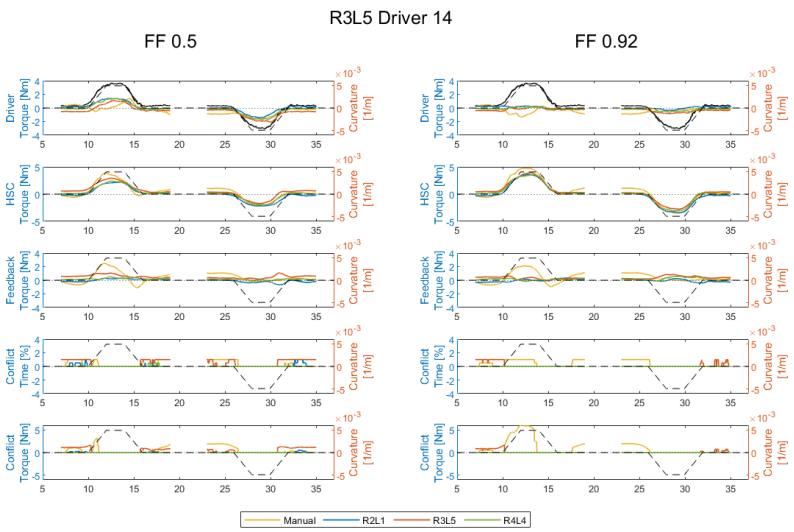

Figure D.1: This figure illustrates the averaged timeseries (across 5 repetitions) of a R3L5 driver subjected to the eight different experiment conditions.

D.2. COMPLETE RESULTS OF DRIVER TORQUE AND CONFLICT TIME

D.2.1. RIGHT CURVES

The **driver torque** T_D, illustrated in Fig. D.3a) & b), is significantly affected by HCR type in the *curve* phase ($F(2.1,61.7) = 81.9$, $p<0.01$). There is no significant interaction between driver

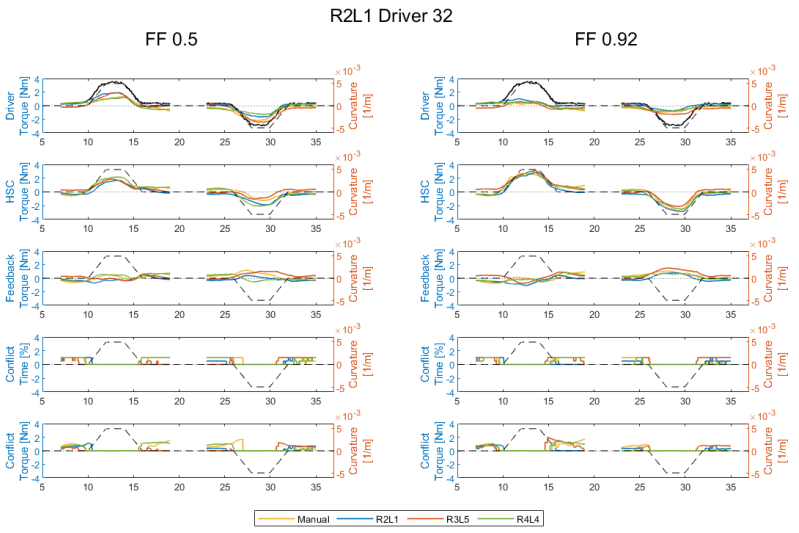

Figure D.2: This figure illustrates the averaged timeseries (across 5 repetitions) of a curve cutter (R2L1) subjected to the eight different experiment conditions.

group and HCR. Pairwise comparisons show that this effect stems from a significant decrease between CL and OD with an average of 0.46 Nm ($p<0.01$), a significant decrease between CL and CC with an average of 0.29 Nm ($p<0.01$), a significant decrease between CL and P with an average of 0.33Nm ($p<0.01$), a significant increase between OD and CC with an average of 0.16 Nm and a significant increase between OD and P with an average increase of 0.12 Nm. A significant effect is not found between CC and P. The highest median driver torque is found for CL and the lowest for both OD and P. For offset-drivers there is a significant difference between class-average and P (opposing H.I), whereas for curve-cutters there is no significant difference between class average and P (supporting H.I). The highest median driver torque (thereby highest driver effort) is found for CL compared to the other HCRs (supporting H.II).

Driver torque is significantly affected by HCR in the *prepositioning* phase ($F(1.9,58.9) = 10.9$, $p<0.01$). A significant interaction between HCR and driver group is found in this phase ($F(1.9,48.9) = 12.4$, $p<0.01$). For curve-cutters, pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 0.13 Nm ($p<0.05$) and a significant decrease between OD and P with an average of 0.23 Nm ($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and CC with and average of 0.32 Nm ($p<0.01$), a significant increase between OD and CC with an average of 0.28 Nm ($p<0.05$) and a significant decrease between CC and P with an average of 0.3 Nm ($p<0.01$). For curve cutters there is no significant difference between class-average and P (supporting H.I), for offset-drivers there is also no significant difference between class-average and P (supporting H.I). For both driver group CL does not result in the highest driver torque (opposing H.II).

Driver torque significantly decreases with increasing K_{LoHS} (FF) gain for the curve phase only ($F(1,30) = 1122.4$, $p<0.01$). Here, an average decrease of 0.66 Nm is found for curve cutters and an average decrease of 0.68 Nm is found for offset-drivers (supporting H.III). A similar significant decrease is not present for the prepositioning phase, due to the contribution of feedforward torque in this phase being relatively small.

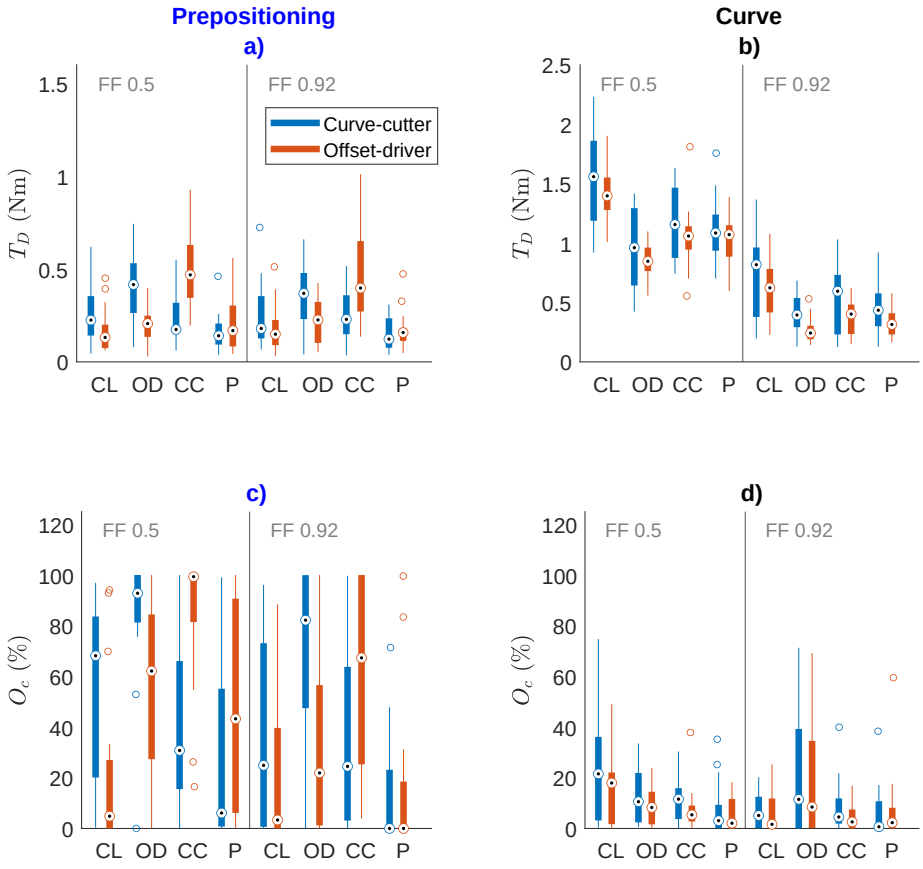

Figure D.3: For **right curves**, the data of the curve-cutters are given across a) to d). The data of the offset-drivers are given across e) to h). The driver torque T_D is given in a) & e), feedback torque T_{SoHF} is given in b) & f), conflict time O_c is given in c) & g) and conflict torque T_c is given in d) & h). These results are exclusively for *Left* curves, presented separately for each level of haptic support, FF = 0.5 and FF = 0.92 and, for the 'Prep' prepositioning phase and 'Curve' in-curve phase separately. The HCRs are Center-Line (CL), Offset-Driving (OD), Curve-Cutting (CC) and Personalised (P).

The **conflict time** O_c, illustrated in Fig. D.3 c) & d), is significantly affected by HCR in the *curve* ($F(1.6,48.5) = 4.29$, $p<0.01$). There is a non-significant interaction between driver group and HCR. Pairwise comparisons show that this stems from a significant decrease between OD and P with an average of 9.1% ($p<0.05$). For curve-cutters is supports H.I, whereas for offset drivers it opposes H.I. The CL does not result in largest conflict time (opposing H.II).

The conflict time is significantly affected by HCR during *prepositioning* ($F(1,30) = 9.2$, $p<0.01$). There is a significant interaction between HCR and driver group ($F(1,30) = 7.8$, $p<0.01$). For curve cutters, pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 30% ($p<0.01$) and a significant decrease between OD and P with an average of 57% ($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and CC with an average of 52% ($p<0.01$) and a significant decrease between CC and

P with an average of 42% ($p<0.01$). Both driver groups achieve the lowest conflict time with P with no significant difference to class-average (supporting H.I). Here, the curve-cutters achieve the highest conflict time with OD and the offset-drivers achieve the highest conflict time with an CC (opposing H.II).

Conflict time significantly decreases with increasing K_{LoHS} (FF) gain in the prepositioning phase only ($F(1.6,48.5) = 4.29$, $p<0.01$). Here, an average decrease of 12.8% is found for curve cutters and 20.1% for R3L5 drivers, when increasing the K_{LoHS} gain (supporting H.III).

LEFT CURVES

The **driver torque**, illustrated in Fig. D.4 a) & b), is significantly affected by HCR in the *curve* ($F(2.4,70.8) = 80.2$, $p<0.01$). There is a significant interaction between driver group and HCR ($F(2.3,70.2) = 12.9$, $p<0.01$). For curve cutters, pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 0.26 Nm ($p<0.01$), a significant decrease between CL and CC with an average of 0.14 Nm ($p<0.01$), a significant decrease between CL and P with an average of 0.11 Nm ($p<0.01$), a significant decrease between OD and CC with an average of 0.39 Nm ($p<0.01$) and a significant decrease between OD and P with an average of 0.37 Nm ($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and OD with an average of 0.29 Nm ($p<0.01$), a significant increase between CL and P with an average of 0.16 Nm ($p<0.01$), a significant decrease between OD and CC with an average of 0.32 Nm ($p<0.01$) and a significant increase between CC and P with an average of 0.2 Nm ($p<0.01$). For curve cutters the highest driver torque is found for OD and the lowest for both CC and P, whereas for R3L5 drivers the highest driver torque is also found for OD and the lowest for CL and CC. For both driver groups there is no significant effect between class-average and P (supporting H.I), nor does CL instigate the largest driver torque (opposing H.II).

Driver torque is significantly affected by HCR in the *prepositioning* phase ($F(1.4,41.4) = 5.3$, $p<0.05$). Contrary to right curves, for left curves there is no interaction effect between HCR and driver group in the prepositioning phase. Pairwise comparisons show that there is a significant decrease between CL and P with an average of 0.18 Nm ($p<0.01$), a significant decrease between OD and P with an average of 0.15 Nm ($p<0.01$) and a significant decrease between CC and P with an average of 0.16 Nm ($p<0.01$). Here, the lowest torque is found for P, with the other HCRs having no significant differences from eachother. This is expected because both an CC and OD preposition toward the outer part of the left curve. Both driver groups experience significant differences between class average and P (opposing H.I) and CL does not instigate the largest driver torque (opposing H.II).

Driver torque significantly decreases with increasing K_{LoHS} (FF) gain for the curve phase only ($F(1,30) = 707.8$, $p<0.01$). Here, an average decrease of 0.62Nm is found for curve cutters and an average decrease of 0.61Nm is found for offset-drivers, supporting H.III. A similar decrease in driver torque due to K_{LoHS} (FF) gain is not present for the prepositioning phase.

The **conflict time**, illustrated in Fig. D.4 c) & d), is significantly affected by HCR in the *curve* ($F(2.4,70.2) = 7.5$, $p<0.01$). There is no significant interaction between driver group and HCR. Pairwise comparisons show that there is a significant decrease between OD and P with on average 5.4 % ($p<0.01$) and a significant decrease between CC and P with on average 8.4 % ($p<0.01$). Conflict time is lowest with P with significant increase for only the class-average HCR for both driver groups (opposing H.I and H.II).

The conflict time is significantly affected by HCR in the *prepositioning* phase ($F(1.6,49.1) = 5.4$, $p<0.05$). There is no significant interaction between driver group and HCR. Pairwise comparisons show that there is a significant decrease between OD and CC with an average of 12 % ($p<0.01$), a significant decrease between OD and P with an average of 38% ($p<0.01$) and a significant decrease between CC and P with an average of 20% ($p<0.05$). Conflict time is lowest with

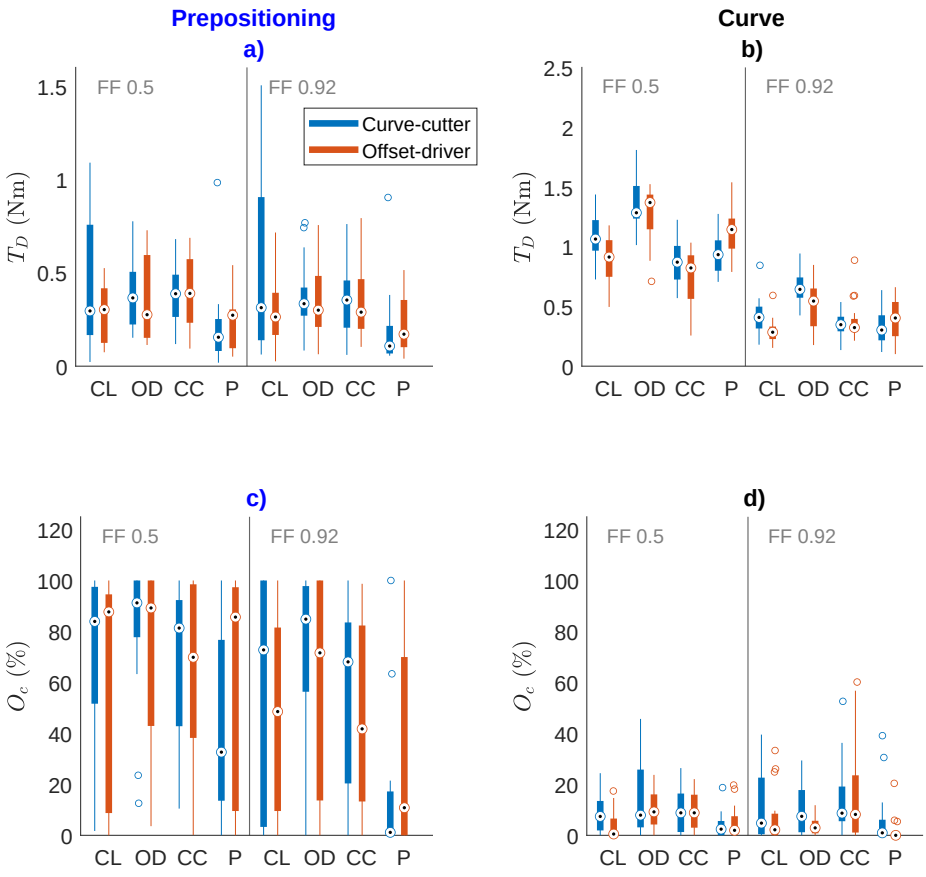

Figure D.4: For **left curves**, the data of the curve-cutters are given across a) to d). The data of the offset-drivers are given across e) to h). The driver torque T_D is given in a) & e), feedback torque T_{SoHF} is given in b) & f), conflict time O_c is given in c) & g) and conflict torque T_c is given in d) & h). These results are exclusively for *Left* curves, presented separately for each level of haptic support, FF = 0.5 and FF = 0.92 and, for the 'Prep' prepositioning phase and 'Curve' in-curve phase separately. The HCRs are Center-Line (CL), Offset-Driving (OD), Curve-Cutting (CC) and Personalised (P).

P with significant increase for only the class-average HCR for both driver groups (opposing H.I and H.II). Here, the highest conflict time is found for OD and the lowest for P.

D.3. COMPLETE RESULTS OF FEEDBACK TORQUE AND CONFLICT TORQUE

D.3.1. RIGHT CURVES

The **feedback torque** T_{SoHF} in the *prepositioning* phase, illustrated in Fig. D.5 a), is significantly affected by HCR in the *prepositioning* phase ($F(1.9,56.6) = 9.7$, $p<0.01$). There is a significant interaction between HCR and driver group ($F(1.8,56.7) = 13.7$, $p<0.01$). For curve-cutters, pairwise

comparisons show that this effect stems from a significant increase between CL and OD with an average of 0.17 Nm ($p<0.05$) and a significant decrease between OD and P with an average of 0.32 Nm($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and CC with an average of 0.4 Nm ($p<0.01$), a significant increase between OD and CC with an average of 0.34 Nm ($p<0.01$) and a significant decrease between CC and P with an average of 0.35 Nm ($p<0.01$). Hereby, the highest feedback torque is achieved with OD for curve-cutters and with CC for offset-drivers. For both driver groups there is no significant difference found between class-average and P (supporting H.I), nor does CL instigate the largest feedback torque (opposing H.II).

The **feedback torque** T_{SoHF} in the *curve*, illustrated in Fig. D.5 b), is significantly affected by HCR type in the ($F(1.6,46.6) = 13.4$, $p<0.01$). There is no significant interaction between driver group and HCR. Pairwise comparisons show that this stems from a significant decrease between CL and OD with and average of 0.21 Nm ($p<0.05$), a significant decrease between CL and CC with an average of 0.22 Nm ($p<0.01$) and a significant decrease between CL and P with an average of 0.27 Nm ($p<0.01$). The lowest feedback torque is found for P, however there is no significant difference between P, CC and OD (supporting H.I). The highest feedback torque is found for CL (supporting H.II). The feedback torque is also significantly affected by driver group in the curve ($F(1,30) = 7.5$, $p<0.01$), curve-cutters experience 0.17 Nm more feedback torque on average than offset-drivers.

Feedback torque significantly increases with increasing K_{LoHS} (FF) gain for the curve phase only ($F(1,30) = 4.2$, $p<0.05$). Here, an average increase of 0.04 Nm is found for curve cutters, whereas an average increase of 0.05 Nm is found for R3L5 drivers (opposing H.III).

The **conflict torque**, illustrated in Fig D.5 c), is significantly affected by HCR during the *prepositioning* phase ($F(1.6,50.9) = 4.99$, $p<0.01$). There is a significant interaction between HCR and driver group ($F(1.9,58.3) = 11.8$, $p<0.01$). For curve cutters, pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 0.4 Nm ($p<0.01$) and a significant decrease between OD and P with an average of 0.69 Nm ($p<0.01$). For offset-drivers, this effect stems from a significant increase between CL and CC with an average of 0.71 Nm ($p<0.01$) and a significant decrease between CC and P with an average of 0.67 Nm ($p<0.01$). For both driver groups there is no significant difference between class average and P (supporting H.I), nor does CL instigate the highest driver torque (opposing H.II).

The **conflict torque**, illustrated in Fig D.5 d), is significantly affected by HCR in the *curve* ($F(1.6,50.9) = 4.99$, $p<0.05$). There is a non-significant interaction between driver group and HCR. Pairwise comparisons show that this stems from a significant decrease between OD and CC with an average of 0.24 Nm ($p<0.05$), this supports H.I & H.II.

Conflict torque significantly decreases with increasing K_{LoHS} (FF) gain in the prepositioning phase only ($F(1.6,48.5) = 4.29$, $p<0.01$), for both driver groups. Here, an average decrease of 0.08Nm is found for curve cutters and 0.12Nm for R3L5 drivers, when increasing K_{LoHS} (supporting H.III).

D.3.2. LEFT CURVES

The **feedback torque** in the *prepositioning* phase, illustrated in Fig. D.6 a), is significantly affected by HCR ($F(1.4,41.4) = 6.4$, $p<0.01$). No significant interaction between driver group and HCR is found. Pairwise comparisons show that there is a significant increase between CL and CC with an average of 0.25 Nm ($p<0.01$), a significant increase between OD and CC with an average of 0.07 Nm ($p<0.01$), a significant decrease between OD and P with an average of 0.15 Nm ($p<0.01$) and a significant decrease between CC and P with an average of 0.22 Nm ($p<0.01$), opposing H.I. Here we see the lowest feedback torque is achieved for the P, whereas the highest is achieved for CC for both driver groups, opposing H.II.

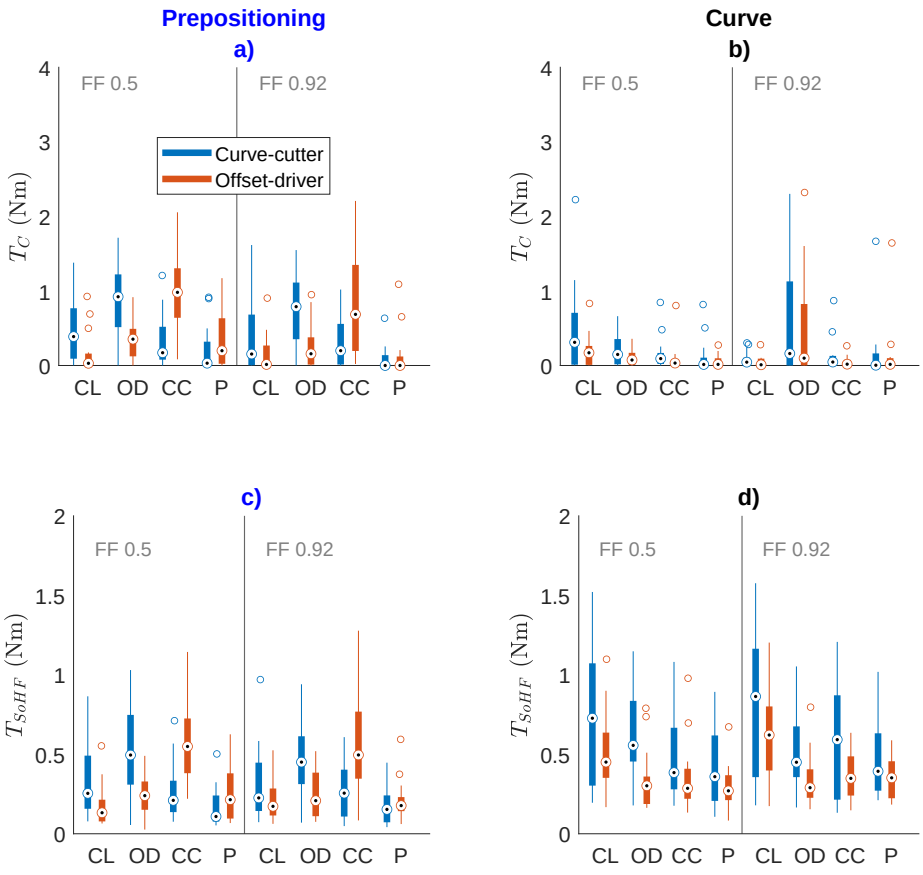

Figure D.5: For **right curves**, the feedback torque T_{SoHF} and conflict torque T_c

The **feedback torque** in the curve, illustrated in Fig. D.6 b), is significantly affected by HCR ($F(1.7,50.6) = 18.8$, $p<0.01$). There is a significant interaction between driver group and HCR ($F(1.6,50.6) = 6.8$, $p<0.01$). For curve-cutters pairwise comparisons show that this effect stems from a significant increase between CL and OD with an average of 0.18 Nm ($p<0.01$), a significant decrease between CL and P with an average of 0.15 Nm ($p<0.01$), a significant decrease between OD and P with an average of 0.32 Nm ($p<0.01$) and a significant decrease between CC and P with an average of 0.24 Nm ($p<0.05$). For offset-drivers, this effect stems from a significant increase between CL and CC with an average of 0.44 Nm ($p<0.01$), a significant increase between OD and CC with an average of 0.32 Nm ($p<0.05$) and a significant decrease between CC and P with an average of 0.43 Nm ($p<0.05$). For curve-cutters, a significant difference between class average and P is found (opposing H.II), whereas for offset-drivers such a significance is not found (supporting H.II). For the curve-cutters, the highest feedback torque is achieved for OD, whereas for offset drivers the highest feedback torque is achieved for CC, opposing H.II.

Feedback torque significantly increases with increasing K_{LoHS} gain for the curve phase only ($F(1,30) = 8.91$, $p<0.01$). Here, an average increase of 0.08Nm is found for curve cutters, whereas an average increase of 0.03Nm is found for R3L5 drivers, opposing H.III.

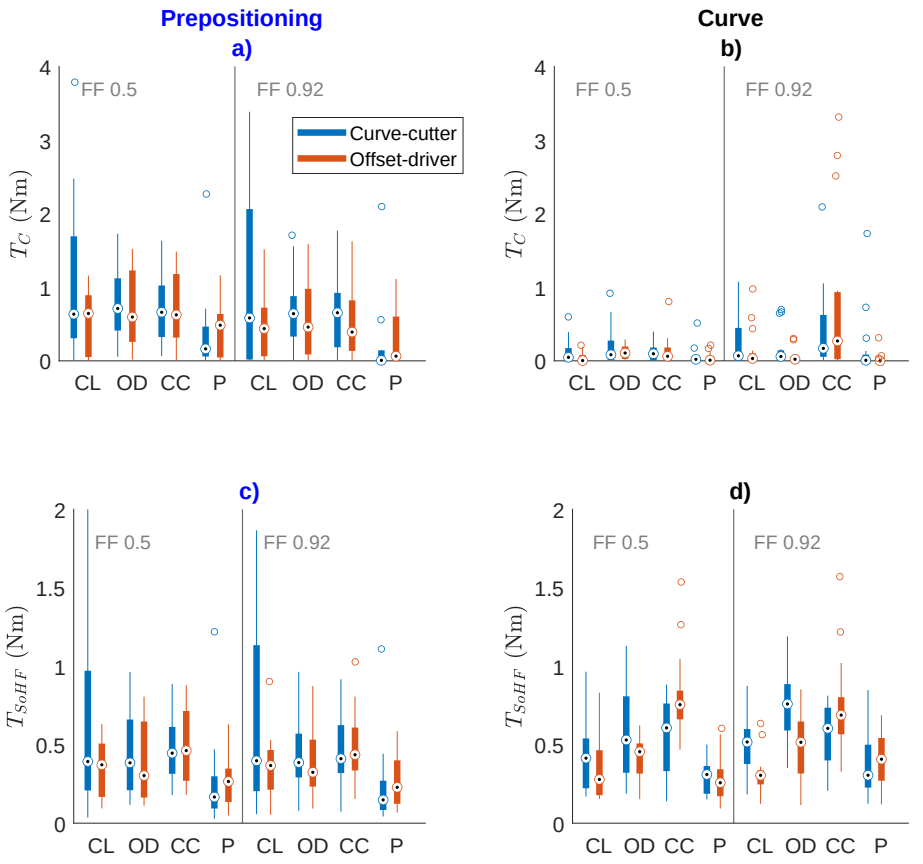

Figure D.6: For **left curves**, the feedback torque T_{SoHF} and conflict torque T_c

The **conflict torque** in the *prepositioning* phase, illustrated in Fig D.6 c), is significantly affected by HCR ($F(1.3,40.8) = 5.4$, $p<0.05$). There is no significant interaction between driver group and HCR. Pairwise comparisons show that there is a significant decrease between CL and P with an average of 0.46 Nm ($p<0.01$), a significant decrease between OD and P with an average of 0.36 Nm ($p<0.01$) and a significant decrease between CC and P with an average of 0.33 Nm ($p<0.01$). For both driver groups the class-average exhibits significantly more conflict torque than P (opposes H.I). The smallest conflict torque is found for P, whereas the other HCRs are not significantly different from one another (opposes H.II).

The **conflict torque** in the *curve*, illustrated in Fig D.6 d), is significantly affected by HCR ($F(1.4,41.4) = 7.4$, $p<0.01$). There is no significant interaction between driver group and HCR. Pairwise comparisons show that there is a significant increase between CL and CC with an average of 0.22 Nm and a significant decrease between CC and P with an average of 0.28 Nm. For both driver groups a significant peak in conflict exists for CC, which opposes H.I for curve cutters, supports H.I for offset-drivers and opposes H.II.

Conflict torque significantly decreases with increasing K_{LoHS} (FF) gain in the prepositioning phase ($F(1,30) = 10$, $p<0.01$), whereas conflict torque increases with increasing gain in the

curve phase (F(1,30) = 8.3, p<0.01). Here, an average decrease in conflict torque of 0.06Nm is found for curve cutters and 0.15Nm for offset-drivers, when increasing K_{LoHS} in the prepositioning phase, supporting H.III. Whereas, an average increase of 0.13Nm is found for curve cutters and 0.16Nm for R3L5 drivers, when increasing K_{LoHS} in the curve phase, opposing H.III.

D.4. SEPARATED L1 AND L2 GROUPS

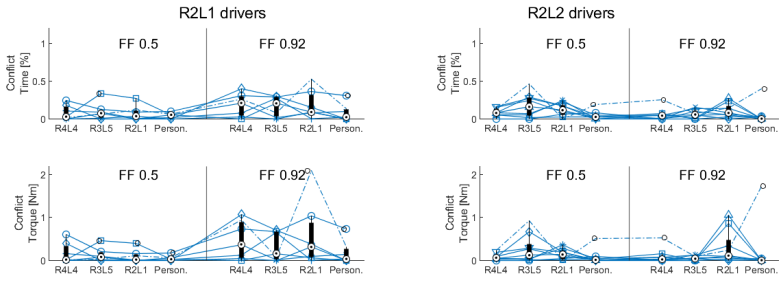

Figure D.7: Side-by-side analysis of conflict times and torques for R2L1 and R2L2 drivers in left curves.

D.5. CONSISTENCY IN STYLE

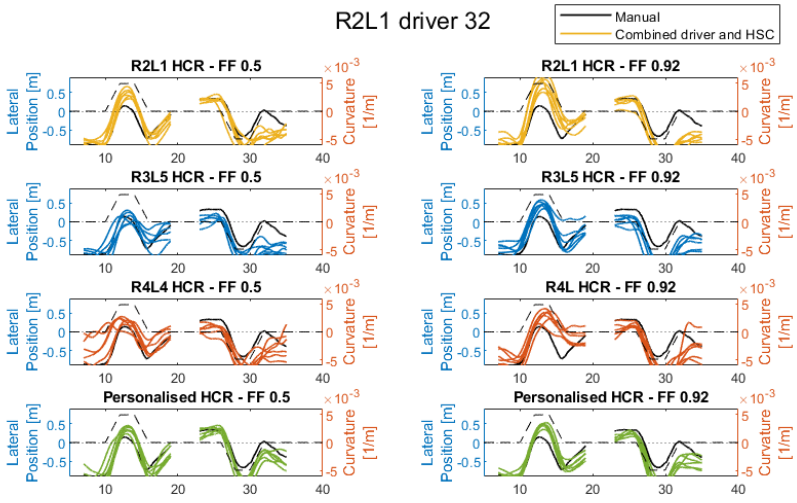

Figure D.8: Driven trajectories by an R2L1 driver subjected to the eight different experiment conditions.

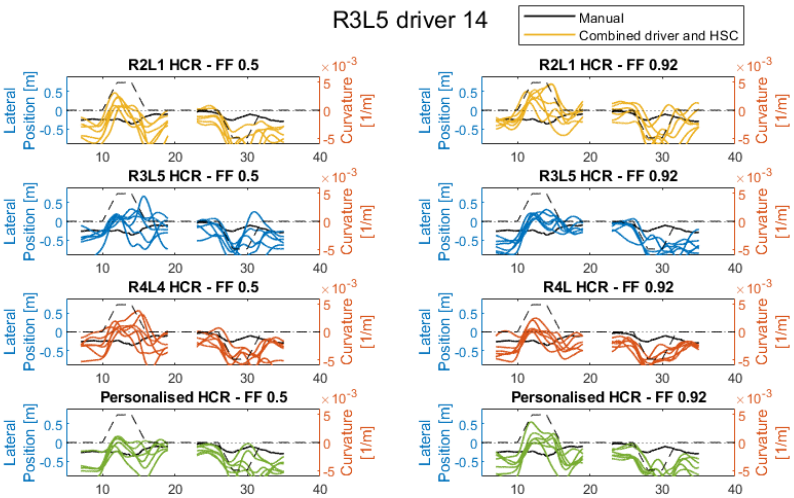

Figure D.9: Driven trajectories by an R2L1 driver subjected to the eight different experiment conditions.

D.5.1. TRAJECTORY RECLASSIFICIATION

This appendix presents the analysis of reclassification of the driven trajectories for both driver groups and per condition, to analyse the adaptations made. The results are summarised in Figures D.10, D.11, D.12 and D.13, each representing a different HCR. The R2L1 HCR is shown in Figure D.10. It presents the drivers trajectory on in the first session of the experiment, separately for the left and right curves and feedforward gains and the reclassification after the driver is subjected to a specific HCR. Figure D.10 shows how the curve cutters, both R2L1 and R2L2 drivers, are 84.38% consistent for right curves and 59.38% for left curves with R2L1 guidance. The R3L5 drivers on the other hand adapt to follow the guidance and end up in more 'efficient' classes, i.e., trajectories that result in a higher TLC. A significant portion, 43.75% for FF 0.5 and 50.00% for FF 0.92, of the R3L5 drivers reclassifies as R2L1 drivers when following this guidance.

The adaptations made when drivers are subjected to the R3L5 HCR is given in Figure D.10. Here it can be seen that the R3L5 drivers most often stick to their own class after reclassification, with 37.50% changing their trajectory to class 2 for left curves. Conversely, 25 % the curve cutters reclassify to a class 5 with a 0.5 feed forward gain.

The adaptations made when drivers are subjected to the R4L4 HCR is given in Fig. D.12. Here it can be seen that the R4L4 has a large effect on both groups of drivers. In right curves, 50 % of curve cutters adapt to a class 3 for the low feedforward gain. In left curves, 31 % of curve cutters adapt to class 3 for 0.92 feedforward gain and 25% adapt to a class 5 for 0.5 feedforward gain. Conversely, some R3L5 drivers however, gain efficiency and reclassify as more 'efficient' classes for left curves with R4L4 guidance.

The adaptations made when drivers are subjected to a personalised HCR seems to induce the most adaptation for left curves. Infact, many of the R3L5 drivers reclassify as class 1 or 2 (25.00% for FF 0.50 and 56.25% for FF 0.92) and the R2L2 drivers, part of the curve cutters, sometimes reclassify as class 1.

Table D.1: Consistency of the different driver groups towards their manual driving class.

Driver Group	Curve	Consistency
R2L1 drivers	Right	65%
	Left	72.50%
R2L2 drivers	Right	67.50
	Left	52.60%
R3L5 drivers	Right	80.68%
	Left	43.18%

Finally, Table D.1 summarises the overall consistency of the different driver groups. It shows how consistent the drivers are in retaining their natural driving style after being subjected to all the different conditions. This as an attempt to see which driver groups are less willing to adapt to the HSC than others. It seems that most drivers, regardless of their manual driving class, change their driving habits when subjected to haptic shared control. Some groups are very adamant about the continuation of their own trajectory, such as R2L1 drivers in left curves, while other groups are more adaptive, such as R3L5 drivers in left curves. Furthermore, the analysis of the left curves of R3L5 drivers showed low conflict values for all conditions, resulting in overall low conflict values. This might be linked to the R3L5 drivers being very willing to adapt their left curve behaviour as their consistency is only 43.18%.

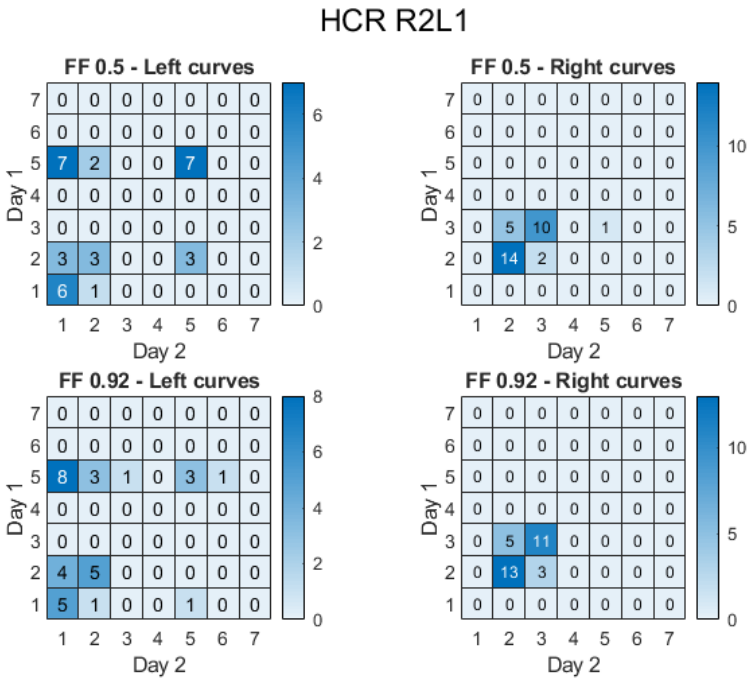

Figure D.10: This figure illustrated the reclassification results of all drivers after driving the R2L1 guidance. On the y-axis lie the classification results for day 1 and on the x-axis for day 2.

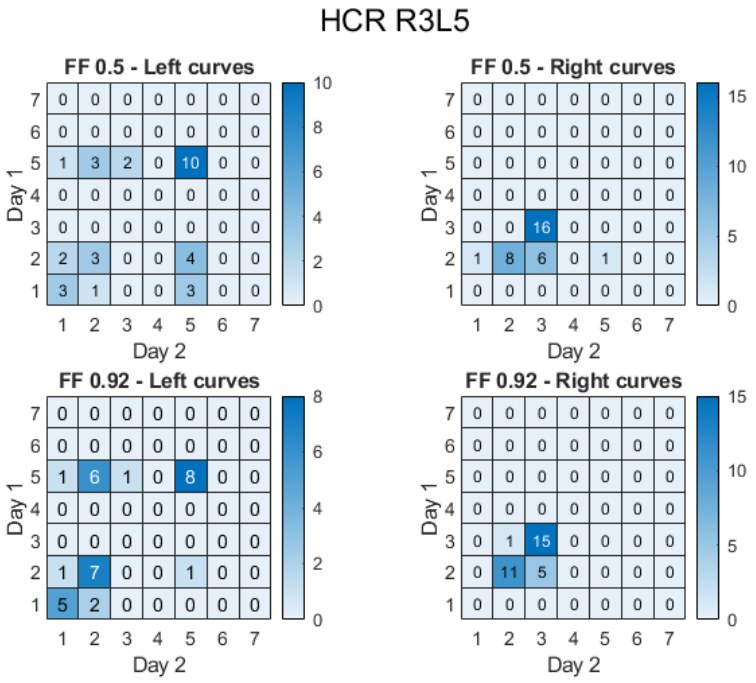

Figure D.11: This figure illustrated the reclassification of all drivers after driving the R3L5 guidance. On the y-axis lie the classification results for day 1 and on the x-axis for day 2.

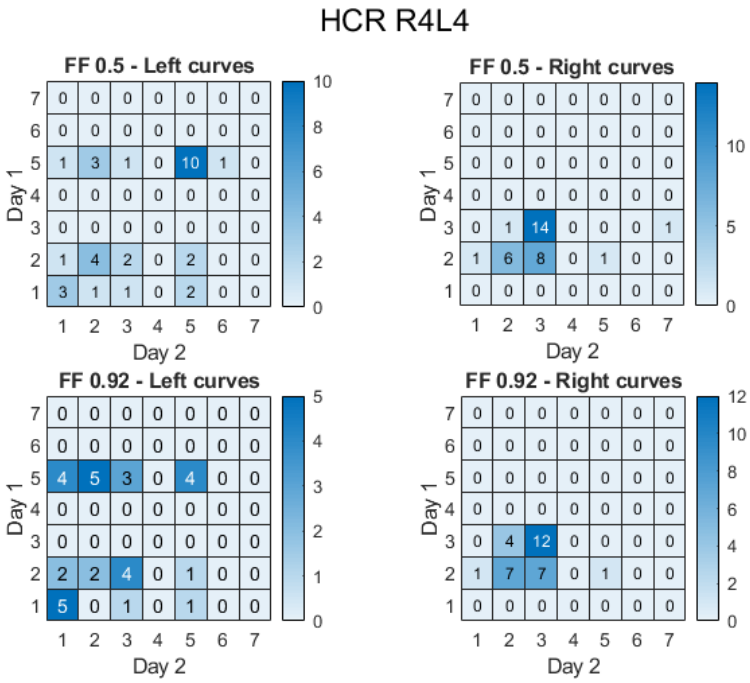

Figure D.12: This figure illustrated the reclassification of all drivers after driving the R4L4 guidance. On the y-axis lie the classification results for day 1 and on the x-axis for day 2.

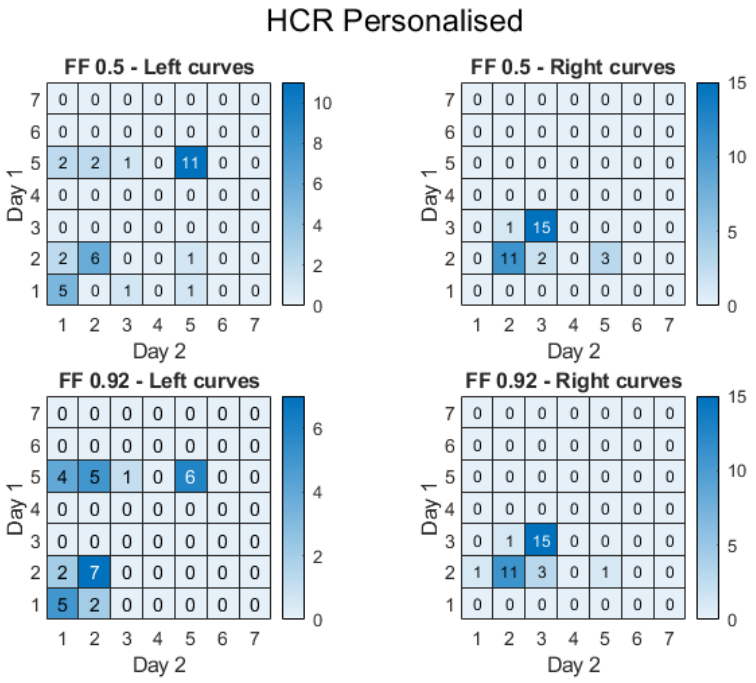

Figure D.13: This figure illustrated the reclassification of all drivers after driving the Personalised guidance. On the y-axis lie the classification results for day 1 and on the x-axis for day 2.

List of Publications

Journal Articles

1. **S. Barendswaard, D. Pool, D.A. Abbink**, *A method to assess individualized driver models: Descriptiveness, Identifiability and Realism*, Transportation Research Part F: Traffic Psychology and Behaviour **61**, 16-29 (2019).
2. **S. Barendswaard, D. Pool, M.M. Van Paassen, M. Mulder**, *Dual-Axis Manual Control: Performance Degradation, Axis Asymmetry, Crossfeed, and Intermittency*, IEEE Transactions on Human-Machine Systems **49**, 2 (2019).
3. **S. Barendswaard, E.R. Boer, D. Pool, D.A. Abbink,** *Comparison of Control-Theoretic Driver Models for Personalised Trajectory Generation using an Assessment Methodology*, IEEE Transactions on Human-Machine Systems, to be published.
4. **S. Barendswaard, E.R. Boer, M.M. Van Paassen, D. Pool, D.A. Abbink**, *A pre-curve prepositioning model to enhance the descriptiveness of specific control-theoretic driver steering models*, IEEE Transactions on Human-Machine Systems, to be published.
5. **S. Barendswaard, E. Ghys, E.R. Boer, D.A. Abbink, M.M. Van Paassen, D. Pool**, *Personalised Haptic Shared Control Approaches to Improve Driver Acceptance in Curve Driving*, IEEE Transactions on Human-Machine Systems, to be published.

Conference Articles

1. **S. Barendswaard, D. Pool, M. Mulder**, *Human Crossfeed in Dual-Axis Manual Control with Motion Feedback* , 13th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Human-Machine Systems, Kyoto, Japan (2016).
2. **K. van der El, S. Barendswaard, D. Pool, M.M. Van Paassen, M. Mulder**, *Effects of Preview Time on Human Control Behavior in Rate Tracking Tasks*, 13th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Human-Machine Systems, Kyoto, Japan (2016).
3. **W. Scholtens, S. Barendswaard, D. Pool, M.M. Van Paassen, D.A. Abbink,** *A New Haptic Shared Controller Reducing Steering Conflicts*, IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2018.
4. **W. Vreugdenhil, S. Barendswaard, D. A. Abbink, C. Borst, S.M. Petermeijer**, *Complementing Haptic Shared Control With Visual Feedback for Obstacle Avoidance*, IFAC-PapersOnLine 52(19) 2019.
5. **S. Barendswaard, D. Pool, E.R. Boer, D.A. Abbink,** *A Classification Method for Driver Trajectories during Curve-Negotiation*, IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2019.
6. **S. Barendswaard, L. van Breugel, B. Schelfaut, J. Sluijter, L. Zuiker, D. Pool, E.R. Boer, D.A. Abbink,** *The Effect of Velocity and Curve Radius on Driver Steering Behaviour before Curve Entry*, IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2019.

ACKNOWLEDGEMENTS

Do you know how my PhD journey started?

I am actually an offspring of aerospace; there I did both my BSc and MSc. After my masters, it was clear to me that I have a strong affinity and capability to continue in research. I got to know from Max that a joint research just started between Control and Simulation and Cognitive Robotics. Max then introduced me to David. That is how the journey began!

Max, this thesis owes you great thanks, without your trust in my qualities and drive I may not have been getting this opportunity. Not forgetting the fact that you also supervised me in the start phase of my PhD, I am incredibly grateful to your continuous fatherly guidance, critical pieces of advice and talks on the wisdom of life.

Now, the mission of my PhD- journey is successfully accomplished! But successfulness of this journey is only made possible by many people's outstanding contributions from Control and Simulations, and Cognitive Robotics. From the bottom of my heart, I would like to say a big thank you for **all** who made this quest successful. I wouldn't be able to name everybody who contributed to this research since there are at least around 200 participants who were involved in the experiments. Nevertheless, special and personal attention will be given to the major players in this acknowledgement.

David, my promotor, I am extremely grateful for your patience and encouragement during the PhD. My learning curve owes you a great thanks, in your unique way of guidance, you taught me to focus on the whole picture, that is, to consider the why before the how. I learnt to spot what I need to convey in presentations and articles at a very early stage. Thank you indeed for the gentle but sharp guidance. David, you gave me a lot of freedom to be innovative and creative and supported me accordingly when I came with new ideas. Also, you let me lecture McRuer for three years in your human-controller course. Thank you for giving me space, trust and motivation.

Daan, my co-promotor and 24-7 supervisor :). You have rooted and shaped my scientific way of thinking, ever since the time of my MSc, your guidance played a crucial role in my academic development throughout the last six years. If I would have to write all that I am indebted to you within my learning curve, even a chapter may not suffice. Your perfectionistic and hard-core never-enough attitude is contagious, and so you taught me to persevere. Besides that, you are very realistic and have guided me to avoid pitfalls. Thank you for always diligently and carefully reading all my work, there is no single sentence in this thesis that has managed to escape your red pen!

Erwin, my co-promotor and skype- supervisor ;). I am immensely grateful for all your dexterous input, always adding a new dimension to my work. Your brilliant and profound understanding of human behaviour to human-machine interaction, motivated me to be more innovative, challenge hypotheses, solve a complicated situation, and break it down into clear and straightforward theories. When I was drowning in the vast amount of information, you asked me sharp questions, which led me to land on 'epiphany' shore. Thank you very much for always having time for me, even at the most unorthodox times like 6 am or Saturdays!

Rene, your contribution and guidance were enormous, on the order of a co-promotor, thank you. Especially the guidelines you gave me on how to tackle DUECA are uncountable. You also inspired me to come with the novel ideas of chapter 6: incorporating prepositioning curvature.

Your FDCA concept shaped chapter 7 and 8. Guru, your concepts from another dimension and the way you think have been an enrichment for my thesis. Thank you.

Harold and **Andries**, your generous support during experimental HMI-lab work are highly appreciated. Besides, you were always on standby for helping with unprecedented software/hardware issues. Thank you for your highly valued assistance.

Clark, I am sincerely grateful for your indirect contribution to this thesis through your brotherly guidance, wise bits of advice and the support you gave from the time of my MSc till now. Our countless conversations and revival brainstorms were inspiring and enriched me to think outside of the box and form multiple perspectives. I learned a lot from you during the supervision of Wilco: making jokes but still being serious about quality work; apparently, you manage to apply the 'human-centred' principles to supervision!

A special thanks to Wietske, Emma and Wilco. Supervising you during your MSc thesis period has been a source of great inspiration and knowledge. **Wietske**, your work is the backbone of Chapter 7, thank you for your tenacious attitude and extraordinary approach to literature scavenging. **Emma**, your work is the backbone of Chapter 8. Thank you for your voluntary perseverance to continue with phase 2 experiments under the most unorthodox conditions during corona time, and, your exceptional presentation skills especially when dealing with so much data. **Wilco**, your work that we published together resulted in my most recommended paper on Research Gate and used in important discussion material. The collaboration with you was insightful and enjoyable, the quality of work you produced was impeccable and your perseverance to do HMI experiments till 10 pm was admirable. Thank you.

As a part of my PhD journey, I got the rewarding experience of supervising two BSc groups. The first group, **Julian**, **Rik**, **Daan** and **Isaac**. The data of the experiment you conducted was used to verify my developed classifiers in Chapter 4. You were an extraordinary and unique group with a high grade of energetic teamwork, including lots of fun and excitement. Supervising you was super enjoyable and inspiring, thank you. Julian, your particular way of briefing subjects was unique, entertaining and memorable. The second group, **Luuk**, **Bart**, **Jim** and **Lourens**. Your work is the backbone of Chapter 5. Working with you was not only motivating but also insightful. Your perseverance to produce high-quality work resulted in a strong foundation that grew into my most viewed publication on Research Gate, thank you.

I am grateful to all my colleagues in both Cognitive Robotics and Control and Simulation. We have had numerous talks, coffees, barbeques, laughs, and trips contributing to creating a healthy, enjoyable and inspiring working environment. A special thanks to my 'comrades', **Sarvesh** and **Timo** for all the splendid times we worked together on assignments, arranging numerous demo's, preparing course material, taking courses such as Pattern Recognition and most of all, the great conversations and fun we had. **Bastiaan**, **Joris**, **Niek** and **Arkady**, thank you for all the great discussions, coffees and laughs. **Tugrul**, our rewarding talks on the philosophy of life and your peculiar motion sickness experiment on a rotating chair was much fun, unique and remarkable, thank you.

Dirk, your expert knowledge in DUECA and your hobby of helping others is impressive and much appreciated. You were a tremendous help in setting up numerous DUECA experiments, without your help, I might have broken down some computers :) You were always available to: reflect upon ideas, a fun chat or a cup of coffee with your delicious home-baked cakes. Thank you, **Kasper**, all the 'summer school' courses, conferences and numerous discussions we had together were invaluable. **Isabel**, thank you for all our energetic and fun work and non-work related conversations during our regular catch-up coffees/lunches. Also, thank you **Malik**, **Noor**, **Mario**, **Daniel**, **Wei**, **Jelmer**, **Paolo**, **Sihao**, **Rolf**, **Sven**, **Diana**, **Rowenna**, **Jork**, **Sherry** and **Tao** for the great company.

I am very grateful to my friends who helped me to recharge my energy: **Sebahat**, **Leyla**, **Hatice**, **Mohamed**, **Yunusi**, **Hanife**, **Bushra**, **Samir**, **Raja**, **Haseeb**, **Isra**, **Abdelrahman**, **Ruwaida**, **Elisabeth** and **Kaoutar**. But especially my 'besties': **Ezgi**, **Fatma**, **Noura**, **Nema** and **Zaid**, thank you for the warm company, dinners, rejuvenating trips and great fun during our free time. This was the motor that continuously charged my motivation and boosted my psychology.

Last but foremost, my mother, she played a crucial supporting role in my life, that of both a mother and father. Mom, I would like to express my deepest gratitude for your unconditional love and care. Your support and loving presence have motivated and encouraged me tremendously throughout the journey of my PhD. There is not a page in this thesis which you have not read, revised and commented on. My gratitude can not be put into words. Thank you for being the secret of my success. Thank you from the bottom of my heart.

Sarah Barendswaard *Den Haag, February 2021*

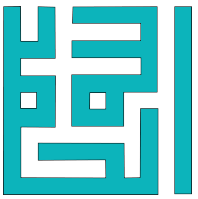

CURRICULUM VITÆ

Sarah BARENDSWAARD

20-07-1992 Born in Sittard, The Netherlands.

EDUCATION

1996–2003	International Schooling International School of Maastricht (1996–2005) International School of The Hague (2005–2010)
2010–2016	Univeristy Education BSc in Aerospace Engineering, TU Delft (2010-2013) MSc in Control and Simulation, Aerospace Engineering, TU Delft (2013-2016)
2016–2021	PhD. Cognitive Robotics, TU Delft Personalisation of Haptic Shared Control

AWARDS

2011	Propadeuse Cum Laude
2016	MSc Cum Laude
2016	Nomination for the NVvL Wittenberg prize, for the best aerospace MSc thesis in the Netherlands

ACADEMIC ACTIVITIES

2016-2017	Chairwoman of the Haptic's Lab meetings
2017-2019	Guest lecturing in the Human-Controller course for 3 years
Nov 2018	Invited to give a talk at Max Planck Institute of Biological Cybernetics
2018-2020	Reviewed 4 IEEE conference articles and 1 IEEE journal article

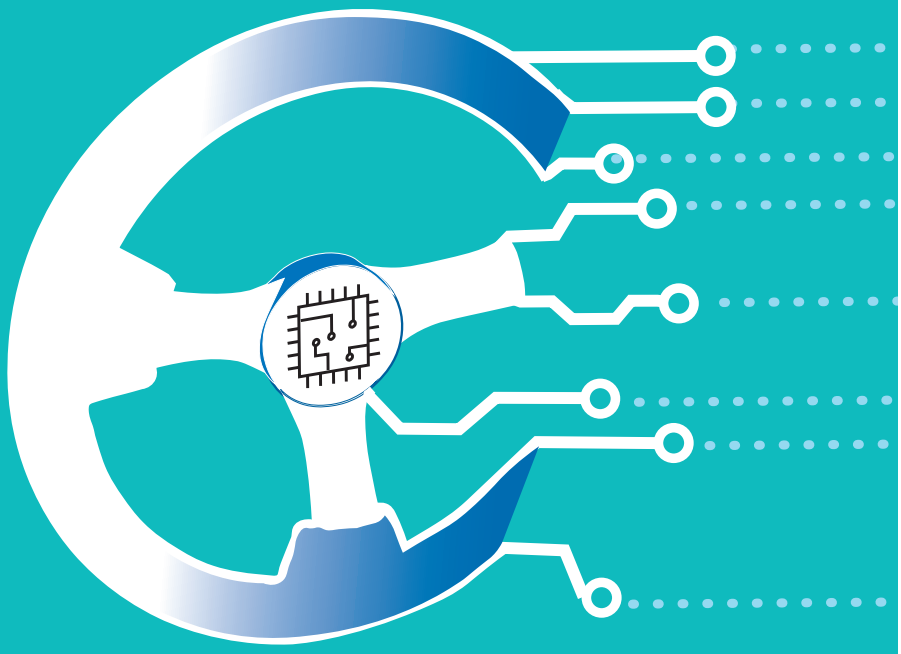

ISBN 9789464212556