

Control of Evasive Manoeuvres for Automated Driving Solving the Edge Cases

Bertipaglia, A.

DOI

[10.4233/uuid:05ba52a6-72ed-4259-9541-bf2b09c9a8b6](https://doi.org/10.4233/uuid:05ba52a6-72ed-4259-9541-bf2b09c9a8b6)

Publication date

2025

Document Version

Final published version

Citation (APA)

Bertipaglia, A. (2025). *Control of Evasive Manoeuvres for Automated Driving: Solving the Edge Cases*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:05ba52a6-72ed-4259-9541-bf2b09c9a8b6>

Important note

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

Copyright

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

Takedown policy

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

CONTROL OF EVASIVE MANOEUVRES FOR AUTOMATED VEHICLES

SOLVING THE EDGE CASES



ALBERTO BERTIPAGLIA

CONTROL OF EVASIVE MANOEUVRES FOR AUTOMATED DRIVING:

SOLVING THE EDGE CASES

CONTROL OF EVASIVE MANOEUVRES FOR AUTOMATED DRIVING:

SOLVING THE EDGE CASES

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on
Monday 24th, February 2025 at 15:00 o'clock

by

Alberto BERTIPAGLIA

Master of Science in
Automotive Engineering,
Politecnico di Torino, Turin, Italy.
born in Novi Ligure, Italy.

This dissertation has been approved by the promotor.

promotor: Prof. dr. ir. R. Happee

promotor: Dr. B. Shyrokau

copromotor: Dr. M. Alirezaei

Composition of the doctoral committee:

Rector Magnificus

chairperson

Prof. dr. ir. R. Happee

Delft University of Technology, promotor

Dr. B. Shyrokau

Delft University of Technology, promotor

Dr. M. Alirezaei

Eindhoven University of Technology, copromotor

Independent members:

Prof. dr. ir. B. De Schutter

Delft University of Technology, The Netherlands

Prof. dr. G. Mastinu

Politecnico di Milano, Italy

Prof. dr. E. Velenis

Cranfield University, United Kingdom

Prof. dr. ir. M. Wisse

Delft University of Technology, The Netherlands

Reserve member:

Dr. G. Papaioannou

Delft University of Technology, The Netherlands



Keywords:

automated vehicles, collision avoidance, handling limits, physics-informed neural networks, unscented Kalman filter, model predictive control, Student-t process, learning-based model predictive control

Printed by:

Gildeprint Drukkerijen

Cover:

Giulia Beck and Alberto Bertipaglia

Style:

TU Delft House Style, with modifications by Alberto Bertipaglia, Moritz Beller and Dr. Andras Pallfy

<https://github.com/Inventitech/phd-thesis-template>

ISBN 978-94-6518-006-9

An electronic version of this dissertation is available at

<http://repository.tudelft.nl/>.

*Research is meant as a tool for knowledge,
not as an object of competition or a means of power.*

Rita Levi-Montalcini

CONTENTS

Table of Contents	vii
List of Abbreviations	xi
Summary	xv
Samenvatting	xix
1 Introduction	1
1.1 Automated Vehicles	4
1.2 Scope of this Thesis	5
1.2.1 Sub-Goal 1: Vehicle Sideslip Angle Estimation	7
1.2.2 Sub-Goal 2: Obstacle Avoidance at The Limit of Handling	7
1.2.3 Sub-Goal 3: Automated Drifting	7
1.3 Thesis Outline and Contributions	8
1.3.1 Chapter 2: An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation	10
1.3.2 Chapter 3: Model Predictive Contouring Control for Vehicle Obsta- cle Avoidance at the Limit of Handling	10
1.3.3 Chapter 4: On the Benefits of Torque Vectoring for Automated Collision Avoidance at the Limits of Handling	11
1.3.4 Chapter 5: A Learning-Based Model Predictive Contouring Control for Vehicle Evasive Manoeuvres	12
1.3.5 Chapter 6: A Nonlinear Model Predictive Control for Automated Drifting with a Standard Passenger Vehicle	13
2 An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation	15
2.1 Introduction	17
2.2 Related Works	19
2.3 UKF-Informed Neural Network	23
2.3.1 Data-driven Component	24
2.3.2 Model-based Component	25
2.3.3 Training Phase	27
2.4 Experiment Setup	29
2.4.1 Experimental Setup and Dataset	29
2.4.2 Key Performance Indicators	30
2.4.3 Baseline Methods	31

2.5	Results	32
2.5.1	Full Dataset Results	32
2.5.2	Robustness Analysis using the Limited Dataset	37
2.5.3	Robustness Analysis to Tyre Model Parameters	40
2.6	Conclusion	41
3	Model Predictive Contouring Control for Vehicle Obstacle Avoidance at the Limit of Handling	43
3.1	Introduction	45
3.2	Related Works	46
3.3	Prediction Model	48
3.4	Model Predictive Contouring Control	49
3.4.1	Cost Function	50
3.4.2	Constraints	51
3.4.3	Obstacle Avoidance Prioritisation	52
3.5	Simulation and Experimental Setup	52
3.5.1	Simulation Setup	52
3.5.2	Experimental Setup	53
3.6	Results	54
3.6.1	Cartesian vs Frenet Reference Frame	54
3.6.2	Simulation Results	55
3.6.3	Experimental Results	57
3.7	Conclusion	58
4	On the Benefits of Torque Vectoring for Automated Collision Avoidance at the Limits of Handling	61
4.1	Introduction	63
4.2	Related Works	64
4.3	Prediction Model	66
4.3.1	Double-Track Vehicle Model	66
4.3.2	Extended Fiala Tyre Model	68
4.4	Model Predictive Contouring Control Using Torque Vectoring	71
4.4.1	Cost Function with Obstacle Avoidance Prioritisation	71
4.4.2	Constraints	73
4.5	Simulation Setup and Experimental Validation	74
4.5.1	Hard Real-Time Implementation and Simulation Setup	74
4.5.2	Experimental Model Validation	76
4.5.3	Robustness and Sensitivity Analysis Setup	77
4.6	Results	80
4.6.1	Controller Performance	80
4.6.2	Robustness Analysis	83
4.6.3	Sensitivity Analysis	86
4.7	Conclusion	87

5	A Learning-Based Model Predictive Contouring Control for Vehicle Evasive Manoeuvres	89
5.1	Introduction	91
5.2	Learning-based Model Predictive Contouring Control.	92
5.3	Student-T Process & Uncertainty Propagation	93
5.4	Results	94
5.5	Conclusions	95
6	A Nonlinear Model Predictive Control for Automated Drifting with a Standard Passenger Vehicle	97
6.1	Introduction	99
6.2	Vehicle Model	100
6.2.1	Single-Track Vehicle Model	100
6.2.2	Tyre Model.	101
6.3	Equilibrium Analysis.	101
6.3.1	Steady-State Equilibrium Locations.	101
6.3.2	Experimental Validation	103
6.4	Proposed Controller Architecture	104
6.5	Simulation Results	105
6.5.1	Simulation Setup	105
6.5.2	Alternating Drifting Manoeuvre	105
6.6	Experimental Results.	107
6.6.1	Experimental Setup.	107
6.6.2	Semi-Automated Drifting Manoeuvre.	108
6.7	Conclusions	109
7	Conclusions & Recommendations	111
7.1	Discussion & Main Findings	113
7.1.1	Sub-Goal 1: Vehicle Sideslip Angle Estimation	113
7.1.2	Sub-Goal 2: Obstacle Avoidance at The Limit of Handling	115
7.1.3	Sub-Goal 3: Automated Drifting	119
7.2	Limitations & Future Works	120
7.2.1	Road Friction Coefficient.	121
7.2.2	One Manoeuvre Does Not Fit All.	122
7.2.3	Can We Avoid Obstacles While Drifting?.	122
7.2.4	Machine Learning	123
7.2.5	Automated Vehicles	124
7.3	General Conclusion	124
A	A Two-Stage Bayesian Optimisation for Automatic Tuning of an Unscented Kalman Filter for Vehicle Sideslip Angle Estimation	127
A.1	Introduction	129
A.2	Related Works	130
A.3	Bayesian Optimisation	132
A.3.1	Surrogate Model	133
A.3.2	Acquisition Function	134

A.4	Two-Stage Bayesian Optimisation	135
A.4.1	Fast Exploration	135
A.4.2	Pure Exploitation	136
A.5	Application: UKF Tuning	137
A.5.1	Dataset	137
A.5.2	Cost Function Formulation	137
A.5.3	Key Performance Indicator	138
A.6	Results	138
A.7	Conclusions	143
	References	145
	Acknowledgments	163
	Curriculum Vitæ	167
	List of Publications	169
	Propositions	171

LIST OF ABBREVIATIONS

ACADO.....	Automatic Control and Dynamic Optimization
AF.....	Acquisition Function
ANFIS.....	Adaptive Neuro-Fuzzy Inference System
ARD.....	Automatic Relevance Determination
AV.....	Automated Vehicles
BFGS.....	Broyden–Fletcher–Goldfarb–Shanno
BO.....	Bayesian Optimisation
BPTT.....	Back-Propagation Through Time
CA.....	Collision Avoidance
CAN.....	Controller Area Network
CBM.....	Confidence Bound Minimisation
CEO.....	Chief Executive Officer
CNN.....	Convolutional Neural Network
CoG.....	Centre of Gravity
CR.....	Collision Rate
DE.....	Deep Ensemble
ECU.....	Electronic Control Unit
EI.....	Expected Improvement
EKF.....	Extended Kalman Filter
EPS.....	Electric Power Steering
FFNN.....	Feed Forward Neural Network
FFO.....	Fruit Fly Optimisation
GA.....	Genetic Algorithm
GNSS.....	Global Navigation Satellite System

GP	Gaussian Process
GPML	Gaussian Processes for Machine Learning
GRU	Gated Recurrent Unit
IMU	Inertial Measurement Unit
KPI	Key Performance Indicator
L-MPCC	Learning-based Model Predictive Contouring Control
LPV	Linear Parameter-Varying
LQR	Linear Quadratic Regulator
LSD	Limited-Slip Differential
LSTM	Long Short-Term Memory
MAE	Maximum Absolute Error
MAE _{nl}	Nonlinear Maximum Absolute Error
ME	Maximum Error
ME _{nl}	Nonlinear Maximum Error
MPC	Model Predictive Control
MPCC	Model Predictive Contouring Control
MSE	Mean Squared Error
MSE _{nl}	Nonlinear Mean Squared Error
mVD	minimum Vehicle Distance
NMPC	Nonlinear Model Predictive Control
NMR	Near Miss Rate
NMVCCS	National Motor Vehicle Crash Causation Survey
NN	Neural Network
NRMSE	Normalised Root Mean Squared Error
ReLU	Rectified Linear Unit
RL	Reinforcement Learning
RMSE	Root Mean Squared Error
RMSE _{nl}	Nonlinear Root Mean Squared Error

RNN	Recurrent Neural Network
RTK	Real-Time Kinematic
SA	Simulated Annealing
SAE	Society of Automotive Engineers
SI	Sensitivity Indices
STP	Student-t Process
TNO	Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek
TSBO	Two-Stage Bayesian Optimisation
TV	Torque Vectoring
UKF	Unscented Kalman Filter
V2E	vehicle-to-edge
V2O	vehicle-to-obstacle

SUMMARY

This thesis addresses the challenge of controlling automated vehicles performing evasive manoeuvres at the limit of handling. Special attention is paid to the development of nonlinear controllers, which can prioritise obstacle avoidance over path tracking objectives while considering vehicle stability constraints, to improve passenger safety. The goal of this work is to develop the entire pipeline for obstacle avoidance controllers, focusing on three aspects: vehicle state estimation, collision avoidance and control beyond the stable handling limits, e.g. drifting. Vice-versa, this work does not aim to enhance motion planning or obstacle detection algorithms. First, a brief introduction to automated vehicles is provided to highlight their advantages and limitations. Subsequently, the introduction presents the operating principles of vehicle collision avoidance controllers, focusing on ensuring that automated vehicles can adeptly avoid obstacles in normal driving as well as the limits of handling. This capability is essential for maximising safety not only in normal driving but especially in emergency driving scenarios.

Active vehicle controls and collision avoidance systems play a crucial role in maintaining vehicle stability and enhancing safety, particularly when avoiding obstacles at the limit of handling. These systems rely on accurate vehicle sideslip angle and yaw rate information to ensure stability and controllability. A high sideslip angle indicates that the vehicle is nearing or has exceeded its handling limits, a condition where traditional control strategies may no longer suffice to prevent a loss of control. While low-cost gyroscopic sensors can directly measure the yaw rate, the vehicle sideslip angle typically needs to be estimated. Therefore, this thesis first focuses on the vehicle sideslip angle estimation. This thesis combines the accuracy of data-driven approaches, such as neural networks, in operational conditions they have been trained on with the robustness and reliability of physics-based models in unseen emergency scenarios. A novel physics-informed data-driven vehicle sideslip angle estimator is proposed, which leverages the physical knowledge from an Unscented Kalman Filter based on a nonlinear single-track vehicle model to enhance the estimation accuracy of a neural network. The hybrid approach employs a mutual relationship between the model-based and data-driven approaches enforced by the end-to-end training. The proposed architecture is developed on a large-scale experimental dataset of 216 manoeuvres, which contains a great diversity of vehicle behaviours. The results show a notable improvement in the accuracy of the proposed architecture compared to the current state-of-the-art hybrid approaches. Furthermore, the architecture demonstrates significant robustness even when trained on a limited dataset.

The thesis then explores how the vehicle dynamics and stability information can be leveraged to improve automated vehicles' ability to perform evasive manoeuvres. Generally, collision avoidance algorithms are based on a hierarchical controller architecture, which separates motion planning, path tracking, and vehicle stability tasks. Although each task can be optimised separately, in an evasive manoeuvre, the three objectives could be in conflict with each other. For instance, the vehicle stability constraints could lead to tracking errors,

potentially causing collisions. Therefore, to avoid potential conflict, a Model Predictive Contouring Control (MPCC) based on a nonlinear single-track vehicle model is proposed. The proposed approach integrates all the controllers' objectives into a single cost function and, if needed, modifies the desired trajectory to keep the vehicle stable and at a safe distance from the obstacle. To address this, the proposed MPCC computes a desired steering angle and a force at the vehicle's centre of gravity, which is used to brake and accelerate. Results in a high-fidelity simulation environment prove an enhanced performance in a double-lane change for vehicle obstacle avoidance at the limit of handling on a high friction surface. The proposed controller's performance is experimentally validated by performing a double-lane change evasive manoeuvre on a low friction surface using a TU Delft experimental Toyota Prius. Furthermore, the validation proves the robustness of the proposed MPCC to different road friction coefficients.

Following this, the thesis analyses how the recent development of new electric power-trains can be used in vehicle collision avoidance controllers to increase automated vehicle safety. In particular, a focus is put on vehicles employing multiple in-wheel electric motors. For this reason, the previously developed MPCC is extended to consider torque vectoring capabilities to improve the vehicle's lateral agility at the limit of handling. This is achieved thanks to the adoption of a double-track vehicle based on an extended Fiala tyre model while maintaining real-time feasibility on a rapid prototyping platform. Results show that the proposed approach successfully avoids obstacles and maintains vehicle stability in a high-fidelity simulation environment validated with experimental data collected at a proving ground. Extended Monte Carlo analyses demonstrate the controller's robustness to variations in vehicle parameters, electric motor delay, road friction, perception accuracy, number of false negatives in the perception stack, and localisation errors. Particularly relevant is the statistical assessment of how each variation impacts performance. The sensitivity analysis reveals that tyre parameters, along with obstacle localisation accuracy, are the most influential factors affecting the controller's performance.

Building on this, the thesis delves into strategies for mitigating the impact of tyre model mismatches on the performance of the MPCC. To address this, a Learning-based Model Predictive Contouring Control approach is proposed, utilising a Student-t Process (STP) to predict mismatches between the prediction model and the vehicle measurements. The selected measurements are the yaw rate, measured by the inertial measurement unit and the lateral tyre forces measured by intelligent (force-sensing) bearings. Unlike traditional approaches that focus on vehicle velocity states, the proposed method directly reduces mismatches in the tyre model, improving the controller's pro-activity. This results in successfully performing evasive manoeuvres at a higher velocity than the current state-of-the-art MPCC. The algorithm uses the STP not only to minimise model mismatches but also to evaluate the associated uncertainties. These are then propagated in the prediction horizon and employed in the MPCC cost function, where they are minimised. By doing so, the proposed controller reduces the vehicle's operating time in unstable conditions. As a result, the proposed approach performs evasive manoeuvres with greater stability than the current state-of-the-art learning-based MPCC.

Subsequently, the thesis aims to prove the capability to control and stabilise a vehicle beyond the stable handling limits. A novel approach to automated drifting with a standard passenger vehicle based on a Nonlinear Model Predictive Control is proposed. The proposed

architecture is split into three components. The first part consists of the offline computed equilibrium maps, which provide the equilibrium points for each vehicle state given the desired sideslip angle and radius of the path. The second is the predictive controller, which minimises the errors between the equilibrium and actual vehicle states. The third is a path-following controller, which reduces the path error, altering the equilibrium curvature path. This thesis experimentally verifies the correspondence between the vehicle state equilibria computed using a nonlinear single-track model and those measured in a proving ground. Experiments with a standard passenger vehicle demonstrate that the proposed approach can achieve and maintain a desired 30 degree sideslip angle in both high and low friction conditions in real-world testing.

In conclusion, this thesis introduced innovative and robust solutions to key challenges in obstacle avoidance control for automated driving. By addressing the complexities of tyre-road nonlinearities, uncertainties in perception and localisation, and the control of vehicles beyond stable handling limits, the proposed framework advances the safety and robustness of automated vehicles.

SAMENVATTING

Dit proefschrift behandelt de uitdaging van het besturen van geautomatiseerde voertuigen die ontwijkende manoeuvres uitvoeren op de limiet van handelbaarheid. Bijzondere aandacht gaat uit naar de ontwikkeling van niet-lineaire regelaars die obstakelvermijding kunnen prioriteren boven baanvolgingsdoelen, terwijl ze rekening houden met stabiliteitsbeperkingen van het voertuig om de veiligheid van passagiers te verbeteren. Het doel van dit werk is het ontwikkelen van de volledige pijplijn voor obstakelvermijdingsregelaars, met focus op drie aspecten: voertuigtoestandschatting, botsingsvermijding en controle voorbij de stabiele limiet van handelbaarheid, bijvoorbeeld driften. Omgekeerd is het niet de bedoeling van dit werk om bewegingsplanning of obstakeldetectie-algoritmen te verbeteren. Eerst wordt een korte introductie gegeven over geautomatiseerde voertuigen om hun voordelen en beperkingen te benadrukken. Vervolgens presenteert de introductie de werkingsprincipes van botsingsvermijdingsregelaars voor voertuigen, met de nadruk op het garanderen dat geautomatiseerde voertuigen behendig obstakels kunnen vermijden tijdens normaal rijden en op de limiet van handelbaarheid. Deze capaciteit is essentieel voor het maximaliseren van de veiligheid, niet alleen tijdens normaal rijden maar vooral in noodsituaties.

Actieve voertuigbesturingen en botsingsvermijdingssystemen spelen een cruciale rol bij het handhaven van voertuigstabiliteit en het verhogen van de veiligheid, vooral bij het vermijden van obstakels op de limiet van handelbaarheid. Deze systemen zijn afhankelijk van nauwkeurige informatie over de zijsliphoeck en giersnelheid van het voertuig om stabiliteit en bestuurbaarheid te waarborgen. Een hoge zijsliphoeck geeft aan dat het voertuig zijn limiet van handelbaarheid nadert of heeft overschreden, een toestand waarin traditionele regelstrategieën mogelijk niet langer volstaan om verlies van controle te voorkomen. Terwijl goedkope gyroscopische sensoren de giersnelheid direct kunnen meten, moet de zijsliphoeck van het voertuig doorgaans worden geschat. Daarom richt dit proefschrift zich eerst op de schatting van de zijsliphoeck van het voertuig. Dit proefschrift combineert de nauwkeurigheid van data-gedreven aanpakken, zoals neurale netwerken, in operationele omstandigheden waarmee ze zijn ontwikkeld, met de robuustheid en betrouwbaarheid van fysisch-gebaseerde modellen in nieuwe noodscenario's. Een nieuw fysica-geïnformeerde, data-gedreven schatter voor de zijsliphoeck van het voertuig wordt voorgesteld, die de fysieke kennis uit een Unscented Kalman Filter, gebaseerd op een niet-lineair single-track voertuigmodel, benut om de schattingsnauwkeurigheid van een neurale netwerk te verbeteren. De hybride aanpak maakt gebruik van een wederzijdse relatie tussen de modelgebaseerde en data-gedreven benaderingen, afgedwongen door de end-to-end optimalisatie. De voorgestelde architectuur is geoptimaliseerd met een grootschalige experimentele dataset met 216 manoeuvres, die een grote diversiteit aan voertuiggedragingen bevat. De resultaten tonen een aanzienlijke verbetering in de nauwkeurigheid van de voorgestelde architectuur vergeleken met de huidige state-of-the-art hybride benaderingen. Bovendien vertoont de architectuur significante robuustheid, zelfs wanneer geoptimaliseerd met een beperkte dataset.

Het proefschrift verkent vervolgens hoe de informatie over voertuigdynamica en stabiliteit kan worden benut om het vermogen van geautomatiseerde voertuigen te verbeteren bij het uitvoeren van ontwijkende manoeuvres. Over het algemeen zijn botsingsvermijdingsalgoritmen gebaseerd op een hiërarchische regelaararchitectuur, die bewegingsplanning, baanvolging en voertuigstabiliteitstaken scheidt. Hoewel elke taak afzonderlijk kan worden geoptimaliseerd, kunnen bij een ontwijkende manoeuvre de drie doelstellingen met elkaar in conflict komen. Bijvoorbeeld, de voertuigstabiliteitsbeperkingen kunnen leiden tot volgfouten, wat mogelijk tot botsingen kan leiden. Daarom wordt, om potentiële conflicten te vermijden, een Model Predictive Contouring Control (MPCC) voorgesteld, gebaseerd op een niet-lineair single-track voertuigmodel. De voorgestelde aanpak integreert alle doelstellingen van de controllers in een enkele kostfunctie en past, indien nodig, het gewenste traject aan om het voertuig stabiel te houden en op veilige afstand van het obstakel. Om dit te bereiken, berekent de voorgestelde MPCC een gewenste stuurhoek en een kracht op het zwaartepunt van het voertuig, die wordt gebruikt om te remmen en te accelereren. Resultaten in een hoogwaardige simulatieomgeving bewijzen een verbeterde prestatie in een dubbele rijstrookwissel voor voertuigobstakelvermijding op de limiet van handelbaarheid op een oppervlak met hoge wrijving. De prestaties van de voorgestelde regelaar worden experimenteel gevalideerd door het uitvoeren van een dubbele rijstrookwissel ontwijkende manoeuvre op een oppervlak met lage wrijving met behulp van een experimentele Toyota Prius van de TU Delft. Bovendien bevestigt de validatie de robuustheid van de voorgestelde MPCC voor verschillende wegwrijvingscoëfficiënten.

Vervolgens analyseert dit proefschrift hoe de recente ontwikkeling van nieuwe elektrische aandrijflijnen kan worden ingezet in botsingsvermijdingsregelaars om de veiligheid van geautomatiseerde voertuigen te verhogen. In het bijzonder wordt de focus gelegd op voertuigen met meerdere in-wiel elektromotoren. Om deze reden wordt de eerder ontwikkelde Model Predictive Contouring Control (MPCC) uitgebreid om rekening te houden met koppelverdeling (torque vectoring) capaciteiten, met als doel de laterale wendbaarheid van het voertuig op de limiet van handelbaarheid te verbeteren. Dit wordt bereikt door de toepassing van een double-track voertuigmodel gebaseerd op een uitgebreid Fiala-bandenmodel, terwijl de real-time uitvoerbaarheid op een rapid prototyping platform behouden blijft. De resultaten tonen aan dat de voorgestelde aanpak obstakels succesvol vermijdt en de voertuigstabiliteit handhaaft in een hoogwaardige simulatieomgeving, gevalideerd met experimentele data verzameld op een testcircuit. Uitgebreide Monte Carlo-analyses demonstreren de robuustheid van de regelaar tegen variaties in voertuigparameters, wegwrijving, perceptienauwkeurigheid, het aantal fout-negatieven in perceptiesysteem en localisatiefouten. Bijzonder relevant is de statistische beoordeling van hoe elke variatie de prestaties beïnvloedt. De gevoeligheidsanalyse onthult dat bandparameters, samen met de nauwkeurigheid van obstakellocalisatie, de meest invloedrijke factoren zijn die de prestaties van de regelaar beïnvloeden.

Voortbouwend hierop verdiept het proefschrift zich in strategieën om de impact van het verschil in bandmodel op de prestaties van de MPCC te verminderen. Om dit aan te pakken, wordt een op leren gebaseerde Model Predictive Contouring Control-benadering voorgesteld, waarbij een Student-t-proces (STP) wordt gebruikt om verschillen tussen het voorspellingsmodel en de voertuigmetingen te voorspellen. De geselecteerde metingen zijn de giersnelheid, gemeten door de inertiaële meeteenheid (IMU), en de laterale bandkrachten, gemeten door intelligente (kracht-sensorische) lagers. In tegenstelling tot tradi-

tionale benaderingen die zich richten op snelheidstoestanden van het voertuig, vermindert de voorgestelde methode direct verschillen in het bandmodel, wat de proactiviteit van de regelaar verbetert. Dit resulteert in het succesvol uitvoeren van ontwijkende manoeuvres bij een hogere snelheid dan de huidige state-of-the-art MPCC. Het algoritme gebruikt de STP niet alleen om modelverschillen te minimaliseren, maar ook om de bijbehorende onzekerheden te evalueren. Deze worden vervolgens gepropageerd in de voorspellingshorizon en toegepast in de kostfunctie van de MPCC, waar ze worden geminimaliseerd. Door dit te doen, vermindert de voorgestelde regelaar de tijd dat het voertuig in onstabiele omstandigheden opereert. Als gevolg hiervan voert de voorgestelde benadering ontwijkende manoeuvres uit met grotere stabiliteit dan de huidige state-of-the-art op leren gebaseerde MPCC.

Daaropvolgend heeft dit proefschrift tot doel de mogelijkheid aan te tonen om een voertuig te besturen en te stabiliseren voorbij de stabiele limieten van handelbaarheid. Een nieuwe benadering voor geautomatiseerd driften met een standaard personenvoertuig, gebaseerd op een niet-lineaire modelvoorspellende regelaar, wordt voorgesteld. De voorgestelde architectuur is opgedeeld in drie componenten. Het eerste deel bestaat uit de offline berekende evenwichtskaarten, die de evenwichtspunten voor elke voertuigtoestand leveren, gegeven de gewenste zijsliphoek en de straal van het pad. De tweede is de voorspellende regelaar, die de fouten tussen de evenwichtstoestand en de werkelijke voertuigtoestand minimaliseert. De derde is een padvolgregelaar, die de padfout vermindert door de evenwichtskromming van het pad aan te passen. Dit proefschrift verifieert experimenteel de overeenkomsten tussen de toestandsevenwichten van het voertuig berekend met behulp van een niet-lineair single-track model en die gemeten op een testcircuit. Experimenten met een standaard personenvoertuig tonen aan dat de voorgestelde benadering een gewenste zijsliphoek van 30 graden kan bereiken en handhaven onder zowel hoge als lage wrijvingsomstandigheden in tests in de praktijk.

Concluderend heeft dit proefschrift innovatieve en robuuste oplossingen geïntroduceerd voor belangrijke uitdagingen in obstakelvermijdingscontrole voor geautomatiseerd rijden. Door de complexiteiten van band-weg niet-lineariteiten, onzekerheden in perceptie en lokalisatie, en de controle van voertuigen voorbij de stabiele limiet van handelbaarheid aan te pakken, verbetert het voorgestelde kader de veiligheid en robuustheid van geautomatiseerde voertuigen.

1

INTRODUCTION

*Consider well the seed that gave you birth:
you were not made to live your lives as brutes,
but to be followers of worth and knowledge*

Canto XXVI, Inferno, The Divine Comedy, Dante Alighieri

THE vision of automated vehicles (AV) has captivated people all over the world, driven by the goal of achieving a "zero fatalities" future. Automated driving technology aims to revolutionise transportation by eliminating human error. Advances in artificial intelligence, motion planning, and vehicle control are rapidly propelling this vision forward, making automated vehicles an increasingly attainable reality. For instance, in 2015, Mr. Elon Musk, CEO of Tesla Motors, predicted his cars would be capable of "complete autonomy" by 2017 [1]. General Motors said in 2018 that it would launch a fleet of cars without steering wheels or pedals in 2019. Waymo, the Alphabet subsidiary widely seen as the industry leader, committed itself to launching a driverless taxi service in Phoenix at the end of 2018 [1]. Together with the increased expectations, from 2014 to 2017, investment in automated vehicle technology increased to 80 billion [2], with accompanying spikes in acquisitions, publications, patents, and media coverage. However, recent history has shown that the arrival of fully automated vehicles was overestimated, and only incremental advancements, like lane-keeping and automatic parking, are currently widespread in the automotive market. A series of self-driving car accidents dampened expectations, and market valuations have adjusted accordingly. Waymo has faced delays, leading to a reduced valuation by 40 % in 2019 [3]. The unforeseen challenges, such as frequent disengagements, revealed the limits of the technology, causing the industry to follow a typical hype cycle, from inflated expectations to disillusionment, see Fig 1.1.

Despite setbacks delaying the widespread deployment of automated vehicles, the mobility community continues to agree that automated driving has the potential to reshape transportation [5]. For instance, in June 2024, Waymo launched its fully operational robotaxi service for the public in San Francisco [6], with plans to expand the service to Austin and Atlanta by 2025 [7]. These recent advancements highlight the interest in automated driving to deliver new levels of safety and convenience to consumers and generate value within the automotive industry.

This chapter first provides an overview of what constitutes an automated vehicle and how such technologies can enhance vehicle safety. It then explores how automated vehicles improve safety during emergency manoeuvres. Finally, the chapter presents this thesis's contributions to advancing the safety of automated vehicles in emergency situations.




Figure 1.1: The Gartner hype cycle [4].

1.1 AUTOMATED VEHICLES

In 2018, the Society of Automotive Engineers (SAE) introduced a classification system for automated driving, known as SAE J3016, which defines six levels of automation [8]. These levels are as follows:

- Level 0 (L0, manual driving): The vehicle lacks any form of driving automation technology, and the driver is responsible for all tasks.
- Level 1 (L1, driver assistance): Some systems assist with specific driving tasks, such as acceleration, braking, or steering. An example is adaptive cruise control.
- Level 2 (L2, partial driving automation): Vehicles at this level provide partial automation, taking over acceleration, braking, and steering in specific scenarios while the driver remains engaged. Examples are the adaptive cruise control plus lane keeping or the Tesla Autopilot [9].
- Level 3 (L3, conditional driving automation): The vehicle can independently manage all driving tasks under certain conditions, but the driver must be ready to take control when requested by the system. An example is the Drive Pilot for highways by Mercedes [10].
- Level 4 (L4, high driving automation): These vehicles can fully control all driving tasks without human intervention, but their operation is typically limited to specific conditions or predefined areas, such as certain weather conditions or geographic zones (e.g., designated routes). For instance, Waymo robotaxi [6].
- Level 5 (L5, full driving automation): This level represents full automation, where the vehicle can operate independently in any conditions or environment without the need for human oversight.

Currently, Level 2 (partial automation) and Level 3 conditional driving automation are the prevailing standards among many automakers, where vehicles still rely on human attention and intervention. However, the widespread interest from diverse societal groups can likely be attributed to the anticipated and promised advantages of fully automated (Level 5) vehicles. Broadly speaking, the most commonly envisioned benefits can be split into three key categories: safety, comfort, and economic impact [11–13]. This thesis will focus on the widely recognised argument that traffic safety is the primary justification for the introduction of self-driving cars [14, 15]. Driver error is the leading cause in approximately 94 % of motor vehicle crashes, according to the National Motor Vehicle Crash Causation Survey (NMVCCS) conducted by the National Highway Traffic Safety Administration, [16]. Currently, vehicle active safety systems, e.g. autonomous emergency braking, lane departure warning, and electronic stability control, help mitigate the human risk factor in vehicle collisions and have been effective in reducing the types of crashes [17]. Only fully automated vehicles (Level 5 automation) could potentially eliminate all the crashes resulting from human error [18]. Unlike humans, AVs are not limited by perceptual shortcomings and are immune to risks such as distraction, inattention, alcohol impairment, or incapacitation [19]. However, AV behaviour will still be programmed by humans, often modelled on idealised human driving behaviour to ensure acceptance by riders and other

road users [20]. This leads to the crucial task of examining the role of human error in crash causation. Using data from the NMVCCS, crashes can be attributed to five driver factors: sensing and perceiving (failure to recognise hazards), predicting (misjudgement of other vehicles' behaviour), planning/deciding (poor decision-making related to traffic laws), execution/performance (inadequate vehicle control), and incapacitation (impairment due to alcohol or other causes) [19]. Given the expected superior perception capabilities of AVs and their immunity to incapacitation, it is expected that a significant portion of crashes caused by these two factors could be prevented by AV technology. However, research indicates that approximately 67 % of crashes may still persist even when the sensing/perception and incapacitation categories are removed from the causes, but were not designed to avoid other types of human error, particularly planning/deciding (41 %), execution/performance (23 %), and predicting (17 %) factors [19].

Therefore, this thesis specifically addresses potential crashes related to the decision-making and execution performance of the driver. Particular attention is placed on the most dangerous scenarios, such as incorrect evasive actions when the vehicle is subjected to a high lateral acceleration [21]. A key contributor to these potential collisions is the complexity of vehicle behaviour in these situations. At the limits of handling, vehicle dynamics become highly nonlinear due to the interaction between the tyres and the road surface. This nonlinearity is further complicated by variable factors like road conditions (e.g., wet, icy, or dry surfaces) and by the coupling between longitudinal and lateral forces, making precise vehicle control extremely challenging. Most drivers, except for highly trained professionals, lack the experience or skill to execute effective evasive actions in emergencies, particularly when the vehicle approaches or exceeds its handling limits. To address this, the thesis focuses on developing an advanced obstacle avoidance controller as part of the automated vehicle control system. This critical component is responsible for vehicle state estimation to assess stability and controllability, vehicle motion control for obstacle avoidance at the limit of handling with a focus on prioritising safety over path tracking and stability, and automated drifting to stabilise and control the vehicle beyond the stable handling limits or at high sideslip angles. By developing these solutions, the thesis aims to bridge the gap in automated vehicle safety during emergency manoeuvres, providing a robust system capable of safely performing evasive actions in the most challenging driving conditions.

1.2 SCOPE OF THIS THESIS

AUTOMATED vehicles have proven they can effectively reduce the number of injury-causing crashes in the first ≈ 35 million km driven fully automated by Waymo [22, 23]. Although the relatively low number of crashes involving automated vehicles makes it difficult to statistically analyse collision types, there is broad consensus that events leading to potential high-severity collisions are rare in real-world driving [24, 25]. However, the primary challenge for AV design lies not only in reducing actual collisions but in effectively managing potential high-severity collision scenarios. These high-risk events, though infrequent, are critical for ensuring safety and demand superior performance from the system. While AVs can already excel in routine driving, high-severity situations push the system to respond to rare but critical events that significantly impact overall safety outcomes [22, 24, 25].

An example of these critical situations is when a vehicle, person, or object may suddenly

appear in front of a car, leaving insufficient distance to avoid a crash through emergency braking alone. In such cases, an evasive manoeuvre involving precise steering is necessary. An abrupt movement of the steering wheel in an attempt to avoid the obstacle can result in loss of control, causing the vehicle to spin or collide with other cars [19, 26]. Automated vehicles face similar challenges when performing evasive manoeuvres because most driving technologies are programmed for moderate lateral accelerations and standard, smooth manoeuvres. However, vehicle dynamics change significantly during extreme manoeuvres, with the tyres entering a nonlinear operating mode that cannot be accurately measured by standard sensing technologies [27]. For this reason, there is strong interest in new sensor technologies that can provide accurate tyre-road force estimation during both regular and extreme manoeuvres. Recent examples include intelligent bearings developed by SKF [28–30], and Intelligent/Cyber tyres by Pirelli [31–33]. Simulations have shown that utilising this technology for proactive wheel slip control in anti-lock braking systems outperforms standard rule-based control by approximately 75 % in various longitudinal manoeuvres [34]. However, the benefits of these technologies in more complex manoeuvres, beyond straight braking, remain unknown, particularly for scenarios involving the coupling of longitudinal and lateral dynamics, such as evasive manoeuvres at the limit of handling. These manoeuvres frequently occur in edge cases, where operating conditions are atypical, further complicating the vehicle’s response. Despite substantial research, a significant gap remains between current vehicle automation technologies and fully safe automation, primarily due to these limitations [21]:

- The absence of accurate, real-time information on vehicle states, particularly sideslip angle and tyre forces, across all possible scenarios, particularly edge-cases.
- The lack of proactive control capabilities in AV systems to effectively respond to the nonlinear behaviour of tyre dynamics.
- The inability to stabilise and control the vehicle beyond the stable handling limits or at high sideslip angles, limiting the range of controllable vehicle states and possible trajectories.

Edge-cases, in particular, present unique challenges because they occur outside the nominal operational envelope of automated driving systems. These scenarios often involve unpredictable or rare conditions that push the limits of the vehicle’s control systems and sensors. Addressing edge cases requires innovative control strategies that go beyond routine driving manoeuvres to manage the added complexity of multiple dynamic factors. While significant advancements have been made in automated driving, its full potential has yet to be realised. Therefore, this thesis aims to unlock that potential by developing innovative control strategies specifically designed to address edge-cases in automated driving.

The objective of this thesis is to develop innovative nonlinear motion control algorithms for obstacle avoidance at the limit of handling in automated vehicles and to evaluate the benefits of integrating tyre force sensing bearing technology.

The design of the innovative motion control algorithms is divided into three sub-goals: vehicle sideslip angle estimation to assess stability and controllability, vehicle motion control for obstacle avoidance at the limit of handling with a focus on prioritising safety over path tracking and stability, and automated drifting to stabilise and control the vehicle beyond the stable handling limits or at high sideslip angles.

1.2.1 SUB-GOAL 1: VEHICLE SIDESLIP ANGLE ESTIMATION

Designing vehicle motion controllers for obstacle avoidance at the handling limit requires accurate real-time estimation of key parameters like sideslip angle. While yaw rate can be reliably measured, sideslip angle estimation is challenging, especially under high-excitation driving conditions. This thesis aims to develop filter architectures for real-time sideslip angle estimation with an error of less than one degree [35]. Existing methods for sideslip estimation fall into three categories [36, 37]: model-based, data-driven, and hybrid approaches. Model-based methods rely on nonlinear vehicle models but struggle under transient and near-steady-state conditions. Data-driven approaches leverage historical data and generally provide higher accuracy but require large, high-quality datasets, and performance degrades when the dataset lacks representation in certain conditions. Hybrid methods combine both approaches, achieving better results but often without creating a mutualistic relationship between them. Therefore, the entire hybrid approach potential for sideslip angle estimation has not been exploited yet. Regarding the use of tyre force measurements, they enhance the accuracy of all these approaches, but further development is needed for optimal integration performance improvement assessment.

To address all these gaps, the objective of **Sub-goal 1** is to design a hybrid model-based and data-driven algorithm for vehicle sideslip angle estimation with intelligent (force sensing) bearings integration, which works for a large diversity of vehicle manoeuvres, e.g. steady-state, transient, low, and high excitation.

1.2.2 SUB-GOAL 2: OBSTACLE AVOIDANCE AT THE LIMIT OF HANDLING

The safety of automated vehicles largely depends on their ability to execute effective evasive manoeuvres to avoid obstacles. At the core of any nonlinear motion control algorithm designed for such manoeuvres is the obstacle avoidance controller. Despite significant advancements in automated vehicle safety systems, nonlinearities in tyre behaviour remain a persistent challenge, particularly during manoeuvres at the limit of handling [38]. The most common solution is a hierarchical controller architecture that separates the tasks of motion planning, path tracking, and vehicle stability control [39]. However, during evasive manoeuvres, these objectives can sometimes conflict with one another [26]. For example, prioritising vehicle stability can introduce path tracking errors, which may increase the risk of a collision, as the vehicle may fail to follow its intended trajectory accurately [26]. Resolving these conflicts, especially in the presence of tyre nonlinearities, is critical to improving the safety and performance of automated vehicles during evasive manoeuvres where multiple control objectives must be balanced simultaneously.

To address all these gaps, the objective of **Sub-goal 2** is to design nonlinear model predictive controls which can prioritise obstacle avoidance even at the limit of handling while compensating for inaccuracies in localisation and perception systems, as well as non-nominal road and vehicle conditions.

1.2.3 SUB-GOAL 3: AUTOMATED DRIFTING

The ability to control a vehicle beyond the stable handling limits or at high sideslip angles significantly extends the range of controllable vehicle states and trajectories. This capability is crucial for enhancing safety during extreme collision avoidance manoeuvres, where a

conventionally driven vehicle would have already reached its handling limits and potentially lost control [40–42]. In such scenarios, maintaining a desired path while operating at high sideslip angles poses a considerable challenge. It demands precise control over the nonlinear interactions between longitudinal and lateral tyre forces, which becomes more complex at higher slip angles due to the nonlinearities in tyre dynamics [43]. Furthermore, to achieve stability in these conditions, it is critical to ensure that the rear tyres remain saturated, providing the necessary lateral force in the front axle to maintain the vehicle's trajectory without exceeding its physical limits [44].

For this reason, the objective of **Sub-goal 3 is to design a controller for automated drifting, enabling the vehicle to induce oversteer while maintaining control and driving the vehicle through the entirety of a corner.**

1.3 THESIS OUTLINE AND CONTRIBUTIONS

THIS section outlines the approach adopted in this thesis to achieve the primary objective and sub-goals outlined in Section 1.2. To facilitate ease of navigation, the thesis is structured into seven chapters and one appendix, as shown in Fig. 1.2. Each chapter and the appendix contribute to the overall aim of the thesis, with specific chapters addressing distinct sub-goals. Chapter 2 addresses Sub-goal 1 by proposing an innovative Unscented Kalman Filter-Informed Neural Network for vehicle sideslip angle estimation. Chapters 3, 4 and 5 focus on Sub-goal 2. In particular, Chapter 3 proposes a novel Model Predictive Contouring Control (MPCC) for evasive manoeuvre at the limit of handling capable of prioritising obstacle avoidance over path tracking and vehicle stability in case of emergency. Chapter 4 provides a detailed evaluation and original guidelines for how the proposed MPCC deals with inaccuracies in the localisation and perception stacks, as well as handling non-nominal road and vehicle parameters. Building on this, Chapter 5 proposes an innovative Learning-based Model Predictive Contouring Control (L-MPCC) for evasive manoeuvres at the limit of handling that minimises the model mismatches and uncertainties online through the adoption of a Student-t Process (STP). Chapter 6 is dedicated to Sub-goal 3, presenting a novel Nonlinear Model Predictive Control (NMPC) designed to bring and stabilise a standard passenger vehicle at a high sideslip angle. Finally, Chapter 7 summarises the thesis results, demonstrating how each component contributes to the overarching objectives. It also discusses the limitations of the work and offers recommendations for future research. The Appendix A describes a novel two-stage Bayesian Optimisation, designed to speed up and improve the tuning of the approaches proposed from Chapter 2 to Chapter 5.

A more comprehensive description of the proposed approaches and their contributions to the existing literature is provided below. It is also important to note that each Sub-Goal of this thesis contains an experimental validation using the vehicles shown in Fig. 1.3. For Chapters 4 and 5, which do not include experimental campaign, it should be emphasised that these have been validated in a high-fidelity simulation environment, itself validated with experimental data.

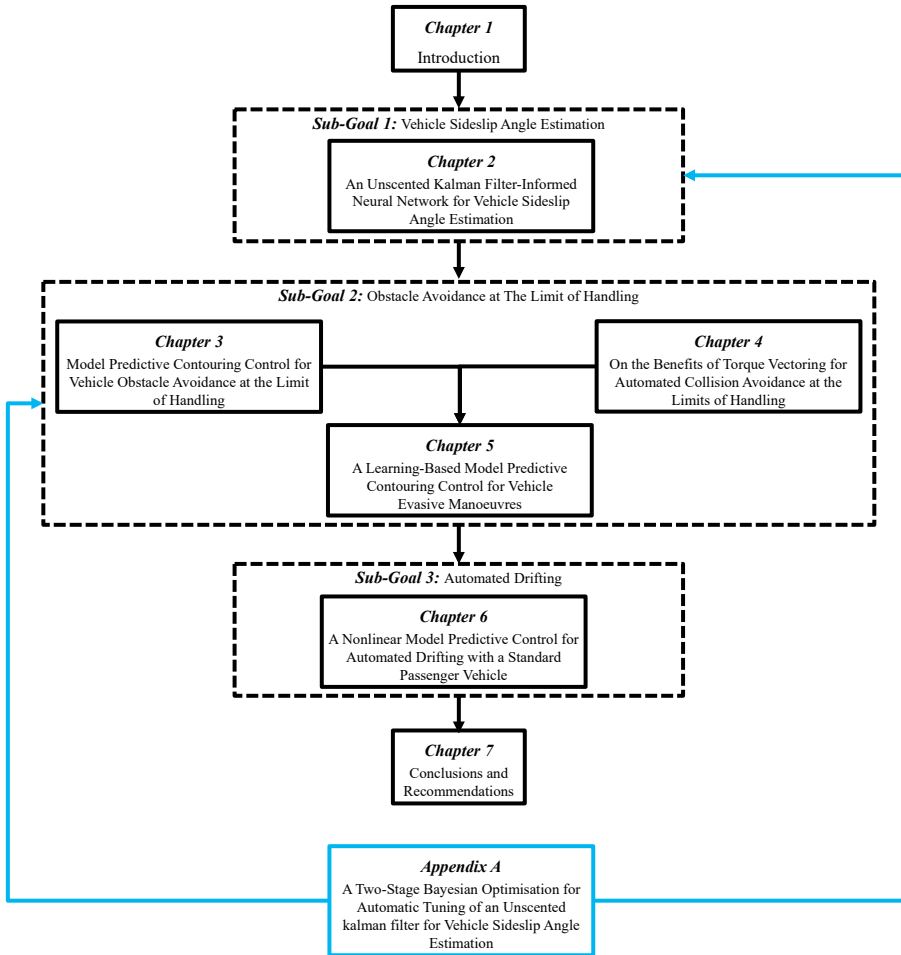


Figure 1.2: Thesis structure.



Figure 1.3: A BMW Series 545i is used at the Automotive Testing Papenburg GmbH to record the vehicle sideslip angle dataset for Sub-goal 1 and Appendix A, see Fig. 1.3a. An experimental TU Delft third-generation Toyota Prius is used to validate the proposed MPCC for obstacle avoidance at the limit of handling, see Fig. 1.3b. Additionally, a BMW M3 Competition is used to experimentally validate the NMPC for automated drifting, see Fig. 1.3c.

1.3.1 CHAPTER 2: AN UNSCENTED KALMAN FILTER-INFORMED NEURAL NETWORK FOR VEHICLE SIDESLIP ANGLE ESTIMATION

Chapter 2 introduces a hybrid method that integrates a mutual relationship between a model-based and a data-driven approach for vehicle sideslip angle estimation. It establishes a mutual relationship between a Convolutional Neural Network (CNN) and an Unscented Kalman Filter (UKF), where the NN's inputs and outputs correspond to the UKF's inputs and measurements. The training process is end-to-end, enabling back-propagation between both the CNN and UKF. The NN assists the UKF by providing pseudo-measurements of the vehicle sideslip angle, particularly in scenarios where the UKF's process model encounters significant mismatches, such as high tyre model nonlinearity. Simultaneously, the UKF imparts physics-based knowledge to the NN, improving its robustness and accuracy, especially in emergency situations that may not be adequately captured in a vehicle sideslip angle estimation dataset.

The contributions of this chapter are threefold. *The first contribution is the mutual hybrid approach in which the CNN is trained end-to-end with the UKF, developing an innovative stochastic, deterministic state estimation method for vehicle sideslip angle estimation. Hence, the proposed hybrid approach enhances the accuracy of the state-of-the-art model-based, data-driven, and hybrid approach for vehicle sideslip angle estimation. The UKF-informed CNN outperforms the state-of-the-art deep ensemble-UKF, achieving a 24.84 % reduction in the nonlinear mean squared error for vehicle sideslip angle estimation on an experimental test set comprising 23 manoeuvres: 2 braking in a turn manoeuvres, 2 skidpads, 5 J-turns, 4 slaloms, 4 lane changes, 2 random steers, 3 spirals, and 1 lap on a track.*

The second contribution of the proposed hybrid architecture is its consideration of the hetero-stochasticity in model dynamics. By estimating the UKF process model uncertainties online through the NN and UKF's end-to-end training, the approach improves estimation accuracy and provides greater robustness, even when trained on a limited dataset. While the improvement compared to model-based methods, tuned with the approach proposed in Appendix A, is modest, the result is significant as it demonstrates the hybrid approach's robustness, showing a 30.04 % lower mean squared error compared to a purely data-driven approach, even with a limited dataset.

The third contribution is the hybrid architecture integrates domain knowledge from the UKF into the neural network, ensuring it adheres to the physical laws governing vehicle dynamics. The UKF-informed CNN achieves a 5.60 % lower average nonlinear maximum error on the test set compared to the hybrid state-of-the-art approaches.

1.3.2 CHAPTER 3: MODEL PREDICTIVE CONTOURING CONTROL FOR VEHICLE OBSTACLE AVOIDANCE AT THE LIMIT OF HANDLING

Chapter 3 presents a Model Predictive Contouring Control approach for obstacle avoidance, utilising a nonlinear single-track vehicle representation with a Fiala tyre model as the prediction component. The MPCC extends previous work on robot motion planning and lap-time optimisation [45, 46] by addressing both obstacle avoidance and vehicle stability. By incorporating lag and contouring errors into the cost function, the MPCC approximates state-of-the-art Frenet-based Model Predictive Control (MPC) while using a Cartesian refer-

ence system. This approach avoids overestimating vehicle-to-obstacle and vehicle-to-edge distances, simplifying the optimisation process. A double lane change manoeuvre with two obstacles is designed to evaluate the controller's performance in avoiding obstacles at the limit of handling.

The contributions of this chapter are threefold. *The first contribution is that the proposed MPCC improves overall obstacle avoidance performance compared to a state-of-the-art MPC [38]. While both controllers can successfully avoid collisions during a double-lane change manoeuvre, the baseline MPC fails to maintain a safe distance from obstacles and road edges. Specifically, the state-of-the-art controller passes just 48 cm from the second obstacle, missing the defined safety distance of 70 cm.*

The second contribution focuses on improved vehicle stability by minimising sideslip angle peaks and increasing the minimum velocity during manoeuvres, thanks to better prioritisation of the obstacle avoidance objective within the MPCC framework. The baseline controller produces a sideslip angle peak of 9 deg and requires counter-steering to return the vehicle to linear behaviour. In contrast, the MPCC with collision avoidance offers a higher stability margin by reducing the sideslip angle peak to 3 deg, eliminating the need for counter-steering. As a result, the vehicle maintains a minimum velocity of approximately 6 km/h, closer to the desired one.

The third contribution is the successful experimental validation of the proposed MPCC, tested on a Toyota Prius performing a double lane change at 30 km/h on a low-friction surface in heavy rain conditions. Additionally, the low road friction in the experimental scenario highlights the robustness of the proposed controller across varying friction conditions.

1.3.3 CHAPTER 4: ON THE BENEFITS OF TORQUE VECTORING FOR AUTOMATED COLLISION AVOIDANCE AT THE LIMITS OF HANDLING

Chapter 4 proposes an MPCC controller based on a nonlinear double-track vehicle model with an extended Fiala tyre model for collision avoidance at the handling limit. The approach, initially introduced in 3, is extended to incorporate torque vectoring capabilities and integrates motion replanning, path tracking, torque vectoring, and vehicle stability tasks, with a focus on prioritising collision avoidance. The MPCC leverages torque vectoring to enhance the vehicle's lateral agility at the limit of handling. The extended Fiala tyre model accurately captures variations in cornering stiffness under combined slip conditions, enabling the MPCC to directly optimise longitudinal tyre forces and yaw moments without requiring a lower-level controller for force allocation. This improves the accuracy of the prediction model and reduces architectural conservativeness while maintaining real-time feasibility on a rapid prototyping platform. The performance and robustness of the proposed MPCC are evaluated in a high-fidelity simulation environment validated with experimental data.

The contributions of this chapter are threefold. *The first contribution is the development of the first real-time feasible MPCC controller augmented with torque vectoring functionalities that can safely avoid vehicle collisions in a double-lane change manoeuvre at the limit of handling. In contrast, existing state-of-the-art approaches [38, 41] (Chapter 3) would result in a collision under similar conditions. By leveraging the enhanced responsiveness provided*

by torque vectoring, the vehicle is driven away from obstacles while maintaining stability, significantly reducing the sideslip angle peak to 7.5 deg, compared to 16 deg in the baseline controller without torque vectoring, and preventing potential crashes. Notably, the proposed controller also demonstrates similar performance in low-friction conditions, such as during heavy rain.

The second contribution is the development and application of an extended Fiala tyre model that accurately captures variations in cornering stiffness in relation to longitudinal and vertical forces, while adjusting the gradient within the tyre's saturation region without disrupting function continuity. This model significantly improves the prediction accuracy, making it highly suitable for MPC applications in scenarios with strong force coupling, such as when torque vectoring is enabled.

The third contribution is the improved robustness of the proposed MPCC with torque vectoring against vehicle and tyre parameter variations, perception inaccuracies, number of false negatives, and localisation errors compared to baseline controllers. The proposed controller demonstrates robustness, showing collision rates of 31.20 % and 38.50 % for vehicle parameter variations and perception inaccuracies, respectively, whereas the baseline controllers always crash into obstacles. Additionally, the percentage of crashes with simulated values near nominal vehicle parameters and common perception inaccuracies is below 1.5 %. A sensitivity analysis further provides numerical insights into which parameters and perception inaccuracies contribute most to performance degradation in obstacle avoidance MPCC controllers.

1.3.4 CHAPTER 5: A LEARNING-BASED MODEL PREDICTIVE CONTOURING CONTROL FOR VEHICLE EVASIVE MANOEUVRES

Chapter 5 addresses the most influential parameters affecting controller performance, as identified in Chapter 4, by proposing a learning-based Model Predictive Contouring Control approach based on a Student-t Process for obstacle avoidance at the limit of handling. The STP predicts the mismatches between the L-MPCC prediction model and the vehicle's yaw rate and lateral tyre forces, measured by intelligent (force sensing) bearings [29]. Unlike traditional methods that focus on vehicle velocity states [47], this approach directly reduces mismatches in the tyre model, enhancing the controller's pro-activity. Additionally, the proposed controller minimises prediction model uncertainties by utilising the STP posterior covariance of tyre forces to propagate lateral state uncertainties along the prediction horizon. This reduces the vehicle's time spent at unstable operating points or unnecessarily close to the handling limit.

The contributions of this chapter are threefold. *The first contribution is the development of the L-MPCC using a Student-t Process, which directly addresses prediction model mismatches in the tyre model, rather than relying on a Gaussian Process that focuses on vehicle velocity states, as typically used in traditional methods [47]. This results in successfully performing evasive manoeuvres at 8.5 % higher velocity than the current state-of-the-art L-MPCC [47].*

The second contribution pertains to the Student-t Process, which enhances outlier resistance compared to the state-of-the-art Gaussian Process. Additionally, the STP posterior covariance is influenced by the observed measurements, leading to higher variance for

operating points that differ from the training set and lower variance for those more closely aligned with the training data [48].

The third contribution is the improvement in vehicle stability, demonstrated by a 76 % reduction in the sideslip angle peak during an evasive manoeuvre compared to the current state-of-the-art L-MPCC based on Gaussian Process [47]. This improvement is achieved by reducing model mismatches and minimising vehicle lateral state uncertainties, which in turn reduces the unnecessary time spent at the limit of handling.

1.3.5 CHAPTER 6: A NONLINEAR MODEL PREDICTIVE CONTROL FOR AUTOMATED DRIFTING WITH A STANDARD PASSENGER VEHICLE

Chapter 6 introduces a novel architecture for automated drifting, composed of three interconnected components. The first component is the offline computation of equilibrium maps, which define the equilibrium points for all vehicle states, given a specific desired sideslip angle and the radius of the intended path. These maps serve as references for initiating and stabilising the vehicle during high sideslip angle manoeuvres. The second component is the proposed Nonlinear Model Predictive Control, which is based on a nonlinear single-track vehicle model and a simplified Pacejka tyre model. A key modification in the NMPC prediction model is the inclusion of the vehicle's sideslip angle as a state variable, which helps simplifying the cost function formulation. The third component is a path-following controller designed to minimise path errors by dynamically adjusting the equilibrium path curvature [49], ensuring the vehicle stays on the desired trajectory while drifting. This integrated architecture enables robust and precise control during automated drifting, extending the path tracking capabilities introduced in Chapters 3, 4 and 5.

The main contributions of this chapter are threefold. *The first contribution is that the proposed architecture is capable of stabilising the vehicle during automated drifting along a desired path, with a maximum lateral path deviation of 1 m while maintaining a sideslip angle of 30 deg, in a high-fidelity simulation environment validated by experimental data.*

The second contribution is the experimental validation of the offline computed steady-state drifting equilibria at a proving ground. The results highlight also the challenges faced by a human driver in maintaining the vehicle at a true equilibrium point due to various disturbances, such as road surface irregularities, which require continuous adjustments to the steering angle and throttle commands.

The third contribution is the experimental validation of the NMPC for automated drifting using a standard passenger vehicle, distinguishing it from previous research, which has either been limited to simulations [50–52] or required extensive hardware modifications [53, 54]. This validation is performed with both high and low friction conditions, successfully stabilising the vehicle at a sideslip angle of 30 degree.



2

2

AN UNSCENTED KALMAN FILTER-INFORMED NEURAL NETWORK FOR VEHICLE SIDESLIP ANGLE ESTIMATION

The key to artificial intelligence has always been the representation.

Jeff Hawkins

This chapter is based on  **A. Bertipaglia**, M. Alirezaei, R. Happee and B. Shyrokau, “An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation,” in *IEEE Transactions on Vehicular Technology*, vol. 73, no. 9, pp. 12731-12746, 2024 [55]. This chapter builds on  **A. Bertipaglia**, M. Alirezaei, R. Happee, and B. Shyrokau, “Model-based vs data-driven estimation of vehicle sideslip angle and benefits of tyre force measurements,” in *International Symposium on Advanced Vehicle Control*, Kanagawa, Japan, 2022 [56].

ABSTRACT

This paper proposes a novel vehicle sideslip angle estimator, which uses the physical knowledge from an Unscented Kalman Filter (UKF) based on a nonlinear single-track vehicle model to enhance the estimation accuracy of a Convolutional Neural Network (CNN). The model-based and data-driven approaches interact mutually, and both use the standard inertial measurement unit and the tyre forces measured by load sensing technology. CNN benefits from the UKF the capacity to leverage the laws of physics. Concurrently, the UKF uses the CNN outputs as sideslip angle pseudo-measurement and adaptive process noise parameters. The back-propagation through time algorithm is applied end-to-end to the CNN and the UKF to employ the mutualistic property. Using a large-scale experimental dataset of 216 manoeuvres containing a great diversity of vehicle behaviours, we demonstrate a significant improvement in the accuracy of the proposed architecture over the current state-of-art hybrid approach combined with model-based and data-driven techniques. In the case that a limited dataset is provided for the training phase, the proposed hybrid approach still guarantees estimation robustness.

2.1 INTRODUCTION

ACTIVE vehicle control systems rely on the sideslip angle and yaw rate information to ensure stability and controllability [57, 58]. Whereas low-cost gyro sensors measure the yaw rate, the vehicle sideslip angle must be estimated. Its direct measurement is possible via optical speed sensors or real-time kinematic positioning-global navigation satellite system (RTK-GNSS), but they are expensive to be installed in passenger vehicles [59]. Hence, the development of filter architectures is required to estimate the sideslip angle in real-time and with the desired accuracy error below one degree in high excitation driving conditions [35]. Sideslip angle estimation is particularly challenging for the following aspects:

- A large diversity of vehicle manoeuvres, e.g. steady-state, transient, low, and high excitation.
- The highly nonlinear behaviour of tyres leads to a substantial limitation due to tyre model accuracy.
- Data collection requires expensive and high calibration sensitive instruments.
- Numerous external disturbances, e.g. bank angle, road slope, road irregularities, and road friction coefficient.

Several approaches have been proposed for vehicle sideslip angle estimation [36, 37]. They are split into three main groups, i.e. model-based, data-driven and hybrid approaches. The model-based approach relies on the physical knowledge of a vehicle model for state estimation. Open-loop deterministic models are insufficient to provide an accurate estimation, so stochastic closed-loop observers, e.g. extended Kalman filter (EKF), unscented Kalman filter (UKF), and particle filters are currently applied to estimate unmeasurable states. EKF and UKF are the industrial state-of-the-art for vehicle sideslip angle estimation because their accuracy can be guaranteed in a specific operating region, and their properties are easily assessed [60]. However, they both struggle in transient and high excitation driving conditions due to the growing nonlinearities in the vehicle model [56]. Another branch of model-based approaches is based on deterministic observers adapted to deal with stochastic noise, e.g. Luenberger observer, the sliding-mode observer and H-infinity methods. Despite their advantage in computational complexities, their need for a higher model fidelity makes them less utilised in industry [61, 62]. Nevertheless, these model-based approaches require extensive system knowledge. The data-driven approach has higher accuracy than the model-based approach when enough quality data are provided in the training phase [56]. Different neural network (NN) architectures have been proposed, e.g. feed forward neural network (FFNN) and recurrent neural network (RNN). However, they all lack interpretability and generalisation capabilities. Thus, a purely data-driven approach is hardly applicable for safety applications in the automotive domain [60]. The third approach, named hybrid, combines the pros of the model-based and data-driven approaches. It improves the model-based accuracy thanks to the NN outputs and, simultaneously, gives the data-driven approach an interpretability thanks to the vehicle model. In the proposed hybrid architectures [60, 63, 64], the model and the neural network work in a unidirectional way. Thus, the model in the hybrid architecture relies on the NN knowledge without backward communication, reducing the approach's potential.

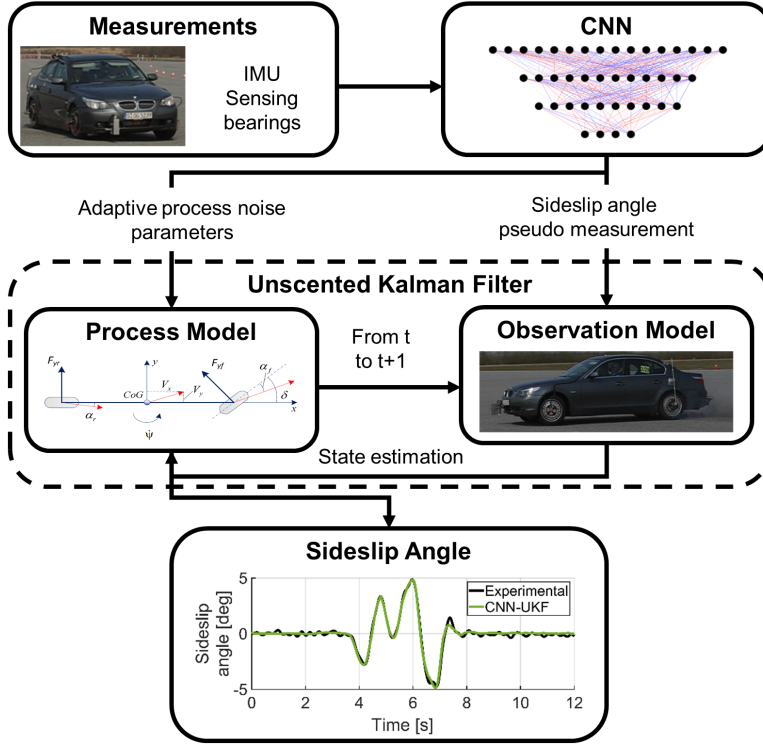


Figure 2.1: Framework overview of the CNN-UKF approach. A CNN provides a sideslip angle pseudo-measurement and the process noise parameters to a UKF based on a single-track vehicle model. The UKF monitors and weights the CNN's estimation through physical laws.

This paper proposes a new hybrid approach for vehicle sideslip angle estimation. Its novelty is in the mutualistic relationship between the model-based approach, characterised by a UKF based on a single-track model, and the data-driven approach, represented by a Convolutional Neural Network (CNN). The proposed approach consists of a sequential hybrid architecture in which the CNN passes the pseudo-measurement of the sideslip angle, the level of distrust of its estimation and the process noise parameters of the vehicle model to the UKF, see Fig. 2.1. A key aspect of the proposed hybrid approach is the training process, which allows the development of a physics-informed NN [65]. Considering that the nonlinear vehicle dynamics are described in a UKF, the physics-informed NN will also be referenced as a UKF-informed NN. The training is end-to-end, so the Back-Propagation Through Time (BPTT) algorithm moves through the CNN, the UKF and backwards. Thus, the CNN is constrained to respect the physical laws of vehicle dynamics. Furthermore, it allows the CNN to estimate variables for which a reference is unavailable, i.e. the process noise parameters and pseudo-measurement level of distrust. This will lead to a high estimation accuracy compared to the state-of-the-art hybrid approaches, which always separate the data-driven component from the model-based one during the training [63, 64, 66–68]. The split proposed in the literature makes the hybrid approach training lighter from a memory

Table 2.1: Overview of the vehicle sideslip angle estimation approaches.

Approach	Features & Authors
Model-based	Kalman filter based on a kinematic model [69–71]
	Kalman filter based on IMU & GNSS [72–78]
	EKF based on a dynamic model [79–85]
	UKF based on a dynamic model [33, 86–90]
	Sliding mode observer [91–93]
	H-infinity observer [94, 95]
	Luenberger observer [96]
	Hybrid - dynamic & kinematic models [59, 97–99]
Data-driven	Online gradient descent [100]
	Modular estimation scheme [101, 102]
	FFNN [103, 104]
	RNN [105–107]
Hybrid	ANFIS [66]
	Kernel-based LPV [108]
	FFNN, ANFIS & UKF [67, 109, 110]
	RNN (GRU) & UKF [68]
	Differentiable EKF [111]
	Kalman filter & FFNN [112]
	Piecewise Affine & Takagy-Sugeno [113]
	FFNN & Kalman in the back-propagation [114]
	KalmanNet [115]
	Kinematic model & RNN (GRU) [60]
	Deep Ensemble Network (LSTM) & EKF [63], UKF [64]

and computational point of view. However, it does not allow the model-based approach to understand when it can rely on the CNN and at the same time, it does not allow the CNN to learn the physical laws of vehicle dynamics. The performance is evaluated using a large-scale, real-world experimental dataset. The dataset contains a great diversity of driving situations, recorded with a constant high friction coefficient.

The paper is organised as follows. Section 2.2 contains a summary of the previous works and the main paper contributions. Section 2.3 describes the CNN and UKF used in the proposed hybrid approach. Section 2.4 describes how the experiments are conducted and evaluated. Section 2.5 reports the obtained results, and Section 2.6 summarises the conclusions and future research paths.

2.2 RELATED WORKS

A summary of the three categories, i.e. *model-based*, *data-driven* and *hybrid*, is presented in Table 2.1.

The first approach is called *model-based* and relies on the laws of physics. The vehicle

behaviour can be described using the geometric constraints, i.e. kinematic model, or considering the forces and moments acting on the vehicle, i.e. dynamic model. The kinematic model requires only geometrical parameters and does not need extensive vehicle parametrisation because its reliability depends mainly on sensing capabilities. The state-of-art kinematic observer [71] is based on a linear parameter varying system, where the states are the vehicle velocities, and the accelerations are the inputs. This approach leads to high accuracy in transient manoeuvres, but the model is not observable in nearly steady-state conditions [59]. Hence to avoid unobservability, a heuristic function is applied to lead the lateral velocity to zero when the vehicle is moving straight or nearly straight [71]. The downside is the amount of data necessary to define the heuristic function. Moreover, despite the performance improvement, it is still susceptible to integration error and sensor drifting. Thus, in recent publications [72–75], the measurements from the Inertial Measurement Unit (IMU) are coupled with those from a Global Navigation Satellite System (GNSS) to increase the amount of information available for the estimator. The velocities measured by the GNSS are integrated into an estimation-prediction framework, which estimates the sideslip angle and partially compensates for the error induced by the low GNSS sampling rate [72]. However, GNSS/IMU fusion kinematic approach still suffers from the low GNSS sampling rate. Furthermore, a high-precision GNSS is too expensive as the standard sensor in passenger vehicles, and signal reception cannot always be assured. Therefore, it is mainly applied to racing [76]. Thus, a solution is to consider dynamic models to rely less on the sensor signal quality. Dynamic models allow a more robust noise computation of the accelerations than kinematic models [59]. However, dynamic models require a more profound knowledge of vehicle parameters and the presence of a tyre model, which is a critical source of uncertainty [33]. EKF and UKF are the state-of-art estimation techniques for the model-based approach, and the process and the observation noises are commonly assumed to be Gaussian and uncorrelated. The EKF uses a first-order Taylor series expansion to linearise around the current mean and covariance. It has excellent accuracy in nearly steady-state conditions and when the vehicle behaves close to linearity, i.e. up to a lateral acceleration of approximately 4 m/s^2 [56]. When the vehicle behaves with strong nonlinearities, UKF assures a better estimation accuracy because it approximates up to the third order of the Taylor series expansion [56]. However, both observers suffer from the mismatches between the physical and modelled tyre behaviour. A possible solution is to combine the pros of dynamic and kinematic models to develop a hybrid kinematic-dynamic observer [97, 98]. This family combines the accuracy in transient manoeuvres of the kinematic models and the better robustness to sensor noise of the dynamic models. The kinematic and the dynamic filters work simultaneously, and the final sideslip angle estimation is a weighted average of the two approaches. The weights are chosen according to the lateral acceleration signal [98]. However, the weighting coefficients' tuning is complex, and the optimum solution varies according to the considered manoeuvres. Another solution to combine dynamic and kinematic models is the development of a modular scheme to estimate in sequential steps tyre forces, longitudinal and lateral velocities [101]. It consists of monitoring the wheel capacities under combined slip at each vehicle corner to estimate the individual forces and velocities. The approach is experimentally validated in different road conditions, but the results do not show its performance when the vehicle is driven at the limit of handling. Thus, the approach's applicability to evasive manoeuvres is limited.

Despite EKF and UKF being the most implemented filters for vehicle sideslip angle estimation [62], also different kinds of observers are proposed in the literature for their advantages in computational complexity and theoretical guarantees of convergence [61], e.g. sliding mode observer [91–93], H-infinity [94, 95], state-dependent Riccati equations [84] and Luenberger observer [96]. Particularly interesting is the combination of an adaptive sliding mode observer to estimate the lateral tyre force with an adaptive compensation algorithm to estimate the sideslip angle [95]. Despite the improved performance compared with EKF in an experimental scenario, there is no comparison with UKF, which is the state-of-the-art in extreme driving conditions. Moreover, these observers fail to perform well in evasive manoeuvres due to increased model mismatches, especially in tyre forces [116].

A solution to enhance the state estimation robustness to tyre model inaccuracies of dynamic model is the introduction of adaptive tyre models [79, 109] or new proprioceptive load-sensing technology, e.g. intelligent bearings or smart tyres [29, 33]. The Kalman filter can use tyre force measurements as an additional feedback to improve the estimation and magnify the Kalman gain, especially in the case of nonlinear vehicle behaviour. The enhanced vehicle safety and the sensor's cost efficiency (lower than 1000 € per vehicle) make the innovative load-sensing technologies candidate to become part of the standard sensor setup for passenger vehicles [33].

A *data-driven* approach reduces extensive requirements of system knowledge compared to the model-based approach. A deep NN with eight hidden layers, each having a different number of long short-term memory (LSTM) cells, is proposed [106]. Despite the increased training time of such a deep NN, the authors state that a smaller network was incapable of reaching the level of accuracy of deeper NN. The issue is that deep NNs are prone to overfit, and their performance strongly lacks generalisation capabilities. To overtake this issue, an NN classification is applied to choose which available NN is most suitable for specific road conditions [105]. Each of the three FFNNs is built with a single hidden layer, and they are trained with three different datasets corresponding to three different road friction conditions, i.e. dry, wet and icy. Moreover, the performance of data-driven approach can be enhanced by the availability of tyre force measurements [56]. In this case, a FFNN with two hidden layers outperforms the accuracy of a more complex RNN architecture based on LSTM cells. A FFNN also exceeds the performance of various model-based approaches, even if it tends to sporadic higher maximum error. However, the data-driven performance is highly dependent on the amount of representative data, and the data-driven approach will lack performance as soon as the dataset contains a lower amount of data in a particular range of the sideslip angle.

Although the data-driven approach generally has a better estimation performance than the model-based approach, it is impossible to guarantee robust performance over vehicle operating conditions. Conversely, a model-based approach based on a dynamic model with tyre force measurements has lower accuracy, but its performance is consistent over the working region [56]. Thus, a *hybrid* model-based and data-driven approach is proposed. We employ two leading typologies, model-to-NN and NN-to-model, as explained below.

The model-to-NN family aims to augment the number of the NN's inputs using the output of a vehicle model. This will transfer some immeasurable physical states to a NN. A kinematic vehicle model can compute the derivative of the sideslip angle, which is used as extra input for the following RNN based on a Gated Recurrent Unit (GRU) cell [60].

The kinematic model provides the NN with a pseudo-measurement that contains significant errors, biases and drift. However, with the extra vehicle model information, the NN reduces the sideslip angle's Mean Squared Error (MSE) of the non-informed NN by 2.7 %, 5.6 %, and 1.2 %, respectively, for dry, wet and snow conditions [60]. The slight improvement shows the benefits of developing a hybrid approach and highlights the importance of providing a more accurate pseudo-measurement.

The NN-to-model family aims to provide a sideslip angle pseudo-measurement to the following EKF/UKF. In this case, the NN output is post-processed by a Kalman filter to improve the sideslip angle estimation. One of the first approaches combines an Adaptive Neuro-Fuzzy Inference System (ANFIS) [66] with a UKF to estimate sideslip angle. The model-based component is employed as a filter to minimise the noise of the NN output and the variance of the estimation mean square error [67]. The ANFIS is trained using synthetic data generated through a high-fidelity simulation environment (CarSim). The ANFIS-UKF improves the performance of only an ANFIS [66] by 21 % on average for five manoeuvres with high friction conditions. However, the presented figures show a maximum value of sideslip angle of only 3 deg in absolute value, which makes the estimation performance easier than in extreme driving conditions. Furthermore, there is no explanation of how the observation noise parameter related to the pseudo-sideslip angle is tuned. This value is essential because it defines the level of distrust the UKF can give to the NN output. A similar approach involves integrating an FFNN with a UKF based on a kinematic vehicle model [110]. Contrary to the ANFIS-UKF approach, the model-based component of the hybrid approach is responsible for filtering the estimation noise and correcting the NN output. This is possible thanks to a proportional feedback correction, which improves the performance of the pseudo-sideslip angle. Although the NN is trained using synthetic measurements, the approach is validated using experimental data and the presented results. Unfortunately, all the results are normalised, so it is impossible to understand if the vehicle was driven at the limit of handling. The presented approach improves the data-driven approach estimation accuracy of 73.3 %. However, there needs to be an analysis of how to decide the distrust level of the pseudo-sideslip angle. Otherwise, when the NN is uncertain due to a lack of data, it will negatively influence the UKF's performance. Furthermore, the kinematic vehicle model is highly susceptible to measurement noise and is not observable in steady-state driving. For this reason, the FFNN is substituted by a deep ensemble NN in more recent publications [63, 64].

Deep Ensemble (DE) of RNN, based on LSTM cells, estimates a sideslip angle pseudo-measurement and its level of distrust, which are then provided as extra measurements to a UKF [63, 64]. The level of distrust is modified through a user-defined linear function before being used by the UKF. This further step is mandatory to scale the NN's distrust level to a meaningful value for the UKF. This hybrid architecture reduces the Root Mean Squared Error (RMSE) by, on average, 8 % vs the RNNs[63]. The extra tuning of the level of distrust can easily lead the approach to overfit. Moreover, the level of distrust is computed through the standard deviation of the sideslip angle pseudo-measurements estimated by the RNNs. This does not lead to a physics-informed NN, so it is still complex to assess the properties of this hybrid approach. The reason is that the DE-RNN is not aware of the performance of the UKF, so the estimated level of distrust is not scaled according to the UKF's accuracy. Vice-versa, a physics-informed NN learns the Kalman filter's precision during the training,

providing the best level of distrust to maximise the hybrid sideslip angle estimation.

This work proposes a hybrid approach employing a mutual relationship between the model-based and data-driven approaches for vehicle sideslip angle estimation. The inputs and outputs of the NN are, respectively, inputs and measurements of the UKF. The end-to-end training enforces the mutual relationship, meaning that the back-propagation algorithm passes through the NN, UKF and vice-versa. The UKF benefits from the CNN when the mismatch in the UKF process model is particularly significant, for instance, due to the high nonlinearity in the tyre model or the non-modelled phenomena. Thus, the CNN provides a vehicle sideslip angle pseudo-measurement to the UKF, which guides it towards an accurate estimation. At the same time, a purely model-based approach, unable to rely on any extra information, would only face a decrease in the estimation accuracy. On the other hand, the CNN gains from the UKF a physics domain knowledge, which helps CNN improve its robustness and accuracy in all the real-life situations that were not fully described by the data used during the training. This is particularly relevant in emergencies that cannot be adequately represented in a vehicle sideslip angle estimation dataset.

The main contribution of this paper is threefold. The first is a mutual hybrid approach in which the CNN is trained end-to-end with the UKF, developing an innovative stochastic, deterministic state estimation method [114] for vehicle sideslip angle estimation. Thus, the UKF observation model has access to CNN's deterministic estimation of a pseudo-measurement of the sideslip angle and its level of distrust. On the other hand, CNN is informed by the following stochastic state estimation filter during the training, so it learns how to estimate the level of distrust independently without requiring extra tuning after its training, as previously proposed in the literature [63, 64]. Hence, the proposed hybrid approach enhances the accuracy (MSE) of the state-of-the-art model-based, data-driven, and hybrid approach for vehicle sideslip angle estimation.

The second contribution is that the proposed hybrid architecture considers the hetero-stochasticity of the model dynamics [114, 117, 118]. Thus, the proposed approach estimates the UKF process model uncertainties online thanks to the CNN and UKF end-to-end training, which helps improve the estimation accuracy. Moreover, it provides higher robustness than the state-of-art, even when a limited dataset is used for the training.

The third contribution is that the proposed hybrid architecture is a UKF-informed NN, which means that the NN incorporates the domain knowledge described in the UKF, and it complies with the vehicle dynamics laws of physics. Thus, the proposed approach has a lower Maximum Error (ME) than a state-of-art model-based [33] and data-driven [56] approach, as well as the state-of-art hybrid approach [63].

2.3 UKF-INFORMED NEURAL NETWORK

THIS section describes the proposed hybrid approach based on a CNN end-to-end trained with a UKF (CNN-UKF). A comparison between the proposed mutualistic hybrid approach and the hybrid unidirectional baseline [63] is represented in Fig. 2.2a and Fig. 2.2b, respectively. The proposed approach develops a UKF-informed NN, where the NN is constrained to respect the vehicle dynamics. At the same time, the baseline (DE-UKF) corresponds to a UKF augmented by the DE outputs. The approach's discretisation is performed through a zero-order hold method [101] due to its good trade-off between simplicity and accuracy. The discretisation works at 100 Hz, the standard frequency for vehicle state estimation.

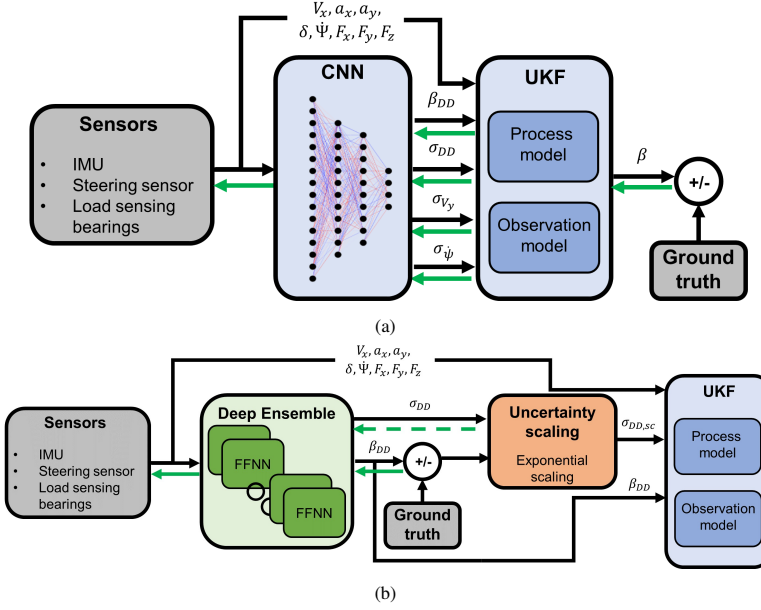


Figure 2.2: Fig. 2.2a shows the proposed hybrid approach architecture (CNN-UKF). Fig. 2.2b shows the baseline hybrid architecture (DE-UKF) proposed by [63]. The black arrows show the flow of information during the online estimation, while the green arrows show the flow of information during the back-propagation. The dashed green arrow represents a term used in the cost function computation but not used by the optimiser to update the NN weights.

2.3.1 DATA-DRIVEN COMPONENT

A straightforward CNN can cope with the complexity of the task because the approach's strength is inside the hybrid architecture. It consists of an input layer, two hidden layers and an output layer.

Seventeen measurements form the input layer (x): longitudinal and lateral accelerations a_x and a_y respectively, longitudinal velocity V_x , road wheel angle δ , yaw rate $\dot{\psi}$, and longitudinal, lateral and vertical tyre forces for each of the four wheels, respectively F_x , F_y and F_z . Before being used, the input measurements are normalised because each input has a different physical meaning and order of magnitude. Thus, all the inputs are mapped onto the interval $[0, 1]$ to speed up and stabilise the training process [119]. A different normalisation method which scales the data to a mean of zero and a standard deviation of one has been tested. Still, the mapping onto the interval $[0, 1]$ produced the best results after the training.

The two hidden layers consist of 200 and 100 neurons and Rectified Linear Unit (ReLU) activation functions. The hidden layers are 2D convolutions with kernel sizes 1×1 , 0 padding, stride equal to 1 and active bias. The CNN uses a dropout regularisation technique equal to 0.2 and a Xavier weight initialisation to avoid overfitting.

The output layer is formed by four neurons corresponding to the pseudo-measurement of the sideslip angle β_{DD} , the level of distrust in the pseudo-measurement σ_{DD} , the uncertainty of the UKF process model lateral velocity σ_{V_y} and the uncertainty of the UKF process model $\sigma_{\dot{\psi}}$. A reference is available only for β_{DD} , but the other three outputs strongly affect the




Figure 2.3: Single-track vehicle model.

estimation of the model-based component, which is used in the training loss function; see Section 2.3.3 for further details. Thus, all four CNN outputs are correctly trained during the end-to-end training. σ_{V_y} , $\sigma_{\dot{\psi}}$ and σ_{DD} are further processed with a sigmoid function to constrain their values inside the meaningful interval $[0, 1]$. This last step assures that the CNN can produce uncertainties which not lead to UKF failure.

2.3.2 MODEL-BASED COMPONENT

A UKF based on a nonlinear single-track model with tyre axle forces computed by the Dugoff tyre is chosen as the model component of this study [56], see Fig. 2.3. The Dugoff tyre is adapted to only assume pure lateral slip condition, reducing the computational complexity [56, 80, 83]. The tyre model parameters (p_t), i.e. tyre cornering stiffness, peak friction coefficient and velocity reduction friction coefficient, are optimised offline using experimental skidpad measurements [33, 86, 88]. The implemented optimisation is a genetic algorithm due to its efficiency with a nonlinear and nonconvex cost function. The vehicle's symmetry is exploited to merge the left and right wheels into a single central axle, which emulates the entire vehicle's behaviour. This model considers only the in-plane dynamics, so the lateral weight transfer, roll and pitch dynamics are ignored. The static weight distribution is considered together with the effect of steady-state longitudinal weight transfer concerning the normal forces on the front and rear axle. A UKF is implemented for its superior estimation accuracy when the vehicle behaves strongly nonlinearly. The vehicle states (x_s) are the V_y and the $\dot{\psi}$, while the vehicle inputs (u_v) are the V_x and the δ . The stochastic process model is responsible for predicting the next time steps of the states according to the following equation:

$$\dot{x}_s(t) = f(x_s(t), u_v(t), p_t) + \omega(t) \quad (2.1)$$

where $f(x_s(t), u_v(t), p_t)$ is the nonlinear single track vehicle model, eq. 2.2, and ω is the vector containing the process noise parameters $[\sigma_{V_y}, \sigma_{\dot{\psi}}]$.

$$\begin{aligned} f(x_s, u_v, p_t) = & \begin{cases} \dot{V}_y = \frac{1}{m} (F_{yf}(x_s, u_v, p_t) \cos(\delta) + F_{yr}(x_s, u_v, p_t)) - V_x \dot{\psi} \\ \dot{\psi} = \frac{1}{I_{zz}} (l_f F_{yf}(x_s, u_v, p_t) \cos(\delta) - l_r F_{yr}(x_s, u_v, p_t)) \end{cases} \end{aligned} \quad (2.2)$$

where m (1970 kg) is the vehicle mass, I_{zz} (3498 kgm²) is the vehicle moment of inertia about the vertical axis, l_f (1.47 m) and l_r (1.41 m) are, respectively, the distance of front and rear axles from the vehicle CoG. F_{yf} and F_{yr} are, respectively, the lateral tyre forces at the front and rear axles. The process noise parameters, σ_{V_y} and $\sigma_{\dot{\psi}}$, are assumed Gaussian and uncorrelated and they capture the uncertainties due to:

- The mismatch between the physical and modelled vehicle behaviour.
- The discretisation error.
- The various operational environments in which the sensors operate.

The filter performance is strongly connected with the process noise parameters, so these are initially tuned using a two-stage Bayesian optimisation (TSBO) [120], see Appendix A. During the estimation, they are computed online by CNN. This is only possible thanks to the mutualistic relationship between CNN and the UKF.

The observation model is responsible for comparing the process model predictions with the available measurements, according to the following equation.

$$y_m(t) = g(x_s(t), u_v(t), p_t) + v(t) \quad (2.3)$$

where $g(x_s(t), u_v(t), p_t)$ is the measurement vehicle model, eq. 2.4, and v is the vector containing the observation noise parameters $[\sigma_{a_{yme}}, \sigma_{\dot{\psi}_{me}}, \sigma_{F_{yyme}}, \sigma_{F_{yyme}}, \sigma_{DD}]$.

$$g(x_s, u_v, p_t) = \begin{cases} a_{yme} = \frac{1}{m} (F_{yf}(x_s, u_v, p_t) \cos(\delta) + F_{yr}(x_s, u_v, p_t)) \\ \dot{\psi}_{me} = \dot{\psi} \\ F_{yyme} = F_{yf}(x_s, u_v, p_t) \\ F_{yyme} = F_{yr}(x_s, u_v, p_t) \\ \beta_{DD} = \text{atan}\left(\frac{V_y}{V_x}\right) \end{cases} \quad (2.4)$$

where a_{yme} , $\dot{\psi}_{me}$, F_{yyme} and F_{yyme} are the vehicle measurements, and β_{DD} is the pseudo-measurement, corresponding to the CNN's output. The observation noise parameters $\sigma_{a_{yme}}$ (0.033 m/s²), $\sigma_{\dot{\psi}_{me}}$ (0.001 rad/s), $\sigma_{F_{yyme}}$ (26 N) and $\sigma_{F_{yyme}}$ (56 N) are the uncertainties of the vehicle measurements and they compensate for the sensor noises. They are tuned by a statistical analysis of the vehicle sensor measurements, which consists of computing the standard deviation of the low-pass measured signal when the steering angle is null and the longitudinal velocity is constant [120], see Appendix A. The variable σ_{DD} is the level of distrust assigned to the pseudo-measurement β_{DD} provided by CNN. The level of distrust computed by the CNN differs from a classic uncertainty measurement because it corresponds to the uncertainty of the pseudo-measurement scaled to match the weight of the noise parameters.

The observability analysis is performed to assess under which conditions it is possible to infer the internal states given the vehicle inputs and measurement. Being the mode highly nonlinear, only the local observability around an operating point can be computed, performing a system linearisation [86]. The observation matrix is built on the Jacobian

matrices of the process and observation models, and it is full rank, equal to two, for all the operating regions in which $\delta \neq k\pi$, $\forall k \in \mathbf{Z}$ and $V_x \neq 0$, where \mathbf{Z} is the integers set. The second condition is always respected because the measurement is considered when V_x is higher than 5 m/s. The steering angle is always inside the range $|\delta| \leq \pi/2$, so the only realistic unobservability happens when $\delta = 0$. However, the vehicle sideslip angle is relevant for lateral dynamics, so it only happens when $\delta \neq 0$.

2.3.3 TRAINING PHASE

The UKF-informed CNN is trained in a supervised way using a labelled dataset. The training is split into two phases: pre-training and end-to-end learning.

PRE-TRAINING

It consists of the back-propagation algorithm applied only to the CNN to speed up and stabilise the following end-to-end training phase. The pre-training loss function is constituted by the sum of σ_{V_y} , $\sigma_{\dot{\psi}}$ MSE losses and by the β_{DD} , σ_{DD} Gaussian negative log-likelihood loss. The MSE loss functions ($MSE_{L, \sigma_{V_y}}$ and $MSE_{L, \sigma_{\dot{\psi}}}$) are represented as:

$$\begin{aligned} MSE_{L, \sigma_{V_y}} &= \frac{1}{N} \sum_{i=1}^N (\hat{\sigma}_{V_y, i} - \sigma_{V_y, i})^2 \\ MSE_{L, \sigma_{\dot{\psi}}} &= \frac{1}{N} \sum_{i=1}^N (\hat{\sigma}_{\dot{\psi}, i} - \sigma_{\dot{\psi}, i})^2 \end{aligned} \quad (2.5)$$

where N is the mini-batch size (256), $\hat{\sigma}_{V_y}$ (0.0007 m/s) and $\hat{\sigma}_{\dot{\psi}}$ (0.002 rad/s) are the initial process model uncertainties tuned by the TSBO for the model-based approach. These losses steer the CNN to predict the process model uncertainties with a meaningful order of magnitude. For σ_{DD} , the Gaussian negative log-likelihood loss function ($NLL_{L, \beta_{DD}}$) represented in the following equation is chosen:

$$NLL_{L, \beta_{DD}}(\beta_{DD}, \sigma_{DD}^2 | \beta_{me}) = \frac{1}{2} \sum_{i=1}^N \left(\log(\max(\sigma_{DD, i}, \epsilon)) + \frac{(\beta_{me, i} - \beta_{DD, i})^2}{\max(\sigma_{DD, i}, \epsilon)} \right) \quad (2.6)$$

where ϵ (10^{-6}) is a constant term for stability, and β_{me} is the sideslip angle ground truth. Thus, this loss function adjusts CNN's weight to maximise the likelihood of the observed data, performing negative log-likelihood minimisation. This process naturally leads to the CNN learning to predict both the mean (β_{DD}) and the variance (σ_{DD}^2). By predicting the variance alongside the mean, the pre-training leads the CNN to provide a point estimate and a measure of confidence or uncertainty associated with each prediction. Thus, it assures a meaningful σ_{DD} , even without its ground truth.

The sideslip angle ground truth is measured through the Corrsys-Datron optical speed sensor installed in the vehicle's front bumper. The sensor reference system is moved to correspond with the vehicle CoG. The measurement is filtered using a zero-phase low-pass filter (bandwidth 5 Hz) because the training phase is sensitive to extreme outliers or noisy references [37]. The cost function is minimised by a mini-batch stochastic gradient descent algorithm based on a standard ADAM optimiser with a learning rate (0.0008). The training procedures' user-defined parameters are optimised through a Bayesian optimisation.

END-TO-END LEARNING

The back-propagation through time (BPTT) algorithm is applied end-to-end to the CNN and UKF. This step creates a mutualistic relationship between the model-based and data-driven approaches. Thus, the UKF is treated as a computation graph unrolled through time, so the CNN-UKF is discriminatively trained over the entire mini-batch length and not on a single step. The procedure to compute the loss function gradient is close to [114], but in the proposed study, a UKF is implemented rather than a linear Kalman filter. The first step is the computation of a loss of a function ($L(\theta)$) that connects the output of the UKF-CNN (β) structure with the available ground truth (β_{me}), given the CNN weights (θ). The training phase minimises the loss function error between the estimated and the measured sideslip angle, allowing to correctly estimate β and all the variables influencing it, so σ_{DD} , σ_{ψ} and σ_{V_y} . The loss function depends on the CNN's weights (θ) and is based on the following equation:

$$L(\theta) = \frac{1}{N} \sum_{i=1}^N (\beta_{me,i} - \beta_{DD,i})^2 + \frac{1}{N} \sum_{i=1}^N (\beta_{me,i} - \beta_i)^2 \quad (2.7)$$

where β_{DD} is the output of the CNN, and β_{me} is the sideslip angle ground truth. The first loss function part, $\frac{1}{N} \sum_{i=1}^N (\beta_{me,i} - \beta_{DD,i})^2$, helps the CNN to estimate the correct pseudo-measurement β_{DD} . The second part of the loss function, $\frac{1}{N} \sum_{i=1}^N (\beta_{me,i} - \beta_i)^2$, is affected by all four CNN outputs and the UKF. The following step to train the proposed UKF-CNN is the computation of the loss function's gradient with respect to the CNN weights ($\nabla_{\theta} L(\theta)$). This is performed following the typical BPTT algorithm. Moving backwards from the loss function, the $\nabla_{\theta} L(\theta)$ is computed by a recursive computation of the loss function gradient with respect to the vehicle states from $t-1$ to t according to:

$$\frac{\partial L}{\partial x_{s,t-1}} = \frac{\partial ukf_{t-1}}{\partial x_{s,t-1}} \frac{\partial L}{\partial ukf_{t-1}} + \frac{\partial x_{s,t}}{\partial x_{s,t-1}} \frac{\partial L}{\partial x_{s,t}} \quad (2.8)$$

where ukf_{t-1} represents all the functions that describe the UKF algorithm, i.e. process model, observation model and Kalman gain computation. The UKF process and observations model, see eq. 2.2 and eq. 2.7 respectively, are nonlinear models with a nondifferentiable point only in the tyre force model. Thus, the Dugoff tyre model, described in eq. 2.9, has been modified to be fully differentiable according to eq. 2.10.

$$\begin{aligned} \mu &= \mu_0 \left(1 - e_r V_x \sqrt{\kappa^2 + \tan(\alpha)^2} \right) \\ \lambda &= \frac{\mu F_z (1 - \kappa)}{2 \sqrt{(C_x \kappa)^2 + (C_y \tan(\alpha))^2}} \end{aligned} \quad (2.9)$$

$$f_{\lambda} = \begin{cases} \lambda(2 - \lambda), & \text{if } \lambda \leq 1 \\ 1 & \end{cases}$$

$$F_y = \frac{C_y \tan(\alpha) f_{\lambda}}{1 - \kappa}$$

$$f_{\lambda} = \frac{2}{e^{-4\lambda} + 1} - 1 \quad (2.10)$$

where μ_0 is the peak friction coefficient, e_r is the friction reduction coefficient, V_x is the longitudinal velocity, κ and α are the longitudinal and lateral slip, C_x is the longitudinal slip stiffness, and C_y is the cornering stiffness. Regarding all the UKF operations, they are differentiable and available in the open-source machine learning platform TensorFlow or PyTorch. The gradient computation continues applying the chain rule to eq. 2.8 and moving backwards, computing the derivative to each CNN weight as for a normal NN. This step is performed automatically by the chosen machine learning platform, PyTorch. The training is based on a mini-batch stochastic gradient descent algorithm (mini-batch size of 256) based on a standard ADAM optimiser with a learning rate equal to 0.0008.

2.4 EXPERIMENT SETUP

THIS section describes how the experiments have been conducted and how the proposed approach has been compared to the baseline methods.

2.4.1 EXPERIMENTAL SETUP AND DATASET

The experiments have been conducted at the Automotive Testing Papenburg GmbH with the test platform based on a BMW Series 545i. The test vehicle was instrumented with the standard IMU, Kistler wheel force transducers and SKF intelligent bearings for each wheel, a dual antenna GNSS from Oxford Technical Solutions and a Corrsys-Datron non-contact optical sensor to measure the sideslip angle (measurement accuracy of $\pm 0.2^\circ$). The high-end optical speed sensor was used to measure the ground truth of the vehicle sideslip angle. The vehicle was equipped with a dSPACE 1007 AutoBox as a real-time control platform. All the equipment was interconnected through the Controller Area Network (CAN) interface, and the sampling rate was set up at 100 Hz. The intelligent bearings demonstrate a similar accuracy to the wheel force transducer [30], the most common sensor technique in research for tyre force measurement. Thus, the tyre forces in the training dataset are taken from the wheel force transducers, making the paper easier to reproduce. The dataset contains 216 manoeuvres corresponding to two hours of driving and consists of standard vehicle dynamics manoeuvres, e.g. double-lane change, slalom, random steer, J-turn, spiral, braking in the turn, and steady-state circular tests, together with recorded laps at the handling track. All manoeuvres were driven on dry asphalt with tyres inflated according to the manufacturer's specifications. The bank angle and the road slope were negligible, and the friction coefficient was approximately constant. Two different electronic stability control settings (On, Off) were used. All the measurements were recorded at 100 Hz, the standard frequency for vehicle state estimation. A statistical outlier removal has been applied to remove extreme outliers. However, particular attention is paid to not deleting edge case measurements, which are the most valuable data. Furthermore, all the manoeuvres were manually inspected to check the outlier removal efficacy. The measurements are considered when V_x is higher than 5 m/s and are filtered using a low-pass zero-phase filter with a cut-off frequency of 5 Hz based on a finite impulse response technique [37].

The log distribution of the sideslip angle and lateral acceleration is represented in Fig. 2.4. The lateral acceleration is almost spread equally in the range $[-10, 10]$ m/s². In contrast, the sideslip angle measurements mainly distribute in the range $[-3, 3]$ deg. The latter is a common phenomenon because it is challenging to perform manoeuvres with a high sideslip




Figure 2.4: Log distribution of sideslip angle and lateral acceleration. Each bin corresponds to 1 deg and 1 m/s²

angle, even when the vehicle has a very high lateral acceleration. Especially in dry road conditions, only a professional driver can induce a high sideslip angle.

A second dataset is selected from the same measurements. It will be referenced as limited dataset because it only contains measurements of when the vehicle has lateral acceleration $|a_y| \leq 7 \text{ m/s}^2$. This simulates the cost and complexity of recording a large number of manoeuvres in which the vehicle is driven at the extreme vehicle behaviour, but not at the handling limits. Such a situation is common in the automotive field because, at the handling limits, the driver can easily lose the vehicle's control. Thus, the limited dataset will be used to analyse the proposed hybrid approach regarding its robustness and generalisation capabilities.

Both datasets are split into three sub-sets: training (75 %), validation (15 %) and test (10 %). The test set contains the same manoeuvres for both the full and limited datasets. It consists of manoeuvres representing the entire driving behaviour, but more focus is paid to highly nonlinear situations. It includes 23 manoeuvres: two braking in the turn, two skidpad, five J-turn, four slalom, four lane change, two random steers, three spiral and one lap track.

2.4.2 KEY PERFORMANCE INDICATORS

The performance of the different approaches is assessed through four key performance indicators (KPIs), which are commonly used in sideslip angle estimation [33, 56, 86].

- The MSE assesses the overall estimation performance.
- The nonlinear MSE (MSE_{nl}) corresponds to the MSE computed only when $|a_y| \geq 4 \text{ m/s}^2$. It measures the estimation performance when the vehicle behaves nonlinearly.
- The absolute maximum error (ME) measures the worst estimation performance.
- The nonlinear ME (ME_{nl}) measures the worst estimation performance in the case of nonlinear vehicle behaviour.

The nonlinear KPIs analyse the hybrid approach performance in the most critical scenarios. The MSE and MSE_{nl} are used to evaluate the estimation accuracy, while ME and ME_{nl} are used to assess temporary high errors in the estimation. The latter is relevant to assess whether the estimation is always coherent with the physical vehicle behaviour.

2.4.3 BASELINE METHODS

The proposed hybrid approach is compared with the state-of-art model-based, data-driven and hybrid approaches. All the considered baselines are adapted and optimised to use the same sensor setup and dataset, ensuring an objective and fair comparison.

The model-based approach is a UKF-based on a single-track model with tyre force measurements, as presented in [56]. The process noise parameters are tuned with the TSBO, and the observation noise parameters associated with the tyre force measurements are adapted online to enhance the observer's performance. The adaptability is related to the reduction of the level of noise coupled with tyre force measurements, this increases the Kalman gain when the vehicle behaves nonlinearly. Thus, the effect of the Kalman gain is magnified during manoeuvres at the handling limit. Otherwise, a magnified Kalman gain when the vehicle behaves linearly could influence the vehicle states to follow the measurement sensor noises. The adaptability is triggered with a hysteresis loop to avoid the chattering phenomenon.

The data-driven approach is a FFNN that uses IMU and tyre force measurements as inputs, as evaluated in [56]. A simple FFNN reaches a better performance than a RNN when the tyre force measurements are included in the input set because the RNN prediction power is insufficient to compensate for the higher numbers of parameters to be trained. The NN is formed by two hidden layers with respectively 250 and 125 neurons each and ReLU activation functions. It uses a dropout regularisation technique (0.2) and a Xavier initialisation to avoid overfitting. An early stopping method with patience equal to 20 is applied for the same reason. The MSE is the loss function minimised by a mini-batch stochastic gradient descent algorithm based on a standard ADAM optimiser with a learning rate (0.001). The mini-batch size is 1024. For the training procedures, user-defined parameters are optimised through a Bayesian optimisation.

The hybrid approach is a deep ensemble-UKF (DE-UKF) [63] adapted to maximise the estimation performance on a dataset with tyre force measurements. The DE is formed by 20 FFNNs trained independently on the same dataset. The FFNNs different estimations are combined in a model averaging. Hence, the final β_{DD} is the mean of the FFNNs estimations, and σ_{DD} is the variance of the different model estimations. Each FFNN is trained using a Gaussian negative log-likelihood cost function optimised through mini-batch stochastic gradient descent based on an ADAM optimiser with a learning rate (0.0008). The epoch's number for each FFNN is 30. DE relies on the stochasticity of neural network training, which allows every FFNN to converge to a different set of parameters. However, the estimation accuracy is low when all models predict incorrectly, and there is no guarantee that the σ_{DD} will be high. This especially happens when the error is in the low sideslip angle range because the NNs estimations tend to be closer. A high level of distrust suggests that the UKF does not rely on the data-driven pseudo-measurement but trusts the estimation of the UKF process model. Vice-versa, when the level of distrust is low, the UKF considers the neural network estimation reliable. σ_{DD} must be scaled before being used by the UKF because the output of the DE does not match the weight of the other noise parameters. Otherwise, the UKF puts too much trust in β_{DD} . The scaling is based on an exponential function (eq. 2.11) which differentiates approximately similar σ_{DD} .

$$\sigma_{DD,sc} = 10^{p_1} \sigma_{DD}^{p_2} \quad (2.11)$$

where p_1 (-4.2690 for the full dataset and -1.4353 for the limited one) and p_2 (0.7901 for the

full dataset and 1.465 for the limited one) are two scaling parameters tuned using a Bayesian Optimisation. The values of p_1 and p_2 change according to the dataset because they strongly influence the DE's estimation performance. If p_1 and p_2 are not re-tuned for the limited dataset, the UKF will put too much trust in the DE, even if it lacks performance.

2

2.5 RESULTS

THIS section demonstrates the performance of the proposed approach. Subsection 2.5.1 analyses how accurate the proposed approach is with respect to the baselines when it is trained using a full dataset. Subsection 2.5.2 shows the results of the robustness analysis when only a limited dataset is available. This demonstrates that the data-driven approach is highly influenced by the amount and quality of the data. Subsection 2.5.3 proves the approaches' robustness to different tyre model parameters.

2.5.1 FULL DATASET RESULTS

The CNN-UKF, the DE-UKF and the data-driven approach have been trained using the full dataset.

The overall comparison is presented in Table 2.2. Both hybrid approaches perform better than the model-based and data-driven approaches considering all four KPIs. This highlights the importance of the hybrid architecture for vehicle sideslip angle estimation. For instance, the model-based approach has a higher MSE and MSE_{nl} than the data-driven approach but a lower average ME and ME_{nl} . The hybrid approaches have the same estimation accuracy (MSE and MSE_{nl}) as the data-driven approach without the average higher ME. The reason is that in a hybrid approach, data-driven estimation is always validated through the model-based approach.

It can be seen that CNN-UKF outperforms the three other approaches for all the proposed KPIs. However, it does not have the same benefits in magnitude for all of them. The overall MSE and ME of DE-UKF and CNN-UKF are comparable. The minor improvements for the linear vehicle behaviour are respectively 1.15 % and 0.20 % in favour of the CNN-UKF. Anyhow, if the performance is evaluated when the vehicle behaves nonlinearly, the CNN-UKF will strongly outperform DE-UKF with an improvement of 24.84 % for the MSE_{nl} and 5.60 % for the ME_{nl} . A possible explanation is that the end-to-end training informs CNN about the vehicle dynamics compensating for the lower amount of data in this operating condition.

Table 2.2: Sideslip angle estimation comparison using the full dataset.

Approaches	MSE [deg ²]	MSE _{nl} [deg ²]	ME [deg]	ME _{nl} [deg]
Model-based	0.161	0.277	1.111	0.991
Data-driven	0.096	0.157	1.293	1.123
DE-UKF	0.087	0.157	0.981	0.822
CNN-UKF	0.086	0.118	0.979	0.776

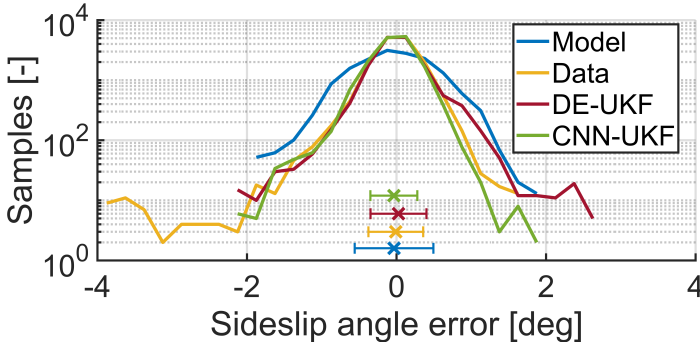


Figure 2.5: Distribution of the sideslip angle error when the vehicle $|a_y| > 4 \text{ m/s}^2$ for every approach in the test set. Each bin is 0.25 deg wide. The x represents the mean and the line between the vertical symbols (| - |) is the standard deviation of the sideslip angle error.

On the contrary, the DE during the training is not aware of the physical vehicle behaviour, so it is subjected to a decay in performance where the dataset has fewer samples. The DE becomes aware of the UKF performance only while tuning the level of distrust scaling parameters. Furthermore, the process model noise parameters are online adapted in the CNN-UKF, allowing the UKF to accommodate better the mismatches between the physical and modelled vehicle behaviour.

Similar conclusions can be stated from the log distribution of the sideslip angle error in the nonlinear operating range, see Fig. 2.5. The data-driven and the hybrid approaches have a similar amount of β error samples in the range $[-1.5, 1.5]$ deg. In contrast, the model-based approach suffers from the lower accuracy of the vehicle model in the nonlinear operating region. However, the data-driven approach and partially the DE-UKF are more prone to high estimation errors (≥ 1.5 deg) than the model-based and CNN-UKF. The latter outperforms all other approaches and has the β error mean closest to zero and the lowest standard deviation. Hence, a UKF coupled with a data-driven approach has the same performance as a data-driven approach in a low error range, but it reduces the sporadic high errors of a purely data-driven approach. Furthermore, the end-to-end training and the process noise parameters adaption allow the CNN-UKF to maximise the hybrid capability especially when the vehicle $|a_y| > 4 \text{ m/s}^2$.

Fig. 2.6 analyses how the estimation performance change for different manoeuvres. The model-based approach has a weak accuracy, especially in braking-in-the-turn, J-turn and skidpad tests. The braking-in-the-turn involves a coupling between the longitudinal and lateral dynamics, which is not modelled in the used single-track vehicle model. In a J-turn manoeuvre, the vehicle is driven at the limits of handling, where the mismatches between the physical and modelled vehicle behaviour are higher. Whereas for skidpad tests, the explanation is that it is a quasi steady-state manoeuvre, so the vehicle yaw acceleration is almost null, and the difference between estimated and measured tyre forces becomes essential for the β estimation. The tyre model is one of the most significant uncertainty sources in the model-based approach. The data-driven approach almost constantly behaves better than the model-based but worse than the hybrid approaches for estimation accuracy. However,

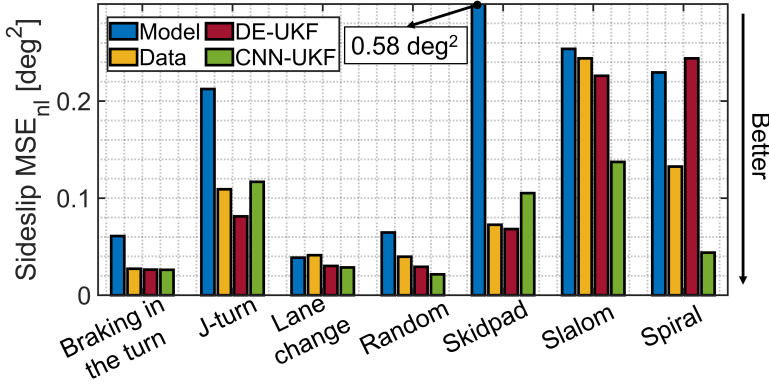


Figure 2.6: Sideslip angle MSE_n comparison for every group of manoeuvres.

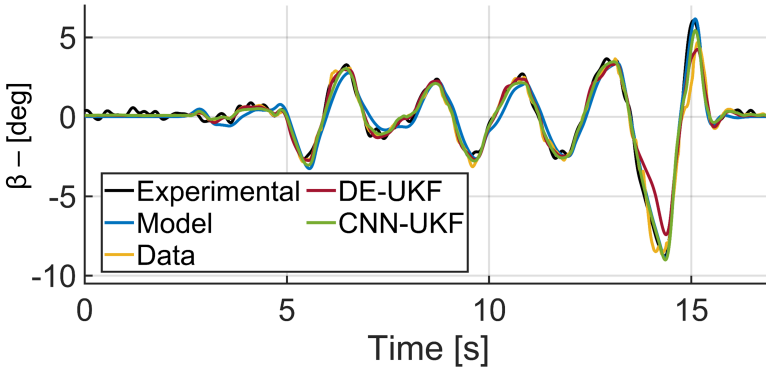


Figure 2.7: Slalom manoeuvre. Comparison of the sideslip angle estimation between all four approaches.

it outperforms the DE-UKF in a spiral manoeuvre, and a possible explanation is that the DE-UKF puts too much trust in the UKF process model. The CNN-UKF outperforms all the other approaches in five out of seven manoeuvres. Particularly relevant is the improvement in the slalom and spiral manoeuvres. The slalom has the highest number of sideslip angle peaks (Fig. 2.7), which are the most difficult moments to estimate sideslip. Spiral manoeuvres are particularly challenging because it has an extra turn respect the J-turn. Fig. 2.7 shows the sideslip angle estimation in a slalom manoeuvre at the handling limits. All four approaches provide a reliable estimation, but the CNN-UKF outperforms the other approaches when the vehicle reaches a β peak of 10 deg at around 14 s. This is a typical situation where a correct estimation of β is essential to help the vehicle control system maintain vehicle stability. Thus, an improved estimation in this condition is particularly relevant for safety. The already mentioned high nonlinearities reduce the accuracy of the model-based approach. The data-driven approach lacks accuracy at 15 s due to the few data in the training set describing this vehicle's operation point. The DE-UKF improves the estimation performance between 5 s and 13 s combining the pros of the model-based and data-driven approach, but it lacks performance at around 14 s. CNN-UKF improves the estimation accuracy not only in

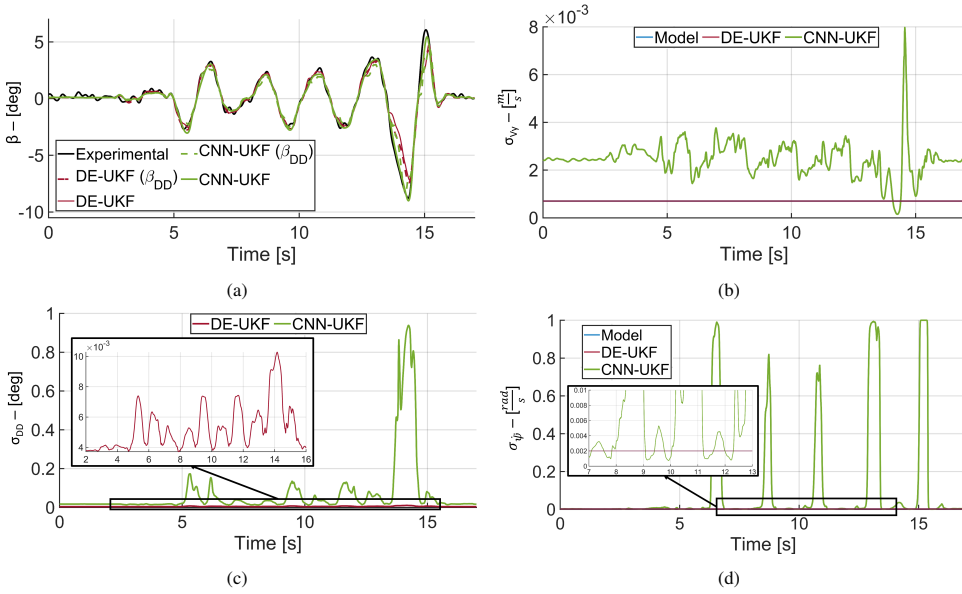


Figure 2.8: Fig. 2.8a shows the estimated and the pseudo-measurement of the sideslip angle. Fig. 2.8b shows the process noise parameter associated with the V_y . Fig. 2.8c shows the level of distrust in the NN for the hybrid approaches. Fig. 2.8d shows the process noise parameter associated with the ψ .

the range of [5, 10] s but also in the highest peak at 14 s, as can be observed in Fig. 2.8.

Fig. 2.8a shows the β and β_{DD} for the hybrid approaches. CNN-UKF and DE-UKF β_{DD} s lack accuracy between 12 s and 15 s, but the CNN-UKF β is accurate because the UKF is correctly weighting the UKF process model's information with the NN's pseudo-measurement. Vice-versa, the UKF of the DE-UKF puts too much trust in β_{DD} . When the β_{DD} error rises, the corresponding level of distrust (Fig. 2.8c) also grows. CNN-UKF and DE-UKF σ_{DD} s have the same order of magnitude in normal driving, but the one related to CNN-UKF rises much more than the DE-UKF. This broader range makes the proposed approach much less confident in the NN when its output is incorrect. This is not possible for the DE-UKF due to its training process. The DE-UKF does not have end-to-end training, so its σ_{DD} cannot match the weight of the other UKF noise parameters. The DE-UKF σ_{DD} nonlinear scaling compensates only partially this issue. Fig. 2.8c clearly demonstrates how the CNN-UKF distrust level range is $[10^{-3}, 1]$, while the range for the DE-UKF is only $[10^{-3}, 10^{-2}]$.

Another explanation for the better performance of the CNN-UKF is related to the online adaptation of the process noise parameters. The adaptive parameters allow the UKF to know the current mismatches between the modelled and physical vehicle behaviour. The process noise parameters of the DE-UKF and model-based approach are constant, so they correspond to a trade-off between the different driving conditions. Vice-versa, the CNN-UKF relies on optimal tuned process noise parameters every instant. Fig. 2.8b and 2.8d show the values of σ_{V_y} and σ_{ψ} , respectively. As expected from the literature [121], both increase with the

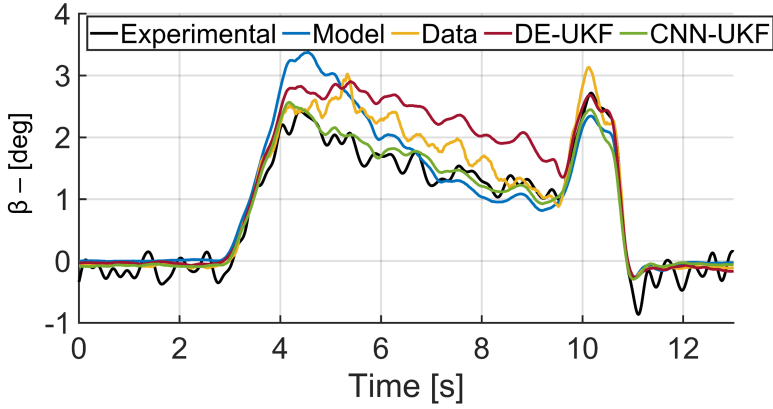


Figure 2.9: Spiral manoeuvre. Comparison of the sideslip angle estimation between all four approaches.

growth of vehicle nonlinearities. This further proves that CNN-UKF behaves according to physical vehicle motion. σ_{V_y} has a peak at 14 s, corresponding to the last vehicle's right turn, where the rear inner tyre is even detached from the ground due to the aggressiveness of the manoeuvre. This extreme condition is created by a transient lateral load transfer (not modelled) which strongly influences lateral tyre force production, resulting in a significant V_y model mismatch. Moreover, the effect of the front axle longitudinal force (F_{xf}) on the lateral velocity is not modelled ($\frac{F_{xf} \sin(\delta)}{m}$). Overall, the constant and the adapted process noise parameter have the same magnitude. Still, the one associated with CNN-UKF is generally bigger (apart from 13 s to 14 s). The reason is that the constant σ_{V_y} was optimised, considering also less aggressive manoeuvres where the vehicle model is more reliable.

The process noise parameter σ_{ψ} rises by two orders of magnitude when the vehicle has a high sideslip angle. At the same time, when β is low, the adapted σ_{ψ} is slightly lower than the constant process noise parameter. A possible explanation is that the mismatches of the modelled ψ are higher than that of V_y . The meaningful adaptability of the process parameter noises shows the CNN-UKF has an insight into vehicle dynamics physics and it can online compensate for it.

Similar conclusions are obtained from the spiral manoeuvre represented in Fig. 2.9. Here, the CNN-UKF approach outperforms the accuracy of all other three approaches, particularly from 5 s to 13 s. The performance of the CNN-UKF is similar to sum of the best estimation between the data-driven approach, from 4 s to 6 s, and the model-based approach, from 6 s to 10 s.

The test set also contains a recording of an entire lap in a racing circuit, where the effect of combined slip is maximal. Fig. 2.10b shows the vehicle's lateral and longitudinal acceleration, and it highlights how the driver is pushing the vehicle at the limit of handling in all the corners, see [1, 7] s, [16, 19] s and [23, 29] s. The sideslip angle estimation performance of the four approaches is represented in Fig. 2.10a. The model-based approach has the lowest performance, especially in the range [1, 7] s. This result is expected because the implemented Dugoff tyre model works in pure slip conditions. A similar conclusion is also visible in the spiral manoeuvre, Fig. 2.9. The data-driven approach performs better

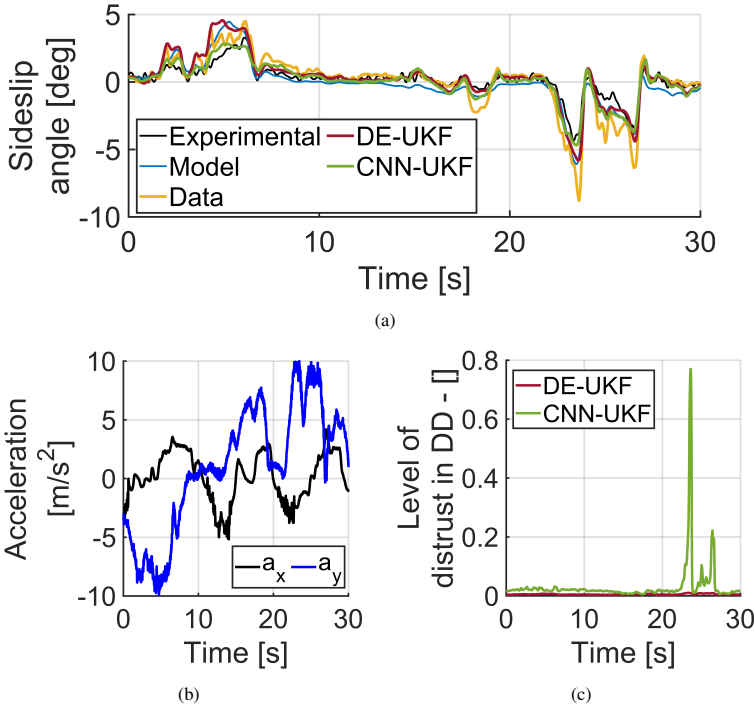


Figure 2.10: Fig. 2.10a compares the sideslip angle estimation between four approaches in a portion of a racing track. Fig. 2.10b shows the recorded lateral and longitudinal acceleration of the vehicle. It highlights the combined slip situation at which the vehicle is driven. Fig. 2.10c shows the level of distrust in the NN for the hybrid approaches.

than a model-based approach. However, it has the maximum absolute error at 23 s and 26 s, where the vehicle performs a cornering while braking. Both proposed hybrid approaches have higher accuracy than the others, but the CNN-UKF has the best performance. A possible explanation is the physic-informed NN architecture, which allows evaluating a very accurate NN level of distrust. Fig. 2.10c shows the NN level of distrust for the DE-UKF and CNN-UKF. While the DE-UKF level of distrust is almost constant along the manoeuvre, the one associated with CNN-UKF has two peaks in correspondence with the data-driven maximum errors. This allows the CNN-UKF to avoid following the high estimation error of the data-driven component. It is further proof of how the CNN-UKF is a physics-informed NN in which the UKF and NN are mutually cooperating to improve the overall estimation of the hybrid approach.

2.5.2 ROBUSTNESS ANALYSIS USING THE LIMITED DATASET

A sideslip angle filter must not only be accurate, but it should be robust to a different amount of qualitative data during the training and tuning phase. Hence, to prove the robustness of the proposed hybrid approach, its estimation performance is compared with the baseline methods when they all have been trained using the limited dataset.

Table 2.3: Sideslip angle estimation comparison using the limited dataset.

Approaches	MSE [deg ²]	MSE _{nl} [deg ²]	ME [deg]	ME _{nl} [deg]
Model-based	0.161	0.277	1.111	0.991
Data-driven	0.223	0.358	1.445	1.284
DE-UKF	0.157	0.270	1.103	0.983
CNN-UKF	0.156	0.269	1.099	0.975

The overall comparison is presented in Table 2.3. Here, the model-based approach shows the same performance as with the full dataset (see Table 2.2), because it is not influenced by the amount of data. As expected the other approaches show a reduced performance with the limited dataset where the MSE is more than doubled while the ME sees a moderate increase. Now the data-driven approach has the worst performance in all four KPIs. The accuracy loss is higher than 30 % for all the indicators, without a particular weakness in one of the proposed KPIs. An explanation is that the dataset does not have representative data of the vehicle driven with $|a_y| \geq 7 \text{ m/s}^2$, so it must generalise much more than with a full dataset. Significantly, the NN must reconstruct the extreme nonlinear vehicle behaviour, the most complex vehicle operating region, without having representative data for these conditions.

The model-based and hybrid approaches' performance is very similar, but DE-UKF and CNN-UKF have the best KPIs. The explanation is that hybrid approaches use the best estimation accuracy of the model and the NN together. Both hybrid approaches strongly rely on the estimation of the UKF process model because they cannot put much trust in the data-driven part. However, the NN still has benefits when the vehicle behaves linearly due to the excellent amount of data in that range. This highlights how the hybrid approaches improve the robustness of both model-based and data-driven approaches. The hybrid approach shows a minor improvement compared with the model-based approach. However, the result is significant because it highlights how the hybrid approach is as robust as a model, even if trained with a limited training dataset. On the contrary, a purely data-driven approach is not robust for using a small training set resulting in poor estimation accuracy.

The CNN-UKF performs slightly better than DE-UKF in all four KPIs. However, the CNN-UKF outperforms the DE-UKF, mainly for the MSE and MSE_{nl}. The main reason is the adaptability of the process noise parameters, which cope with the change of vehicle model mismatches in the various vehicle operating points. However, the improvement in accuracy is not enough to be considered significant ($< 5 \%$).

Fig. 2.11 shows the sideslip angle error log distribution in the nonlinear operating range. All the approaches which rely on a model highly outperform the data-driven approach. The latter have β error samples in the range $[-3.8, 4] \text{ deg}$, while the other approaches have β errors between $[-1.8, 1.8] \text{ deg}$. This proves that the data-driven approach is highly prone to high estimation errors when trained with a limited dataset. The performance of the model-based and hybrid approaches is very similar. They also share an equal error distribution. The data-driven approach slightly outperforms the other approaches in the very low error range $[-0.3, 0.6] \text{ deg}$. This explains why the hybrid approaches are more accurate overall than the




Figure 2.11: Distribution of the sideslip angle error when the vehicle $|a_y| > 4 \text{ m/s}^2$ for every approach in the test set. Each bin is 0.25 deg wide. The x represents the mean and the line between the vertical symbols (| – |) is the standard deviation of the sideslip angle error. Results based on the limited dataset, see Fig. 2.12 the for best results.




Figure 2.12: J-turn manoeuvre. Comparison of the sideslip angle estimation between all four approaches using the limited dataset.

model-based one, despite mainly relying on it.

Fig. 2.12 shows the sideslip angle estimation in a J-turn manoeuvre at the handling limits. The model-based and hybrid approaches behave almost identically, and all strongly outperform the purely data-driven approach. The only visible differences are between [1.5, 3] s, where the CNN-UKF captures slightly better the conclusion of the peak and between [4, 7] s where the DE-UKF is closer to the β reference.

The major difference between DE-UKF and CNN-UKF is visible from the comparison of the β_{DD} , see Fig. 2.13. The β_{DD} computed by the CNN-UKF is highly outperforming the one estimated by the DE-UKF. The explanation is that the CNN-UKF is trained end-to-end, so the output of the CNN has physical information that the NN uses to increase its accuracy. The DE is trained independently, performing similarly to the purely data-driven approach. An higher accuracy of β_{DD} implies that the following UKF can rely on a better sideslip angle pseudo-measurement. This proves the benefits of using a physical informed-NN. Despite this, the β estimation of DE-UKF and CNN-UKF is similar because the model-based approach still outperforms both β_{DD} . Thus, both hybrid approaches mainly rely on the UKF. Due to the high chances of dealing with a limited dataset, the hybrid approach is fundamental to improving vehicle sideslip angle estimation.

However, the performance of the proposed CNN-UKF approach is still influenced by the amount and quality of data in the training set. Thus, it still represents a limitation of




Figure 2.13: J-turn manoeuvre. Comparison of the estimated β and β_{DD} between the hybrid approaches using the limited dataset.

the proposed approach that must be addressed in the future. This highlights the importance of defining standards procedure to collect valuable and broad datasets. Regardless, the proposed CNN-UKF allows the introduction of possible solutions for lack of data, e.g., weakly-supervised learning during the end-to-end training, which allows for using data recorded without expensive sensors.

2.5.3 ROBUSTNESS ANALYSIS TO TYRE MODEL PARAMETERS

An essential property of vehicle sideslip angle estimators is the robustness to vehicle parameter variations. The variation of cornering stiffness strongly influences the sideslip angle estimation, because it can vary from its nominal value due to numerous factors, e.g. tyre pressure, temperature, and wear [33]. Thus, to prove the robustness of the proposed approach, Fig. 2.14 shows how the MSE_{nl} (2.14b) and the MSE (Fig. 2.14a) of the vehicle sideslip angle vary at the variation of the axle cornering stiffness. In particular, the front and rear axle cornering stiffnesses are changed by $\pm 10\%$ [33]. Overall, it is visible that the model-based approach is the most sensible to the variation of the inner model, while the data-driven approach, which does not have a physical model, is not affected. Both hybrid approaches are less influenced by the variation in the physical model than the purely model-based approach. However, Fig. 2.14 shows that the CNN-UKF performance is more influenced by the model mismatch than the DE-UKF. A possible explanation is that the CNN-UKF varies the level of distrust in the NN during the manoeuvre according to how it learned during the training, giving more trust to the UKF at some specific moments. When it faces a different model mismatch than previously learned, it is more influenced by it than the DE-UKF, which does not vary the level of distrust to the NN, always prioritising the data-driven side of the approach. Despite this, the CNN-UKF consistently has a lower MSE_{nl} than the DE-UKF, see Fig. 2.14b and only when both axle cornering stiffness of the vehicle model are increased by 10 % as, on average, a higher MSE than the DE-UKF. Thus, it can be concluded that hybrid approaches are more robust to parameter uncertainties than purely model approaches and that the DE-UKF performance is less affected by the internal model accuracy than CNN-UKF.

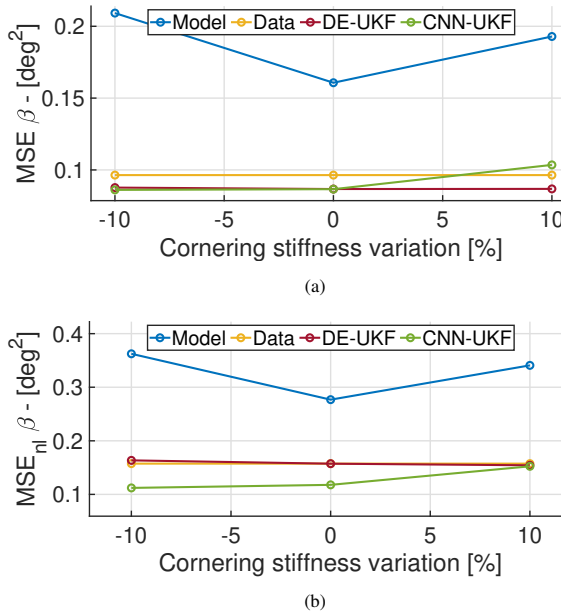


Figure 2.14: 2.14a shows how the axle cornering stiffness influences the vehicle sideslip angle MSE of all four approaches. 2.14b shows how the axle cornering stiffness influences the vehicle sideslip angle MSE_{nl} of all four approaches.

2.6 CONCLUSION

THE paper presents a novel hybrid approach to vehicle sideslip angle estimation, which involves utilising the physical knowledge from a UKF based on a single-track vehicle model to enhance the estimation accuracy of a CNN. Using a large-scale experimental dataset of 216 manoeuvres, it has been shown that the hybrid approach is more accurate than purely model-based or data-driven approaches. Moreover, the CNN-UKF is slightly reducing the MSE of the DE-UKF. However, when the MSE_{nl} is compared, the CNN-UKF outperforms the DE-UKF by 25 %, providing a much higher accuracy in the most critical operating region for active vehicle control systems. The CNN-UKF, thanks to the end-to-end training, is forcing the CNN to comply with the vehicle physics, reducing the ME and ME_{nl} of all other approaches. When a limited dataset is provided, the proposed hybrid approach has a minor improvement in the estimation robustness over the model-based and the DE-UKF approach for all the KPIs. The CNN-UKF is highly outperforming the estimation of a purely data-driven approach. Future works involve testing the generalisation capability of the CNN-UKF utilising a dataset with different levels of road grip, e.g. wet, snow or icy roads.

ACKNOWLEDGEMENT

The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484).



3

3

MODEL PREDICTIVE CONTOURING CONTROL FOR VEHICLE OBSTACLE AVOIDANCE AT THE LIMIT OF HANDLING

Compromise consists of reconciling divergent claims through mutual concessions. For this, it is necessary that both parties have a valid claim and something of value to offer each other.

The Virtuous Egoist, Ayn Rand

This chapter will be extended with more experimental data and submitted for publication as:  **A. Bertipaglia, M. Alirezai, R. Happee, and B. Shyrokau**, “Experimental Validation of Model Predictive Contouring Control for Vehicle Obstacle Avoidance at the Limit of Handling”. This chapter builds on  **A. Bertipaglia, M. Alirezai, R. Happee, and B. Shyrokau**, “Model predictive contouring control for vehicle obstacle avoidance at the limit of handling,” in *Symposium on the Dynamics of Vehicle on Roads and on Tracks*, Ottawa, Canada, 2023 [41].

ABSTRACT

This paper proposes a nonlinear Model Predictive Contouring Control (MPCC) for obstacle avoidance in automated vehicles driven at the limit of handling. The proposed controller integrates motion planning, path tracking and vehicle stability objectives, prioritising obstacle avoidance in emergencies. The controller's prediction model is a nonlinear single-track vehicle model with the Fiala tyre to capture the vehicle's nonlinear behaviour. The MPCC computes the optimal steering angle and brake torques to minimise tracking error in safe situations and maximise the vehicle-to-obstacle distance in emergencies. Furthermore, the MPCC is extended with the tyre friction circle to fully exploit the vehicle's manoeuvrability and stability. First, the performance is compared with a state-of-the-art Model Predictive Control (MPC) in a high-fidelity simulation environment. The double lane change scenario results demonstrate a significant improvement in successfully avoiding obstacles and maintaining vehicle stability at the limit of handling. Afterwards, the proposed MPCC is validated by performing a double lane change with the TU Delft experimental Toyota Prius.

3.1 INTRODUCTION

AUTOMATED vehicles' safety relies heavily on their ability to effectively avoid obstacles through evasive manoeuvres. Nevertheless, tyre nonlinearities pose a significant challenge in this regard [38]. A hierarchical controller architecture typically separates motion planning, path tracking, and vehicle stability tasks [39]. Although each task can be optimised separately, in an evasive manoeuvre, the three objectives could be in conflict with each other [26]. Therefore, to avoid potential conflict, we integrate motion planning, path tracking, and vehicle stability into a single controller for obstacle avoidance at high speed.

Recent studies highlight that vehicle stability constraints could lead to tracking errors, potentially causing collisions [26, 42]. A potential solution is integrating an obstacle avoidance controller with objectives including motion planning, path tracking and vehicle stability [122]. Model Predictive Control (MPC) based on a nonlinear single-track vehicle model can integrate all the controllers' objectives and modify the desired trajectory to keep the vehicle stable and at a safe distance from the object [26]. In order to run the MPC controller in real-time, the nonlinear single-track vehicle model is linearised into an affine time-varying model. The longitudinal and lateral control is considered separately. However, the affine model diminishes the model's accuracy and limits the control capabilities. Furthermore, the model fidelity is particularly affected when the lateral and longitudinal dynamics are coupled [38]. Thus, to address these limitations, a nonlinear vehicle model and nonlinear optimisation must be adopted [38, 57]. The vehicle and the obstacles are represented as a set of circles, and their distance is constantly measured in the cost function of the nonlinear MPC [38]. Vehicle kinematics is described using the Frenet coordinate system because it allows an easy determination of the vehicle location relative to a reference line. Despite these advantages, when the trajectory has a curvature, the vehicle-to-obstacle (V2O) distance in the Frenet coordinate system is over-estimated with respect to the distance in the Cartesian coordinate system [41]. Furthermore, the computation of the travelled distance of the vehicle with respect to the reference line at every time step requires an additional optimisation [46].

This paper proposes a Model Predictive Contouring Control (MPCC) based on a nonlin-



Figure 3.1: TU Delft experimental automated vehicle (third-generation Toyota Prius) used in the experimental validation.

ear single-track vehicle model for obstacle avoidance. The MPCC, using a Cartesian frame, aims to approximate the MPC based on the Frenet reference system, thanks to the introduction of the lag and contouring error in the cost function. Thus, it avoids overestimating the V2O distance due to the Frenet coordinates, and it eliminates the additional optimisation to compute the travelled distance of the vehicle with respect to the reference line. We exhaustively assess the controller performance in a high-fidelity simulation environment designing a double-lane change for vehicle obstacle avoidance. Furthermore, we validate the proposed controller in an experimental test, see Fig. 3.1, under heavy rain conditions, successfully performing the evaluated double-lane change at 30 km/h.

This paper is organised as follows: Section 3.2 reviews previous work and outlines the main contributions of this paper. Section 3.3 introduces the prediction model, and Section 3.4 details the proposed controller. Section 3.5 describes the experimental setup, with results summarised in Section 3.6. Finally, Section 3.7 presents the key findings and suggests directions for future work.

3.2 RELATED WORKS

A substantial body of literature addresses vehicle control for obstacle avoidance in automated driving at the handling limit. The common obstacle avoidance architecture is often referred to as hierarchical, as it divides the task into three separate control layers: motion planning, which generates a collision-free trajectory; path tracking, which computes the actuator commands; and vehicle stability, where the previously computed commands are adjusted to ensure that the vehicle remains within the stability envelope [123]. A typical approach involves using a high-level motion planning algorithm to compute a trajectory via nonlinear MPC based on a simplified point-mass vehicle model. The friction circle then constrains the point-mass model's accelerations to determine the vehicle handling limit. At the low level, a linear-time-varying MPC controller optimises the actuator commands to follow the planned trajectory. To minimise tracking error, the vehicle model must more accurately reflect vehicle dynamics; so, the low-level controller utilises a nonlinear single-track vehicle model with Pacejka tyre formulations rather than a simple point-mass model [39].

In this architecture, the performance of each controller is optimised independently, and potential conflicts between their objectives are not thoroughly analysed [39, 124–126]. For instance, the simplified point-mass vehicle model can generate unfeasible trajectories near the handling limit that the path-tracking controller cannot accurately follow. For instance, a single-track or a double-track vehicle model has handling limits different from those of a point-mass model. The difference can be negligible at low velocity, but it becomes relevant when the vehicle's absolute lateral acceleration exceeds 4 m/s^2 . Another notable issue is that the vehicle stability controller may modify the inputs computed by the path-tracking controller, thereby increasing tracking errors and potentially leading to a collision with an obstacle, even though it successfully maximises its objective of maintaining vehicle stability [124].

As a result, path-tracking and vehicle stability objectives are often integrated into a single controller [42, 57, 127, 128]. For instance, a possible solution consists of a path-tracking nonlinear MPC (NMPC), which can ensure stability by constraining the vehicle's sideslip angle [57]. The control inputs are the steering angle and the brake torque at each wheel. The NMPC prediction model is based on a double-track vehicle model with

the Dugoff tyre model, incorporating braking actuator dynamics via a first-order transfer function. The Dugoff-modified cornering stiffness is provided to the controller as an external parameter at every iteration, and its dynamic is excluded from the prediction model to reduce computational load. This simplification is accurate enough for the controller to follow the reference trajectory while staying inside the linear tyre working range. When the controller works in the tyre nonlinear region, particularly near the handling limit, the prediction model mismatches become relevant, limiting the NMPC's overall performance. The NMPC works in real-time with a sampling interval of 0.035 s [57]. Regardless, none of these approaches accounts for the consequences that tracking errors can have on obstacle avoidance performance.

To address these limitations, an alternative obstacle avoidance architecture is proposed, often referred to as "integrated" because it combines motion planning (re-planning), path tracking, and vehicle stability into a single controller [26, 38, 129]. This approach considers potential path-tracking errors that could lead to a collision and incorporates them into the trajectory re-planning process. Additionally, the controller can prioritise collision avoidance by temporarily relaxing stability constraints in emergencies, ensuring a more robust response.

One possible solution is to use an MPC to compute the optimal steering angle for obstacle avoidance, while a simple feedforward-feedback longitudinal controller calculates the required braking or acceleration force [26]. The MPC optimisation is split into two parts: a short prediction horizon is used for vehicle stability, while a longer horizon is employed for obstacle avoidance. The complex nonlinear coupling between longitudinal and lateral dynamics is ignored, and the single-track vehicle model in the MPC is linearised into an affine time-varying model. This allows the controller to work with a quadratic cost function and operate in real-time with a sampling interval of 0.01 s. Despite successful experimental validation in a hairpin manoeuvre, the controller's performance is limited by the lack of accuracy in modelling the longitudinal and lateral dynamic coupling [38]. An alternative approach uses an NMPC, which simultaneously optimises the steering angle, longitudinal force, and optimal brake distribution between the front and rear axles [38]. Due to the highly nonlinear dynamics involved, the prediction model cannot be linearised, so a nonlinear interior point solver is used to solve the NMPC. The controller's time step is increased to 0.05 s to ensure real-time performance. The vehicle and obstacles are modelled as circles, and the distance between them is continuously measured within the cost function. The controller prioritises obstacle avoidance over path tracking by considering the distances between the vehicle and obstacles and the vehicle and road edges to avoid passing dangerously close to them. The vehicle kinematics are modelled using a Frenet reference frame, simplifying the calculation of the vehicle's position relative to a reference line. The V2O distance is also measured using this coordinate system. However, when the trajectory has curvature, the Frenet system tends to overestimate the V2O distance compared to the Cartesian coordinate system [41].

This paper proposes a Model Predictive Contouring Control (MPCC) approach based on a nonlinear single-track vehicle model with a Fiala tyre model for obstacle avoidance. Originally introduced for robot motion planning at low speeds [45] and lap-time optimisation for scaled vehicles [46], the MPCC is extended here to address obstacle avoidance and vehicle stability. Unlike traditional MPCs that utilise the Frenet reference system, the MPCC operates in a Cartesian frame, incorporating lag and contouring errors into the cost function.

This approach prevents the overestimation of the vehicle-to-obstacle (V2O) distance, a common issue with Frenet coordinates, and eliminates the need for additional optimisation to compute the vehicle's distance travelled along the reference line. The MPCC simultaneously optimises the steering angle, longitudinal force, and brake distribution between the front and rear axles. Despite the increased complexity of the prediction model, the controller achieves improved accuracy while maintaining real-time feasibility on a rapid prototyping platform.

The contributions of this paper are threefold. The first is improving the overall obstacle avoidance performance of the proposed MPCC over a state-of-the-art MPC [38]. Both controllers can successfully avoid a collision between the vehicle and the obstacle in a double-lane change. However, the baseline MPC cannot keep the vehicle outside the unsafe area close to obstacles or road edges. The second contribution encompasses improved vehicle stability by minimising the peaks in sideslip angle and increasing the minimum velocity during manoeuvres due to better prioritisation of the obstacle avoidance objective within the MPCC framework. The third contribution is the successful experimental validation of the proposed MPCC, tested on a Toyota Prius performing a double-lane change at 30 km/h on a low-friction surface under heavy rain conditions. Additionally, the experimental scenario with a low road friction coefficient demonstrates the robustness of the proposed controller to different friction conditions.

3

3.3 PREDICTION MODEL

A nonlinear single-track vehicle model (Fig. 3.2) is used in the proposed MPCC. Only the in-plane dynamics are considered, ignoring the lateral weight transfer and the roll and pitch dynamics. The vehicle position is described using a Cartesian reference frame by the states ($x = [X, Y, \psi, v_x, v_y, r, \theta, \delta, F_x]$): longitudinal position (X), lateral position (Y), and the heading angle (ψ) of the vehicle centre of gravity (CoG) relative to an inertial frame. The velocity states are the longitudinal and lateral velocity at the CoG, respectively (v_x) and (v_y), and the yaw rate (r). Furthermore, an additional state corresponding to the vehicle travelled distance (θ) is introduced, and the MPCC cost function uses it to compute the vehicle position relative to the reference line. The steering angle (δ) and the longitudinal force (F_x) correspond to the integral of the control inputs. The implemented state derivatives, (\dot{x}), correspond to the following equations:

$$\begin{cases} \dot{X} = v_x \cos(\psi) - v_y \sin(\psi) \\ \dot{Y} = v_x \sin(\psi) + v_y \cos(\psi) \\ \dot{\psi} = r \\ \dot{v}_x = \frac{-F_{yf}(x, u_v) \sin(\delta) + F_{xf}(x, u_v) \cos(\delta) + F_{yr}(x, u_v) - F_{drag}}{m} + r v_y \\ \dot{v}_y = \frac{F_{yf}(x, u_v) \cos(\delta) + F_{xf}(x, u_v) \sin(\delta) + F_{yr}(x, u_v)}{m} - r v_x \\ \dot{r} = \frac{l_f F_{yf}(x, u_v) \cos(\delta) + l_f F_{xf}(x, u_v) \sin(\delta) - l_r F_{yr}(x, u_v)}{I_{zz}} \\ \dot{\theta} = \sqrt{v_x^2 + v_y^2} \end{cases} \quad (3.1)$$

where F_{xi} and F_{yi} are, respectively, the longitudinal and lateral tyre forces, i stands for front (f) or rear (r), l_f and l_r are the distance between the front and rear axle to the CoG, I_{zz} is the vehicle inertia around the z-axis, m corresponds to the vehicle mass and F_{drag} is the aerodynamic drag resistance.




Figure 3.2: A nonlinear single-track vehicle model.

The vehicle model inputs (u_v) are the road-wheel angle rate ($\dot{\delta}$), the total longitudinal force rate applied at the CoG (\dot{F}_x) and the brake repartition between the front and rear axle (λ_b).

The control input rates are integrated into the prediction model before being applied to the vehicle. The inputs governing the longitudinal dynamics are F_x and λ_b and not F_{xf} and F_{xr} . The reason is that the vehicle must be able to accelerate and brake, but it is forbidden to accelerate and brake simultaneously [38]. Thus, the constraint ($F_{xf}F_{xr} \geq 0$) is commonly applied. However, it introduces a saddle point when both F_{xf} and F_{xr} are equal to zero, not guaranteeing the Hessian to be positive-definite [38]. For this reason, the constraint is implicitly formulated inside the prediction model as follows:

$$F_{xf} = \begin{cases} \lambda_b F_x & \text{if } F_x \leq 0 \\ \lambda_d F_x, & \text{otherwise} \end{cases} \quad F_{xr} = \begin{cases} (1 - \lambda_b) F_x & \text{if } F_x \leq 0 \\ (1 - \lambda_d) F_x, & \text{otherwise} \end{cases} \quad (3.2)$$

where λ_d represents whether the vehicle is front- or rear-wheel driven.

A Fiala tyre model computes the lateral tyre force for each axle, while the longitudinal force is defined as an input of the system [38]. The nonlinear coupling between F_{xi} and F_{yi} is captured according to the friction circle [57]. The tyre parameters are optimised by performing quasi-steady-state circular driving in a high-fidelity simulation based on a Delft-Tyre model 6.1.

The vehicle and obstacles are represented as circles so that their Euclidean distance can constantly be computed as follows:

$$D_{V2O} = \sqrt{(X - X_{obs})^2 + (Y - Y_{obs})^2} - r_{obs} - r_{veh} \quad (3.3)$$

where X , Y and X_{obs} , Y_{obs} are, respectively, the longitudinal and lateral position of the vehicle and obstacle centre, and r_{veh} and r_{obs} are the radius of the vehicle and obstacle circles. Thus, the vehicle and the obstacle collide when the V2O distance D_{V2O} is lower than zero. The controller aims to keep D_{V2O} always positive and above a user-defined safety distance.

3.4 MODEL PREDICTIVE CONTOURING CONTROL

This section describes how the proposed MPCC controller integrates motion planning, path tracking, and vehicle stability objectives.




Figure 3.3: A representation of the contouring e_{Con} and lag error e_{Lag} . θ and θ_s are the vehicle travelled distance and the distance with respect to the reference line, respectively.

3.4.1 COST FUNCTION

The MPCC employs an iterative approach to solve an optimal control problem, enabling the vehicle to operate at its handling limit while avoiding obstacles. The cost function objectives encompass several aspects: tracking a reference longitudinal and lateral position, maintaining a desired velocity, dynamically adjusting the indicated trajectory to ensure a safe distance from obstacles while maintaining stability, and guaranteeing the physical feasibility of input signals. The proposed cost function (J) is:

$$J = \sum_{i=1}^N \left(q_{e_{Con}} e_{Con,i}^2 + q_{e_{Lag}} e_{Lag,i}^2 + q_{e_{Vel}} e_{Vel}^2 + q_{\delta} \delta_i^2 + q_{F_x} F_{x_i}^2 + \right. \\ \left. + q_{\lambda_b} e_{\lambda_b,i}^2 + \sum_{j=1}^{N_{obs}} (q_{e_{V2O}} e_{V2O,j,i}^2) + \sum_{j=1}^{N_{edg}} (q_{e_{V2E}} e_{V2E,j,i}^2) \right) \quad (3.4)$$

where N is the length of the prediction horizon, N_{obs} is the number of obstacles in the road, $N_{edg} = 2$ is the number of road edges, and parameters q_* are the weights of the respective quadratic errors, defined below. The weights are fine-tuned to optimise controller performance by minimising longitudinal velocity error and sideslip angle peaks [120] and enabling the vehicle to avoid obstacles while maintaining a safe distance from obstacles and road edges.

The reference trajectory is tracked through the introduction of the contouring error (e_{Con}) and the lag error (e_{Lag}) [45, 46], see Fig. 3.3. e_{Con} corresponds to the projection of the vehicle position over the desired trajectory. It is computed as a function of the vehicle travelled distance with respect to the reference line (θ_s). However, based on the Cartesian coordinate frame, the controller cannot determine the distance θ_s in the prediction model of Eq. 3.1. Vice versa, it corresponds to a state when a Frenet reference system is employed [38, 57]. Thus, in the MPCC, θ_s is approximated by the total vehicle travelled distance θ computed as in Eq. 3.1. To ensure the validity of this approximation, the norm between the two distances, called lag error, must be minimised. The errors are approximated as follows [45, 46]:

$$\bar{e}_{Con} = \sin(\Psi_t(\theta))(X - X_t(\theta)) - \cos(\Psi_t(\theta))(Y - Y_t(\theta)) \\ \bar{e}_{Lag} = -\cos(\Psi_t(\theta))(X - X_t(\theta)) - \sin(\Psi_t(\theta))(Y - Y_t(\theta)) \quad (3.5)$$

where X_t , Y_t , and Ψ_t are the desired longitudinal, lateral position and heading angle. It should be noted that \bar{e}_{Lag} is minimised by following the reference trajectory and modifying the vehicle velocity. Thus, if the desired velocity is unfeasible for the given trajectory, the controller will reduce it to keep the \bar{e}_{Lag} close to zero.

The velocity is tracked by computing e_{Vel} , which corresponds to the difference between the vehicle's velocity and the desired one.

The controller dynamically adjusts the reference trajectory to ensure a safe distance from obstacles. This is possible by evaluating the V2O distance. The error function ($e_{V2O} = D_{V2O} - D_{S_{ft,O}}$) is computed as the difference between a user-defined safety distance between the vehicle and the obstacle ($D_{S_{ft,O}}$), and the V2O distance D_{V2O} , see eq. 3.3. It should be noted that when $D_{V2O} \geq D_{S_{ft,O}}$, q_{V2O} is automatically set at zero, while when $D_{V2O} \leq D_{S_{ft,O}}$ the error is computed to keep the vehicle always at a safe distance from the object. A similar error is introduced to keep the vehicle away from the road edges ($e_{V2E} = D_{V2E} - D_{S_{ft,E}}$). $D_{S_{ft,E}}$ is the safety distance between the vehicle and the road edge, and D_{V2E} is the distance between the vehicle and the road edge.

The other cost terms focus on the smoothness and feasibility of control inputs. Thus, the cost terms are introduced to the steering angle rate and the longitudinal force rate applied at the vehicle CoG. Furthermore, the controller minimises the error (e_{λ_b}) between the actual brake repartition and the ideal one, similar to the ideal brake force distribution. The ideal brake repartition, $\frac{F_{zf}}{F_{zr}} / \frac{F_{zf} + F_{zr}}{F_{zr}}$, corresponds to the percentage of the total braking force that should be applied to the front axle to lock the front and the rear axle simultaneously. The error e_{λ_b} helps the controller provide the ideal braking repartition. It also allows a variation depending on various factors, e.g., the tyre saturation in the front and rear axle or the value of the steering angle.

3.4.2 CONSTRAINTS

The cost function is constrained based on actuator limitations, vehicle stability and road track width. The road-wheel steering angle, the total longitudinal force and their respective rates are limited by upper and lower constraints.

The vehicle stability is enforced by restricting the total available tyre force at each axle using the tyre friction circle. At first, the longitudinal force (F_x) is restricted considering $|F_x| \leq sf\mu F_z$. Given the difficulties of estimating the road friction coefficient (μ), the tyre friction circle is multiplied by a safety coefficient (sf) equal to 0.95, which limits the available longitudinal force. The Fiala tyre in the prediction model limits the maximum lateral tyre force according to the tyre friction circle at a given F_x .

The vehicle is enforced to stay on the track using the following constraint:

$$\left\| \begin{bmatrix} X \\ Y \end{bmatrix} - \begin{bmatrix} X_{cen} \\ Y_{cen} \end{bmatrix} \right\|^2 \leq \left(\frac{W_t}{2} \right)^2 \quad (3.6)$$

where X_{cen} and Y_{cen} are the longitudinal and lateral location of the centre of the track, and W_t is the width of the road track [46].

3.4.3 OBSTACLE AVOIDANCE PRIORITISATION

The cost function combines different objectives and needs to prioritise collision avoidance during an evasive manoeuvre. Thus, the weights associated with the V2O or vehicle-to-edge (V2E) distance vary dynamically, according to:

$$q_{V2O} = \begin{cases} P_k, & \text{if } D_{V2O} < 0 \\ P_k e^{-\frac{2D_{V2O}^2}{D_{Sty,O}^2}}, & \text{elseif } 0 \leq D_{V2O} \leq D_{Sty,O} \\ 0, & \text{otherwise} \end{cases} \quad (3.7)$$

where P_k represents the maximum value that q_{V2O} can reach. This value prioritises collision avoidance as the V2O distance decreases. It increases up to P_k in a Gaussian shape curve. This prioritisation approach is more gradual than the step growth from [38], which results in faster and more robust convergence of the MPCC solver. It is important to note that prioritising collision avoidance may result in a larger tracking error and deviation from the vehicle's desired velocity.

3.5 SIMULATION AND EXPERIMENTAL SETUP

THIS section is divided into two subsections. Subsection 3.5.1 describes how the proposed MPCC is evaluated in a high-fidelity simulation environment during a double-lane change at the limit of handling. Subsection 3.5.2 outlines the experimental validation of the proposed MPCC.

3.5.1 SIMULATION SETUP

The proposed control is firstly tested on the SCALEXIO dSPACE real-time platform, based on a multi-core DS6001 processor (2.8 GHz quad-core, 1 GB DDR2 SD RAM). The MPCC is set up in a separate core from the vehicle plant. The prediction model is discretised using Runge-Kutta 2 for its proper trade-off between accuracy and simplicity [38]. A 0.05 s sampling time and 50 steps prediction horizon is selected to make the controller real-time implementable. The optimisation problem is solved using the nonlinear interior point solver of FORCESPro [130]. The optimisation Hessian is approximated using the Broyden–Fletcher–Goldfarb–Shanno algorithm, and a user-defined initial Hessian matrix is provided to speed up the solving time. All the other solver parameters are kept as defaults. The platform successfully solves the nonlinear optimisation with an average time of 15.6 ms and max solving time of 18.6 ms. However, there is no mathematical guarantee that the solver will converge in time. The vehicle plant runs in a separate core at 1000 Hz. It is a high-fidelity BMW Series 545i vehicle model based on an IPG CarMaker simulation platform. Its parameters are determined through experimental inertia measurements, while the suspensions are characterised using a Kinematics & Compliance test rig. The tyre dynamics are based on a Delft-Tyre 6.1. The actuator dynamics are included through a second-order transfer function to increase the simulation accuracy [57]. Including actuator delays is particularly important for vehicle stability, as these delays contribute significantly to the onset of unstable limit cycles [131–133].

A double lane change manoeuvre with two obstacles is considered to assess the proposed controller's capabilities in avoiding obstacles at the limit of handling. The manoeuvre




Figure 3.4: Toyota Prius experimental setup 3.4a and the vehicle during the experimental validation 3.4b.

contains multiple obstacles because it is crucial to assess how the re-planning to avoid the first obstacle influences the vehicle's capacity to avoid the second one. A coarse trajectory is provided, corresponding to the desired path from the behaviour planner. The desired path indicates which side of the obstacle the collision avoidance controller needs to follow. Thus, the proposed controller follows the desired trajectory while also making necessary adjustments when the trajectory is too close to an obstacle or deemed infeasible. Considering the initial vehicle position and orientation, the desired trajectory is determined as a function of the road curvature and the distance along the reference line. The desired vehicle velocity is constant along all the manoeuvres [38].

3.5.2 EXPERIMENTAL SETUP

The experimental validation of the proposed controller is conducted on a third-generation Toyota Prius, currently available at the Intelligent Vehicle group at Delft University of Technology, as shown in Fig. 3.1. An Oxford Technical Solutions RT3000 V3 Global Navigation Satellite System (GNSS) aided inertial measurement system is used to provide vehicle position and orientation, as well as linear and angular velocities and accelerations at the vehicle's CoG. Due to physical constraints in the trunk, it was not possible to install the SCALEXIO dSPACE real-time platform in the Toyota Prius. As a result, the controller is adapted to run on a dSPACE 1007 AutoBox, as shown in 3.4a. The lower computational power of the dSPACE 1007 AutoBox, compared to the SCALEXIO dSPACE, required modifications to the proposed MPCC to ensure real-time performance on this rapid prototyping platform. Specifically, the MPCC prediction horizon is shortened to 30 steps from the 50 used in the simulation setup. Additionally, the optimisation problem is solved using the FORCESPro nonlinear sequential quadratic programming solver [130].

The third-generation Toyota Prius does not permit control over the brake distribution between the front and rear axles, which is one of the command inputs in the MPCC. Therefore, the proposed controller is reformulated to use the steering angle rate and the longitudinal force rate at the vehicle's CoG as inputs. Specifically, the commanded inputs are

the steering wheel angle and the longitudinal force, with factory brake distribution applied during braking and front-wheel drive used during acceleration. The MOVE box, developed by the Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek (TNO), serves as a gateway between the application platform and the vehicle's electronic control unit (ECU). The MOVE box interprets high-level commands received via the controller area network (CAN) from the dSPACE AutoBox and translates them into low-level commands tailored for the third-generation Toyota Prius.

3

The MOVE box incorporates safety mechanisms to ensure that internal safety limits are not exceeded and that the driver can always regain manual control. For longitudinal control, a maximum deceleration of -6 m/s^2 is permitted. Stricter rules apply to lateral control: steering torque is restricted to 1.5 Nm , and the steering rate is limited to 3.4 rad/s , with this limit further reduced as vehicle speed increases. Additionally, the steering wheel angle is restricted to approximately 5.2 rad at vehicle speeds up to $\approx 5 \text{ km/h}$ and to about $\approx 1 \text{ rad}$ at speeds of around $\approx 18 \text{ km/h}$.

In the experimental scenario, a double-lane change manoeuvre with two obstacles is considered. However, due to the limitations of the experimental setup, such as the constraints of the MOVE box, the manoeuvre is performed at 30 km/h . Additionally, the test is conducted on a low-friction surface under heavy rain conditions. A coarse trajectory, representing the desired path from the behaviour planner, is provided and computed under the assumption of a high-friction scenario. The MPCC prediction model is also tuned for high-friction conditions. Therefore, this experimental test is designed not only to validate the proposed controller but also to demonstrate its robustness to varying road friction coefficients.

3.6 RESULTS

THIS section is divided into three subsections. Subsection 3.6.1 compares the assumptions of using a Cartesian reference system versus a Frenet reference system. Subsection 3.6.2 analyses and compares the performance of the MPCC with a baseline MPC for collision avoidance at the limit of handling [38]. Finally, Subsection 3.6.3 presents the experimental validation of the proposed MPCC, conducted using a Toyota Prius.

3.6.1 CARTESIAN VS FRENET REFERENCE FRAME

The proposed MPCC controller computes the V2O and V2E distances using a Cartesian reference system. Vice versa, the baseline is an MPC built on a Frenet reference frame. The difference in the coordinate system implies a disagreement on how distance is measured. Fig. 3.5a and Fig. 3.5b show a vehicle driven on a 20 m radius circular road in the Cartesian and Frenet reference system. An obstacle is located on one side of the circle, with a normal displacement from the reference line equal to the radius of the obstacle. Assuming the vehicle drives on the reference line, the V2O distance is computed using both coordinates at every instance. Fig. 3.5c shows how the two distances vary depending on the vehicle's distance driven on the reference line. The distance in Frenet coordinate frame [38] is similar to the one computed in the Cartesian coordinate frame only when the vehicle is close to the obstacle. Vice versa, when the vehicle still needs to drive along the reference line, the V2O distance in the Frenet reference frame is overestimated compared with the Cartesian distance.

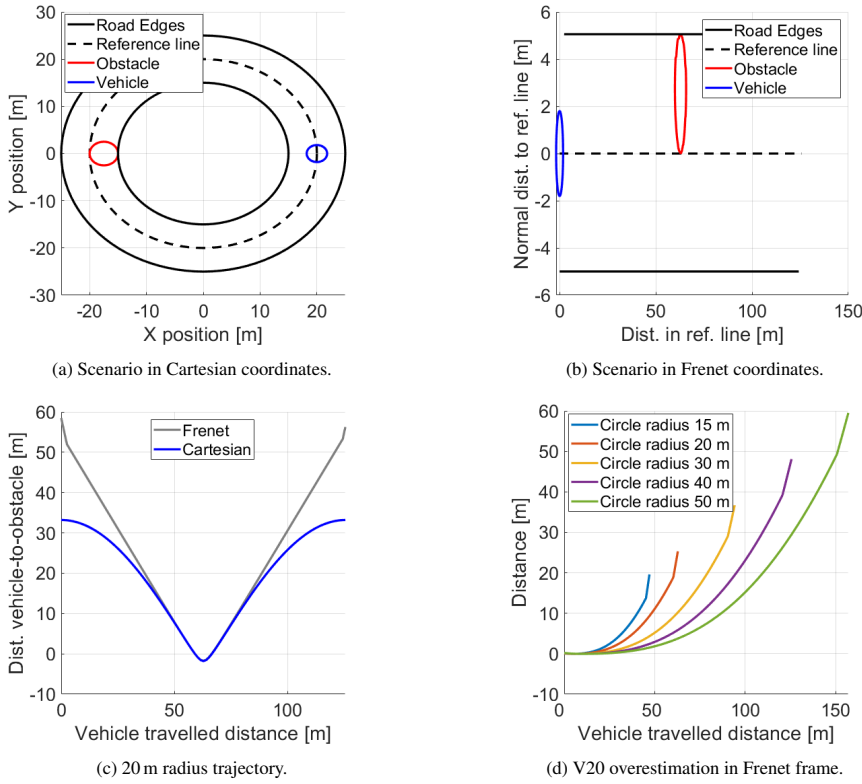


Figure 3.5: Fig. 3.5a and Fig. 3.5b show the scenario in Cartesian and Frenet coordinates. Fig. 3.5c - the V2O distance in a 20 m radius circular trajectory using Frenet coordinates. Fig. 3.5d - effect of trajectory curvature on the V2O distance overestimation.

Thus, the MPCC can prioritise the obstacle avoidance objective before the MPC based on the Frenet reference frame, improving the chances of safely avoiding the obstacle while keeping the vehicle stable. The difference between the Frenet and the Cartesian V2O distances depends on the road's curvature and the obstacle's normal distance from the reference line. Fig. 3.5d shows how the difference in V2O distance between Frenet and Cartesian changes with the road curvature. It can be observed that the higher the road curvature, the greater will be the overestimation.

3.6.2 SIMULATION RESULTS

Fig. 3.6a shows the vehicle trajectories obtained by three different controllers, the proposed MPCC with and without collision avoidance (CA) objective and the baseline controller for CA based on an MPC [38]. The MPCC without CA causes the vehicle to collide with the first obstacle of the double-lane change, as visible from the negative V2O distance, see Fig. 3.6d. The collision happens because the controller prioritises vehicle stability and path tracking over obstacle avoidance, even when the desired trajectory passes dangerously

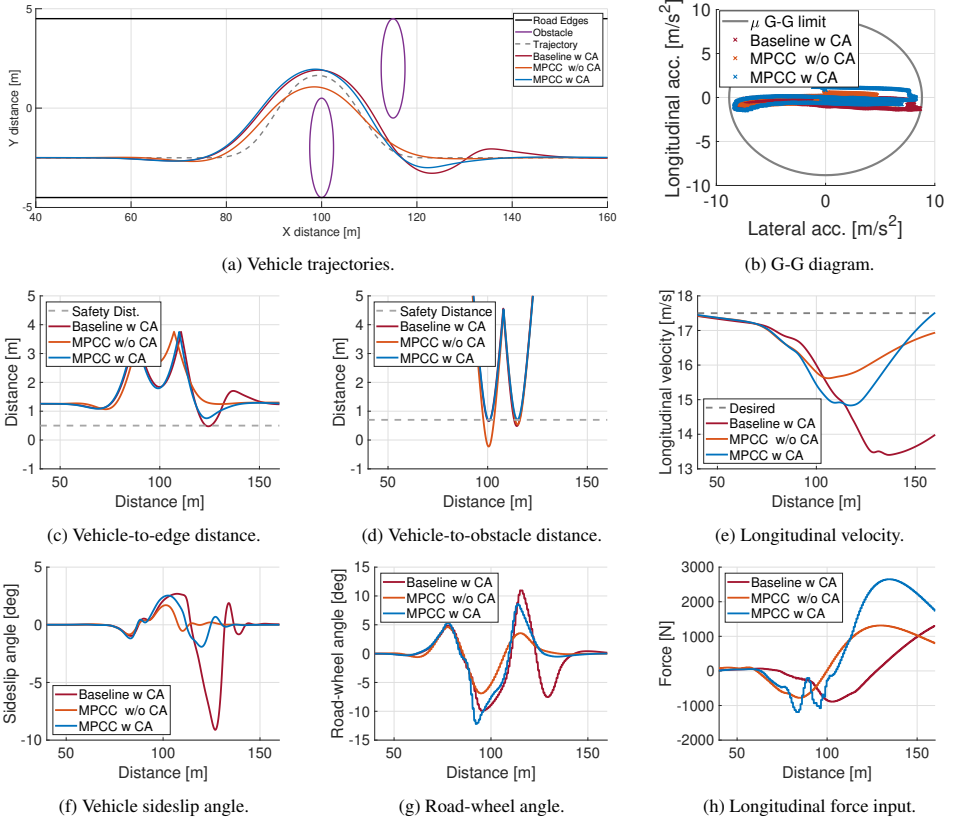


Figure 3.6: States and control inputs for the double-lane change manoeuvre.

close to an obstacle. This situation can happen when the path tracker and the vehicle stability controller affect the tracking performance without taking into account the position of obstacles. Alternatively, when the desired trajectory is unfeasible due to a mismatch in the handling limit of the motion planner and path-tracking prediction model. Vice versa, the MPCC and the baseline, both with CA prioritisation, successfully avoid both obstacles, see Fig. 3.6d. Despite the successful avoidance, the replanning around the first obstacle generated by the baseline negatively affects the rest of the manoeuvre. Thus, the vehicle trajectory has two noticeable overshoots at 125 m and 135 m, which causes the vehicle to enter inside the unsafe area near the second obstacle and to the right road edge, respectively Fig. 3.6d and Fig. 3.6c. The MPCC with CA can avoid both obstacles and keep the vehicle outside the unsafe area. This shows that CA prioritisation comes at the cost of increasing the path and velocity tracking errors and decreasing vehicle stability, pushing the vehicle to the boundaries of the handling limit.

For this reason, the baseline creates a sideslip angle peak of 9 deg at 130 m, see Fig. 3.6f. The baseline needs to counter-steer to return the vehicle to linear behaviour; see Fig. 3.6g. On the contrary, the MPCC with CA has a higher stability margin than the baseline

by reducing the vehicle sideslip angle peaks and allowing the vehicle to drive through the manoeuvres at a higher velocity, see Fig. 3.6e. Considering Fig. 3.6h, the MPCC with CA begins braking 20 m ahead of the baseline MPC, so the vehicle velocity at the entry of the corner is lower, and the vehicle can release the brake to fully exploit the lateral tyre force. However, the baseline applies hard braking during the central part of the corner, reducing the tyre's capability of generating lateral tyre force. Thus, the vehicle exceeds its handling limit obliging the controller to brake for a longer time and to stabilise the vehicle.

The MPCC applies braking before the baseline for two reasons: the Cartesian reference frame employed by the MPCC does not overestimate the V2O distance. The second reason is that the MPCC performs the tracking computing the contouring and the lag error. When the vehicle needs to replan the desired trajectory to avoid an obstacle, the MPCC controller tends to reduce the velocity because it is the only way to decrease the lag error and, consequently, the contouring error. However, the baseline MPC does not compute the lag error, so it decreases the vehicle velocity later when it faces a contouring error during the corner. However, it should be noted that the lag error in the MPCC brings complexity to the cost function weights' tuning due to the strong coupling between velocity and path tracking. Considering the available road friction coefficient, Fig. 3.6b shows that all the controllers reach the maximum lateral acceleration. However, the controllers do not fully exploit the maximum braking capabilities.

3.6.3 EXPERIMENTAL RESULTS

Fig. 3.7a shows the trajectory followed by the experimental vehicle during the double-lane change manoeuvre on a low-friction surface. Despite significant model mismatch caused by heavy rain, which reduced the road friction coefficient to approximately 0.5, the vehicle successfully avoided collisions with obstacles and remained within the road boundaries, as shown in Fig. 3.7e. The experimental results align with the simulation's conclusions, as the MPCC with CA successfully replanned the trajectory around the first obstacle without compromising the vehicle's performance near the second obstacle. However, unlike in the simulations, the vehicle entered the unsafe area near the first obstacle at approximately 63 m. This discrepancy could be due to two factors: first, the prediction model mismatch is greater than in the simulations due to the different road friction conditions; second, software constraints within the MOVE box, as indicated by the active flag in Fig. 3.7h, prevented the actuators from fully applying the desired MPCC steering angle and steering rate. This effect becomes particularly significant as the vehicle's longitudinal velocity increases, because the software constraints tighten exponentially with a linear increase in velocity.

The controller keeps the vehicle close to the desired velocity, as shown in Fig. 3.7c, with slight acceleration even during the collision avoidance manoeuvres (see Fig. 3.7h). As mentioned earlier, the proposed MPCC with CA successfully avoids both obstacles. However, this comes at the cost of increased path tracking error and reduced vehicle stability, pushing the vehicle to the limits of handling and achieving lateral accelerations of up to $\approx 4 \text{ m/s}^2$ (see Fig. 3.7b). In terms of vehicle stability, during the double-lane change, the vehicle reaches a sideslip angle peak of -4° (Fig. 3.7d). Nevertheless, the MPCC eliminates the need for counter-steering and, as observed in the simulations, does not negatively affect the controller's collision avoidance performance. The increased path tracking error, particularly the peak error around 70 m, is not solely due to the prioritisation of collision avoidance and

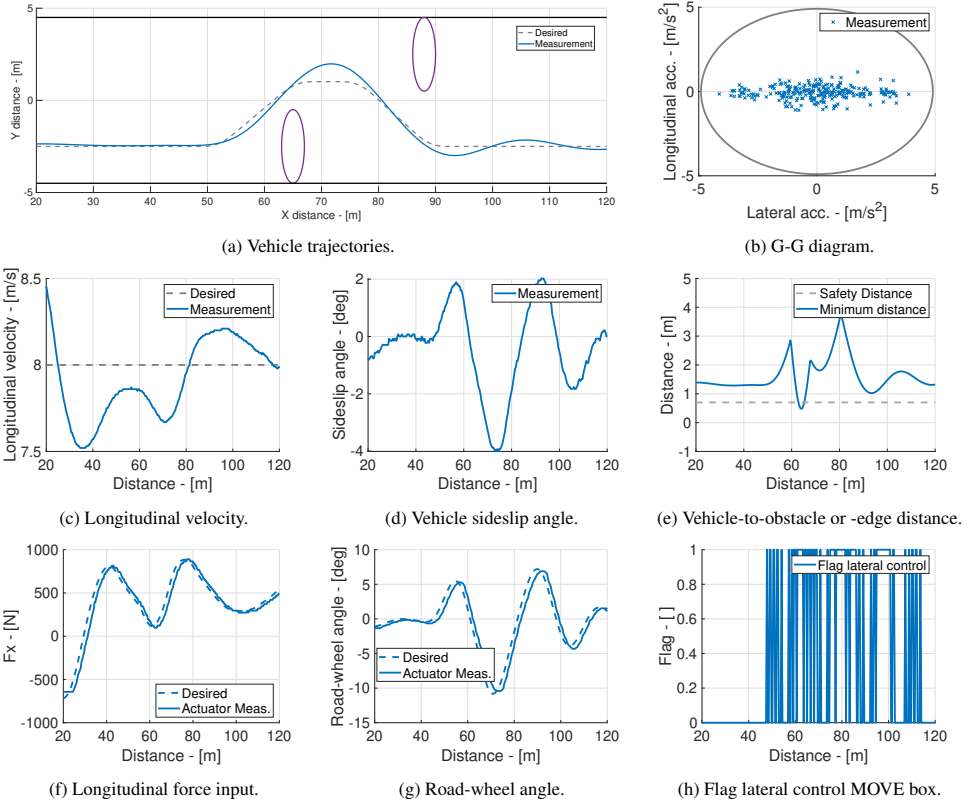


Figure 3.7: States and control inputs for the experimental double-lane change manoeuvre with low road friction coefficient.

the vehicle operating at the handling limit, as seen in the simulation results. It can also be attributed to the software limitations of the MOVE box.

Overall, the experimental test successfully validated the proposed controller, demonstrating the benefits of prioritising collision avoidance over path tracking and vehicle stability during a manoeuvre. Additionally, the heavy rain conditions at the proving ground provided an opportunity to evaluate the robustness of the proposed MPCC with CA under varying road friction coefficients. However, the strict software constraints of the MOVE box prevented a full evaluation of the controller near the limit of handling at a high friction coefficient. This limitation can only be addressed with a different experimental setup, for instance, substituting the MOVE box with a driving robot.

3.7 CONCLUSION

THIS paper proposed an innovative approach for obstacle avoidance in automated vehicles driven at the handling limit. A nonlinear MPCC integrates motion planner, path tracking and vehicle stability objectives into a single controller, prioritising obstacle avoidance in

emergency situations. In a double-lane change manoeuvre, the MPCC successfully avoids two obstacles, shortly re-planning the target trajectory from the behaviour planner. At the same time, the same controller without CA prioritisation collides with the obstacles. The CA prioritisation comes with decreased path-tracking performance and increased vehicle sideslip angle peaks up to 3 deg. However, the vehicle remains stable and manoeuvrable along the double-lane change. The state-of-the-art baseline also avoids the two obstacles but cannot keep the vehicle outside the unsafe area close to the two obstacles. Furthermore, it loses vehicle stability reaching a sideslip angle peak equal to 9 deg. The proposed controller is experimentally validated under heavy rain conditions, successfully performing a double lane change at 30 km/h with two obstacles using a third-generation Toyota Prius. The robustness of the MPCC with CA is further confirmed through successful validation on a low road friction surface. Future work will focus on implementing the proposed MPCC in a different experimental setup that allows for validation near the limit of handling, including scenarios with high friction coefficients.

ACKNOWLEDGEMENT

The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484).



4

4

ON THE BENEFITS OF TORQUE VECTERING FOR AUTOMATED COLLISION AVOIDANCE AT THE LIMITS OF HANDLING

*A statistical analysis, properly conducted, is a delicate dissection of uncertainties,
a surgery of suppositions.*

Facts from Figures, Michael J. Moroney

This chapter is based on  **A. Bertipaglia**, D. Tavernini, U. Montanaro, M. Alirezai, R. Happee, A. Sorniotti, B. Shyrokau, "On the Benefits of Torque Vectoring for Automated Collision Avoidance at the Limits of Handling," in *IEEE Transactions on Vehicular Technology*, 2025 (Under review). This chapter builds on  **A. Bertipaglia**, D. Tavernini, U. Montanaro, M. Alirezai, R. Happee, A. Sorniotti, B. Shyrokau, "Model Predictive Contouring Control for Vehicle Obstacle Avoidance at the Limit of Handling Using Torque Vectoring," in *IEEE International Conference on Advanced Intelligent Mechatronics*, USA, 2024 [134].

ABSTRACT

This paper presents a novel approach integrating motion replanning, path tracking and vehicle stability for collision avoidance using nonlinear Model Predictive Contouring Control. Employing torque vectoring capabilities, the proposed controller is able to stabilise the vehicle in evasive manoeuvres at the limit of handling. A nonlinear double-track vehicle model, together with an extended Fiala tyre model, is used to capture the nonlinear coupled longitudinal and lateral dynamics. The optimised control inputs are the steering angle and the four longitudinal wheel forces to minimise the tracking error in safe situations and maximise the vehicle-to-obstacle distance in emergency manoeuvres. These optimised longitudinal forces generate an additional direct yaw moment, enhancing the vehicle's lateral agility and aiding in obstacle avoidance and stability maintenance. The longitudinal tyre forces are constrained using the tyre friction cycle. The proposed controller has been tested on rapid prototyping hardware to prove real-time capability. In a high-fidelity simulation environment validated with experimental data, our proposed approach successfully avoids obstacles and maintains vehicle stability. It outperforms two baseline controllers: one without torque vectoring and another one without collision avoidance prioritisation. Furthermore, we demonstrate the robustness of the proposed approach to vehicle parameter variations, road friction, perception, and localisation errors. The influence of each variation is statistically assessed to evaluate its impact on the performance, providing guidelines for future controller design.

4.1 INTRODUCTION

ENSURING automated vehicles can adeptly avoid obstacles at the limit of handling, which is essential for maximising safety in real-world driving scenarios. Nonetheless, the highly nonlinear characteristics of tyres, mainly when longitudinal and lateral forces interact, present a notable challenge [57, 135]. A prevalent strategy entails developing collision avoidance controllers that optimise steering angle and total longitudinal brake force. However, a pure braking force can potentially lead to dangerous situations. For example, a vehicle must accelerate after performing an evasive manoeuvre to avoid a rear-end collision or quickly steer to the original lane during a double-lane change manoeuvre. Furthermore, purely braking and steering commands may not act fast enough to prevent a collision in an emergency manoeuvre [136]. Thus, we develop a collision avoidance controller which includes a positive acceleration force in addition to the typical inputs such as braking force and steering angle. Moreover, we extend the controller with torque vectoring capabilities to enhance vehicle lateral agility and yaw rate control.

Vehicle obstacle avoidance controllers are typically based on a hierarchical architecture that splits motion planning, path tracking, and vehicle stability objectives into separate controllers. This simplifies the cost function formulation of each layer and allows the adoption of a different prediction model for each controller (Fig. 4.1a). For instance, the motion planner can adopt a simple kinematic model to extend the prediction horizon or include a probabilistic approach to deal with uncertainties [137]. However, the three different objectives might conflict with each other, e.g. the vehicle stability controller introduces an unwanted tracking error to keep the vehicle stable, or the path tracking cannot perfectly track the motion planner trajectory due to the different complexity in the prediction models. Thus, various Model Predictive Control (MPC) algorithms have recently been introduced to integrate motion planning, path tracking and vehicle stability objectives into

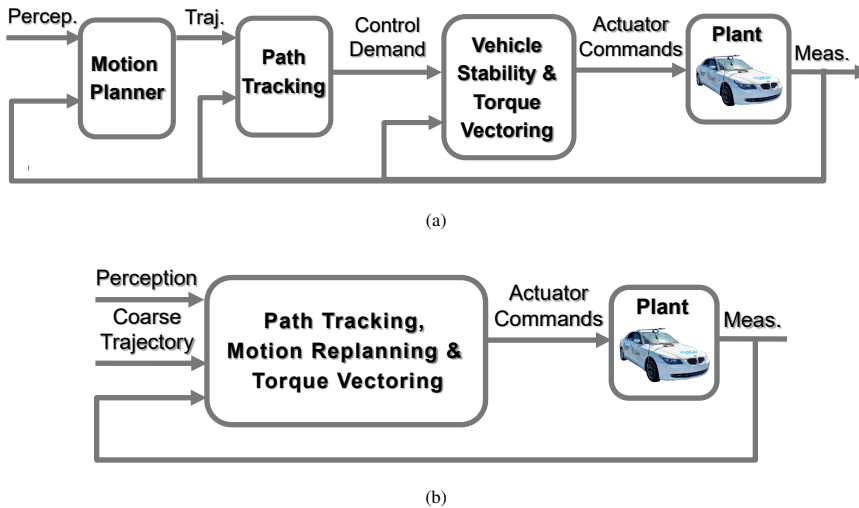


Figure 4.1: Hierarchical architecture for vehicle obstacle avoidance (Fig. 4.1a); proposed architecture (Fig. 4.1b).

a single controller for vehicle collision avoidance during emergency manoeuvres [38, 41, 42, 138] (See Chapter 3). Despite the improved performance of integrated vehicle obstacle avoidance controllers, they have yet to be enabled to utilise torque vectoring capabilities (Fig. 4.1b). Emerging control methodologies incorporating torque vectoring have gained significant appeal in light of the recent development of new electric powertrains, notably those employing multiple in-wheel electric motors [139–141].

Thus, this paper proposes a vehicle obstacle avoidance approach consisting of Model Predictive Contouring Control (MPCC), which combines motion planning, path tracking, and vehicle stability tasks into a single cost function, prioritising vehicle collision avoidance in case of an emergency. Furthermore, the approach employs torque vectoring capabilities to improve the vehicle's lateral agility at the limit of handling. The torque vectoring is enabled in the optimal control problem formulation thanks to the adoption of a double-track vehicle model based on an extended Fiala tyre model. The proposed tyre model captures the effect of longitudinal force on the cornering stiffness and the diminishing tyre saturation region. This improves the accuracy of the prediction model, reducing the vehicle model mismatch with the plant while maintaining the real-time feasibility on a rapid prototyping platform.

This paper is organised as follows: Section 4.2 summarises the previous works and the main paper contributions, highlighting the differences and improvements compared with our previous conference paper [134]. Section 4.3 presents the prediction model. Section 4.4 describes the proposed controller. Section 4.5 shows the experimental setup, summarising the results in Section 4.6. Section 4.7 concludes the essential findings and future works.

4.2 RELATED WORKS

A sizeable amount of literature exists on obstacle avoidance controllers. For the sake of brevity, a selection of the most relevant research is presented herein.

One line of research has focused on solutions represented by a three-layer control framework: a Nonlinear Model Predictive Control (NMPC) for path tracking, a stability controller to compute the reference yaw rate, and an optimal tyre force allocation algorithm for torque vectoring [142]. NMPC is based on a single-track formulation with a linear tyre model. Given the significant model mismatch intrinsic in the simplistic tyre model, the stability controller layer calculates a desired steady-state yaw rate to ensure vehicle stability. Meanwhile, the torque vectoring layer optimally distributes tyre forces across each wheel based on the pre-calculated desired yaw rate and longitudinal force. Simulation results showcase enhancements in lateral stability and reductions in path tracking error. However, the limited accuracy of the linear tyre model during emergency manoeuvres significantly reduces performance. Additionally, separating the path tracking layer from the reference yaw rate computation diminishes the advantages of torque vectoring. For instance, even a simple Linear Quadratic Regulator controller with integrated path tracking and torque vectoring achieves a higher entry speed and vehicle agility in a double-lane change than a split controller configuration [143].

An integrated MPC [136] is proposed to address these challenges. The MPC incorporates steering and differential braking strategies to facilitate collision avoidance [136]. The approach splits the longitudinal and lateral dynamics, with the MPC utilising a linearised brush tyre model to calculate the desired lateral tyre force and additional differential braking moment. However, the accuracy of the linearised tyre model diminishes significantly at

handling limits and neglects the interaction between longitudinal and lateral dynamics. Moreover, the prediction relies on a single-track vehicle model, limiting the maximum stabilising yaw moment by a portion of the tyre force capacity. This limitation is necessary to safeguard the vehicle's lateral force capacity from unmodeled dynamics, i.e. combined slip and lateral weight transfer, resulting in a more conservative controller approach.

A similar approach relies on a single-track vehicle model to control the steering angle, the braking force distribution and an additional yaw moment [51]. The controller can drive the vehicle at the handling limit and even maintain the vehicle in a drifting equilibrium. Nonetheless, due to the constrained prediction horizon of 1 s necessary for real-time execution and the cost function formulation, there is no possibility for trajectory replanning in emergency scenarios, when obstacles might suddenly change their position unless a new reference trajectory is provided to the path tracking algorithm. A similar limitation is recognisable also for a dynamic inversion-based path tracking controller recently proposed to control the vehicle beyond and within the stability limits [144].

A Model Predictive Contouring Control (MPCC) based on a nonlinear single-track vehicle model has been proposed for vehicle collision avoidance at the limit of handling [41] (See Chapter 3). By adopting a Cartesian reference frame, the MPCC delineates the vehicle's kinematics, ensuring precise measurement of the vehicle-to-obstacle (V2O) distance, a factor prone to overestimation when employing the Frenet reference frame [38]. Additionally, it circumvents the need for supplementary optimisation to compute the distance travelled by the vehicle to the reference line [41, 46]. Thus, its safety performance is superior to controllers based on hierarchical architecture [39, 42, 122] and even surpasses an integrated architecture reliant on the Frenet reference system [38]. However, the controller does not consider the torque vectoring capabilities, and the single-track vehicle model cannot accurately capture the effect of lateral load transfer on vehicle handling.

This paper extends the design and the analysis of the solution presented in [134], proposing an MPCC controller based on a nonlinear double-track vehicle with an extended Fiala tyre model for collision avoidance at the limit of handling. The approach is extended to consider the torque vectoring capabilities, and it integrates motion replanning, path tracking, torque vectoring and vehicle stability tasks, prioritising vehicle collision avoidance. The additional complexity brought by the longitudinal tyre force optimisation requires a reduction of the model mismatch, specifically in the tyre. Thus, the Fiala tyre model is extended to accurately reflect the variability of cornering stiffness with vertical and longitudinal forces and adjustments to its saturation region to capture the tyre behaviour in the combined slip conditions. The added complexity allows the MPCC to optimise directly the longitudinal tyre forces and the yaw moment without using a lower controller for force allocation. This reduces architecture conservativeness [145], improving the collision avoidance performance. Despite the controller complexity, the proposed MPCC runs in real-time on a rapid prototyping platform, contrary to our conference paper [134], and is extensively evaluated in a high-fidelity simulation environment validated with experimental data.

The contributions of this paper are threefold. The first is the development of the first real-time feasible MPCC controller augmented with torque vectoring functionalities that can safely avoid vehicle collisions in a double-lane change manoeuvre at the limit of handling. In contrast, existing state-of-the-art approaches [38, 41] (Chapter 3) would result in a collision under similar conditions or not work in real-time [134]. Leveraging the improved

responsiveness induced by torque vectoring, the vehicle is directed away from the obstacle, ensuring stability and preventing potential crashes. Unlike our previous conference paper [134], the first contribution is extended to an analysis of the proposed controller in low friction conditions ($\mu = 0.5$), e.g. in heavy rain. Even in these conditions, the proposed controller can successfully avoid the obstacles without losing stability. Furthermore, it is tested at different initial velocities, highlighting its higher performance in a broader set of initial conditions compared to the baselines.

The second contribution involves developing and applying an extended Fiala tyre model capable of capturing variations in cornering stiffness relative to longitudinal and vertical forces and adjusting the gradient within the tyre's saturation region without losing the continuity of the function. It significantly enhances the accuracy of the prediction model, so its applicability for the controller in scenarios with high levels of force coupling, e.g. when torque vectoring is enabled. Due to the increased accuracy provided by the extended Fiala tyre model, the wheel dynamics can be neglected in the prediction model. This simplification allows for an increased sampling time and reduced computational effort.

In contrast to our conference paper [134], the third contribution is the robustness evaluation of the proposed MPCC with torque vectoring against variations in vehicle and tyre parameters, perception inaccuracies, false negatives, and localisation errors, compared to baseline controllers. Notably, the collision rate of the proposed controller is 31.20 % and 38.50 % for vehicle parameter and perception inaccuracies, respectively, compared to nearly ~100 % for the baseline controllers. Additionally, the sensitivity analysis offers numerical insights into which parameters and perception inaccuracies have the most significant impact on the performance degradation of obstacle avoidance MPCC controllers. Specifically, the lateral peak friction coefficient and the obstacle localisation inaccuracies in the lateral direction emerge as the most influential factors affecting controller performance.

4.3 PREDICTION MODEL

THIS section presents the prediction model implemented in the proposed MPCC controller. At first, the double-track vehicle model is described. Second, the extended Fiala tyre model with the proposed improvements is presented.

4.3.1 DOUBLE-TRACK VEHICLE MODEL

The proposed MPCC integrates a nonlinear double-track vehicle model, chosen over the single-track vehicle model [38, 136, 146] due to its ability to capture lateral weight transfer and its superior accuracy at the vehicle's handling limits [135]. However, it is simplified by neglecting roll and pitch dynamics due to their relatively small contribution to load transfer for the considered vehicle class. The prediction model comprises twelve states denoted by $x = [X, Y, \psi, v_x, v_y, r, \theta, \delta, F_{x,fl}, F_{x,fr}, F_{x,rl}, F_{x,rr}]$. These states describe the vehicle's position and orientation in a Cartesian reference system: longitudinal position (X), lateral position (Y), and the heading angle (ψ) of the vehicle's centre of gravity (CoG) relative to an inertial frame. Additionally, longitudinal and lateral velocities at the CoG (v_x and v_y , respectively) and yaw rate (r) are included. The MPCC requires knowing the vehicle's travelled distance (θ), utilised by the cost function to compute the vehicle's position relative to the reference line; thus, θ is incorporated as an extra state [41] (See Chapter 3). The road

Table 4.1: Vehicle parameters description.

Parameters	Symbol	Value
Vehicle mass	m	1997 kg
Vehicle inertia around the z-axis	I_{zz}	3198 kgm ²
Distance between the front axle to CoG	l_f	1.430 m
Distance between the rear axle to CoG	l_r	1.455 m
Front axle track width	t_f	1.540 m
Rear axle track width	t_r	1.576 m
Air density	ρ	1.204 kg/m ³
Drag coefficient	C_{d1}	0.25
Rolling resistance	C_{d0}	45 N
Vehicle frontal area	A_f	2.4 m ²

wheel angle (δ) and the longitudinal forces at the front left ($F_{x,fl}$), front right ($F_{x,fr}$), rear left ($F_{x,rl}$), and rear right ($F_{x,rr}$) wheels are controlled via their derivatives, leading to five additional equations of motion. The state derivatives are computed as follows:

$$\begin{cases}
 \dot{X} = v_x \cos(\psi) - v_y \sin(\psi) \\
 \dot{Y} = v_x \sin(\psi) + v_y \cos(\psi) \\
 \dot{\psi} = r \\
 \dot{v}_x = \frac{1}{m} \left((F_{x,fl} + F_{x,fr}) \cos(\delta) - (F_{y,fl} + F_{y,fr}) \sin(\delta) + \right. \\
 \quad \left. + F_{x,rl} + F_{x,rr} - F_{res} \right) + r v_y \\
 \dot{v}_y = \frac{1}{m} \left((F_{x,fl} + F_{x,fr}) \sin(\delta) + (F_{y,fl} + F_{y,fr}) \cos(\delta) + \right. \\
 \quad \left. + F_{y,rl} + F_{y,rr} \right) - r v_x \\
 \dot{r} = \frac{1}{I_{zz}} \left((F_{y,fl} + F_{y,fr}) \cos(\delta) l_f - (F_{y,rl} + F_{y,rr}) l_r + \right. \\
 \quad \left. + (F_{x,fl} + F_{x,fr}) \sin(\delta) l_f + \frac{l_f}{2} (F_{y,fl} - F_{y,fr}) \sin(\delta) + \right. \\
 \quad \left. + \frac{l_f}{2} (F_{x,fr} - F_{x,fl}) \cos(\delta) + \frac{l_r}{2} (F_{x,rr} - F_{x,rl}) \right) \\
 \dot{\theta} = \sqrt{v_x^2 + v_y^2}
 \end{cases} \quad (4.1)$$

where $F_{x,ij}$ and $F_{y,ij}$ are the longitudinal and lateral tyre forces, i stands for front (f) or rear (r), and j stands for left (l) or right (r). All other vehicle parameters are reported in Table 4.1. Moreover, F_{res} is the aerodynamic drag and the rolling resistance computed according to the following equation:

$$F_{res} = \frac{1}{2} \rho A_f C_{d1} v_x^2 + C_{d0} \quad (4.2)$$

The vehicle model inputs are the rates of the previously mentioned road wheel angle and the rates of longitudinal forces applied to each wheel, i.e. $(u_v = [\dot{\delta}, \dot{F}_{x,fl}, \dot{F}_{x,fr}, \dot{F}_{x,rl}, \dot{F}_{x,rr}])$. The rates are used as inputs to apply constraints representing actuator dynamics and create further responses. The decision to use longitudinal forces as inputs, rather than motor torques or longitudinal slips, simplifies the vehicle model by eliminating the need to account for wheel or motor dynamics in the MPC prediction model. This not only reduces computational

complexity but also facilitates the implementation of constraints, allowing tyre forces to be constrained using the tyre friction circle [38, 147].

The vehicle and obstacles are represented as circles so the MPCC controller can constantly monitor the V2O distance. A similar approach is implemented for the vehicle-to-edge (V2E) of the road distance. Their Euclidean distance is computed as follows:

$$D_{V2O} = \sqrt{(X - X_{obs})^2 + (Y - Y_{obs})^2} - r_{obs} - r_{veh} \quad (4.3)$$

where X , Y and X_{obs} , Y_{obs} are, respectively, the longitudinal and lateral position of the vehicle and obstacle centre, and r_{veh} and r_{obs} are the radii of the vehicle and obstacle circles. The proposed MPCC controller aims to keep D_{V2O} above a user-defined safety distance. The vehicle and the obstacles will collide if D_{V2O} is lower than zero.

4

4.3.2 EXTENDED FIALA TYRE MODEL

The lateral tyre forces for each wheel of the double-track vehicle model are captured by an extended Fiala tyre model. The classic Fiala tyre model is modified to capture the variation of cornering stiffness depending on the longitudinal and vertical force [135], and the saturation region is adapted to include a negative gradient. The latter allows the prediction model not to overestimate the maximum lateral force when the tyre works with a high lateral slip angle, e.g. driving at the limit of handling or drifting. The extended Fiala tyre model is defined as follows:

$$F_y(\alpha, F_x, F_z) = \begin{cases} -C_{ym}(F_x, F_z) \tan \alpha + \frac{C_{ym}^2(F_x, F_z) \tan \alpha \tan |\alpha|}{3F_{y,max}} + \frac{C_{ym}^3(F_x, F_z) \tan \alpha^3}{27F_{y,max}^2}, & |\alpha| \leq \alpha_{thr} \\ \frac{2C_{ym}(F_x, F_z)(\zeta-1) \tan \alpha}{3} - \frac{C_{ym}^2(F_x, F_z)(\zeta-1) \tan \alpha \tan |\alpha|}{9F_{y,max}} + \\ -F_{y,max}\zeta \operatorname{sign}(\alpha), & |\alpha| > \alpha_{thr} \end{cases} \quad (4.4)$$

where α is the tyre slip angle, C_y is the tyre cornering stiffness, which is a function of the vertical (F_z) and longitudinal (F_x) tyre force, $F_{y,max}$ is the maximum lateral tyre force, α_{thr} is the tyre slip threshold corresponding to the peak of the tyre lateral force, and ζ is a parameter defined between 0 and 1 which characterises the gradient of the saturation region. When $\alpha \leq \alpha_{thr}$, apart from the effect of F_x , which is active also for low lateral slip, the extended Fiala model is formulated as the classic Fiala tyre model [148], while the saturated region ($\alpha > \alpha_{thr}$) is modified to have a gradient that better captures the maximum lateral force reduction with large slip angles. At the same time, the proposed model still keeps the advantages of the classical Fiala tyre, so it is fully continuous and differentiable when $\alpha = \alpha_{thr}$. A gradient different from zero in the saturation region helps numerical optimisation algorithms based on the gradient calculation to avoid derivative vanishing and to optimise the road wheel angle when the tyre works in the saturation region [149]. Furthermore, the proposed solution has a positive gradient when $\zeta \in [1, 2]$ and a negative one when $\zeta \in [0, 1]$.

To further reduce the tyre model mismatch, the C_y is not considered constant, but it is firstly adapted depending on the vertical force [135] as follows:

$$C_y(F_z) = c_1 F_{z0} \sin \left(2 \operatorname{atan} \left(\frac{F_z}{c_2 F_{z0}} \right) \right) \quad (4.5)$$

Table 4.2: Tyre parameters description.

Parameters	Symbol	Value
Lateral cornering stiffness effect	c_1	49.3
Lateral cornering stiffness peak effect	c_2	3.5
Long. and Lat. tyre force coupling effect	c_3	4.1
Saturation region gradient	ζ	0.87
Nominal vertical tyre force (as in .tir property file)	F_{z0}	4300 N
Tyre friction coefficient	μ	0.95

where c_1 and c_2 are tunable parameters, and F_{z0} is the nominal vertical load. Second, the tyre model considers the coupling effect between longitudinal and lateral axle forces, so the previously computed C_y is further modified to capture the F_x dependency in accordance with literature [135, 150] as follows:

$$C_{ym}(F_x, F_z) = \frac{1}{2}(\mu F_z - F_x) + \left(1 - \left(\frac{|F_x|}{\mu F_z}\right)^{c_3}\right)^{1/c_3} \left(C_y(F_z) - \frac{1}{2}\mu F_z\right) \quad (4.6)$$

where c_3 is a user-defined parameter, which is in the range of 2 to 8 [150], and μ is the friction coefficient. The $C_{ym}(F_x, F_z)$ is used to compute the tyre slip threshold (α_{thr}) as follows:

$$\alpha_{thr} = \frac{3F_{y,max}}{C_{ym}(F_x, F_z)} \quad (4.7)$$

The maximum lateral tyre force ($F_{y,max}$) is limited by the tyre friction circle, defined as follows:

$$F_{y,max} = \sqrt{(\mu F_z)^2 - F_x^2} \quad (4.8)$$

All the parameters used in the extended Fiala tyre model are reported in Table 4.2, and their value is obtained by performing a nonlinear optimisation known as Two-Stage Bayesian Optimisation [120]. The parametrisation process involves two steps: first, the tyre models are optimised to minimize discrepancies in lateral forces compared to a high-fidelity Delft-Tyre 6.2 model, using nominal loads from 1000 N to 8000 N representative of double lane change manoeuvres at the limit of handling. Second, the tyre parameters are further optimised against experimental data from skidpad manoeuvres in clockwise and counterclockwise directions. Notably, the optimisation and evaluation manoeuvres differ to avoid overfitting and ensure robust model validation. Fig. 4.2 shows how the proposed extended Fiala tyre model captures the effect of the normal load on the tyre cornering stiffness and how the model mismatch between the Delft tyre model and the proposed one is reduced not only in the linear region but also around the peak lateral force area. Particularly relevant is the saturated region, which is well matched by the extended Fiala tyre model while it is overestimated by the classic Fiala tyre model with a constant saturation region [38, 135, 148] or by the simplified Magic Formula when subjected to an increased vertical force of 7300 N. A model mismatch in the tyre's large slip angle working area is particularly detrimental for obstacle avoidance controllers at the limit of handling. Fig. 4.2 shows that the cornering stiffness of the proposed extended Fiala model captures the longitudinal and lateral force

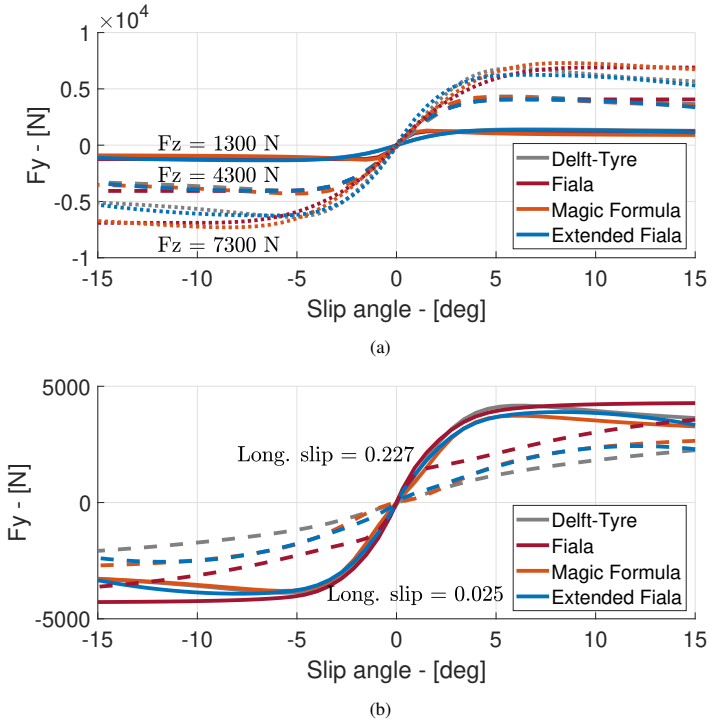


Figure 4.2: Lateral force of the high-fidelity Delft tyre model, the classic Fiala, the simplified Magic Formula and the extended (proposed) Fiala tyre model for pure slip conditions (a); lateral forces of the models in the combined slip conditions (b).

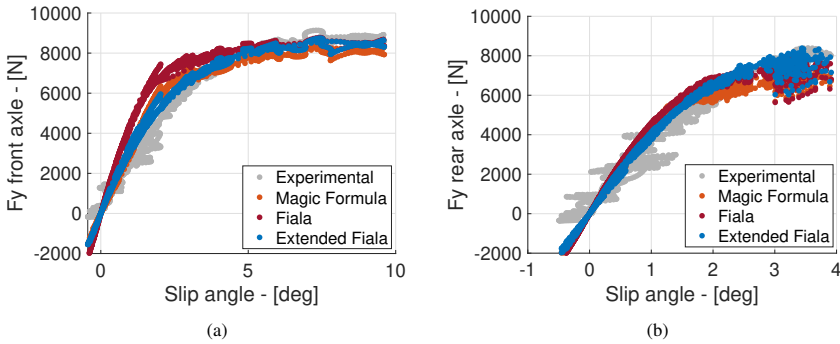


Figure 4.3: Front (a) and rear (b) lateral forces for the classic Fiala, the simplified Magic Formula, the extended (proposed) Fiala tyre model and experimental measurements in a quasi-steady-state circular driving test.

coupling more accurately, which is important for the scope of this work. The reason is that the proposed MPCC uses torque vectoring capabilities, which implies using a longitudinal force coupled with a lateral one.

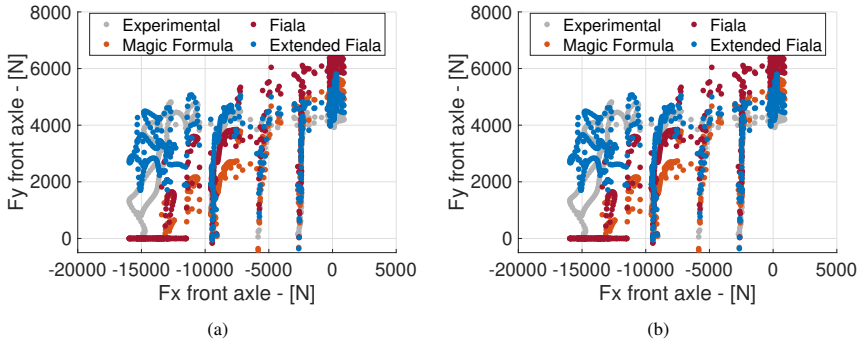


Figure 4.4: Front (a) and rear (b) lateral forces for the classic Fiala, the simplified Magic Formula, the extended (proposed) Fiala tyre model and experimental measurements in four braking-in-a-turn manoeuvres with different longitudinal accelerations.

4

Furthermore, the optimised tyre model is experimentally validated by performing a quasi-steady-state circular driving test and four braking in a turn manoeuvres with different longitudinal accelerations ranging from a minimum of -1.5 m/s^2 to a maximum of -8.5 m/s^2 . Fig. 4.3 shows how the extended Fiala tyre model captures the linear and nonlinear tyre working regions, slightly outperforming the simplified Magic Formula in these conditions. Furthermore, the proposed extended Fiala tyre model outperforms both the classic Fiala and simplified Magic Formula models in four experimentally recorded braking-in-turn manoeuvres, as shown in Figure 4.4. It is particularly noteworthy that as the longitudinal force increases, the model mismatch between experimental data and both the classic Fiala and simplified Magic Formula becomes more pronounced, whereas the extended Fiala tyre model continues to maintain its accuracy.

4.4 MODEL PREDICTIVE CONTOURING CONTROL USING TORQUE VECTORING

THIS section explains how the cost function and the constraints of the proposed MPCC are formulated. Subsection 4.4.1 focuses on describing the MPCC cost function and how it is designed to prioritise obstacle avoidance over path tracking in case of emergency. Subsection 4.4.2 explains how the MPCC constraints are defined to improve safety, taking into account the vehicle actuators' limitations, and avoid redundant torque vectoring utilisation.

4.4.1 COST FUNCTION WITH OBSTACLE AVOIDANCE PRIORITISATION

The proposed MPCC is based on iterative optimisation of a nonlinear cost function ($J = J_{\text{track}} + J_{\text{inp}} + J_{\text{obs}}$) [41], which is responsible for ensuring path tracking (J_{track}), minimising the control effort (J_{inp}) and obstacle avoidance prioritisation (J_{obs}) in case of emergency.

The path tracking properties in the cost function are defined as follows:

$$J_{track} = \sum_{i=1}^N (q_{e_{con}} e_{con,i}^2 + q_{e_{Lag}} e_{Lag,i}^2 + q_{e_{vel}} e_{vel}^2) \quad (4.9)$$

where N is the length of the prediction horizon, e_{con} is the contouring error, e_{Lag} is the lag error, e_{vel} is the velocity error and q_* are the weights of the respective quadratic errors. The path is followed by minimising the contouring (e_{con}) and the lag error (e_{Lag}) [41, 46]. e_{con} represents the vehicle position projection onto the desired trajectory, depending on the vehicle's travelled distance related to the reference line (θ_s). However, contrary to MPC or NMPC controllers [57] based on a Frenet reference system, the θ_s is unavailable for an MPCC based on a Cartesian reference frame. Thus, θ_s is approximated by the vehicle total travelled distance (θ), and the approximation accuracy is ensured by the lag error minimisation, defined as the norm between the two distances. Mathematically, e_{con} and e_{Lag} are defined as follows:

$$\begin{aligned} e_{con} &= \sin(\Psi_t(\theta))(X - X_t(\theta)) - \cos(\Psi_t(\theta))(Y - Y_t(\theta)) \\ e_{Lag} &= -\cos(\Psi_t(\theta))(X - X_t(\theta)) - \sin(\Psi_t(\theta))(Y - Y_t(\theta)) \end{aligned} \quad (4.10)$$

X_t , Y_t , and Ψ_t are the desired longitudinal and lateral positions and heading angle. Despite the added nonlinearities and complexities of e_{con} and e_{Lag} , they allow an approximation of the Frenet reference frame with a Cartesian reference frame, which means that V2O distance is never overestimated [41]. Thus, the MPCC based on a Cartesian frame is more pre-emptive than an NMPC based on a Frenet reference frame in prioritising collision avoidance over path tracking [41]. Thus, the vehicle is more prone to stay inside a safe and stable working area even when it needs to avoid a collision at the limit of handling. Furthermore, the e_{Lag} is equal to zero only when the vehicle follows the reference trajectory perfectly. As soon as the desired velocity is unfeasible for the planned trajectory, the MPCC will modify the desired velocity to minimise the e_{Lag} . The reference velocity is tracked by minimising the quadratic error between the vehicle velocity (v_x) and the desired one (v_{des}). For what concerns the weights, they are firstly empirically tuned to reduce the path tracking error and the vehicle sideslip angle peaks [55, 56] (See Chapter 2). Second, they are fine-tuned using a two-stage Bayesian optimisation [120] (See Appendix A). It is important to highlight that the $q_{e_{vel}}$ is tuned to have a low weight in magnitude because the controller must allow the vehicle to slow down in case of obstacle avoidance prioritisation [38, 41].

The minimisation of the control inputs is ensured in the cost function as follows:

$$\begin{aligned} J_{inp} &= \sum_{i=1}^N (q_{\delta} \dot{\delta}_i^2 + q_{\dot{F}_x} \dot{F}_{x,FL,i}^2 + q_{\dot{F}_x} \dot{F}_{x,FR,i}^2 + \\ &\quad + q_{\dot{F}_x} \dot{F}_{x,RL,i}^2 + q_{\dot{F}_x} \dot{F}_{x,RR,i}^2) \end{aligned} \quad (4.11)$$

where $\dot{\delta}$ is the road wheel angle rate, and $\dot{F}_{x,FL}$, $\dot{F}_{x,FR}$, $\dot{F}_{x,RR}$, and $\dot{F}_{x,RR}$ are the rate of the longitudinal forces applied to each of the vehicle's four wheels. These cost terms are added in order to make the control inputs smooth, and they are considered equal to simplify the tuning.

The motion replanning for obstacle avoidance prioritisation is defined as follows:

$$J_{obs} = \sum_{i=1}^N \left(\sum_{j=1}^{N_{obs}} (q_{e_{V2O}} e_{V2O,j,i}^2) + \sum_{j=1}^{N_{edg}} (q_{e_{V2E}} e_{V2E,j,i}^2) \right) \quad (4.12)$$

where N_{obs} and N_{edg} are the number of obstacles and road edges, e_{V2O} and e_{V2E} are the difference between the V2O and V2E distances and the user defined safety distances between the obstacles ($D_{S_{ft,O}}$) and the road edges ($D_{S_{ft,E}}$). When the vehicle is at a safe distance from obstacles or road edges, it does not interfere with the path tracking properties of the MPCC. On the other hand, the e_{V2O} and e_{V2E} errors allow the MPCC controller to dynamically perform a short trajectory replanning when the vehicle passes close to the obstacles. The obstacle avoidance prioritisation is due to the dynamically varying weights associated with e_{V2O} and e_{V2E} [41]. The weights, here reported only q_{V2O} for compactness, vary as follows:

$$q_{V2O} = \begin{cases} P_k, & \text{if } D_{V2O} < 0 \\ P_k e^{-\frac{2D_{V2O}^2}{D_{S_{ft,O}}^2}}, & \text{elseif } 0 \leq D_{V2O} \leq D_{S_{ft,O}} \\ 0, & \text{otherwise} \end{cases} \quad (4.13)$$

where P_k denotes the upper limit of the achievable value for q_{V2O} . The magnitude of q_{V2O} increases with a Gaussian-shaped curve with the decrease of the V2O distance, and it is zero when V2O is above $D_{S_{ft,O}}$.

4.4.2 CONSTRAINTS

The constraints are designed to accommodate actuator limitations, vehicle stability, and path tracking and to avoid redundant torque vectoring utilisation. The actuators' limitations are applied to δ , $F_{x,FL}$, $F_{x,FR}$, $F_{x,RR}$, and $F_{x,RL}$ and their respective rates. The values implemented are reported in Table 4.3.

The vehicle stability is enforced using the tyre friction circle, as described in the extended Fiala tyre model, and as follows [38, 41]:

$$F_{x,ij} \leq S_f \mu F_{z,ij} \quad (4.14)$$

where S_f is a safety factor that limits the applicable longitudinal force considering the tyre road friction coefficient uncertainty (μ), and the subscripts ij represent the front-rear axle and left-right side.

The following inequality forces the vehicle to stay inside the road boundaries:

$$\left\| \begin{bmatrix} X \\ Y \end{bmatrix} - \begin{bmatrix} X_{cen} \\ Y_{cen} \end{bmatrix} \right\|^2 \leq \left(\frac{W_t}{2} \right)^2 \quad (4.15)$$

where X_{cen} and Y_{cen} are the longitudinal and lateral locations of the track's centre, and W_t is the road width [46].

The MPCC is constrained not to use redundant torque vectoring while driving in a straight to avoid excessive tyre wear and energy consumption as follows:

$$\begin{aligned} |F_{x,FL} - F_{x,FR}| &\leq |F_{z,FL} - F_{z,FR}| T_s \\ |F_{x,RL} - F_{x,RR}| &\leq |F_{z,RL} - F_{z,RR}| T_s \end{aligned} \quad (4.16)$$

Table 4.3: Upper and lower constraints [41, 57], where i stands for front or rear, and j for left or right.

Symbol	Boundaries	Symbol	Boundaries
$\dot{\delta}$	$\pm 90 \text{ deg/s}$	δ	$\pm 18 \text{ deg}$
$\dot{F}_{x,ij}$	$\pm 7200 \text{ N/s}$	$F_{x,ij}$	$\pm 3600 \text{ N}$

where T_s is a user-defined parameter which works as a safety coefficient, allowing a difference in the longitudinal forces higher than the normal load difference between the right and left sides of the vehicle.

4.5 SIMULATION SETUP AND EXPERIMENTAL VALIDATION

THIS section is split into three subsections. Subsection 4.5.1 describes how the proposed MPCC is implemented in a real-time rapid prototyping platform and how it is assessed in high and low friction conditions. Subsection 4.5.2 explains how the high-fidelity and prediction models are validated using experimental data. Subsection 4.5.3 presents how the MPCC robustness is evaluated and how the controller sensitivity to vehicle parameter variations and perception inaccuracies is computed.

4.5.1 HARD REAL-TIME IMPLEMENTATION AND SIMULATION SETUP

The proposed MPCC is tested using the dSPACE SCALEXIO real-time platform, which operates on a multi-core DS6001 processor (2.8 GHz quad-core, 1 GB DDR2 SD RAM). In this setup, the MPCC is allocated to a distinct core separate from the vehicle plant. The prediction model is discretised using the Runge-Kutta 2 method, chosen for its balanced trade-off between computational effort and accuracy [38]. A control integration interval of 0.05 s is used, with 10 integrator nodes and a prediction horizon of 30 steps to ensure the real-time feasibility of the controller implementation. The optimisation problem is solved using the nonlinear interior-point solver available in FORCESPro [130], with the Hessian matrix approximation based on the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm. Rather than initializing the Hessian with an identity matrix, an efficient initial estimate is employed by averaging optimised BFGS matrices from multiple manoeuvres, reducing computational effort without compromising performance. To further decrease computational time, especially on embedded platforms, the variational derivative method is utilised for differentiation method instead of the conventional chain rule [130]. The maximum number of iterations is set to 100 as extensive testing revealed that the solver consistently finds an optimal solution within this limit. The platform demonstrates effective execution of the nonlinear optimisation, achieving an average solving time of 38.9 ms, with minimum and maximum solving times of 30.7 ms and 41.8 ms, respectively. The computational solving time corresponds to the total time required by dSPACE to compute the parameters, such as the reference trajectory sent to the MPCC, solve the optimization problem, and transmit the optimized control inputs to the vehicle plant. Although the maximum number of iterations is limited to 100 to partially control the solving time, it is important to note that there is no formal guarantee of solver convergence within this predetermined time frame [38]. However,

no convergence issues were observed in the performed test. The vehicle plant operates on a separate core at 1000 Hz. It is a high-fidelity representation of a BMW Series 545i vehicle built upon an IPG CarMaker platform, modified to include four in-wheel electric motors. Experimental measurements conducted on a proving ground are used for validation purposes to appropriately select the vehicle parameters. The suspension settings are tuned based on Kinematics & Compliance test rig measurements, while the tyre dynamics are described using the experimentally validated Delft-Tyre 6.2 model. Steering dynamics are accounted for using a second-order transfer function to enhance the fidelity of the vehicle model [151]. The electric motors exhibit faster dynamics than conventional powertrains, and they are modelled with a first-order transfer function and a delay adjusted based on measurements conducted on a powertrain rig provided by the electric motor manufacturer. The electric motor has a time constant of ~ 25 ms, and the initial inverter reaction can be neglected [152, 153]. The control inputs for the MPCC, representing longitudinal forces mapped to the desired torque, are subsequently passed to a low-level control system that tracks these inputs [152, 153]. It is important to highlight that, although the MPCC prediction model does not account for wheel and electric motor dynamics, these dynamics are incorporated in the vehicle plant model. The torque ripple is modelled as a Fourier series as follows [154]:

$$T_{rip} = As_1 \sin(p_1 \omega_r) + Ac_1 \cos(p_1 \omega_r) + As_2 \sin(2p_1 \omega_r) + Ac_2 \cos(2p_1 \omega_r) \quad (4.17)$$

where ω_r is the rotor velocity, p_1 is the least common multiple of the electric motor's stator slots and pole numbers, and As_1 , As_2 , Ac_1 , and Ac_2 are coefficients dependent on the electric motor. The torque ripple (T_{rip}) is added to the electric motor torque calculated by the first-order transfer function.

The proposed MPCC controller with torque vectoring and collision avoidance properties (MPCC TV+CA) is compared with two baselines:

- MPCC CA: it uses the state-of-the-art MPCC controller for path tracking and motion replanning [41], but without torque vectoring capabilities. The goal is to evaluate the benefits of torque vectoring in terms of added stability and agility during evasive manoeuvres, specifically for obstacle avoidance at the limit of handling;
- MPCC TV: it uses the same MPCC with torque vectoring as the proposed approach but excludes the collision avoidance prioritization term J_{obs} , see Eq. 4.12, from the cost function. The goal is to demonstrate that the proposed approach provides advantages not only in torque vectoring but also in enhancing overall vehicle safety by prioritising collision avoidance.

All the controllers are evaluated on a double lane change manoeuvre with two obstacles with high and low friction conditions, respectively 1 and 0.5. Additionally, a split μ condition is analysed to assess the controllers' ability to adapt to varying friction levels, thereby preventing excessive sideslip angles while manoeuvring to avoid obstacles. The split μ condition is designed with a high friction coefficient on the right lane ($\mu = 1$) and a low friction coefficient on the left lane ($\mu = 0.5$). A coarse trajectory function of the road distance and the road curvature is provided to the MPCC, simulating the output of a behavioural planner [41, 136]. The desired velocity is assumed constant along the manoeuvre. The controller must track the reference trajectory and perform an online trajectory replanning




Figure 4.5: Experimental model validation of the vehicle characteristics for a skidpad manoeuvre with 40 m radius: (Fig. 4.5a) understeer gradient, (Fig. 4.5b) yaw rate and (Fig. 4.5c) sideslip angle.




Figure 4.6: Experimental model validation of the vehicle characteristics for a brake-in-turn manoeuvre with a $|8| \text{ m/s}^2$ longitudinal and $|4| \text{ m/s}^2$ lateral acceleration: (Fig. 4.6a) G-G diagram, 4.6b) yaw rate and (Fig. 4.6c) sideslip angle.




Figure 4.7: Experimental model validation of the vehicle characteristics for a double lane change at 70 km/h: (Fig. 4.7a) lateral acceleration, (Fig. 4.7b) yaw rate and (Fig. 4.7c) sideslip angle.

when the vehicle passes dangerously close to one of the obstacles. This can happen due to different handling limits of the behaviour planner and the path tracking controller, as well as the intervention of the stability controller. The scenario contains two obstacles to demonstrate that the short trajectory replanning around the first obstacle does not interfere with the vehicle's capacity to avoid the subsequent obstacle. The key performance indicators (KPIs) are the V2O and the V2E distances, which must be higher than zero to avoid a collision and preferably higher than 0.5 m to avoid a near-miss collision. A test is considered failed when the vehicle collides with an obstacle or goes outside the road edges.

4.5.2 EXPERIMENTAL MODEL VALIDATION

The computationally efficient vehicle model with the extended Fiala tyre used within the MPCC and the high-fidelity vehicle model with Delft Tyre model are validated with experimental data collected on a proving ground. The test platform was equipped with the conventional Inertial Measurement Unit (IMU), Kistler wheel force transducers, and SKF

intelligent bearings installed on each wheel. Additionally, it had a dual antenna GNSS system and a Corrsys-Datron non-contact optical sensor dedicated to measuring the sideslip angle, boasting a measurement accuracy of approximately ± 0.2 deg. Interconnection of all equipment was facilitated through the Controller Area Network (CAN) interface, with the sampling rate configured to operate at 100 Hz.

The accuracy of the models during a 40 m radius skidpad test, a brake-in-turn manoeuvre with longitudinal and lateral accelerations of 8 m/s^2 and 4 m/s^2 respectively, and a double lane change at 70 km/h are depicted in Figs. 4.5, 4.6, 4.7. In the quasi-steady-state behaviour, both the prediction models and the high-fidelity model demonstrate high accuracy across both linear and nonlinear regions of tyre operation. The only minor discrepancy observed is an overestimation of the sideslip angle by the prediction model, remaining below 1 deg. This overestimation is more pronounced when the Magic Formula tyre model is used, particularly when the absolute value of lateral acceleration exceeds 4 m/s^2 . The high-fidelity model exhibits the highest alignment with experimental data during the brake-in-turn manoeuvre. The prediction model based on the extended Fiala tyre model also achieves high accuracy in predicting yaw rate and vehicle sideslip angle, albeit with a slight overestimation of lateral acceleration. In contrast, the prediction model using the Magic Formula tends to overestimate lateral acceleration and shows a differing trend, underestimating the vehicle's yaw rate during combined longitudinal and lateral forces due to a greater mismatch in the tyre model under such conditions. Given that torque vectoring inputs increase the duration of coupled longitudinal and lateral forces, these results emphasize the importance of the proposed extended Fiala tyre model. Regarding transient vehicle behaviour, all models can capture even the lateral acceleration and the yaw rate peaks recorded with the experimental vehicle. The most significant difference is noticeable in Fig. 4.7c, which shows that the high-fidelity model overestimates the measured vehicle sideslip angle at most 0.5 deg, while the prediction model based on the Magic Formula slightly underestimates the peaks of the sideslip angle by a similar margin. However, this minor discrepancy does not interfere with vehicle model validation. The superior accuracy of the prediction model based on the proposed extended Fiala tyre model, compared to the one based on the Magic Formula, underscores the importance of the proposed tyre model. Additionally, it validates the design choice to neglect wheel dynamics in the prediction model, which facilitates increased discretisation time and reduces the computational time of the MPCC, making the real-time implementation feasible on an embedded platform.

4.5.3 ROBUSTNESS AND SENSITIVITY ANALYSIS SETUP

The collision avoidance performance and robustness of the proposed controller are evaluated by considering variations in vehicle parameters and perception inaccuracies. Two different Monte Carlo analyses are conducted to simplify the evaluation of the results, focusing on a double lane change manoeuvre with two obstacles and high friction conditions ($\mu = 1$). The first only assesses the performance by altering the vehicle and tyre parameters, and the second evaluates the effects of obstacle and vehicle localisation inaccuracies. The first Monte Carlo analysis considers the variation of the additional mass (Δm) and corresponding moment of inertia (I_{zz}), electric motor time delay (τ_{em}), and the scaling factors of the front and rear cornering stiffness (K_{yF} , K_{yR}), lateral peak friction coefficients (μ_{yF} , μ_{yR}), longitudinal stiffness shape factors (K_{x_F} , K_{x_R}) and tyre relaxation lengths of the vehicle plant (Rel),




Figure 4.8: Gaussian distribution of additional mass Δm (Fig. 4.8a); normal multivariate distribution of the scaling factors for the cornering stiffness, lateral peak friction coefficients, and longitudinal stiffness shape factors (Fig. 4.8b).

4

leaving the vehicle prediction model unchanged. The additional mass variation simulates the change in the vehicle mass due to different amounts of fuel, number of passengers and pieces of luggage. Thus, the nominal additional mass correspondent to 160 kg is varied by $\pm 10\%$ of the total vehicle mass according to a Gaussian distribution [51], and its changes are associated with the corresponding variation of I_{zz} . The Δm distribution is visible in Fig. 4.8a, and the same approach is used for all the parameters varied according to a Gaussian distribution. A variation of $\pm 15\%$ is associated with the electric motor delay, and it simulates the effect of different load conditions, power supply and environmental factors [155]. The tyre cornering stiffnesses, lateral peak friction coefficients, longitudinal stiffness shape factors and tyre relaxation lengths of the vehicle are modified by $\pm 15\%$ over their nominal value to simulate the effect of tyre wear, temperature variations, suspension kinematics, bushing wear and other effects [51, 156]. Considering that tyre wear is not always homogeneous in the four wheels and that the wear in the suspension bushings can affect the characteristics of the front and rear axles differently [156], the parameters variations of the front and rear axles are varied according to a multivariate normal distribution, see Fig. 4.8b. This can model the different front and rear axle variations, but it avoids the unrealistic scenario of a simultaneous increased front cornering stiffness equal to 15 % and a -15% reduction in the rear axle cornering stiffness. A similar concept is applied for scaling the μ_{yF} , μ_{yR} and K_{xF} , K_{xR} . Specifically, the prediction model can accommodate a maximum estimation error of $\pm 15\%$ for the friction coefficient, distributed according to a multivariate normal distribution, as depicted in Fig. 4.8b. The same variation is considered for the front and rear axles regarding the relaxation length. The first Monte Carlo analysis consists of 11 000 different simulations generated by a random combination of the varied parameters.

The second Monte Carlo analysis considers the variation of obstacle localisation error in the longitudinal ($\sigma_{Obs,X}$) and lateral direction ($\sigma_{Obs,Y}$) [157, 158], vehicle localisation error in the longitudinal ($\sigma_{Veh,X}$) and lateral direction ($\sigma_{Veh,Y}$) [159, 160] and the percentage of false negative obstacle detections [161]. The obstacle localisation inaccuracies in X and Y directions are modelled as random Gaussian noises in which the standard deviations are varied according to a multivariate normal distribution, see Fig. 4.9a. The longitudinal and

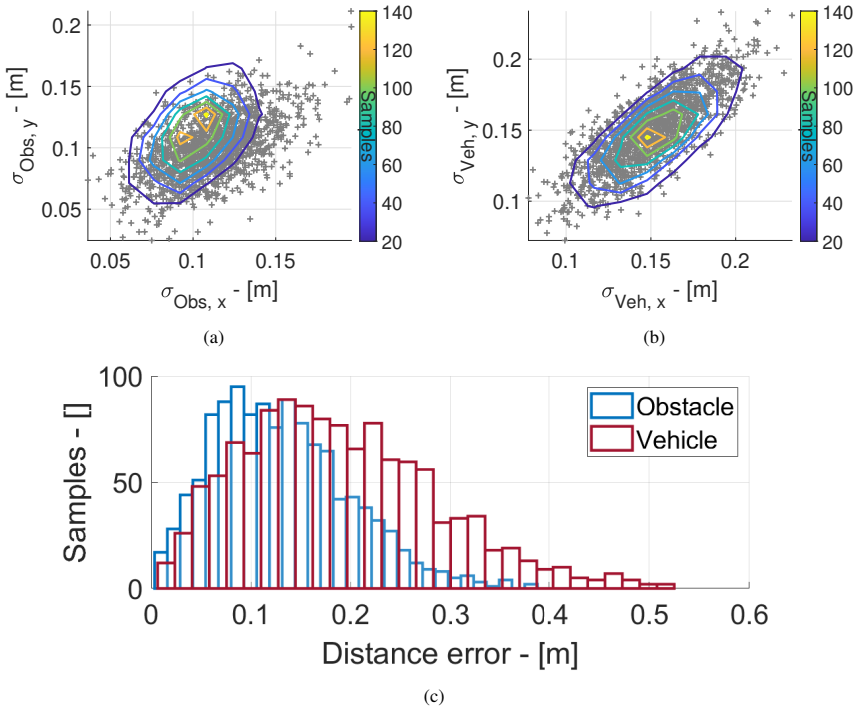


Figure 4.9: Multivariate distribution of the standard deviations for lateral and longitudinal errors of obstacle (Fig. 4.9a) and vehicle (Fig. 4.9b) localisation; (Fig. 4.9c) distributions of the vehicle and obstacle localisation distance errors generated by the nominal standard deviations.

lateral inaccuracies are not always considered equal because the performance of the obstacle detection algorithm is influenced by the obstacle aspect ratio bias and camera perspective, which can influence in which direction the perception algorithm is more accurate [157]. Fig. 4.9c shows that the most probable standard deviation combination creates a mean distance error equal to 0.13 m for vehicle obstacle localisation, corresponding to the mean average translation error in vehicle detection for the leaderboard works using the nuScenes dataset [162]. The same approach is applied to vehicle localisation which errors in the longitudinal and lateral directions are treated as random Gaussian noises. The noises are varied according to a multivariate normal distribution; see Fig. 4.9b. In this situation, the nominal standard deviations create a mean distance error equal to 0.19 m, see Fig. 4.9c. The mean value is chosen comparing the performance of different localisation techniques for automated vehicles [163]. Furthermore, the analysis considers the percentage of false negatives modelled through the parameter (λ) of a Poisson distribution, which defines the number of times a random event occurs given a period. It is varied according to a Gaussian distribution with a mean equal to 0.5 and a standard deviation of 0.15. Thus, it corresponds to a probability equal to 16.3 % of having a single false negative during the simulation, 1.7 % of two and 0.2 % of three, corresponding to the performance of the leaderboard works [162]. It is essential to specify that we only model the probability that this event occurs,

but where and for how many times it happens is not regulated in the simulation scenario, but it depends on the Poisson distribution. This Monte Carlo analysis consists of 3000 simulations generated by a random combination of the varied parameters. The lower number of simulations compared to the first Monte Carlo analysis is due to the lower number of assessed parameters.

The controller performance is evaluated by computing the Collision Rate (CR), defined as the percentage of simulations in which the vehicle collides with an obstacle or goes out from the road edges. This is measured through the minimum distance between the V2O and V2E, represented by the symbol (mVD). Furthermore, the robustness is evaluated through the Near Miss Rate (NMR), which measures the percentage of simulations in which the vehicle has a mVD lower than 0.5 m.

The mVD is also used to perform a sensitivity analysis and mathematically evaluate the standing of which vehicle parameters and perception inaccuracies influence the controller performance most. The proposed sensitivity analysis is based on the PAWN approach [164], and it computes the density-based sensitivity (SI) indices. The PAWN approach assesses the impact of input variables, e.g. vehicle parameters, on the cumulative probability distribution of the output, e.g. mVD , considering the entire distribution rather than just specific moments. Thus, it is less sensitive to skewed or multi-modal distribution than the classical Variance-based sensitivity analysis [165]. The SI indices are dimensionless absolute measures ranging between zero and one. Thus, the non-influential inputs will have a value of zero, and, vice versa, the most influential ones will have a value closer to one.

4

4.6 RESULTS

THIS section is split into three subsections. Subsection 4.6.1 demonstrates the improved performance of the proposed MPCC over the baselines in high and low friction conditions with the high-fidelity model based on nominal parameters. Subsection 4.6.2 proves the robustness of the proposed approach to varying vehicle parameters and perception inaccuracies in the high-fidelity vehicle model. Subsection 4.6.3 presents the sensitivity analysis results aimed at evaluating the parameters that most significantly influence the controller's performance.

4.6.1 CONTROLLER PERFORMANCE

Fig. 4.10a shows the trajectories of the high-fidelity vehicle model with nominal parameters attained by the three different controllers. The proposed MPCC TV+CA is the only controller to drive through the double lane change manoeuvre without colliding with either of the two obstacles. The MPCC TV is unaware of the obstacles' location, so its path tracking and vehicle stability objectives lead the vehicle to collide with the first obstacle located at 99 m. On the other hand, the MPCC CA can successfully avoid the first obstacle, replanning the initial trajectory in a very similar way to the MPCC TV+CA. However, it cannot avoid a collision with the second obstacle due to a loss of stability despite harsh braking from ~105 m to ~115 m. Vice versa, the MPCC TV+CA stabilises the vehicle between 100 m and 120 m, reducing the vehicle sideslip angle from a peak higher than 15 deg to a peak equal to 7.5 deg. This is possible thanks to the extra yaw moment, up to ~2000 Nm, at the vehicle CoG generated by torque vectoring. Furthermore, it is relevant to notice that thanks to the

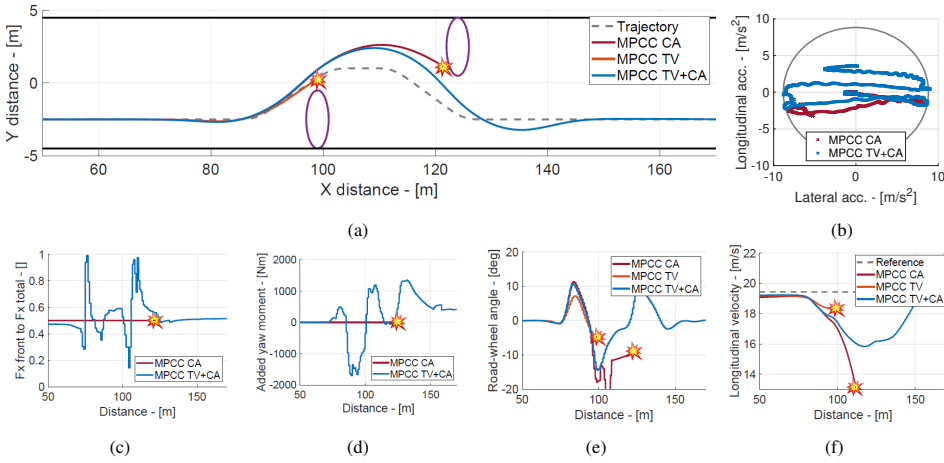


Figure 4.10: States and control inputs of the controllers in a double lane change manoeuvre with high friction conditions ($\mu = 1$).

lower sideslip angle, the vehicle controller MPCC TV+CA can drive through the double lane change at a higher speed than the MPCC CA, without performing any harsh braking after avoiding the first obstacle. Fig. 4.10f shows that the minimum speed of the vehicle driven by the MPCC TV+CA is 16 m/s, while the MPCC CA cannot avoid the second obstacle despite reducing the speed to 13 m/s. The MPCC TV+CA also optimises the front and rear longitudinal force repartition (Fig. 4.10c). The controller moves the brake repartition to the front axle during the hard braking at 90 m. However, the front-rear ratio is restored to 50 % or even moved to the rear axle, when the vehicle enters the corner, and the front axle has a high road wheel angle (Fig. 4.10e). It is relevant to notice that in both the front and rear axle, the added yaw moment generated by the torque vectoring is substantial in magnitude only during the manoeuvre and at its exit to stabilise the vehicle. On the contrary, the added inequality constraints prevent the MPCC TV+CA from using torque vectoring when the vehicle is driving straight. Regarding the road wheel angle computed by the MPCC CA, it is visible how it already exceeds the constraints at ~ 120 m because the controller cannot converge to an optimal solution. Fig. 4.10b shows that both MPCC controllers with CA capabilities reach the maximum lateral acceleration, which can be generated with the available road friction coefficient. On the contrary, the maximum braking capability is not fully exploited by any controller. This work demonstrates that the TV can be integrated into an MPCC with CA avoidance prioritisation and that the TV is essential to stabilise the vehicle while avoiding obstacles at the limit of handling. However, the strong coupling between longitudinal and lateral dynamics brings complexity to the MPCC. For this reason, it was essential to improve the accuracy of the prediction model, e.g. including an extended version of the Fiala tyre model.

The proposed controller and the two baselines are also tested in a double lane change manoeuvre with a low friction condition ($\mu = 0.5$) and a 55 km/h desired and initial velocity, corresponding to a heavy rain scenario. Fig. 4.11a shows the three vehicle trajectories, and it is visible how only the newly developed controller can successfully avoid the two obstacles.




Figure 4.11: Vehicle trajectories and G-G diagram in a double lane change manoeuvre with low friction conditions ($\mu = 0.5$).

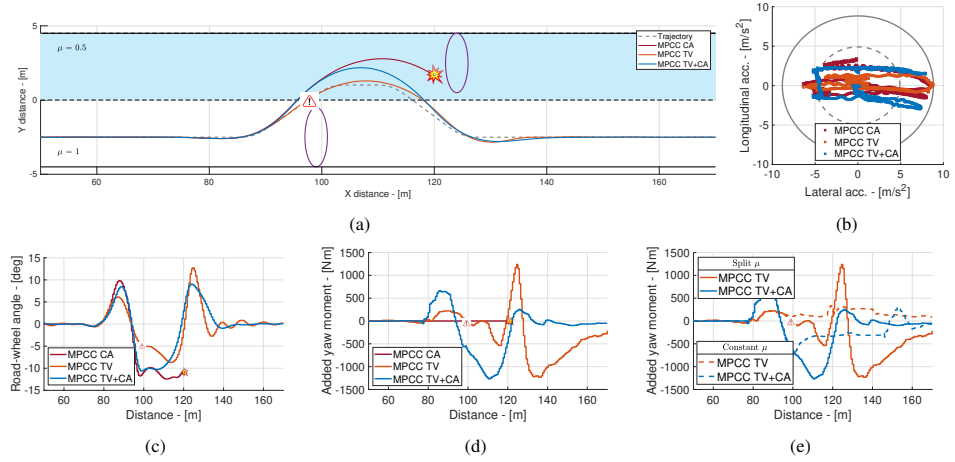


Figure 4.12: States and control inputs of the controllers in a double lane change manoeuvre with split μ conditions. The white area has $\mu = 1$, and the light blue area has $\mu = 0.5$.

This further proves the robustness of the proposed controller to different friction conditions. Fig. 4.11b shows how this is achieved while the proposed controller drives the vehicle at the handling limits. To evaluate the controllers' capacity to react to a sudden change in the friction coefficient, we tested the proposed controllers and baselines during a double lane change manoeuvre under a split μ condition with a desired initial velocity of 55 km/h. This manoeuvre featured a high-friction surface on the right lane ($\mu = 1$) and a low-friction surface on the left lane ($\mu = 0.5$), as shown in Fig. 4.12a. An analysis of the trajectories, shown in Fig. 4.12a, indicates that only the proposed controller and the baseline MPCC TV successfully avoided the two obstacles, guiding the vehicle at the limit of handling (see Fig. 4.12b). However, while the baseline MPCC TV approaches the unsafe area near the first obstacle, generating a near-miss event, the proposed approach maintains a safer distance, highlighting the importance of collision avoidance prioritisation. The baseline without collision avoidance prioritization successfully navigates around the first obstacle but fails to avoid the second one due to the sudden drop in road friction. The friction drop, together with the saturation of the steering actuator, leads to an instability which can only be mitigated by controllers that incorporate torque vectoring. The additional yaw moment applied at the




Figure 4.13: Analysis of the mVD at the variation of desired and initial vehicle velocity in a double lane change.

vehicle's centre of gravity (Fig. 4.12d) stabilizes the vehicle during the manoeuvre, allowing it to avoid obstacles effectively. Conversely, the MPCC CA attempts to manage the sudden change in friction primarily through adjustments in the road-wheel angle (Fig. 4.12c) and additional braking force. A particularly notable aspect is the comparison of the additional yaw moment generated by controllers with torque vectoring in the split μ condition versus scenarios with a constant friction coefficient. Fig. 4.12e demonstrates that both controllers with torque vectoring generate additional yaw moments to stabilize the vehicle under the split μ condition, which was unnecessary during the high-friction scenario. This confirms the ability of the proposed controllers to effectively react to sudden changes in the tyre-road friction condition.

4.6.2 ROBUSTNESS ANALYSIS

The performance of all the controllers is also evaluated in the same double lane change with high friction conditions ($\mu = 1$), varying the initial and desired vehicle velocities to determine the point at which each controller begins to fail. Fig. 4.13 shows that the proposed controller effectively avoids the obstacles in the simulated scenario with velocities up to 73 km/h, a 7.5 % improvement over the MPCC CA. In contrast, the MPCC TV successfully executes the manoeuvre only at velocities up to 60 km/h. This analysis further highlights the advantages of incorporating collision avoidance prioritisation and torque vectoring.

The robustness of the proposed controller MPCC TV+CA and the baseline MPCC CA, as explained in section 4.5.2, are compared at varying velocities in Fig. 4.14. The controller without collision avoidance is not analysed in this section because it already shows a low performance in nominal conditions; see Fig. 4.13. Both the analysed controllers show an exponential growth of the collision rate with the increase of vehicle velocity; see Fig. 4.14a. However, the proposed controller is substantially more robust to vehicle parameter variations for velocities higher than 60 km/h. On the other hand, they show a similar CR, corresponding to $\sim 5\%$ for lower velocities because the vehicle behaves less nonlinearly and the scenario allows a higher margin of error. Thus, the controller task is simplified, and both can achieve an optimum level of robustness. Vice versa, vehicle dynamics become highly nonlinear at high velocities, and the scenario becomes more challenging. In this situation, both controllers reduce their robustness. However, the MPCC TV+CA can achieve adequate robustness thanks to the extra added yaw moment, which increases the vehicle responsiveness.

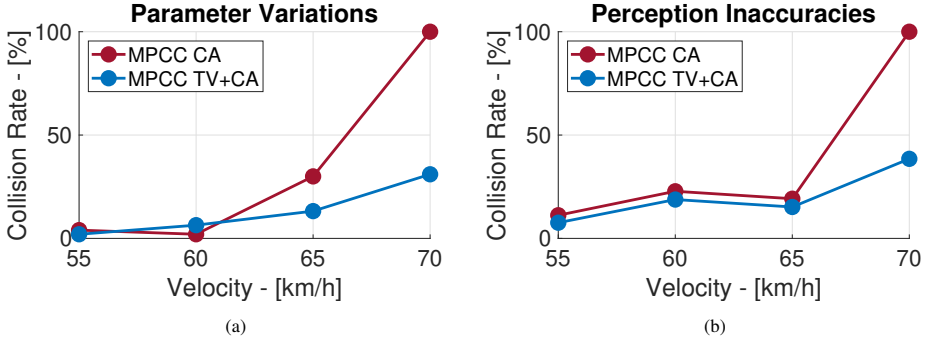


Figure 4.14: Robustness evaluation to vehicle parameters (Fig. 4.14a) and perception inaccuracies (Fig. 4.14b) of the analysed controllers at varying velocities.

4

Fig. 4.14b shows the controllers' robustness to vehicle perception inaccuracies at different velocities. Similar to the above case, the proposed approach shows a higher robustness than the controller without torque vectoring. However, it is noticeable that the MPCC TV+CA and the MPCC CA show similar robustness to velocities up to 65 km/h. A possible explanation is that the most significant cause of failure up to this velocity is the obstacle lateral error noise, which leads to an impossible scenario for both MPCCs. Thus, the major difference between MPCC TV+CA and MPCC TV for perception inaccuracies robustness happens at a velocity equal to 70 km/h, where the torque vectoring allows a higher accuracy even in nominal conditions.

Considering the scope of the paper, which is to assess the performance of the proposed controller at the limit of handling, we will focus on the robustness analysis of the MPCC TV+CA during the double lane change at 70 km/h. Table 4.4 presents the results of the controller robustness analysis against variations in vehicle parameters. The proposed MPCC TV+CA is the only one with a CR lower than 100 % among the controllers evaluated. Although the vehicle parameters variations impact the proposed controller's performance, a CR equal to 31 % and an NMR of 35 % demonstrate the MPCC TV+CA robustness. The slight difference between NMR and CR suggests that the controller can avoid obstacles only a few times when it enters the danger zone. However, the vehicle collides with a velocity on average 29 % lower than its reference velocity. Thus, it can still mitigate the consequences of collisions. Particularly interesting are the collision patterns at the variation of the vehicle parameters; see Table 4.5. Approximately 50 % of collisions occur due to the vehicle driving off-road on the right side. This indicates that the controller can successfully drive the vehicle around the obstacles, but it compromises its ability to rejoin the original trajectory. This behaviour may stem from increased model mismatch caused by parameter variations, impacting the controller's efficacy. Further proof is the 23 % of collision with the second obstacle, which indicates how the vehicle can safely avoid the first collision but cannot avoid the second one. Although the collision rate is at 31 %, as depicted in Fig. 4.15a, it is evident that the trend in collision rates is exponentially decreasing with the occurrence of vehicle parameter configurations with higher probability. This observation underscores the robustness of the proposed controller to vehicle parameter variations, as it

Table 4.4: Monte Carlo analysis results related to vehicle parameter variations and perception inaccuracies for the proposed controller.

Configuration	Collision Rate [%]	Near Miss Rate [%]
Parameter Variations	31.20	34.80
Perception Inaccuracies	38.50	49.50

Table 4.5: Different collision patterns observed in the Monte Carlo analysis for the proposed controller.

	Obstacle 1 [%]		Obstacle 2 [%]		Edge Left [%]		Edge Right [%]	
	CR	NMR	CR	NMR	CR	NMR	CR	NMR
Parameter Variations	3.21	36.49	22.76	38.51	23.72	26.72	50.32	51.72
Perception Inaccuracies	54.67	94.48	1.38	4.17	4.67	9.82	39.27	30.55

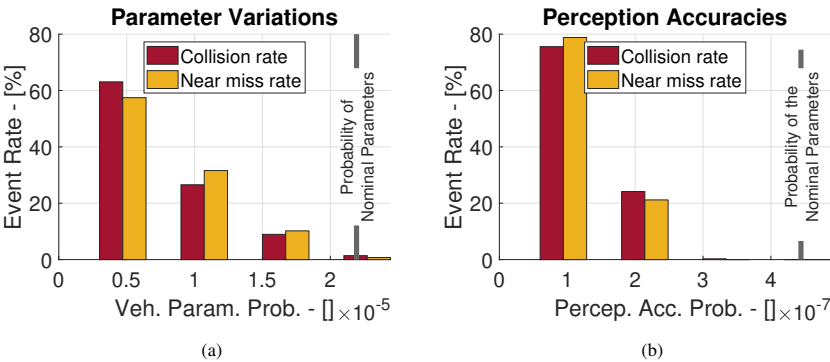


Figure 4.15: Collision and Near Miss Rates change to the probabilities variation of parameter configurations for parameter variations (Fig. 4.15a) and perception accuracies (Fig. 4.15b)

highlights that the 90 % of collisions occurs for configurations where parameters are situated in the tails of the Gaussian distribution, so less probable configurations in a real-life scenario. The configurations with the highest probability have only ~1 % of collisions.

Table 4.4 also presents the results of the controller robustness to perception inaccuracies. The proposed MPCC TV+CA demonstrates robustness comparable to that under vehicle parameter variations, achieving a CR of 38.5 % and an NMR of 49.5 %. This underscores the robustness of the proposed approach not only against vehicle parameter variations but also against uncertainties in vehicle and obstacle localisation. In contrast to the previous scenario, when the controller fails, it leaves the roadway only 44 % of the time, see Table

4.5. Particularly noteworthy is that in 55 % of collision cases, the first obstacle encountered results in a collision. This suggests that the MPCC TV+CA, influenced by uncertainties in the vehicle and obstacle localisation, can fail to promptly replan its trajectory around obstacles. In comparison, under vehicle parameter variations, the MPCC TV+CA primarily fails after successfully avoiding the first and second obstacles. An explanation for this could be the increased error in obstacle localisation, which causes the MPCC TV+CA to be unaware of its proximity to obstacles, leading to hazardous situations. Furthermore, erroneous obstacle localisation can mislead the MPCC TV+CA into believing that the road ahead is obstructed, preventing the vehicle from passing safely with adequate distance.

Figure 4.15b shows the trend in collision and near miss rates as the probability of perception inaccuracies decreases. As expected, the 75 % of collisions happens when the proposed controller is affected by the lowest probability of perception accuracies. Once again, the exponential decrease in the number of collisions and near-miss rate events with the increase of perception accuracy probability highlights the robustness of the proposed controller.

4

4.6.3 SENSITIVITY ANALYSIS

Fig. 4.16a shows the SI indices for the PAWN sensitivity analysis evaluating the vehicle parameter variations. The box plot area represents the uncertainty of these indices, which has been minimised due to the extensive number of simulations conducted. It is visible that the nine analysed parameters can be split into three influential groups. The most influential group consists of only one parameter: the lateral peak friction coefficient for the rear axle (μ_{yR}). This suggests that μ_{yR} influences the performance of the proposed controller the most, which can be explained by the fact that a decrease in μ_{yR} can lead the vehicle to an oversteer behaviour, which is particularly difficult to control and stabilise even for a human driver [166]. The second group is composed of Δm , the front and rear axle cornering stiffness (K_{yF} and K_{yR}) and the front axle lateral peak friction coefficient. Once again, the parameters K_{yF} , K_{yR} and μ_{yF} can strongly vary the vehicle behaviour, changing it from understeer to oversteer and vice versa. A similar conclusion can also be stated for Δm because it is not placed in the vehicle CoG, so it moves the CoG location. The least influential parameters are the front and rear axle longitudinal stiffness shape factors (K_{x_F} and K_{x_R}), the tyre relaxation length (Rel) and the electric motor delay (τ). The relative non-influential effect on the proposed controller performance can be explained by the fact that the analysis scenario involves primarily the vehicle lateral dynamics and not the longitudinal one, which is most affected by the K_{x_F} and K_{x_R} . The analysis of the effect of Rel and τ suggests that the controller can deal adequately with an additional delay, even though delays can significantly contribute to the onset of limit cycles [131–133]. This capability is further evidenced by the successful implementation of a similar MPCC controller for a vehicle with hydraulic brakes, which exhibit even slower dynamics (See Chapter 3) [41].

Fig. 4.16b shows the SI indices for vehicle perception uncertainties. The analysis reveals that the five examined parameters can be categorised into three distinct levels of influence. The most influential one consists of the lateral obstacle localisation ($\sigma_{Obs,Y}$) error which strongly limits the controller's capacity to avoid obstacles. A possible explanation is that the increased $\sigma_{Obs,Y}$ can lead to an impossible scenario for the MPCC. For instance, a further step of the first or second obstacle towards the road centre can make collision avoidance




Figure 4.16: Sensitivity analysis results for vehicle parameter variations (Fig. 4.16a) and vehicle perception inaccuracies (Fig. 4.16b).

unfeasible. The second group consists of $\sigma_{Obs,X}$. Interestingly, controller failure can occur not only due to the magnitude of noise but, more critically, due to the direction and timing of misleading obstacle localisation. This means that vehicle obstacle localisation errors are the most influencing factors for controller performance, and they must be carefully considered during the controller design. The third group include the vehicle localisation errors and the false negative probability. It is expected that $\sigma_{Veh,X}$, $\sigma_{Veh,Y}$ could have a similar effect on the MPCC performance than $\sigma_{Obs,X}$ and $\sigma_{Obs,Y}$. However, the MPCC feedback on the vehicle localisation can partially compensate for their effect. The relatively lower influence of the false negative percentage (coefficient λ) can be attributed to its probabilistic nature. Since we cannot control when a false negative occurs within a scenario, it may happen when the vehicle is too far from the obstacle to be influenced by it. However, controlling the timing of false negatives is unrealistic and, thus, has not been analysed in this study.

4.7 CONCLUSION

THIS paper presents a novel approach to vehicle obstacle avoidance. It is based on a nonlinear Model Predictive Contouring Control, which employs torque vectoring capabilities to stabilise the vehicle, performing evasive manoeuvres at the limit of handling. The proposed controller with torque vectoring and collision avoidance successfully avoids the two obstacles in the Double Lane Change manoeuvre with high and low friction conditions, and it has proven to work on a rapid prototyping platform in real-time with an average solving time of 38.9 ms. The first baseline controller without collision avoidance collides with the

4 first obstacle due to its lack of motion replanning capability. The second baseline controller without torque vectoring cannot stabilise the vehicle after avoiding the first obstacle, and it collides with the second one. The proposed controller produces an extra yaw moment up to ~ 2000 Nm in magnitude, increasing the vehicle's lateral agility to avoid the obstacles and reduce the vehicle sideslip angle peak down to 7.5 deg rather than the 16 deg for the second baseline controller without torque vectoring. However, the proposed controller requires a more accurate prediction model, so an extended Fiala tyre model has been developed to describe the tyre behaviour in combined slip conditions. Using extended Monte Carlo analyses, the proposed controller shows robustness to vehicle parameter variations and perception inaccuracies with a collision rate of respectively 31.20 % and 38.50 %. As further proof, the percentage of crashes occurring with simulated values near the nominal vehicle parameters and the most probable vehicle perception inaccuracies is below 1.5 %. Furthermore, the results of sensitivity analysis are statistically assessed to evaluate the sensitivity indices to each analysed parameter. The lateral peak friction coefficient of the rear axle and the obstacle localisation error in the lateral direction have a higher effect on the controller performance. Future works involve the experimental validation of the proposed controller and the enhancement of the algorithm to better deal with the most influential parameters.

ACKNOWLEDGEMENT

The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484). European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie actions, under grant agreement Nr. 872907.

5

A LEARNING-BASED MODEL PREDICTIVE CONTOURING CONTROL FOR VEHICLE EVASIVE MANOEUVRES

5

*Uncertainty is the only certainty there is,
and knowing how to live with insecurity is the only security.*

A Mathematician Plays the Market, John Allen Paulos

ABSTRACT

This paper presents a novel Learning-based Model Predictive Contouring Control (L-MPCC) algorithm for evasive manoeuvres at the limit of handling. The algorithm uses the Student-t Process (STP) to minimise model mismatches and uncertainties online. The proposed STP captures the mismatches between the prediction model and the measured lateral tyre forces and yaw rate. The mismatches correspond to the posterior means provided to the prediction model to improve its accuracy. Simultaneously, the posterior covariances are propagated to the vehicle lateral velocity and yaw rate along the prediction horizon. The STP posterior covariance directly depends on the variance of observed data, so its variance is more significant when the online measurements differ from the recorded ones in the training set and smaller in the opposite case. Thus, these covariances can be utilised in the L-MPCC's cost function to minimise the vehicle state uncertainties. In a high-fidelity simulation environment, we demonstrate that the proposed L-MPCC can successfully avoid obstacles, keeping the vehicle stable while driving a double-lane change manoeuvre at a higher velocity than an MPCC without STP. Furthermore, the proposed controller yields a significantly lower peak sideslip angle, improving the vehicle's manoeuvrability compared to an L-MPCC with a Gaussian Process.

5.1 INTRODUCTION

A crucial safety element of automated driving is to prove the capacity to avoid obstacles at the limit of handling. A common solution is based on Nonlinear Model Predictive Control (NMPC), which optimises the steering angle and the longitudinal force of a vehicle. However, in such scenarios, where longitudinal and lateral dynamics are coupled, the uncertainties and inaccuracies due to the tyre's nonlinear behaviour pose a particularly challenging problem [41, 134]. Therefore, we focus on developing a Student-t Process (STP) combined with a Model Predictive Contouring Control (MPCC), which improves the prediction model, reducing the tyre model mismatches and minimising the vehicle lateral state uncertainties. An L-MPCC based on a sparse Gaussian Process (GP) has recently been proposed for lap-time optimisation [47]. It leverages the capacity of the GP to predict the mismatches between the prediction model and the measured vehicle states: longitudinal and lateral velocity and yaw rate. Furthermore, the vehicle is constrained inside the track, tightening the track boundaries with the vehicle position uncertainty. The latter is first computed as the GP's posterior covariance, then open-loop propagated along the prediction horizon using successive linearisation similar to an Extended Kalman Filter (EKF) [47]. However, the propagated position uncertainty can increase exponentially over the prediction horizon, strongly limiting or eliminating the allowable driving area. This results in a very conservative controller. A possible alternative is to consider the uncertainty in the vehicle velocity states rather than in the vehicle position [129]. For instance, the NMPC cost function can be extended by the vehicle lateral velocity and yaw rate variance. Thus, the latter can be minimised to reduce the operating time in the unstable region of the vehicle. However, these variances are computed and propagated by linearising the prediction model with a constant tyre model uncertainty optimised offline. This simplification does not consider the remarkably different accuracy of the tyre model in different operating regions, and therefore, it does not reduce the prediction model mismatches in all regions.

We propose an L-MPCC based on an STP, which predicts the mismatches between the prediction model and the vehicle yaw rate and the lateral tyre forces measured by intelligent (force-sensing) bearings [29]. The proposed controller minimises the vehicle's operating time at unstable working points, thanks to the STP posterior covariance of the tyre forces. The latter ones are used to compute and propagate the vehicle lateral state uncertainties along the prediction horizon.

The contributions of this paper are threefold. The first is the development of the L-MPCC based on an STP, which directly reduces the prediction model mismatches in the tyre model rather than in the vehicle velocity states typically used [47]. This results in successfully performing evasive manoeuvres at 8.5 % higher velocity than the current state-of-the-art MPCC [41]. The second contribution is related to the STP, which improves the outlier resistance of the state-of-the-art GP. Furthermore, the STP posterior covariance depends on the observed measurements, providing a higher variance than a GP for operating points different than in the training set and a lower one in the opposite situation [48]. The third contribution is improving the vehicle stability by reducing the sideslip angle peak of a 76 % during an evasive manoeuvre, thanks to the reduction of the model mismatches and the minimisation of the vehicle lateral state uncertainties. Thus, the controller decreases the time the vehicle spends in operating points close to the vehicle's handling limits.

5.2 LEARNING-BASED MODEL PREDICTIVE CONTOURING CONTROL

THE proposed L-MPCC controller is based on an MPCC for obstacle avoidance at the limit of handling [41]. The prediction model is a nonlinear single-track vehicle model. The Cartesian reference system describes the vehicle kinematics, as the contouring formulation requires. The vehicle dynamics, the longitudinal (v_x) and lateral (v_y) velocity, and the vehicle yaw rate (r) are described as follows:

$$\begin{cases} \dot{v}_x = \frac{(F_{x,F} \cos(\delta) - (F_{y,F} + \Delta F_{y,F}) \sin(\delta) + F_{x,R} - F_{drag})}{m} + (r + \Delta r) v_y \\ \dot{v}_y = \frac{(F_{x,F} \sin(\delta) + (F_{y,F} + \Delta F_{y,F}) \cos(\delta) + F_{y,R} + \Delta F_{y,R})}{m} - (r + \Delta r) v_x \\ \dot{r} = \frac{((F_{y,F} + \Delta F_{y,F}) \cos(\delta) l_f - (F_{y,R} + \Delta F_{y,R}) l_r + F_{x,F} \sin(\delta) l_f)}{I_{zz}} \end{cases} \quad (5.1)$$

where the road-wheel angle (δ) and the longitudinal force at the front ($F_{x,F}$) and rear axle ($F_{x,R}$). The lateral front and rear tyre forces, respectively $F_{y,F}$ and $F_{y,R}$, are computed using validated Fiala tyre model [134]. However, during an evasive manoeuvre, the prediction model inaccuracies can increase significantly, so the model mismatches of the front ($\Delta F_{y,F}$) and rear ($\Delta F_{y,R}$) lateral tyre forces and the yaw rate (Δr) are computed by an STP.

The proposed L-MPCC cost function is responsible for ensuring path tracking, maintaining the physical feasibility of the control inputs, prioritising obstacle avoidance in case of collision risk, and minimising the uncertainties of the prediction model, thus limiting the operating time in nonlinear regions. The cost function (J) is defined as follows:

$$\begin{aligned} J = \sum_{i=1}^N \left(\sum_{j=1}^{N_{Obs}} (q_{e_{Obs,ji}} e_{Obs,ji}^2) + \sum_{k=1}^{N_{Edg}} (q_{e_{Edg,ki}} e_{Edg,ki}^2) + q_{\delta} \dot{\delta}_i^2 + q_{F_x} \dot{F}_{x,i}^2 + \right. \\ \left. + q_{e_{Con}} e_{Con,i}^2 + q_{e_{Lag}} e_{Lag,i}^2 + q_{e_{vel}} e_{vel}^2 \right) + \sum_{j=1}^{N_{Prob}} (q_{\sigma_r} \sigma_{r,j}^2 + q_{\sigma_{v_y}} \sigma_{v_y,j}^2) \end{aligned} \quad (5.2)$$

where N is the length of the prediction horizon, N_{Obs} and N_{Edg} are respectively the number of obstacles and road edges, and $q_{e_{Obs}}$, e_{Obs} , $q_{e_{Edg}}$ and e_{Edg} are used to prioritise obstacle avoidance and keeping the vehicle away from the road edges in case of emergency [41]. The tracking performance is ensured by the minimisation of the contouring (e_{Con}), lag (e_{Lag}) and velocity e_{vel} errors with their respective cost terms, i.e. $q_{e_{Con}}$, e_{Lag} , $q_{e_{vel}}$. The lateral velocity and yaw rate uncertainties are minimised through the σ_{v_y} and σ_r parameters of the cost function. The σ_{v_y} and σ_r are propagated open-loop, so they can increase rapidly, potentially overpowering the other elements of the cost function. To avoid the problem, they are propagated not over the entire prediction horizon, $N = 30$, but only for the reduced horizon $N_{Prob} = 20$ [47]. Furthermore, the cost parameters, $q_{\sigma_{v_y}}$ and q_{σ_r} , are tuned not to exceed the cost related to obstacle avoidance prioritisation. Therefore, the proposed controller pushes the vehicle to work at operating points close to normal driving. However, it allows the vehicle to drive at operating points close to the handling limits if it helps avoid a collision.

5.3 STUDENT-T PROCESS & UNCERTAINTY PROPAGATION

This study uses an STP as a stochastic process with a Laplace inference to estimate model mismatches and uncertainties for two reasons. First, the Student-t distribution allows the algorithm to define the level of Kurtosis, reducing the influence of the outliers and improving the accuracy of the predictions [168, 169]. The robustness to outliers is an essential property for L-MPCC, which relies on online measurements with sensor noise. For instance, a single outlier can highly influence the GP prediction; pushing the posterior means far away from the level of the other observations, which is not the case in an STP [169]. The second reason is that the STP posterior covariance directly depends on the observed measurements and not only on the location of the observed measurements. This implies that the covariance increases when the measurements vary more than expected, i.e., a higher difference between training and test sets, and vice-versa decreases when the difference is lower [48, 168]. Once again, this is an essential property for an L-MPCC, which can work in conditions with a high discrepancy from the training scenario. Regarding the kernel selection, this work opts for the automatic relevance determination Matérn 5/2 function due to its generalisation capabilities and interpretability properties. The STP is implemented using the GPML Matlab Code version 4.2 [170].

We predict the $\Delta F_{y,F}$, $\Delta F_{y,R}$ and Δr model mismatches, rather than the errors in v_x and v_y as usually done in literature [47] because it allows the reduction of the mismatches in a proactive way. It also directly targets the source of the errors, i.e., the tyre forces, and not the states that depend on them. This is possible thanks to the availability of tyre lateral forces at the front and rear axle, measured by the intelligent bearings [29], the vehicle yaw rate and the control inputs F_x and δ . Thus, we assume that the discrepancies are independent of the vehicle position [47] and that the velocity states, not measurable through the standard sensor setup for passenger vehicles, are entirely dependent on the tyre forces. The training data is derived from the algebraic difference between the measurements and the nominal model predictions.

The training set is composed of 2 double-lane changes at respectively 55 km/h and 80 km/h, and a double-lane change with collision avoidance prioritisation at 55 km/h. The manoeuvres are selected to identify the mismatches in the linear and nonlinear regions of the vehicle. The test set comprises a double-lane change at 60 km/h and the same manoeuvre with collision prioritisation. Despite the high prediction model accuracy and the reduction of the model mismatches, some unmodeled effects can still influence the prediction model. These are captured by the STP posterior covariance for the front and rear lateral tyre force, respectively $\sigma_{F_{y,F}}$ and $\sigma_{F_{y,R}}$. Thus, the tyre forces are not only evaluated as a single point but also as a distribution [129]. The STP variance and degrees of freedom are transformed and simplified into a Gaussian variance to be able to propagate it through the prediction horizon and relate them to the lateral states related to vehicle stability, i.e. v_y and r , using a procedure similar to an EKF [47, 129].

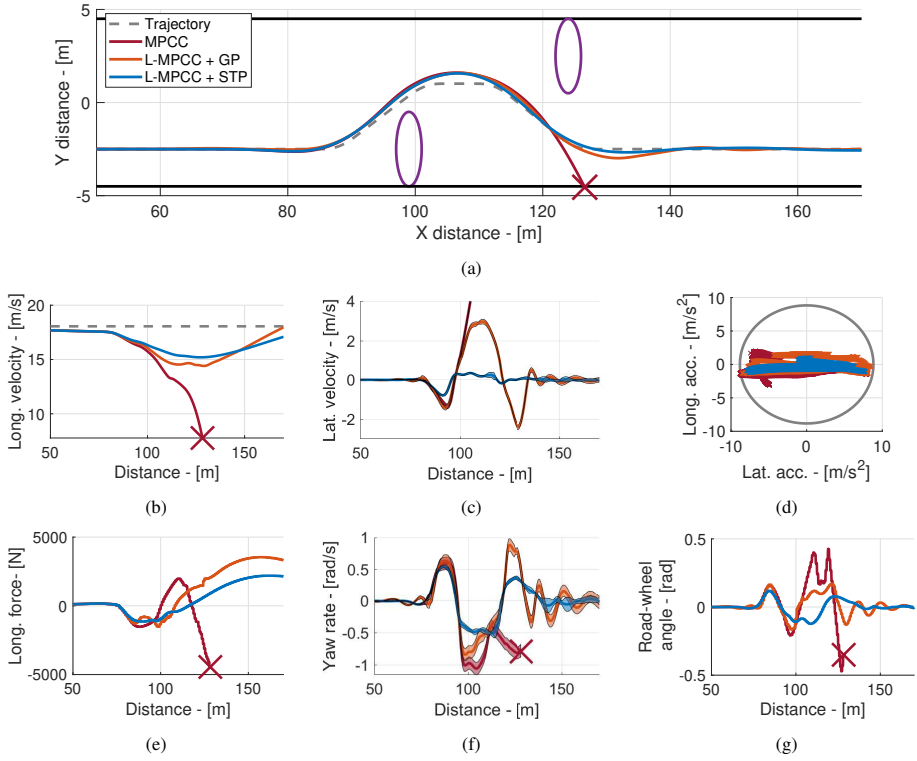


Figure 5.1: States and control inputs in a double-lane change.

5.4 RESULTS

FIGURE 5.1 shows the performance of the evaluated controllers in a double-lane change with collision avoidance prioritisation at 65 km/h. Fig. 5.1a shows that only the learning-based controllers can successfully avoid the two obstacles and stay inside the road track. Vice versa, the MPCC baseline [41] leaves the track on the right side at 130 m. The improvement is due to reduced model mismatches in the L-MPCC. For instance, the proposed L-MPCC+STP reduces the root mean square of the $F_{y,F}$, $F_{y,R}$ and r mismatches by respectively 33.14 %, 24.85 % and 60.61 % compared to the MPCC and further 20.14 %, 11.85 % and 61.33 % with respect to the L-MPCC+GP. This proves the importance of implementing an STP rather than a GP to predict the model mismatches. The enhanced performance of the proposed controller is also visible from the reduction of the vehicle lateral velocity peak; see Fig. 5.1c. Furthermore, the L-MPCC+STP can reduce the lateral velocity and, consequently, the sideslip angle peak by 76 %, maintaining an overall higher velocity during the manoeuvre than the L-MPCC+GP. Figs. 5.1c and 5.1f show that the vehicle sideslip angle reduction and the increased operating time in the stable vehicle behaviour are not only due to the lower model mismatches but also to the reduced uncertainties for the vehicle lateral states. However, it has a lower vehicle longitudinal acceleration when the obstacle avoidance manoeuvre is over to achieve the desired velocity, Fig. 5.1e.

5.5 CONCLUSIONS

THIS paper presented a novel Learning-based Model Predictive Contouring Control based on a Student-t Process for evasive manoeuvres with an online minimisation of model mismatches and uncertainties. In a high-fidelity simulation environment, we demonstrate that our proposed controller successfully avoids obstacles, keeping the vehicle stable while driving a double-lane change manoeuvre at a higher velocity compared to a non-learning-based baseline. Furthermore, the proposed controller reduces the peak of vehicle sideslip angle by 76 %. The performance enhancement is due to the properties of the Student-t Process, which can further reduce the prediction model mismatches and better capture the vehicle state uncertainties. Future work involves the proposed controller implementation on rapid prototyping hardware to prove its real-time capability.

ACKNOWLEDGEMENT

The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484). European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie actions, under grant agreement Nr. 872907.

6

A NONLINEAR MODEL PREDICTIVE CONTROL FOR AUTOMATED DRIFTING WITH A STANDARD PASSENGER VEHICLE

6

True stability results when presumed order and presumed disorder are balanced. A truly stable system expects the unexpected, is prepared to be disrupted, waits to be transformed.

Even Cowgirls Get the Blues, Tom Robbins

ABSTRACT

This paper presents a novel approach to automated drifting with a standard passenger vehicle, which involves a Nonlinear Model Predictive Control to stabilise and maintain the vehicle at high sideslip angle conditions. The proposed controller architecture is split into three components. The first part consists of the offline computed equilibrium maps, which provide the equilibrium points for each vehicle state given the desired sideslip angle and radius of the path. The second is the predictive controller minimising the errors between the equilibrium and actual vehicle states. The third is a path-following controller, which reduces the path error, altering the equilibrium curvature path. In a high-fidelity simulation environment, we validate the controller architecture capacity to stabilise the vehicle in automated drifting along a desired path, with a maximal lateral path deviation of 1 m. In the experiments with a standard passenger vehicle, we demonstrate that the proposed approach is capable of bringing and maintaining the vehicle at the desired 30 degree sideslip angle in both high and low friction conditions.

6.1 INTRODUCTION

THE ability to control the vehicle beyond the stable handling limits or at a high sideslip angle extends the number of controllable vehicle states and possible trajectories. Thus, it improves vehicle safety for the collision avoidance manoeuvres in which a conventionally driven vehicle would have reached its handling limit [40–42]. However, driving along a desired path while sustaining a large sideslip angle is particularly challenging because it requires exploiting the coupled nonlinearities in the tyre force response [43]. Furthermore, it is necessary to keep the rear tyres saturated [44]. Thus, in this paper, we focus on designing a nonlinear model predictive control (NMPC), which can stabilise a vehicle in a drifting motion while remaining on a desired path. Furthermore, we aim to implement and evaluate the proposed controller in a real-world experiment on a standard passenger vehicle without hardware modification (Fig. 6.1).

Different control techniques for automated drifting have recently been proposed in the literature. A possible solution is applying a linear quadratic regulator (LQR) controller based on the single-track model combined with a Fiala tyre model [172]. The control inputs are the steering angle and rear wheel speed, computing from the required rear longitudinal force component. The internal model has also been upgraded with a double-track vehicle model based on a Pacejka tyre model, to reduce the model mismatches [40]. However, all the proposed LQR controllers only stabilise the vehicle in the drifting equilibrium while an external path-tracking controller executes the path following control.

Thus, other approaches are proposed that simultaneously ensure the path-following properties and control the vehicle into a desired equilibrium drifting state [53, 54]. These controllers use the lateral error with respect to a reference path to control the vehicle along the desired trajectory. At the same time, the sideslip error with respect to a reference sideslip angle brings the vehicle into a state of drifting. The steering angle and rear drive torque are treated as control inputs, and they are determined through a function of imposed error dynamics, a nonlinear model inversion followed by a wheel speed control. The imposed error dynamics convert the lateral and sideslip errors into a desired yaw rate and yaw acceleration. A numerical approach of nonlinear model inversion is applied to determine the desired steering and throttle angles. The steering angle is directly applied to the system, while the throttle angle is mapped to a desired wheel speed. Experimental verification of these controllers has shown that the approach can successfully track a reference path and sideslip with good performance. However, the vehicle powertrain layout includes independent electric engines simulating the behaviour of a fully locked differential, which is uncommon for standard passenger vehicles.



Figure 6.1: BMW M3 Competition in automated drifting.

Several solutions based on Model Predictive Control (MPC) are proposed in the literature [173–175]. For instance, a simplified linear vehicle can be implemented as an MPC prediction model [175]. The MPC modifies a human driver input such that a vehicle is stabilised on the handling limits. As modified linear vehicle and tyre models are used, the optimisation becomes convex, reducing the computational effort and increasing real-time feasibility. However, the controller is not designed to bring a vehicle into a steady-state drift fully. Thus, another solution is linearising the nonlinear single-track vehicle with a Fiala tyre model around defined state and input variables [173]. As a steady-state drift requires a reference equilibrium state, the system will change according to the values of equilibria through equilibrium-varying simulation. The approach shows that the prediction model states can be chosen arbitrarily, as the sideslip angle is not chosen as a state, but tracking the reference drifting state is still possible. However, the approach still suffers from a computational point of view, making the controller not real-time feasible. Thus, in this paper, we develop a real-time NMPC based on a nonlinear single-track vehicle with a simplified Pacejka tyre model, using the sideslip angle as a vehicle state. The proposed controller aims to bring the vehicle to a high sideslip angle and stabilise the vehicle in this unstable equilibrium, considering high and low friction conditions.

The main contribution is the experimental validation of the proposed NMPC for automated drifting with a standard passenger vehicle, contrary to previous works in the literature that are limited to simulations [50, 51] or experimental demonstrators with heavy hardware modifications [53, 54].

6

6.2 VEHICLE MODEL

SUBSECTION 6.2.1 describes the nonlinear single-track model, and subsection 6.2.2 explains the tyre model.

6.2.1 SINGLE-TRACK VEHICLE MODEL

The single-track vehicle model simplifies the dynamics of a four-wheel vehicle model by lumping the tyres on the left and right sides together in the centre axis of the vehicle and ignoring the lateral load transfer. The kinematics of the single-track vehicle model is described using the Frenet reference system as follows:

$$\begin{cases} \dot{\psi} &= r - K \frac{V \cos(\psi + \beta)}{1 - e_y K} \\ \dot{e}_y &= V \sin(\psi + \beta) \\ \dot{s}_x &= \frac{V \cos(\psi + \beta)}{1 - e_y K} \end{cases} \quad (6.1)$$

where ψ is the heading angle, s_x is the travelled distance with respect to the desired path, and e_y is the normal distance from the desired path. The relation between the absolute vehicle velocity V , path curvature K and yaw rate r describes the desired trajectory.

The dynamics of the nonlinear single-track vehicle model are defined in Eq. 6.2, where the vehicle states are the sideslip angle β , the yaw rate, and the absolute vehicle velocity. The absolute velocity describes the actual velocity of the vehicle in space, combining both the lateral and longitudinal velocity components with respect to the vehicle frame. The sideslip angle is the angle between the vehicle's longitudinal axis and the absolute velocity, which is, therefore, the main indication for a vehicle being in a drift at large sideslip angles

[55, 56, 120]. Using the absolute velocity, a path following property is established, as path curvature can be expressed as a function of absolute velocity divided by the yaw rate. The suspension dynamics are not modelled to reduce the computational effort, so the longitudinal weight transfer is considered in a quasi-static way.

$$\begin{cases} \dot{V} = \frac{1}{m}(F_{xF} \cos(\delta - \beta) - F_{yF} \sin(\delta - \beta) + \\ \quad + F_{xR} \cos \beta + F_{yR} \sin \beta) \\ \dot{\beta} = \frac{1}{mV}(F_{yF} \cos(\delta - \beta) - F_{xF} \sin(\delta - \beta) + \\ \quad - F_{xR} \sin \beta + F_{yR} \cos \beta) - r \\ \dot{r} = \frac{a(F_{xF} \sin \delta + F_{yF} \cos \delta) - bF_{yR}}{I_z} \end{cases} \quad (6.2)$$

6.2.2 TYRE MODEL

The vehicle and road surface interaction is essential for the drifting motion. The transition between conventional driving and drifting requires a large rear lateral force to establish a state of drift. Similarly, a high optimal lateral force must be maintained to remain in a steady-state drift. This implies that the knowledge of the tyres should be as accurate as possible, capturing most characteristics with the lowest complexity. This work combines the standard Magic Formula (5.2) and its simplified version. Results on computational capacity and tracking performance show that the simplified Magic Formula can be applied at the front wheels. In contrast, the standard formula is desired to better capture rear tyre characteristics. The simplified model can capture those characteristics as front wheels are considered within the tyre friction limits during a drift. Regarding rear tyres, they work beyond the stable handling limits, so tyre characteristics should be captured as well as possible.

6.3 EQUILIBRIUM ANALYSIS

SUBSECTION 6.3.1 describes how the locations of the desired steady-state drift equilibria are computed, and subsection 6.3.2 shows the experimental validation of the drifting equilibria.

6.3.1 STEADY-STATE EQUILIBRIUM LOCATIONS

The steady-state drifting equilibria are computed, imposing the derivative of the nonlinear single-track vehicle model states to zero. Based on this condition, the differential equations are solved using nonlinear least-squares, allowing the incorporation of constraints on respective states. An initial guess allows the solver to converge to a desired equilibrium, e.g. a high sideslip drifting equilibrium. Furthermore, constraints on desired velocities, wheel speeds, and wheel slip make finding equilibria significantly faster. Nevertheless, as can be mathematically derived, the proposed model contains more variables than equations, which means that specific equilibrium values must be user-defined, considering a value that describes a drifting motion. Thus, the system of equations is solved using a set of desired sideslip angles β^{eq} with a set of experimentally confirmed feasible path radii R^{eq} for drifting motion. The optimisation variables are described by \mathbf{x}_{eq} (Eq.6.3), which allows the

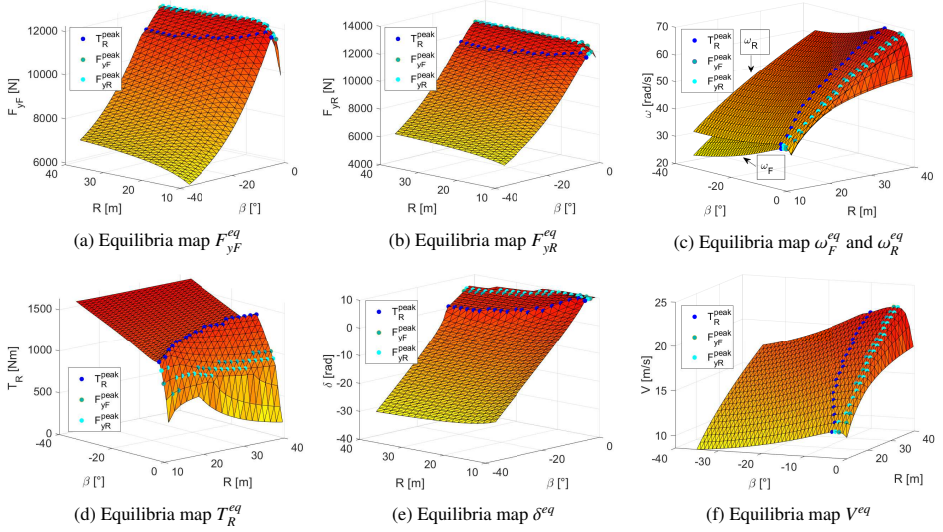


Figure 6.2: Steady-state drifting equilibria maps. The locations of the peak drive torques in \bullet , rear lateral peak forces in \bullet and the front lateral peak forces in \bullet .

6

computation of all the respective tyre force components.

$$\mathbf{x}^{eq} = [V^{eq} \quad \beta^{eq} \quad r^{eq} \quad \omega_F^{eq} \quad \omega_R^{eq} \quad \delta^{eq} \quad T_i^{eq}]^T \quad (6.3)$$

The results of the equilibrium points are considered non-unique, as a drifting equilibrium is controlled through a balancing relation between the front and rear force magnitudes. Thus, the variations in rear-drive torques and their resultant slip coefficient may vary, as the front axle is subject to varying poses and forces. Fig. 6.2 shows the equilibria maps, which have the locations of the peak drive torques in \bullet , rear lateral peak forces in \bullet and the front lateral peak forces in \bullet . For the velocity state equilibrium, see Fig. 6.2f, which increases with the desired path radius at a constant sideslip angle. In contrast, a constant radius causes a decrease in sideslip angle due to the force balance required between the front and rear axles. As a result, the vehicle's yaw rate increases with the path radius. Fig. 6.2a and 6.2b show that the lateral tyre forces are at their highest at relatively small sideslip angles. Thus, towards these peaks, the vehicle is considered in conventional driving with the tyre working in the linear range. This can also be deduced by the similar wheel speeds (Fig. 6.2c) and positive steering angles (Fig. 6.2e). Fig. 6.2c shows that the front and rear wheel speeds start to deviate in magnitude at higher sideslip angles, and the steering angle approaches zero, indicating that the vehicle is entering a drifting mode. However, the rear longitudinal force has not reached its peak force (Fig. 6.2b). This means that the applied rear torque results in a reasonable amount of slip such that the rear tyres become saturated in the lateral direction with a reasonable sideslip angle. However, since the peak friction force in the longitudinal direction is not reached, we consider the state towards this point a transient equilibrium, i.e., a slight drift. Thus, tyres are not fully saturated, nor is countersteering taking place. Once the peak in the longitudinal slip is reached, the difference in wheel speeds becomes




Figure 6.3: Comparison of the computed and measured vehicle states in equilibrium.

more significant than the rear axle longitudinal velocity. Thus, the tyres are saturated in both lateral and longitudinal directions, and a direct change in steering angle is required (Fig. 6.2e) to balance the front and rear forces. Thus, larger sideslip angles describe fully drifting equilibria where the proportional behaviour in the steering angle and the deviation in the wheel speeds (Fig. 6.2c) result from a more significant required rear slip and a decline in front wheel speed due to a lower yaw rate.

6.3.2 EXPERIMENTAL VALIDATION

The steady-state drifting equilibria are experimentally validated in a proving ground and Fig. 6.3 shows a comparison of the computed and measured vehicle state equilibria. The results show the complexity of a human driver remaining at an actual equilibrium point due to many disturbances on the vehicle, e.g., road surface. Real road surfaces are considered non-homogeneous, where a $\sim 5\%$ variation in the friction coefficient is measured for the respective proving ground, which a driver needs to compensate for by changing the steering angle or throttle commands. Thus, the measured steering angle and drive torque are never constant. The vehicle's slightly tilted position due to the surface's conical shape is also unmodelled in the single-track vehicle, causing variations in the vertical axle loads. In addition, the driver has to stabilise the desired drifting state while stabilising on a desired path. Thus, the experimental data are not in equilibrium for all time instants. However, the comparison does show reasonable similarities in the order of magnitude. The comparison

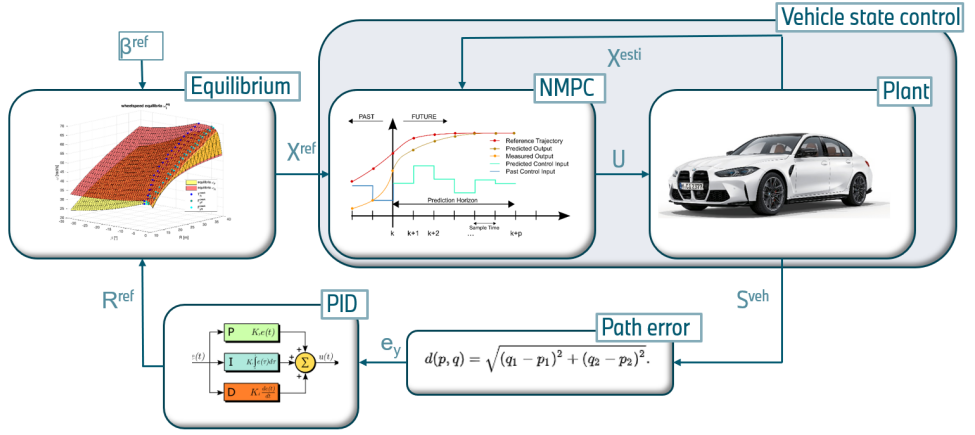


Figure 6.4: Architecture of the proposed controller.

between computed and measured drifting states shows that many variations of drifting equilibria exist around an actual drift equilibria.

6

6.4 PROPOSED CONTROLLER ARCHITECTURE

THE proposed controller architecture is shown in Fig. 6.4, designed to stabilise the vehicle into a high sideslip angle and follow a desired path. The NMPC is used to compute the optimal control inputs to reach any feasible vehicle reference state. Thus, the NMPC is deployed to stabilise the vehicle in a high sideslip state and for conventional driving. Thus, it can bring the vehicle from conventional driving towards a drifting motion. The NMPC is based on the following online optimisation:

$$\min_{\mathbf{u} \in \mathbb{U}, \mathbf{x} \in \mathbb{X}} \sum_{k=0}^N J(\mathbf{x}_k, \mathbf{u}_k) \quad (6.4a)$$

$$\text{s.t.} \quad \mathbf{x}_0 = \mathbf{x}_{\text{init}} \quad (6.4b)$$

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), \quad k = 0, \dots, N-1 \quad (6.4c)$$

where J is the objective, which includes the quadratic errors between the defined equilibria and the vehicle sideslip angle and yaw rate states; furthermore, it includes a penalty on the steering rate $\dot{\delta}$. The latter is possible by introducing δ as an extra vehicle state correspondent to the integration of the $\dot{\delta}$. Thus, two benefits are counted: first, it is possible to add a constraint on $\dot{\delta}$, modelling the physical limitation of the steering actuator, and second, the control input δ becomes smoother. The $f(\mathbf{x}_k, \mathbf{u}_k)$ corresponds to the nonlinear single-track vehicle model. The optimisation problem in Eq. 6.4 is solved using the ACADO toolkit [176] with the Sequential Quadratic Programming solver at 50 Hz with a prediction horizon $N = 25$.

The desired equilibria are provided to the NMPC through dynamic referencing, which is implemented to influence the vehicle's direction. The offline computed equilibria are selected dynamically using s_x to alter the path curvature such that the lateral error e_y of a

desired path is decreased. This is established by altering the desired path curvature in relation to the lateral error through a path-following controller that determines a compensating factor on the path curvature. The altered path curvature functions as an input for a lookup table that contains a complete set of equilibria, which, therefore, picks the optimal equilibrium that functions as a reference for the controller.

6.5 SIMULATION RESULTS

SUBSECTION 6.5.1 describes how the proposed controller is evaluated in a high-fidelity environment, and subsection 6.5.2 shows the proposed controller validation.

6.5.1 SIMULATION SETUP

The proposed controller is first evaluated in a simulation environment. The vehicle plant is a high-fidelity model of a BMW M3 Competition. It is a rear-wheel drive vehicle with a limited-slip differential (LSD), allowing for the locking rate modification. The vehicle plant is based on a 17 degrees of freedom model, experimentally parameterised and validated by BMW. Magic Formula 5.2 is used to model the tyre dynamics. An alternating drifting manoeuvre on a high-friction surface is defined as a testing scenario. The manoeuvre is initialised with a high yaw rate state within friction limits, i.e. conventionally turning. Furthermore, measurement noises and friction coefficient variations are included in the simulation. Friction coefficient variation is based on surface measurements at the BMW proving grounds of Aschheim, showing that seemingly homogeneous surfaces tend to have friction coefficient fluctuation of approximately 0.05. The amplitude and frequency of simulated measurement noises are based on experimental vehicle measurements.

6.5.2 ALTERNATING DRIFTING MANOEUVRE

The simulated vehicle states in the alternating drifting manoeuvre are shown in Fig. 6.5. Once the drifting reference is initialised, the proposed controller can successfully track the initial left-hand drifting manoeuvre and perform a nearly instant transition into a right-hand drift with optimal tracking performance. In this case, the response of a 30 % LSD is reached as the left and right wheel speeds become similar, implying a locked differential state. Furthermore, the results highlight that the nonlinear single-track model captures the drifting vehicle dynamics, as the desired equilibrium is based upon that description. The control inputs are shown in Fig. 6.6. Despite the same NMPC tuning for the complete simulation showing some erratic behaviour in the linear driving range, the solver remained stable. The results highlight that the steering angle remains smooth, as the δ is constrained in the optimisation.

The simulator's response shows powertrain dynamics and systematic delays, as the single-track model calculation of the left and right wheel drive torque is not of equal peak magnitude. This implies that the rapid increase of desired drive torque is tracked slowly. However, the drifting states are reached as the solver computes the equilibrium drive torque.




Figure 6.5: Vehicle states in the alternating drifting manoeuvre on a high friction surface.




Figure 6.6: Control inputs in the alternating drifting manoeuvre on a high friction surface.




Figure 6.7: Desired and applied control inputs during automated drifting experiment.

6.6 EXPERIMENTAL RESULTS

SUBSECTION 6.6.1 describes how the proposed controller is evaluated in the proving ground, and subsection 6.6.2 shows the experimental validation of the proposed NMPC for automated drifting.

6.6.1 EXPERIMENTAL SETUP

The experimental validation is conducted with a BMW M3 Competition vehicle. The auto-generated real-time ACADO c-code is implemented on a dSPACE Autobox DS1007 platform, which also acts as an interface for desired Flexray and CAN bus communication protocols to send and receive information from the vehicle and desired external devices, e.g. GNSS or IMU. Measured and estimated state signals at 10 ms are found on the vehicle Flexray. Regarding the control actuators, the human driver inputs provided to the driving assistance systems interface are replaced by the NMPC-optimised control inputs. However, the automated lane keeping, commonly available in a passenger vehicle, is not designed to be active in fully automated mode, so applied steering torque and angle through the driving assistance systems interface are greatly limited by safety features.

When the proposed controller is validated, ignoring these limitations, it can only successfully bring the vehicle to a high sideslip angle condition but not stabilise it. Fig. 6.7 shows a deviation between the desired and measured steering angle for all performed tests reaching a high sideslip angle, where the front axle aligns with the direction of travel rather than compensates for the vehicle state. A possible explanation is the lack of applicable torque from the standard electric power steering motor that should be able to apply sufficient torque to overcome the self-aligning moment occurring at the front tyres. The e-motor can provide a maximum of ~ 30 Nm of torque and is designed to function for several low-torque ADAS systems, such as automatic lane change. Thus, the steering angle cannot be decreased as the NMPC desires and, as a result, the actual steering angle remains on its limit, which causes the vehicle to spin, as the front axle is not performing the desired compensation for the deviations in absolute velocity and vehicle sideslip angle. Regarding these limitations, only




Figure 6.8: Vehicle states in semi-automated drifting manoeuvre.

software/hardware modifications or a steering robot can solve the steering angle limitation. Therefore, the proposed controller is validated with a semi-automated drifting manoeuvre. The proposed NMPC computes the optimal control inputs, but only the desired drive torque is provided to the vehicle while a human driver assists in the steering. In this experiment, we can validate the NMPC drive torque and compare the actual and desired steering input of the NMPC. The experiment is performed on a dry surface and a watered skid pad, respectively, with high and low friction conditions.

6.6.2 SEMI-AUTOMATED DRIFTING MANOEUVRE

The vehicle states in a semi-automated drifting manoeuvre are shown in Fig. 6.8. The vehicle successfully enters the desired drifting mode, and the high sideslip angle is stabilised for multiple rounds at 30 deg. Thus, the NMPC driving torque is sufficient and solved optimally to bring the vehicle into drifting mode and maintain its steady-state drifting position. However, it is visible that the tracking performance of the velocity and front wheel speed is less accurate than the sideslip angle, yaw rate, and rear wheel speed. A possible explanation is that the controller tuning is optimised to bring the vehicle into a high sideslip angle, established with a high yaw rate and accurate rear wheel speed tracking. On the contrary, the front axle aligns with the direction of travel, and the front wheel speed and the absolute velocity are proportional under constant path curvature. Fig. 6.7 shows the optimised and actuated control inputs. The comparison between steering angles demonstrates that the steering angle provided by the human driver is very similar to the desired steering angle optimised by the NMPC solver. The maximum variation in magnitude is ~ 5 deg,




Figure 6.9: Actuated and optimised control inputs in a semi-automated drifting manoeuvre.

demonstrating the NMPC capacity to stabilise the vehicle in drifting mode with appropriate actuators.

6.7 CONCLUSIONS

THIS paper presented a novel approach to automated drifting with a standard passenger vehicle, focused on a Nonlinear Model Predictive Control to stabilise and maintain the vehicle at high sideslip angle conditions. In this work, we experimentally verified the correspondence of the vehicle state equilibria computed with a nonlinear single-track model with the one measured in a proving ground. The experimental verification of the controller showed that using a standard production vehicle without significant hardware and software modifications is not possible due to the limited steering torque provided by the standard interface of electric power steering. However, a semi-automated drifting manoeuvre demonstrated the controller's capacity to bring a real vehicle into drifting mode and stabilise it at a high vehicle sideslip angle of 30 deg. Future works involve performing the experimental validation of the proposed controller in a fully automated mode, using a steering robot as an actuator.

ACKNOWLEDGEMENT

The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484).

7

CONCLUSIONS & RECOMMENDATIONS

*We shall not cease from exploration
And the end of all our exploring
Will be to arrive where we started
And know the place for the first time.*

Little Gidding, Four Quartets, Thomas Stearns Eliot

THE preceding chapters explored the development of vehicle sideslip angle estimation algorithms and obstacle avoidance controllers for evasive manoeuvres at the limit of handling in automated vehicles. In this chapter, I present a comprehensive analysis of the research outcomes, emphasising the key findings along with their limitations. Furthermore, I conclude the thesis by proposing directions for future work based on the new questions and insights that have arisen from this study.

7.1 DISCUSSION & MAIN FINDINGS

AUTOMATED vehicles have proven their effectiveness in reducing the number of injury-causing crashes [22, 25], confirming the relevance of the technology for the entire society. However, despite remarkable technological advancements, automated vehicles still fall short of ensuring safety in the most hazardous situations, such as when a vehicle, pedestrian, or object suddenly appears in front of the car, leaving insufficient distance to prevent a collision through emergency braking alone [26]. In such instances, an evasive manoeuvre involving precise coordination of steering and braking is necessary. Furthermore, any abrupt control actions aimed at avoiding the obstacle can lead to loss of vehicle control, potentially resulting in a spin or collisions with other vehicles [19, 26].

Therefore, **the objective of this thesis is to develop innovative nonlinear motion control algorithms for obstacle avoidance at the limit of handling in automated vehicles and to evaluate the benefits of integrating tyre force sensing technology.**

The design of the innovative motion control algorithms was split into three sub-goals: vehicle sideslip angle estimation to assess stability and controllability, vehicle motion control for obstacle avoidance at the limit of handling with a focus on prioritising safety over path tracking and stability, and automated drifting to stabilise and control the vehicle beyond the friction limit or at high sideslip angles.

7.1.1 SUB-GOAL 1: VEHICLE SIDESLIP ANGLE ESTIMATION

The design of vehicle motion controllers capable of safely avoiding obstacles at the limit of handling demands a comprehensive understanding of the vehicle's stability and controllability envelope. This requires precise information on the yaw rate and sideslip angle, which are critical to active vehicle control systems [57, 58]. While low-cost gyroscopic sensors can reliably measure yaw rate, the sideslip angle must be estimated in real-time. Consequently, the first sub-goal of this thesis was to design a hybrid model-based and data-driven algorithm for vehicle sideslip angle estimation with intelligent (force sensing) bearings integration, which works for a large diversity of vehicle manoeuvres, e.g. steady-state, transient, low, and high excitation. Numerous approaches for sideslip angle estimation have been proposed [36, 37], which can be categorised into three main groups: model-based, data-driven, and hybrid approaches. The model-based approach relies on a nonlinear vehicle and tyre model. However, it faces limitations in transient and high-excitation driving scenarios, where the tyre's nonlinear behaviour becomes more pronounced. Furthermore, its efficiency and robustness depend on the process and observation noise parameters, which must be accurately tuned for optimal performance [177, 178]. To address these challenges, I proposed a new tuning methodology called Two-Stage Bayesian Optimisation [179], based on a Student-t Process, as explained in Appendix A. This methodology accelerates the optimisation process

while improving the accuracy of the Unscented Kalman Filter tuning, as demonstrated on an experimental dataset. Additionally, I demonstrated that the Unscented Kalman Filter outperforms the Extended Kalman Filter in terms of accuracy. However, when additional tyre force measurements are utilised, their performance becomes comparable [56]. Data-driven approaches, on the other hand, utilise historical data to model relationships between input and output variables. These methods generally offer higher accuracy than model-based approaches when sufficient, high-quality training data are available, as I also demonstrated in Chapter 2 [56]. This becomes even more pronounced when the data-driven approaches utilise the tyre force measurements [56]. However, their performance declines when the dataset lacks adequate representation in specific ranges of sideslip angles. The data-driven approach primarily relies on historical data collected to establish relationships between input, internal, and output variables. Overall, it has a higher accuracy than the model-based approach when enough quality data are provided in the training phase. However, its performance is highly dependent on the amount of representative data, and it will lack performance as soon as the dataset contains a lower amount of data in a particular range of the sideslip angle. Hybrid approaches aim to combine the strengths of both model-based and data-driven methods. While this leads to improved performance [63, 64], the literature to date has primarily focused on integrating these approaches without creating a mutualistic relationship between them. The entire hybrid approach potential for sideslip angle estimation was not exploited.

Chapter 2 presented a novel hybrid approach that establishes a mutual relationship between model-based and data-driven approaches for vehicle sideslip angle estimation. This approach leverages physical insights from an Unscented Kalman Filter based on a nonlinear single-track vehicle model to enhance the accuracy of a Convolutional Neural Network (CNN). The performance of the different approaches was evaluated not only using traditional key performance indicators (KPIs) for sideslip angle estimation, such as the mean squared error (MSE) to assess overall estimation accuracy and the absolute maximum error (ME) to examine the worst-case performance. Additionally, I introduced two nonlinear metrics: the nonlinear MSE (MSE_{nl}) and the nonlinear ME (ME_{nl}), which specifically measure estimation accuracy when the vehicle exhibits nonlinear behaviour, such as when the absolute value of the lateral acceleration exceeds 4 m/s^2 . These nonlinear KPIs offer a more detailed assessment of the hybrid approach's performance in the most critical driving scenarios, directly aligning with the primary objective of this work: addressing obstacle avoidance at the limit of handling. Using a large-scale experimental dataset consisting of 216 manoeuvres, the hybrid approach demonstrated superior accuracy compared to purely model-based or data-driven methods. Moreover, the CNN-UKF approach slightly reduced the MSE compared to the Deep Ensemble UKF (DE-UKF). However, when comparing the MSE_{nl} , the CNN-UKF outperformed the DE-UKF by 25 %, providing significantly higher accuracy in the most critical operating regions for active vehicle control systems. This result is particularly aligned with the main scope of this work. Through end-to-end training, the CNN is complying with the vehicle's physical dynamics, leading to reductions in both the maximum error and nonlinear maximum error compared to all other approaches. The latter results are particularly relevant to assess that the estimation is always coherent with the physical vehicle behaviour.

A new vehicle state estimation algorithm must enhance performance under normal conditions while also proving its robustness. The performance of the proposed CNN-UKF

was tested on a limited dataset, which only included measurements where the vehicle's absolute lateral acceleration is below 7 m/s^2 . This scenario simulates the practical challenges of collecting a large number of manoeuvres at the handling limits. Such conditions are common in the automotive field, as driving at the limits of handling can result in the driver losing control of the vehicle. In this context, the hybrid CNN-UKF approach showed a slight improvement in estimation robustness compared to the model-based and DE-UKF approaches across all KPIs. However, the CNN-UKF significantly outperformed the purely data-driven approach in terms of estimation accuracy.

The main findings on Sub-Goal 1 can be summarised as follows:

- **The benefits of measured tyre forces are quantified for vehicle sideslip angle estimation using physics model-based and data-driven approaches with experimental data.**
- **End-to-end training between a CNN and a differentiable UKF enables the development of a hybrid approach that establishes a mutual relationship between model-based and data-driven methods for vehicle sideslip angle estimation.**
- **The hybrid CNN-UKF approach ensures adherence to the physical laws of vehicle dynamics, improving MSE and absolute ME estimation across a diverse range of vehicle manoeuvres, such as steady-state, transient, low, and high excitation, all of which recorded on a proving ground.**
- **The online estimation of UKF process model uncertainties thanks to CNN and UKF end-to-end training improves accuracy and robustness, even when trained with only a limited dataset.**
- **The proposed hybrid approach offers greater robustness in vehicle sideslip angle estimation than data-driven approaches, even when trained on a limited dataset.**

7.1.2 SUB-GOAL 2: OBSTACLE AVOIDANCE AT THE LIMIT OF HANDLING

The safety of automated vehicles heavily depends on their ability to execute evasive manoeuvres, even at the handling limit. Central to any nonlinear motion control algorithm for these manoeuvres is the obstacle avoidance controller, which must balance competing objectives while remaining robust to real-world variations and uncertainties. For this reason, the second sub-goal of this thesis was to design nonlinear model predictive controls which can prioritise obstacle avoidance even at the limit of handling while compensating for inaccuracies in localisation and perception systems, as well as non-nominal road and vehicle conditions. Despite recent advancements in automated vehicle safety systems, tyre nonlinearities continue to present a significant challenge [38]. The most common approach is a hierarchical controller architecture, which separates motion planning, path tracking, and vehicle stability tasks [39]. However, during evasive manoeuvres, these objectives can conflict with each other [26]. For instance, vehicle stability constraints could lead to tracking errors, potentially causing collisions [26, 42]. To address this, I proposed a unified controller that integrates motion planning, path tracking, and vehicle stability for high-speed obstacle avoidance.

Chapter 3 introduced a Model Predictive Contouring Control (MPCC) based on a nonlinear single-track vehicle model for obstacle avoidance at the limit of handling. The nonlinear MPCC integrates motion planner, path tracking and vehicle stability objectives into a single controller, prioritising obstacle avoidance in emergency situations. In a high-fidelity simulation validated by experimental data, the MPCC successfully avoided two obstacles during a double-lane change after re-planning the target trajectory from the behaviour planner. In contrast, a controller without collision avoidance prioritisation failed, resulting in a collision. The collision avoidance prioritisation slightly reduced path-tracking accuracy and increased the sideslip angle peak to 3 deg, but the vehicle remained stable and manoeuvrable throughout the manoeuvre. This demonstrates that prioritising collision avoidance increases path and velocity tracking errors and reduces vehicle stability, pushing the vehicle closer to the handling limits, but it is essential to maintain a safe distance from obstacles. The state-of-the-art baseline, which employs a Model Predictive Control (MPC) approach using a Frenet reference system, also avoids the two obstacles but fails to keep the vehicle outside the unsafe zone near the obstacles, resulting in a near-miss event. Additionally, it compromises vehicle stability, with the sideslip angle peaking to 9 deg. The improved performance of the MPCC, which operates in a Cartesian frame, is attributed to its ability to approximate the MPC based on the Frenet reference system by incorporating lag and contouring errors into the cost function. This prevents the overestimation of the vehicle-to-obstacle distance caused by Frenet coordinates and eliminates the need for additional optimisation to compute the vehicle's distance travelled relative to the reference line. The proposed approach was also experimentally validated on a Toyota Prius performing a double-lane change at 30 km/h on a low-friction surface in heavy rain conditions. The experimental validation demonstrated the controller's real-time implementability on an embedded platform, such as the dSPACE AutoBox 1007, as well as its capacity to handle measurement noise, actuator delays and constraints. Additionally, the low road friction in the experimental scenario highlighted the robustness of the proposed controller across varying friction conditions.

The successful validation of the proposed MPCC proved the importance of prioritising collision avoidance properties rather than path tracking and vehicle stability performance in case of emergency. However, the assumptions of driving a vehicle at the limit of handling and of a perfect obstacle perception may raise concerns about the robustness of the proposed approach. To address this, the proposed approach's robustness is analysed to vehicle and tyre parameter variations, perception inaccuracies, the number of false negatives in the perception stack, and localisation errors. Additionally, the computational complexity of the MPCC highlights potential challenges in adapting it to emerging control methodologies that incorporate torque vectoring. These methodologies have gained significant attention with the development of new electric powertrains, especially those using multiple in-wheel electric motors [139–141].

For this reason, chapter 4 extended the MPCC for obstacle avoidance, combining motion planning, path tracking, and vehicle stability as described in Chapter 3 by incorporating torque vectoring to enhance the vehicle's lateral agility at the handling limits. Torque vectoring is integrated into the optimal control problem through the use of a double-track vehicle model with an extended Fiala tyre model. To address robustness concerns, the enhanced MPCC with torque vectoring was evaluated under varying vehicle and tyre parameter variations, perception accuracies, number of false negatives, and localisation errors. Addi-

tionally, a sensitivity analysis was performed to quantify which parameters and perception inaccuracies most significantly degrade the performance of the obstacle avoidance MPCC controllers.

The proposed controller with torque vectoring and collision avoidance successfully avoided two obstacles during a double-lane change manoeuvre under both high and low friction conditions. It was implemented on a real-time rapid prototyping platform (dSPACE SCALEXIO), with an average solving time of 38.9 ms. The first baseline controller, lacking collision avoidance, collided with the first obstacle due to its inability to replan motion and stay away from the area near the obstacles. The second baseline controller, without torque vectoring but with collision avoidance prioritisation, failed to stabilise the vehicle after avoiding the first obstacle, leading to a collision with the second one. The proposed controller generated an additional yaw moment of up to ~ 2000 Nm, significantly enhancing lateral agility and reducing the vehicle's sideslip angle peak to 7.5 deg, compared to 16 deg for the baseline without torque vectoring. However, the controller requires a more accurate prediction model, leading to the development of an extended Fiala tyre model to capture tyre behaviour under combined slip conditions. In particular, the proposed tyre model captures the effect of longitudinal and vertical force on the cornering stiffness and the diminishing tyre saturation region.

Extended Monte Carlo analyses demonstrated the controller's robustness to vehicle parameter variations and perception inaccuracies, showing collision rates of 31.20 % and 38.50 %, respectively, while the baseline controllers consistently crashed into obstacles. Notably, the percentage of crashes occurring near nominal vehicle parameters and typical perception inaccuracies was below 1.5 %, indicating that most crashes occurred with parameter configurations that deviated significantly from the nominal values and were statistically less probable.

A sensitivity analysis identified different influential groups for vehicle and road parameter variations and perception and localisation inaccuracies. Regarding the first group, the rear axle's lateral peak friction coefficient had the greatest influence. This is likely because a reduction in lateral friction can induce oversteer, which is particularly challenging to control and stabilise, even for a human driver [166]. The least influential parameters included the front and rear axle longitudinal stiffness shape factors, tyre relaxation length, and electric motor delay. Their minimal effect can be attributed to the fact that the analysis primarily focused on lateral dynamics, while the longitudinal parameters mainly affect longitudinal dynamics. Moreover, the controller was designed to function also with hydraulic brakes, which have even slower dynamics than electric motors, see Chapter 3 [41]. Regarding perception and localisation inaccuracies, obstacle localisation errors in the lateral direction had the highest impact on controller performance. Increased lateral error can create an impossible scenario for the MPCC, such as shifting an obstacle too close to the road centre, making collision avoidance infeasible. The lower influence of the number of false negatives in the perception stack is due to its probabilistic nature; if a false negative occurs when the vehicle is too far from the obstacle, it may not affect the outcome. However, since controlling the timing of false negatives is unrealistic, it was not analysed in this study.

Chapters 3 and 4 successfully introduced an innovative nonlinear model predictive control that prioritises obstacle avoidance, even at the handling limits. Additionally, Chapter 4 demonstrated the robustness of the proposed approach to variations in vehicle and tyre

parameters, perception inaccuracies, false negatives, and localisation errors when compared to baseline controllers, while also identifying the most influential parameters affecting controller performance. This achievement fulfils the objective of Sub-goal 2. The next step to complete the overall objective of this thesis is to analyse the potential benefits of incorporating tyre force sensing technology and addressing the most influential parameters impacting controller performance.

To address this, Chapter 5 proposed a novel Learning-based Model Predictive Contouring Control (L-MPCC) based on a Student-t Process (STP) for evasive manoeuvres, which minimises model mismatches and uncertainties online. The STP predicts the mismatches between the prediction model and the actual vehicle yaw rate and lateral tyre forces, measured through intelligent (force sensing) bearings [29, 180]. Unlike traditional approaches that focus on vehicle velocity states [47], this method directly reduces mismatches in the tyre model, improving the controller's pro-activity. Additionally, the proposed controller minimises the vehicle's operating time at unstable points, leveraging the STP's posterior covariance of tyre forces. These uncertainties are propagated along the prediction horizon and minimised within the MPCC cost function. In a high-fidelity simulation environment, validated with experimental data, the controller successfully avoided obstacles and maintained vehicle stability during a double-lane change manoeuvre at a higher velocity than a non-learning-based baseline. This performance enhancement is attributed to the properties of the Student-t Process, which better captures vehicle state uncertainties than a Gaussian Process (GP). Furthermore, the STP offers improved outlier resistance and adapts its variance based on observed data. Thus, it can produce a higher variance for unseen operating points and a lower variance for familiar ones [48]. Notably, the controller reduced the peak vehicle sideslip angle by 76 % compared to a non-learning-based MPCC, a result of reduced model mismatches and minimised lateral state uncertainties.

The main findings on Sub-Goal 2 can be summarised as follows:

- **A MPCC that integrates motion planning, path tracking, and vehicle stability into a single cost function, prioritising collision avoidance in emergencies, improves safety for obstacle avoidance at the limit of handling compared to a hierarchical approach.**
- **The distances between the vehicle and obstacles, as well as the vehicle and road edges, are more accurately captured in a Cartesian frame, whereas they tend to be overestimated in the commonly used Frenet reference frame.**
- **The proposed MPCC improves safety and vehicle stability compared to the state-of-the-art MPC based on a Frenet reference frame largely due to the minimisation of lag error. This feature encourages the vehicle to brake proactively when motion replanning is required by the controller.**
- **The MPCC for obstacle avoidance at the limit of handling can be extended with torque vectoring capabilities while maintaining the real-time feasibility on a rapid prototyping platform.**
- **Torque vectoring improves vehicle responsiveness, ensuring stability and preventing potential crashes in both high and low friction conditions.**

- The extended Fiala tyre model accurately captures variations in cornering stiffness with respect to longitudinal and vertical forces, while adjusting the gradient within the tyre's saturation region, significantly improving the accuracy in scenarios with strong force coupling.
- The proposed MPCC is robust to variations in vehicle and road parameters, as well as to false negatives, perception errors, and localisation inaccuracies.
- Among the analysed parameters, the lateral peak friction coefficient of the rear axle and the obstacle localisation error in the lateral direction have the greatest impact on the controller's performance.
- Intelligent (force sensing) bearings contribute to the performance improvement of an L-MPCC based on an STP, directly reducing of prediction model mismatches in the tyre model.
- An L-MPCC based on an STP can successfully avoid obstacles at higher speeds compared to one based on a GP.
- A Student-t Process enhances outlier resistance and improves the estimation of posterior covariance compared to the state-of-the-art Gaussian Process.
- The STP posterior covariance is influenced by observed measurements, resulting in higher variance for operating points that deviate from the training set and lower variance for those closely aligned with the training data.
- The uncertainty minimisation of the L-MPCC based on an STP enhances vehicle stability during evasive manoeuvres by reducing unnecessary time spent near the vehicle's handling limits.

7.1.3 SUB-GOAL 3: AUTOMATED DRIFTING

The ability to control a vehicle beyond the stable limits of handling or at high sideslip angles significantly expands the range of controllable vehicle states and trajectories. This capability is crucial for enhancing safety during extreme collision avoidance manoeuvres, where a conventionally driven vehicle would likely have reached its handling limits and lost control [40, 42]. In such scenarios, maintaining a desired path while operating at high sideslip angles presents a considerable challenge. It demands precise management of the nonlinear interactions between longitudinal and lateral tyre forces, which become increasingly complex at higher slip angles due to tyre dynamics nonlinearities [43]. For this reason, the third sub-goal of this thesis was to design a controller for automated drifting, enabling the vehicle to induce oversteer while maintaining control and driving the vehicle through the entirety of a corner.

Various control techniques for automated drifting have recently been proposed in the literature. One solution is the application of a Linear Quadratic Regulator controller based on the single-track model combined with a Fiala tyre model [58]. Despite the high accuracy in stabilising the vehicle at a high sideslip, the controller is unable to perform the transition phase between low and high sideslip angles due to inherent nonlinearities. To address this limitation, other approaches have proposed the use of a numerical approach of nonlinear

model inversion to determine the desired steering and throttle angles [49, 52–54]. Experimental validation has shown successful tracking of both the reference path and sideslip with good performance. However, the powertrain layout in these experiments, which includes independent electric motors simulating a fully locked differential, is uncommon in standard passenger vehicles. A similar conclusion applies to the use of high negative camber and toe-out in the front axle.

Chapter 6 presented a nonlinear model predictive control (NMPC) capable of stabilising a vehicle in a drifting motion while following a desired path. The NMPC was implemented and evaluated through experimental tests using a standard passenger vehicle, without any hardware modifications, at a proving ground. The real-time NMPC is based on a nonlinear single-track vehicle model with a simplified Pacejka tyre model, using sideslip angle as a vehicle state. The experimental validation confirmed that the vehicle state equilibria computed using the nonlinear single-track model corresponded with those measured on a proving ground. Across 420 experiments on various surfaces, the proposed NMPC consistently transitioned the vehicle from conventional driving to drifting. However, the vehicle was only stabilised at a high sideslip angle in a few experiments due to limitations of the standard electric power steering system, which could not robustly actuate the commanded angle because of software and hardware constraints. As a result, the proposed controller was validated through a semi-automated drifting manoeuvre. In this setup, the NMPC computed the optimal control inputs, with the desired drive torque applied to the vehicle, while a human driver assisted with steering (without overriding the steering input). This setup allowed validation of the NMPC's drive torque and enabled comparison between the actual and desired steering inputs. The semi-automated drifting manoeuvre successfully demonstrated the controller's ability to induce drifting and stabilise the vehicle at a high sideslip angle of 30 deg in both high- and low-friction conditions.

7

The main findings on Sub-Goal 3 can be summarised as follows:

- **The proposed NMPC architecture stabilises the vehicle during automated drifting along a desired path, achieving a maximum lateral deviation of 1 m while maintaining a 30 deg sideslip angle in a high-fidelity simulation environment validated by experimental data.**
- **The offline computed steady-state drifting equilibria are experimentally validated at a proving ground using a standard passenger vehicle driven by a professional driver.**
- **The proposed NMPC can successfully induce and stabilise the vehicle at a high sideslip angle of 30 deg in semi-automated mode under both high and low friction conditions.**

7.2 LIMITATIONS & FUTURE WORKS

THIS section aims to comprehensively present and explain the limitations and underlying assumptions of this thesis, with the objective of offering clear recommendations for future research and development.

7.2.1 ROAD FRICTION COEFFICIENT

A common assumption throughout this thesis is the availability of an approximate value for the road friction coefficient. This assumption is widespread in the literature, as it significantly reduces the challenges in the controller design [38, 47, 54]. Typically, the road friction coefficient is estimated by an independent system layer, separate from the obstacle avoidance controller [33, 181]. This thesis highlights the importance of having an accurate estimation of the road friction coefficient during obstacle avoidance manoeuvres. Chapter 4, for instance, statistically analysed the impact of various road and vehicle parameters on the performance of the obstacle avoidance controller. The analysis revealed that the road friction coefficient, particularly the rear axle's lateral peak friction coefficient, has the most significant influence on controller performance. In light of this, the proposed controllers and estimators were tested under both high friction conditions (simulating dry roads) and low friction conditions (simulating heavy rain). Chapters 3, 4 and 5 incorporate a safety factor to account for potential inaccuracies in road friction coefficient estimation. This safety factor has been demonstrated to enhance the robustness of the obstacle avoidance controller, as discussed in Chapter 4. Nevertheless, the reliance on an accurate road friction coefficient remains a notable limitation of this work, and it is one that must be addressed in future research. Future work could address this limitation by integrating road friction estimation directly into the UKF-informed neural network presented in Chapter 2. One possible approach would involve modelling the road friction coefficient dynamics using a random walk model [33, 79]. Alternatively, the advent of intelligent (force sensing) bearings offers the potential to estimate road friction by the tyre aligning moments and lateral forces [182, 183]. Unfortunately, the implementation and testing of these ideas were not feasible in this thesis due to the lack of an experimental dataset featuring different road friction conditions, such as dry, wet, snowy, and icy surfaces. As a result, Chapter 2 was limited to evaluate the robustness of the proposed UKF-informed NN to a few manoeuvres recorded during light rain conditions.

Another promising line of research involves modifying the obstacle avoidance controller itself. Chapters 3, 4, and 5 focus on tracking a reference trajectory and, in case of emergency, prioritising the obstacle avoidance, which maximises the vehicle to obstacle distance. An alternative approach could involve maximising the vehicle's forces according to an acceleration reference [184]. Despite the clever solution, it is currently limited to situations where the vehicle must avoid a single obstacle, which differs from the more complex, multi-obstacle avoidance scenarios addressed in this thesis.

Another direction for future work would be to build upon the approach presented in Chapter 5. Specifically, the proposed Student-t Process, which currently estimates model mismatches related to tyre nonlinearities and unmodeled dynamics, could be extended to estimate mismatches caused by varying road friction coefficients. This solution would require online training of the Student-t Process to adapt to sudden changes in road friction conditions. However, a significant challenge lies in determining how quickly this online adaptation could be achieved to effectively compensate for rapid changes in friction. Further investigation is required to assess the feasibility of such an implementation in real-time scenarios.

Regarding automated drifting, sensitivity analysis on road friction variations has shown that controllers can still minimise path-tracking errors, although the vehicle is stabilised around a different equilibrium point. This effect is particularly evident when the road friction

coefficient is overestimated [52]. The ability to maintain path tracking during drifting, despite variations in road friction, underscores the potential of further research in automated drifting for collision avoidance.

7.2.2 ONE MANOEUVRE DOES NOT FIT ALL

In Chapters 3, 4 and 5, the benefits of the proposed obstacle avoidance controllers were evaluated using a double-lane change manoeuvre with two obstacles. This scenario was chosen because it aligns with the International Standard for evaluating vehicle obstacle avoidance performance [185]. While this manoeuvre effectively represents a range of emergency situations, it does not account for all potential edge cases that a vehicle may encounter during obstacle avoidance. To improve the robustness of the proposed controllers, future work should investigate their performance in a wider variety of scenarios beyond the standard double-lane change. For example, manoeuvres such as hairpin turns [26] or curved roads with multiple obstacles [136] could offer additional insights into their capabilities. However, simply increasing the number of simulated scenarios may not fully resolve this limitation, as it is crucial to avoid redundant scenarios [186]. Addressing this issue requires identifying unknown unsafe scenarios that might not be captured by standard manoeuvres [187]. Furthermore, the development of a dissimilarity metric to compare different known unsafe scenarios is essential for ensuring a comprehensive and meaningful evaluation of the controllers' performance [186].

Another key assumption of Chapters 3, 4, and 5 is that obstacles appear during the manoeuvre but remain static afterwards. The next step will be the inclusion of moving obstacles in the scenario. Although this might seem straightforward, the presence of moving obstacles could, in certain cases, simplify the emergency scenario. This is because the location of static obstacles is carefully designed to push the vehicle to its handling limit during the evasive manoeuvres. In contrast, realistic moving obstacles, which may attempt to avoid collisions autonomously, could reduce the severity of emergency situations. On the other hand, incorporating unrealistic moving obstacles that intentionally try to collide with the vehicle could lead to unfeasible scenarios, which may not provide meaningful insights. Once again, it is essential to conduct a thorough analysis of known-unsafe scenarios. This approach will ensure a comprehensive evaluation of the proposed controllers' performance in dynamic environments without resorting to overly contrived or unrealistic situations.

7.2.3 CAN WE AVOID OBSTACLES WHILE DRIFTING?

This thesis is built on the assumption that controlling a vehicle beyond stable limits of handling or at high sideslip angles increases the number of controllable vehicle states and potential trajectories, thereby enhancing vehicle safety during collision avoidance manoeuvres [40]. Chapters 3 and 4 have demonstrated that during an emergency manoeuvre, an automated vehicle must prioritise obstacle avoidance over path tracking and vehicle stability. In fact, the vehicle was able to avoid obstacles effectively only by allowing a higher sideslip angle than would have been achieved with a controller that prioritises stability and tracking. Chapter 5 showed that a similar controller, which prioritises collision avoidance, can achieve comparable evasive manoeuvres with a lower peak sideslip angle when the cost function is designed to minimise model uncertainties. These findings underscore the complexity not only of controlling a vehicle at high sideslip angles but also of generating

an optimal reference for collision avoidance that incorporates drifting. Furthermore, these controllers have proven their ability to bring a vehicle to a high sideslip angle, but they were not designed to stabilise the vehicle to track a trajectory with a high sideslip angle. In contrast, Chapter 6 proposed a controller capable of tracking a desired path while operating at high sideslip angles, but it lacks the ability to prioritise obstacle avoidance during emergencies. This highlights two limitations of this work. The first is the possibility of combining the controllers developed in Chapters 3, 4 and 5 with the one presented in Chapter 6. The second is the possibility of designing a motion planner capable of providing a reference that pushes the vehicle into drifting as a deliberate strategy for avoiding obstacles. While the first limitation has a straightforward solution, which consists of combining the cost functions and tuning the prediction model to make it real-time feasible, addressing the second limitation is more complex and challenging. Among the existing studies on drifting motion planning, the optimisation based methods are the most popular. Most of these methods focus on lap-time optimisation [188, 189], with only a few specifically targeting collision avoidance [190, 191]. The main complexity lies in the significant computational resources required to solve such complex tasks in real-time. A potential solution could involve using the controller proposed in Chapter 3 not only for controlling the vehicle during obstacle avoidance but also as a motion planner in scenarios where drifting is required for collision avoidance. Subsequently, the controller developed in Chapter 6 could be employed to track the resulting trajectory with high sideslip angles. This combined approach could offer a more efficient way to manage the complexity of the problem. However, it introduces a hierarchical architecture, which may lead to conflicting objectives, as demonstrated in Chapter 3.

7.2.4 MACHINE LEARNING

Chapters 2 and 5 proposed solutions that integrate machine learning algorithms with vehicle dynamics models, creating hybrid approaches that are both physical consistent and data efficient. These models improve accuracy in standard driving conditions by leveraging data-driven insights while mitigating performance degradation in unforeseen scenarios through the inclusion of physics-based constraints. However, despite their data efficiency, hybrid methods still require extensive training on experimental data, which is both complex and costly to obtain. This challenge is particularly pronounced in the context of obstacle avoidance, as the dataset must capture the full range of vehicle behaviours, from conventional driving to manoeuvres at the limit of handling. The lack of datasets that contain such scenarios, especially those accounting for varying road friction coefficients, complicates the development of advanced data-driven or hybrid approaches for obstacle avoidance in automated vehicles.

To address this limitation, future work should focus on exploring new methods for constructing optimal training sets tailored to a specific machine learning algorithm and target model. The goal is to efficiently learn the target model with a minimal number of iterations and the smallest possible dataset [192]. This area of research is commonly referred to as machine teaching [193, 194], which represents the inverse problem of machine learning. In machine teaching, the objective is to identify an optimal training sequence that strategically guides the learner toward the desired outcome [195]. While traditional machine learning emphasises the development of new algorithms to enhance accuracy, machine teaching concentrates on improving the effectiveness of the learning process. It focuses on how the

machine teacher curates and presents the most relevant training examples at the appropriate stages of the learner's progression [196].

In the context of this thesis, the focus must be on determining how to obtain the optimal training set for a given machine learning algorithm. Specifically, the research should aim to identify the most effective sequence of manoeuvres or generate an optimal manoeuvre profile that enhances the training of data-driven or hybrid approaches. This approach ensures that the training process is both efficient and representative of the full range of vehicle behaviours, especially in challenging scenarios such as manoeuvres at the limit of handling. Furthermore, it would improve the interpretability of models by understanding how specific training data influences learning outcomes.

7.2.5 AUTOMATED VEHICLES

Despite setbacks that have delayed the widespread deployment of automated vehicles, Level 4 robotaxi services are already fully operational for the public in San Francisco [6], with plans to expand these services to other American cities [7]. These advancements mark a significant milestone in automated driving research, demonstrating the feasibility of the technology. However, this should not be considered the final destination but rather a new starting point for further research. While the reduced number of crashes [197] underscores the safety of automated vehicles, I believe, especially after my personal experience with a robotaxi ride in San Francisco, that the current technology remains limited in executing emergency manoeuvres to avoid obstacles. This is particularly concerning given that the risk of serious injury increases exponentially with vehicle speed [198]. Although many current robotaxi services operate at relatively low speeds, primarily in urban or sub-urban settings, this limitation in evasive manoeuvre capability becomes even more critical as automated systems begin to operate at higher speeds, such as on highways [199]. An example is the Drive Pilot for highways by Mercedes [10] and the Level 3 Honda SENSING Elite [200]. These concerns emphasise the need for future research into automated vehicle collision avoidance at the limit of handling. Advancements in this area could have a profound societal impact by reducing the number of serious injuries in vehicle crashes. As automated vehicles become more prevalent, there is also the potential for road users, including human drivers, to overtrust the technology, crossing streets or merging into traffic with reduced caution. Early signs of this overtrust can already be observed in incidents involving Tesla's Autopilot [201] and the fatal Uber robotaxi crash in 2018 [202]. As automated driving systems evolve, particularly for higher speed scenarios, ensuring the vehicles' ability to perform effective evasive manoeuvres will be vital for public safety.

7.3 GENERAL CONCLUSION

This thesis developed innovative nonlinear motion control algorithms for obstacle avoidance at the limit of handling in automated vehicles and analysed the benefits of introducing tyre force sensing technology. Where feasible, results were derived from experiments conducted on a proving ground; otherwise, they were obtained from a high-fidelity simulation environment validated with experimental data.

In Chapter 2, I introduced the first mutual hybrid approach for vehicle sideslip angle estimation, where a neural network is trained end-to-end with the Unscented Kalman

Filter. Using a large-scale experimental dataset, I demonstrated significant improvements in the accuracy of the proposed architecture when utilising tyre force sensing technology, surpassing the current state-of-the-art.

In Chapter 3, I presented a nonlinear Model Predictive Contouring Control (MPCC) algorithm for obstacle avoidance in automated vehicles operating at the limit of handling. This approach integrates motion planning, path tracking, and vehicle stability objectives, prioritising obstacle avoidance in emergencies. Through a double-lane change scenario in both high-fidelity simulations and real experiments at low speeds, the proposed controller showed significant improvements in successfully avoiding obstacles while maintaining vehicle stability.

In Chapter 4, I extended the MPCC algorithm by incorporating torque vectoring capabilities, demonstrating its benefits for vehicle stability and obstacle avoidance performance. The robustness of the approach was validated against variations in vehicle parameters, road friction, and perception and localisation errors. A statistical analysis was conducted to assess the influence of these variations on performance, providing valuable guidelines for future controller design.

In Chapter 5, I introduced a novel Learning-based MPCC algorithm for evasive manoeuvres at the limit of handling. This algorithm employs a Student-t Process (STP) to minimise the tyre model mismatches and uncertainties in real-time, utilising tyre force measurements. In high-fidelity simulations, I demonstrated the superior performance of this learning-based MPCC with STP over existing state-of-the-art methods, quantifying the advantages of innovative tyre force sensing technology.

Finally, in Chapter 6, I presented a novel approach to automated drifting using a standard passenger vehicle. Through experiments conducted on a proving ground, I demonstrated that the proposed method successfully brings and maintains the standard passenger vehicle at the desired 30 deg sideslip angle in both high and low friction conditions.

In conclusion, this thesis presented robust and novel approaches to addressing key challenges in obstacle avoidance control for automated driving. By tackling the nonlinearities in tyre-road interactions, uncertainties in the perception and localisation pipeline, and the challenge of controlling vehicles beyond stable handling limits, the proposed framework brings us one step closer to achieving safer, more reliable automated vehicles.



A TWO-STAGE BAYESIAN OPTIMISATION FOR AUTOMATIC TUNING OF AN UNSCENTED KALMAN FILTER FOR VEHICLE SIDESLIP ANGLE ESTIMATION

*Optimization hinders evolution. Everything should be built top-down, except the first time.
Simplicity does not precede complexity, but follows it.*

Epigrams on Programming, Alan Jay Perlis

ABSTRACT

This paper presents a novel methodology to auto-tune an Unscented Kalman Filter (UKF). It involves using a Two-Stage Bayesian Optimisation (TSBO), based on a student-t Process, to optimise the process noise parameters of a UKF for vehicle sideslip angle estimation. Our method minimises performance metrics, given by the average sum of the states' and measurement estimation error for various vehicle manoeuvres covering a wide range of vehicle behaviour. The predefined cost function is minimised through a TSBO which aims to find a location in the feasible region that maximises the probability of improving the current best solution. Results on an experimental dataset show the capability to tune the UKF in 79.9 % in less time than using a genetic algorithm (GA) and the overall capacity to improve the estimation performance in an experimental test dataset of 9.9 % to the current state-of-the-art GA.

A.1 INTRODUCTION

AN accurate and real-time estimation of the vehicle sideslip angle is essential to strengthen the performance of active vehicle control systems. The state-of-the-art estimation of sideslip angle mostly relies on model-based approaches using an Unscented Kalman Filter (UKF). Despite its efficiency and robustness, the process and observation noise parameters must be accurately tuned to achieve good performance [177, 178]. The noise parameters selection is particularly relevant because they need to capture the following aspects:

- Process and observation model mismatch with respect to the real vehicle dynamics.
- Discretisation error because the time-step in the Kalman filter influences the noise parameters to achieve the best state estimation [203].
- Different working conditions of the sensors installed in the vehicle.

If the tuning is obtained through numerical optimisation, the cost function is highly non-convex, non-smooth and nonlinear, calling for robust optimisation methods [177]. Several optimisation algorithms for UKF tuning [204] have been studied: GA, Sequential Quadratic Programming, Nelder-Mead, Artificial Bee Colony, Fruit Fly Optimisation (FFO) and Differential Evolution. The comparison demonstrates that each algorithm leads to different optimum locations, illustrating the importance of the optimisation method. However, a comparison regarding the amount of data and the optimisation time has been omitted. Furthermore, only stochastic optimisation algorithms have been compared that can be inefficient if the UKF takes a long simulation time. For vehicle sideslip angle estimation, large-scale training sets representing vehicle behaviour [33] make the cost function evaluation very time-consuming. In different engineering applications, several solutions have been proposed, e.g. Reinforcement Learning (RL) [205] or Bayesian Optimisation (BO) [177]. However, the majority of these are tested only on simulation data or toy examples. A more detailed analysis of previous works is in Section A.2.

This paper proposes a new UKF tuning methodology using Two-Stage Bayesian Optimisation (TSBO) [179], based on a student-t Process (STP) for vehicle sideslip angle estimation. The proposed methodology to tune the process noise parameters reduces the optimisation time and improves the optimum localisation. Both process and observation noise are assumed Gaussian, zero mean and uncorrelated, as for the majority of the model-based vehicle state estimation filters [206]. The methodology framework is represented in Fig. A.1, a more detailed explanation will be presented in Section A.4. The observation noise parameters are tuned by performing a statistical analysis of the vehicle sensor measurements [33]. The process noise parameters' tuning is based on optimising a cost function defined as the sum of the Normalised Root Mean Squared Errors (NRMSE) of the vehicle states and measurements for 8 manoeuvres. The 8 manoeuvres which compose the training set and the 23 forming the test set are selected from a real-world experimental dataset composed of 216 manoeuvres recorded at the Automotive Testing Papenburg GmbH. The manoeuvres are chosen to cover a wide range of vehicle motions. The cost function is optimised using the modified TSBO [179]. It consists of a fast exploration and a pure exploitation stage to find the optimum, reducing the number of simulations required. Differently from the TSBO proposed in [179], the surrogate model is a STP to enhance the robustness and the ability to

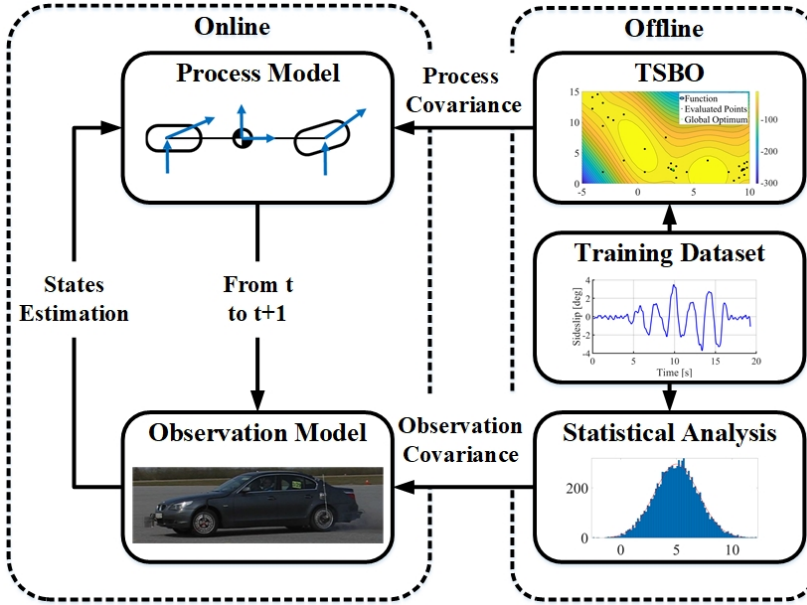


Figure A.1: Framework utilising a Two-Stage Bayesian Optimisation to tune the process noise parameters of a UKF.

capture heavy-tailed NRMSE. The methodology is experimentally validated using a UKF, but can also be applied to an Extended Kalman Filter (EKF).

A.2 RELATED WORKS

DESPITE the extensive use of UKF for vehicle sideslip angle estimation, the tuning of the process and observation noise parameters is rarely investigated in the literature. The observation noise parameters are usually tuned through statistical analysis, obtained after sensor calibration tests [33]. The only exception is when the observation noise parameters are tuned together with the process noise parameters [88]. This rarely happens because the filter performance depends mainly on the ratio between the eigenvalues of the observation and process covariance matrices. Thus, it is more convenient to fix the observation noise parameters, which can be determined through sensor calibration, and optimise the process noise parameters. The methodologies for process noise parameters tuning can be split into three different categories: *manual tuning*, *metaheuristic* optimisation and *data-driven* techniques. An overview of optimisation algorithms for UKF tuning is shown in Table A.1.

Manual tuning is a trial-and-error approach driven by the user experience. The idea is to perform a grid search in the parameter space and reach the best tuning iteratively. Despite the oversimplicity of the concept, it is still very often used, especially for vehicle sideslip angle estimation [79], because the performance of Kalman filters is stable for a range of process noise parameters settings [33]. The main drawback is the inefficiency of the approach. Moreover, there is no proof to reach optimal performance. This leads to an

Table A.1: Overview of the optimisation algorithm most used to tune the noise parameters of nonlinear Kalman filter.

Authors	Process noise tuning	Dataset
van Aalst, 2018 [79]	Manual Tuning	Experimental
Abbeel, 2005 [178]	Coordinate Ascent Algorithm	Experimental
Powell, 2002 [207]	Downhill Simplex Algorithm	Simulated
Mazzilli, 2021 [33], Oshman, 2000 [208]	Multi-Objective GA, GA	Experimental
Acosta, 2019 [204]	Fruit Fly Optimisation	Simulated
Heidfeld, 2021 [88]	Simulated Annealing	Experimental
Tang, 2021 [205]	Reinforcement Learning	Experimental
Escoriza, 2021 [209]	Recurrent Neural Network	Experimental
Chen, 2018 [177], Chen, 2019 [210]	BO based on Gaussian & student-t Process	Simulated
Current Paper	TSBO with student-t Process	Experimental

unreliable comparison when different filter architectures are analysed.

Thus, Kalman filter tuning with numerical optimisation techniques is widely used. Gradient-based optimisations are the most straightforward algorithms which can be implemented [29], but the nonconvex and nonlinear cost functions let them easily be trapped into local minima. Thus, gradient-free optimisation algorithms are introduced. One of the simplest algorithms is the coordinate descent algorithm [178], which consists of a cycling increase or decrease of each parameter of the process covariance. A similar optimisation technique is the downhill simplex algorithm [207], which consists of evaluating the cost function at the simplex vertices of the predefined sample space. After this, the algorithm should converge to the optimum through a series of movements: reflection, expansion, and contraction. Interested readers can find more information on both algorithms in [178] and [207], respectively. These algorithms may easily fall in local optima when the cost function is nonsmooth, so more advanced metaheuristic optimisations have been proposed.

Metaheuristics are procedures that can provide an acceptable solution to an optimisation problem with incomplete information about the cost function. Examples of metaheuristic algorithms are: the already cited downhill simplex algorithm [207], GA [208] or the Multi-Objective GA [33], Simulated Annealing (SA) [88] and FFO [204]. A comparison of different metaheuristic algorithms for Kalman filter tuning [204] shows that GA and FFO found the best cost function optimum. This explains why GA is the current state-of-the-art for UKF tuning for vehicle sideslip angle estimation, which is applied in many recent publications [33]. GA is an evolutionary algorithm that tries to emulate Darwin's theory of natural evolution. Due to the high number of generations required, GA can be very time-consuming when the cost function is costly to evaluate.

Data-driven optimisation algorithms such as BO are based on the creation of a surrogate model, Gaussian Processes GP [177], or STP [210], using as inputs the tuning parameters and as output the value of the cost function. The surrogate model does not only approximate the

cost function, but also it is associated with a model probability distribution. The acquisition function takes this information to find the new parameters with the highest probability of being the best new optimum. Standard BO fits the surrogate model on parameters randomly taken from the sample space, but this can lead to a non-optimal surrogate model. Furthermore, it will require many cost-function evaluations if the sample space is broad. Other data-driven approaches can be implemented; for instance, Reinforcement Learning RL can be applied to choose the parameters of the process noise, or RL can compensate for the error of the Kalman filter [205] due to not-optimal process noise parameters. A similar concept is the so-called "KalmanNet" [209], which uses a Recurrent Neural Network RNN to compute the Kalman gain, reducing the problems of non-optimal Kalman tuning. Despite the potential of RL and RNN, the amount of data required to train the Neural Networks makes them impractical for the tuning/improvement of the Kalman filter for sideslip angle estimation.

The main contributions of this paper are twofold. The first is the development of a new TSBO based on a STP for tuning the process noise parameters of UKF for vehicle sideslip angle estimation. The TSBO in this paper differs from a recent TBSO related to electronic design [179] because the surrogate model is based on a STP to increase its robustness and ability to reach a better global optimum without increasing the required computational power.

The second contribution is that the developed TSBO reduces the number of simulations required to tune the Kalman filter with respect to the state-of-the-art GA [33, 204] and it improves the accuracy of the tuning respect GA. The tuning performance is tested and validated using an experimental dataset for vehicle sideslip angle estimation.

A.3 BAYESIAN OPTIMISATION

FILTER tuning consists of minimising a nonlinear, nonconvex and nonsmooth cost function due to the reasons already expressed in Section A.1. A minimisation problem is:

$$\begin{aligned} \min_{q \in Q} J(q) \\ \text{where } Q \subseteq \mathbb{R}^d \end{aligned} \quad (\text{A.1})$$

where q is the parameter vector of dimension d , $J(q)$ is the cost function, and $Q \subseteq \mathbb{R}^d$ is the solution space. The function $J(q)$ is a "black-box" function because it is not accessible, but its outputs are observed based on some given inputs. Thus, "black-box" stochastic optimisation algorithms are considered. It is time-consuming to perform a dense sampling of the solution space Q , so the algorithm should focus on sampling the subset of Q with the highest probability to contain the cost function optimum q^* . BO can deal with incomplete and sparse knowledge of the solution thanks to its probabilistic approach. It is based on the Bayes' theorem:

$$P(\text{beliefs} \mid \text{data}) = \frac{P(\text{beliefs}) \times P(\text{data} \mid \text{beliefs})}{P(\text{data})} \quad (\text{A.2})$$

where data are the available observations of the function $J(q)$, beliefs are the beliefs of the shape of $J(q)$, and $P(\text{beliefs} \mid \text{data})$, $P(\text{data} \mid \text{beliefs})$, $P(\text{beliefs})$, and $P(\text{data})$ are the posterior, likelihood, and the marginal probabilities respectively. Thus, BO aims to find q^* learning the shape of $J(q)$ through Bayesian inference. The TSBO overcomes the standard




Figure A.2: GP and sTP model of the cost function varying the process noise parameter related with the vehicle lateral velocity. The STP allows more outliers, and it can follow the measurements better than the GP. The shadow area is the variance of the STP and the uncertainty for the GP.

BO due to its ability to define a better ratio between exploration (function shape learning) and exploitation (approaching the optimum).

BO can be split into two steps: creating a stochastic process of the cost function, called "Surrogate Model", and building an "Acquisition Function" that uses the surrogate model to approach the optimum q^* .

A.3.1 SURROGATE MODEL

BO approximates the "black-box" cost function through a stochastic process called the surrogate model. It represents a probabilistic prior over the space of functions. The prior is updated as soon as the cost function is evaluated thanks to Bayes' Theorem. Different priors are eligible for being the surrogate model; a partial list comprises GP, STP, Bayesian neural networks and polynomial chaos expansion. GP are the current state-of-art because it is a non-parametric model [179], and the Bayesian update step is analytical. Despite this, in this paper, a STP prior is selected for the surrogate model due to two reasons. The first one is that the student-t distribution allows defining the level of Kurtosis, so it allows much more likely outliers [48, 168] than GP. The second reason is that the observations directly influence the predictive covariance [48, 168]. Both properties are beneficial for the tuning of UKF because the cost function can have discontinuities, for instance, due to the noise. Fig. A.2 shows the explained differences between STP and GP modelling in an example that models the cost function varying the process noise parameter related with lateral velocity. It is visible how STP reduces the influence of the outliers, allowing the mean to have a better fitting.

The STP is defined as:

$$f(q) = STP(m(q), k(q, q'), \nu) \quad (\text{A.3})$$

where $m(q)$ represents the mean which is assumed equal to constant zero mean function because there is no prior [211], ν are the degrees of freedom of the student-t distribution, which defines the level of Kurtosis, and $k(q, q')$ is the Kernel function between the inputs pair q and q' . The choice of the Kernel is essential to have a reliable model fitting. This paper

chooses the automatic relevance determination Matérn 5/2 function, due to its generalisation capabilities and interpretability properties:

$$k_{q,q'} = k(q, q') = \sigma_f^2 \left(1 + \frac{\sqrt{5}r}{\sigma_d} + \frac{5r^2}{3\sigma_d^2} \right) e^{-\frac{\sqrt{5}r}{\sigma_d}} \quad (\text{A.4})$$

where r is the Euclidean distance between the two points q and q' calculated as:

$$r = \sqrt{(q - q')^T (q - q')} \quad (\text{A.5})$$

where σ_f and σ_d are two hyperparameters, which are trained to minimize the negative log marginal likelihood of the STP. It is essential to highlight that STP should be accurate only when there are higher chances of finding the cost function optimum. Thus, every time a new observation is available, the surrogate model is retrained. STP are implemented using the GPML Matlab Code version 4.2 [170].

A.3.2 ACQUISITION FUNCTION

The Acquisition Function (AF) is responsible for moving the BO towards the optimum region and choosing the next sample point to be evaluated. Thus, it is responsible for defining the ratio between exploration, moving the search towards the area where the surrogate model is uncertain, and exploiting, moving the search towards the area with a higher chance of improving the current optimum. Ideally, the more focus is given to the exploration stage, the higher the chances of finding the global optimum; however, it increases the computational time due to a larger number of cost function evaluations.

The next point to be evaluated (q_{t+1}) is usually obtained through an additional optimisation of the AF. This paper evaluates the AF on some candidate points obtained through the process that will be explained in Section A.4 to avoid further optimisation. The AF is evaluated on the candidate points [48, 170] according to:

$$\hat{\mu}_{STP}(q_{t+1}) = k_{q,q}^T K_{\tilde{q},\tilde{q}}^{-1} f(\tilde{q}) \quad (\text{A.6})$$

where \tilde{q} are the sampling points already evaluated, $\hat{\mu}_{STP}$ is the posterior mean of a STP for a new sampling point q and k is defined in Eq. A.4.

$$\hat{\sigma}_{STP}^2(q_{t+1}) = \left(\frac{\nu + f(\tilde{q})^T K_{\tilde{q},\tilde{q}}^{-1} J(\tilde{q}) - 2}{\nu + |D| - 2} \right) \times (k(q, q) - k_{q,\tilde{q}}^T K_{\tilde{q},\tilde{q}}^{-1} k_{q,\tilde{q}}) \quad (\text{A.7})$$

where $\hat{\sigma}_{STP}$ is the posterior covariance, and D is the following set of samples:

$$D = [(\tilde{q}_1, J(\tilde{q}_1)), (\tilde{q}_2, J(\tilde{q}_2)), (\tilde{q}_3, J(\tilde{q}_3)), \dots] \quad (\text{A.8})$$

Two AFs are selected: the Expected Improvement (EI) and the Confidence Bound Minimisation (CBM) [212].

The EI is defined as:

$$EI_{STP}(q) = \hat{\sigma}_{STP}(q) \left(\frac{\nu}{\nu - 1} \right) \left(1 + \frac{z_s^2}{\nu} \right) \phi(z_s) + [\hat{y} - \hat{\mu}_{STP}] \Phi(z_s) \quad (\text{A.9})$$

where $\phi(z_s)$ and $\Phi(z_s)$ are the probability density function and the cumulative distribution function of an univariate standard student-t's random variable, z_s , and \hat{y} is the current optimum sampled by the BO.

The CBM is formulated as:

$$CBM_{STP}(q) = \hat{\sigma}_{STP}(q) \sqrt{\beta} + [f^* - \hat{\mu}_{STP}] \quad (\text{A.10})$$

where the β defines the ratio between exploitation and exploration, and f^* is the best known optimum, representing the prior knowledge of the optimum.

A.4 TWO-STAGE BAYESIAN OPTIMISATION

THIS paper proposes an evolution of the BO to tune the process noise parameters of the UKF. It is called TSBO, code by [179], and aims to reduce the number of function evaluations and improve the ratio between exploitation and exploration. The trade-off improvement is obtained thanks to the subdivisions of the BO in two stages. The first, "fast exploration", aims to find the sample space region where the optimum is contained. The second, "pure exploitation", aims to fine-tune the restricted sample space obtained from the first stage. The first stage allows starting the surrogate model training from a single evaluation at the centre of the sample space. This allows the user not to evaluate a random set of process noises to fit the first surrogate model.

A.4.1 FAST EXPLORATION

The fast exploration aims to narrow the sample space $Q \subseteq \mathbb{R}^d$ where the BO search the optimum as fast as possible. It consists of subdividing the sample space into 2^d hyper-rectangles. Every hyper-rectangle centre point is a candidate point, which the AF evaluates to choose the next sampling point q_{t+1} where the cost function is evaluated. The following sampling point is not obtained as in the standard BO through an auxiliary optimisation of the AF, but it is the candidate point with the maximum value of the AF between the other candidate points. The cost function is only evaluated at points q_{t+1} . A new set of hyper-rectangles is generated starting from the hyper-rectangle enclosing the best candidate point. The two available AFs, EI and CBM, are alternatively used to evaluate the candidate points. After a predefined number of iterations MAX_{AF} , the AF which would have had the greatest gain to the optimisation is selected as the final AF.

The fast exploration stage is repeated until the Euclidean distance between the new best sampling point $q_{t+1, max}$ and the previous best sampling point $q_{t, max}$ is below a user-defined threshold TR_{FE} for a number of times MAX_{FE} :

$$n = \begin{cases} n + 1, & \text{if } \|q_{t+1, max} - q_{t, max}\| < TR_{FE} \\ 0, & \text{otherwise} \end{cases} \quad (\text{A.11})$$

Every time the cost function is evaluated and a new measurement is available, the surrogate model is retrained to improve the fitting with the black-box function. Algorithm 1 sums up the fast exploration stage of the TSBO.

Algorithm 1 Fast Exploration stage of the TSBO

Require: A sample space Q and an evaluated sampling point at the centre of Q
 Fit the STP Surrogate model
 Subdivide Q in 2^d hyper-rectangles
while $n \leq MAX_{FE}$ **do**
 Compute candidate points at the centre of hyper-rectangles
 Evaluate AF at the candidate points
 Pick the candidate points with the maximum AF
 if $\|q_{t+1max} - q_{tmax}\| < TR_{FE}$ **then**
 $n + 1 \leftarrow n$
 else if $\|q_{t+1max} - q_{tmax}\| \geq TR_{FE}$ **then**
 $n \leftarrow n$
 end if
 Evaluate the cost function in the sampling points
 Train the STP of the Surrogate with the new observation
 Subdivide the hyper-rectangle which were enclosing the best candidate points in 2^d
 hyper-rectangles
end while

Algorithm 2 Pure Exploitation stage of the TSBO

Require: $n = MAX_{FE}$
while $n_{iter} \leq MAX_{PE}$, and the output from Algorithm 1 **do**
 Compute candidate points at the centre of hyper-rectangles
 Evaluate AF at the candidate points
 Pick the candidate points with the maximum AF
 Evaluate the cost function in the sampling points
 if $n_{iter} \leq MAX_{SM}$ **then**
 Train the STP of the Surrogate with the new observation
 end if
 Subdivide the hyper-rectangle which were enclosing the best candidate points in 3^d
 hyper-rectangles
end while

A.4.2 PURE EXPLOITATION

The pure exploitation stage follows a procedure similar to fast exploration, but it specialises in refining the optimum position. The first step is the computation of the sample space \tilde{Q}^d around the optimum computed in the fast exploration stage:

$$\tilde{Q}^d = \begin{bmatrix} (1 - \alpha) q_1^* & (1 + \alpha) q_1^* \\ \vdots & \vdots \\ (1 - \alpha) q_d^* & (1 + \alpha) q_d^* \end{bmatrix} \quad (\text{A.12})$$

where α is a hyper-parameter that defines how wide the sample space should be, and q_i^* is the optimum sampling point obtained in the fast exploration stage. The new sample space

\tilde{Q}^d is divided into 3^d hyper-rectangles to have more candidate points to evaluate. The higher number of candidate points helps the pure exploitation stage to converge faster to the global optimum of the function. The rest of the procedure is the same of the fast exploration. The only difference is that when the number of evaluated sampling points n_{iter} is higher than a predefined threshold MAX_{SM} , the STP is no more retrained. Algorithm 2 summarises the pure exploitation stage of the TSBO, which is iterated until n_{iter} is lower than a user-defined threshold MAX_{PE} .

A.5 APPLICATION: UKF TUNING

THE proposed TSBO based on STP is tested on the tuning of the process noise parameters of a UKF for vehicle sideslip angle estimation. In this paper, the single-track model with tyre axle forces computed by the Dugoff tyre is chosen as the vehicle model. The states are the longitudinal velocity at the Centre-of-Gravity CoG V_x , the lateral velocity at the CoG V_y and the yaw rate $\dot{\psi}$ while the vehicle measurements are the V_x , the lateral acceleration at the CoG a_y and the ψ . The vehicle sideslip angle β_s is obtained as $\beta_s = \arctan \frac{V_y}{V_x}$:

A.5.1 DATASET

All the manoeuvres that compose the training set and test set are taken from a real-world experimental dataset composed of 216 manoeuvres, which correspond to 2 hrs of driving. It puts together standard vehicle dynamics manoeuvres, e.g. double-lane change, slalom, random steer, J-turn, spiral, braking in a turn, and steady-state circular tests, together with recorded laps at the Papenburg track. All manoeuvres were performed on dry asphalt with tyres inflated according to the vehicle specifications.

All manoeuvres are recorded using the BMW test platform instrumented by the inertial measurement unit, wheel force transducers for all four wheels, GPS and an optical sensor from Corrsys-Datron to measure the sideslip angle. The high-end optical speed sensor is not present in consumer vehicles and will be used as a reference for TBSO tuning and validation.

The UKF should be tuned for different working conditions so various vehicle manoeuvres are jointly considered in the training set. The training set is composed of 8 manoeuvres: 1 braking in a turn, 1 skidpad, 2 J-turn, 2 slaloms and 2 lane changes.

A test set composed of 23 manoeuvres is taken to prove the generalisation capability of the approach. All 23 manoeuvres of the test set are different from the training set, even if some of them have similar characteristics. The manoeuvres are: 2 braking in a turn, 2 skidpad, 5 J-turn, 4 slalom, 4 lane change, 2 random steer, 1 lap track and 3 spiral. The characteristics of the manoeuvres cover a wide range of vehicle driving behaviours and different settings of the Electronic Stability Control.

A.5.2 COST FUNCTION FORMULATION

In this paper, the UKF tuning aims to improve the accuracy of the vehicle state estimation; therefore, this paper focuses on an objective function defined by minimising the Normalised Root Mean Squared Error NRMSE rather than maximising the negative log-likelihood [178].

The minimisation problem is described as:

$$\min_{q \in Q} J(q)$$

$$\text{where } Q = \begin{bmatrix} q_{1,min} & q_{1,max} \\ \vdots & \vdots \\ q_{d,min} & q_{d,max} \end{bmatrix} \quad (\text{A.13})$$

where J is the cost function defined as:

$$J = W_1 E_\beta + W_2 E_\psi + W_3 E_{a_y}$$

$$\text{where } E_j = \sqrt{\frac{\sum_{i=1}^{N_{man}} (NRMS E_{j,i})^2 N_{s,i}}{\sum_{i=1}^{N_{man}} N_{s,i}}} \quad (\text{A.14})$$

where W_1, W_2, W_3 are the weights of the cost function, N_s is the length of the manoeuvre measured in sampling points, N_{man} is the number of manoeuvres considered in the training set, and $NRMS E$ is computed as follow:

$$NRMS E = \sqrt{\frac{\sum_{k=1}^{N_s} (\hat{X}_k - X_k)^2}{N_s}} \frac{1}{\max(|X|)} \quad (\text{A.15})$$

where X and \hat{X} are the vectors of the measured and estimated states respectively.

A.5.3 KEY PERFORMANCE INDICATOR

The accuracy of the estimation of sideslip angle is evaluated through the use of the root mean squared error RMSE, maximum absolute error MAE, the RMSE and the MAE when the absolute vehicle acceleration a_y is greater than 4 m/s^2 , which are respectively abbreviated with $RMSE_{nl}$ and MAE_{nl} . The latter is used to analyse how the UKF tuning is especially relevant when the vehicle behaves in a nonlinear way.

A.6 RESULTS

THE performance in terms of minimisation time and optimum localisation of the TSBO based on STP is compared with the current state-of-the-art technique GA and with a TSBO based on GP. All optimisations perform in the same sample space Q where all the process noise parameters for V_x, V_y and $\dot{\psi}$ vary from 10^{-10} to 1.

The optimisation parameters for TSBO based on STP and GP are summarised in Table A.2. The parameters for GA are summarised in Table A.3.

The optimum cost and optimisation time for the TSBO based on STP, TSBO based on GP and GA are summarised in Table A.4 and Fig. A.3.

Fig. A.3 shows how both TSBOs approach the global optimum faster than GA. Where TSBO converges after 40 evaluations of the cost function, the GA needs at least 220 evaluations, resulting in TSBO being on average 79 % faster than GA. Thus, TSBOs are particularly convenient when the cost function evaluation is computationally heavy, for instance, due to numerous manoeuvres considered in the training set. Despite the faster optimisation,

Table A.2: User-defined parameters for TSBO optimisations.

Parameters	TSBO - STP	TSBO - GP
Likelihood function	student-t 's	Gaussian
Inference	Variational Bayes	Exact
Kernel function	Matérn ARD 5/2	Matérn ARD 5/2
Mean function	Constant	Constant
ν	15	NA
MAX_{FE}	15	15
MAX_{PE}	40	40
MAX_{SM}	38	38
TR_{FE}	$0.01 \parallel q^* \parallel$	$0.01 \parallel q^* \parallel$
β	0.01	0.01
α	0.15	0.15
W_1, W_2, W_3	5, 1, 1	5, 1, 1

Table A.3: User-defined parameters for GA optimisation.

Parameters	GA
Population size	15
Max generations	15
Élite count	0.75
Crossover fraction	0.8
W_1, W_2, W_3	5, 1, 1

TSBO based on STP reaches an optimum 16.4 % lower than the GA optimum. The proposed algorithm outperforms the state-of-the-art approach both in terms of optimisation time and optimisation accuracy. TSBO based on GP is 4.4 % faster than TSBO based on STP. The rise in optimisation time is due to the surrogate model's training phase, which is slower for the STP due to the more optimisation parameters. The better optimisation accuracy, 21 %, of the STP than the GP is due to the higher Kurtosis of the student-t's distribution which makes the nonparametric model capable of dealing with outliers. A much higher optimisation accuracy compensates for the slight higher computational time.

Table A.5 presents the RMSE, MAE, $RMSE_{nl}$ and MAE_{nl} of the sideslip angle for the various optimised UKFs. The reported values are the average of all manoeuvres composing the training set. Table A.5 highlights how the cost function's minimisation is related to the KPIs used to compare the estimation accuracy because TSBO based on STP outperforms the other algorithms. On average, the RMSE of the sideslip angle optimised with tSBO based on STP is 15.4 % lower than the one with GA, and the MAE is 15 % lower. Particularly important is the improvement of the $RMSE_{nl}$ and of the MAE_{nl} because when the $|a_y| \geq 4 \text{ m/s}^2$ the sideslip angle estimation becomes more critical for vehicle stability control. The

Table A.4: Numerical results of the optimum cost function and optimisation time using different algorithms for UKF tuning.

Algorithm	Total cost [-]	Optim. time [s]
Manual tuning	1.714	NA
GA	0.378	1599
TSBO - GP	0.400	306
TSBO - STP	0.316	320




Figure A.3: Performance comparison of TSBO based on STP, TSBO based on GP and GA.

UKF optimised by the TSBO based on STP reduces the $RMSE_{nl}$, and the MAE_{nl} of the UKF optimised by GA of 24.6 % and 12.5 %, respectively. The accuracy of the UKF tuned by GA is higher than the one optimised by TSBO based on GP.

Table A.6 reports the same information as Table A.5, but for the test set, to check that the numerical optimisation algorithms do not overfit the training set leading to worse performance in other manoeuvres. Table A.6 shows that TSBO based on STP still behaves better than all other algorithms for the 4 KPIs, proving better performance in tuning the process noise UKF parameters.

On average, the RMSE of the sideslip angle optimised with tSBO based on STP is 9.9 % lower than the one with GA, and the MAE is 17.6 % lower. The improvement of the $RMSE_{nl}$ and of the MAE_{nl} is particularly significant because when the $|a_y| \geq 4 \text{ m/s}^2$, the sideslip angle estimation becomes more critical for vehicle stability control. The UKF optimised with the TSBO based on STP reduces the $RMSE_{nl}$, and the MAE_{nl} of the UKF optimised with GA of 10.6 % and 9.8 %, respectively. The improvement in the test set is lower than the one obtained in the training set but still particularly relevant, especially for what concerns the MAE.

Fig. A.4 shows the relative improvement of the UKF tuned by TSBO based on STP and GA in 6 different manoeuvres of the test set. The relative improvement is measured with the

Table A.5: Comparison of the average KPIs for the UKF tuned by manual tuning, GA, TSBO based on GP and TSBO based on STP. Results of the training set.

KPI	Manual	GA	TSBO - GP	TSBO - STP
RMSE [deg]	1.028	0.500	0.634	0.423
MAE [deg]	1.196	0.900	1.252	0.765
RMSE _{nl} [deg]	0.716	0.568	0.855	0.428
MAE _{nl} [deg]	0.818	0.671	0.961	0.587

Table A.6: Comparison of the average KPIs for the UKF tuned by manual tuning, GA, TSBO based on GP and TSBO based on STP. Results of the test set.

KPI	Manual	GA	TSBO - GP	TSBO - STP
RMSE [deg]	0.577	0.383	0.431	0.345
RMSE [deg]	0.672	0.529	0.689	0.436
RMSE _{nl} [deg]	1.328	0.860	1.030	0.769
MAE _{nl} [deg]	1.047	0.746	0.902	0.673

RMSE of the sideslip angle, a_y and $\dot{\psi}$. In 5 out of 6 manoeuvres the TSBO reaches higher performance than the GA. The spiral manoeuvre is the only one which shows a degraded accuracy in the RMSE of the sideslip angle of 8 %. Performance degradation in this particular manoeuvre is considered acceptable due to the various driving conditions of the test set. The relative improvement is significant, especially for the sideslip angle and for the a_y , while the $\dot{\psi}$ does not have this. The reason is that the $\dot{\psi}$ has a lower maximum absolute error compared with the other variables, so the relative improvement is less noticeable. Fig. A.4 shows how TSBO can enhance the performance of the UKF for both transient (J-turn, slalom) and steady state (skidpad) manoeuvres. Significant is the improvement in the skidpad; see also Fig. A.5. The reason is that during a steady-state manoeuvre, the vehicle yaw acceleration is almost null, so the difference between estimated and measured tyre forces becomes essential for the V_y estimation. The best way to reduce the tyre force error is to optimise the process noise parameters.

The improvement in sideslip angle estimation is also noticeable for a J-turn and a slalom manoeuvre, see respectively Figs. A.6b and A.6a. Fig. A.6a shows how the improvement in the tuning of the UKF is noticeable at the sideslip angle peaks of the slalom manoeuvre. The improvement at the peaks is more remarkable when the sideslip angle is above 3 deg, while when it is between 0 and 3 deg, the difference is not visible due to the lower absolute error of both UKFs. Fig. A.6b shows the sideslip angle estimation in J-turn, and it highlights another time how the UKF tuned by TSBO based on STP reduces the estimation error, especially when the sideslip angle is above 4 deg.




Figure A.4: Relative improvement of the UKF tuned by TSBO based on STP with respect to the UKF tuned by the state-of-the-art GA. The improvement is measured through the sideslip angle's RMSE for 6 different manoeuvres chosen from the test set.




Figure A.5: Sideslip angle estimation in a Skidpad manoeuvre.





Figure A.6: Sideslip angle estimation in a slalom manoeuvre (Fig. A.6a) and in a J-turn manoeuvre (Fig. A.6b).

A.7 CONCLUSIONS

This paper presented a two-stage Bayesian optimisation approach based on the student-t process to tune the process noise parameters of an unscented Kalman filter for sideslip angle estimation. The proposed approach reaches the cost function optimum 79.9 % faster than the state-of-the-art GA and the optimum is 16.4 % better than the one obtained by GA. The obtained results are tested on an experimental test set composed of transient and quasi-steady state manoeuvres. Furthermore, two-stage Bayesian optimisation creates a physical representation of the black-box cost function thanks to surrogate model training. The performance of the two-stage Bayesian optimisation based on the student-t process is also compared with the two-stage Bayesian optimisation based on the Gaussian process. The algorithm based on the student-t process is 4.4 % slower than the one based on the Gaussian process, but it reaches a 21 % better optimum. As future research, the plan is to tune not only the process noise parameters of the unscented Kalman filter but also the parameters defining the location of the sigma points.

ACKNOWLEDGEMENT

The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484).

REFERENCES

- [1] Leaders. Driverless cars are stuck in a jam. <https://www.economist.com/leaders/2019/10/12/driverless-cars-are-stuck-in-a-jam>, October 2019.
- [2] Cameron F. Kerry and Jack Karsten. Gauging investment in self-driving cars. <https://www.brookings.edu/articles/gauging-investment-in-self-driving-cars/>, October 2017.
- [3] Jennifer Elias. Alphabet’s waymo valuation cut 40 <https://www.cnbc.com/2019/09/27/waymo-valuation-cut-40percent-by-morgan-stanley-to-105-billion.html>, September 2019.
- [4] Omayma Alqatawneh, Alex Coles, and Ertu Unver. Towards the adoption of self-driving cars. In *Smart Trends in Computing and Communications*, pages 455–461, 2020.
- [5] Johannes Deichmann, Eike Ebel, Kersten Heineke, Ruth Heuss, Martin Kellner, and Fabian Steiner. Autonomous driving’s future: Convenient and connected. <https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/autonomous-drivings-future-convenient-and-connected>, January 2023.
- [6] The Waymo Team. Waymo one is now open to everyone in san francisco. <https://waymo.com/blog/2024/06/waymo-one-is-now-open-to-everyone-in-san-francisco/>, June 2024.
- [7] The Waymo and Uber Teams. Waymo and uber expand partnership to bring autonomous ride-hailing to austin and atlanta. <https://waymo.com/blog/2024/09/waymo-and-uber-expand-partnership/>, September 2024.
- [8] SAE International. Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. https://doi.org/10.4271/J3016_202104, 2021.
- [9] The Tesla Team. The bigger picture on autopilot safety. <https://www.tesla.com/blog/bigger-picture-autopilot-safety>, 2024.
- [10] Mercedes-Benz Research & Development North America. Introducing DRIVE PILOT: An automated driving system for the highway. <https://group.mercedes-benz.com/dokumente/innovation/>

- [sonstiges/2023-03-06-vssa-mercedes-benz-drive-pilot.pdf](#), 2023.
- [11] Rasheed Hussain, JooYoung Lee, and Sherali Zeadally. Autonomous Cars: Social and Economic Implications. *IT Professional*, 20(6):70–77, 2018.
 - [12] Il Bae, Jaeyoung Moon, and Jeongseok Seo. Toward a Comfortable Driving Experience for a Self-Driving Shuttle Bus. *Electronics*, 8(9), 2019.
 - [13] Alexandre Moreira Nascimento, Lucio Flavio Vismari, Caroline Bianca Santos Tancredi Molina, Paulo Sergio Cugnasca, João Batista Camargo, Jorge Rady de Almeida, Rafia Inam, Elena Fersman, Maria Valeria Marquezini, and Alberto Yukinobu Hata. A Systematic Literature Review About the Impact of Artificial Intelligence on Autonomous Vehicle Safety. *IEEE Transactions on Intelligent Transportation Systems*, 21(12):4928–4946, 2020.
 - [14] Esko Lehtonen, Fanny Malin, Tyron Louw, Yee Mun Lee, Teemu Itkonen, and Satu Innamaa. Why would people want to travel more with automated cars? *Transportation Research Part F: Traffic Psychology and Behaviour*, 89:143–154, 2022.
 - [15] Kareem Othman. Exploring the implications of autonomous vehicles: A comprehensive review. *Innovative Infrastructure Solutions*, 7(2):165, 2022.
 - [16] Santokh Singh. Critical reasons for crashes investigated in the national motor vehicle crash causation survey. Technical report, National Highway Traffic Safety Administration, 2015.
 - [17] Jessica B. Cicchino. Effectiveness of forward collision warning and autonomous emergency braking systems in reducing front-to-rear crash rates. *Accident Analysis & Prevention*, 99:142–152, 2017.
 - [18] US Department of Transportation. Preparing for the future of transportation: Automated vehicles 3.0. Technical report, US Department of Transportation Washington, DC, 2018.
 - [19] Alexandra S. Mueller, Jessica B. Cicchino, and David S. Zubry. What humanlike errors do autonomous vehicles need to avoid to maximize safety? *Journal of Safety Research*, 75:310–318, 2020.
 - [20] Chongfeng Wei, Richard Romano, Natasha Merat, Yafei Wang, Chuan Hu, Hamid Taghavifar, Foroogh Hajiseyedjavadi, and Erwin R. Boer. Risk-based autonomous vehicle motion control with considering human driver’s behaviour. *Transportation Research Part C: Emerging Technologies*, 107:1–14, 2019.
 - [21] Diomidis I. Katzourakis, Joost C.F. de Winter, Mohsen Alirezaei, Matteo Corno, and Riender Happee. Road-departure prevention in an emergency obstacle avoidance situation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 44(5):621–629, 2013.

-
- [22] Kristofer D. Kusano, John M. Scanlon, Yin-Hsiu Chen, Timothy L. McMurphy, Ruoshu Chen, Tilia Gode, and Trent Victor. Comparison of waymo rider-only crash data to human benchmarks at 7.1 million miles. *arXiv preprint arXiv:2312.12675*, 2023.
 - [23] John M. Scanlon, Kristofer D. Kusano, Laura A. Fraade-Blanar, Timothy L. McMurphy, Yin-Hsiu Chen, and Trent Victor. Benchmarks for retrospective automated driving system crash rate analysis using police-reported crash data. *arXiv preprint arXiv:2312.13228*, 2023.
 - [24] Nick Webb, Dan Smith, Christopher Ludwick, Trent Victor, Qi Hommes, Francesca Favaro, George Ivanov, and Tom Daniel. Waymo’s safety methodologies and safety readiness determinations. *arXiv preprint arXiv:2011.00054*, 2020.
 - [25] John M. Scanlon, Kristofer D. Kusano, Tom Daniel, Christopher Alderson, Alexander Ogle, and Trent Victor. Waymo simulated driving behavior in reconstructed fatal crashes within an autonomous vehicle operating domain. *Accident Analysis & Prevention*, 163:106454, 2021.
 - [26] Joseph Funke, Matthew Brown, Stephen M. Erlien, and J. Christian Gerdes. Collision avoidance and stabilization for autonomous vehicles in emergency scenarios. *IEEE Transactions on Control Systems Technology*, 25(4):1204–1216, 2016.
 - [27] Anders Lie, Claes Tingvall, Maria Krafft, and Anders Kullgren. The effectiveness of ESP (electronic stability program) in reducing real life accidents. *Traffic Injury Prevention*, 5(1):37–41, 2004.
 - [28] Stijn Kerst, Barys Shyrokau, and Edward Holweg. A semi-analytical bearing model considering outer race flexibility for model based bearing load monitoring. *Mechanical Systems and Signal Processing*, 104:384–397, 2018.
 - [29] Stijn Kerst, Barys Shyrokau, and Edward Holweg. A model-based approach for the estimation of bearing forces and moments using outer ring deformation. *IEEE Transactions on Industrial Electronics*, 67(1):461–470, 2020.
 - [30] Stijn Kerst, Barys Shyrokau, and Edward Holweg. Reconstruction of wheel forces using an intelligent bearing. *SAE International Journal of Passenger Cars-Electronic and Electrical Systems*, 9(2016-01-0092):196–203, 2016.
 - [31] Hojong Lee and Saied Taheri. Intelligent tires? a review of tire characterization literature. *IEEE Intelligent Transportation Systems Magazine*, 9(2):114–135, 2017.
 - [32] Hassan Hariri, Jiseok Kim, Woo Soo Kim, Luc G. Frechette, and Patrice Masson. Performance validation of printed strain sensors for active control of intelligent tires. *Applied Acoustics*, 123:73–84, 2017.
 - [33] Victor Mazzilli, Davide Ivone, Stefano De Pinto, Leonardo Pascali, Michele Contrino, Giulio Tarquinio, Patrick Gruber, and Aldo Sorniotti. On the benefit of smart tyre technology on vehicle state estimation. *Vehicle System Dynamics*, 60(11):1–26, 2021.

- [34] Francesco Pretagostini, Laura Ferranti, Giovanni Berardo, Valentin Ivanov, and Barys Shyrokau. Survey on wheel slip control design strategies, evaluation and application to antilock braking systems. *IEEE Access*, 8:10951–10970, 2020.
- [35] Havard Fjaer Grip, Lars Imsland, Tor A. Johansen, Jens C. Kalkkuhl, and Avshalom Suissa. Vehicle sideslip estimation. *IEEE Control Systems Magazine*, 29(5):36–52, 2009.
- [36] Jizheng Liu, Zhenpo Wang, Lei Zhang, and Paul Walker. Sideslip angle estimation of ground vehicles: a comparative study. *IET Control Theory and Applications*, 14(20):3490–3505, 2020.
- [37] Marco Viehweger, Cyrano Vaseur, Sebastiaan van Aalst, Manuel Acosta, Enrico Regolin, Angel Alatorre, Wim Desmet, Frank Naets, Valentin Ivanov, Antonella Ferrara, et al. Vehicle state and tyre force estimation: demonstrations and guidelines. *Vehicle System Dynamics*, 59(5):675–702, 2021.
- [38] Matthew Brown and J Christian Gerdes. Coordinating tire forces to avoid obstacles using nonlinear model predictive control. *IEEE Transactions on Intelligent Vehicles*, 5(1):21–31, 2019.
- [39] Paolo Falcone, Francesco Borrelli, H. Eric Tseng, Jamal Asgari, and Davor Hrovat. A hierarchical model predictive control framework for autonomous ground vehicles. In *American Control Conference*, pages 3719–3724, 2008.
- [40] Efstathios Velenis, Diomidis Katzourakis, Emilio Frazzoli, Panagiotis Tsiotras, and Riender Happee. Steady-state drifting stabilization of rwd vehicles. *Control Engineering Practice*, 19(11):1363–1376, 2011.
- [41] Alberto Bertipaglia, Mohsen Alirezaei, Riender Happee, and Barys Shyrokau. Model predictive contouring control for vehicle obstacle avoidance at the limit of handling. In *Symposium on the Dynamics of Vehicle on Roads and on Tracks*, 2023.
- [42] Daan Lenssen, Alberto Bertipaglia, Felipe Santafo, and Barys Shyrokau. Combined path following and vehicle stability control using model predictive control. In *SAE Technical Paper*, 2023.
- [43] Mujahid Abdulrahim. On the dynamics of automobile drifting. In *SAE Technical Paper*, 2006.
- [44] Rami Y. Hindiyeh and J. Christian Gerdes. A controller framework for autonomous drifting: Design, stability, and experimental validation. *Journal of Dynamic Systems, Measurement, and Control*, 136(5):051015, 2014.
- [45] Bruno Brito, Boaz Floor, Laura Ferranti, and Javier Alonso-Mora. Model predictive contouring control for collision avoidance in unstructured dynamic environments. *IEEE Robotics and Automation Letters*, 4(4):4459–4466, 2019.

-
- [46] Alexander Liniger, Alexander Domahidi, and Manfred Morari. Optimization-based autonomous racing of 1: 43 scale rc cars. *Optimal Control Applications and Methods*, 36(5):628–647, 2015.
 - [47] Juraj Kabzan, Lukas Hewing, Alexander Liniger, and Melanie N. Zeilinger. Learning-based model predictive control for autonomous racing. *IEEE Robotics and Automation Letters*, 4(4):3363–3370, 2019.
 - [48] Brendan D. Tracey and David Wolpert. Upgrading from gaussian processes to student-t processes. In *AIAA Non-Deterministic Approaches Conference*, page 1659, 2018.
 - [49] Erik Wachter, Mohsen Alirezaei, Fredrik Bruzelius, and Antoine Schmeitz. Path control in limit handling and drifting conditions using state dependent riccati equation technique. *Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 234(2-3):783–791, 2020.
 - [50] Manuel Acosta, Stratis Kanarachos, and Michael E. Fitzpatrick. On full magv lateral dynamics exploitation: Autonomous drift control. In *International Workshop on Advanced Motion Control*, 2018.
 - [51] Pietro Stano, Davide Tavernini, Umberto Montanaro, Manuela Tufo, Giovanni Fiengo, Luigi Novella, and Aldo Sorniotti. Enhanced active safety through integrated autonomous drifting and direct yaw moment control via nonlinear model predictive control. *IEEE Transactions on Intelligent Vehicles*, 9(2):4172–4190, 2024.
 - [52] Mart Baars, Hans Hellendoorn, and Mohsen Alirezaei. Control of a scaled vehicle in and beyond stable limit handling. *IEEE Transactions on Vehicular Technology*, 70(7):6427–6437, 2021.
 - [53] Jonathan Y. Goh and J. Christian Gerdes. Simultaneous stabilization and tracking of basic automobile drifting trajectories. In *IEEE Intelligent Vehicles Symposium*, 2016.
 - [54] Jonathan Y. Goh, Tushar Goel, and J. Christian Gerdes. Toward automated vehicle control beyond the stability limits: drifting along a general path. *Journal of Dynamic Systems, Measurement, and Control*, 142(2):021004, 2020.
 - [55] Alberto Bertipaglia, Mohsen Alirezaei, Riender Happee, and Barys Shyrokau. An unscented kalman filter-informed neural network for vehicle sideslip angle estimation. *IEEE Transactions on Vehicular Technology*, 73(9):12731–12746, 2024.
 - [56] Alberto Bertipaglia, Denis de Mol, Mohsen Alirezaei, Riender Happee, and Barys Shyrokau. Model-based vs data-driven estimation of vehicle sideslip angle and benefits of tyre force measurements. In *International Symposium on Advanced Vehicle Control*, 2022.
 - [57] Nishant Chowdhri, Laura Ferranti, Felipe Santafé Iribarren, and Barys Shyrokau. Integrated nonlinear model predictive control for automated driving. *Control Engineering Practice*, 106:104654, 2021.

- [58] Marsie Trego Peterson, Tushar Goel, and J. Christian Gerdes. Exploiting linear structure for precision control of highly nonlinear vehicle dynamics. *IEEE Transactions on Intelligent Vehicles*, 8(2):1852–1862, 2023.
- [59] Yi-Wen Liao and Francesco Borrelli. An adaptive approach to real-time estimation of vehicle sideslip, road bank angles, and sensor bias. *IEEE Transactions on Vehicular Technology*, 68(8):7443–7454, 2019.
- [60] Torben Gräber, Stefan Lupberger, Michael Unterreiner, and Dieter Schramm. A hybrid approach to side-slip angle estimation with recurrent neural networks and kinematic vehicle models. *IEEE Transactions on Intelligent Vehicles*, 4(1):39–47, 2018.
- [61] Junjian Qi, Ahmad F. Taha, and Jianhui Wang. Comparing kalman filters and observers for power system dynamic state estimation with model uncertainty and malicious cyber attacks. *IEEE Access*, 6:77155–77168, 2018.
- [62] Daniel Chindamo, Basilio Lenzo, and Marco Gadola. On the vehicle sideslip angle estimation: a literature review of methods, models, and innovations. *Applied Sciences*, 8(3):355, 2018.
- [63] Dongchan Kim, Kyushik Min, Hayoung Kim, and Kunsoo Huh. Vehicle sideslip angle estimation using deep ensemble-based adaptive kalman filter. *Mechanical Systems and Signal Processing*, 144:106862, 2020.
- [64] Dongchan Kim, Gihoon Kim, Seungwon Choi, and Kunsoo Huh. An integrated deep ensemble-unscented kalman filter for sideslip angle estimation with sensor filtering network. *IEEE Access*, 9:149681–149689, 2021.
- [65] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [66] Beatriz L Boada, María Jesús L Boada, Antonio Gauchía, Ester Olmeda, and Vicente Díaz. Sideslip angle estimator based on anfis for vehicle handling and stability. *Journal of Mechanical Science and Technology*, 29(4):1473–1481, 2015.
- [67] Beatriz L Boada, María Jesús López Boada, and V Diaz. Vehicle sideslip angle measurement based on sensor data fusion using an integrated anfis and an unscented kalman filter algorithm. *Mechanical Systems and Signal Processing*, 72:832–845, 2016.
- [68] Philipp Maximilian Sieberg, Sebastian Blume, Nele Harnack, Niko Maas, and Dieter Schramm. Hybrid state estimation combining artificial neural network and physical model. In *IEEE Intelligent Transportation Systems Conference*, pages 894–899, 2019.
- [69] Ari J. Tuononen. Vehicle lateral state estimation based on measured tyre forces. *Sensors*, 9(11):8761–8775, 2009.

-
- [70] Anil Kunnappillil Madhusudhanan, Matteo Corno, and Edward Holweg. Vehicle sideslip estimator using load sensing bearings. *Control Engineering Practice*, 54:46–57, 2016.
 - [71] Donald Selmanaj, Matteo Corno, Giulio Panzani, and Sergio M. Savaresi. Vehicle sideslip estimation: A kinematic based approach. *Control Engineering Practice*, 67:1–12, 2017.
 - [72] Wei Liu, Xin Xia, Lu Xiong, Yishi Lu, Letian Gao, and Zhuoping Yu. Automated vehicle sideslip angle estimation considering signal measurement characteristic. *IEEE Sensors Journal*, 21(19):21675–21687, 2021.
 - [73] Xin Xia, Ehsan Hashemi, Lu Xiong, Amir Khajepour, and Nan Xu. Autonomous vehicles sideslip angle estimation: Single antenna gnss/imu fusion with observability analysis. *IEEE Internet of Things Journal*, 8(19):14845–14859, 2021.
 - [74] Wenfeng Li, Zhengchao Xie, Pak Kin Wong, Yunfeng Hu, Ge Guo, and Jing Zhao. Event-triggered asynchronous fuzzy filtering for vehicle sideslip angle estimation with data quantization and dropouts. *IEEE Transactions on Fuzzy Systems*, 30(8):2822–2836, 2022.
 - [75] Xiaolin Ding, Zhenpo Wang, and Lei Zhang. Event-triggered vehicle sideslip angle estimation based on low-cost sensors. *IEEE Transactions on Industrial Informatics*, 18(7):4466–4476, 2022.
 - [76] Alexander Wischnewski, Tim Stahl, Johannes Betz, and Boris Lohmann. Vehicle dynamics state estimation and localization for high performance race cars. *IFAC-PapersOnLine*, 52(8):154–161, 2019.
 - [77] Rui Song and Yongchun Fang. Vehicle state estimation for INS/GPS aided by sensors fusion and SCKF-based algorithm. *Mechanical Systems and Signal Processing*, 150:107315, 2021.
 - [78] Jong-Hwa Yoon and Huei Peng. Robust vehicle sideslip angle estimation through a disturbance rejection filter that integrates a magnetometer with gps. *IEEE Transactions on Intelligent Transportation Systems*, 15(1):191–204, 2013.
 - [79] Sebastiaan van Aalst, Frank Naets, Boulaid Boulkroune, Wouter De Nijs, and Wim Desmet. An adaptive vehicle sideslip estimator for reliable estimation in low and high excitation driving. *IFAC-PapersOnLine*, 51(9):243–248, 2018.
 - [80] Marco Gadola, Daniel Chindamo, Matteo Romano, and Fabrizio Padula. Development and validation of a kalman filter-based model for vehicle slip angle estimation. *Vehicle System Dynamics*, 52(1):68–84, 2014.
 - [81] Cyrano Vaseur, Sebastiaan Van Aalst, and Wim Desmet. Robust vehicle state and tire force estimation: Highlights on effects of road angles and sensor performance. In *IEEE Intelligent Vehicles Symposium*, pages 28–33, 2021.

- [82] Liang Li, Gang Jia, Xu Ran, Jian Song, and Kaihui Wu. A variable structure extended kalman filter for vehicle sideslip angle estimation on a low friction road. *Vehicle System Dynamics*, 52(2):280–308, 2014.
- [83] Federico Cheli, Davide Ivone, and Edoardo Sabbioni. Smart tyre induced benefits in sideslip angle and friction coefficient estimation. In *Sensors and Instrumentation*, volume 5, pages 73–83, 2015.
- [84] Robbin van Hoek, Mohsen Alirezaei, Antoine Schmeitz, and Henk Nijmeijer. Vehicle state estimation using a state-dependent riccati equation. *IFAC-PapersOnLine*, 50(1):3388–3393, 2017.
- [85] Takayuki Ando, Hirotaka Mukumoto, Keiji Aoki, Shogo Okazaki, Tomohiko Nagao, Hitoshi Aoyama, Michiharu Yamamoto, and Kimihiko Nakano. Localization using global magnetic positioning system for automated driving bus and intervals for magnetic markers. *IEEE Transactions on Intelligent Vehicles*, 8(1):878–887, 2023.
- [86] Moustapha Doumiani, Alessandro Correa Victorino, Ali Charara, and Daniel Lechner. Onboard real-time estimation of vehicle lateral tire-road forces and sideslip angle. *IEEE/ASME Transactions on Mechatronics*, 16(4):601–614, 2010.
- [87] Sergiy Antonov, Achim Fehn, and Andreas Kugi. Unscented kalman filter for vehicle state estimation. *Vehicle System Dynamics*, 49(9):1497–1520, 2011.
- [88] Hannes Heidfeld and Martin Schünemann. Optimization-based tuning of a hybrid UKF state estimator with tire model adaption for an all wheel drive electric vehicle. *Energies*, 14(5):1396, 2021.
- [89] Kanwar Bharat Singh. Vehicle sideslip angle estimation based on tire model adaptation. *Electronics*, 8(2):199, 2019.
- [90] Jakob Bechtoff and Rolf Isermann. Cornering stiffness and sideslip angle estimation for integrated vehicle dynamics control. *IFAC-PapersOnLine*, 49(11):297–304, 2016.
- [91] Joanny Stephant, Ali Charara, and Dominique Meizel. Virtual sensor: Application to vehicle sideslip angle and transversal forces. *IEEE Transactions on Industrial Electronics*, 51(2):278–289, 2004.
- [92] Joanny Stéphant, Ali Charara, and Dominique Meizel. Evaluation of a sliding mode observer for vehicle sideslip angle. *Control Engineering Practice*, 15(7):803–812, 2007.
- [93] Jie Zhang, Hai Wang, Jinchuan Zheng, Zhenwei Cao, Zhihong Man, Ming Yu, and Long Chen. Adaptive sliding mode-based lateral stability control of steer-by-wire vehicles with experimental validations. *IEEE Transactions on Vehicular Technology*, 69(9):9589–9600, 2020.
- [94] Hui Zhang, Guoguang Zhang, and Junmin Wang. Sideslip angle estimation of an electric ground vehicle via finite-frequency \mathcal{H}_∞ approach. *IEEE Transactions on Transportation Electrification*, 2(2):200–209, 2016.

-
- [95] Shuo Cheng, Liang Li, Bingjie Yan, Congzhi Liu, Xiangyu Wang, and Jigen Fang. Simultaneous estimation of tire side-slip angle and lateral tire force for vehicle lateral stability control. *Mechanical Systems and Signal Processing*, 132:168–182, 2019.
 - [96] Nenggen Ding, Wen Chen, Yipeng Zhang, Guoyan Xu, and Feng Gao. An extended luenberger observer for estimation of vehicle sideslip angle and road friction. *International Journal of Vehicle Design*, 66(4):385–414, 2014.
 - [97] Olga Galluppi, Matteo Corno, and Sergio M. Savaresi. Mixed-kinematic body sideslip angle estimator for high performance cars. In *IEEE European Control Conference*, pages 941–946, 2018.
 - [98] Elvis Villano, Basilio Lenzo, and Aleksandr Sakhnevych. Cross-combined UKF for vehicle sideslip angle estimation with a modified dugoff tire model: design and experimental results. *Meccanica*, 56(11):2653–2668, 2021.
 - [99] Giseo Park. Vehicle sideslip angle estimation based on interacting multiple model kalman filter using low-cost sensor fusion. *IEEE Transactions on Vehicular Technology*, 71(6):6088–6099, 2022.
 - [100] Wuwei Chen, Dongkui Tan, and Linfeng Zhao. Vehicle sideslip angle and road friction estimation using online gradient descent algorithm. *IEEE Transactions on Vehicular Technology*, 67(12):11475–11485, 2018.
 - [101] Ehsan Hashemi, Mohammad Pirani, Amir Khajepour, Alireza Kasaiezadeh, Shih-Ken Chen, and Bakhtiar Litkouhi. Corner-based estimation of tire forces and vehicle velocities robust to road conditions. *Control Engineering Practice*, 61:28–40, 2017.
 - [102] Milad Jalali, Ehsan Hashemi, Amir Khajepour, Shih-ken Chen, and Bakhtiar Litkouhi. Integrated model predictive control and velocity estimation of electric vehicles. *Mechatronics*, 46:84–100, 2017.
 - [103] Hideaki Sasaki and Takatoshi Nishimaki. A side-slip angle estimation using neural network for a wheeled vehicle. In *SAE Technical Paper*, 2000.
 - [104] Stefano Melzi and Edoardo Sabbioni. On the vehicle sideslip angle estimation through neural networks: Numerical and experimental results. *Mechanical Systems and Signal Processing*, 25(6):2005–2019, 2011.
 - [105] Angelo Bonfitto, Stefano Feraco, Andrea Tonoli, and Nicola Amati. Combined regression and classification artificial neural networks for sideslip angle estimation and road condition identification. *Vehicle System Dynamics*, 58(11):1766–1787, 2020.
 - [106] Jyotishman Ghosh, Andrea Tonoli, and Nicola Amati. A deep learning based virtual sensor for vehicle sideslip angle estimation: experimental results. In *SAE Technical Paper*, 2018.
 - [107] Jizheng Liu, Zhenpo Wang, and Lei Zhang. A time-delay neural network of sideslip angle estimation for in-wheel motor drive electric vehicles. In *IEEE Vehicular Technology Conference*, pages 1–5, 2020.

- [108] Valentina Breschi, Simone Formentin, Gianmarco Rallo, Matteo Corno, and Sergio M. Savaresi. Vehicle sideslip estimation via kernel-based lpy identification: Theory and experiments. *Automatica*, 122:109237, 2020.
- [109] Manuel Acosta, Stratis Kanarachos, and Michael E. Fitzpatrick. Robust virtual sensing for vehicle agile manoeuvring: A tyre-model-less approach. *IEEE Transactions on Vehicular Technology*, 67(3):1894–1908, 2017.
- [110] Tommaso Novi, Renzo Capitani, and Claudio Annicchiarico. An integrated artificial neural network–unscented kalman filter vehicle sideslip angle estimation based on inertial measurement unit measurements. *Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 233(7):1864–1878, 2019.
- [111] Nicola A. Piga, Ugo Pattacini, and Lorenzo Natale. A differentiable extended kalman filter for object tracking under sliding regime. *Frontiers in Robotics and AI*, 8, 2021.
- [112] Woongsun Jeon, Ankush Chakrabarty, Ali Zemouche, and Rajesh Rajamani. Simultaneous state estimation and tire model learning for autonomous vehicle applications. *IEEE/ASME Transactions on Mechatronics*, 26(4):1941–1950, 2021.
- [113] Qian Zhang, Houhua Jing, Zhiyuan Liu, Yongfeng Jiang, and Mingqin Gu. A novel pwa lateral dynamics modeling method and switched ts observer design for vehicle sideslip angle estimation. *IEEE Transactions on Industrial Electronics*, 69(2):1847–1857, 2021.
- [114] Tuomas Haarnoja, Anurag Ajay, Sergey Levine, and Pieter Abbeel. Backprop KF: Learning discriminative deterministic state estimators. *Advances in neural information processing systems*, 29:4383–4391, 2016.
- [115] Guy Revach, Nir Shlezinger, Xiaoyong Ni, Adria Lopez Escoriza, Ruud JG Van Sloun, and Yonina C Eldar. Kalmannet: Neural network aided kalman filtering for partially known dynamics. *IEEE Transactions on Signal Processing*, 70:1532–1547, 2022.
- [116] Nan Xu, Yanjun Huang, Hassan Askari, and Zepeng Tang. Tire slip angle estimation based on the intelligent tire technology. *IEEE Transactions on Vehicular Technology*, 70(3):2239–2249, 2021.
- [117] Kristian Kersting, Christian Plagemann, Patrick Pfaff, and Wolfram Burgard. Most likely heteroscedastic gaussian process regression. In *IEEE International conference on Machine learning*, pages 393–400, 2007.
- [118] Maria Bauza and Alberto Rodriguez. A probabilistic data-driven model for planar pushing. In *IEEE International Conference on Robotics and Automation*, pages 3008–3015, 2017.
- [119] Jason Brownlee. *Better deep learning: train faster, reduce overfitting, and make better predictions*. Machine Learning Mastery, 2018.

-
- [120] Alberto Bertipaglia, Barys Shyrokau, Mohsen Alirezaei, and Riender Happee. A two-stage bayesian optimisation for automatic tuning of an unscented kalman filter for vehicle sideslip angle estimation. In *IEEE Intelligent Vehicles Symposium*, pages 670–677, 2022.
 - [121] Manuel Acosta, Stratis Kanarachos, and Michael E. Fitzpatrick. Optimized tire force estimation using extended kalman filter and fruit fly optimization. In *IEEE Conference of the Industrial Electronics Society*, pages 4074–4079, 2017.
 - [122] Yiqi Gao, Andrew Gray, Ashwin Carvalho, Eric Tseng, and Francesco Borrelli. Robust nonlinear predictive control for semiautonomous ground vehicles. In *American Control Conference*, 2014.
 - [123] Qingjia Cui, Rongjun Ding, Chongfeng Wei, and Bing Zhou. Path-tracking and lateral stabilisation for autonomous vehicles by using the steering angle envelope. *Vehicle System Dynamics*, 59(11):1672–1696, 2021.
 - [124] Edwin K. Liebemann, Klaus Meder, J. Schuh, and Gero Nenninger. Safety and performance enhancement: The bosch electronic stability control (ESP). In *SAE Technical Paper*, 2004.
 - [125] Yiqi Gao, Theresa Lin, Francesco Borrelli, Eric Tseng, and Davor Hrovat. Predictive control of autonomous ground vehicles with obstacle avoidance on slippery roads. In *Dynamic Systems and Control Conference*, volume 1, pages 265–272, 2010.
 - [126] Yanchuan Xu, Huarong Zheng, Weimin Wu, and Jun Wu. Robust hierarchical model predictive control for trajectory tracking with obstacle avoidance. *IFAC-PapersOnLine*, 53(2):15745–15750, 2020.
 - [127] Hongqing Chu, Dele Meng, Siwen Huang, Mengjian Tian, Jia Zhang, Bingzhao Gao, and Hong Chen. Autonomous high-speed overtaking of intelligent chassis using fast iterative model predictive control. *IEEE Transactions on Transportation Electrification*, 10(1):1244–1256, 2024.
 - [128] Hung Duy Nguyen, Dongryul Kim, Young Seop Son, and Kyoungseok Han. Linear time-varying MPC-based autonomous emergency steering control for collision avoidance. *IEEE Transactions on Vehicular Technology*, 72(10):12713–12727, 2023.
 - [129] Matthew Brown and J. Christian Gerdes. Robust stabilization and collision avoidance through minimizing open-loop velocity uncertainty. In *IEEE Intelligent Vehicles Symposium*, pages 259–264, 2020.
 - [130] Andrea Zanelli, Alexander Domahidi, Juan Jerez, and Manfred Morari. Forces nlp: an efficient implementation of interior-point methods for multistage nonlinear nonconvex programs. *International Journal of Control*, pages 1–17, 2017.
 - [131] Gianpiero Mastinu, Fabio Della Rossa, Giorgio Previati, Massimiliano Gobbi, and Marco Fainello. Global stability of road vehicle motion with driver control. *Nonlinear Dynamics*, 111(19):18043–18059, 2023.

- [132] Giampiero RM Mastinu, Giorgio Previati, Fabio Della Rossa, Massimiliano Gobbi, and Marco Fainello. How drivers lose control of the car. *SAE International Journal of Vehicle Dynamics, Stability, and NVH*, 8(10-08-01-0007), 2024.
- [133] Fabio Della Rossa, Giampiero Mastinu, and Carlo Piccardi. Bifurcation analysis of an automobile model negotiating a curve. *Vehicle system dynamics*, 50(10):1539–1562, 2012.
- [134] Alberto Bertipaglia, Davide Tavernini, Umberto Montanaro, Mohsen Alirezaei, Riemer Happee, Aldo Sorniotti, and Barys Shyrokau. Model predictive contouring control for vehicle obstacle avoidance at the limit of handling using torque vectoring. In *IEEE International Conference on Advanced Intelligent Mechatronics*, pages 1468–1475, 2024.
- [135] John K. Subosits and J. Christian Gerdes. Impacts of model fidelity on trajectory optimization for autonomous vehicles in extreme maneuvers. *IEEE Transactions on Intelligent Vehicles*, 6(3):546–558, 2021.
- [136] Reza Hajiloo, Mehdi Abroshan, Amir Khajepour, Alireza Kasaiezadeh, and Shih-Ken Chen. Integrated steering and differential braking for emergency collision avoidance in autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 22(5):3167–3178, 2021.
- [137] Siyu Teng, Xuemin Hu, Peng Deng, Bai Li, Yuchen Li, Yunfeng Ai, Dongsheng Yang, Lingxi Li, Zhe Xuanyuan, Fenghua Zhu, and Long Chen. Motion planning for autonomous driving: The state of the art and future perspectives. *IEEE Transactions on Intelligent Vehicles*, 8(6):3692–3711, 2023.
- [138] Ehsan Hashemi and Amir Khajepour. Integrated path-tracking and combined-slip force controls of autonomous ground vehicles with safe constraints adaptation. *IEEE Transactions on Intelligent Vehicles*, 9(3):4265–4274, 2024.
- [139] Wolfgang Degel, Stefan Lupberger, Dirk Odenthal, and Naim Bajcinca. Scalable slip control with torque vectoring including input-to-state stability analysis. *IEEE Transactions on Control Systems Technology*, 31(3):1250–1265, 2023.
- [140] Leonardo De Novellis, Aldo Sorniotti, Patrick Gruber, and Andrew Pennycott. Comparison of feedback control techniques for torque-vectoring control of fully electric vehicles. *IEEE Transactions on Vehicular Technology*, 63(8):3612–3623, 2014.
- [141] Pak Kin Wong and Di Ao. A novel event-triggered torque vectoring control for improving lateral stability and communication resource consumption of electric vehicles. *IEEE Transactions on Intelligent Vehicles*, 9(1):2046–2060, 2024.
- [142] Yue Ren, Ling Zheng, and Amir Khajepour. Integrated model predictive and torque vectoring control for path tracking of 4-wheel-driven autonomous vehicles. *IET Intelligent Transport Systems*, 13(1):98–107, 2019.

-
- [143] Christoforos Chatzikomis, Aldo Sorniotti, Patrick Gruber, Mattia Zanchetta, Dan Willans, and Bryn Balcombe. Comparison of path tracking and torque-vectoring controllers for autonomous electric vehicles. *IEEE Transactions on Intelligent Vehicles*, 3(4):559–570, 2018.
 - [144] Yunchang Yu, Wenfei Ji, Runfeng Li, Ao Lu, and Guangyu Tian. Vehicle motion control beyond and within the stability limits for 4wd electric vehicles. *IEEE Transactions on Intelligent Vehicles*, 9(1):2348–2363, 2024.
 - [145] Victor Mazzilli, Stefano De Pinto, Leonardo Pascali, Michele Contrino, Francesco Bottiglione, Giacomo Mantriota, Patrick Gruber, and Aldo Sorniotti. Integrated chassis control: Classification, analysis and future trends. *Annual Reviews in Control*, 51:172–205, 2021.
 - [146] Dhanaraja Kasinathan, Alireza Kasaiezadeh, Andy Wong, Amir Khajepour, Shih-Ken Chen, and Bakhtiar Litkouhi. An optimal torque vectoring control for vehicle applications via real-time constraints. *IEEE Transactions on Vehicular Technology*, 65(6):4368–4378, 2015.
 - [147] Yixiao Liang, Yinong Li, Amir Khajepour, and Ling Zheng. Holistic adaptive multi-model predictive control for the path following of 4wd autonomous vehicles. *IEEE Transactions on Vehicular Technology*, 70(1):69–81, 2021.
 - [148] Trey P. Weber and J. Christian Gerdes. Modeling and control for dynamic drifting trajectories. *IEEE Transactions on Intelligent Vehicles*, 9(2):3731–3741, 2024.
 - [149] Vincent A. Laurence and J. Christian Gerdes. Long-horizon vehicle motion planning and control through serially cascaded model complexity. *IEEE Transactions on Control Systems Technology*, 30(1):166–179, 2022.
 - [150] Hans Pacejka. *Tire and vehicle dynamics*. Elsevier, 2005.
 - [151] Chuanyang Yu, Yanggu Zheng, Barys Shyrokau, and Valentin Ivanov. MPC-based path following design for automated vehicles with rear wheel steering. In *IEEE International Conference on Mechatronics*, pages 1–6, 2021.
 - [152] Víctor Vidal, Pietro Stano, Gaetano Tavolo, Miguel Dhaens, Davide Tavernini, Patrick Gruber, and Aldo Sorniotti. On pre-emptive in-wheel motor control for reducing the longitudinal acceleration oscillations caused by road irregularities. *IEEE Transactions on Vehicular Technology*, 71(9):9322–9337, 2022.
 - [153] Alberto Parra, Davide Tavernini, Patrick Gruber, Aldo Sorniotti, Asier Zubizarreta, and Joshué Pérez. On nonlinear model predictive control for energy-efficient torque-vectoring. *IEEE Transactions on Vehicular Technology*, 70(1):173–188, 2021.
 - [154] Alessandro Scamarcio, Carmine Caponio, Mario Mihalkov, Petar Georgiev, Javad Ahmadi, Kai Man So, Davide Tavernini, and Aldo Sorniotti. Predictive anti-jerk and traction control for v2x connected electric vehicles with central motor and open differential. *IEEE Transactions on Vehicular Technology*, 72(6):7221–7239, 2023.

- [155] Dzmityr Savitski, Valentin Ivanov, Barys Shyrokau, Jasper De Smet, and Johan Theunissen. Experimental study on continuous ABS operation in pure regenerative mode for full electric vehicle. *SAE International Journal of Passenger Cars-Mechanical Systems*, 8(2015-01-9109):364–369, 2015.
- [156] Carlo Lugaro, Mohsen Alirezaei, Ioannis Konstantinou, and Abhijeet Behera. A study on the effect of tire temperature and rolling speed on the vehicle handling response. In *SAE Technical Paper*, 2020.
- [157] Yinpeng Dong, Caixin Kang, Jinlai Zhang, Zijian Zhu, Yikai Wang, Xiao Yang, Hang Su, Xingxing Wei, and Jun Zhu. Benchmarking robustness of 3d object detection to common corruptions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2023.
- [158] Cheng Zhang, Hai Wang, Long Chen, Yicheng Li, and Yingfeng Cai. Mixedfusion: An efficient multimodal data fusion framework for 3-d object detection and tracking. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2023.
- [159] Antonio Artuñedo, Jorge Villagra, Jorge Godoy, and Maria Dolores del Castillo. Motion planning approach considering localization uncertainty. *IEEE Transactions on Vehicular Technology*, 69(6):5983–5994, 2020.
- [160] Andreas Pfrunder, Paulo V. K. Borges, Adrian R. Romero, Gavin Catt, and Alberto Elfes. Real-time autonomous ground vehicle navigation in heterogeneous environments using a 3d lidar. In *IEEE International Conference on Intelligent Robots and Systems*, 2017.
- [161] Andras Palffy, Ewoud Pool, Srimannarayana Baratam, Julian F. P. Kooij, and Darius M. Gavrilă. Multi-class road user detection with 3+1d radar in the view-of-delft dataset. *IEEE Robotics and Automation Letters*, 7(2):4961–4968, 2022.
- [162] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [163] Yongqiang Lu, Hongjie Ma, Edward Smart, and Hui Yu. Real-time performance-focused localization techniques for autonomous vehicle: A review. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):6082–6100, 2022.
- [164] Francesca Pianosi and Thorsten Wagener. Distribution-based sensitivity analysis from a generic input-output sample. *Environmental Modelling & Software*, 108:197–207, 2018.
- [165] Arnald Puy, Samuele Lo Piano, and Andrea Saltelli. A sensitivity analysis of the pawn sensitivity index. *Environmental modelling & Software*, 127:104679, 2020.
- [166] Mohsen Alirezaei, Sven Jansen, Antoine Schmeitz, and Anil Kunnappillil Madhusudhanan. Collision avoidance system using state dependent riccati equation technique:

-
- An experimental robustness evaluation. In *International Symposium on Advanced Vehicle Control*, pages 125–132, 2016.
- [167] Alberto Bertipaglia, Mohsen Alirezaei, Riender Happee, and Barys Shyrokau. A learning-based model predictive contouring control for vehicle evasive manoeuvres. In *International Symposium on Advanced Vehicle Control*, pages 632–638, 2024.
 - [168] Amar Shah, Andrew Wilson, and Zoubin Ghahramani. Student-t processes as alternatives to gaussian processes. In *Artificial Intelligence and Statistics*, pages 877–885, 2014.
 - [169] Jarno Vanhatalo, Pasi Jylänki, and Aki Vehtari. Gaussian process regression with student-t likelihood. *Advances in Neural Information Processing Systems*, 22, 2009.
 - [170] Carl Edward Rasmussen and Hannes Nickisch. Gaussian processes for machine learning toolbox. *Journal of Machine Learning Research*, 11:3011–3015, 2010.
 - [171] Stan Meijer, Alberto Bertipaglia, and Barys Shyrokau. A nonlinear model predictive control for automated drifting with a standard passenger vehicle. In *IEEE International Conference on Advanced Intelligent Mechatronics*, pages 284–289, 2024.
 - [172] Mincheol Park and Yeonsik Kang. Experimental verification of a drift controller for autonomous vehicle tracking: A circular trajectory using lqr method. *International Journal of Control, Automation and Systems*, 19:404–416, 2021.
 - [173] Szilard Czibere, Adam Domina, Adam Bardos, and Zsolt Szalay. Model predictive controller design for vehicle motion control at handling limits in multiple equilibria on varying road surfaces. *Energies*, 14(20), 2021.
 - [174] Vivian Zhang, Sarah M. Thornton, and J Christian Gerdes. Tire modeling to enable model predictive control of automated vehicles from standstill to the limits of handling. In *International Symposium on Advanced Vehicle Control*, 2018.
 - [175] Craig Earl Beal and J. Christian Gerdes. Model predictive control for vehicle stabilization at the limits of handling. *IEEE Transactions on Control Systems Technology*, 21(4):1258–1269, 2013.
 - [176] Boris Houska, Hans J. Ferreau, and Moritz Diehl. ACADO Toolkit – An Open Source Framework for Automatic Control and Dynamic Optimization. *Optimal Control Applications and Methods*, 32(3):298–312, 2011.
 - [177] Zhaozhong Chen, Christoffer Heckman, Simon Julier, and Nisar Ahmed. Weak in the NEES? auto-tuning kalman filters with bayesian optimization. In *IEEE International Conference on Information Fusion*, pages 1072–1079, 2018.
 - [178] Pieter Abbeel, Adam Coates, Michael Montemerlo, Andrew Y Ng, and Sebastian Thrun. Discriminative training of kalman filters. In *Robotics: Science and Systems*, volume 2, page 1, 2005.

- [179] Hakki Mert Torun, Madhavan Swaminathan, Anto Kavungal Davis, and Mohamed Lamine Faycal Bellaredj. A global bayesian optimization algorithm and its application to integrated system design. *IEEE Transactions on Very Large Scale Integration Systems*, 26(4):792–802, 2018.
- [180] Stijn Kerst, Barys Shyrokau, and Edward Holweg. Anti-lock braking control based on bearing load sensing. *EuroBrake*, pages 4–6, 2015.
- [181] Kyoungseok Han, Eunjae Lee, Mooryong Choi, and Seibum B. Choi. Adaptive scheme for the real-time estimation of tire-road friction coefficient and vehicle velocity. *IEEE/ASME Transactions on Mechatronics*, 22(4):1508–1518, 2017.
- [182] Gurkan Erdogan, Lee Alexander, and Rajesh Rajamani. Estimation of tire-road friction coefficient using a novel wireless piezoelectric tire sensor. *IEEE Sensors Journal*, 11(2):267–279, 2011.
- [183] Yahui Liu, Tong Li, Yiyong Yang, Xuwu Ji, and Jian Wu. Estimation of tire-road friction coefficient based on combined APF-IEKF and iteration algorithm. *Mechanical Systems and Signal Processing*, 88:25–35, 2017.
- [184] Victor Fors, Björn Olofsson, and Lars Nielsen. Autonomous wary collision avoidance. *IEEE Transactions on Intelligent Vehicles*, 6(2):353–365, 2021.
- [185] International Organization for Standardization. Passenger cars — test track for a severe lane-change manoeuvre: Part 2: Obstacle avoidance. <https://www.iso.org/standard/57253.html>, 2011.
- [186] Bharath Bhaskar Mahadikar, Nishant Rajesh, Kevin Tom Kurian, Erjen Lefeber, Jeroen Ploeg, Nathan van de Wouw, and Mohsen Alirezaei. Formulating a dissimilarity metric for comparison of driving scenarios for automated driving systems. In *IEEE Intelligent Vehicles Symposium*, pages 1091–1098, 2024.
- [187] Tajinder Singh, Edwin van Hassel, Akshay Sheorey, and Mohsen Alirezaei. A systematic approach for creation of SOTIF’s unknown unsafe scenarios: An optimization based method. In *SAE Technical Paper*, 2024.
- [188] Zlatan Ajanović, Enrico Regolin, Barys Shyrokau, Hana Čatić, Martin Horn, and Antonella Ferrara. Search-based task and motion planning for hybrid systems: Agile autonomous vehicles. *Engineering Applications of Artificial Intelligence*, 121:105893, 2023.
- [189] Wangjia Weng, Cheng Hu, Zhouheng Li, Hongye Su, and Lei Xie. An aggressive cornering framework for autonomous vehicles combining trajectory planning and drift control. In *IEEE Intelligent Vehicles Symposium*, pages 2749–2755, 2024.
- [190] Daofei Li, Jiajie Zhang, and Siyuan Lin. Planning and control of drifting-based collision avoidance strategy under emergency driving conditions. *Control Engineering Practice*, 139:105625, 2023.

-
- [191] LWP Luuk Voesten. Optimal trajectory generation for a collision avoidance maneuver within and beyond stable limits of handling (drifting). Master's thesis, Eindhoven University of Technology, 2021.
 - [192] Eduardo Mosqueira-Rey, David Alonso-Ríos, and Andrés Baamonde-Lozano. Integrating iterative machine teaching and active learning into the machine learning loop. *Procedia Computer Science*, 192:553–562, 2021.
 - [193] Patrice Simard, Saleema Amershi, David M. Chickering, Alicia Edelman Pelton, Soroush Ghorashi, Christopher Meek, Gonzalo Ramos, Jina Suh, Johan Verwey, Mo Wang, et al. Machine teaching: A new paradigm for building machine learning systems. *arXiv preprint arXiv:1707.06742*, 2017.
 - [194] Gonzalo Ramos, Christopher Meek, Patrice Simard, Jina Suh, and Soroush Ghorashi. Interactive machine teaching: a human-centered approach to building machine-learned models. *Human–Computer Interaction*, 35(5-6):413–451, 2020.
 - [195] Matthew Riemer, Michele Franceschini, Djallel Bouneffouf, and Tim Klinger. Generative knowledge distillation for general purpose function compression. In *Workshop on Teaching Machines, Robots, and Humans*, volume 5, page 30, 2017.
 - [196] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *IEEE International Conference on Machine Learning*, pages 41–48, 2009.
 - [197] Luigi Di Lillo, Tilia Gode, Xilin Zhou, Margherita Atzei, Ruoshu Chen, and Trent Victor. Comparative safety performance of autonomous-and human drivers: A real-world case study of the waymo driver. *Heliyon*, 10(14), 2024.
 - [198] Sam D. Doecke, Matthew R.J. Baldock, Craig N. Kloeden, and Jeffrey K. Dutschke. Impact speed and the risk of serious injury in vehicle crashes. *Accident Analysis & Prevention*, 144:105629, 2020.
 - [199] United Nations Economic Commission for Europe. UN Regulation No. 157: Automated Lane Keeping Systems. <https://unece.org/sites/default/files/2023-12/R157e.pdf>, January 2021.
 - [200] Honda Global. For a Joyous Future Where Everyone Can Go Anywhere Freely With Total Peace of Mind . https://global.honda/en/tech/Automated_drive_safety_and_driver_assistive_technologies/, 2022.
 - [201] Brad Templeton. Ntsb releases report on 2018 fatal silicon valley tesla autopilot crash. <https://www.forbes.com/sites/bradtempleton/2020/02/13/ntsb-releases-report-on-2018-fatal-silicon-valley-tesla-autopilot-crash/>, February 2020.
 - [202] Rory Cellan-Jones. Uber's self-driving operator charged over fatal crash. <https://www.bbc.com/news/technology-54175359>, September 2020.

- [203] Zhaozhong Chen, Christoffer Heckman, Simon Julier, and Nisar Ahmed. Time dependence in kalman filter tuning. In *IEEE International Conference on Information Fusion*, pages 1–8, 2021.
- [204] Manuel Acosta and Stratis Kanarachos. Optimized vehicle dynamics virtual sensing using metaheuristic optimization and unscented kalman filter. In *Evolutionary and Deterministic Methods for Design Optimization and Control With Applications to Industrial and Societal Problems*, pages 275–290, 2019.
- [205] Yujie Tang, Liang Hu, Qingrui Zhang, and Wei Pan. Reinforcement learning compensated extended kalman filter for attitude estimation. In *IEEE International Conference on Intelligent Robots and Systems*, pages 6854–6859, 2021.
- [206] Hannes Heidfeld, Martin Schünemann, and Roland Kasper. UKF-based state and tire slip estimation for a 4WD electric vehicle. *Vehicle System Dynamics*, 58(10):1479–1496, 2020.
- [207] Thomas Powell. Automated tuning of an extended kalman filter using the downhill simplex algorithm. *Journal of Guidance, Control, and Dynamics*, 25(5):901–908, 2002.
- [208] Yaakov Oshman and Ilan Shaviv. Optimal tuning of a kalman filter using genetic algorithms. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2000.
- [209] Adrià López Escoriza, Guy Revach, Nir Shlezinger, and Ruud J. G. van Sloun. Data-driven kalman-based velocity estimation for autonomous racing. In *IEEE International Conference on Autonomous Systems*, pages 1–5, 2021.
- [210] Zhaozhong Chen, Nisar Ahmed, Simon Julier, and Christoffer Heckman. Kalman filter tuning with bayesian optimization. *arXiv preprint arXiv:1912.08601*, 2019.
- [211] Hakki Torun, Jose Hejase, Junyan Tang, Wiren Beckert, and Madhavan Swaminathan. Bayesian active learning for uncertainty quantification of high speed channel signaling. In *IEEE Conference on Electrical Performance of Electronic Packaging and Systems*, pages 311–313, 2018.
- [212] Conor Clare, Glenn Hawe, and Sally McClean. Expected regret minimization for bayesian optimization with student’s-t processes. In *International Conference on Artificial Intelligence and Pattern Recognition*, pages 8–12, 2020.

ACKNOWLEDGMENTS

One of my favourite authors, Luigi Pirandello, wrote in *The Pleasure of Honesty*, "Anyone can be heroic from time to time, but a gentleman is something you have to be all the time." In my opinion, Pirandello expressed that heroism is often defined by a single commendable act while being a gentleman requires constant effort and attention. I believe this distinction can also apply to the journey of a PhD. Achieving a PhD is undeniably challenging, but it is attainable, and it requires only three journal papers. However, maintaining the humility, responsibility, and resilience of a dedicated student is far more demanding. It requires recognising your mistakes, acknowledging your limitations, enduring high levels of stress without ever placing the burden on other, and understanding that success in a PhD is not a solitary achievement but is the result of a collective effort supported by many.

With this in mind, I want to begin by expressing my deepest gratitude to all those who have supported me throughout this journey. First and foremost, I extend my sincere thanks to my supervisory team: **Dr. Barys Shyrokau**, **Dr. Mohsen Alirezaei**, and **Prof. Riender Happee**. I am not only grateful for your invaluable feedback and suggestions but also for the genuine care you show toward your students.

Dr. Barys Shyrokau, it is impossible to fully express all the reasons I am grateful to you. You have guided me through every stage of this journey, helping me grow into an (almost) independent researcher. I hope one day to be the kind of mentor you have been to me. I am especially thankful for the opportunity to work on this project and for your patience in teaching me. I still remember the various mistakes I made early on in my PhD, but you always provided me with the tools and the advice to ensure I would not repeat them. Honestly, I do not think there is something more valuable than that.

Dr. Mohsen Alirezaei, despite our meetings being mostly online, you consistently grounded my research and helped me focus on the most impactful aspects. Your guidance has been crucial in shaping my work, ensuring I addressed real-world problems and not just theoretical contributions. I am also deeply grateful for your constant availability, especially in reviewing and providing feedback on this thesis, even at the last moment.

Prof. Riender Happee, I will always remember what you said during my first interview for this position. You told me that a PhD needs to learn. Despite sometimes being given for granted, that statement showed me the importance you give to teaching your students. I understood the importance you give not only to what we can achieve now but to what we will be capable of in the future, provided we maintain a constant desire to learn. For this and many other reasons, I am deeply thankful for your supervision and support throughout this PhD.

Working on a PhD is made possible not only by the supervisory team but also by those who fund the project. I would like to extend my gratitude to the Netherlands Organisation for Scientific Research (**NWO**) for supporting this research through the **EVOLVE** project.

Especially, I'm deeply grateful to all the users of the project for their constant effort in supporting my research, with data, models, infrastructure and knowledge transfer.

SKF, particularly through **Marco Cambiano** and **Henk Mol**, has played a crucial role in making my research possible. I am deeply grateful for your willingness to share the extensive experimental dataset that formed the basis of many of my publications. Additionally, your insights into the intelligent load-sensing bearings and the models you have shared with me have been invaluable. I look forward to the potential fruitful collaboration you have suggested.

TASS International (Siemens), particularly through **Dr. Mohsen Alirezaei**, has been instrumental in enhancing the supervision of my PhD and providing me with the opportunity to utilise their software.

TNO, particularly through **Dr. Antoine Schmeitz**, for your constant feedback and guidance on all of my PhD publications. Without the tyre models you shared, I would not have been able to fully validate and enhance the accuracy of the vehicle plant model. I am also deeply grateful for your interest in my research, which made every user committee meeting an opportunity to refine and improve my work.

Toyota Motor Europe, particularly through **Giovanni Berardo** and **Felipe Santafe**, your willingness and availability have contributed enormously to my PhD. I am especially grateful for your support in hosting and supervising master students in collaboration with **TU Delft**.

Volvo Car Corporation, particularly through **Dr. Yang Derong**, your collaboration and support in conducting the experimental validation of the approaches I proposed during my PhD have been and will continue to be, invaluable. I look forward to carrying out the experimental validation in Sweden.

The complexity of the EVOLVE project would not have been resolved without the continuous efforts of another team at **TU Delft**, with whom I have closely collaborated over the last four years. For this, I extend my sincere gratitude to all those who worked alongside me and provided constant inspiration. In particular, I would like to thank **Leila Gharavi** for her consistent collaboration throughout the PhD, especially in the supervision of master students. Additionally, I would like to thank **Prof. Simone Baldi** and **Prof. Bart De Schutter** for the insightful meetings over the years and for their valuable feedback on my research.

As I previously mentioned, it is impossible to fully express all the reasons why I am grateful to **Dr. Barys Shyrokau**, but for sure, I can mention other examples. One is for giving me the opportunity to participate in the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie actions within the OWHEEL project. I had the chance to experience two enriching secondments, the first at the **University of Surrey** and the second one at the **Vilnius Gediminas Technical University**.

At the **University of Surrey**, I had the opportunity to collaborate with automotive experts who enhanced my knowledge of vehicle control and supported my research with valuable feedback. In particular, I want to express my gratitude to **Dr. Davide Tavernini** for his constant feedback and guidance in the development of Chapter 4 of this thesis. His support was essential not only for Chapter 4 but also for ensuring that all the administrative processes during my secondment went smoothly. Furthermore, I extend my deepest thanks to **Dr.**

Umberto Montanaro and **Prof. Aldo Sorniotti** for closely following my secondment progress and always being available to assist with the development of our joint publication. My experience at Surrey was not only about work but also about making friends, building connections, and integrating into Guildford's life. For that, I am especially grateful to **Dr. Aymen Alshaw** and **Dr. Megan Pegram**, who immediately made me feel at home.

Although I had previously visited the United Kingdom, this was my first opportunity to visit Lithuania, where I had the pleasure of collaborating with **Vilnius Gediminas Technical University**. I am deeply thankful to **Dr. Viktor Skrickij** not only for hosting me and supervising during my visit but especially for the insightful conversations we had about work, automotive research, and Lithuanian culture and history. I would also like to express my gratitude to **Dr. Eldar Šabanovič** for our stimulating discussions on machine learning, and I hope we can collaborate on some of those ideas in the future. Furthermore, I am grateful to **Paulius Kojis** for his friendship and for showing me around Vilnius. Best of luck with your PhD, and I hope we can meet again in Vilnius. Lastly, I would like to extend my gratitude to all the wonderful people I met in Lithuania for your warm welcome. It was an incredible experience to celebrate the Act of Independence of Lithuania with you, ačiū!

So far, I have thanked many companies and universities, but not yet the institution that has hosted me all these years. For this, I would like to express my sincere gratitude to **Delft University of Technology** for creating an environment that emphasises a willingness to learn and think critically. In particular, I want to thank the members of the **Intelligent Vehicles** permanent staff who I have not yet mentioned. A special thanks to **Dr. Georgios Papaioannou** for the enjoyable conversations and for sharing the (dis)adventures of Ottawa with me. (In a parallel universe, we are still waiting for our luggage at the airport). **Dr. Julian F. P. Kooij**, thank you for your enthusiasm and passion for research. I would also like to thank **Dr. Holger Caesar** and **Prof. Dariu M. Gavrila** for their support. I am also extremely grateful to **Ronald Ensing** and **Dr. Mario Garzon Oviedo** for your essential help in setting up the Toyota Prius and for your patience and assistance during my experiments.

Last but not least, I would like to thank all the secretaries of the Cognitive Robotics department for their constant support with bureaucracy, reimbursements, and business trips. A special thanks goes to **Hanneke Hustinx** for always being available to assist me. I am especially grateful for your help in ensuring I received my Postdoc Researcher contract on time.

I began this section by noting that a PhD is not a solitary achievement but the result of a collective effort supported by many. These 'many' are not only supervisors, colleagues, and industrial partners but, most importantly, **friends**.

First and foremost, I want to thank the members of the Societas: **Hidde Boekema**, **Vishrut Jain**, and **Dr. Wilbert Tabone**, with whom I've spent most of my free time over the past few years. Your friendship has been, and continues to be, essential in helping me achieve this PhD and in my growth not only as an engineer but as a person. As Francesco De Sanctis said, before we are engineers, we are human beings, and our conversations constantly reminded me to keep my mind and interests open. It is likely this, along with our many differences, that brought us together as friends. Our friendship, however, wasn't just about philosophical discussions. Its beauty lay in our ability to seamlessly shift from high-minded topics to the lightest subjects. We've shared many adventures together, from the 'dry' (sorry,

Wilbert) lands of Groningen to the warmth of Malta, from the 4 % Irish beers of Dublin to the 8 % Belgian beers in Ghent. While I will cherish all our memories (even those lost on the 'battlefield'), what excites me most are the ones we will create in the future. In thea veritas!

I would also like to express my deepest gratitude to **Dr. Tuğrul Irmak** and **Monika Šalandová** for their wonderful friendship over the past four years. Your story and the challenges you've overcome could easily be the plot of a film, showing your perseverance and strength in facing difficulties together. Even more remarkable is how you used those challenges to grow closer and embark on new, exciting careers. I will never forget your support during my fight for the green pass, the dinners we shared, and your beautiful wedding in Telč. P.S. Yes, Tuğrul, I will also remember the time you almost broke my heart in the mountains, but I have to admit, it makes for a great story now.

Remarkable was also the wonderful atmosphere in the office, where I had the pleasure of sitting for four years. However, such a great environment isn't the result of mere coincidence; it's created by the presence of good people with whom it was a joy to spend time. Among this amazing group were **Hide Boekema**, **Dr. Wilbert Tabone**, **Dr. András Pálffy**, **Dr. Xiaolin He**, **Dr. Zimin Xia**, **Felix Fent** and the new entry **Dr. Guopeng Li**. A special thanks goes to **Dr. András Pálffy** for all your advice and support throughout my PhD, and of course, for the fun times we shared at the pubs. I would also like to thank **Dr. Xiaolin He** for his generosity and for always being willing to share new Chinese delicacies with all of us in the office.

The wonderful atmosphere extended beyond the office and was present throughout the entire Intelligent Vehicles group. For this, I want to thank all the amazing colleagues I haven't mentioned yet: **Dr. Oscar de Groot**, for your help with the motion planner; **Dr. Yanggu Zheng**, for the most unexpected limoncello I've ever had; **Mubariz Zaffar**, for your contagious humour; **Varun Kotian**, for your RGB vision; **Shiming Wang**, for your kind heart; **Chrysovalanto Messiou**, for your cheerfulness; **Emma Schippers**, for your support during my final thesis sprint; **Marcell Kégl**, for the fun nights in Rotterdam; **Dr. Javier Pérez Fernández**, for being an honorary member of the Societas; **Dr. Soyeon Kim**, for all the fun nights and for explaining me better the film *Past Lives*; **Ted de Vries Lentsch**, for your strength; **Dr. Yancong Lin**, for always beating me to the office; and **Dr. Farzam Tajdari**, **Jetze Schuurmans**, **Dr. Charlotte Croucher**, **Dr. Jork Stapel** and **Dr. Marko Cvetković**.

A very special thanks also goes to **Giulia Beck** for her support and kindness in these months, and for helping me design the cover of my thesis. Your artistic talent is a source of inspiration.

I am deeply grateful to my parents and my brother for their constant support and for all they have done for me, not only during my PhD but throughout my entire journey until this moment.

*Alberto Bertipaglia
Delft, October 2024*

CURRICULUM VITÆ

Alberto BERTIPAGLIA

JUNE 02, 1995 | Born in Novi Ligure, Italy.

EDUCATION

- | | |
|-----------|---|
| 2020-2024 | PhD candidate at COGNITIVE ROBOTICS, Delft University of Technology
<i>Visitor:</i> Centre for Automotive Engineering, University of Surrey.
Transport Engineering, Vilnius Gediminas Technical University.
<i>Thesis:</i> Control of Evasive Manoeuvres for Automated Driving: Solving the Edge Cases. |
| 2017-2020 | MSc. in AUTOMOTIVE ENGINEERING, Politecnico di Torino , <i>Graduated cum Laude</i> .
<i>Erasmus:</i> Oxford Brookes University. |
| 2014-2017 | BSc. in AUTOMOTIVE ENGINEERING, Politecnico di Torino |

PROFESSIONAL EXPERIENCE

- | | |
|-----------|---|
| 2024-2025 | Postdoc at COGNITIVE ROBOTICS, DELFT UNIVERSITY OF TECHNOLOGY
Develop motion planning and decision making methods that adapt to slippery road conditions, to provide feasible and low-risk trajectories. |
| 2019-2020 | Algorithm Developer at AVL LIST GMBH, Graz, Austria
Algorithm development to assess the life of vehicle suspension components in a virtual multi-axial suspension test rig. |
| 2018-2019 | Vehicle Dynamics member at OXFORD BROOKES RACING TEAM, Oxford, United Kingdom
Development of the in house-built quasi-static lap time simulator. |
| 2017-2018 | Head of Vehicle Dynamics and Race Strategy at H2POLITO, Turin, Italy
Development of the control of a parallel hybrid vehicle model, and race strategy optimisation to minimise the fuel consumption during the race. |
| 2016-2017 | Vehicle Dynamics and Race Strategy member at H2POLITO, Turin, Italy
Vehicle handling study to assess the feasibility of the optimum race strategy. |

LIST OF PUBLICATIONS

PUBLICATIONS INCLUDED IN THIS THESIS

1. **Alberto Bertipaglia**, Mohsen Alirezaei, Riender Happee and Barys Shyrokau, "An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation," IEEE Transactions on Vehicular Technology, vol. 73, no. 9, pp. 12731-12746, 2024. (Chapter 2)
2. **Alberto Bertipaglia**, Mohsen Alirezaei, Riender Happee, and Barys Shyrokau, "Experimental Validation of Model Predictive Contouring Control for Vehicle Obstacle Avoidance at the Limit of Handling," (Under preparation). (Chapter 3)
3. **Alberto Bertipaglia**, Davide Tavernini, Umberto Montanaro, Mohsen Alirezaei, Riender Happee, Aldo Sorniotti, and Barys Shyrokau, "On the Benefits of Torque Vectoring for Automated Collision Avoidance at the Limits of Handling," IEEE Transactions on Vehicular Technology, 2025 (Under review). (Chapter 4)
4. **Alberto Bertipaglia**, Mohsen Alirezaei, Riender Happee, and Barys Shyrokau, "A Learning-Based Model Predictive Contouring Control for Vehicle Evasive Manoeuvres," in International Symposium on Advanced Vehicle Control, Milan, Italy, 2024. (Chapter 5)
5. Stan Meijer, **Alberto Bertipaglia** and Barys Shyrokau, "A Nonlinear Model Predictive Control for Automated Drifting with a Standard Passenger Vehicle," in IEEE International Conference on Advanced Intelligent Mechatronics, Boston, USA, 2024. (Best paper finalist) (Chapter 6)
6. **Alberto Bertipaglia**, Barys Shyrokau, Mohsen Alirezaei, and Riender Happee, "A two-stage bayesian optimisation for automatic tuning of an unscented kalman filter for vehicle sideslip angle estimation," in IEEE Intelligent Vehicles Symposium, pp. 670-677, 2022. (Appendix A)

PUBLICATIONS NOT INCLUDED IN THIS THESIS

1. **Alberto Bertipaglia**, and Barys Shyrokau, "A hybrid neural network Kalman filter for vehicle state estimation", Benelux Workshop on Systems and Control, Rotterdam, The Netherlands, 2021, (Extended abstract).
2. **Alberto Bertipaglia**, Dennis de Mol, Mohsen Alirezaei, Riender Happee, and Barys Shyrokau, "Model-based vs data-driven estimation of vehicle sideslip angle and benefits of tyre force measurements," in International Symposium on Advanced Vehicle Control, Kanagawa, Japan, 2022. (Foundations for Chapter 2)

3. **Alberto Bertipaglia**, Mohsen Alirezai, Riender Happee, and Barys Shyrokau, "Model predictive contouring control for vehicle obstacle avoidance at the limit of handling," in Symposium on the Dynamics of Vehicle on Roads and on Tracks, Ottawa, Canada, 2023. (Foundations for Chapter 3)
4. Daan Lenssen, **Alberto Bertipaglia**, Felipe Santafo, and Barys Shyrokau, "Combined path following and vehicle stability control using model predictive control," SAE Technical Paper, 2023.
5. **Alberto Bertipaglia**, Davide Tavernini, Umberto Montanaro, Mohsen Alirezai, Riender Happee, Aldo Sorniotti, Barys Shyrokau, "Model Predictive Contouring Control for Vehicle Obstacle Avoidance at the Limit of Handling Using Torque Vectoring," in IEEE International Conference on Advanced Intelligent Mechatronics, Boston, USA, 2024. (Foundations for Chapter 4)
6. Petar Velchev, **Alberto Bertipaglia**, Felipe Santafo, and Barys Shyrokau, "Vehicle Motion Control of Automated Driving Using Model Predictive Control," EuroBrake, Mainz, Germany, 2024, (Oral presentation).

PREVIOUS PUBLICATIONS

1. Massimiliana Carello, **Alberto Bertipaglia**, Alessandro Messana, Andrea Giancarlo Airale, and Lorenzo Sisca, "Modeling and optimization of the consumption of a three-wheeled vehicle," SAE Technical Paper, 2019.

Propositions accompanying the dissertation
CONTROL OF EVASIVE MANOEUVRES FOR AUTOMATED DRIVING:
SOLVING THE EDGE CASES

by

Alberto BERTIPAGLIA

1. Algorithms that are based entirely on machine learning do not match the safety in emergency scenarios of those that integrate physics-based knowledge.
This proposition pertains to Chapter 2.
2. Automated vehicles outperform average human drivers and have the potential to surpass even professional drivers in evasive manoeuvres.
This proposition pertains to Chapter 3 and Chapter 4.
3. Incorporating tyre force information into learning-based Model Predictive Control improves vehicle's stability and performance during evasive manoeuvres.
This proposition pertains to Chapter 5.
4. Automated drifting improves collision avoidance systems by maximising safety.
This proposition pertains to Chapter 6.
5. Technological advancements reduce, but not eliminate, uncertainties; so, their quantification must be a fundamental aspect of designing collision avoidance algorithms.
6. There is a significant difference between automotive and robotics research in collision avoidance for automated vehicles.
7. Earning a PhD can open doors to a fulfilling career in academia or research, but the competition and delayed rewards often lead to disillusionment and career shifts.
8. Watching films with complex characters and moral dilemmas enhances psychological resilience by allowing viewers to navigate emotional ambiguity and uncertainty.
9. Algorithmic censorship and recommendation systems pose a greater threat to freedom than government regulation, as they operate invisibly and with limited liability.
10. If you ignore your automated vehicle's maintenance alerts, it will schedule its own appointment and drive there.

These propositions are regarded as opposable and defensible, and have been approved as such by the promoters Prof. Dr. Riender Happee, Dr. Barys Shyrokau and Dr. Mohsen Alirezai.

CONTROL OF EVASIVE MANOEUVRES FOR AUTOMATED VEHICLES

SOLVING THE EDGE CASES

This thesis addresses the challenge of controlling automated vehicles performing evasive manoeuvres at the limit of handling. Special attention is paid to the development of nonlinear controllers, which can prioritise obstacle avoidance over path tracking objectives while considering vehicle stability constraints, to improve passenger safety. The thesis develops the entire pipeline for obstacle avoidance controllers, focusing on three aspects: vehicle state estimation, collision avoidance and control beyond the stable handling limits, e.g. drifting.

