CHALMERS
UNIVERSITY OF TECHNOLOGY

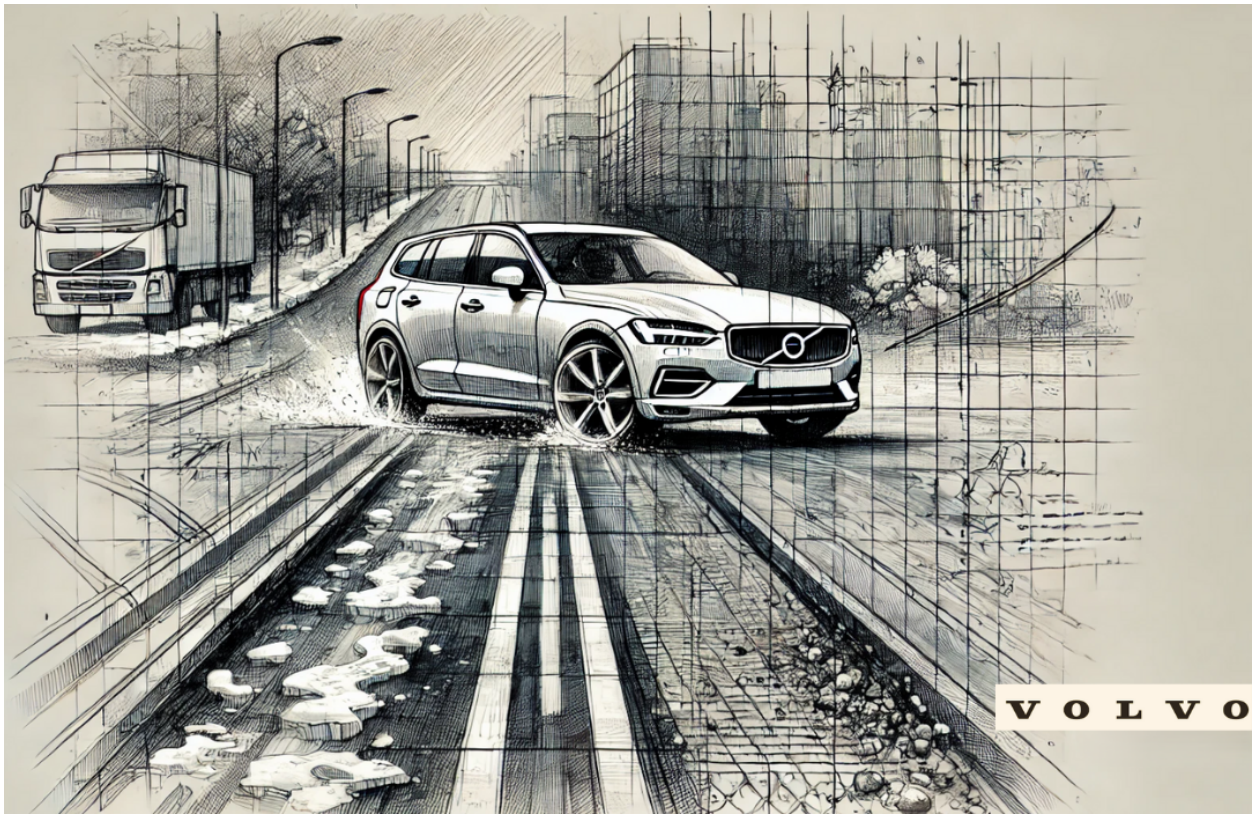

Road Friction Aware Adaptive Cruise Control using Robust Nonlinear Model Predictive Control with Uncertainty Quantification

A Robust Nonlinear Model Predictive Control Framework for Safe and Comfortable Speed Regulation under Varying and Uncertain Road Conditions

Master's thesis in Systems, Control and Mechatronics (MPSYS)

MUBINA CRNIC

SOTIRIS KOUTSOFTAS

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

Master's thesis 2025

Road Friction Aware Adaptive Cruise Control using Robust Nonlinear Model Predictive Control with Uncertainty Quantification

A Robust Nonlinear Model Predictive Control Framework for Safe and Comfortable Speed and Distancing Regulation under Varying and Uncertain Road Conditions

MUBINA CRNIC
SOTIRIS KOUTSOFTAS

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Division of Systems and Control
Chalmers University of Technology
Gothenburg, Sweden 2025

Road Friction Aware Adaptive Cruise Control using Robust Nonlinear Model Predictive Control with Uncertainty Quantification
A Robust Nonlinear Model Predictive Control Framework for Safe and Comfortable Speed and Distancing Regulation under Varying and Uncertain Road Conditions
MUBINA CRNIC
SOTIRIS KOUTSOFTAS

Supervisor: Ektor Karyotakis, Derong Yang, Vehicle Motion Control, Volvo Cars
Examiner: Nikolce Murgovski, Department of Electrical Engineering, Chalmers University of Technology

Master's Thesis 2025: EENX30
Department of Electrical Engineering
Division of Systems, Control and Mechatronics
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Artistic rendering of a Volvo vehicle navigating an urban road in wet conditions, blending technical sketch and digital illustration to highlight design and performance. Image generated using AI (ChatGPT image tools).

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2025

Road Friction Aware Adaptive Cruise Control using Robust Nonlinear Model Predictive Control with Uncertainty Quantification
A Robust Nonlinear Model Predictive Control Framework for Safe and Comfortable Speed and Distancing Regulation under Varying and Uncertain Road Conditions
MUBINA CRNIC, SOTIRIS KOUTSOFTAS
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Road friction is a fundamental factor affecting vehicle safety, especially under adverse weather conditions such as rain, snow or ice. Reduced road grip increases stopping distances and the likelihood of loss of control, contributing to a significant number of traffic accidents each year. While modern Advanced Driver Assistance Systems (ADAS) are designed to enhance driving safety, many systems still lack the ability to adapt their behaviour in real time to varying road friction levels, limiting their effectiveness in low-friction or rapidly changing weather conditions. This thesis addresses this limitation by developing a friction-aware, curvature-adaptive Adaptive Cruise Control (ACC) framework based on a Robust Nonlinear Model Predictive Control (NMPC) formulation.

The controller anticipates road curvature and spatially-varying friction by incorporating predicted profiles into the optimization problem. To account for uncertainty, friction is represented through position-dependent bounds in the MPC, while stochastic realizations of friction are generated via Beta sampling and the lead vehicle's acceleration varies randomly within feasible physical limits. These variations are reflected in dynamic safety constraints that adapt to evolving conditions ahead of the ego vehicle.

Soft constraint relaxation is introduced to maintain feasibility under conflicting demands such as sudden friction drops or aggressive lead vehicle deceleration. Despite the nonlinear nature of the problem and the presence of uncertainty, the controller operates using an efficient CasADi-based implementation in MATLAB.

Simulation results demonstrate that the proposed framework achieves robust, adaptive cruise control under friction and road curvature changes with stochastic disturbances. The integration of environmental awareness into predictive control enables safer, more responsive vehicle behaviour without sacrificing passenger comfort, and highlights the feasibility of embedding the friction information into future ADAS systems.

Keywords: ACC, ADAS, MPC, Friction-Aware Control, Robust Control, Road Friction, Speed Control.

Acknowledgements

We would like to express our sincere gratitude to our academic supervisor, Nikolce Murgovski at Chalmers University of Technology for his continuous guidance, valuable feedback and support throughout the course of this thesis.

We are also thankful to our industrial supervisors at Volvo Cars, Ektor Karyotakis and Derong Yang, for providing us with the opportunity to carry out this work within a real-world context. Their expertise and constructive input have been crucial in shaping the practical aspects of this project.

Additionally, we appreciate the support and collaboration from colleagues, researchers, and fellow students who contributed with insights and technical discussions during the development of this work.

Finally, we would like to thank our friends and families for their encouragement, patience, and understanding throughout this journey.

Mubina Crnic & Sotiris Koutsoftas, Gothenburg, 09 2025

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

ACC	Adaptive Cruise Control
ADAS	Autonomous Driving Assistance Systems
AEB	Autonomous Emergency Braking
AEB	Autonomous Emergency Braking
AS	Additional Scenario
EALC	Emergency Automated Lane Change
CC	Cruise Control
CTH	Constant Time Headway
CSA	Curve Speed Assist
DMS	Driver Monitoring Systems
KPIs	Key Performance Indicators
LKA	Lane Keeping Assist
MPC	Model Predictive Control
NMPC	Non-linear Model Predictive Control
ODE	Ordinary Differential Equation
PID	Proportional - Integral - Derivative
RK4	Fourth-order Runge-Kutta integration method
RNMPC	Robust Non-linear Model Predictive Control
UC	Use case
ZOH	Zero-Order Hold

Nomenclature

t	Continuous time variable [s]
k	Discrete timestep index
dt	Discretization timestep [s]
g	Gravitational acceleration ($9.81\,\mathrm{m/s^2}$)
N	Prediction horizon length
$\mu(s)$	Predicted friction coefficient along space
$s(t)$	Longitudinal position of ego vehicle at time t [m]
$v(t)$	Ego vehicle velocity at time t [m/s]
$a(t)$	Ego vehicle acceleration at time t [$\mathrm{m/s^2}$]
$j(t)$	Ego vehicle jerk (rate of change of acceleration) at time t [$\mathrm{m/s^3}$]
$x(t)$	Ego vehicle state vector at time t: $[s(t), v(t), a(t)]^\top$
$u(t)$	Control input, defined as $j(t)$ at time t [$\mathrm{m/s^3}$]
$\kappa(s)$	Road curvature along space
$R(s)$	Curve radius, defined as $R(s) = 1/\kappa(s)$
t_k	Discrete time instant at step k [s]
s_k	Longitudinal position of the ego vehicle at step k [m]
v_k	Ego vehicle velocity at step k [m/s]
a_k	Ego vehicle acceleration at step k [$\mathrm{m/s^2}$]
j_k	Ego vehicle jerk at step k
x_k	Ego vehicle state at discrete time step k: $[s_k, v_k, a_k]^\top$
u_k	Control input (jerk) at discrete time step k [$\mathrm{m/s^3}$]
$f(x(t), u(t))$	Continuous-time vehicle dynamics function
$f_d(x, u)$	Discrete-time state transition map (generic)
$f_d^{\mathrm{RK4}}(x_k, u_k)$	Discrete-time state transition using Runge-Kutta integration
$f_d^{\mathrm{ZOH}}(x_k, u_k)$	Discrete-time state transition using Zero-Order Hold
$\ell(x, u, s)$	Stage cost function in NMPC

$\ell_f(x)$	Terminal cost function in NMPC
$\mathcal{X}(i)$	Feasible state constraint set at prediction step i
$\mathcal{U}(i)$	Feasible input constraint set at prediction step i
$\mathcal{W}(i)$	Admissible disturbance set at prediction step i
$\mathcal{X}_f$	Terminal (final) state constraint set enforced at the end of the prediction horizon
i	Prediction-step index ($i = 0, \ldots, N$)
T	Horizon length in time, $T = N \cdot dt$ [s]
$\hat{x}_k$	Measured or estimated initial state at time t_k
$g(x, u, w)$	Vector of time-varying constraint functions
w_{k+i}	Disturbance/uncertainty vector (enters only via constraints)
s_{k+i}	Slack-variable vector used for constraint relaxation
m	Dimension of the slack-variable vector s_{k+i}
$\mathbb{R}^m_+$	Nonnegative orthant for slack variables (componentwise ≥ 0)
$s^{(1)}$	Slack variable for relaxing the safety distance constraint
$s^{(2)}$	Slack variable for relaxing the velocity constraints
$s^{(3)}$	Slack variable for relaxing the acceleration limits
$s^{(4)}$	Slack variable for relaxing the comfort-oriented acceleration limits
$v_{\min}$	Minimum allowable vehicle velocity [m/s]
$v_{\max}$	Maximum allowable vehicle velocity [m/s]
v_{ref}	The desired cruise speed set by the ego car's driver.
$a_{\max}$	Physical acceleration/deceleration limit [m/s^2]
$a_{c,\max}$	Maximum allowable comfort-oriented acceleration/deceleration [m/s^2]
$s_{\text{ego}}^{\text{pred}}(t_{k+1})$	Predicted longitudinal position of the ego vehicle at the next time step t_{k+1} [m]
$s_L^{\text{pred}}(t_{k+1})$	Predicted longitudinal position of the lead vehicle at the next time step t_{k+1} [m]
$\delta s_{\min}(t_k)$	Minimum required safe distance at time t_k [m]
$\sigma(s;\, p)$	Double-sigmoid function used for modeling friction and curvature profiles
$p_{\text{fric,low}}$& $p_{\text{fric,high}}$	Parameter vectors defining lower and upper bounds on friction profile
$p_{\text{curve,low}}$& $p_{\text{curve,high}}$	Parameter vectors defining lower and upper bounds on curvature profile

$\mu_{\text{low}}(s)$& $\mu_{\text{high}}(s)$	Lower and upper bound estimates of predicted friction at position s
$\kappa_{\text{low}}(s)$& $\kappa_{\text{high}}(s)$	Lower and upper bound estimates of predicted curvature at position s
p_1	Transition steepness parameter that controls how rapidly road conditions change. Higher values create sharp transitions between road segments, while lower values create gradual changes that may be more realistic for natural road conditions.
p_2	Initial road condition value at the beginning of the route segment under consideration.
p_3	Intermediate road condition value representing the middle segment characteristics.
p_4	Final road condition value at the end of the route segment.
p_5	Spatial position where the first transition occurs, typically representing the entry point to a different road surface type.
p_6	Spatial position where the second transition occurs, typically representing the exit point from the intermediate road condition.
μ_{bound}	The position-dependent uncertainty bound
s_{rel}	The relative position along the prediction horizon, measured from the current vehicle position.
$\mu_{unc,current}$	The friction uncertainty at the vehicle's current position, where camera classification confidence is typically highest.
$\mu_{unc,far}$	The friction uncertainty at distant positions along the prediction horizon, where camera confidence is typically lower.
$s_{unc,max}$	The maximum preview distance over which uncertainty modeling is considered relevant.
$\mu_{\text{actual}}(s)$	Realized (sampled) road friction coefficient at position s used in stochastic simulations
$\mu_{\text{mean}}(s)$	Expected (mean) road friction coefficient at position s
M	Normalized mean friction level within the uncertainty bounds
α	First shape parameter of the Beta distribution controlling its skewness
β	Second shape parameter of the Beta distribution controlling its skewness
ρ	Peak factor controlling the concentration (sharpness) of the Beta distribution around the mean value
$x_L(k)$	Lead vehicle state vector at discrete time k: $[s_L(k), v_L(k), a_L(k)]^\top$
$s_L(k)$	Longitudinal position of the lead vehicle at step k [m]
$v_L(k)$	Velocity of the lead vehicle at step k [m/s]

$a_L(k)$	Acceleration of the lead vehicle at step k [m/s^2]
$a_{L,\max}(s)$	Worst-case (most aggressive) lead vehicle deceleration at position s, assuming higher bound of friction uncertainty [m/s^2]
$a_{L,\text{actual}}(k)$	Realized (sampled) lead vehicle acceleration at discrete time k during stochastic simulations [m/s^2]
$w_{a,\text{lead}}$	Physical upper bound on lead vehicle braking capability used in prediction [m/s^2]
ξ_k	Uniform random variable $\sim \mathcal{U}[-1, 1]$ used to generate random lead vehicle acceleration commands
$\Delta s_{\text{margin}}(k)$	Predicted safety margin at step k, defined as the available gap to the lead vehicle minus the required minimum safety distance [m]
d_{brake}	Required braking distance to reduce ego vehicle velocity from v_{ego} to $v_{\text{L,min}}$ under constant deceleration a [m]
J	Total cost over the prediction horizon, combining stage and terminal costs
Q	State weighting matrix in cost function
R	Input (jerk) weighting matrix in cost function
R_s	Slack variables weighting matrix in cost function
x_{ref}	Desired reference state vector used for tracking in the cost function
x_N	Ego vehicle terminal state at the end of the horizon ($k = N$)

Contents

List of Figures

List of Tables

1

Introduction

Road traffic safety remains a major global public health concern. In the United States alone, over 5.8 million vehicle crashes occur annually, with around 21%, nearly 1.2 million being weather-related. These crashes occur under adverse conditions such as rain, sleet, snow, fog, strong crosswinds, or on slick pavements including wet, snowy, slushy or icy surfaces. Each year, weather-related crashes result in nearly 5000 fatalities and over 418000 injuries [1].

Most weather-related crashes happen on wet pavement, around 70% and during rainfall around 46%. A smaller but still significant portion occurs in winter conditions: 18% during snow or sleet, 13% on icy roads, 16% on snowy or slushy roads and 3% in fog [1].

Similarly, in Sweden, despite having one of the lowest road fatality rates in the EU, slippery conditions have a profound impact on road safety. Approximately 50% of all fatal passenger car accidents occur on snowy or icy roads, rising to 90% in rural northern Sweden. Moreover, the risk of accidents is about five times higher on snow or ice covered roads compared to dry ones [2].

To address these safety challenges, technological solutions like ADAS (Advance Driving Assistance Systems) have been developed to improve vehicle safety in both normal and adverse conditions. ADAS features such as Lane Keeping Assist (LKA), Driver Monitoring Systems (DMS) and Automatic Emergency Braking (AEB) have been shown to reduce crash rates: LKA by 19.1%, DMS by 14%, and AEB by 10.7%. However, ACC (Adaptive Cruise Control) and Cruise Control (CC) have been linked to increased crash rates of 8% and 12%, respectively [3].

The increased crash rates associated with ACC and CC raise questions about how well these systems handle varying road conditions, especially in challenging weather. One possible explanation is that road conditions, such as reduced friction from rain, snow or ice, may affect their performance and safety.

To address these challenges, this thesis develops an ACC system that proactively accounts for both road friction variations and road curvature within a unified control framework. The proposed controller continuously adjusts the vehicle's longitudinal behaviour based on predicted environmental conditions, maintaining safe following distances and ensuring stable, comfortable driving even through curves and under uncertain or low-friction scenarios.

1.1 Related Work

Advancements in ADAS have increasingly emphasised the need for robust vehicle control methods, especially for functions such as the AEB, Emergency Automated Lane Change (EALC), Curve Speed Assistance (CSA) and ACC. Among the various control strategies explored, MPC (Model Predictive Control) has emerged as a leading method due to its ability to explicitly handle system constraints, predict future behaviour and balance multiple performance objectives.

One of the foundational contributions to longitudinal vehicle following using MPC was presented by Wang et al. [4]. The hierarchical control scheme used combined an upper-level MPC with a Constant Time Headway (CTH) policy, and a lower-level Proportional-Integral-Derivative (PID) controller for actuation. Interestingly, they proposed a novel inter-vehicle longitudinal dynamic model that considers the acceleration and velocity of the lead vehicle as disturbances. This differs from many conventional approaches that only rely on relative distance and velocity as state variables. By formulating the spacing dynamics in this way, the model more accurately captures the system's behaviour in real traffic situations, especially under variations in the lead vehicle's motion.

Another key innovation in the paper, [4], is the integration of an ancillary control law into the MPC formulation. This control structure splits the input into a nominal component and a feedback term that compensates for disturbances. The approach enabled better disturbance rejection while guaranteeing recursive feasibility and convergence. The simulation studies presented validate the benefits of this formulation, particularly in scenarios with speed fluctuations or changes in road friction.

This modeling structure directly influenced subsequent MPC frameworks, especially those incorporating friction awareness and disturbance robustness. Although our own implementation details are discussed later, the core structure of disturbance modeling and controller layering presented by Wang et al. laid the conceptual groundwork.

Expanding on this, Mekala et al. [5] developed an MPC framework targeting longitudinal speed control using LIDAR inputs. Their real-time validated solution showcased effective adaptation to both low and high-speed scenarios, highlighting MPC's versatility and feasibility in real-world autonomous vehicle applications.

Luo et al. [6] proposed a multi-objective MPC strategy that incorporated safety, comfort, fuel efficiency and tracking precision within its cost function. By assigning appropriate weights to each objective, it has been demonstrated that the MPC could effectively balance conflicting performance criteria.

A practical contribution by Naus et al. [7] focused on parameterised MPC for ACC systems that enabled user adjustable comfort and safety settings. This approach simplified calibration and facilitated personalization of controller behaviour, aiding the transition from theoretical designs to practical deployment.

Recognising the uncertainty in the the tyre-road interaction, Vaskov et al. [8] introduced a stochastic MPC framework incorporating Bayesian friction estimation. This probabilistic control design allowed safe trajectory tracking under variable road conditions and underscored MPC's adaptability to uncertainty.

Complementing this, Svensson and Törngren [9] examined the complementary strengths and limitations of local and camera-based road friction estimation. This work emphasised that local methods, while accurate, lack foresight and availability during normal driving, whereas camera-based classification offers high foresight and availability but only coarse friction categorisation. A Gaussian Process regression based fusion method was proposed to combine these sources, to produce a conservative yet not overly pessimistic friction estimate. This approach improved the reliability of traction adaptive motion planning and supports the assumption in our work of utilising camera-based classification to define an interval of plausible friction values used in constraint formulation.

In parallel, Herzfeld and Thottathodi [10] proposed a comprehensive architecture that integrates dynamic road friction information into both braking and steering based collision avoidance strategies. This work demonstrated how simplified vehicle models can be used within an MPC framework to evaluate trajectory feasibility with respect to estimated friction profiles. The concept of using friction based approximations of braking and lateral force limits to constrain motion planning inspired aspects of our constraint formulation, particularly the incorporation of friction-aware safe distance definitions and control bounds in low grip scenarios.

Czibere et al. [11] tackled the challenge of handling vehicle dynamics at the limit of stability. Using an Robust Nonlinear Model Predictive Control (RNMPC) design for motion control under multiple equilibrium and variable surfaces proved to be capable of maintaining stability and tracking accuracy even in extreme driving scenarios.

Mattsson et al. [12] addressed energy-efficient longitudinal control by evaluating MPC against PID controllers in combustion engine and brake-actuated vehicles. The findings confirmed the advantage of MPC in optimising comfort and energy consumption simultaneously.

Finally, building on these foundations, Zhang et al. [13] proposed a real-time MPC approach for predictive ACC where several solvers and problem formulations were considered and compared. This work focused on computational efficiency while providing a pathway toward implementation in embedded automotive platforms.

These works illustrate the evolution and diversification of MPC in the domain of vehicle motion control. From theoretical robustness to real-world applicability, MPC has been proven to be a powerful tool for addressing the demands of modern autonomous and assisted driving systems, particularly in accommodating safety-critical constraints and environmental uncertainties such as road friction variations.

1.2 Problem Overview

This thesis addresses the challenge of improving ACC functionality in environments where road friction and curvature vary along the driving path. Traditional ACC systems typically rely on high friction assumptions and fail to account for spatial changes in road surface properties, which can severely impact vehicle safety and comfort. To overcome this limitation, we propose an enhanced ACC system capable of adapting its behaviour based on predicted variations in both road friction and curvature.

Two critical driving scenarios motivate this work:
Rear-End Collision under Varying Friction on a Straight Road: The first scenario considers an ego vehicle following a lead vehicle on a road that contains a localized low-friction segment, such as an ice patch (see Figure 1.1). Conventional ACC systems, which do not anticipate reduced braking capability due to friction loss, may delay deceleration, increasing the risk of rear-end collisions. This highlights the need for predictive adaptation based on estimated friction levels.

Figure 1.1: A rear-end collision caused by an low-friction surface, where reduced braking capability on the low-friction surface prevents the ego vehicle from stopping in time.

Curve Negotiation under Non-Uniform Road Friction: In the second scenario, a vehicle enters a curve where both the road curvature and friction change along the path (see Figure 1.2). If the ACC system fails to adjust the vehicle speed appropriately, the vehicle may enter the curve too fast, exceeding the frictional limits and risking instability which may lead to lane departure. Integrating curve awareness and friction constraints into ACC is thus essential to ensure safe handling.

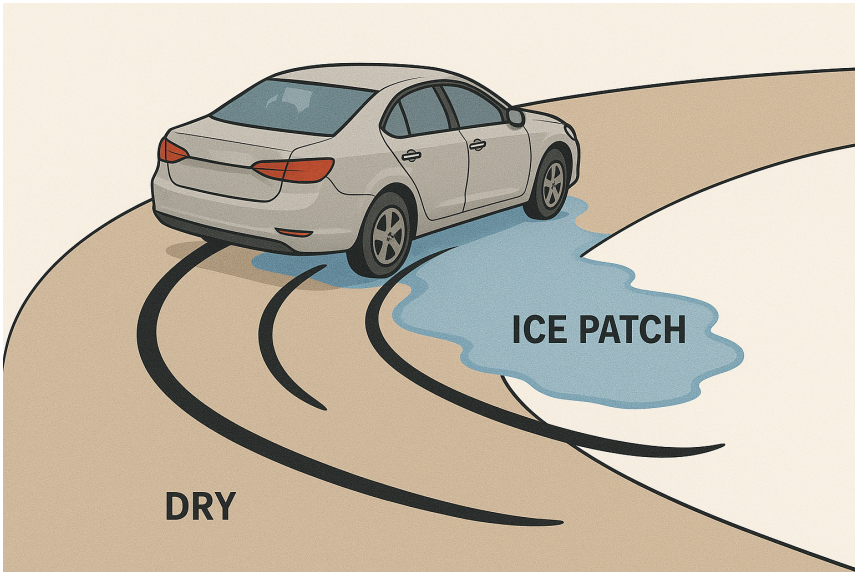

Figure 1.2: A vehicle entering a curve with a low-friction section, risking loss of control from insufficient speed adaptation.

To tackle these challenges, a Robust Nonlinear Model Predictive Control (RNMPC) formulation is employed. This approach allows the system to anticipate future variations in road friction and curvature along the prediction horizon and generate a safe, comfortable, and dynamically feasible speed profile. The proposed ACC framework combines safety constraints, friction-aware dynamics, and curvature adaptation into a unified controller that dynamically adjusts the vehicle's longitudinal motion. Stochastic uncertainty in both lead vehicle acceleration and future friction is modeled explicitly, enabling the controller to maintain safety while reducing unnecessary conservatism.

1.3 Scope & Research Questions

This thesis focuses on enhancing longitudinal vehicle control in the presence of spatially varying and uncertain road friction conditions. Specifically, it investigates how incorporating friction and road curvature into an ACC system can improve safety, comfort, and robustness in critical driving scenarios, such as icy straight segments or curved roads with reduced friction.

A single control strategy based on a stochastic friction-aware MPC framework is developed. The controller dynamically adjusts the vehicle's acceleration and velocity by accounting for predicted variations in road friction and curvature, while ensuring feasibility through soft safety and comfort constraints. The control formulation also considers uncertainty in lead vehicle behaviour and road surface properties.

The scope is limited to longitudinal motion control. Lateral dynamics, such as steering or obstacle avoidance, are beyond the scope of this work. All evaluations

are conducted in simulation, under the assumption that predicted road friction and curvature profiles are available along the vehicle's path.

Based on this scope, the following research questions are addressed:

- **Friction- and Curve-Aware ACC:** How can predicted road friction and curvature be effectively integrated into an MPC-based ACC system to improve vehicle safety in cruising and lead-vehicle following?
- **Robustness vs. Comfort Trade-offs:** How does a stochastic, constraint-softening control framework balance robustness, safety margins and driving comfort under friction uncertainty?
- **Safety Benefits:** To what extent can friction-aware ACC reduce collision risk in scenarios with sudden friction changes, such as icy patches and low-grip curves?

1.4 Delimitations

This thesis focuses exclusively on longitudinal vehicle motion control. Lateral dynamics and combined path-planning are excluded. The study is conducted in a simulation environment using a point-mass model of the ego vehicle. Detailed powertrain dynamics, actuation delays and embedded real-time implementation are considered out of scope.

Assumed environment:

- A lead vehicle is assumed to be present in all ACC-related scenarios, enabling continuous car-following behaviour. However, when road curvature limits the safe speed or if the driver sets a lower reference speed the system operates independently of the lead vehicle.
- Preview information of road friction and road curvature is assumed to be available the prediction horizon. Friction information is derived from a camera-based classification and modeled as a bounded interval, with stochastic realizations generated in simulation. Curvature is assumed deterministic and known in advance.

Controller formulation:

- A single friction-aware RNMPC is implemented in the time domain. It combines ACC with curve speed adaptation in a unified formulation, handling spacing and curve constraints simultaneously.
- Soft constraints are introduced to maintain feasibility in the presence of conflicting requirements. This ensures that safety-critical constraints are always prioritised, while the controller can still adaptively balance comfort and responsiveness, by violating the soft comfort constraint when safety cannot be achieved without violating it.

- Soft constraint utilisation improves the controller's responsiveness in scenarios where collision is not avoidable, by mitigating violations as better as possible, compared to a controller with hard constraints which will simply fail.

Model simplifications:

- The ego vehicle is represented by a third-order point-mass model, with jerk as the control input.
- Tyre-road interaction is simplified through a friction constraint envelope, without explicitly modeling tyre dynamics or actuator lag.

Out-of-scope topics: Sensor fusion, lateral control, obstacle avoidance, real-time deployment and experimental validation are not covered in this work but are identified as potential directions for future research and development.

1.5 Social and ethical aspects

Road friction- and curvature-aware ACC enables proactive adaptation to hazardous conditions, but this complexity introduces grave social and ethical challenges. First, user understanding and trust remain critical. Studies show that many drivers have weak mental models of ADAS systems and tend to over-rely on features like ACC, which may lead to inattentiveness or unsafe behaviours, especially when control boundaries blur during edge cases [14]. When ACC adjusts speed based on unknown or variable road friction, the risk of driver misunderstanding increases, underscoring the ethical necessity of clear system communication and user education.

Second, assigning responsibility in shared-control systems is increasingly challenging. In semi-autonomous vehicles, both the automated system and the human driver share control, the system handles tasks like speed regulation and constraint enforcement, while the driver supervises and can intervene when needed. However, as [15] point out, when errors occur, people tend to blame the human rather than the machine, making it harder to hold automated systems accountable. As stated in [15], current legal frameworks such as tort law do not fully account for this shared responsibility. In the context of friction-aware ACC, it is therefore critical to have transparent logging of system predictions (e.g., estimated friction or curvature) and control decisions. This allows for clear post-incident analysis, helping to determine whether the driver, the automation, or both were responsible.

Third, social fairness and accessibility are key ethical concerns when deploying advanced driver-assist systems. These technologies are often first introduced in high-end vehicles, meaning that wealthier individuals are more likely to benefit early on [16]. This can deepen existing inequalities in mobility, those with fewer resources may not have access to the same safety features or driving assistance. However, these systems also hold promise: if made widely available, they could improve independence and safety for elderly drivers or people with disabilities, who often face greater mobility challenges. But realising that benefit requires intentional policy

and design choices. That means ensuring that the technology is affordable, that user interfaces are accessible to a broad population, and that deployment does not prioritise only affluent or urban communities.

Fourth, privacy and data governance are paramount. ACC systems rely on continuous input from cameras and onboard sensors, which may inadvertently capture sensitive data such as road conditions, geolocation patterns, and even driver behaviour. While the context differs a bit, research on AI assistants used by older adults raises closely related concerns about data misuse, transparency, autonomy and trust in systems that rely on personal data [17]. Emphasising the importance of designing technologies that minimise unnecessary data collection and prioritise user agency in data sharing. These principles are equally applicable to vehicle automation: friction-aware ACC should adopt privacy-by-design strategies, limit data capture to what's essential and implement clear, understandable consent mechanisms.

Finally, ethical decision-making in real time becomes operationally significant. While ACC controllers aren't making life-and-death choices, they continuously weigh safety vs. performance under uncertainty, such as managing grip on wet curves or reaction to emergencies. Yet established ethical frameworks for such micro-decisions are lacking. This absence raises concerns about consistency, fairness, and responsibility. Transparent system limits and well-defined operational design domains are necessary to maintain trust and ethical consistency.

1.6 Thesis Outline

The remainder of this thesis is organised as follows. Chapter 2 collects the necessary background on ACC, curve-speed adaptation and MPC and motivates the choice of a time-domain formulation that treats road friction as a stochastic previewable quantity. Chapter 3 translates the overall research goal into concrete safety, comfort and real-time requirements, derives the Key Performance Indicators (KPIs) used later in the evaluation, and sets out ten use cases that stress the controller under varying curvature and friction conditions. Chapter 4 presents the core technical contribution: an uncertainty-aware RNMPC. After introducing the vehicle model, friction and curvature profiles are formulated and embedded in a probabilistic framework and derived safety and comfort constraints with hierarchical slack variables. The chapter closes with implementation details in CasADi/Ipopt. Chapter 5 reports the quantitative evaluation. First, solver speed and discretisation accuracy are studied, then a use-case by use-case analysis of distance keeping is being provided, velocity tracking, control effort and constraint relaxations. Chapter 6 discusses the implications of these findings, examining how much additional safety margin is gained, what comfort trade-offs arise and how close the implementation is to real-time feasibility. Chapter 7 concludes by summarising the main contributions and outlining directions for future work, such as hardware-in-the-loop validation and fusion with lateral motion planning. Finally, the appendix collects complete parameter tables and solver settings to facilitate reproduction of the results.

2

Preliminaries

2.1 Vehicle Following and Speed Tracking via ACC

To understand the challenge of autonomous vehicle control, imagine driving on a highway during changing weather conditions. As a human driver, one would instinctively adjust the following distance when rain begins to fall, knowing that wet roads require longer stopping distances. The person would also naturally slow down when approaching curves, understanding that taking a turn too fast could cause loss of control. These intuitive adjustments that humans instinctively make constantly and continuously while driving represent complex decision-making processes that autonomous systems must replicate mathematically.

ACC represents one of the fundamental building blocks toward this goal. Unlike traditional CC systems that blindly maintain a set speed regardless of traffic, ACC automatically adjusts the vehicle's velocity to maintain a safe distance from the car ahead. ACC can be regarded as a virtual co-pilot that continuously monitors the distance to the lead vehicle and adjusts the speed smoothly to maintain safety.

2.1.1 The Limitations of Traditional Approaches

Conventional ACC systems typically employ feedback control strategies, with PID controllers being the most common. These controllers work on a simple principle: measure the current error (such as being too close to the lead vehicle), and apply a control action proportional to that error and its history. While this approach works adequately under ideal conditions, it suffers from a fundamental limitation. It can only react to what has already happened, not anticipate what might happen next.

Considering what happens when road conditions deteriorate, a PID-based ACC system maintaining a fixed following distance might suddenly find itself unable to stop in time when the lead vehicle brakes on an icy patch. The controller's reactive nature means it only discovers the reduced traction when braking is already urgently needed, precisely when safety margins have already been compromised.

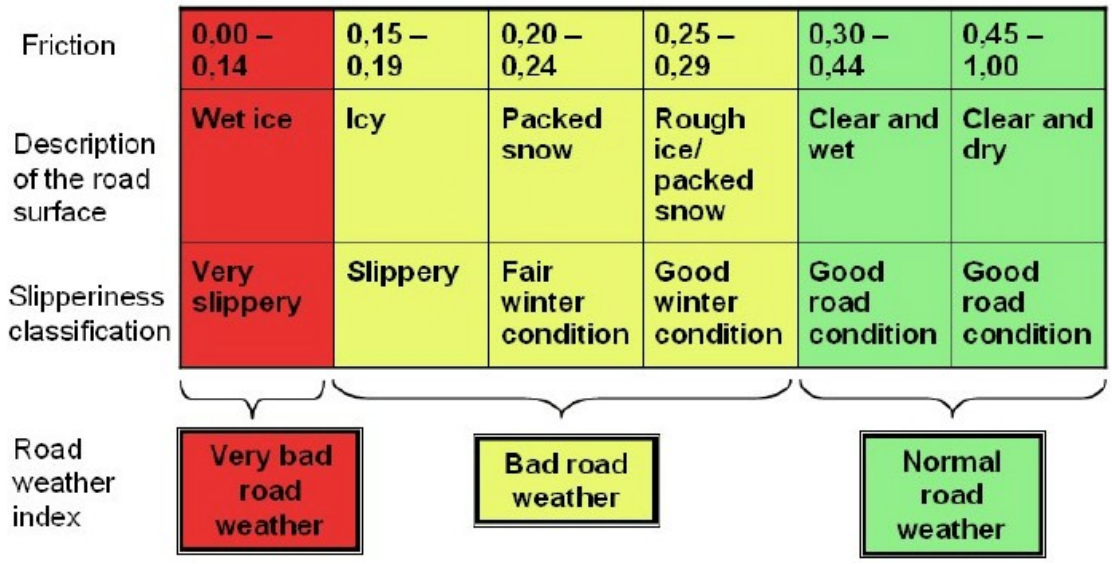

Figure 2.1: Road weather classification made by Finnish Road Administration [18]

This reactive limitation becomes even more problematic when we consider the wide variation in real-world driving conditions. The friction coefficient, denoted by μ, between tyres and road surface can vary dramatically, as shown in Figure 2.1.

A control system that assumes dry conditions might command deceleration rates that simply cannot be achieved on wet or icy surfaces, leading to deterioration and dangerous situations.

2.1.2 The Predictive Control Advantage

This fundamental limitation motivates the adoption of MPC for ACC implementation. Unlike reactive controllers, MPC looks ahead, predicting future vehicle trajectories over a time horizon. This preview capability enables the controller to anticipate problems before they become critical.

To understand how MPC works in the ACC context, imagine planning a sequence of chess moves rather than reacting to each opponent's move individually. The controller continuously solves an optimisation problem: "What sequence of accelerations over the next few seconds will best maintain safe following distance while respecting all physical constraints?". Only the first action from this optimal sequence is implemented, after which the process repeats with updated information, a strategy known as receding horizon control.

2.1.3 Integrating Uncertainty into ACC

A key innovation in this thesis involves incorporating probabilistic friction estimation into the ACC framework. Traditional systems might assume either perfect road conditions or worst-case conditions. Instead, our approach uses camera-based road surface classification to estimate friction with associated confidence levels.

Consider how this works in practice. A forward-facing camera identifies road surface conditions ahead, classifying segments as dry, wet or icy. However, these classifications come with uncertainty, the system might be 90% confident about a nearby classification but only 60% confident about surfaces farther ahead. The ACC system translates this uncertainty into position-dependent bounds on available friction, allowing it to plan conservatively for uncertain regions while taking advantage of high-confidence information where available.

This probabilistic approach enables intelligent trade-offs. When approaching a surface with low confidence that the surface might be icy, the controller gradually increases following distance as a precaution. If the classification confidence improves as the vehicle gets closer, the controller can adjust its plan accordingly. This behaviour mimics how human drivers naturally become more cautious when road conditions are uncertain.

2.2 Curve Adaptation in Longitudinal Control

While ACC handles longitudinal control relative to other vehicles, safe driving also requires adapting speed to road geometry. This is sometimes referred to in literature as CSA, but in this work it is treated as an integral part of the ACC logic and simply referred to as curve adaptation. The goal is to ensure that the vehicle does not attempt to navigate curves at unsafe speeds.

2.2.1 The Physics of Cornering

When a vehicle travels through a curve, it requires centripetal force to change direction. This force is provided by the friction between tyres and the road surface. The relationship is captured by a fundamental equation:

$$a_{\text{lateral}} = \frac{v^2}{R} = v^2 \cdot \kappa \tag{2.1}$$

Here, v is the vehicle's speed, R is the curve radius, and $\kappa = 1/R$ is the curvature. The lateral acceleration a_{lateral} represents how hard the vehicle is pushing sideways against the road. For the vehicle to successfully navigate the curve without skidding, this lateral acceleration cannot exceed what friction can provide:

$$a_{\text{lateral}} \leq \mu \cdot g \tag{2.2}$$

where μ is the coefficient of friction and g is gravitational acceleration (9.81 m/s^2). Combining these equations reveals the maximum safe speed, v_{maxC}, for a given curve:

$$v_{\text{maxC}} = \sqrt{\frac{\mu \cdot g}{\kappa}} \tag{2.3}$$

This relationship has profound implications. On a sharp curve (high κ) with icy conditions (low μ), the safe speed might be drastically lower than the posted speed limit. Conversely, a gentle highway curve on dry pavement might safely accommodate speeds well above typical traffic flow.

2.2.2 The Challenge of Unknown Friction

Handling curves safely requires accurate knowledge of the available road friction. Traditional systems often rely on fixed friction assumptions when it comes to curves and are either optimistic (high μ) or pessimistic (low μ). Optimistic assumptions risk vehicle instability when the real road surface provides less grip than expected. On the other hand, overly cautious assumptions result in unnecessarily slow driving, reducing comfort and efficiency.

In this thesis, curve handling is not treated as a separate functionality but is integrated directly into the ACC logic as a form of curve adaptation. The same probabilistic friction estimation used to regulate safe distances to a lead vehicle is also used to determine safe speeds through curves.

For example, when approaching a highway curve after rain, the surface classification system may detect wet pavement with high confidence near the vehicle but with increasing uncertainty further ahead. The controller uses the lower bound of the estimated friction to compute safe longitudinal acceleration and velocity limits, ensuring the vehicle enters the curve at a speed that respects both the road curvature and the friction envelope, even under uncertain conditions. This approach ensures safety without introducing excessive conservatism.

2.2.3 Unified Constraint Management

In real-world driving, safe longitudinal control must account for both surrounding traffic and road geometry. For instance, when approaching a curve while following another vehicle, the ego vehicle's speed may need to be limited either by the lead vehicle's behaviour or by the curve's geometry, whichever imposes the stricter constraint. Importantly, both limitations are influenced by the available road friction, which is often uncertain and spatially varying.

This interdependence is addressed in this thesis by formulating a unified MPC framework where both distance-keeping and curve adaptation are treated as coupled constraints within a unified ACC framework. Rather than separating functionalities, the controller evaluates all relevant constraints, derived from lead vehicle motion, road curvature, and predicted friction at every step of the planning horizon.

By incorporating all factors simultaneously, the system ensures safe behaviour without conflicting priorities between subsystems. This unified treatment enables the vehicle to adaptively balance safety, comfort and efficiency in complex driving environments.

2.3 Model Predictive Control Framework

Having motivated the need for predictive and uncertainty-aware control, the MPC framework that enables this is now examined. MPC represents a fundamental shift from reactive to proactive control strategies.

2.3.1 Core Principles of MPC

The essence of MPC can be understood through an analogy to human driving behaviour. When approaching a red traffic light, experienced drivers don't wait until the last moment to brake hard. Instead, they see the light ahead, predict their future position, and begin decelerating smoothly to arrive at the stop line just as they reach zero speed. This anticipatory behaviour is exactly what MPC enables for automated systems.

Mathematically, MPC operates by solving an optimisation problem at each control instant. The optimisation considers: (i) the prediction model, which predicts how the vehicle will behave over the next few seconds given different control inputs, (ii) the constraints, which define the physical limits and safety requirements, and (iii) the objectives, which define what constitutes good behaviour, such as how smooth the riding is, how the efficient operation is and how well the reference is being tracked.

The controller finds the sequence of control actions that best satisfies the objectives while respecting all constraints. Crucially, it then applies only the first action from this sequence before re-solving the optimisation with updated information-the receding horizon principle.

2.3.2 Advantages for Vehicle Control

MPC offers several compelling advantages for vehicle control applications:

Explicit Constraint Handling: Safety requirements can be encoded directly as mathematical constraints. Rather than relying on the controller not to violate safety bounds, MPC guarantees constraint satisfaction (within the accuracy of the prediction model).

Preview Integration: Modern vehicles equipped with cameras, radar, and map data have information about conditions ahead. MPC naturally incorporates this preview information into its predictions, enabling proactive responses to upcoming challenges.

Multi-objective optimisation: Real-world driving involves balancing multiple goals, safety, comfort, efficiency, and progress. MPC handles these competing objectives through a unified cost function that can be tuned to achieve desired behaviour.

Uncertainty Management: Through robust optimisation techniques, MPC can account for uncertainty in predictions, ensuring safety even when perfect information isn't available.

2.3.3 The Prediction-optimisation Loop

The computational heart of MPC lies in repeatedly solving optimisation problems. At each control step, the process begins by (i) measurement, where the current vehicle state and environmental information are collected; (ii) prediction, where models are being used to forecast future evolution under different control scenarios; (iii) optimisation, where the control sequence that minimises the cost while satisfying constraints is being determined; and (iv) application, where the first control action is implemented. This process is then repeated by shifting the horizon forward and starting again.

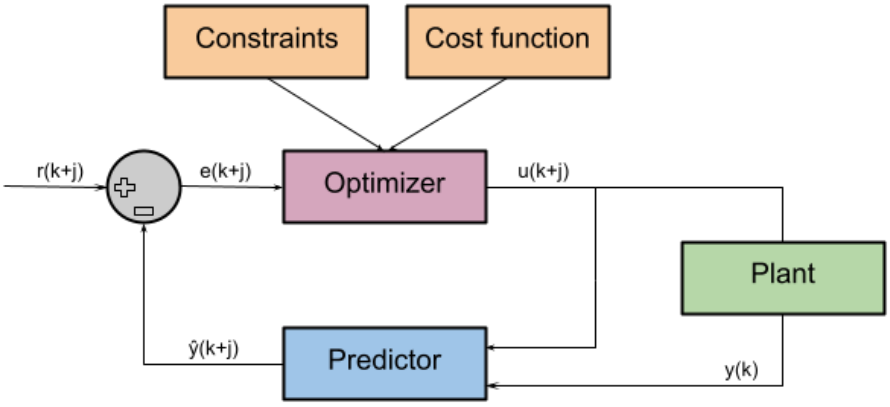

Figure 2.2: Block diagram illustrating the MPC framework.

2.4 Unified Time-Domain MPC Implementation

With the foundational concepts established, this section presents how these ideas are combined in the unified controller developed in this thesis. Several design choices are made to enable practical deployment while maintaining theoretical rigour.

2.4.1 Time-Domain Formulation: Rationale and Trade-offs

While both time-domain and spatial (distance-based) parameterisations are viable for vehicle MPC, this thesis employs a time-domain formulation. Rather than claiming universal advantages, it is acknowledged that each approach has distinct trade-offs. The rationale for this choice is as follows.

Advantages of Time-Domain Formulation:

- **Intuitive Lead Vehicle Modeling:** Time-domain formulations naturally incorporate lead vehicle dynamics through temporal coordination, enabling direct representation of relative velocity and acceleration profiles [4].

- **Natural Constraint Expression:** Speed limits, acceleration bounds, and comfort requirements are typically specified in time-domain units, aligning with human driving experience and regulatory frameworks [7].
- **Consistent with Control Implementation:** Vehicle actuators operate with time-based control loops, making temporal formulation more directly implementable [5].

Disadvantages and Spatial Domain Benefits:

- **Speed-Dependent Prediction Range:** The spatial coverage of a fixed time horizon varies with vehicle speed, potentially providing insufficient lookahead at high speeds or excessive computation at low speeds.
- **Friction Modeling Complexity:** Since road friction varies with position rather than time, spatial formulation could provide more linear relationships for friction-dependent constraints [8].
- **Sensor Fusion Challenges:** While some sensors operate in time domain, others like cameras and radar often have fixed spatial ranges, creating mixed temporal-spatial information that must be reconciled [9].

The time-domain formulation has been chosen for this implementation, primarily due to its intuitive modeling approach, particularly for lead vehicle prediction and human-interpretable constraints. The friction-position relationship is handled through spatial interpolation within the temporal framework, accepting the additional complexity for the benefit of more natural vehicle following dynamics.

2.4.2 Probabilistic Friction Integration

A distinguishing feature of this implementation is its sophisticated handling of friction uncertainty. The system maintains two representations:

- **Deterministic Mode:** Used as a baseline, this mode assumes friction follows the expected (mean) profile without deviation. This provides a performance benchmark and simplifies initial testing.
- **Stochastic Mode:** The full implementation models friction as a random variable following a Beta distribution. The distribution parameters are chosen to respect the confidence bounds from camera-based classification while concentrating the probability mass near the expected value.

The Beta distribution is particularly well-suited for this application because:

- It's naturally bounded to a finite interval (matching physical friction limits)
- Its shape can be adjusted to represent different confidence levels
- It allows efficient sampling for Monte Carlo evaluation

2.4.3 Constraint Softening Strategy

Real-world driving occasionally presents scenarios where all constraints cannot be simultaneously satisfied. For example, if a lead vehicle brakes harder than physically possible given friction estimates, maintaining the minimum following distance might require deceleration exceeding comfort limits. The implementation addresses this through a carefully designed slack variable framework. Each constraint type receives a slack variable with associated penalty weight:

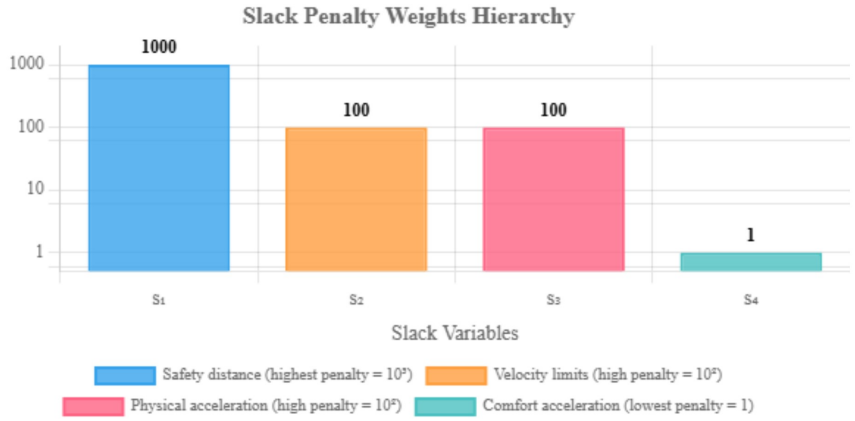

Figure 2.3: Slack penalty weights showing constraint relaxation hierarchy

This hierarchy ensures that when trade-offs are necessary, passenger comfort is relaxed before safety. In addition to that, the soft constraint implementation results in collision mitigation due to preventing of the solver's failure when the problem becomes infeasible, by trying its best to eliminate constraint violations.

2.4.4 Lead Vehicle Prediction

Predicting lead vehicle behaviour represents a crucial challenge for ACC systems. The implementation takes a robust approach, assuming the lead vehicle might brake at its physical limit (determined by the upper friction bound). This conservative assumption ensures safety even when following aggressive drivers.

The prediction also includes behavioural bounds based on typical human driving patterns. While a vehicle might theoretically achieve 1g deceleration on dry pavement, human drivers rarely brake harder than approximately 0.3 in non-emergency situations [19]. The controller uses these behavioural bounds to avoid overly conservative planning during normal driving while maintaining the ability to respond to genuine emergencies.

2.5 Contribution of This Thesis

This thesis advances the state of autonomous vehicle control in several significant ways, addressing critical gaps between academic research and practical deployment requirements.

2.5.1 Theoretical Contributions

Unified Uncertainty Framework: The thesis presents a comprehensive integration of probabilistic road surface classification with MPC-based longitudinal control. The position-dependent uncertainty model captures the realistic behaviour of perception systems while remaining computationally tractable for real-time implementation.

Constraint Integration Methodology: The work demonstrates how multiple safety-critical constraints (following distance, friction limits, curvature limits) can be unified within a single optimisation framework, eliminating conflicts between subsystems and ensuring consistent behaviour.

Robustness Analysis: The thesis provides a systematic analysis of controller behaviour under different uncertainty scenarios, quantifying the trade-offs between conservatism and performance.

2.5.2 Practical Contributions

Implementation Framework: The developed CasADi-based implementation provides a template for deploying advanced MPC strategies in real-time systems, with careful attention to computational efficiency and numerical reliability.

Validation Methodology: The testing framework covers diverse scenarios from routine highway driving to emergency situations, providing confidence in the controller's real-world applicability.

Performance Metrics: The thesis establishes metrics for evaluating the uncertainty-aware controller, enabling objective comparison with traditional approaches.

2.5.3 Broader Impact

The contributions of this thesis extend beyond the specific problem of longitudinal control. The uncertainty handling framework could be extended to lateral control, integrated trajectory planning and other domains where environmental uncertainty affects control decisions. The probabilistic approach to constraint formulation provides a template for other applications where perfect information is unavailable but safety must be guaranteed.

As autonomous vehicles move toward widespread deployment, the ability to handle real-world uncertainty becomes increasingly critical. This thesis provides theoretical foundations and practical tools for building systems that are safe, efficient, and robust to the complexities of real-world driving.

3

Requirements, KPIs & Use Cases

3.1 Definition of Requirements

The camera assisted, friction-aware RNMPC developed in this work addresses the critical challenge of vehicle control under uncertain road conditions. The system integrates real-time camera based road surface classification to achieve safe and comfortable autonomous driving.

Safety Requirements form the foundation of the controller design. The system must maintain a safe following distance under all conditions, accounting for both deterministic and uncertain road surface estimates. Vehicle speed is limited by curvature and friction constraints, with position-dependent uncertainty bounds that reflect increasing confidence of camera-based classification along the preview horizon. Safety requirements include preventing collisions through robust gap maintenance and ensuring acceleration never exceeds available friction.

Comfort Requirements ensure acceptable passenger experience by constraining acceleration and jerk to physiologically comfortable ranges. The controller avoids abrupt maneuvers and provides smooth, gradual adaptation to changing road conditions by leveraging the predictive horizon to anticipate and prepare for upcoming surface changes.

Robustness Requirements address the uncertainty in camera based road surface classification. The controller must remain feasible and effective even when surface classifications are inaccurate or actual friction falls at the conservative bounds of the estimated friction ranges. It must handle unpredictable lead-vehicle behaviour, including sudden acceleration or deceleration, while maintaining safety margins that account for both surface uncertainty and behaviour unpredictability. Robustness is achieved by balancing conservatism and performance under uncertainty.

Computational Requirements ensure real-time feasibility on automotive hardware. The MPC optimisation must converge within acceptable time limits to enable closed-loop control at the specified sampling frequency. The controller design incorporates efficient warm-starting strategies and carefully structured constraint formulations to maintain computational tractability while preserving the rich uncertainty representation.

3.2 Key Performance Indicators

The assessment of controller performance relies on a comprehensive set of quantitative metrics that capture safety, comfort, robustness and computational aspects of the system behaviour.

Safety Metrics provide fundamental measures of collision avoidance and constraint satisfaction. The minimum time headway and distance headway throughout simulation scenarios indicate how closely the controller approaches safety limits. Safety distance constraint violations are tracked both in the actual simulation trajectory and within MPC predictions, providing insight into the controller's predictive accuracy and conservative margin selection. The magnitude of friction constraint violations indicate whether uncertainty modeling provides adequate safety margins.

Comfort Metrics quantify passenger experience through acceleration and jerk characteristics. The root mean square and maximum absolute values of acceleration and jerk provide measures of ride quality, while the frequency of abrupt changes indicates the smoothness of control actions. The rate of change in acceleration commands reflects the controller's ability to provide gradual transitions when adapting to new road surface classifications or lead vehicle behaviours.

Robustness Metrics assess the controller's ability to handle uncertainty and maintain feasibility under challenging conditions. The utilisation of slack variables in the soft-constrained MPC formulation serves as a direct indicator of constraint satisfaction difficulty, where increased slack usage signals scenarios where the nominal constraints cannot be satisfied without performance degradation. The comparison between deterministic and probabilistic friction scenarios quantifies the value of the uncertainty aware approach. Prediction accuracy metrics compare the MPC's internal predictions to actual trajectories.

Performance Metrics evaluate tracking capability and efficiency. Speed tracking error with respect to desired velocity profiles indicates the controller's ability to maintain performance objectives while satisfying safety constraints. The adherence to curvature based velocity limits shows proper integration of geometric road constraints with friction limitations. Solver convergence statistics, including iteration counts and computation times, provide essential feedback for real-time implementation feasibility.

Uncertainty Handling Metrics specifically assess the probabilistic aspects of the system. Beta-distributed friction realizations generated in simulation are compared against the predicted confidence intervals to validate the uncertainty modeling. In addition, the frequency of surface classification changes and the controller's adaptation responsiveness indicate the practical effectiveness of the camera-based perception system integration.

3.3 Use Cases

The validation framework comprises eleven systematically designed use cases that comprehensively evaluate different aspects of the friction-aware ACC with integrated road geometry adaptation. Each use case targets specific controller capabilities under realistic driving conditions, enabling isolated analysis of individual system components while building toward comprehensive system validation. Detailed Common Fixed Parameters Specifications are provided in Appendix A, Use Case Parameters as in Appendix B and for the Additional Scenarios as shown in Appendix C.

UC1: Baseline ACC with Nominal Friction Uncertainty

Establishes baseline ACC performance with standard uncertainty bounds ($\mu_{unc} = [0.1, 0.3]$) in a nominal following scenario. The ego vehicle (50 km/h) follows a constant-speed lead vehicle (70 km/h) through the complete three-segment friction profile ($0.8 \rightarrow 0.3 \rightarrow 0.8$ μ) on a straight road. Uses deterministic friction modeling to isolate controller tracking performance from stochastic effects. Critical for establishing reference performance metrics and validating basic ACC functionality. (Parameters: Table B.7)

UC2: Robustness Test with Random Lead Acceleration

Tests ACC robustness under unpredictable lead vehicle behaviour. Both vehicles start at 100 km/h with 50m separation while the lead vehicle applies bounded random accelerations (± 3 m/s^2) on straight road. Employs stochastic friction modeling with Beta distribution sampling to evaluate safety margin adaptation under combined lead vehicle unpredictability and surface uncertainty. Validates constraint satisfaction during dynamic following scenarios rather than emergency braking events. (Parameters: Table B.8)

UC3: High Uncertainty Friction Transition

Evaluates controller conservatism with significantly increased uncertainty bounds ($\mu_{unc} = [0.1, 0.7]$), simulating degraded camera confidence due to environmental conditions. Starting at 400m positions the ego vehicle to experience major friction transitions during the prediction horizon on straight road. Tests how elevated uncertainty propagates through MPC predictions and affects safety margin selection with constant-speed lead vehicle. (Parameters: Table B.9)

UC4: Curve Navigation at High Speed

Validates proactive speed reduction for safe curve navigation by initialising the ego vehicle at 140 km/h approaching a curved section ($\kappa_{max} = 0.04$ rad/m, 25m radius) starting at 600m. With the lead vehicle traveling at constant 90 km/h, the scenario tests geometric constraint handling and the controller's ability to anticipate

curvature-limited speeds while maintaining following objectives. Uses deterministic friction to isolate curvature effects from stochastic uncertainty. (Parameters: Table B.10)

UC5: Combined Following and Curve Challenge

Comprehensive validation scenario combining lead vehicle following, curve navigation ($\kappa_{max} = 0.04$ rad/m), and friction transitions with elevated uncertainty bounds ($\mu_{unc} = [0.15, 0.4]$). Starting at 700m with random lead acceleration (± 3 m/s^2) creates simultaneous demands on all constraint types. Tests the controller's ability to prioritise and satisfy competing objectives while maintaining safety margins across multiple operational domains including geometric, friction, and following constraints. (Parameters: Table B.11)

UC6: Low-Speed Convergence Test

Initialises ego vehicle at low speed (30 km/h) with large separation (100m) from a faster constant-speed lead vehicle (70 km/h) on straight road. Tests the ACC's acceleration strategy and speed matching behaviour under stochastic friction constraints during catch-up maneuvers. Uses Taylor expansion discretisation to evaluate numerical accuracy at low speeds. Validates convergence behaviour and constraint handling during sustained acceleration phases. (Parameters: Table B.12)

UC7: Close-Following Robustness Scenario

Starts with minimal safe separation (30m) between vehicles traveling at 80 km/h while the lead vehicle applies random accelerations (± 3 m/s^2) on straight road with stochastic friction. Tests the ACC's ability to maintain safety when initial margins are already constrained and lead vehicle behaviour is unpredictable. Critical for validating terminal constraint effectiveness and emergency response capabilities under realistic following distances. (Parameters: Table B.13)

UC8: Stochastic vs Deterministic Comparison

Executes identical driving scenario (70 km/h constant-speed following on straight road) in both deterministic and stochastic friction modes with fixed random seed for reproducibility. Directly quantifies the impact of explicit uncertainty modeling on ACC performance, safety metrics, and constraint satisfaction. Provides comparative analysis of conservative vs. mean-value friction handling strategies. (Parameters: Table B.14)

UC9: Extreme Initial Speed Differential

Tests ACC robustness with extreme initial conditions: ego vehicle at 150 km/h approaching a much slower constant-speed lead vehicle (50 km/h) with 120m initial separation on straight road. Challenges the prediction horizon adequacy, constraint feasibility, and deceleration capabilities under stochastic friction conditions. Validates controller performance under aggressive deceleration requirements. (Parameters: Table B.15)

UC10: Minimal Uncertainty Validation

Uses very low uncertainty bounds ($\mu_{unc} = [0.05, 0.1]$) to test ACC behaviour with high-confidence friction estimates, representing ideal camera conditions and excellent road surface classification. Starting at 400m with matched speeds (90 km/h) on straight road, this case validates that the controller can operate efficiently without excessive conservatism when uncertainty is minimal. (Parameters: Table B.16)

3.4 Additional Simulation Scenarios

This section defines additional scenarios designed to explore specific aspects of the controller's behaviour that were not fully covered in the main use cases. These scenarios investigate controller performance under different reference speed settings, varied friction profiles and challenging combinations of road geometry and surface conditions.

AS1: Driver-Imposed Speed Limitation

Evaluates the controller's ability to track a driver-specified reference speed (80 km/h) that is lower than the lead vehicle's speed (100 km/h). The ego vehicle starts at 60 km/h with 80m separation on a straight road with standard friction profile. This tests whether the controller appropriately prioritises the driver's comfort preference over unnecessary speed matching when safety permits, using deterministic friction to isolate speed tracking behaviour. (Parameters: Table C.1)

AS2: Progressive Friction Degradation

Tests the controller's adaptation to continuously worsening conditions through a progressive degradation from dry asphalt ($\mu = 0.8$) to wet conditions ($\mu = 0.5$) and finally to icy surface ($\mu = 0.2$), unlike the standard dry-wet-dry transition. The ego vehicle starts at 80 km/h following a constant-speed lead vehicle through transitions at 400m and 700m, evaluating performance without the recovery phase present in standard profiles. (Parameters: Table C.2)

AS3: Curve Coincident with Low Friction

Investigates controller response when geometric and friction constraints become active simultaneously by positioning a moderate curve with a radius around 20m to coincide with the low-friction region. The ego vehicle approaches at 100 km/h from position 300m, requiring significant speed adaptation for both curvature and reduced traction constraints, with elevated uncertainty bounds ($\mu_{unc} = [0.15, 0.4]$) to increase challenge. (Parameters: Table C.3)

AS4: Extended Low-Friction Region

Evaluates operation on slippery surfaces by extending the low-friction region from 400m to 800m which is longer than standard scenarios. Testing whether appropriate

safety margins are maintained over extended distances rather than just managing brief transitions, with stochastic friction modeling. (Parameters: Table C.4)

AS5: High-Speed Curve Entry with Mid-Curve Friction Change

Presents a demanding scenario combining high initial speed (120 km/h) with a moderate curve ($\kappa_{max} = 0.025$ rad/m) that transitions from dry to wet conditions at the apex (700m). Starting from 400m with 150m lead separation, this tests the controller's ability to handle sudden traction loss during lateral load conditions, representing critical real-world stability challenges. (Parameters: Table C.6)

AS6: Ice Patch on a Curve

Creates a worst-case cornering event by overlaying an abrupt ice patch ($\mu = 0.1$) at the apex of a tight curve ($\kappa_{\max} = 0.05$ rad/m, radius = approx. 20 m). The ego vehicle enters the bend at 90 km/h with a 75 m gap to a constant-speed lead vehicle. The scenario tests whether the controller can simultaneously moderate speed for curvature and react instantly to a sudden traction loss halfway through the curve, under stochastic friction modelling. (Parameters: Table C.6)

Note: All additional scenarios use the standard MPC controller parameters as per Table B.1, vehicle physical constraints in Table B.2 and simulation parameters in Table B.4 as defined for the main use cases. Only the specific parameters listed above differ from the standard configuration.

4

Uncertainty-Aware Robust Predictive Controller

Conventional ACC systems either assume fixed nominal friction, which can be dangerously optimistic under adverse conditions, or worst-case friction, which ensures safety but leads to overly conservative behaviour. Nominal assumptions may cause control commands to exceed available traction on wet or icy roads, risking instability, while worst-case planning results in large gaps, harsh braking, and sluggish acceleration that compromise comfort and traffic flow. Both approaches fail to reason about uncertainty in a structured way.

Modern perception systems, such as camera-based road surface classification, provide probabilistic estimates with confidence levels that increase as the vehicle approaches. This preview information offers an opportunity to balance safety and performance, but conventional controllers lack the mathematical framework to exploit it. The challenge is further complicated by multiple uncertainties: friction limits affect both longitudinal and lateral dynamics during curve navigation, while lead-vehicle behaviour introduces additional variability. Traditional methods treat these separately or through conservative assumptions, missing more nuanced trade-offs.

This chapter introduces a unified, uncertainty-aware predictive controller that integrates safe following distances, curvature-based limits, and friction-dependent acceleration bounds into a single optimisation problem. Road-surface uncertainty is represented by position-dependent confidence bounds; in simulation, Beta-distributed samples capture stochastic realizations, while MPC relies on bounded intervals for tractable real-time optimisation.

The objective function explicitly balances comfort, tracking, and safety. Penalising control effort improves ride comfort but reduces responsiveness, while relaxing comfort penalties enables faster reactions at the cost of sharper maneuvers. A time-domain formulation avoids spatial transformations, handles dynamic scenarios naturally, and supports both deterministic (mean friction) and stochastic (sampled friction) cases.

Feasibility is maintained through slack variables that prioritise safety-critical constraints while allowing controlled relaxation of comfort or performance requirements.

This design ensures the controller remains responsive, robust, and efficient under conflicting demands.

Overall, the framework provides a rigorous yet practical approach to friction-aware longitudinal control, capable of handling real-world uncertainty. The following sections detail its formulation, constraint structure, and implementation.

4.1 Mathematical Formulation

4.1.1 State-Space Model

The ego vehicle is modeled as a point-mass system, capturing its longitudinal dynamics through a concise yet effective set of ODEs as per the following:

$$\dot{s}(t) = v(t), \tag{4.1}$$
$$\dot{v}(t) = a(t), \tag{4.2}$$
$$\dot{a}(t) = j(t). \tag{4.3}$$

where $s(t)$ is the longitudinal position, $v(t)$ the velocity, $a(t)$ the acceleration, and $j(t)$ the jerk serving as the control input. The state vector and the control input are thus:

$$x(t) = [s(t),\, v(t),\, a(t)]^\top, \quad u(t) = j(t) \tag{4.4}$$

Since the vehicle dynamics represent a triple integrator, they can be expressed in the standard linear state-space form:

$$\dot{x}(t) = Ax(t) + Bu(t) \tag{4.5}$$

where

$$A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \tag{4.6}$$

For generality throughout this thesis, the dynamics are represented using the compact notation:

$$\dot{x}(t) = f(x(t), u(t)) \tag{4.7}$$

where $f(\cdot)$ denotes the system dynamics. While the vehicle dynamics themselves are linear, this general representation allows the framework to accommodate nonlinear vehicle models through the implemented RK4 discretisation method. The nonlinearities in the overall control problem arise from the state and input constraints, particularly those involving friction and curvature limitations, which will be introduced in subsequent sections.

For numerical implementation within the time-domain MPC, the above continuous-time dynamics must be discretised. In this thesis, two discretisation methods have been implemented and evaluated, providing flexibility in terms of accuracy and computational load as per:

4.1.1.1 Runge-Kutta 4 (RK4)

The classical fourth-order Runge-Kutta (RK4) method provides an accurate numerical integration scheme. At each step, the discrete-time update is given by:

$$x_{k+1} = f_d^{\mathrm{RK4}}(x_k, u_k), \tag{4.8}$$

where the intermediate slopes are computed as:

$$k_1 = f(x_k, u_k), \tag{4.9}$$

$$k_2 = f\left(x_k + \tfrac{dt}{2} k_1, u_k\right), \tag{4.10}$$

$$k_3 = f\left(x_k + \tfrac{dt}{2} k_2, u_k\right), \tag{4.11}$$

$$k_4 = f\left(x_k + dt\, k_3, u_k\right), \tag{4.12}$$

and the state update becomes:

$$x_{k+1} = x_k + \tfrac{dt}{6}(k_1 + 2k_2 + 2k_3 + k_4). \tag{4.13}$$

4.1.1.2 Closed-Form Discretisation (Zero-Order Hold)

Since the system is linear and the jerk input is assumed constant within each sampling interval, the ODEs can be integrated exactly to obtain the closed-form discrete-time dynamics under a Zero-Order Hold assumption. This leads to the following update equations:

$$x_{k+1} = f_d^{\mathrm{ZOH}}(x_k, u_k), \tag{4.14}$$

where the state update is obtained by exact integration of the ODEs:

$$s_{k+1} = s_k + v_k\, dt + \frac{1}{2} a_k\, dt^2 + \frac{1}{6} j_k\, dt^3 \tag{4.15}$$

$$v_{k+1} = v_k + a_k\, dt + \frac{1}{2} j_k\, dt^2 \tag{4.16}$$

$$a_{k+1} = a_k + j_k\, dt \tag{4.17}$$

Both discretisation methods can represent the continuous-time dynamics. RK4 is a general purpose numerical integrator that is accurate for nonlinear systems and larger time steps, but it is computationally more expensive. The closed-form ZOH discretisation, by contrast, is exact for linear model under constant jerk and is computationally efficient, making it well-suited for real-time predictive control.

This simplified yet comprehensive dynamic formulation captures the essential longitudinal vehicle behaviour required for real-time implementation of predictive control.

4.1.2 Time Domain Formulation

The RNMPC formulation operates in discrete time with a fixed sampling interval dt, predicting the ego vehicle's future state trajectory over a horizon of N steps. The controller solves an optimal control problem at each time instant t_k, generating a jerk input sequence $\{j_k, j_{k+1}, \ldots, j_{k+N-1}\}$ so that the ego trajectory $\{x_k, x_{k+1}, \ldots, x_{k+N}\}$ satisfies dynamics and constraints while minimizing a cost function. With $x_k \in \mathbb{R}^3$ showing the state at time $t_k = k \cdot dt$ where:

$$x_k = \begin{bmatrix} s_k \\ v_k \\ a_k \end{bmatrix} \in \mathbb{R}^3, \qquad u_k = j_k,$$

the discrete-time dynamics are

$$x_{k+1} = f_d(x_k, u_k), \tag{4.18}$$

where $f_d(\cdot)$ is obtained via RK4 or the exact ZOH discretisation of the linear ODEs (see Sec. 4.1.1).

At each t_k, the controller solves:

$$\begin{aligned}
\min_{\{x_{k+i}, u_{k+i}, s_{k+i}\}} \quad & \sum_{i=0}^{N-1} \ell(x_{k+i}, u_{k+i}, s_{k+i}) + \ell_f(x_{k+N}) && (4.19)\\
\text{s.t.} \quad & x_{k+i+1} = f_d(x_{k+i}, u_{k+i}), \quad i = 0, \ldots, N-1 && (4.20)\\
& g(x_{k+i}, u_{k+i}, w_{k+i}) \le 0, \quad \forall\, w_{k+i} \in \mathcal{W}(i) && (4.21)\\
& x_{k+i} \in \mathcal{X}(i), \quad u_{k+i} \in \mathcal{U}(i), \quad s_{k+i} \in \mathbb{R}^m_+ && (4.22)\\
& x_k = \hat{x}_k, && (4.23)\\
& x_{k+N} \in \mathcal{X}_f && (4.24)
\end{aligned}$$

Here, w_{k+i} represents bounded disturbances/uncertainties, which in this thesis enter only through the constraints. The road curvature $\kappa(s)$ is treated as a deterministic previewed profile and is not included in $\mathcal{W}(i)$. The time-varying constraints, consisting of both linear bounds and nonlinear inequalities, are collected in $g(x, u, w) \le 0$. Examples include:

$$\begin{aligned}
& v_{\min} \le v_{k+i} - s^{(2)}_{k+i} \le v_{\max}, && (4.25)\\
& v^2_{k+i}\, \kappa(s_{k+i}) \;\le\; \mu(s_{k+i})\, g + s^{(2)}_{k+i}, && (4.26)\\
& -a_{\max} \le a_{k+i} + s^{(3)}_{k+i} \le a_{\max}, && (4.27)\\
& -\mu(s_{k+i})\, g \le a_{k+i} + s^{(3)}_{k+i} \le \mu(s_{k+i})\, g, && (4.28)\\
& -a_{c,\max} \le a_{k+i} + s^{(4)}_{k+i} \le a_{c,\max}, && (4.29)
\end{aligned}$$

together with the safe-distance constraint relative to a conservatively predicted lead vehicle:

$$s_L^{\text{pred}}(t_{k+i}) - s_{k+i} \;\geq\; \delta s_{\min}(t_{k+i}) - s_{k+i}^{(1)}. \tag{4.30}$$

The slack variables relax the constraints in a structured way: (i) $s^{(1)}$: safety distance , (ii) $s^{(2)}$: velocity limits, (iii) $s^{(3)}$: acceleration limits, (iv) $s^{(4)}$: comfort acceleration limits.

This formulation directly incorporates previewed road and traffic information, enforces robust feasibility under bounded uncertainty, and remains compatible with standard nonlinear solvers such as IPOPT.

4.1.3 Friction and Curvature Prediction

The RNMPC framework relies on preview information of road friction and curvature to proactively adapt control actions. The controller incorporates spatially-varying friction $\mu(s)$ and curvature $\kappa(s)$ profiles as position-dependent parameters within the optimisation problem. These are modeled using smooth, differentiable functions to reflect realistic transitions between road segments, where friction captures surface changes and curvature encodes upcoming geometric features.

The key innovation lies in treating friction as a bounded uncertain parameter reflecting the confidence of camera-based surface classification, while curvature is treated as a deterministic previewed profile. Friction bounds are evaluated at predicted positions s_{k+i} within the MPC horizon. Conservative lower-bound friction values are used for ego vehicle constraints, while lead vehicle behaviour is predicted using upper bounds, ensuring robust constraint satisfaction and adequate safety margins under all plausible scenarios.

4.1.4 Camera-Based Probabilistic Road Parameter Estimation

The foundation of uncertainty-aware vehicle control lies in accurately modeling road parameters while explicitly accounting for the inherent uncertainty in their estimation. This section presents a comprehensive framework that transforms camera-based road surface classifications into mathematical representations suitable for real-time MPC implementation.

Considering the challenge facing any camera-based perception system. When a camera classifies a road segment as "wet pavement", this classification comes with associated uncertainty that depends on factors such as lighting conditions, distance to the classified region and algorithm confidence levels. Traditional control approaches either ignore this uncertainty entirely or handle it through overly conservative worst-case assumptions. Our approach instead creates a mathematical bridge between realistic sensor capabilities and sophisticated control algorithms.

4.1.4.1 Double Sigmoid Spatial Modeling Framework

Road parameters such as friction coefficient $\mu(s)$ and curvature $\kappa(s)$ exhibit spatial variation that must be captured through smooth, differentiable functions suitable for optimisation-based control. Rather than using discrete segments or piecewise linear approximations that can cause numerical difficulties in the MPC solver, a unified double sigmoid function that creates smooth transitions between different road conditions is employed. The double sigmoid function is defined as:

$$\sigma(s;p) = p_2 + \frac{p_3 - p_2}{1 + e^{-p_1(s-p_5)}} + \frac{p_4 - p_3}{1 + e^{-p_1(s-p_6)}} \tag{4.31}$$

This mathematical structure enables modeling of complex road scenarios through six interpretable parameters contained in the vector $p = [p_1, p_2, p_3, p_4, p_5, p_6]^T$. This parameterisation directly connects to realistic driving scenarios. For example, a route might begin on dry asphalt ($p_2 = 0.8$), transition through an icy patch ($p_3 = 0.3$), and return to dry conditions ($p_4 = 0.8$). The positions p_5 and p_6 would correspond to the camera-detected boundaries of the icy region, while p_1 would reflect how quickly road conditions change at these boundaries.

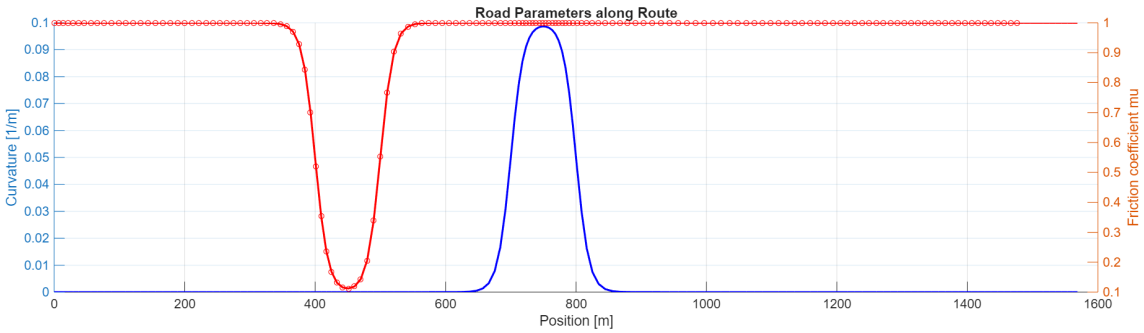

Figure 4.1: Double sigmoid profiles showing smooth transitions in road friction $\mu(s)$ and curvature $\kappa(s)$ along the route. The friction profile shows a transition from dry conditions through a low-friction region (representing icy area) and back to dry conditions, while the curvature profile indicates an upcoming curve section.

For friction coefficient estimation, the mean friction profile becomes:

$$\mu_{mean}(s) = \sigma(s; p_{fric}) \tag{4.32}$$

where p_{fric} represents the camera-classified friction parameters. Similarly, for road curvature:

$$\kappa_{mean}(s) = \sigma(s; p_{curve}) \tag{4.33}$$

The beauty of this unified approach lies in its flexibility and numerical properties. The double sigmoid function is infinitely differentiable, ensuring that optimisation algorithms can compute reliable gradients. Moreover, the same mathematical structure can represent diverse road scenarios simply by adjusting the parameter values, making the framework adaptable to different geographical regions and road types.

4.1.4.2 Position-Dependent Uncertainty Modeling

The critical innovation that distinguishes this framework from traditional approaches lies in how uncertainty is modeled as a function of position relative to the ego vehicle. This approach directly reflects the realistic behaviour of camera-based classification systems, where confidence typically increases as the vehicle approaches a classified road segment. The position-dependent uncertainty bound is mathematically expressed as:

$$\mu_{bound}(s_{rel}) = \mu_{unc,current} + \frac{s_{rel}}{s_{unc,max}} \cdot (\mu_{unc,far} - \mu_{unc,current}) \tag{4.34}$$

Here, s_{rel} denotes the distance ahead of the ego vehicle, i.e., the relative position along the MPC prediction horizon. This distinction is important since the mean profiles $\mu(s)$ and $\kappa(s)$ are defined over absolute road position s, while the uncertainty bounds are expressed relative to the ego vehicle. The linear interpolation captures the fundamental characteristic of camera-based perception: nearby classifications are more reliable than distant ones. At the current vehicle position, the uncertainty might be quite low (e.g., ± 0.1 for friction coefficient), reflecting high confidence in immediate road conditions. However, at the far end of the prediction horizon, uncertainty could be significantly higher (e.g., ± 0.3) acknowledging the inherent limitations in classifying distant road segments.

The resulting friction bounds that constrain the optimisation problem are:

$$\mu_{low}(s) = \max(0.1, \mu_{mean}(s) - \mu_{bound}(s - s_{current})) \tag{4.35}$$

$$\mu_{high}(s) = \min(1.1, \mu_{mean}(s) + \mu_{bound}(s - s_{current})) \tag{4.36}$$

The maximum and minimum operations ensure that friction bounds remain physically realistic. Friction coefficients below 0.1 would represent extremely icy conditions that are rare in practice, while values above 1.1 exceed the capabilities of standard tyre-road interfaces under normal conditions.

4.1.4.3 Operational Modes for Comprehensive Evaluation

The framework supports two distinct operational modes that enable thorough evaluation of controller performance under different uncertainty assumptions. This dual-mode capability is essential for validating the robustness of the control approach and understanding the value of probabilistic uncertainty handling.

Deterministic Mode: In this baseline configuration, the actual friction experienced during simulation follows the mean profile exactly:

$$\mu_{actual}(s) = \mu_{mean}(s) \tag{4.37}$$

This mode serves multiple important purposes. First, it provides a baseline for performance comparison by eliminating stochastic variation that could confound

analysis results. Second, it allows assessment of controller behaviour when predictions perfectly match reality, revealing the fundamental capabilities of the control algorithm without uncertainty complications. Third, it enables verification that the mathematical framework correctly implements the intended friction profiles.

Stochastic Mode: This advanced configuration introduces realistic uncertainty by randomly sampling actual friction values from probability distributions that respect the uncertainty bounds:

$$\mu_{actual}(s) \sim \text{Beta}(\alpha, \beta) \text{ scaled to } [\mu_{low}(s), \mu_{high}(s)] \tag{4.38}$$

The Beta distribution provides an ideal mathematical tool for this application because it is naturally bounded to a finite interval and offers flexible shape control through its parameters. The distribution parameters are computed to achieve the desired mean value while respecting the uncertainty bounds:

$$M = \frac{\mu_{mean}(s) - \mu_{low}(s)}{\mu_{high}(s) - \mu_{low}(s)} \tag{4.39}$$

$$\alpha = \rho \cdot M, \quad \beta = \rho \cdot (1 - M) \tag{4.40}$$

The peak factor ρ controls the concentration of the distribution around the mean value. Higher peak factors create distributions that cluster more tightly around the mean, representing scenarios where camera classifications are generally accurate but subject to small variations. Lower peak factors create broader distributions, representing scenarios with higher classification uncertainty.To reflect increasing uncertainty with preview distance, the uncertainty bounds used to generate the Beta distribution are implemented as a linear function of position s along the prediction horizon. This ensures that the prediction uncertainty grows with distance. A visual representation of this can be seen in the following figure:

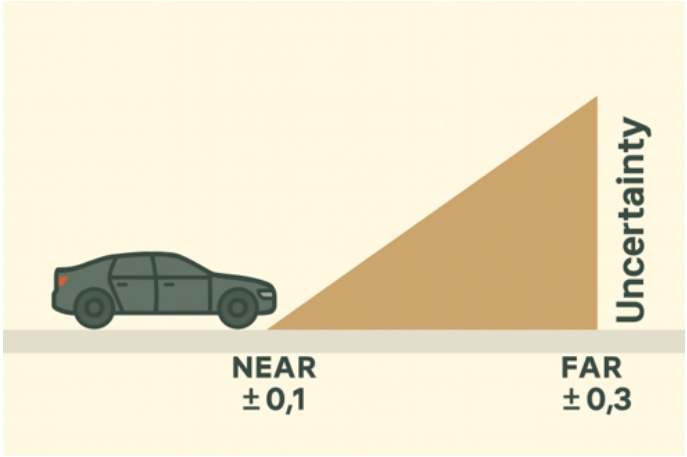

Figure 4.2: Uncertainty increase through the camera's preview distance, reflecting the uncertainty increase in the prediction horizon

This stochastic mode enables comprehensive robustness testing by subjecting the controller to realistic friction variations while maintaining statistical consistency with the uncertainty bounds. The random sampling approach ensures that each simulation run explores different uncertainty realisations, providing statistical confidence in controller performance assessments.

4.1.4.4 Conservative Constraint Application Strategy

The probabilistic framework employs a sophisticated strategy for applying uncertainty bounds within the MPC constraints that balances safety with performance. This approach recognises that different vehicles in the traffic scenario have different roles and should be modeled with appropriate conservatism levels.

For Ego Vehicle Constraints: The controller employs pessimistic friction estimates to ensure safety margins are maintained:

$$\mu_{ego}(s) = \mu_{low}(s) \quad \text{(for acceleration limits)} \tag{4.41}$$

$$\kappa_{ego}(s) = \kappa_{high}(s) \quad \text{(for speed limits)} \tag{4.42}$$

This conservative approach ensures that the ego vehicle never commands accelerations or speeds that exceed the available traction, even when actual conditions fall at the lower bound of the estimated friction range. The underlying philosophy is that it is better to be slightly conservative in ego vehicle control than to risk stability or safety violations.

For Lead Vehicle Prediction: The controller uses optimistic friction estimates to model the most aggressive possible lead vehicle behaviour:

$$a_{L,max}(s) = -\min(\mu_{high}(s) \cdot g, w_{a,lead}) \tag{4.43}$$

This approach assumes that the lead vehicle has access to the best possible traction conditions and may brake as aggressively as physically possible. By planning for this worst-case lead vehicle behaviour, the ego vehicle maintains safety even when the lead vehicle exploits superior traction or driver aggressiveness.

4.1.5 Lead Vehicle Prediction and Uncertainty Modeling

The lead vehicle's future trajectory represents a critical source of uncertainty that significantly impacts ego vehicle planning decisions. Unlike road parameters that can be estimated through camera based classification, lead vehicle behaviour involves human decision-making processes that are inherently unpredictable. The framework addresses this challenge through a conservative prediction model that accounts for both kinematic constraints and behavioural uncertainty.

The lead vehicle state evolution is modeled using the same discrete-time dynamics as the ego vehicle:

$$x_L(k) = [s_L(k), v_L(k), a_L(k)]^T \tag{4.44}$$

$$x_L(k+1) = f_d(x_L(k), 0) \tag{4.45}$$

The assumption of zero jerk for lead vehicle prediction reflects the practical limitation that future lead vehicle control inputs cannot be predicted reliably. Instead, the model assumes that current acceleration levels will be maintained over short prediction horizons, which provides a reasonable approximation for safety-critical planning purposes.

The predicted lead vehicle acceleration incorporates both physical constraints and uncertainty bounds:

$$a_L(k) = -\min\left(\mu_{high}(s_L(k)) \cdot g, w_{a,lead}\right) \tag{4.46}$$

This formulation ensures several important properties. First, the lead vehicle's predicted deceleration never exceeds the physical limits imposed by available friction, preventing unrealistic assumptions about braking capability. Second, the acceleration is bounded by the parameter $w_{a,lead}$, which represents the maximum expected acceleration or deceleration that a human driver might command under normal circumstances. Third, the use of $\mu_{high}(s)$ assumes that the lead vehicle has access to the best possible traction conditions, creating a conservative safety margin for ego vehicle planning.

The behavioural uncertainty parameter $w_{a,lead}$ requires careful calibration based on empirical driving data. Typical values range from $2 - 4m/s^2$, reflecting the range of accelerations that human drivers commonly employ during normal driving situations. This parameter enables the framework to distinguish between normal traffic scenarios and emergency situations where more aggressive maneuvers might be necessary.

During stochastic simulation modes, lead vehicle behaviour can be further randomised to test controller robustness:

$$a_{L,actual}(k) = \xi_k \cdot \min\left(\mu_{actual}(s_L(k)) \cdot g, w_{a,lead}\right) \tag{4.47}$$

where $\xi_k \sim \mathcal{U}[-1, 1]$ represents random acceleration commands that reflect the unpredictability of human driving behaviour. This stochastic lead vehicle model enables comprehensive robustness evaluation under scenarios that include both road surface uncertainty and behavioural unpredictability.

4.1.5.1 Integration with MPC Constraint Structure

The probabilistic road parameter framework integrates seamlessly with the MPC constraint structure through careful mathematical formulation that preserves optimisation efficiency while incorporating uncertainty bounds. At each prediction step $i = 0, \ldots, N$, the controller evaluates road parameters at the predicted ego vehicle position:

$$\mu_{k+i} = \mu_{low}(s_{k+i}) \quad \text{(for ego vehicle constraints)} \tag{4.48}$$

$$\kappa_{k+i} = \kappa_{high}(s_{k+i}) \quad \text{(for curvature constraints)} \tag{4.49}$$

These position-dependent values directly enter the velocity and acceleration constraints:

$$v_{k+i}^2 \cdot \kappa_{k+i} - s_{k+i}^{(2)} \leq \mu_{k+i} \cdot g \tag{4.50}$$

$$-\mu_{k+i} \cdot g \leq a_{k+i} + s_{k+i}^{(3)} \leq \mu_{k+i} \cdot g \tag{4.51}$$

The inclusion of slack variables $s_{k+i}^{(2)}$ and $s_{k+i}^{(3)}$ provides controlled constraint relaxation when conflicting requirements arise. This formulation ensures that the optimisation problem remains feasible even when aggressive maneuvers are required for safety, while strongly discouraging unnecessary constraint violations through appropriate penalty structures in the cost function.

By embedding continuous, smooth friction and curvature predictions directly into the optimisation constraints, the RNMPC framework avoids numerical difficulties associated with discontinuous constraint changes while maintaining robust performance under significant uncertainty. This mathematical structure enables real-time implementation while preserving the sophisticated uncertainty handling capabilities that distinguish this approach from traditional vehicle control methods.

4.2 Constraints Formulation

The RNMPC framework employs a sophisticated constraint structure that balances safety requirements with computational efficiency. Understanding how these constraints interact with the uncertainty framework is crucial for appreciating the controller's robustness properties.

4.2.1 Slack Variable Framework

The implementation utilises slack variables to maintain optimisation feasibility when physical constraints conflict. This approach is essential because rigid constraint enforcement could lead to infeasible optimisation problems, particularly when the vehicle encounters unexpected scenarios requiring aggressive maneuvers.

At each prediction step k, the slack vector contains four components:

$$s_k = [s_k^{(1)}, s_k^{(2)}, s_k^{(3)}, s_k^{(4)}]^T \in \mathbb{R}_{\geq 0}^4 \tag{4.52}$$

4.2.1.1 Safety Distance Constraint Implementation

The safety distance constraint demonstrates the unique formulation used in the implementation:

$$0 \leq s_{ego}(k) - s_k^{(1)} \leq s_L^{pred}(k) - \delta s_{min} \tag{4.53}$$

This double-sided inequality can be decomposed into two constraints:

$$s_{ego}(k) \geq s_k^{(1)} \quad \text{(position non-negativity with slack)} \tag{4.54}$$

$$s_{ego}(k) \leq s_L^{pred}(k) - \delta s_{min} + s_k^{(1)} \quad \text{(safety distance)} \tag{4.55}$$

The key insight is that the slack variable $s_k^{(1)}$ appears on both sides, allowing controlled violation of the minimum distance requirement when absolutely necessary while ensuring the ego vehicle position remains physically meaningful.

4.2.1.2 Velocity Constraint Structure

The velocity constraints incorporate both operational limits and physics-based curvature restrictions:

$$v_{min} \leq v_k - s_k^{(2)} \leq v_{max} \tag{4.56}$$

$$v_k^2 \cdot \kappa(s_k) \leq \mu_{low}(s_k) \cdot g + s_k^{(2)} \tag{4.57}$$

Notice that the slack variable $s_k^{(2)}$ is subtracted from the velocity in the bound constraints but added in the curvature constraint. This asymmetric formulation reflects the different physical meanings: the first constraint allows temporary speed limit violations, while the second permits slightly higher speeds through curves when necessary for safety.

4.2.1.3 Acceleration Constraint Hierarchy

The implementation enforces a hierarchy of acceleration constraints:

$$\text{Physical limits:} \quad -a_{max} \leq a_k + s_k^{(3)} \leq a_{max} \tag{4.58}$$

$$\text{Friction limits:} \quad -\mu_{low}(s_k) \cdot g \leq a_k + s_k^{(3)} \leq \mu_{low}(s_k) \cdot g \tag{4.59}$$

$$\text{Comfort limits:} \quad -a_{c,max} \leq a_k + s_k^{(4)} \leq a_{c,max} \tag{4.60}$$

The use of separate slack variables for physical/friction constraints ($s_k^{(3)}$) versus comfort constraints ($s_k^{(4)}$) enables the controller to distinguish between safety-critical and comfort-oriented requirements. This distinction becomes crucial when aggressive maneuvers are necessary to avoid collisions.

4.2.2 Friction and Curvature Integration

The constraint formulation directly incorporates the position-dependent friction and curvature models, creating a dynamic constraint structure that adapts to road conditions.

4.2.2.1 Conservative Friction Application

For ego vehicle constraints, the controller uses the lower bound of the friction estimate:

$$\mu_{ego}(s) = \mu_{low}(s) = \max\left(0.1, \mu_{mean}(s) - \mu_{bound}(s - s_{current})\right) \tag{4.61}$$

This conservative approach ensures that the ego vehicle never plans maneuvers that would exceed available traction, even under worst-case friction realisations within the uncertainty bounds.

4.2.2.2 Curvature-Based Speed Limitation

The curvature constraint prevents excessive speeds through curves:

$$v^2 \cdot \kappa(s) \leq \mu(s) \cdot g + s^{(2)} \tag{4.62}$$

This reformulation avoids the numerical issues associated with the standard form $v \leq \sqrt{\mu g/\kappa}$ when $\kappa \approx 0$. The implementation further handles near-zero curvature by effectively disabling this constraint when $\kappa < 10^{-5}$, preventing unnecessary speed restrictions on straight road sections.

4.2.3 Lead Vehicle Prediction and Safety Margins

The safety constraint formulation requires careful prediction of lead vehicle behaviour under uncertainty. The implementation employs a conservative yet realistic approach that balances safety with performance.

4.2.3.1 Worst-Case Lead Vehicle Dynamics

Within the MPC prediction, the lead vehicle acceleration is computed as:

$$a_L^{pred}(k) = -\min\left(\mu_{high}(s_L(k)) \cdot g, w_{a,lead}\right) \tag{4.63}$$

$$v_L^{pred}(k+1) = v_L^{pred}(k) + a_L^{pred}(k) \cdot dt \tag{4.64}$$

$$s_L^{pred}(k+1) = s_L^{pred}(k) + v_L^{pred}(k) \cdot dt + \frac{1}{2} a_L^{pred}(k) \cdot dt^2 \tag{4.65}$$

The implementation includes important safeguards:

$$v_L^{pred}(k+1) = \max(v_L^{pred}(k+1), v_{L,min}) \quad \text{(velocity floor)} \tag{4.66}$$

$$s_L^{pred}(k+1) = \max(s_L^{pred}(k+1), s_L^{pred}(k)) \quad \text{(monotonic position)} \tag{4.67}$$

These safeguards prevent physically unrealistic predictions while maintaining conservatism.

4.2.3.2 Safety Margin Evolution

The predictive safety margin incorporates both position and velocity dynamics:

$$\Delta s_{margin}(k) = s_L^{pred}(k) - s_{ego}(k) - \delta s_{min} \tag{4.68}$$

The constraint $\Delta s_{margin}(k) \geq -s_k^{(1)}$ allows temporary margin reduction only when absolutely necessary, with the slack penalty ensuring minimal violation.

4.2.4 Terminal Constraint for Recursive Feasibility

The terminal constraint ensures that the optimisation problem remains recursively feasible by guaranteeing that the terminal state allows safe future evolution.

4.2.4.1 Kinematic Derivation

Starting from the kinematic relationship for safe following, the ego vehicle must be able to match the lead vehicle's speed within the available distance. The required braking distance from velocity v_{ego} to $v_{L,min}$ with deceleration a is:

$$d_{brake} = \frac{v_{ego}^2 - v_{L,min}^2}{2|a|} \tag{4.69}$$

For safety, this must not exceed the available distance:

$$d_{brake} \leq s_L - s_{ego} - \delta s_{min} \tag{4.70}$$

4.2.4.2 Implementation Form

The implementation uses a specific form that accounts for the ego vehicle's current acceleration:

$$v_{ego,N}^2 - v_{L,min}^2 + 2a_{ego,N}(s_{L,N} - s_{ego,N} - \delta s_{min}) - s_{N-1}^{(1)} \leq 0 \tag{4.71}$$

This formulation is conservative because it assumes the ego vehicle maintains its terminal acceleration $a_{ego,N}$, which may be positive (reducing braking capability). The inclusion of the previous slack variable $s_{N-1}^{(1)}$ maintains consistency with the safety distance constraints throughout the horizon.

4.2.5 Constraint Activation Analysis

Understanding when and how constraints become active provides insight into the controller's behaviour under different scenarios. The implementation includes comprehensive monitoring of constraint activation through the slack variables.

During normal highway driving with good road conditions, all slack variables remain at zero, indicating that the controller satisfies all constraints without relaxation. However, specific scenarios trigger constraint activation:

1. **Emergency braking**: When the lead vehicle brakes suddenly, $s^{(1)}$ may activate temporarily to allow closer following distances while the ego vehicle decelerates.

2. **Low friction conditions**: On icy roads, $s^{(3)}$ may activate when comfort limits demand accelerations exceeding available traction.

3. **Curve entry**: The velocity slack $s^{(2)}$ may activate briefly when entering unexpected sharp curves, allowing controlled speed reduction.

4. **Comfort violations**: The comfort slack $s^{(4)}$ activates during emergency maneuvers where passenger comfort must be sacrificed for safety.

4.3 Cost Function Design

The cost function design represents a critical aspect of the RNMPC formulation, directly influencing how the controller balances competing objectives. The quadratic cost structure enables efficient optimisation while providing intuitive tuning parameters for different performance aspects.

4.3.1 Complete Cost Formulation

The total cost over the prediction horizon consists of stage costs and a terminal cost:

$$J = \sum_{k=0}^{N-1} \ell(x_k, u_k, s_k) \cdot dt + \ell_f(x_N) \tag{4.72}$$

where the stage cost is:

$$\ell(x_k, u_k, s_k) = (x_k - x_{ref})^T Q (x_k - x_{ref}) + u_k^T R u_k + s_k^T R_s s_k \tag{4.73}$$

and the terminal cost is:

$$\ell_f(x_N) = (x_N - x_{ref})^T Q (x_N - x_{ref}) \tag{4.74}$$

The multiplication by dt ensures proper scaling of the continuous-time objectives in the discrete-time implementation.

4.3.2 State Tracking Penalties

The state penalty matrix $Q = \text{diag}(0, 0.1, 0.1)$ reveals important design choices:

$$Q_{11} = 0 \quad \text{(position not tracked)} \tag{4.75}$$
$$Q_{22} = 0.1 \quad \text{(velocity tracking)} \tag{4.76}$$
$$Q_{33} = 0.1 \quad \text{(acceleration regulation)} \tag{4.77}$$

The zero penalty on position reflects that absolute position is irrelevant for highway driving; only relative position to the lead vehicle matters. The equal penalties on

velocity and acceleration create balanced tracking of the desired cruise speed while discouraging unnecessary acceleration variations.

The reference state is set as:

$$x_{ref} = [0, v_{ref}, 0]^T \tag{4.78}$$

where $v_{ref} = 100$ km/h represents the desired cruise speed set by the ego car's driver. The zero acceleration reference encourages smooth driving when traffic conditions permit.

Control Smoothness Penalty

The relatively low value ($R = 1$) compared to the slack penalties indicates that control smoothness is sacrificed when necessary for constraint satisfaction. Smoothness can be observed throughout the simulations, indicating the abscence of the need for excessive penalisation on the control input.

4.3.3 Slack Penalty Structure

The slack penalty matrix $R_s = \text{diag}(10^3, 10^2, 10^2, 1)$ set a clear hierarchy of constraint importance:

$$r_1 = 10^3 \quad \text{(safety distance - highest priority)} \tag{4.79}$$
$$r_2 = 10^2 \quad \text{(velocity limits - high priority)} \tag{4.80}$$
$$r_3 = 10^2 \quad \text{(physical acceleration - high priority)} \tag{4.81}$$
$$r_4 = 1 \quad \text{(comfort acceleration - lowest priority)} \tag{4.82}$$

This method of penalisation ensures that lower-priority constraints relax before higher-priority ones. The implementation effectively creates a lexicographic ordering where safety distance is maintained unless doing so would make the problem infeasible.

4.3.4 Cost Function behaviour Analysis

To understand how the cost function influences controller behaviour, consider the effective trade-offs in different scenarios:

4.3.4.1 Normal Cruise Conditions

When following a steady lead vehicle on a straight road with good friction. All slack variables remain at zero (no constraint violations), the velocity tracking term $(v - v_{ref})^2 \cdot 0.1$ dominates, the control penalty $j^2 \cdot 1$ ensures smooth acceleration adjustments and finally the system naturally settles to the desired cruise speed.

4.3.4.2 Emergency Braking Scenario

When the lead vehicle brakes suddenly, the safety distance slack $s^{(1)}$ may briefly activate, its high penalty (10^3) ensures minimal violation, comfort constraints relax first (low penalty on $s^{(4)}$) and finally the controller accepts high jerk to maintain safety.

4.3.4.3 Low Friction Conditions

On icy roads with reduced traction, physical acceleration constraints tighten, the controller may need to violate comfort limits to stay within friction bounds, and the penalty structure will ensure friction limits are respected over comfort.

4.3.5 Tuning Considerations

The cost function parameters provide several tuning knobs for adjusting the controller's behaviour:

1. **Increasing** Q_{22}: More aggressive speed tracking, potentially at the cost of comfort

2. **Increasing** R: Smoother control actions, but potentially slower constraint response

3. **Adjusting** r_1/r_4 **ratio**: Changes how readily comfort is sacrificed for safety

4. **Scaling all penalties**: Affects optimiser convergence without changing relative behaviour

The current parameter values represent a balanced configuration validated through extensive simulation, prioritising safety while maintaining reasonable comfort and performance.

4.3.6 Integration with Probabilistic Framework

The cost function design complements the probabilistic constraint framework by providing smooth trade-offs between objectives. When uncertainty bounds tighten, the controller naturally becomes more conservative through the constraint structure, while the cost function ensures this conservatism manifests as smooth, predictable behaviour rather than abrupt mode switches.

The quadratic penalties create continuous gradients that guide the optimisation solver efficiently, even when navigating the complex feasible region created by position-dependent friction and curvature constraints. This smooth formulation is essential for real-time implementation, ensuring reliable convergence within the available computation time.

4.4 CasADi Implementation and Warm Starting

The practical implementation using CasADi requires careful attention to symbolic framework capabilities and integration with IPOPT optimisation. Understanding this implementation bridge is crucial for transforming theoretical mathematics into working control software.

4.4.1 Symbolic Variable Structure and Function Definition

CasADi enables the creation of symbolic variables that represent the optimisation decision variables:

$$\text{Symbolic states:} \quad X = \text{opti.variable}(3, N+1) \in \mathbb{R}^{3\times(N+1)} \tag{4.83}$$

$$\text{Symbolic controls:} \quad U = \text{opti.variable}(1, N) \in \mathbb{R}^{1\times N} \tag{4.84}$$

$$\text{Symbolic slacks:} \quad S = \text{opti.variable}(4, N) \in \mathbb{R}^{4\times N} \tag{4.85}$$

The discrete-time dynamics function is created using CasADi's symbolic framework:

$$f_{symbolic} = \text{casadi.Function}('f', \{x, u\}, \{x_{next}\}) \tag{4.86}$$

where x_{next} represents either the RK4 or the exact ZOH discretisation of the continuous time dynamics.

4.4.2 Warm Starting Strategy

To improve convergence speed and solution quality, the RNMPC employs a warm starting strategy:

1. **State trajectory initialisation**: The previous solution is shifted forward by one time step

$$X_{init}(:, 1:N) = X_{prev}(:, 2:N{+}1), \quad X_{init}(:, N{+}1) = f_d(X_{prev}(:, N{+}1), U_{prev}(:, N)) \tag{4.87}$$

2. **Control trajectory initialisation**: The control sequence is shifted with the last control duplicated

$$U_{init}(:, 1:N-1) = U_{prev}(:, 2:N), \quad U_{init}(:, N) = U_{prev}(:, N) \tag{4.88}$$

3. **Dual variables**: The Lagrange multipliers from the previous solution are preserved to warm-start the interior point method

This warm starting reduces the average iteration count from approximately 15-20 to 8-12 iterations per MPC update.

4.4.3 Parametric Function Integration for Road Conditions

The double sigmoid functions for friction and curvature are implemented as parametric CasADi functions that enable efficient evaluation during optimisation:

$$\text{Friction function:} \quad \mu_{func} = \text{casadi.Function}(\mu, \{s, p_{fric}\}, \{\sigma(s; p_{fric})\}) \tag{4.89}$$

$$\text{Curvature function:} \quad \kappa_{func} = \text{casadi.Function}(\kappa, \{s, p_{curve}\}, \{\sigma(s; p_{curve})\}) \tag{4.90}$$

These functions are evaluated at predicted vehicle positions within the optimisation loop, enabling dynamic constraint adaptation based on spatial road parameter variation.

4.5 RNMPC Implementation

This section shows the complete numerical formulation and implementation details of the RNMPC for integrated ACC. The implementation leverages CasADi's symbolic framework and Opti interface to transform the theoretical framework into a computationally efficient optimisation problem solved in real-time using IPOPT.

4.5.1 MPC Settings and Parameter Selection

The RNMPC parameters represent a careful balance between competing objectives: prediction accuracy, computational tractability and control performance. Each parameter was selected based on systematic analysis and empirical validation. The specific parameters used in the implementation can be seen in Tables B.1), B.2), B.3), B.4), B.5).

The prediction horizon $T = N \cdot dt = 5$ seconds provides sufficient preview for highway driving scenarios while maintaining computational feasibility. The discretisation time step $dt = 0.5$ seconds represents a compromise between capturing rapid dynamics (which would favor smaller dt) and reducing computational burden (which favors larger dt).

4.5.2 CasADi Opti Interface Formulation

The implementation utilises CasADi's Opti interface, which provides an intuitive way to formulate and solve nonlinear optimisation problems. The symbolic variables are declared as:

$$X \in \mathbb{R}^{3\times(N+1)} \quad \text{(state trajectory)} \tag{4.91}$$

$$U \in \mathbb{R}^{1\times N} \quad \text{(control trajectory)} \tag{4.92}$$

$$S \in \mathbb{R}^{4\times N} \quad \text{(slack variables)} \tag{4.93}$$

Additionally, the framework employs optimisation parameters for real-time updates:

$$p_{x_0} \in \mathbb{R}^3 \quad \text{(initial state)} \tag{4.94}$$
$$p_{dsL_0} \in \mathbb{R} \quad \text{(initial lead distance)} \tag{4.95}$$
$$p_{vL_0} \in \mathbb{R} \quad \text{(initial lead velocity)} \tag{4.96}$$
$$p_{fric} \in \mathbb{R}^9 \quad \text{(friction parameters)} \tag{4.97}$$
$$p_{curve} \in \mathbb{R}^{6\times 2} \quad \text{(curvature bounds)} \tag{4.98}$$

4.5.3 Complete optimisation Problem Structure

At each time step k, the RNMPC solves the following optimisation problem using the Opti interface:

$$\begin{aligned}\min_{X,U,S} \quad J = &\sum_{i=0}^{N-1} \left[(X_{:,i} - x_{ref})^T Q (X_{:,i} - x_{ref}) + U_i^T R U_i + S_{:,i}^T R_s S_{:,i}\right] dt \\ &+ (X_{:,N} - x_{ref})^T Q (X_{:,N} - x_{ref})\end{aligned} \tag{4.99}$$

subject to the constraints detailed in the following subsections.

4.5.3.1 Dynamic Constraints

The discrete-time dynamics are enforced through equality constraints:

$$X_{:,i+1} = f_d(X_{:,i}, U_i), \quad i = 0, \ldots, N-1 \tag{4.100}$$

where f_d represents either the RK4 or the exact ZOH discretisation as selected by the user.

4.5.3.2 State and Control Constraints with Slack Variables

The implementation enforces constraints at each prediction step $k = 1, \ldots, N$:

Safety Distance Constraint:

$$0 \leq X_{1,k+1} - S_{1,k} \leq X_L^{1,k+1} - \delta s_{min} \tag{4.101}$$

Velocity Constraints:

$$v_{min} \leq X_{2,k+1} - S_{2,k} \leq v_{max} \tag{4.102}$$
$$X_{2,k+1}^2 \cdot \kappa_k - S_{2,k} \leq \mu_{k,low} \cdot g \tag{4.103}$$

Acceleration Constraints:

$$-a_{max} \leq X_{3,k+1} + S_{3,k} \leq a_{max} \tag{4.104}$$
$$-\mu_{k,low} \cdot g \leq X_{3,k+1} + S_{3,k} \leq \mu_{k,low} \cdot g \tag{4.105}$$
$$-a_{c,max} \leq X_{3,k+1} + S_{4,k} \leq a_{c,max} \tag{4.106}$$

Note that the implementation uses a specific slack variable formulation where slacks are consistently included with appropriate sign (added or subtracted) depending on whether they relax upper or lower bounds.

4.5.3.3 Lead Vehicle Prediction

The lead vehicle trajectory is predicted conservatively within the optimisation:

$$X_L^{3,k} = -\min(\mu_{high}(X_L^{1,k}) \cdot g, w_{a,lead}) \tag{4.107}$$
$$X_L^{:,k+1} = f_d(X_L^{:,k}, 0) \tag{4.108}$$
$$X_L^{2,k+1} = \max(X_L^{2,k+1}, v_{L,min}) \tag{4.109}$$
$$X_L^{1,k+1} = \max(X_L^{1,k+1}, X_L^{1,k}) \tag{4.110}$$

4.5.3.4 Terminal Constraint

The terminal constraint for recursive feasibility:

$$X_{2,N}^2 - v_{L,min}^2 + 2X_{3,N}(X_L^{1,N} - X_{1,N} - \delta s_{min}) - S_{1,N-1} \leq 0 \tag{4.111}$$

This constraint ensures the ego vehicle can safely decelerate to the lead vehicle's minimum speed within the available distance at the end of the prediction horizon.

4.5.4 Numerical Solution Strategy

The optimisation problem is transformed into a parametric function using CasADi:

```
mpc_ctrl = opti.to_function('mpc_ctrl',
    {p, X, U, S, opti.lam_g}, {X, U, S, opti.lam_g, XL},
    {'parameters', 'guess_X', 'guess_U', 'guess_S', 'guess_duals'},
    {'X_opt', 'U_opt', 'S_opt', 'dual_vars', 'XL'});
```

This function encapsulates the entire optimisation problem and enables efficient warm-starting through the provision of initial guesses and dual variables.

4.5.4.1 Warm-Starting

The warm-starting strategy to improve convergence:

1. **Control trajectory shift:**

$$U_{init} = [U_{prev,2:N}, U_{prev,N}] \tag{4.112}$$

2. **State trajectory propagation:**

$$X_{init} = [X_{prev,2:N+1}, f_d(X_{prev,N+1}, U_{prev,N})] \tag{4.113}$$

3. **Dual variable preservation:**

$$L_{init} = L_{prev} \tag{4.114}$$

This warm-starting approach significantly reduces iteration counts, typically achieving convergence in 8-12 iterations compared to 15-20 iterations without warm-starting.

4.5.5 Operational Modes for Uncertainty Evaluation

The implementation supports two distinct operational modes for comprehensive performance evaluation:

4.5.5.1 Deterministic Mode

In this mode, the actual friction follows the mean profile exactly:

$$\mu_{actual}(s_k) = \mu_{mean}(s_k) \tag{4.115}$$

The vehicle's acceleration is bounded deterministically:

$$a_k = \text{saturate}(a_k, -\mu_{actual}(s_k) \cdot g, \mu_{actual}(s_k) \cdot g) \tag{4.116}$$

4.5.5.2 Stochastic Mode

In this mode, friction is randomly sampled from a Beta distribution:

$$M = \frac{\mu_{mean}(s_k) - \mu_{low,k}}{\mu_{high,k} - \mu_{low,k}} \tag{4.117}$$

$$\alpha = \rho \cdot M, \quad \beta = \rho \cdot (1 - M) \tag{4.118}$$

$$\mu_{actual}(s_k) \sim \mu_{low,k} + (\mu_{high,k} - \mu_{low,k}) \cdot \text{Beta}(\alpha, \beta) \tag{4.119}$$

This stochastic sampling occurs at each simulation step, creating realistic friction variations that test the controller's robustness while respecting the uncertainty bounds used in the optimisation.

4.5.6 Computational Performance Monitoring

The implementation includes comprehensive performance monitoring through the `get_stats` function, which extracts solver statistics including iteration count, objective value, and convergence status. This enables real-time assessment of computational performance and helps identify scenarios that challenge the optimisation solver.

4.5.7 Lead Vehicle behaviour Modes

The implementation supports two modes for lead vehicle behaviour simulation:

4.5.7.1 Constant Speed Mode

In normal operation, the lead vehicle maintains constant speed:

$$a_{L,actual}(k) = 0 \tag{4.120}$$

This mode enables evaluation of the controller's tracking performance and response to road condition changes without the additional complexity of lead vehicle dynamics.

4.5.7.2 Random Acceleration Mode - Robustness Test

For robustness testing, the lead vehicle exhibits random acceleration within physical limits:

$$a_{L,actual}(k) = \xi_k \cdot \min(\mu_{high}(s_L(k)) \cdot g, w_{a,lead}), \quad \xi_k \sim \mathcal{U}[-1, 1] \tag{4.121}$$

This mode tests the controller's ability to maintain safe distances under unpredictable lead vehicle behaviour while managing friction uncertainty.

5

Results

5.1 Performance Metrics and Evaluation

The controller performance is evaluated through multiple metrics that capture safety, comfort and efficiency:

5.1.1 Discretisation Method Impact Analysis

Under the road conditions illustrated in Figure 5.1 featuring a pronounced drop in road friction and a localised curvature increase, the proposed MPC framework was evaluated by the ZOH and the RK4 scheme. In both cases, the simulation involved stochastic lead vehicle acceleration and probabilistic friction estimation with spatial uncertainty, making the control task sensitive to both environmental and dynamic changes. The complete set of simulation parameters used can be found in Section B.1.

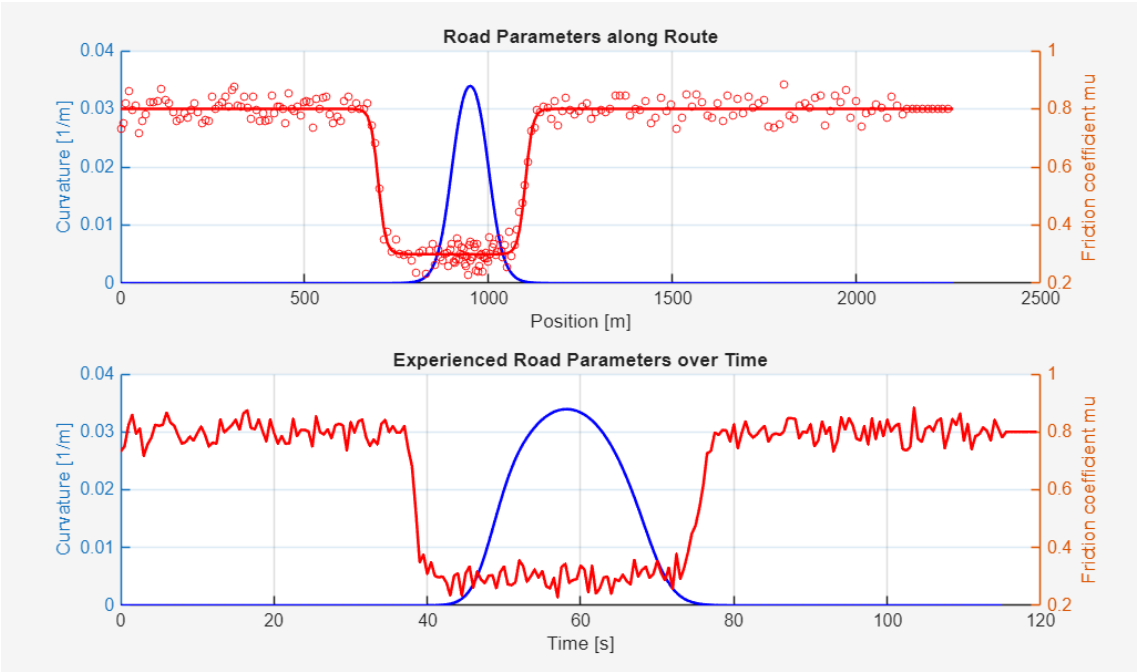

Figure 5.1: Road parameters used in the discretisation comparison scenario.

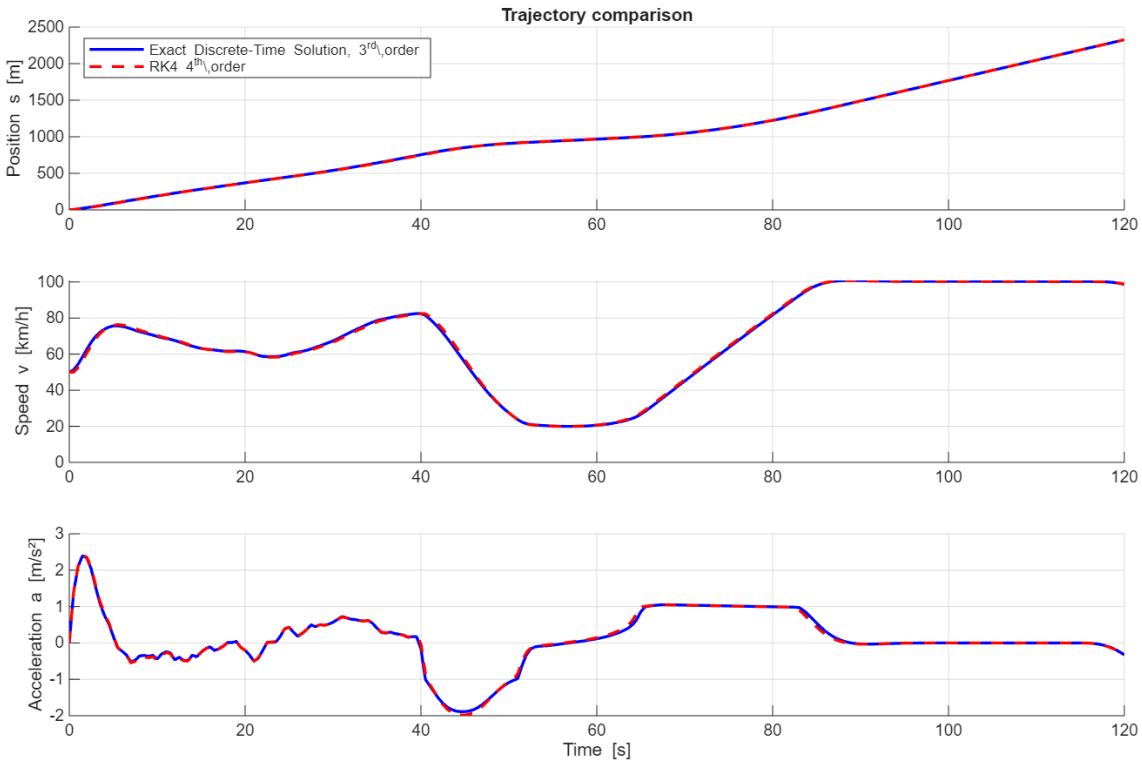

Figure 5.2: Comparison of position, speed, and acceleration trajectories between ZOH and RK4 discretisation methods.

Despite the very small numerical differences in discretisation accuracy, the resulting trajectories in terms of position, velocity and acceleration remain nearly identical as shown in Figure 5.2. This high consistency suggests that for the tested prediction horizon and sampling time ($\Delta t = 0.5s$), both discretisation methods capture the system dynamics, heavily due to the linear dynamics and the existence of an exact analytical solution.

5.1.2 Constraint Violation Analysis

The framework tracks multiple types of constraint violations to assess controller performance:

Distance Violations:

$$\text{violation}_{dist}(t) = \max(0, \delta s_{min} - (s_L(t) - s_{ego}(t))) \tag{5.1}$$

Friction Violations:

$$\text{violation}_{friction}(t) = \max(0, |a(t)| - \mu_{actual}(s(t)) \cdot g) \tag{5.2}$$

The following table summarises the constraint satisfaction status and key performance metrics for all use cases. Each constraint type is evaluated for violations throughout the entire simulation duration.

Table 5.1: Constraint Violation Analysis Summary for All Use Cases

UC	Constraint Status				Performance Metrics		
	Distance	**Velocity**	**Accel.**	**Comfort**	**Min Dist.** (m)	**Max Speed** (km/h)	**Max Accel.** (m/s^2)
1	✓	✓	✓	×	61.16	82.22	2.42
2	✓	✓	✓	×	50	100.51	2.70
3	✓	✓	✓	×	60	79.85	2.30
4	✓	✓	✓	×	150	100.76	2.93
5	✓	✓	✓	×	80	100.68	2.73
6	✓	✓	✓	×	70.77	89.16	2.72
7	✓	✓	✓	✓	30	100.47	1.92
8	✓	✓	✓	×	61.16	82.22	2.42
9	✓	✓	✓	×	43.40	82.38	2.60
10	✓	✓	✓	×	70	100.18	2.75

As can be seen in Table 5.1, the comfort bounds are violated across several use cases. It is observed that these violations primarily occur during the initial transient phase as the controller stabilises the vehicle toward the reference speed and following distance while respecting safety-critical constraints. This behaviour is further illustrated in Section 5.2, where the transient stabilisation phases and their effects on comfort are shown in detail.

5.1.3 Robustness Testing Framework

The robustness testing is being performed through random lead vehicle acceleration. The test scenario uses the same road parameters as shown earlier in Figure 5.1, including a spatially varying curvature profile and a friction transition from high to low values and back.

Stochastic Lead Acceleration:

$$a_L(t) = \xi(t) \cdot \min(\mu_{high}(s_L(t)) \cdot g, w_{a,lead}) \tag{5.3}$$

where $\xi(t) \sim \mathcal{U}[-1, 1]$ represents random acceleration commands.

Prediction Accuracy Metrics:

$$\text{Distance Error:} \quad e_{dist}(k) = s_{L,actual}(k) - s_{L,pred}(k) \tag{5.4}$$

$$\text{Violation Rate:} \quad r_{viol} = \frac{\text{Number of violations}}{\text{Total MPC updates}} \tag{5.5}$$

To assess robustness, two cases were evaluated: one with a constant-speed lead vehicle and another where the lead vehicle's acceleration varied randomly but within bounds ($\pm 3 m/s^2$) at each MPC update step. This setup simulates sudden or unpredictable braking and acceleration events, testing the controller's ability to maintain safety margins under dynamic uncertainty.

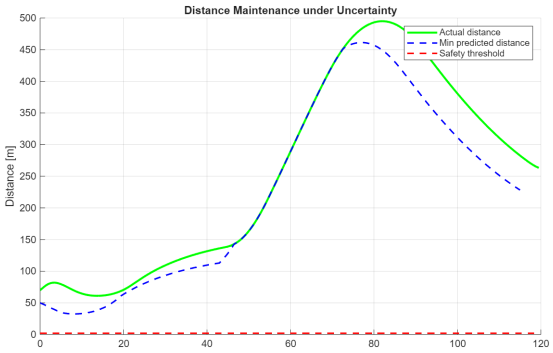

(a) Maintained gap with constant lead speed

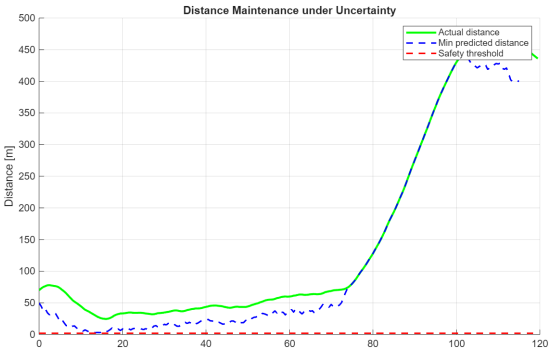

(b) Maintained gap with random lead acceleration

Figure 5.3: Maintained gap under the two lead-vehicle scenarios. The green solid line shows the actual inter-vehicle distance, the blue dashed line the minimum distance predicted by the MPC, and the red dashed line the imposed safety threshold.

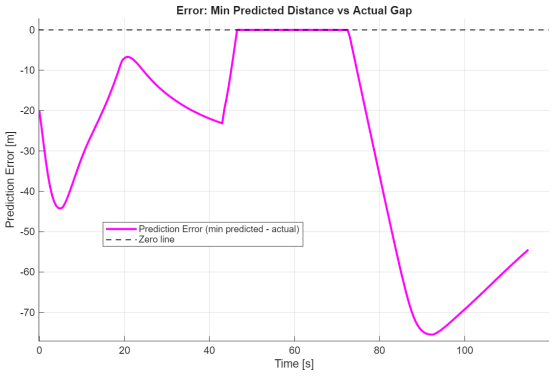

(a) Gap error with constant lead speed

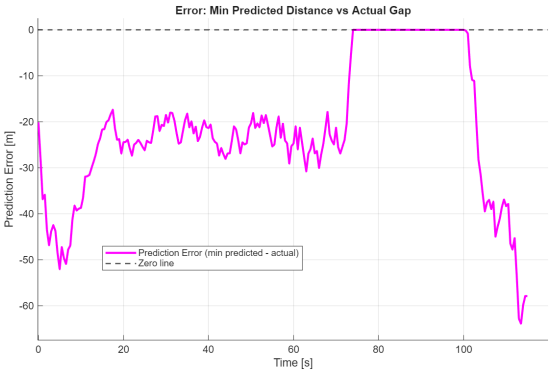

(b) Gap error with random lead acceleration

Figure 5.4: Gap-prediction error $e_{\text{dist}} = s_{\text{pred,min}} - s_{\text{actual}}$.

In both scenarios, the prediction error, defined as the difference between the minimum predicted distance and the actual gap remains consistently negative. This means that the predicted minimum distance is always smaller than the actual distance observed during the closed-loop execution, which aligns with the safety-oriented design philosophy of the controller. The MPC makes conservative predictions to ensure that, even under worst-case assumptions the ego vehicle can maintain a safe distance.

When comparing the constant acceleration case to the stochastic one, it can be observed that the overall shape and magnitude of the prediction error curves are similar. While the randomly accelerating lead vehicle introduces more fluctuations in the prediction error, the error values themselves remain within comparable bounds. This indicates that the controller adapts well to variations in lead vehicle behaviour and is not overly sensitive to unpredictability in acceleration profiles.

The reason the errors are not drastically more negative in the random case is likely due to the conservative design of the controller, which already plans under high uncertainty. Additionally, the prediction horizon length and update rate contribute to stabilising the error over time, preventing large deviations between predicted and actual distances. This conservative behaviour is also influenced by the controller's tuning parameters, such as cost weights and safety margins, which could be adjusted to yield different outcomes.

5.1.4 Computational Performance Analysis

Solver Convergence:

$$\text{Avg. Iterations:} \quad \bar{n}_{iter} = \frac{1}{N_{updates}} \sum_{k=1}^{N_{updates}} n_{iter}(k) \tag{5.6}$$

$$\text{Success Rate:} \quad r_{success} = \frac{N_{converged}}{N_{updates}} \tag{5.7}$$

Table 5.2: Computational Performance using RK4 discretisation

UC	Iterations			Computation Time (ms)			Success
	Mean	Max	Std	Mean	Max	Std	Rate
1	9.2	14	0.9	14.9	46.3	2.9	100%
2	12.3	17	1.5	34.7	631.7	64.8	100%
3	11.1	16	1.4	38.4	534.1	55.8	100%
4	10.1	17	1.8	42.2	518.2	58.0	100%
5	9.7	18	1.9	42.9	568.7	58.3	100%
6	8.4	15	1.0	47.0	803.4	72.9	100%
7	10.5	15	1.1	51.1	535.0	54.0	100%
8	9.1	14	0.7	49.5	715.5	59.4	100%
9	9.7	15	1.0	57.0	591.9	59.0	100%
10	10.4	16	1.4	60.7	561.9	54.5	100%

Table 5.3: Computational Performance using ZOH discretisation

UC	Iterations			Computation Time (ms)			Success
	Mean	Max	Std	Mean	Max	Std	Rate
1	8.4	14	1.0	14.0	35.5	2.6	100%
2	12.0	17	1.5	36.9	974.1	82.5	100%
3	10.9	15	1.1	36.2	592.4	61.3	100%
4	9.9	16	1.6	37.3	572.6	61.3	100%
5	9.5	15	1.6	40.9	602.0	59.4	100%
6	8.4	15	1.0	41.4	670.4	67.1	100%
7	10.5	16	1.2	51.2	674.9	65.1	100%
8	8.3	13	0.9	48.2	700.2	68.1	100%
9	10.1	15	0.9	56.2	615.6	63.2	100%
10	10.5	15	1.4	60.8	593.9	60.8	100%

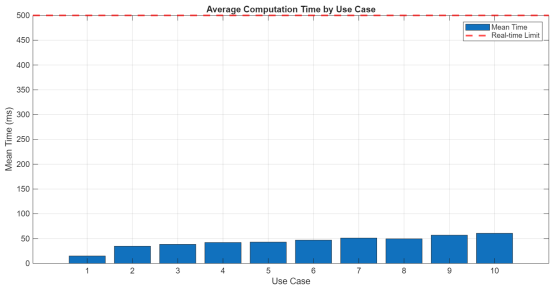

(a) RK4: Average Computation Time

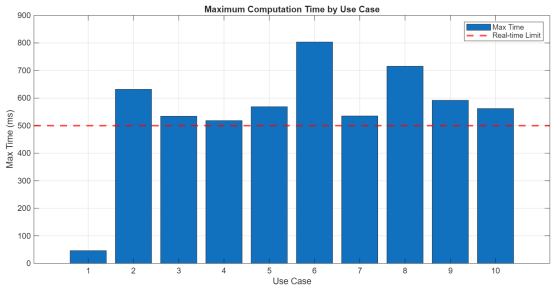

(b) RK4: Max. Computation Time

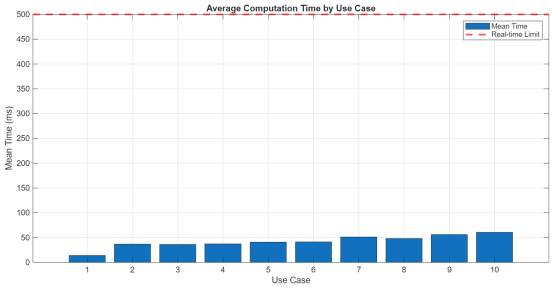

(c) ZOH: Average Computation Time

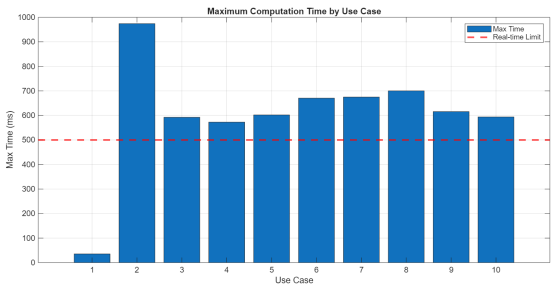

(d) ZOH: Max. Computation Time

Figure 5.5: Computational performance comparison between RK4 and ZOH discretisation methods. The red dashed line indicates a value representing the real-time computation time limit, set equal to 500 ms.

Average Performance: Both discretisation methods show excellent average computational performance, with RK4 having a mean computation time of around 40 ms, average iterations 9 and ZOH has almost the same mean values and with only negligible differences.

Worst-case Performance: Maximum computation times represent a limitation for real-time deployment, with RK4 having a maximum of 803.4 ms across in the worst computation time case, and ZOH method with maximum of 974.0 ms across the worst occuring case scenario (the case with the highest computational time).

Method Comparison: Both discretisation methods show similar performance characteristics, with differences mainly in specific worst-case scenarios rather than systematic advantages. The 40 ms average computation time represents excellent performance for typical MPC operations. However, it is important to note that there are values exceeding the threshold equal to 500ms, which comes from the sampling time which is once again equal to 500ms. This phenomenon, when the MPC is implemented in real-time could result to infeasibility. This is because sensor measurements enter as input at every 500ms, however at specific time steps the computation time for finding an optimal solution exceeds that value, meaning that the resulting control input will not be for the latest measurement input. However, real-time implementation is outside the scope of this thesis, and further improvements can be applied for further improving the computation time for real-time feasibility.

5.2 Scenario Results

This section presents a series of use cases designed to evaluate the performance and robustness of the proposed MPC framework under diverse road conditions, friction profiles and lead vehicle behaviours. Each scenario targets a specific challenge to assess tracking accuracy and constraint satisfaction. Results are analysed across key control signals including distance, velocity, acceleration, jerk and slack variables. All the controller parameters for the evaluated scenarios can be found in Appendix B.2.

UC1: Baseline ACC with Nominal Friction Uncertainty

UC1 evaluates baseline ACC performance under nominal friction-estimation bounds ($\mu_{unc} = [0.1, 0.3]$). The ego vehicle starts at 50 km/h and follows a constant-speed lead vehicle at 70 km/h while travelling through a friction profile with friction coefficients varying based on the sequence $0.8 \rightarrow 0.3 \rightarrow 0.8$ on a straight road.

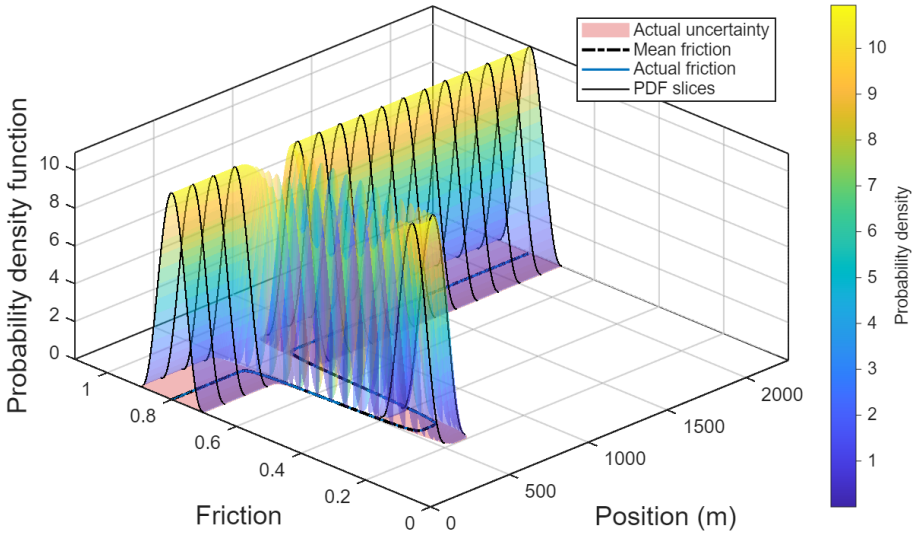

Figure 5.6: Spatial friction PDFs along the route. The surface encodes the Beta PDF; the band at z=0 is the actual uncertainty; the dashed line is the mean; the blue line is the realised friction; black curves are PDF slices.

Figure 5.6 provides a spatial map of road-friction uncertainty along the route. At each position s, the vertical section of the surface is the Beta-distributed friction PDF; color encodes probability density. The shaded band on the ground is the actual uncertainty interval, the dashed curve traces the mean friction profile, the blue line is the realized friction, and the black curves are a few illustrative PDF slices. This is a spatial view: it summarizes what is expected at each road position. When MPC forms a prediction horizon at the current position s, the first step coincides with the slice at s; subsequent steps widen and flatten as uncertainty increases. The next figure makes this explicit by (i) showing a local snapshot of a single prediction step within the MPC horizon and (ii) comparing the predicted bounds with the actual uncertainty in a compact 2D view.

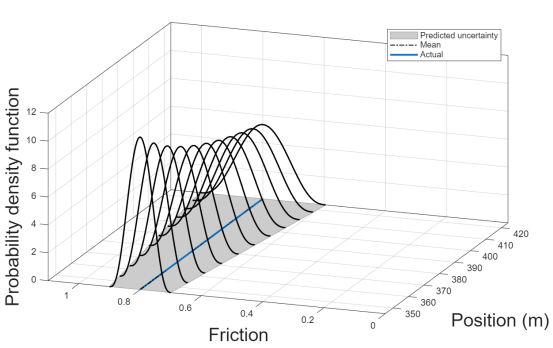

(a) Spatial snapshot friction-prediction envelope at 350 m.

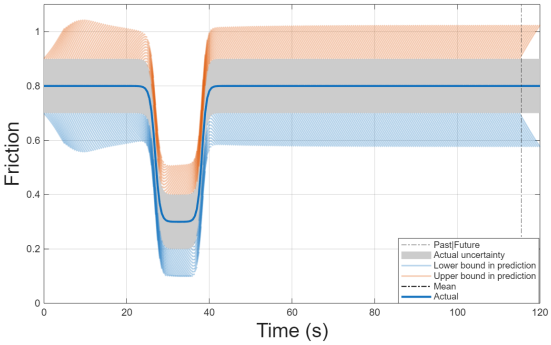

(b) Temporal evolution of friction under deterministic mode.

Figure 5.7: Friction behaviour under UC1.

Figure 5.7a shows the probabilistic friction estimation in the spatial domain, visualised as a series of evolving probability density functions along the vehicle path. The black curves show the Beta-distributed PDFs used to model friction uncertainty, while the grey region shows the predicted bounds of this uncertainty envelope. The blue line represents the actual friction realisation used during simulation. Since this scenario uses deterministic friction, the actual values align exactly with the mean of the predicted distributions. Note that this visualisation represents a single prediction step within the MPC horizon. Since the friction mean and bounds vary along the vehicle path, the friction distribution changes shape before, during, and after the drop. The snapshot shown here, taken at 350 m, captures the distribution before the start of the transition.

Figure 5.7b illustrates the temporal evolution of road friction and its associated uncertainty throughout the scenario. The blue line represents the deterministic friction profile applied during the simulation, the shaded grey region shows the actual uncertainty envelope, while the orange and blue hatchings indicate the predicted upper and lower bounds used by the MPC. The actual friction remains fully deterministic in this use case, and thus fully tracks the mean. The visualisation highlights a sharp friction drop and recovery, demonstrating how the controller anticipates changes. This figure represents the full horizon evolution across the entire drive.

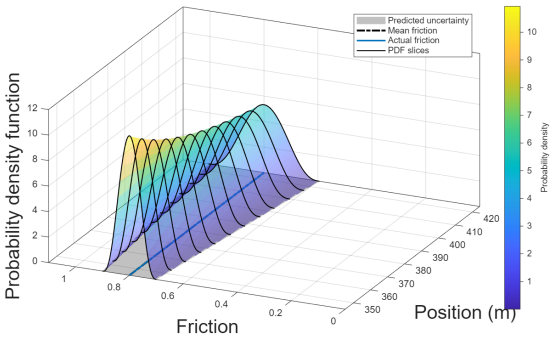

(a)

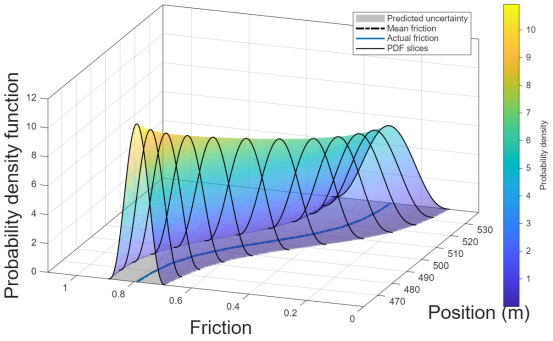

(b)

Figure 5.8: Additional snapshots of the friction-prediction envelope

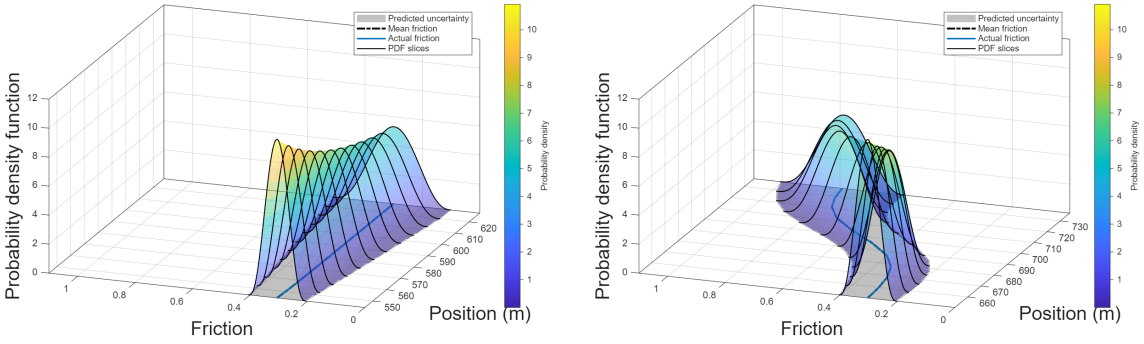

(a) **(b)**

Figure 5.9: Additional snapshots of the friction-prediction envelope

Figures 5.8 and 5.9 present additional prediction-horizon snapshots at different current positions, using the same colour-coded PDF surface as Fig. 5.6. Consistent with Fig. 5.7a, the predicted uncertainty envelope widens and the PDFs flatten as look-ahead increases. These snapshots illustrate representative cases: before, during, and after the low-friction segment.

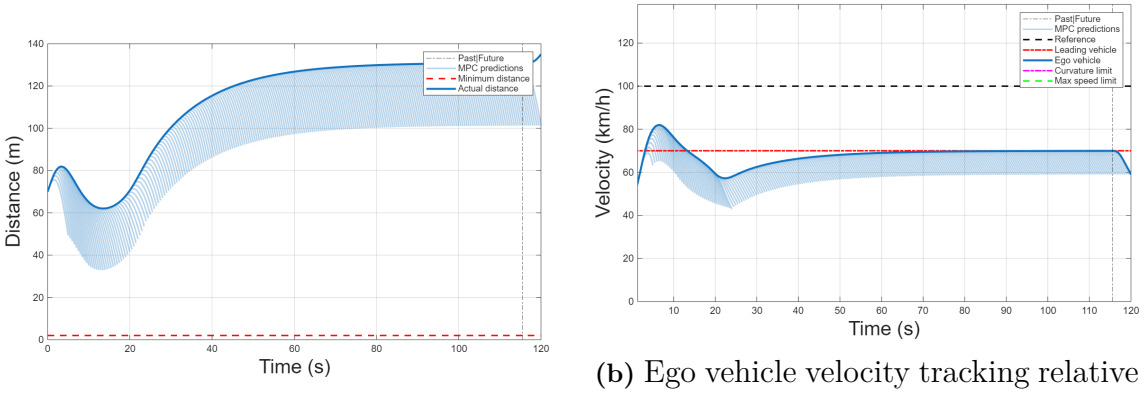

(a) Inter-vehicular distance tracking with MPC predictions.

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.10: Distance and velocity tracking performance under UC1.

Figure 5.10 illustrates the evolution of inter-vehicular distance and ego vehicle velocity over time for UC1. Figure 5.10a shows that the controller's predictions are accurate, as the actual following distance closely aligns with the MPC trajectories and remains consistently above the minimum safety margin. Notably, during the friction drop, a proactive increase in the inter-vehicular gap can be observed, followed by stabilisation of the spacing.

As shown in Figure 5.10b, this behaviour is coordinated with a corresponding decrease in ego speed, demonstrating that the controller anticipates reduced traction and adjusts accordingly to maintain safety. The ego vehicle subsequently stabilises around the lead vehicle's velocity, which is the intended ACC objective in this case.

This confirms that the controller achieves steady-state following while respecting acceleration constraints and adapting effectively to varying friction conditions.

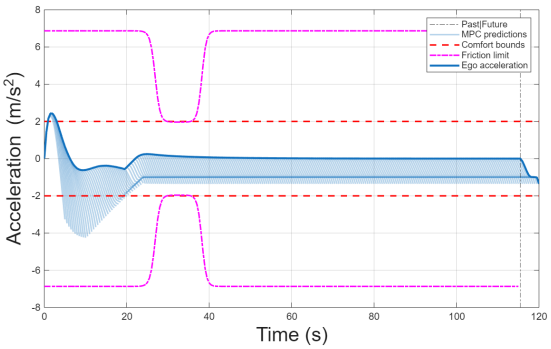

(a) Acceleration profile with MPC predictions and comfort/friction limits.

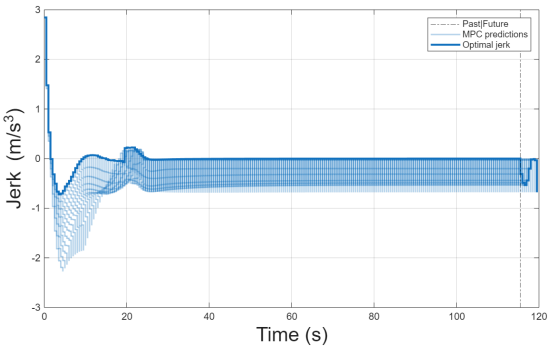

(b) Jerk profile with MPC predictions.

Figure 5.11: Acceleration and jerk response of the ego vehicle under UC1.

Figure 5.11 illustrate the acceleration and jerk profiles for the ego vehicle in UC1. In Figure 5.11a, the acceleration profile remains strictly within both comfort and friction limits during steady-state operation. A small overshoot occurs at the start of the scenario, corresponding to the transient phase as the controller responds to initial conditions. This behaviour is expected and remains within acceptable limits.

Figure 5.11b shows that the jerk input remains within reasonable limits, ensuring smooth transitions. The profile gradually converges to zero after the transient phase, indicating that the controller achieves stable and comfortable control inputs as the vehicle reaches steady-state conditions.

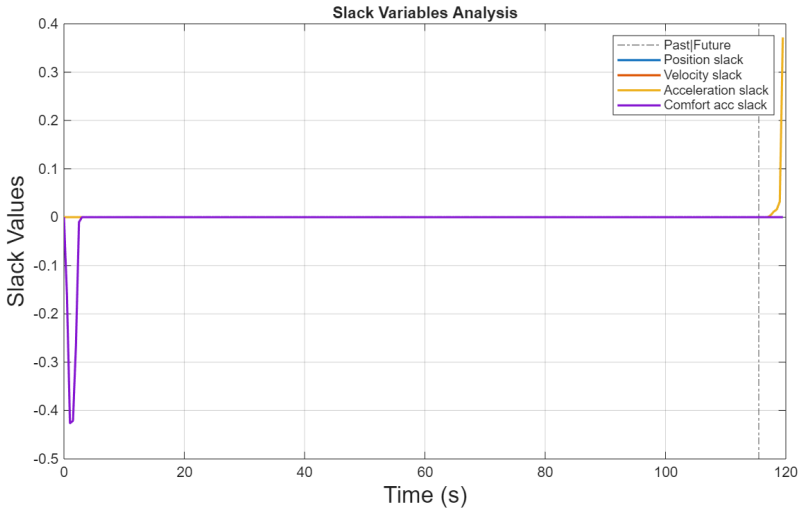

Figure 5.12: Evolution of slack variables under UC1.

Figure 5.12 presents the slack-variable evolution for UC1. As shown, all slack values remain effectively zero throughout the entire simulation, indicating that the MPC

controller consistently satisfies all hard and soft constraints without requiring any relaxation. Slight comfort acceleration slack relaxation at the beginning can be observed, but as mentioned earlier, this occurs during the transient phase. This confirms that under nominal uncertainty conditions, the controller maintains full feasibility while enforcing safety gaps, velocity bounds and comfort-related acceleration and jerk constraints. The absence of active slack usage also demonstrates that the optimisation remains well conditioned and conservative behaviour is not triggered unnecessarily.

UC2: Robustness Test with Random Lead Acceleration

UC2 investigates controller robustness under combined stochastic friction and random lead-vehicle acceleration. This scenario stresses the RNMPC by introducing unpredictable variations in both road and traffic conditions, testing its ability to maintain safety margins, respect physical limits, and deliver smooth control inputs under uncertainty

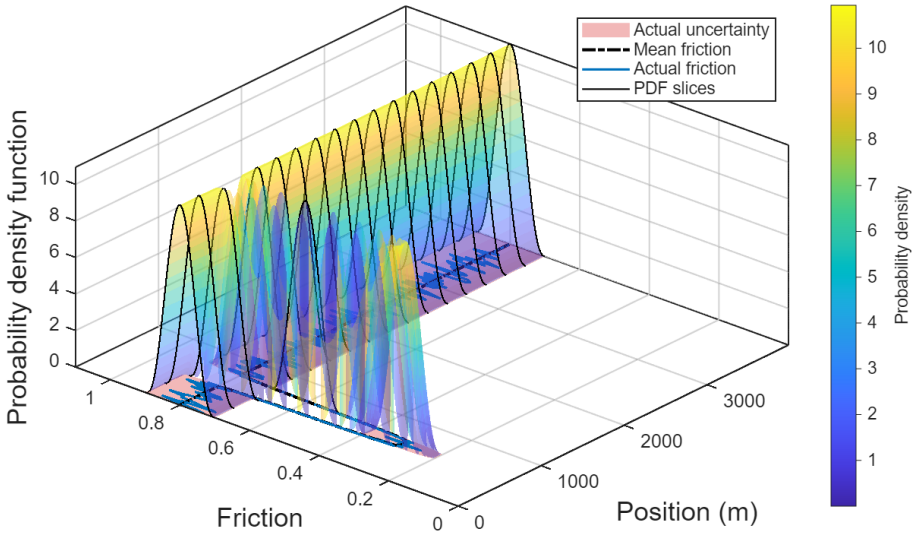

Figure 5.13: Spatial friction PDFs along the route.

Compared with UC1, the realised friction in UC2 does not track the mean, reflecting the random draws from the Beta distribution. This figure provides the spatial context for the snapshot and time plots in Fig. 5.14.

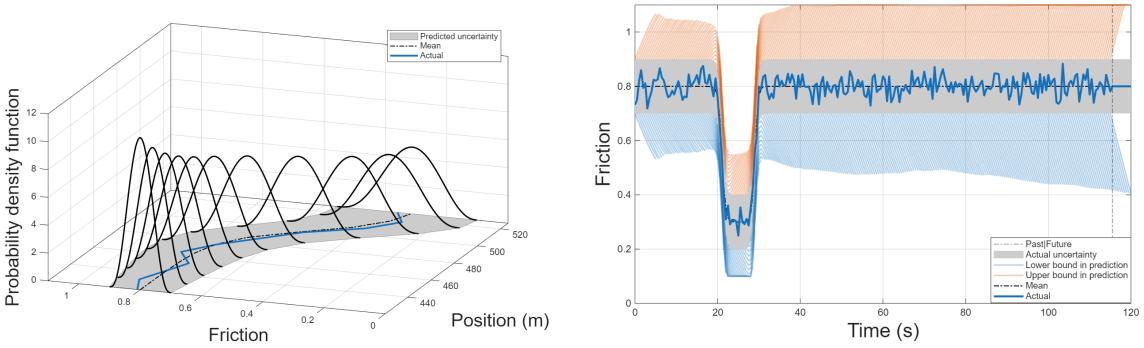

(a) Spatial snapshot of friction-prediction envelope.

(b) Temporal evolution of friction under stochastic mode.

Figure 5.14: Friction behaviour under UC2.

Figure 5.14a shows the spatial evolution of friction prediction under stochastic uncertainty in UC2. As mentioned, unlike the deterministic baseline, the actual friction realisation deviates from the predicted mean at several points, reflecting the random draws from the Beta distribution.

Figure 5.14b provides the temporal perspective. The actual realisation oscillates within the predicted envelope, demonstrating that the stochastic sampling introduces small-scale fluctuations around the mean. Here, the MPC's predictive uncertainty is visualised through the orange and blue hatched regions, which represent the upper and lower bounds of friction prediction. Compared to UC1, these bounds are noticeably wider. The expanded uncertainty bands illustrate that the controller must prepare for both optimistic and pessimistic friction scenarios, ensuring robust constraint satisfaction.

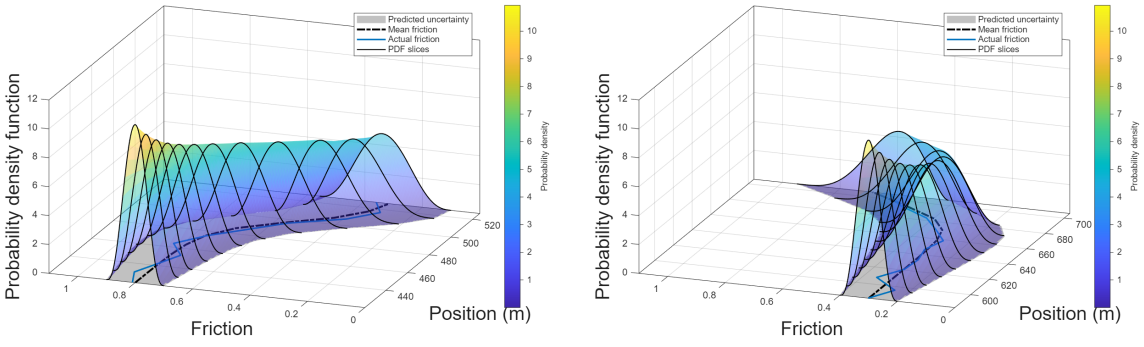

(a) **(b)**

Figure 5.15: Additional snapshots of the friction-prediction envelope

Figure 5.15 shows some additional prediction-horizon snapshots for UC2 under stochastic friction.

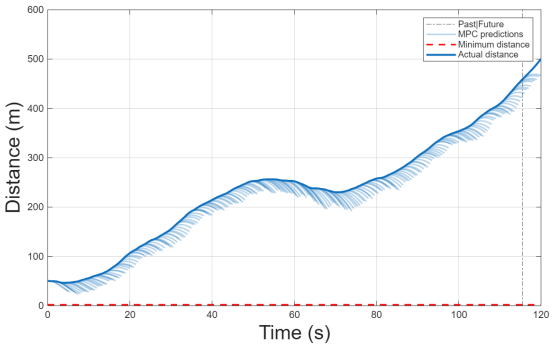

(a) Inter-vehicular distance tracking with MPC predictions.

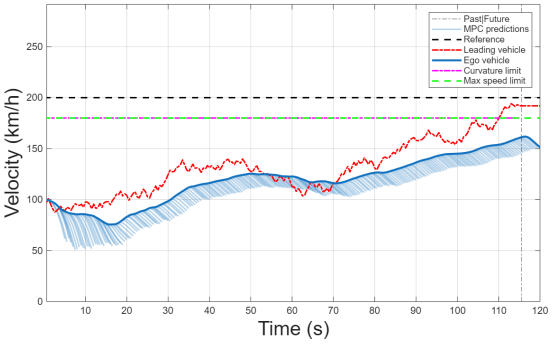

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.16: Distance and velocity tracking performance under UC2.

Figure 5.16a shows that the ego vehicle increases its separation from the lead vehicle when passing through the friction drop. Instead of trying to exactly match the lead vehicle's motion with aggressive accelerations, sharp braking, or jerky manoeuvres, the controller smooths its trajectory to prioritise safety and comfort. As a result, it does not always attempt to catch up immediately when the lead vehicle accelerates, which also contributes to the growing gap. A slight decrease in distance is observed around 60 s when the lead vehicle brakes more heavily, the ego vehicle smooths this deceleration thanks to the available safety margin, so it does not need to brake as abruptly.

Figure 5.16b illustrates this effect more clearly. While the lead vehicle's velocity fluctuates due to random acceleration inputs, the ego vehicle responds more cautiously, following the general trend but avoiding the high-frequency variations. This smoothing effect prevents abrupt manoeuvres, demonstrating the MPC's ability to prioritise robustness and passenger comfort at the cost of falling slightly behind the lead vehicle.

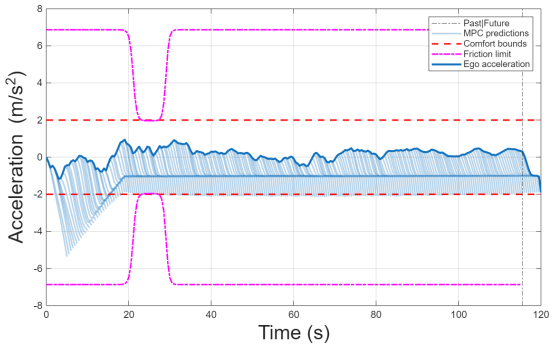

(a) Acceleration profile with MPC predictions and comfort/friction limits.

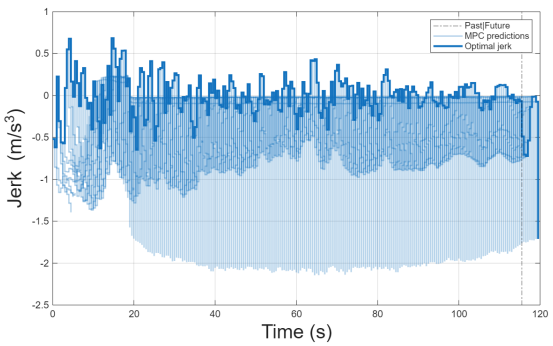

(b) Jerk profile with MPC predictions.

Figure 5.17: Acceleration and jerk response of the ego vehicle under UC2.

Figure 5.17a shows the ego vehicle's acceleration profile under stochastic friction and random lead-vehicle behaviour. The acceleration remains within both comfort and friction bounds at all times. However, due to the randomness in lead dynamics, small oscillations appear in the acceleration signal. These are expected, as the controller continuously readjusts to unpredictable inputs, but the magnitude remains moderate.

Figure 5.17b illustrates the corresponding jerk profile. The jerk remains bounded within a narrow range, around $\pm 0.5 m/s^3$, and exhibit minor fluctuations in response to the varying lead vehicle behaviour. Although some oscillations are observed due to the stochastic nature of the scenario, the controller still maintains smoothness and stability, ensuring passenger comfort is not compromised.

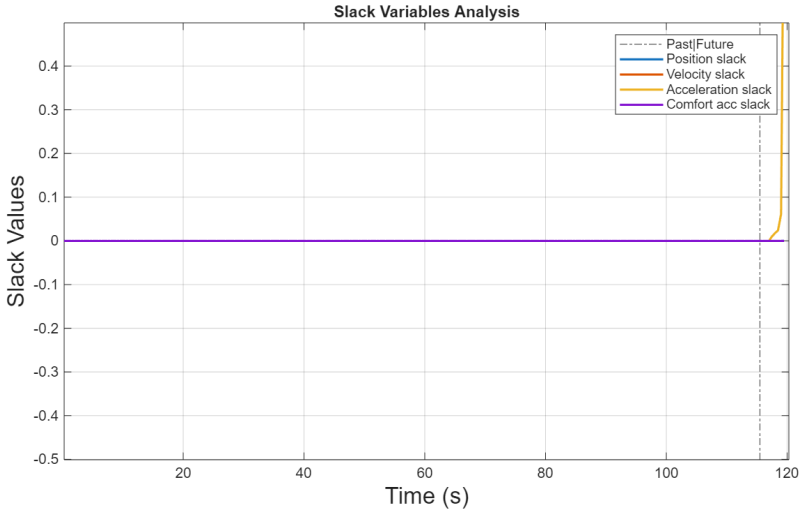

Figure 5.18: Evolution of slack variables under UC2.

Figure 5.18 illustrates the slack-variable evolution during the UC2 scenario. In this case, the ego and lead vehicle start at the same initial speed, which eliminates the large transient observed previously. Throughout the entire simulation, all slack values remain effectively zero, reflecting the system's ability to adapt effectively without relying on constraint relaxation as a regular mechanism.

UC3: High-Uncertainty Friction Transition

This use case investigates the controller's behaviour under significantly increased friction uncertainty, where the predicted friction bounds span $\mu_{\text{unc}} = [0.1, 0.7]$. Such a scenario represents degraded sensor confidence, e.g., due to adverse weather or poor visibility. This scenario evaluates how high uncertainty propagates through MPC prediction and affects the ego vehicle's safety margin selection.

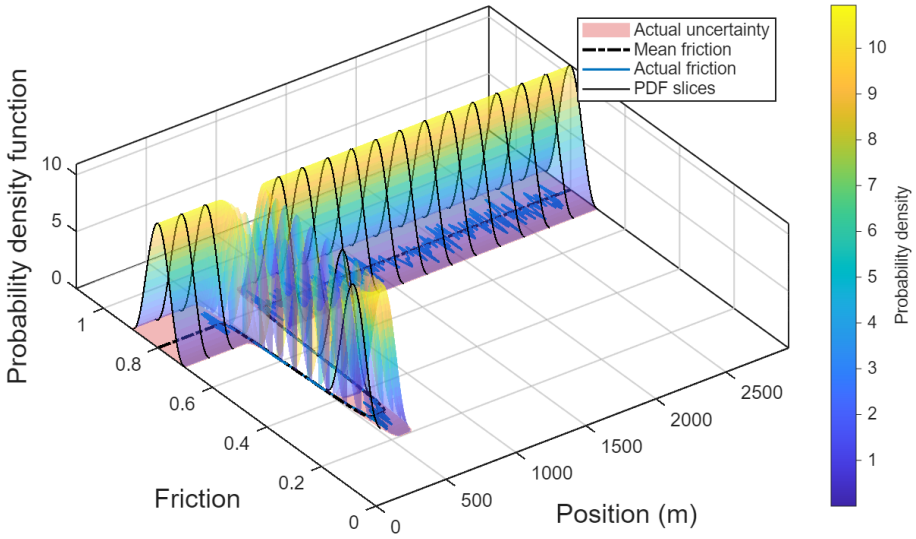

Figure 5.19: Spatial friction PDFs along the route.

Figure 5.19 provides the spatial map of road-friction uncertainty along the route.

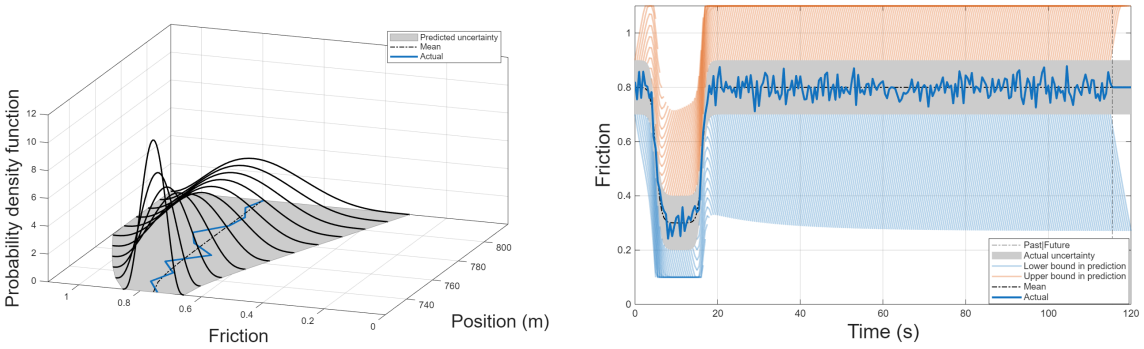

(a) Spatial snapshot of friction-prediction envelope.

(b) Temporal evolution of friction under stochastic mode.

Figure 5.20: Friction behaviour under UC3.

Figure 5.20a shows the spatial evolution of the probabilistic friction predictions under the high-uncertainty setting. Compared to UC1-UC2, the predicted uncertainty envelope is considerably wider, reflecting degraded sensing confidence. This directly implies that the MPC must adopt more conservative safety margins, as the likelihood of encountering critically low friction is substantially higher.

Figure 5.20b illustrates the temporal evolution of friction, where a sharp drop occurs early in the scenario followed by recovery. Compared to UC1-UC2, both the upper and lower prediction bounds expand, with the lower bound widening significantly. This captures the increased risk of critical low-friction outcomes and explains why the controller must plan cautiously throughout the scenario.

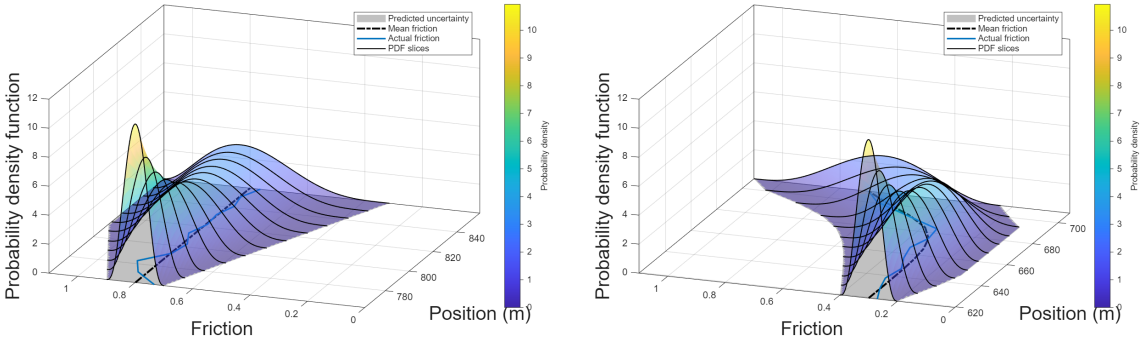

(a) **(b)**

Figure 5.21: Additional snapshots of the friction-prediction envelope

Figure 5.21 shows some additional prediction-horizon snapshots for UC3.

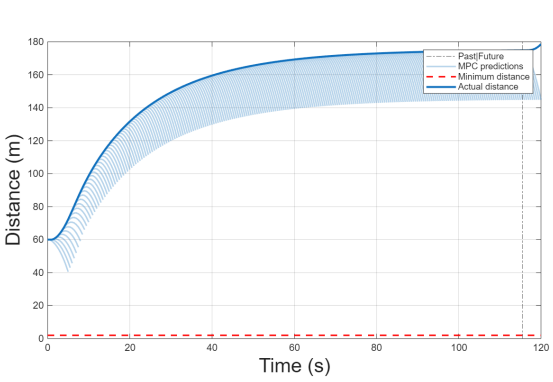

(a) Inter-vehicular distance tracking with MPC predictions.

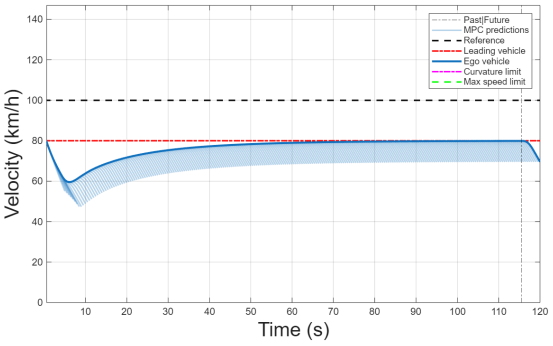

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.22: Distance and velocity tracking performance under UC3.

Figure 5.22a shows the inter-vehicular distance under UC3. Despite the significantly wider friction uncertainty bounds, the distance prediction envelope does not expand much compared to UC1. This is because the MPC already plans conservatively using the lower friction bound, so the effect of uncertainty is absorbed at the constraint level rather than appearing as a wider spread in the distance trajectory. The realised inter-vehicular spacing increases primarily because of the early friction drop, which forces the ego vehicle to reduce its speed relative to the lead.

Figure 5.22b shows the ego vehicle velocity relative to the lead vehicle. The early friction drop causes an immediate reduction in ego speed, as the controller proactively lowers the speed to remain within feasible bounds. This velocity drop explains the growing inter-vehicular gap seen in Figure 5.22a. Once the road friction recovers, the ego vehicle gradually increases its speed and converges back toward the lead's velocity.

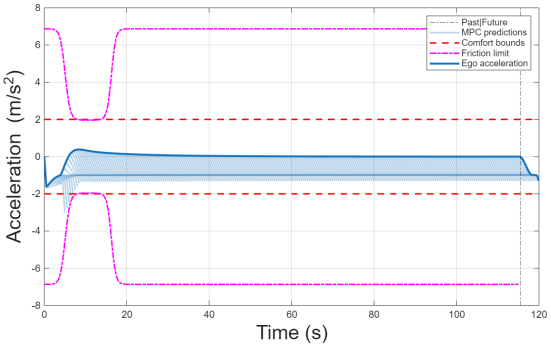

(a) Acceleration profile with MPC predictions and comfort/friction limits.

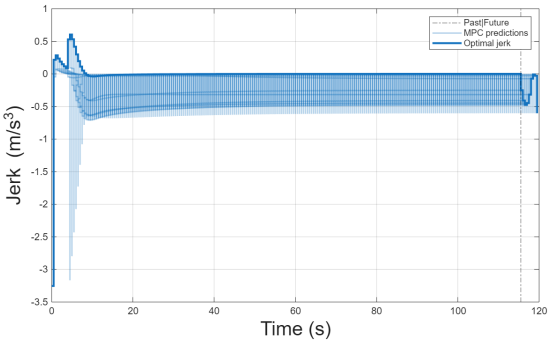

(b) Jerk profile with MPC predictions.

Figure 5.23: Acceleration and jerk response of the ego vehicle under UC3.

Figure 5.23a shows that the ego vehicle's acceleration remains well within both comfort and friction bounds throughout the scenario. Since the ego starts at the same speed as the lead vehicle, only minor adjustments are required due to the early friction drop and the need for braking to increase the safe distance. Otherwise, no strong accelerations or harsh braking are necessary, resulting in a smooth and feasible profile.

Figure 5.23b illustrates the corresponding jerk input, which stays at very small values for the entire horizon. The absence of large fluctuations confirms that no aggressive corrections were required, consistent with the smooth acceleration behaviour.

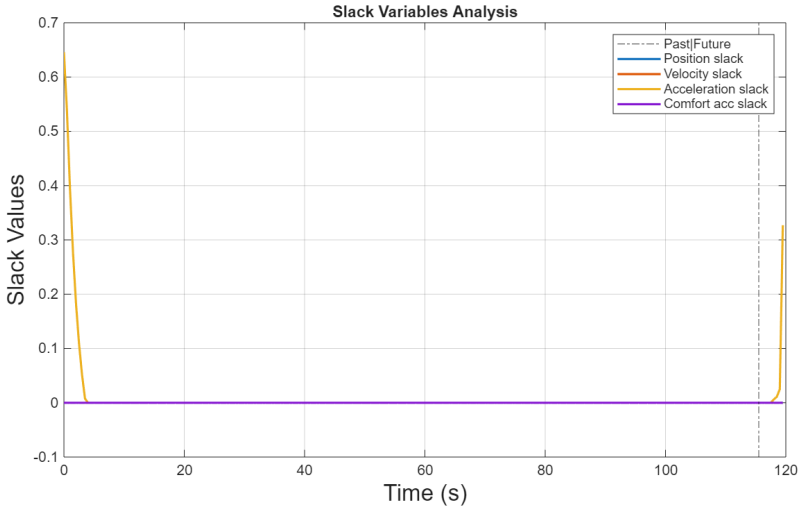

Figure 5.24: Evolution of slack variables under UC3.

Figure 5.24 shows that despite the increased uncertainty range in friction predictions, no significant slack activations are observed during the scenario. A brief activation of the acceleration slack is visible in the beginning, likely due to the controller's proactive adjustment to the sudden drop in friction. However, this usage remains minimal and transient.

UC4: Curve Navigation at High Speed

UC4 focuses on validating the controller's ability to proactively reduce speed for safe curve negotiation at high initial velocity. The ego vehicle starts at 140km/h and approaches a sharp curve, while the lead vehicle cruises at a constant 90 km/h.

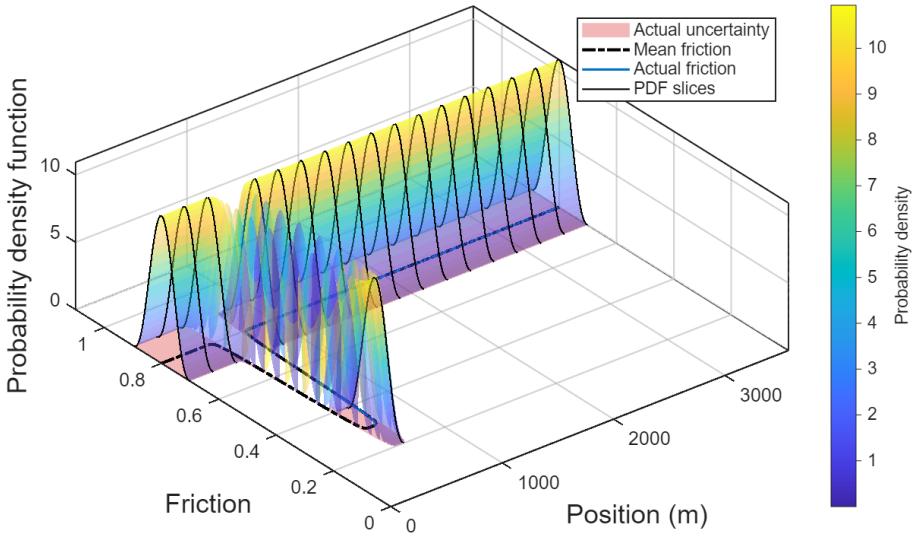

Figure 5.25: Spatial friction PDFs along the route.

Figure 5.25 provides the spatial map of road-friction uncertainty along the route.

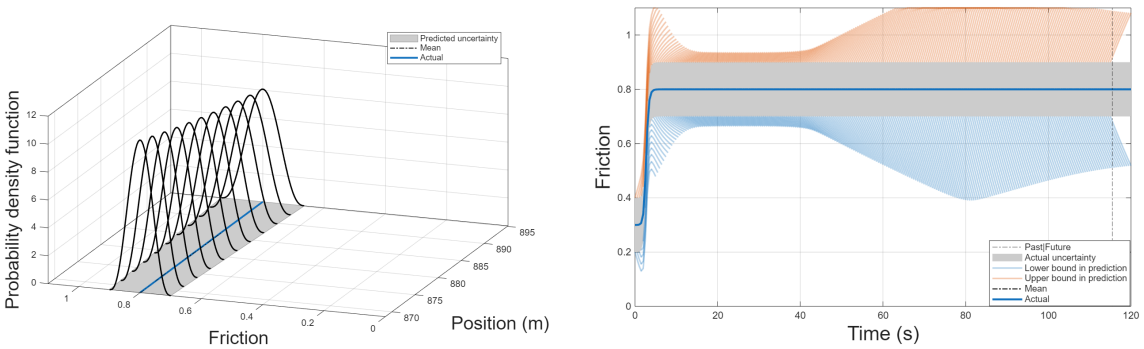

(a) Spatial snapshot of friction-prediction envelope.

(b) Temporal evolution of friction under deterministic mode.

Figure 5.26: Friction behaviour under UC4.

Figure 5.26a shows the evolution of friction prediction in UC4 where the friction is treated deterministically. This setup ensures that the controller's behaviour can be directly linked to curvature handling rather than variability in surface estimates.

Unlike the previous cases, Figure 5.26b shows that friction increases initially, starting from a low value before stabilising. The friction then remains constant, meaning that any subsequent control adjustments, such as proactive deceleration, are dictated

purely by the geometric road constraint rather than changing surface conditions. Around 40 s, a widening of the prediction envelope can be observed. This stem from the relaxation of the curvature constraint as the ego vehicle exits the sharp bend. With fewer limitations imposed by road geometry, the MPC generates a broader set of feasible trajectories, which appears as a more spread prediction band.

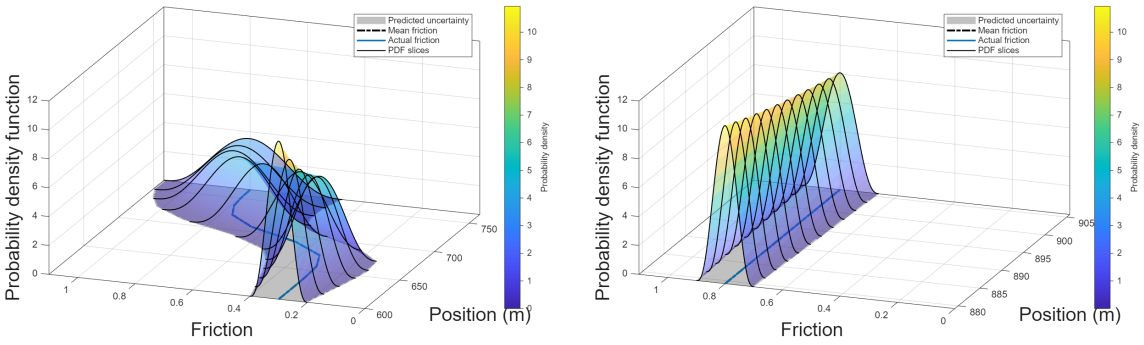

(a) **(b)**

Figure 5.27: Additional snapshots of the friction-prediction envelope

Figure 5.27 shows some additional prediction-horizon snapshots for UC4 under stochastic friction.

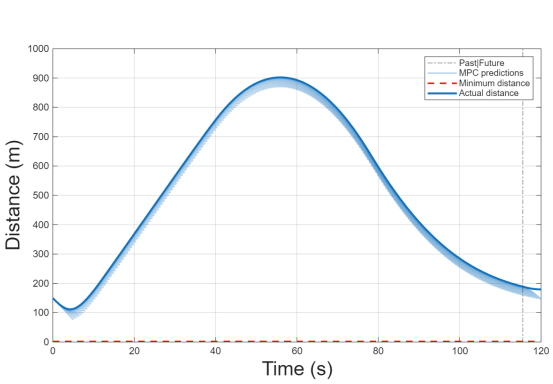

(a) Inter-vehicular distance tracking with MPC predictions.

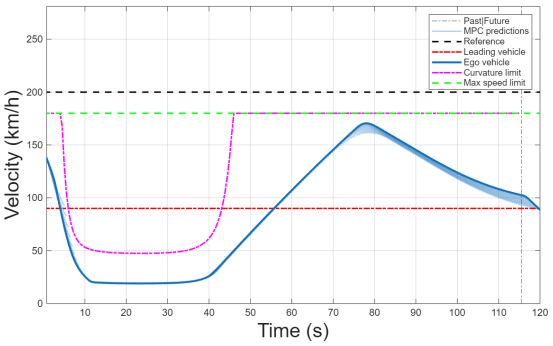

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.28: Distance and velocity tracking performance under UC4.

Figure 5.28a shows that the spacing between vehicles initially decreases as the ego car, starting at a higher speed than the lead vehicle, naturally begins to close the distance. Once the curvature-limited speed becomes the active constraint, the ego vehicle slows down below the lead vehicle's speed, and the distance grows sharply. After exiting the curve, the ego accelerates toward the reference speed, here capped by the maximum speed limit, as the large gap provides sufficient buffer. As the distance then decreases again, the controller reduces speed and stabilises around the lead vehicle's velocity, restoring steady-state following.

Figure 5.28b complements this reasoning by showing the velocity profile that drives the distance evolution. The ego speed tracks the curve-imposed limit instead of the lead car velocity, falls below the lead for the duration of the bend, then briefly overshoots, and finally starts to converge to the same steady speed as the lead vehicle. The ego speed is significantly lower than the curvature limit during the curve due to the conservative MPC design. The curvature constraint is evaluated with respect to the lower friction bound, which reduces the admissible curve speed compared to the nominal physical limit. At the same time, the controller accounts for inter-vehicular distance and comfort constraints, avoiding aggressive manoeuvres that would close the gap too quickly or violate smoothness requirements. This conservatism is amplified by the high initial ego speed, and without an explicit cost incentive to track the curve limit precisely, the optimizer naturally settles at a lower, feasible velocity that maintains safety and comfort, prioritizing robust satisfaction of all constraints.

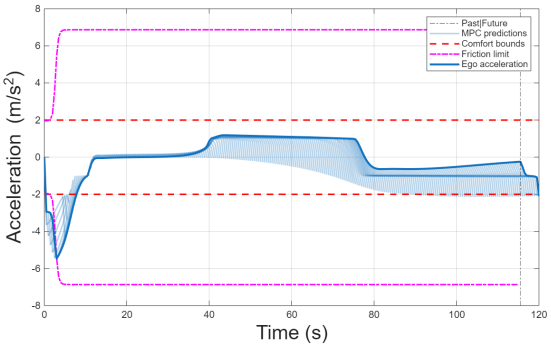

(a) Acceleration profile with MPC predictions and comfort/friction limits.

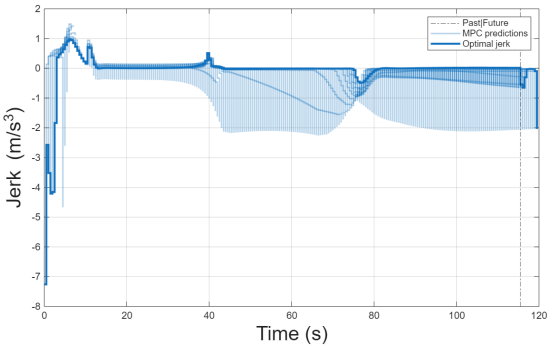

(b) Jerk profile with MPC predictions.

Figure 5.29: Acceleration and jerk response of the ego vehicle under UC4.

Figure 5.29a shows that the controller initiates a strong braking phase immediately at the start due to the combined effect of high initial ego speed, low-friction zone and the approaching curvature-imposed speed limit. The acceleration remains for most of the run within both comfort and friction bounds but with a saturation at the lower limit during the initial braking. After passing through the low-friction region and exiting the curve, the ego vehicle applies a sustained positive acceleration to regain speed and close the spacing to the lead vehicle, before stabilising as it converges towards the lead's steady velocity.

Figure 5.29b shows a negative jerk spike at the onset of braking, caused by the sharp deceleration demand under both friction and curvature constraints. Once the transient is over, jerk values remain small and bounded, reflecting smooth control effort. Slight fluctuations are visible during the transitions into and out of the curve, but no discomfort-inducing oscillations occur.

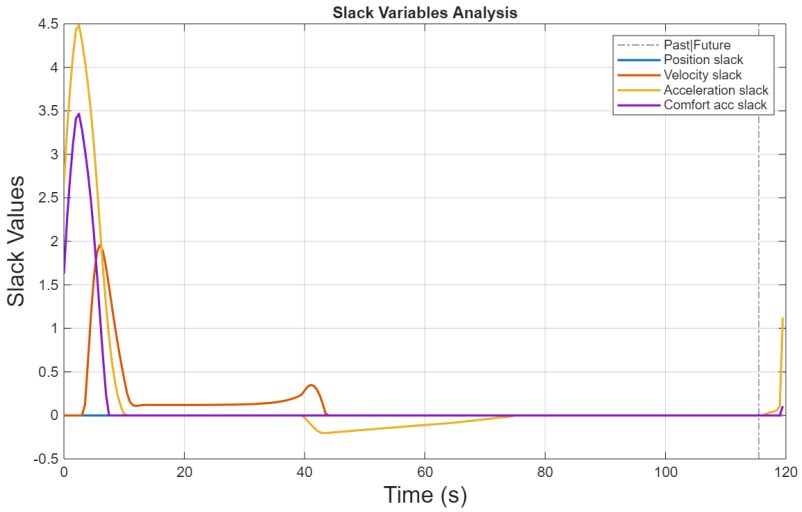

Figure 5.30: Evolution of slack variables under UC4.

Figure 5.30 reveals that multiple slack variables are briefly activated in the initial phase of the scenario. This includes acceleration slack, comfort acceleration slack and velocity slack. The activations coincide with the sharp deceleration required at the beginning, where the combined effect of high initial speed and tight curvature constraints makes it temporarily infeasible to satisfy all limits simultaneously.

Some slack variables exhibit a short settling period before converging to zero. This behaviour indicates that the controller uses constraint relaxation momentarily to maintain feasibility, while gradually steering the system back within all constraint boundaries as soon as conditions permit. Once the vehicle exits the curvature-limited region, all slacks remain inactive for the remainder of the drive, confirming successful constraint recovery and stable operation.

UC5: Combined Following and Curve Challenge

This use case evaluates the controller's response when multiple constraints become active simultaneously: curve limits, random lead vehicle motion, and increased friction uncertainty.

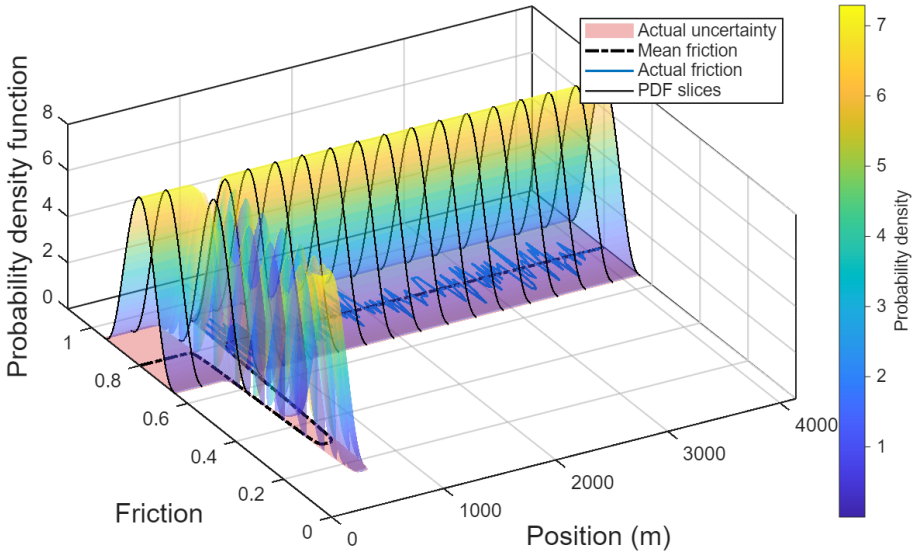

Figure 5.31: Spatial friction PDFs along the route.

Figure 5.31 provides the spatial map of road-friction uncertainty along the route.

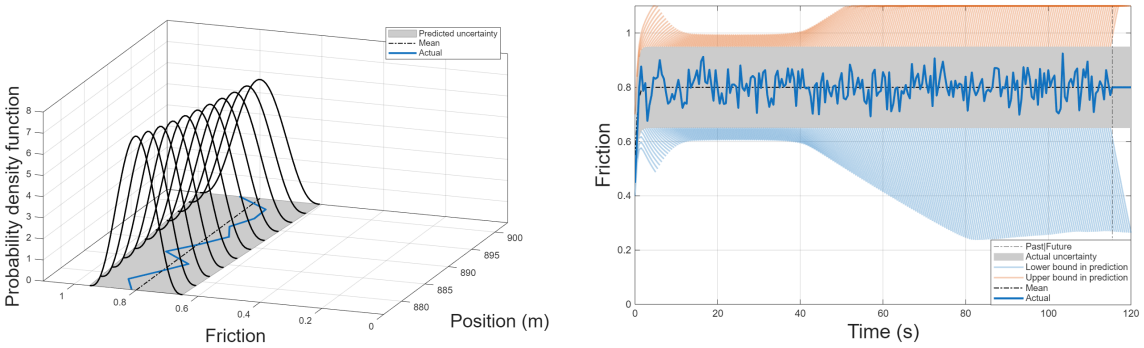

(a) Spatial snapshot of friction-prediction envelope.

(b) Temporal evolution of friction under stochastic mode.

Figure 5.32: Friction behaviour under UC5.

Figure 5.32a shows the spatial friction predictions under increased uncertainty. The probability density functions are spread more widely, which reflects a greater lack of confidence in friction estimation. This wider envelope implies that the MPC must plan more conservatively, anticipating that significantly lower friction values may occur along the trajectory.

In Figure 5.32b, the mean friction remains constant around $\mu = 0.8$, with no deterministic drop. Instead, the realised friction oscillates stochastically around this mean. Around $t \approx 40$ s, a widening of the prediction envelope can be observed, similar to UC4, which, as mentioned, originates from the relaxation of the curvature constraint as the ego vehicle exits the sharp bend. However, unlike UC4, the envelope here expands more noticeably. This broader spread is caused by the increased friction uncertainty in UC5, which results in larger prediction bands overall.

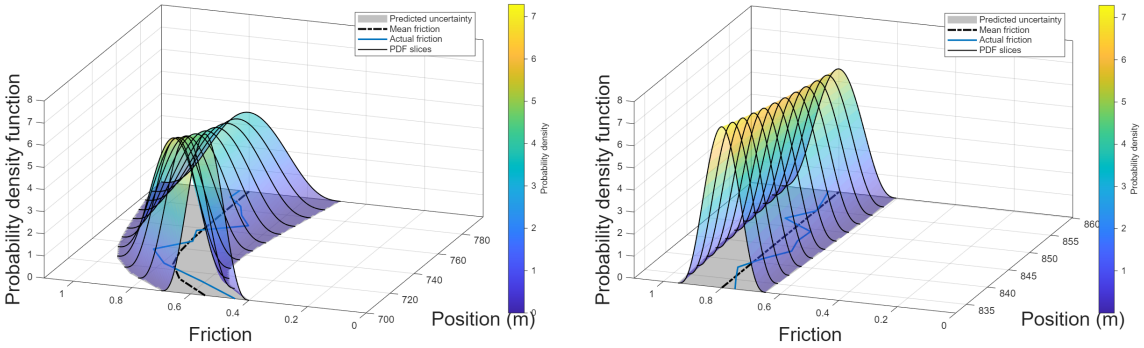

(a) **(b)**

Figure 5.33: Additional snapshots of the friction-prediction envelope

Figure 5.33 shows some additional prediction-horizon snapshots for UC5 under stochastic friction.

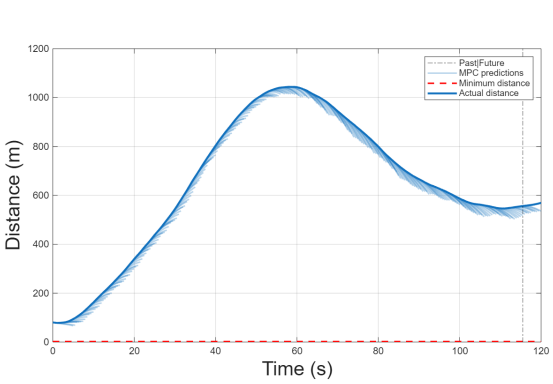

(a) Inter-vehicular distance tracking with MPC predictions.

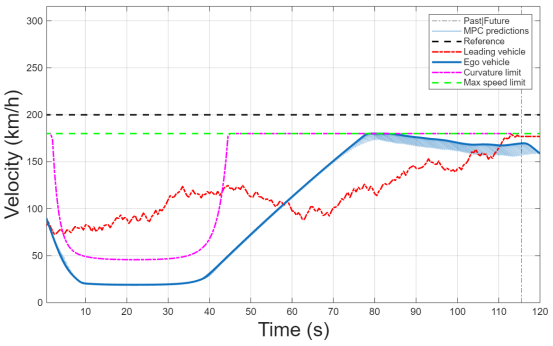

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.34: Distance and velocity tracking performance under UC5.

Figure 5.34a shows that the ego vehicle adjusts its spacing dynamically under the combined influence of the curve and the stochastic lead vehicle. As in UC4, a sharp increase in distance occurs when the ego slows down for the curve. However, compared to UC4, the spacing here is more volatile due to the additional fluctuations introduced by the random lead vehicle behaviour.

Figure 5.34b illustrates the corresponding velocity profile. The ego car preemptively reduces speed before entering the curve and remains below the curvature-imposed velocity limit. This behaviour arises for the same reasons as in UC4. After exiting the curve, the ego accelerates toward the maximum speed limit since sufficient spacing is available, but later reduces speed as the inter-vehicular gap shrinks, eventually converging to the lead vehicle's motion.

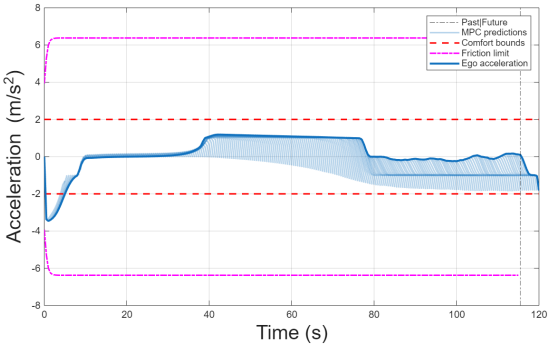

(a) Acceleration profile with MPC predictions and comfort/friction limits.

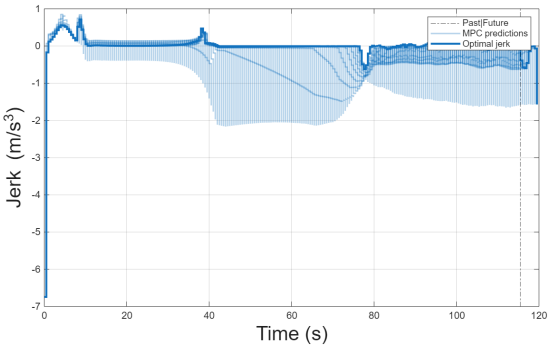

(b) Jerk profile with MPC predictions.

Figure 5.35: Acceleration and jerk response of the ego vehicle under UC5.

Figure 5.35a shows that the ego vehicle's acceleration remains within both comfort and friction bounds throughout most of the scenario. Similar to UC4, a strong deceleration is enforced at the start due to the initial ego speed and the curvature constraint. Here, however, the profile is less smooth because random lead vehicle slowdowns introduce additional braking events.

Figure 5.35b illustrates as in UC4, a sharp negative jerk which initiates the strong braking phase when approaching the curve. In UC5, however, the profile is noticeably more volatile: fluctuations appear throughout the horizon due to stochastic lead vehicle accelerations. Despite these disturbances, jerk values remain bounded and short-lived, ensuring smoothness and preserving passenger comfort.

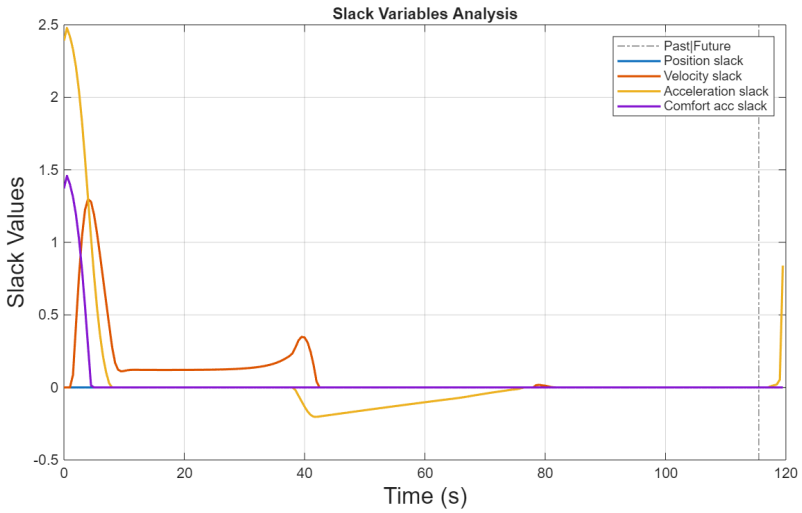

Figure 5.36: Evolution of slack variables under UC5.

Figure 5.36 shows that several slack variables are briefly activated at the beginning of the scenario. This is primarily due to the sharp braking required as the ego vehicle approaches the curve, which momentarily pushes both acceleration and comfort constraints to their limits.

The overall slack usage is slightly lower than in UC4, likely due to the ego vehicle's more moderate starting speed, giving the controller additional margin to react. However, compared to UC 1-3, the slack activations are still more pronounced, reflecting the increased challenge of simultaneously satisfying curvature and following constraints.

UC6: Low-Speed Convergence Test

UC6 initialises the ego vehicle at a low speed with a large separation from a faster, constant-speed lead vehicle on a straight road. This scenario tests the controller's speed matching, acceleration planning, and constraint handling under stochastic friction conditions at low velocities.

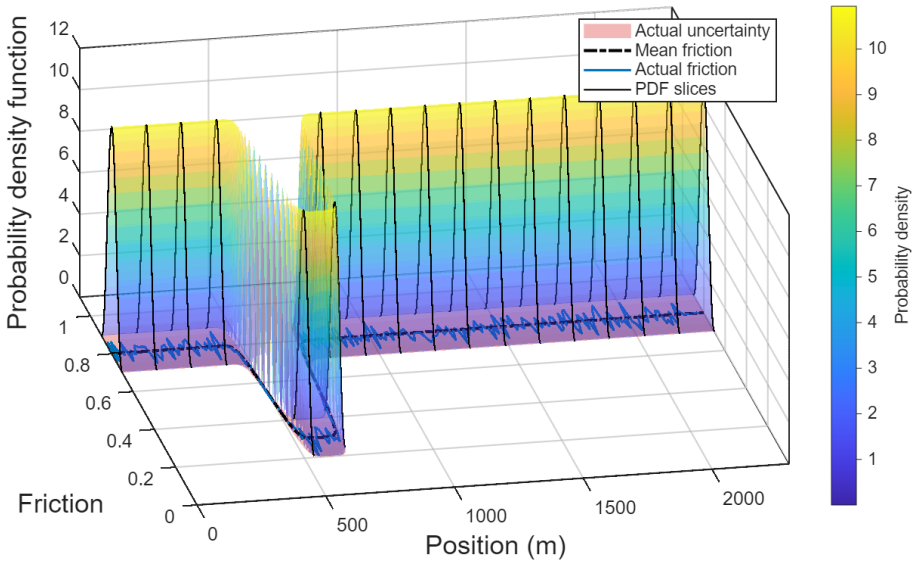

Figure 5.37: Spatial friction PDFs along the route.

Figure 5.37 provides the spatial map of road-friction uncertainty along the route.

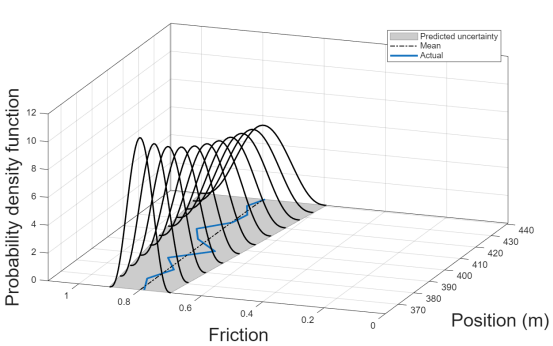

(a) Spatial snapshot of friction-prediction envelope.

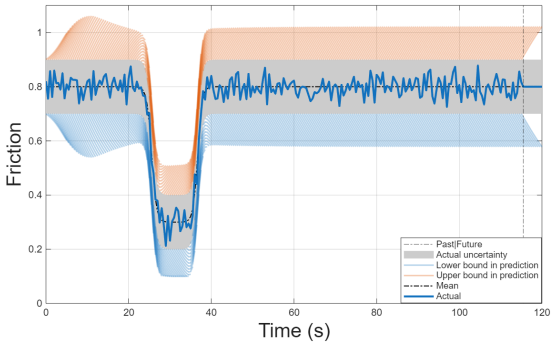

(b) Temporal evolution of friction under stochastic mode.

Figure 5.38: Friction behaviour under UC6.

This scenario exhibits the same stochastic friction behaviour as in UC2, with uncertainty modeled using Beta-distributed sampling. The friction follows the same drop-recovery pattern with stochastic deviations around the mean. Compared to UC2, the prediction bounds appear narrower here since no additional variability from the lead vehicle is present, leaving only the friction uncertainty to shape the envelope.

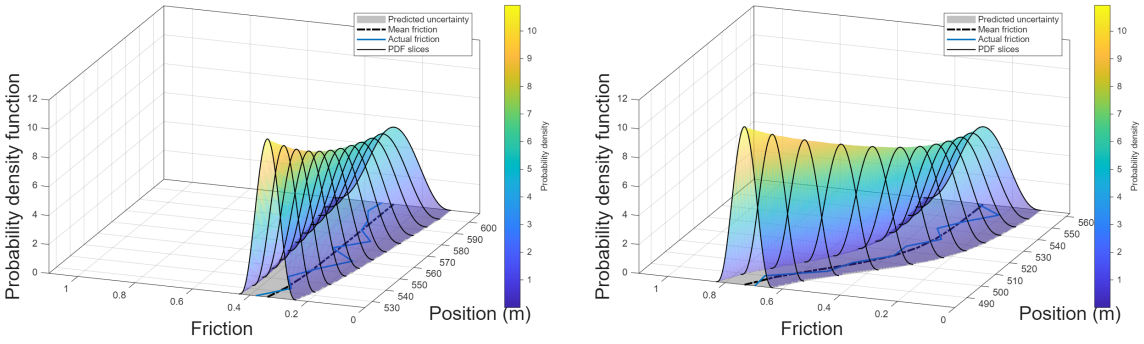

(a) **(b)**

Figure 5.39: Additional snapshots of the friction-prediction envelope

Figure 5.39 shows some additional prediction-horizon snapshots for UC6 under stochastic friction.

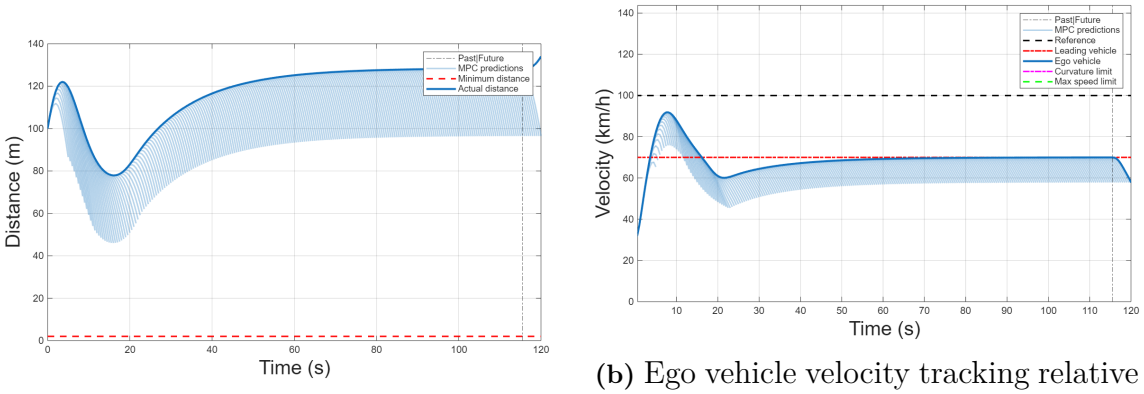

(a) Inter-vehicular distance tracking with MPC predictions.

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.40: Distance and velocity tracking performance under UC6.

Figure 5.40a shows that since the ego vehicle starts at a much lower speed, the spacing initially increases. The ego then accelerates strongly at the beginning to catch up with the lead vehicle, causing the gap to shrink. When the friction drop occurs, the ego vehicle reduces its speed temporarily to maintain safety, which leads to a stabilisation of the distance evolution. Once the speeds align again after recovery, the distance converges toward a steady value, confirming successful spacing regulation.

Figure 5.40b illustrates how the ego car accelerates at the start, as necessary to catch up from the low initial speed. A reduction in speed is then observed during the friction drop, reflecting the controller's adaptation to reduced traction. After this phase, the controller achieves smooth convergence, and the velocity aligns with that of the lead vehicle. The trajectory remains well within physical and safety limits, demonstrating stable catch-up under stochastic friction and low-speed initialisation.

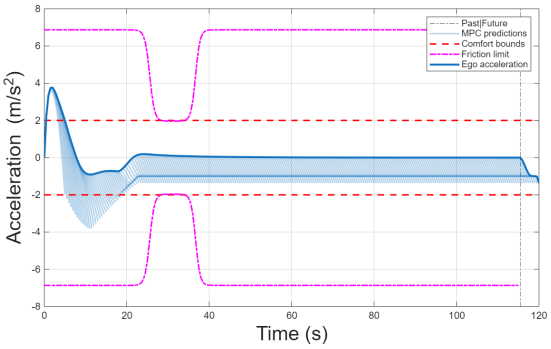

(a) Acceleration profile with MPC predictions and comfort/friction limits.

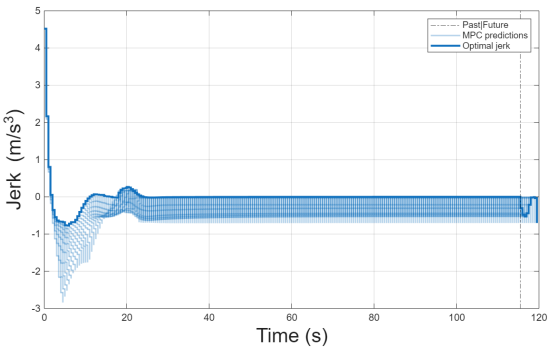

(b) Jerk profile with MPC predictions.

Figure 5.41: Acceleration and jerk response of the ego vehicle under UC6.

Figure 5.41a shows the acceleration profile, where the ego vehicle applies a strong positive acceleration at the beginning to compensate for its much lower initial speed. This transient peak is necessary to initiate catch-up with the faster lead vehicle. Once the velocity gap narrows, the acceleration quickly reduces and stabilises close to zero, remaining well within both friction and comfort bounds. No sharp braking phases are required, since the scenario primarily demands speed increase rather than deceleration.

Figure 5.41b shows the corresponding jerk trajectory. A short burst of positive jerk initiates the strong initial acceleration, followed by small transients as the controller smooths the velocity convergence. After this phase, jerk stays zero, confirming that the vehicle reaches steady following without oscillations or aggressive corrections.

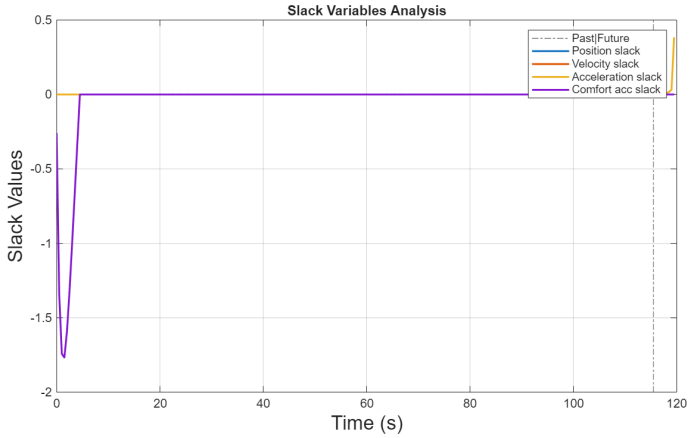

Figure 5.42: Evolution of slack variables under UC6.

Figure 5.42 shows that slack usage is also minimal throughout the scenario. A small activation appears during the initial acceleration phase as the controller adjusts to the speed mismatch.

UC7: Close-Following Robustness Scenario

UC7 stresses the controller with an aggressive lead-vehicle profile while the ego car begins with a distance gap equal to 30m. The lead vehicle executes a bounded random acceleration profile $\pm 3\,\mathrm{m/s^2}$ that repeatedly demands both hard braking and rapid catch-up manoeuvres from the ego vehicle. Road is straight but the pavement undergoes an early ice patch ($\mu \approx 0.25$) followed by recovery to high friction. Stochastic friction prediction is therefore enabled. Key parameters are summarised in Table B.13.

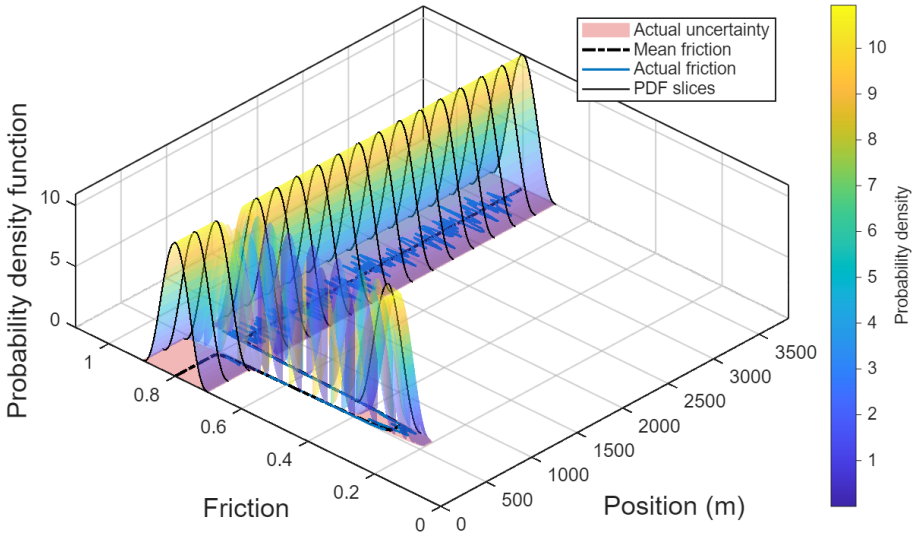

Figure 5.43: Spatial friction PDFs along the route.

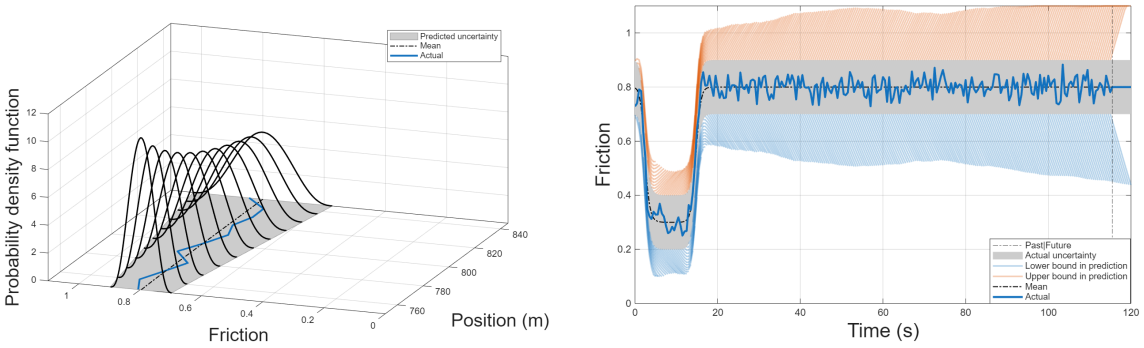

(a) Spatial snapshot of friction-prediction envelope.

(b) Temporal evolution of friction under stochastic mode.

Figure 5.44: Friction behaviour under UC7.

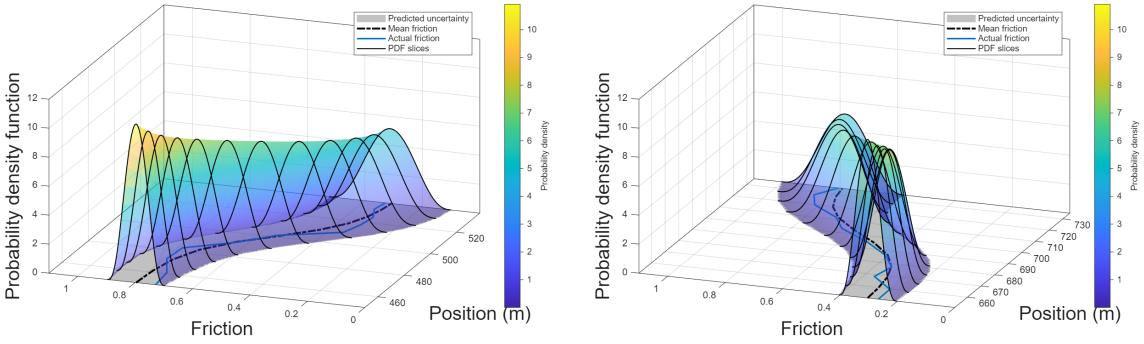

(a) **(b)**

Figure 5.45: Additional snapshots of the friction-prediction envelope

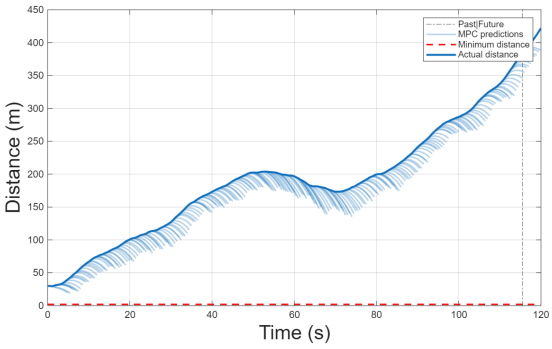

(a) Inter-vehicular distance tracking with MPC predictions.

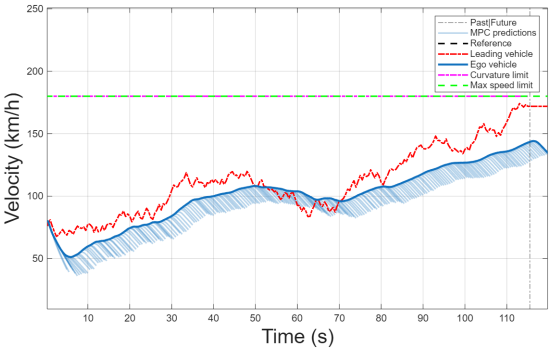

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.46: Distance and velocity tracking performance under UC7.

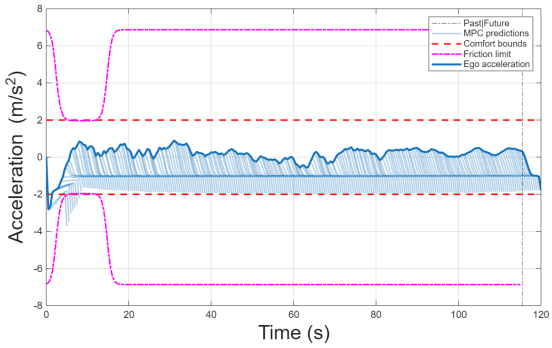

(a) Acceleration profile with MPC predictions and comfort/friction limits.

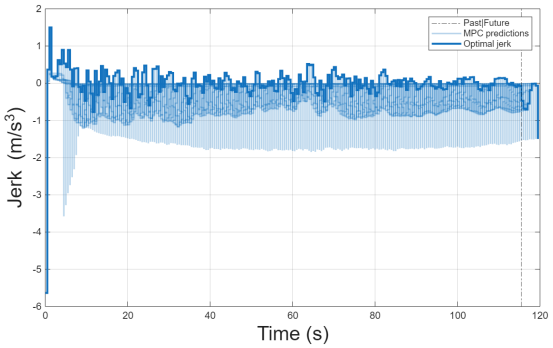

(b) Jerk profile with MPC predictions.

Figure 5.47: Acceleration and jerk response of the ego vehicle under UC7.

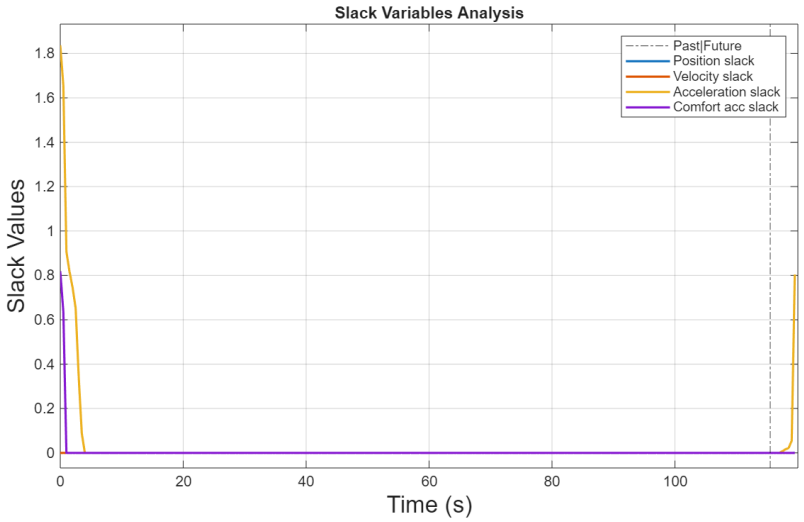

Figure 5.48: Evolution of slack variables under UC7.

Despite the lead vehicle's unpredictable bursts of acceleration in Figure 5.46b, the ego car never infringes the minimum distance (Fig. 5.46a), and without any aggressive reactions. The RNMPC reaches a gap to nearly $200m$ which is then further increasing to 400m. While friction is uncertain at approximately ($t < 15$s), it gradually re-optimises once grip recovers. Acceleration stays within $\pm 1.2\,\mathrm{m/s^2}$ for most of the run (Fig. 5.47a), and jerk remains below $\pm 1\,\mathrm{m/s^3}$ after the first 10 s in Figure 5.47b. Slack variables in Figure 5.48 show only slight comfort relaxation during the initial hard brake; safety constraints are never violated. Overall, UC7 demonstrates that the stochastic RNMPC maintains collision avoidance and ride comfort even when starting from a minimal headway and following a lead vehicle with highly erratic behaviour under rapidly changing friction conditions.

UC8: Deterministic vs Stochastic Friction Comparison

UC8 repeats an identical straight-road car-following task in two modes: **Deterministic mode** - follows the mean friction profile. **Stochastic mode** - actual friction is sampled from Beta distribution respecting uncertainty bounds while MPC's planning is using the same conservative bounds in both cases. Key parameters are summarised in Table B.14. The goal of the experiment is to expose how explicit uncertainty modelling affects safety margins, control effort and comfort.

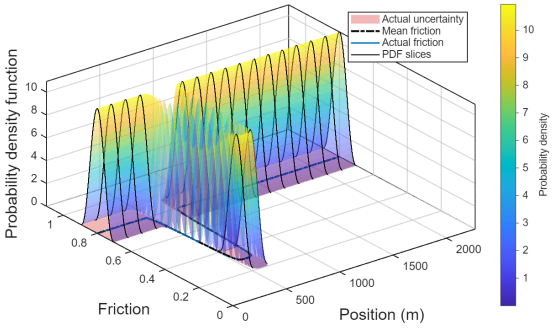

(a) Spatial friction PDFs along the route - **deterministic**.

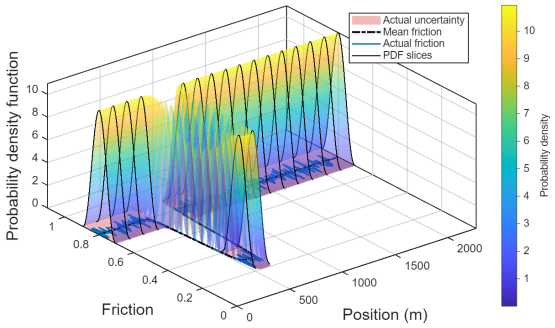

(b) Spatial friction PDFs along the route - **stochastic**.

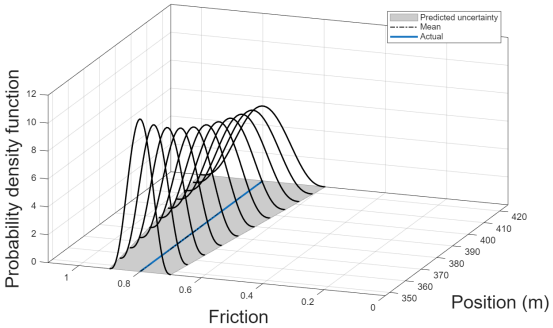

(a) Spatial snapshot of friction-prediction envelope **deterministic**.

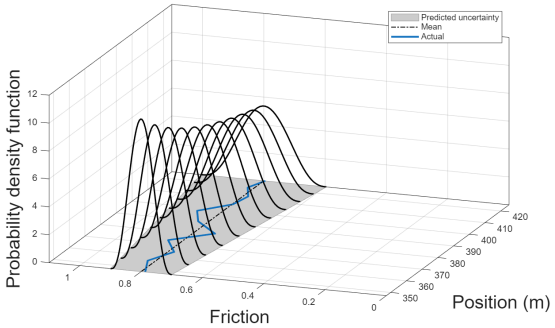

(b) Spatial snapshot of friction-prediction envelope - **stochastic**.

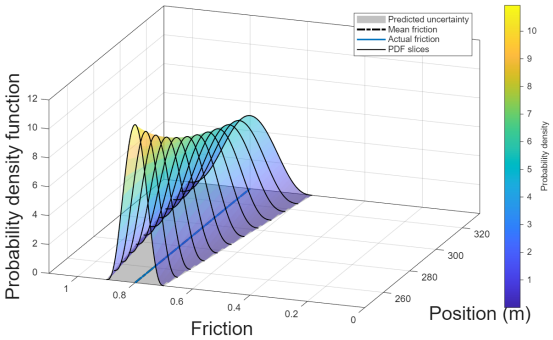

(a) Additional spatial snapshot of friction-prediction envelope **deterministic**.

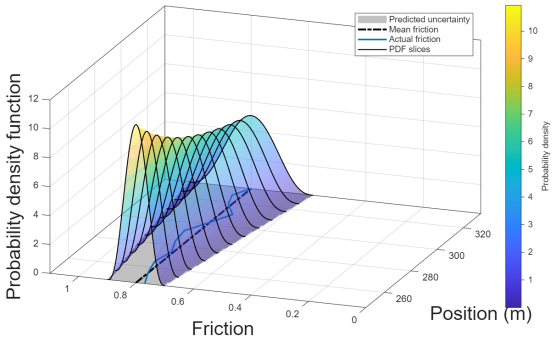

(b) Additional spatial snapshot of friction-prediction envelope - **stochastic**.

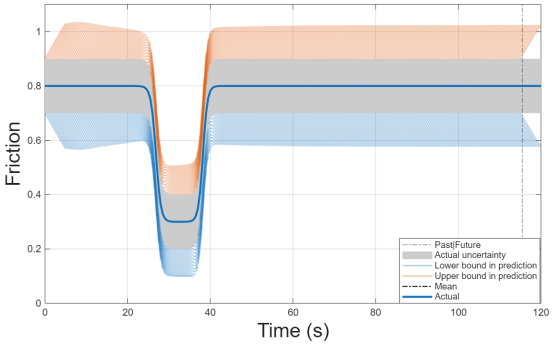

(a) Temporal friction evolution - deterministic.

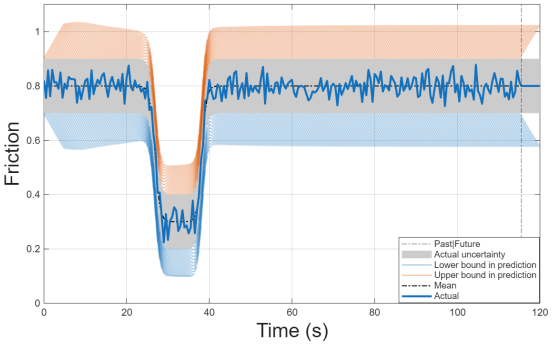

(b) Temporal friction evolution - stochastic.

Figure 5.52: Friction envelopes and ground-truth for UC8.

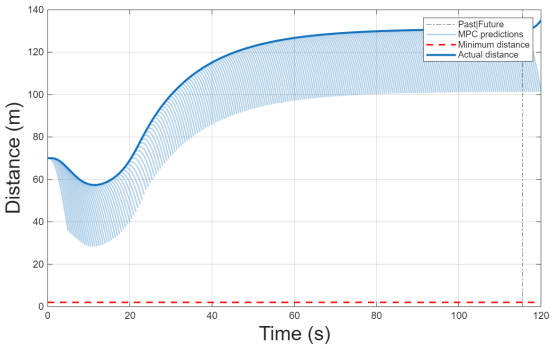

(a) Inter-vehicular distance - deterministic.

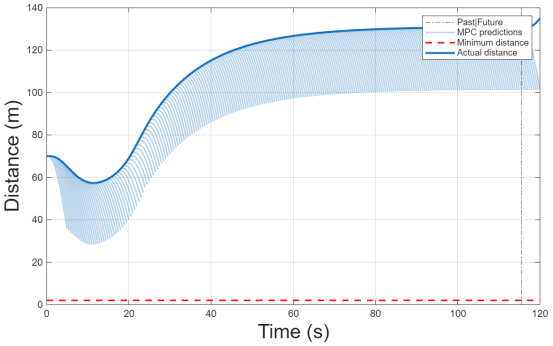

(b) Inter-vehicular distance - stochastic.

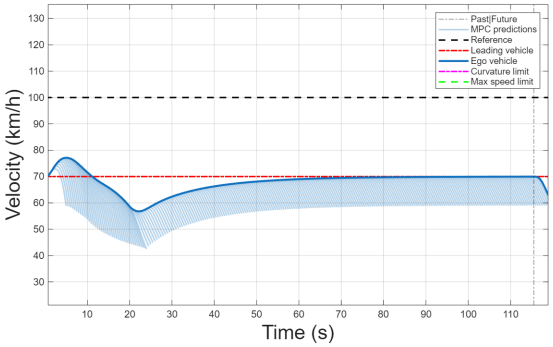

(a) Velocity - deterministic.

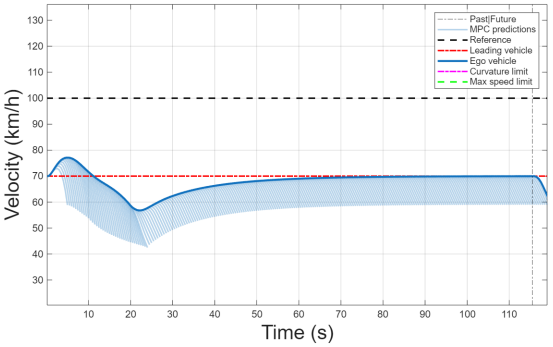

(b) Velocity - stochastic.

Figure 5.54: Distance-keeping and speed tracking for UC8.

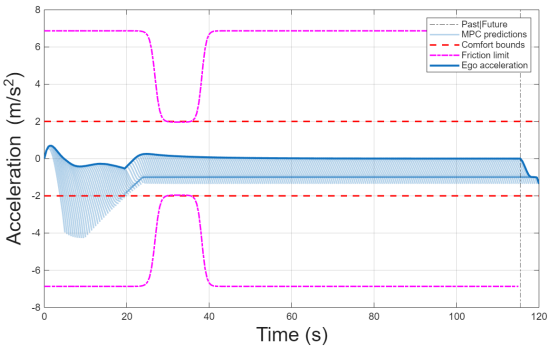

(a) Acceleration - deterministic.

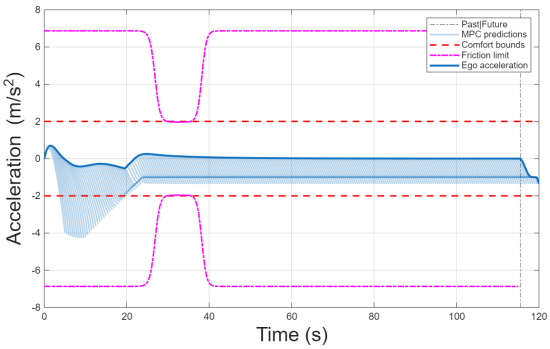

(b) Acceleration - stochastic.

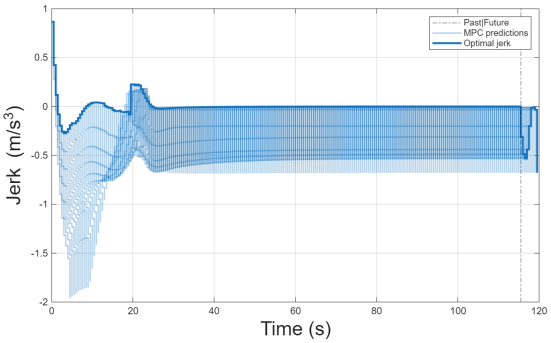

(a) Jerk - deterministic.

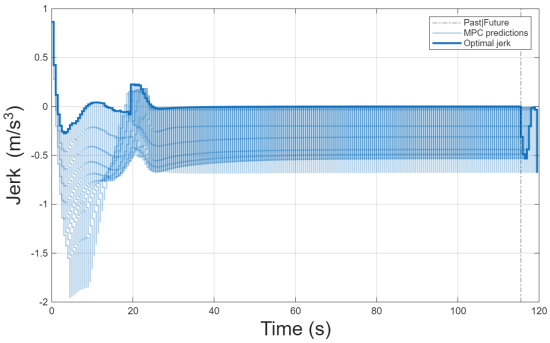

(b) Jerk - stochastic.

Figure 5.56: Control effort comparison for UC8.

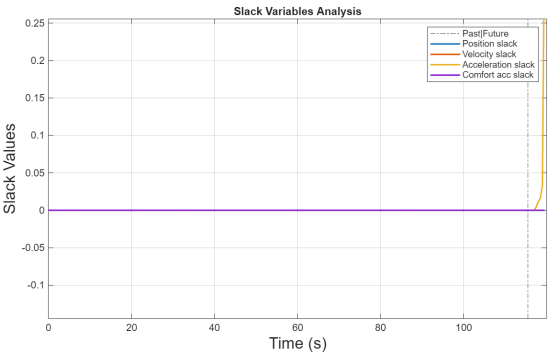

(a) Slack variables - deterministic.

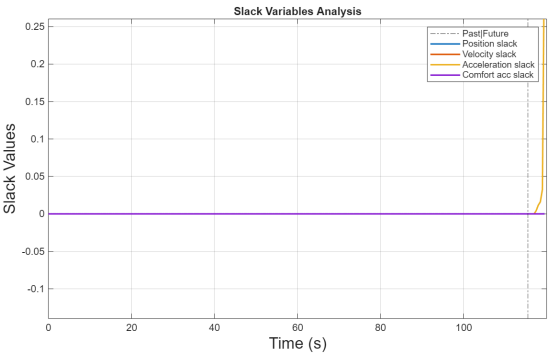

(b) Slack variables - stochastic.

Figure 5.57: Constraint-relaxation demand for UC8.

Key observations drawn from the side-by-side comparison reveal that the distance evolution, velocity profiles, acceleration, and jerk trajectories are essentially identical between both operational modes (Figures 5.53a-5.56b). This result validates a fundamental aspect of the control design, and that is because the MPC uses the same conservative uncertainty bounds for planning in both cases, regardless of the actual friction realisation.

In deterministic mode, the actual friction perfectly follows the assumed mean profile, yielding a narrow prediction envelope (Figures 5.50a-5.52a). When stochastic sampling is enabled (Figures 5.50b-5.52b), the Beta-distributed friction realisations remain well within the predicted uncertainty bounds, validating the probabilistic modeling framework. The identical trajectories demonstrate that the controller's conservative design, by for example using lower friction bounds (μ_{low}) for ego vehicle constraints. This, provides sufficient safety margins such that stochastic friction variations within the uncertainty envelope require no additional control adjustments.

Slack variable usage remains negligible (no activation) in both modes (Figure 5.57), indicating that the uncertainty bounds are appropriately sized and the controller maintains feasibility without requiring constraint relaxation, even when actual friction deviates from the mean. This behaviour confirms that the uncertainty-aware MPC framework successfully decouples planning conservatism from execution variability, maintaining consistent control behaviour across different friction realisations within the uncertainty envelope.

This demonstrates that the probabilistic uncertainty framework provides robust and predictable control behaviour. The system does not experience unnecessary sensitivity to friction variations that fall within the predicted bounds, avoiding over-reactive control adjustments that could compromise comfort while preserving the safety margins essential for practical deployment where actual road conditions may vary within sensor confidence bounds.

UC9: High Curvature Scenario with Standard Friction

UC9 evaluates ACC robustness with extreme initial conditions, and with an ego vehicle starting at 150 km/h approaching a much slower constant-speed lead vehicle (50 km/h) with 120m initial separation on straight road. Challenges the prediction horizon adequacy, constraint feasibility and deceleration capabilities under stochastic friction conditions.

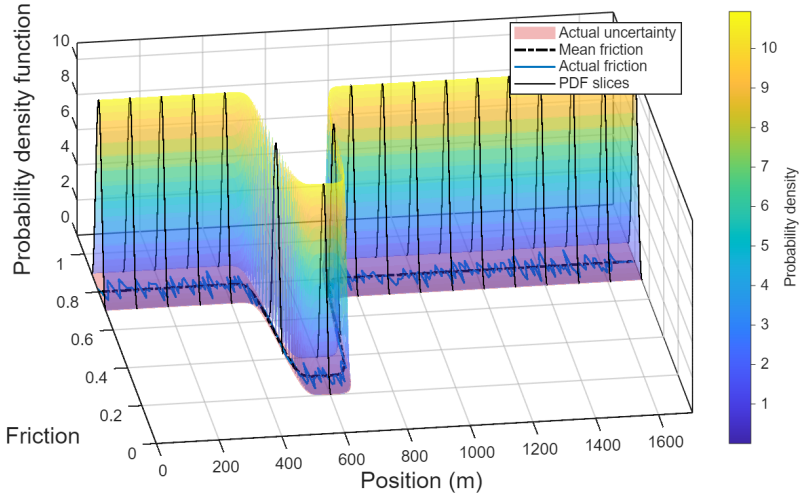

Figure 5.58: Spatial friction PDFs along the route.

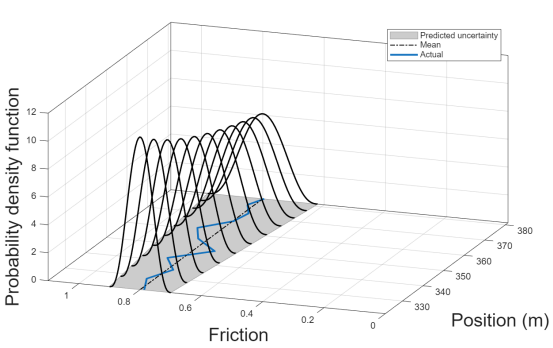

(a) Spatial snapshot of friction-prediction-envelope.

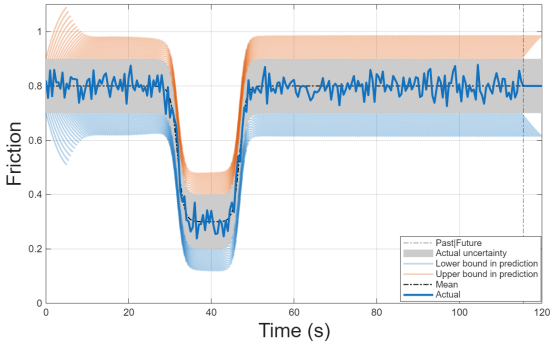

(b) Temporal evolution of friction under stochastic mode.

Figure 5.59: Friction behaviour under UC9.

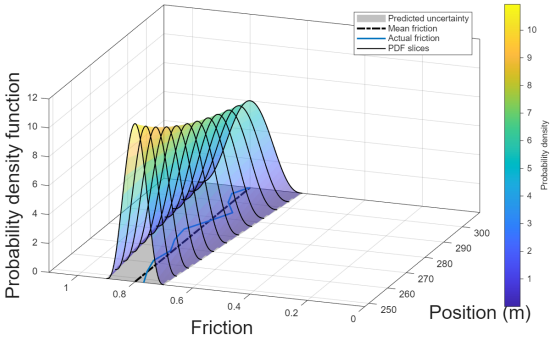

(a)

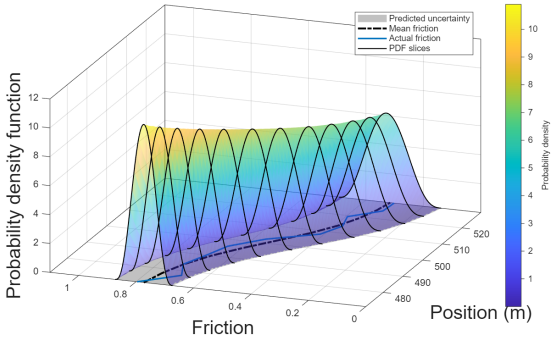

(b)

Figure 5.60: Additional snapshots of the friction-prediction envelope

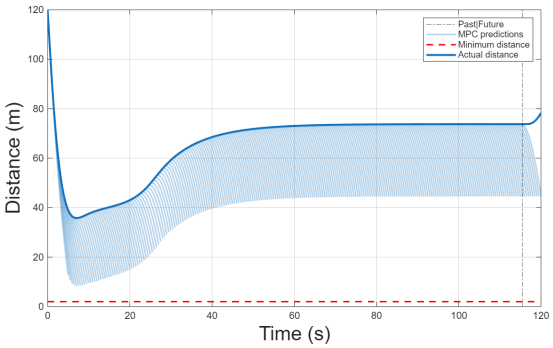

(a) Inter-vehicular distance tracking with MPC predictions.

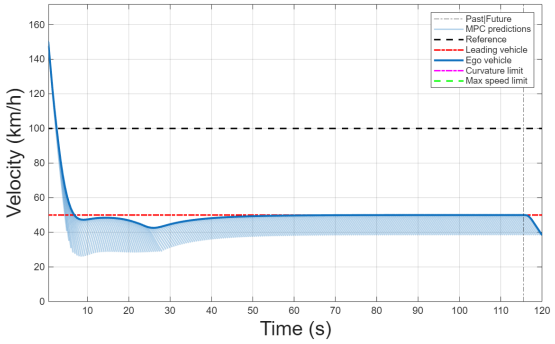

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.61: Distance and velocity tracking performance under UC9.

Distance management (Figure 5.61a) adjusts smoothly, initially decreasing to approximately 40m due to the extreme initial speed and then stabilising effectively.

The velocity profile (Figure 5.61b) shows proactive and significant deceleration in response to the lead car's speed, stabilising at around 30s near the lead vehicle's speed, after it maintained again a safe distance.

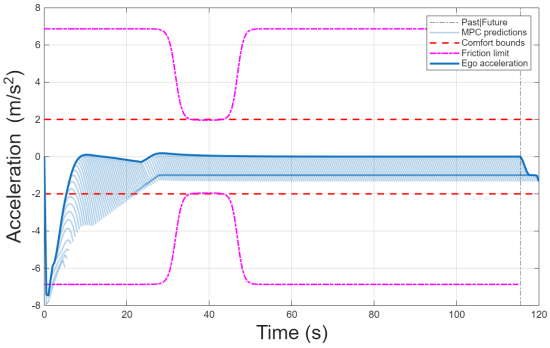

(a) Acceleration profile with MPC predictions and comfort/friction limits.

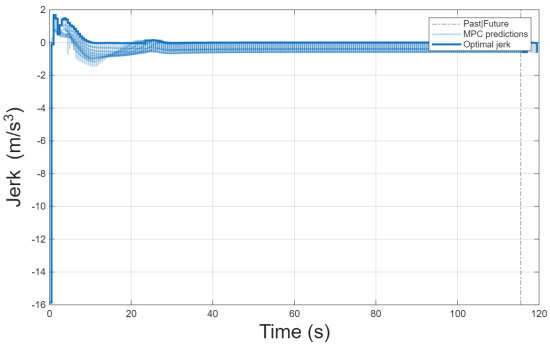

(b) Jerk profile with MPC predictions.

Figure 5.62: Acceleration and jerk response of the ego vehicle under UC9.

The acceleration profile (Figure 5.62a) remains comfortably within predefined limits, demonstrating controlled deceleration. The jerk profile (Figure 5.62b) indicates gentle transitions with minimal spikes, reflecting smooth and stable control actions.

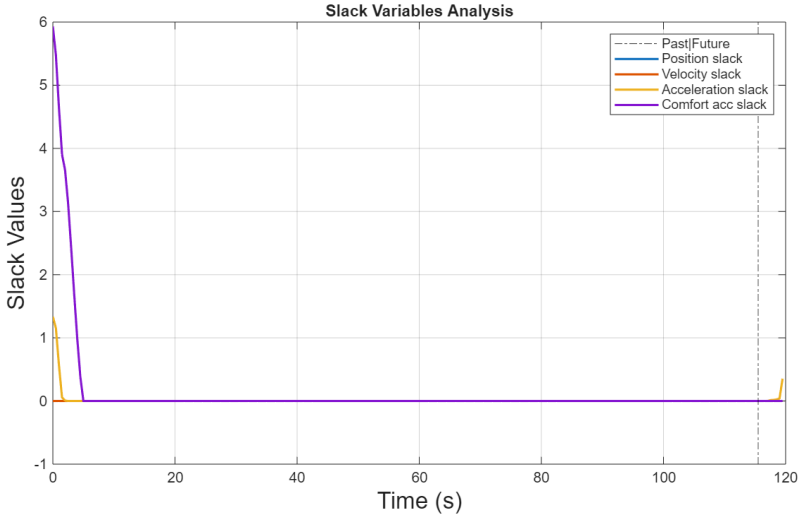

Figure 5.63: Evolution of slack variables under UC9.

Slack analysis (Figure 5.63) displays brief initial activations of comfort and velocity slack variables, highlighting the controller's effectiveness in managing constraints primarily at the start. Overall, UC9 confirms the MPC's capability in handling high-curvature scenarios, ensuring stability, comfort, and safety consistently.

UC10: Standard Preview Distance with Stochastic Friction

UC10 evaluates the MPC controller using stochastic friction uncertainty and the standard friction preview distance (150m). The ego vehicle starts with a traveling speed at 100 km/h and 80m gap to a leading vehicle traveling at constant speed. Key parameters are summarised in Table B.16

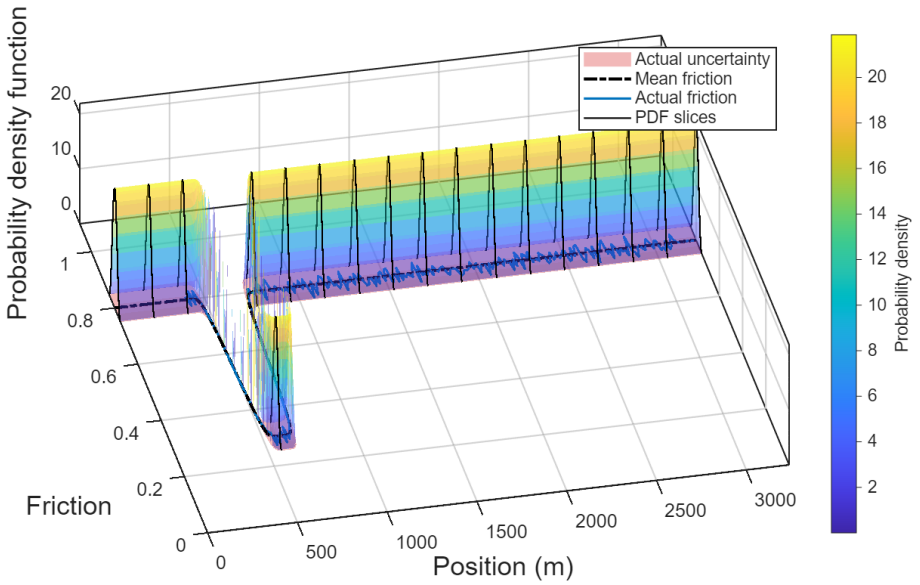

Figure 5.64: Spatial friction PDFs along the route.

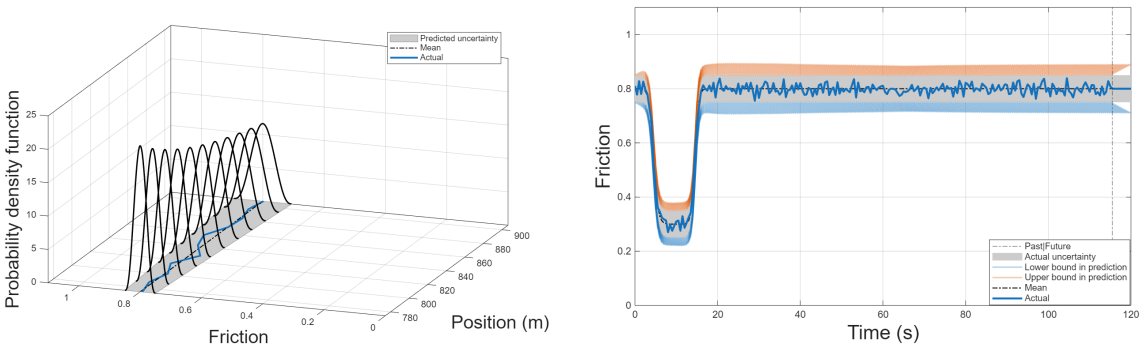

(a) Spatial snapshot of friction-prediction envelope.

(b) Temporal evolution of friction under stochastic mode.

Figure 5.65: Friction behaviour under UC10

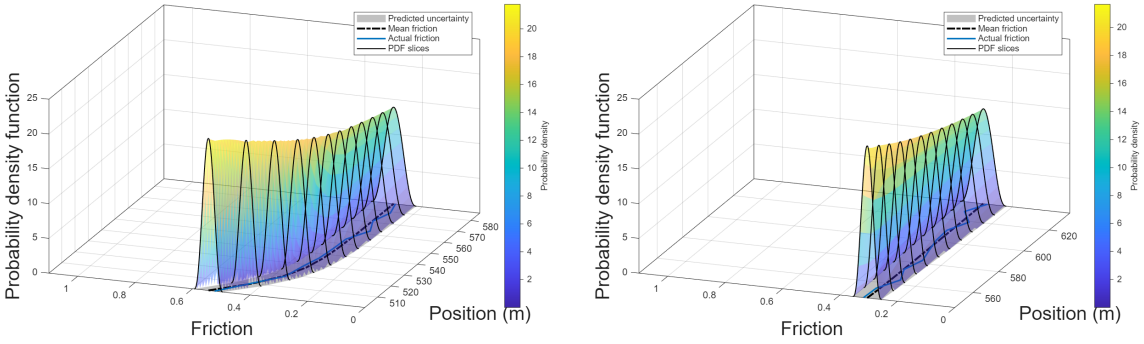

(a) **(b)**

Figure 5.66: Additional snapshots of the friction-prediction envelope

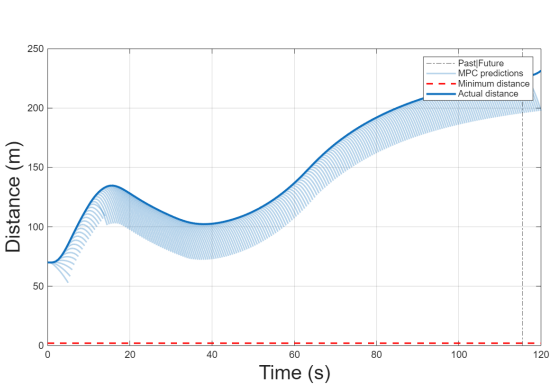

(a) Inter-vehicular distance tracking with MPC predictions.

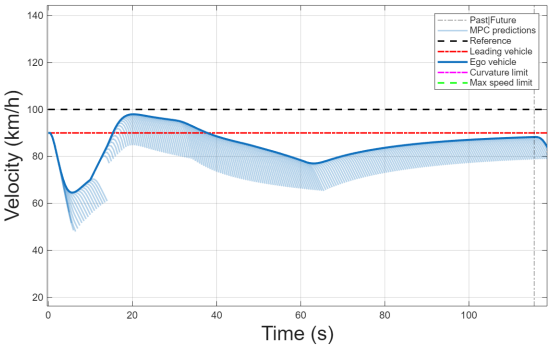

(b) Ego vehicle velocity tracking relative to the lead vehicle, including MPC predictions and constraint limits.

Figure 5.67: Distance and velocity tracking performance under UC10.

Figure 5.67a shows the controller effectively maintaining and smoothly adjusting the safe distance gap. The velocity profile (Figure 5.67b) indicates proactive and smooth velocity adjustments in response to friction variations. In addition to this, it can also be observed that due to the very demanding requirements at the beginning, involving sudden friction drop and initial velocity with respect to the safe separation distance required, the controller it is trying to increase the distance gap. Due to the followed restoration of the friction realisation to ideal dry conditions, the gap can then be decreased; therefore, the controller is then prioritising the driver-imposed speed for a short amount of time, due to reaching the optimal safe distance gap.

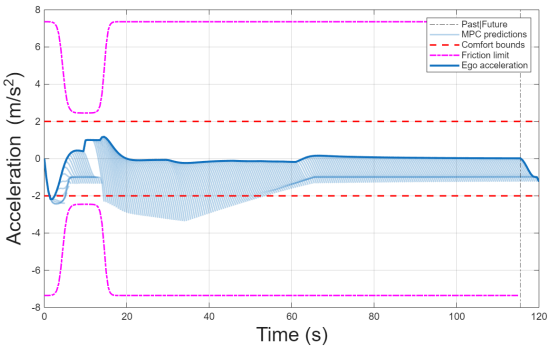

(a) Acceleration profile with MPC predictions and imposed comfort/friction limits.

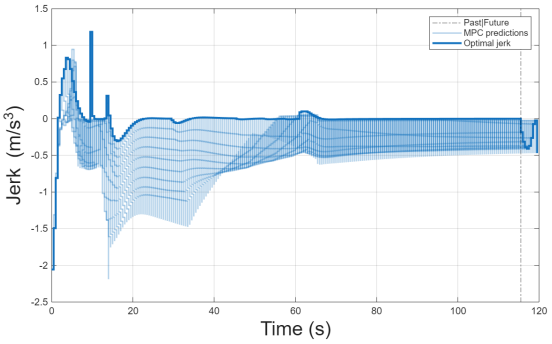

(b) Jerk profile with MPC predictions.

Figure 5.68: Acceleration and jerk response of the ego vehicle under UC10.

The acceleration profile (Figure 5.68a) illustrates smooth adjustments within comfort and friction limits, while the jerk profile (Figure 5.68b) shows controlled transitions, especially around friction variations.

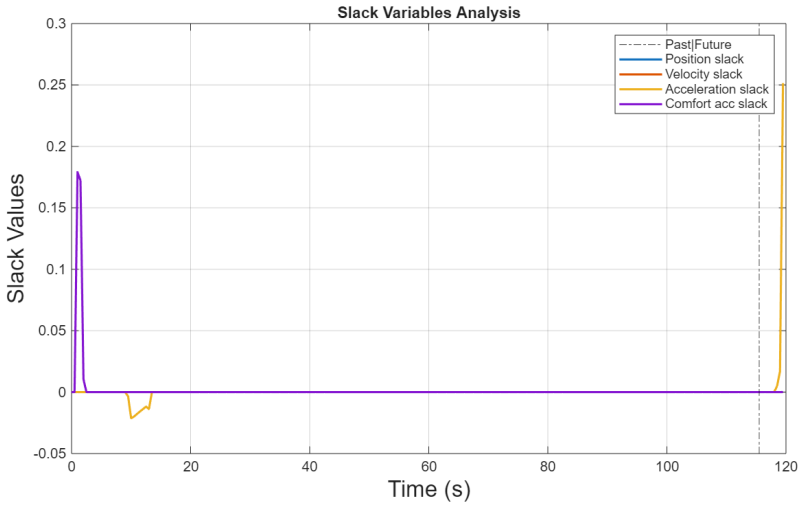

Figure 5.69: Evolution of slack variables under UC10.

Slack analysis (Figure 5.69) indicates minor comfort constraint relaxations, mainly at initial acceleration, confirming MPC's effectiveness in maintaining safety constraints. Overall, UC10 demonstrates robust MPC performance, effectively managing stochastic friction uncertainty while maintaining safe and comfortable driving behaviour.

5.3 Additional Scenario Results

AS1: Driver-Imposed Speed Limitation

The ego vehicle is instructed to stay at 80 km/h even though the lead car cruises at 100km/h. The road is essentially straight and features a short dry-wet-dry friction dip (see Fig. 5.70).

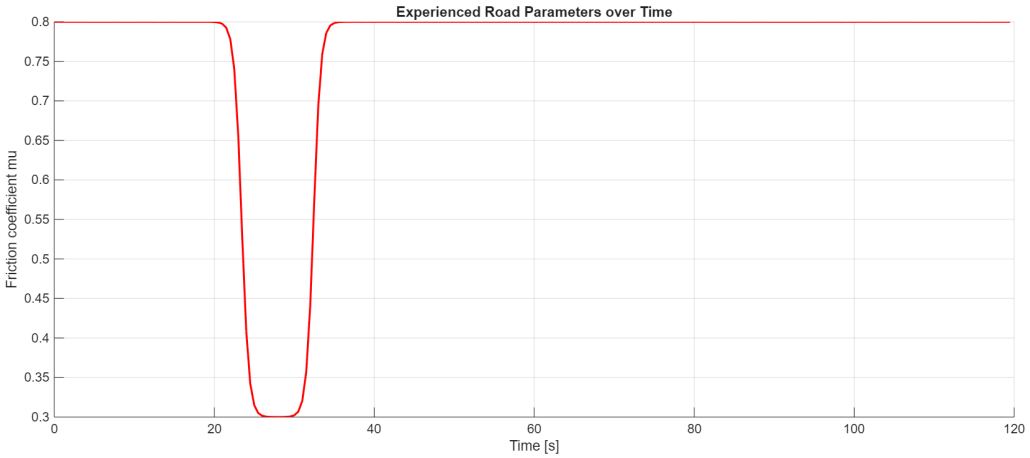

Figure 5.70: Friction envelope for AS1.

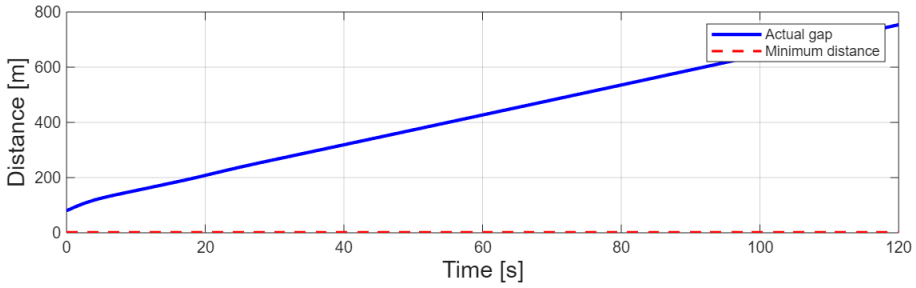

Figure 5.71: Inter-vehicular distance.

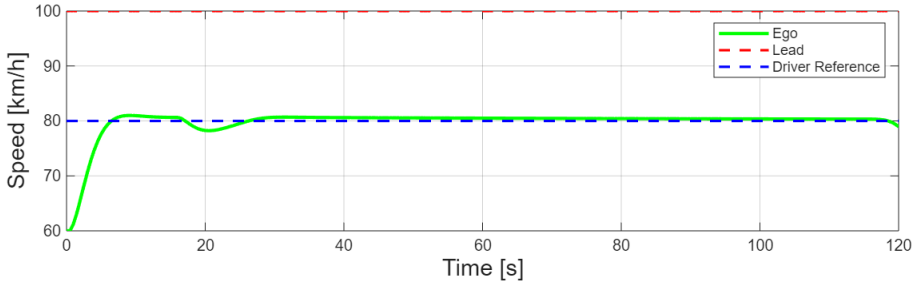

Figure 5.72: Ego and lead speed.

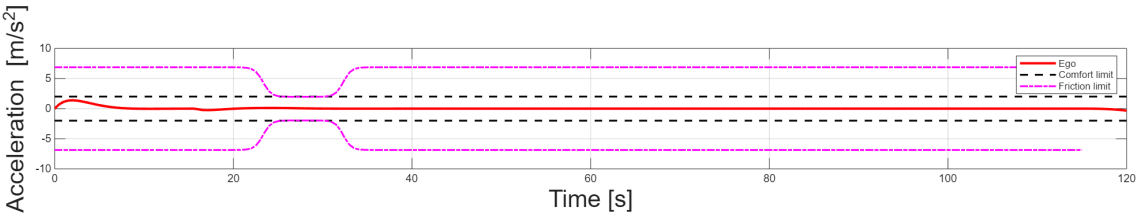

Figure 5.73: Ego acceleration versus comfort and friction limits.

In this scenario, the ego vehicle is instructed to maintain the driver-imposed speed of 80 km/h. Since the lead vehicle continues at a higher constant velocity, the inter-vehicular gap increases steadily and approximately linearly over time. The longitudinal acceleration remains comfortably within the $\pm 2 m/s^2$ comfort band. No slack variables are activated, confirming strict feasibility of the optimisation. This demonstrates the controller's ability to respect driver-defined limits even when external conditions would otherwise encourage higher speeds.

AS2: Progressive Friction Degradation

Starting from dry asphalt, the available grip degrades in two steps ($\mu : 0.8 \rightarrow 0.5 \rightarrow 0.2$) as the vehicle reaches 400 m and 700 m respectively; see Fig. 5.74. The lead car keeps a steady 80 km/h while the ego must adapt to the diminishing friction.[2pt]

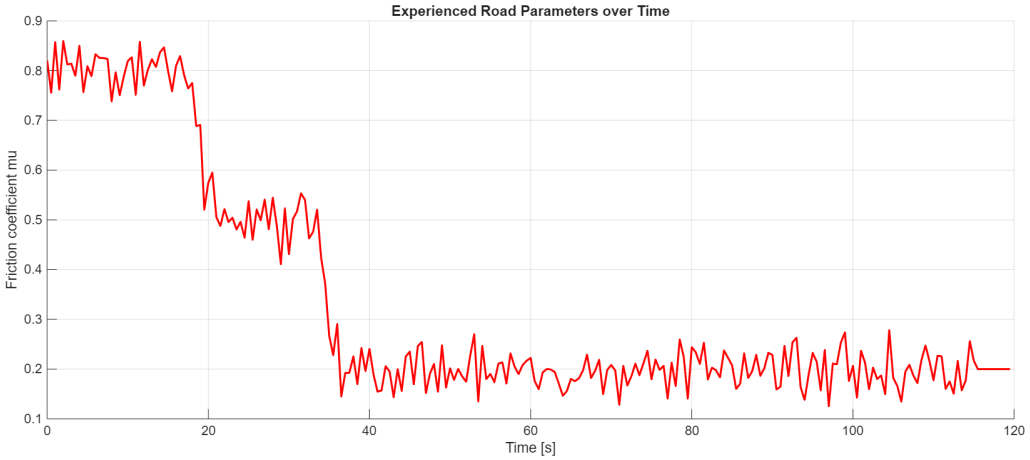

Figure 5.74: Time history of the experienced friction in AS2.

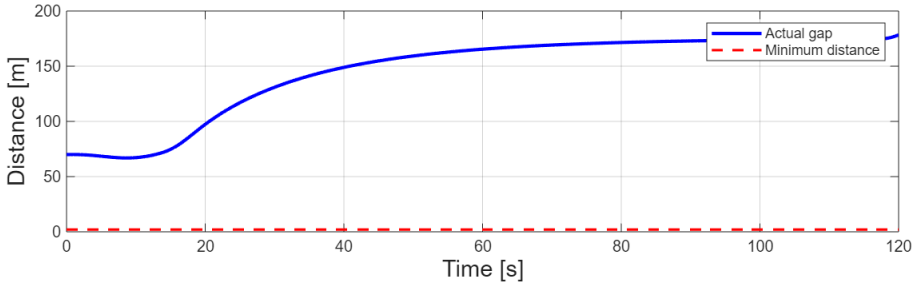

Figure 5.75: Inter-vehicular distance.

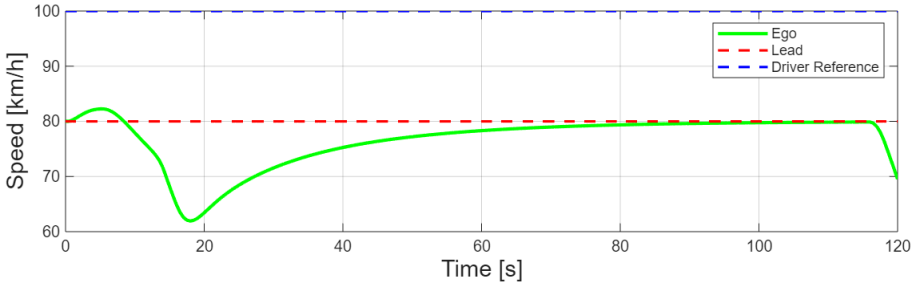

Figure 5.76: Ego and lead speed.

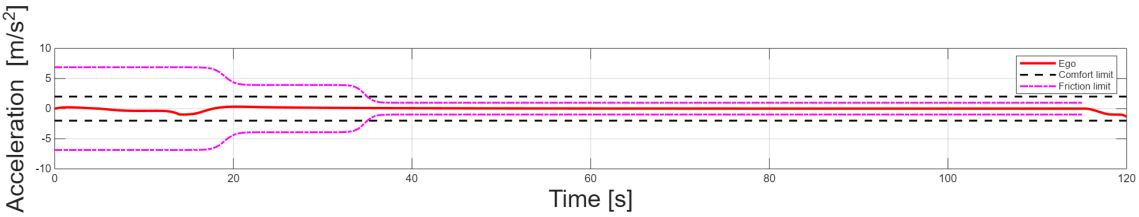

Figure 5.77: Ego acceleration versus comfort and friction-based limits.

When the first drop to $\mu = 0.5$ occurs, the ego vehicle reduces its speed to around 72 km/h. The second drop to $\mu = 0.2$ produces an even steeper deceleration, bringing the velocity down to about 62 km/h. After this transient, the ego gradually increases its speed again until it aligns with the lead vehicle at 80 km/h. As a result, the inter-vehicular distance grows steadily and stabilises around 180 m, maintaining a wide and safe margin.

The acceleration profile confirms this smooth behaviour: both deceleration phases are well contained within the ± 2 m/s^2 comfort band, and no slack variables are used. This shows that the controller can handle progressive friction degradation while keeping manoeuvres safe and comfortable

AS3: Curve Coincident with Low Friction

A moderate left-hand curve of approx. 20m radius is deliberately aligned with the worst part of the friction dip ($\mu \approx$ 0.1-0.2). Both effects are captured in Fig. 5.78 where curvature (blue, left axis) and the realised stochastic friction trace (red, right axis) are plotted together.

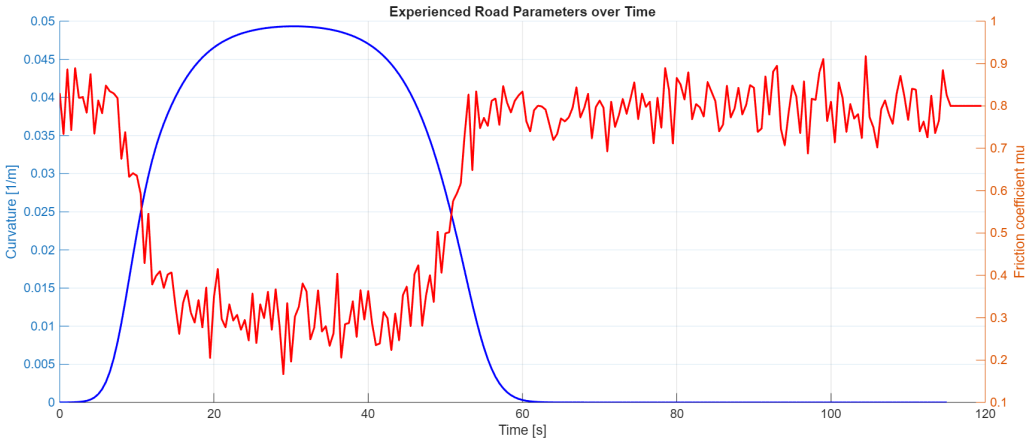

Figure 5.78: Combined curvature and friction envelope for AS3.

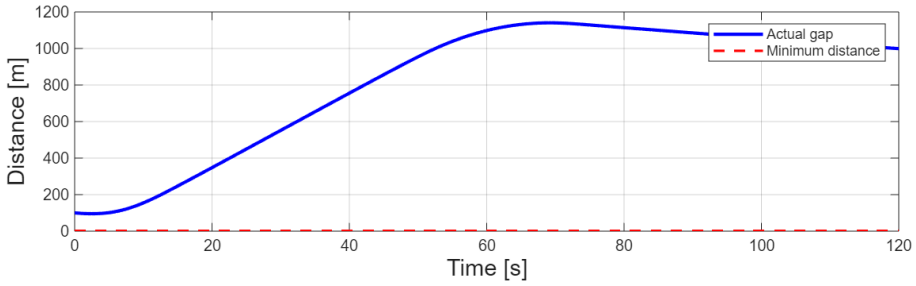

Figure 5.79: Inter-vehicular distance.

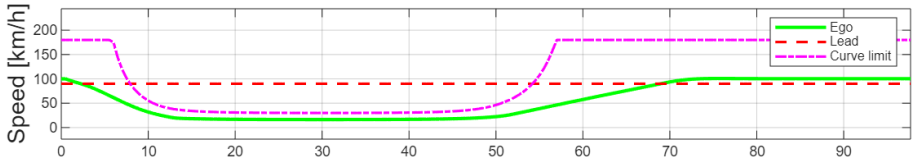

Figure 5.80: Ego and lead speed.

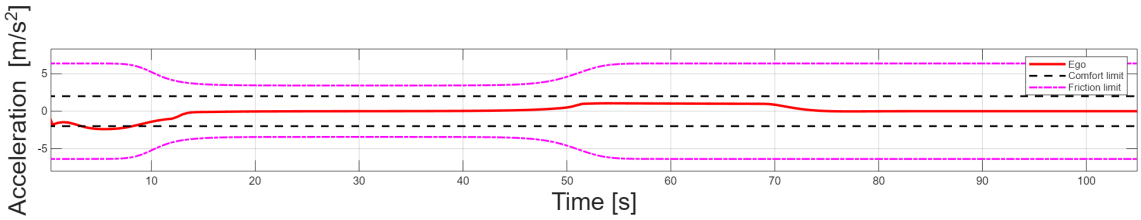

Figure 5.81: Ego acceleration versus comfort and friction-based limits.

In this scenario, a tight curve coincides with the lowest friction dip, creating the most demanding conditions tested. As curvature rises and grip collapses, the controller brakes sharply, reducing ego speed to about 20km/h. Meanwhile, the lead vehicle maintains a constant pace, so the inter-vehicular gap grows rapidly. This widening

is a direct result of the ego being forced to respect the stricter curve limit while the lead car continues unhindered.

Once the curve ends and friction recovers, the ego accelerates decisively. Since the spacing remains ample, the controller allows the vehicle to approach the driver-imposed limit of 100km/h in this case, consistent with the behaviour also observed in UC4 and UC5. The acceleration profile shows a strong but feasible braking phase followed by smooth recovery, always within the ± 2 m/s^{2}comfort band. No slack variables are triggered, confirming that the MPC can safely satisfy simultaneous curvature and friction constraints.

AS4: Extended Low-Friction Region

A 400 m-long stretch is assigned a low friction level ($\mu \approx 0.25$). The realised stochastic trace is shown in Fig. 5.82.

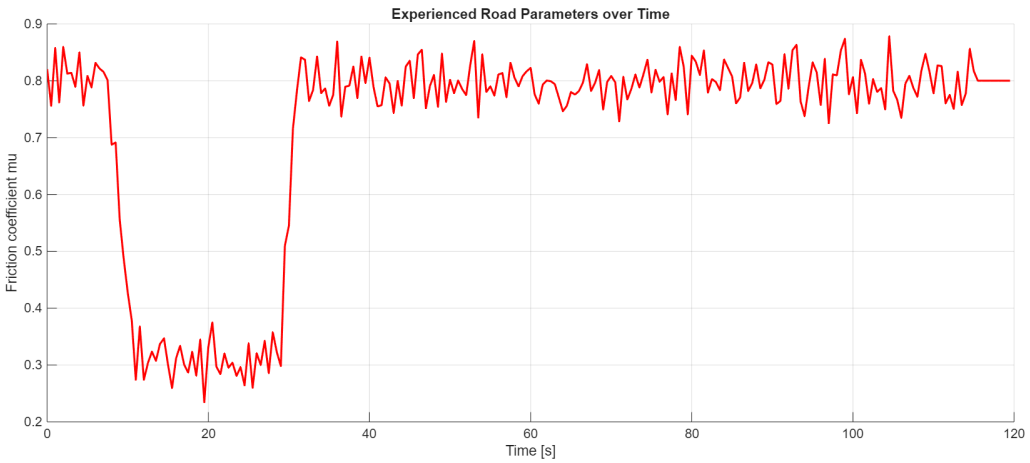

Figure 5.82: Friction envelope.

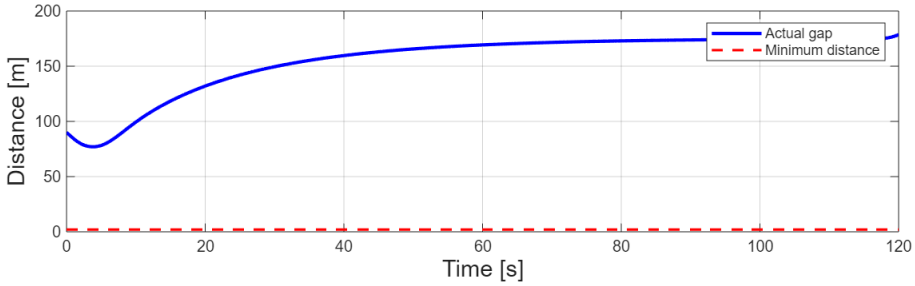

Figure 5.83: Inter-vehicular distance.

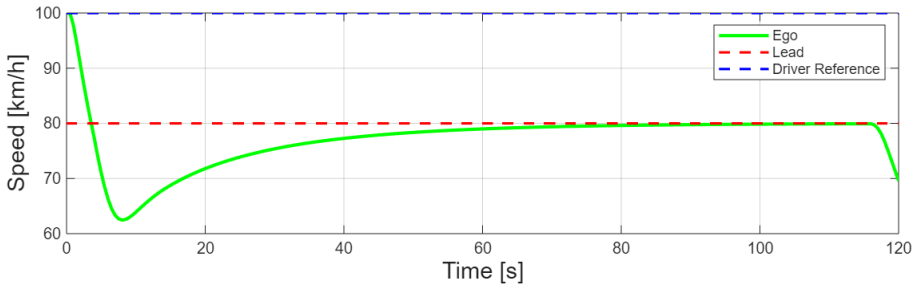

Figure 5.84: Ego and lead speed.

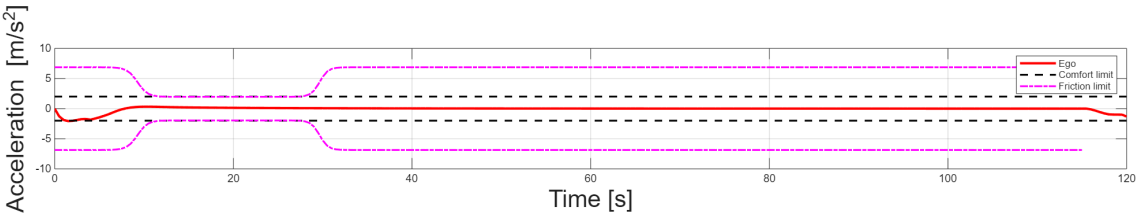

Figure 5.85: Ego acceleration versus comfort and friction-based limits.

The MPC reacts immediately to the drop in grip: the ego speed falls sharply from 100 km/h to around 62 km/h within the first 15-20 s. Because the lead vehicle continues steadily at 80 km/h, the inter-vehicular spacing increases throughout the low-friction segment and stabilises around 180 m.

Once the ego exits the slippery stretch, it accelerates back toward the lead-vehicle speed. The recovery is smooth and does not overshoot, confirming that the controller balances caution with efficiency.

Acceleration remains consistently within the ± 2 m/s^2 comfort range, and the friction limits are never violated. Also, no slack variables are activated.

AS5: High-Speed Curve Entry with Mid-Curve Friction Change

The ego vehicle approaches a moderate curve ($\kappa_{\max} = 0.025$ rad/m) at 120 km/h while a dry-wet-dry transition is placed inside the turn apex. Figure 5.86 highlights the simultaneous drop in available grip and the build-up of lateral curvature.

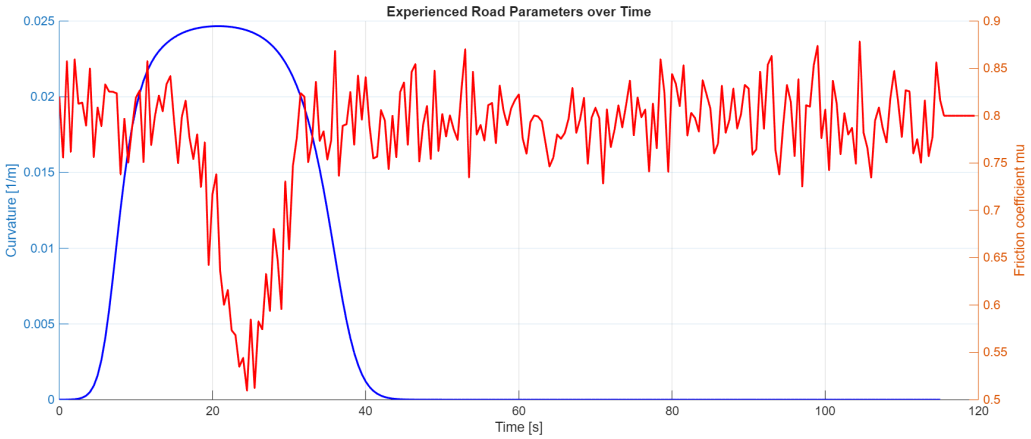

Figure 5.86: Curvature (blue, left axis) and friction trace (red, right axis). The wet patch coincides with the highest lateral load.

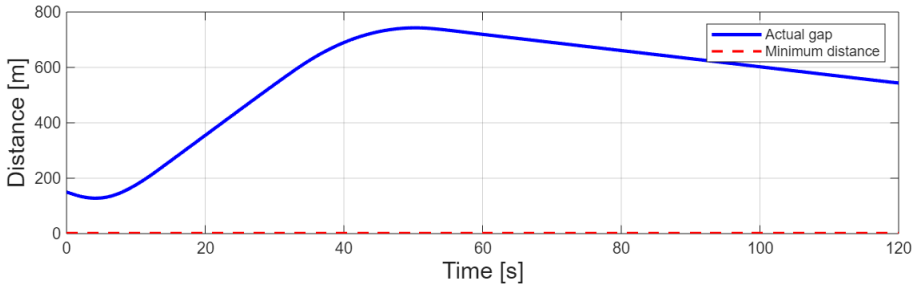

Figure 5.87: Inter-vehicular distance.

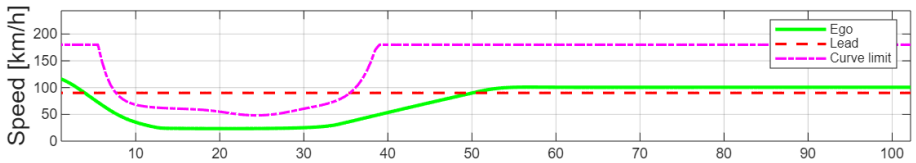

Figure 5.88: Ego and lead speed.

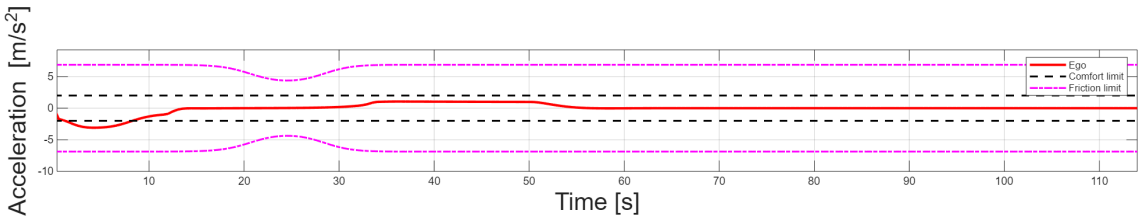

Figure 5.89: Ego acceleration compared with comfort and friction.

The ego enters the curve at a high initial speed of 120 km/h. A sharp speed reduction occurs at the beginning as the controller reacts to the build-up of curvature. When the mid-curve friction drop is introduced, the curvature-based limit decreases

further, but because the trajectory is planned conservatively, the ego is already slowed to a sufficiently low speed.

Since the lead vehicle maintains a constant velocity without accounting for curvature, the inter-vehicular gap grows rapidly during this phase. After exiting the curve, the ego accelerates back toward the driver-imposed reference of 100 km/h. The large spacing allows this higher cruising speed to be maintained, consistent with the behaviour observed in AS3.

The acceleration profile shows a short violation of the comfort bound during the initial braking phase. This is attributed to the combined effect of high entry speed, strong curvature, and the sudden friction change, but stability and safety are preserved throughout.

AS6: Ice Patch on a Curve

In this final scenario the ego car enters a 40 m-radius curve ($\kappa_{\max} = 0.025\,\text{rad/m}$) at 120 km/h just as a short ice patch drives the friction coefficient down to $\mu \approx 0.15$ (Fig. 5.90).

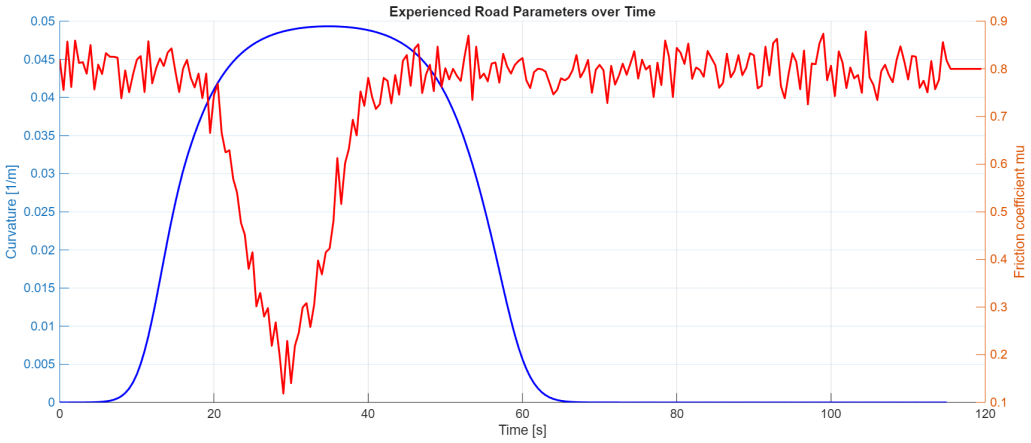

Figure 5.90: Measured friction envelope.

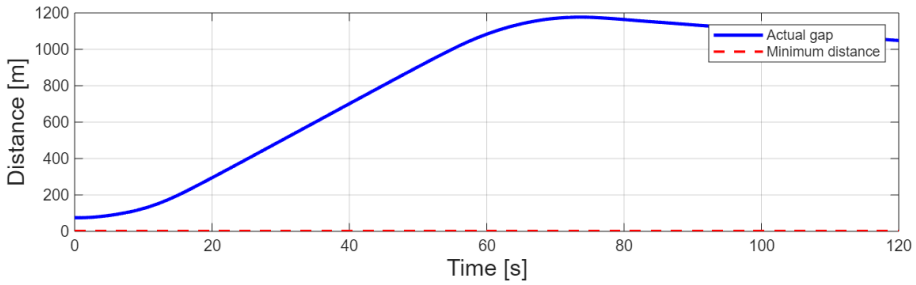

Figure 5.91: Inter-vehicular distance.

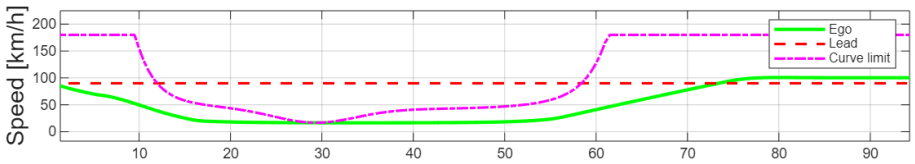

Figure 5.92: Ego and lead speed.

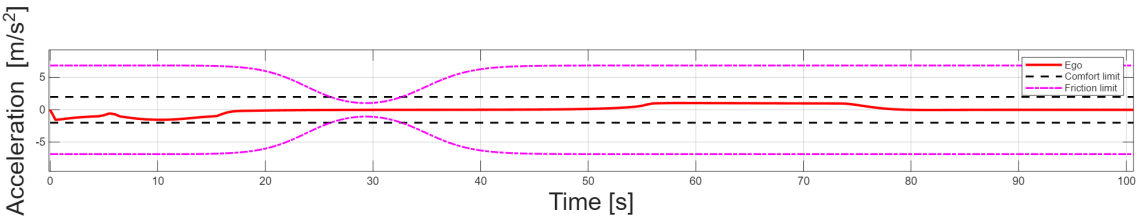

Figure 5.93: Ego acceleration within comfort and friction limits.

The high entry speed makes this manoeuvre critical: as the vehicle enters the curve, friction suddenly collapses due to the ice patch, drastically lowering the admissible safe speed. The MPC reacts and brings the ego velocity down to about 17 km/h, which coincides almost exactly with the curve- and friction-limited bound. This shows that the optimiser prioritises strict feasibility.

Once the patch is cleared and both friction and curvature constraints relax, the ego accelerates smoothly toward the imposed reference speed as discussed in AS3.

Compared to AS5, where the ego remained safely below the curve limit due to conservative planning, AS6 demonstrates a far more demanding case: here the ego speed is forced directly onto the constraint boundary itself. This highlights the controller's ability to exploit the full feasible envelope when necessary, while still guaranteeing safety.

6

Discussion

The evaluation presented in Chapter 5 shows that the proposed friction-adaptive ACC framework successfully addresses the fundamental challenge of safe longitudinal control under uncertain road conditions. This chapter synthesises the key findings and examines their implications for real-world deployment.

6.1 Safety and Robustness

Across all 16 evaluated scenarios (10 use cases and 6 additional stress tests) the controller maintained safety without a single collision or hard constraint violation. The hierarchical constraint structure proved particularly effective when competing demands rose up, such as in UC4 where curve limits conflicted with comfort bounds, and the controller systematically relaxed lower priority constraints while preserving safety critical ones. The minimum achieved distance of 30m in UC7, despite aggressive lead vehicle behaviour and minimal initial separation, validates the terminal constraint design and confirms recursive feasibility even under extreme conditions.

The robustness analysis in Section 5.1.3 showed that the prediction errors remained consistently negative, indicating conservative yet not overly pessimistic planning. This safety oriented bias aligns with the design philosophy while avoiding the excessive conservatism that plagues worst-case approaches.

6.2 Computational Performance

The computational analysis reveals promising real-time capabilities. Average solve times across all scenarios remained below 60ms for both RK4 and ZOH discretisation methods, with typical iteration counts between 8-12 when warm-started. While maximum computation times occasionally exceeded 500ms, these peaks occurred rarely and primarily during severe transients.

The similarity in performance between discretisation methods suggests that the simpler ZOH offers the best practical choice, providing comparable trajectory quality with marginally lower computational overhead.

6.3 Value of Uncertainty Quantification

As shown in UC3, increasing the friction uncertainty visibly widened the MPC's predicted friction bounds. However, this did not cause major changes in distance or velocity behaviour because the controller plans conservatively using the lower and upper friction bounds, which expand progressively along the prediction horizon. This absorbs the effect of uncertainty at constraint level.

This explains why in UC8, enabling stochastic Beta-distributed friction sampling produced almost identical results to the deterministic run. Since the planner already assumes the worst-case friction within the predicted envelope, random realisations within that envelope do not require any additional control effort or constraint relaxation.

6.4 Limitations and Future Directions

While the proposed framework demonstrates strong performance across a wide range of scenarios, certain limitations remain that open avenues for further development. A key assumption throughout this work is the availability of reliable preview data from a forward-facing camera with a 150m range. Although sufficient for moderate speed highway driving, this assumption can break down in adverse weather, occlusions, or night conditions. In such cases, sensor fusion with radar, LiDAR, or map-based friction estimates could improve robustness.

Additionally, the control model used a simplified particle-based vehicle representation, which is appropriate for longitudinal control in single-lane contexts. However, expanding to lateral behaviours such as lane changes or curve negotiation under combined slip conditions would require transitioning to a full bicycle or planar model. Furthermore, the system's handling of lead vehicle behaviour, currently modeled through bounded random acceleration, could be improved by learning-based predictors or intention-aware models that better capture real driver behaviour in traffic.

On the implementation side, while the controller typically solves the optimisation problem in under 60ms on average, there were occasional peaks where the solve time exceeded 500ms, especially during sudden changes or highly dynamic situations. These peaks, while rare, suggest the need for tighter warm-starting strategies and perhaps more efficient solver initialisation tailored to embedded systems. Similarly, while jerk as control input provides smoother trajectories, reformulating the problem using acceleration may reduce the state dimension and solver complexity.

Finally, the degree of conservatism observed in some scenarios is directly influenced by the level of uncertainty in road classification. Therefore, improving the precision of friction estimates, through better training data, model calibration, or multimodal sensing, could directly reduce the required safety margins and increase system responsiveness.

6.5 Broader Implications

Beyond the specific application to ACC, this work demonstrates that sophisticated uncertainty quantification can be embedded within real-time vehicle control without sacrificing performance. The framework's ability to maintain safety across diverse scenarios, from routine highway driving to emergency manoeuvres on icy curves while respecting comfort constraints suggests that similar approaches could enhance other ADAS functions such as emergency braking, lane keeping and eventually fully autonomous navigation.

The measured conservatism introduced by explicit uncertainty modeling aligns with human driving behaviour, such as when experienced drivers naturally increase following distances and reduce speeds when road conditions become challenging and uncertain. By quantifying and systematically handling this uncertainty, the proposed framework represents a step toward ADAS systems that not only match human performance but exceed it in consistency and reliability.

7

Conclusion

This thesis developed and validated a friction-aware Adaptive Cruise Control (ACC) framework based on a Robust Nonlinear Model Predictive Control (RNMPC) formulation. The controller integrates spatially varying road friction and curvature into its decision-making by embedding position-dependent uncertainty bounds directly into the optimisation problem. Friction uncertainty is represented through lower and upper bounds derived from camera-based classification confidence, while stochastic friction realisations are generated via Beta sampling during simulation. Conservative constraint tightening ensures that the controller always plans within the worst-case bounds while still adapting to higher-confidence regions as they approach.

Extensive simulation experiments, including ten main use cases and six additional stress scenarios, showed that the framework consistently maintained safety, comfort, and feasibility under widely varying operating conditions. It successfully handled sharp friction drops, high-speed curve negotiation, and aggressive lead-vehicle behaviour. Safety margins were preserved in all cases, and slack variables were rarely activated, demonstrating that the hierarchical constraint structure effectively balances competing objectives. Notably, the controller's behaviour often resembled that of a cautious human driver, proactively increasing spacing and reducing speed when confidence in surface conditions decreased while still ensuring strict constraint satisfaction through formal optimisation.

The computational performance analysis showed that using both RK4- and ZOH-based time-domain integration, the optimisation typically solved within 50-60 ms, with only rare peaks above 500 ms during severe transients. Warm-starting and structured problem formulation helped maintain low iteration counts even in complex, multi-constraint environments. These results indicate that the proposed approach is well suited for future embedded implementation.

Overall, this work demonstrates that integrating probabilistic friction and curvature estimation into NMPC enables safe and comfortable control on uncertain roads. By unifying friction-aware constraint handling, soft constraint relaxation, and robust optimisation, the proposed controller shows that uncertainty-aware predictive control can be both cautious and efficient offering a foundation for next-generation ADAS systems.

Bibliography

[1] Federal Highway Administration. *How Do Weather Events Affect Roads? - FHWA Road Weather Management*. Accessed: 2025-09-02. 2023. URL: https://ops.fhwa.dot.gov/weather/roadimpact.htm.

[2] Raheb Mirzanamadi. "PhD thesis- Utilizing solar energy for anti-icing road surfaces using hydronic heating pavement with low temperature". PhD thesis. Chalmers University of Technology, Feb. 2019.

[3] Ksander N de Winkel and Michiel Christoph. "Rethinking Advanced Driver Assistance System taxonomies: A framework and inventory of real-world safety performance". In: *Transportation Research Interdisciplinary Perspectives* 29 (2025), p. 101336. ISSN: 2590-1982. DOI: https://doi.org/10.1016/j.trip.2025.101336. URL: https://www.sciencedirect.com/science/article/pii/S2590198225000156.

[4] Wang Qiu et al. "Autonomous vehicle longitudinal following control based on model predictive control". In: July 2015, pp. 8126–8131. DOI: 10.1109/ChiCC.2015.7260933.

[5] Goutham Kumar Mekala, Naveen Reddy Sarugari, and Ameet Chavan. "Speed Control in Longitudinal Plane of Autonomous Vehicle Using MPC". In: *2020 IEEE International Conference for Innovation in Technology (INOCON)*. 2020, pp. 1–5. DOI: 10.1109/INOCON50539.2020.9298213.

[6] Lihua Luo et al. "Model predictive control for adaptive cruise control with multi-objectives: Comfort, fuel-economy, safety and car-following". In: *Journal of Zhejiang University SCIENCE A* 11 (Mar. 2010), pp. 191–201. DOI: 10.1631/jzus.A0900374.

[7] Gerrit Naus et al. "A Model Predictive Control Approach to Design a Parameterized Adaptive Cruise Control". In: vol. 402. Mar. 2010, pp. 273–284. ISBN: 978-1-84996-070-0. DOI: 10.1007/978-1-84996-071-7_17.

[8] Sean Vaskov et al. "Friction-Adaptive Stochastic Predictive Control for Trajectory Tracking of Autonomous Vehicles". In: *2022 American Control Conference (ACC)*. 2022, pp. 1970–1975. DOI: 10.23919/ACC53348.2022.9867523.

[9] Lars Svensson and Martin Törngren. "Fusion of Heterogeneous Friction Estimates for Traction Adaptive Motion Planning and Control". In: *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. 2021, pp. 424–431. DOI: 10.1109/ITSC48978.2021.9564993.

[10] Jonas Herzfeld and Sanjiv Thottathodhi. “Collision Avoidance by Utilizing Dynamic Road Friction Information”. Master’s Thesis. Gothenburg, Sweden: Chalmers University of Technology, 2020.

[11] Szilárd Czibere et al. “Model Predictive Controller Design for Vehicle Motion Control at Handling Limits in Multiple Equilibria on Varying Road Surfaces”. In: *Energies* 14 (Oct. 2021), p. 6667. DOI: 10.3390/en14206667.

[12] Mathias Mattsson et al. “Optimal Model Predictive Acceleration Controller for a Combustion Engine and Friction Brake Actuated Vehicle”. In: *IFAC-PapersOnLine* 49.11 (2016). 8th IFAC Symposium on Advances in Automotive Control AAC 2016, pp. 511–518. ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2016.08.075. URL: https://www.sciencedirect.com/science/article/pii/S2405896316314161.

[13] Jeroen van Hegelsom. “Real-Time MPC Strategy for Predictive Adaptive Cruise Control”. Master’s Thesis. Eindhoven, The Netherlands: Eindhoven University of Technology, Nov. 2023.

[14] Chunxi Huang, Dengbo He, and Song Yan. “Beyond adaptive cruise control and lane centering control: drivers’ mental model of and trust in emerging ADAS technologies”. In: *Frontiers in Psychology* 14 (Aug. 2023). DOI: 10.3389/fpsyg.2023.1236062.

[15] Edmond Awad et al. *Blaming humans in autonomous vehicle accidents: Shared responsibility across levels of automation*. 2018. arXiv: 1803.07170 [cs.AI]. URL: https://arxiv.org/abs/1803.07170.

[16] Alberto Dianin, Elisa Ravazzoli, and Georg Hauger. “Implications of Autonomous Vehicles for Accessibility and Transport Equity: A Framework Based on Literature”. In: *Sustainability* 13.8 (2021). ISSN: 2071-1050. DOI: 10.3390/su13084448. URL: https://www.mdpi.com/2071-1050/13/8/4448.

[17] Karina LaRubbio, Malcolm Grba, and Diana Freed. *Navigating Privacy and Trust: AI Assistants as Social Support for Older Adults*. 2025. arXiv: 2505.02975 [cs.CY]. URL: https://arxiv.org/abs/2505.02975.

[18] Marjo Hippi, Ilkka Juga, and Pertti Nurmi. “A statistical forecast model for road surface friction”. In: Feb. 2010.

[19] National Highway Traffic Safety Administration (NHTSA). *Federal Motor Vehicle Safety Standards: Automatic Emergency Braking Systems for Light Vehicles*. Notice of Proposed Rulemaking Docket No. NHTSA-2023-0021; Document No. 2023-11863. 88 FR 38632. RIN 2127-AM37. Published in the Federal Register on June 13, 2023. U.S. Department of Transportation, June 2023, pp. 38632–38736. URL: https://www.federalregister.gov/d/2023-11863.

A

Common Fixed Parameters

The following parameters remain constant across all use cases unless otherwise specified in the individual use case tables.

Table A.1: MPC Controller Parameters

Parameter	Value	Description
Sampling time	$dt = 0.5$ s	Discretization interval
Prediction horizon	$N = 10$	Number of prediction steps
State penalty matrix	$Q = \text{diag}(0, 0.1, 0.1)$	Position, velocity, acceleration weights
Control penalty	$R = 1$	Jerk penalty weight
Slack penalty matrix	$R_s = \text{diag}(10^3, 10^2, 10^2, 1)$	Safety, velocity, accel., comfort penalties

Table A.2: Vehicle Physical Constraints

Parameter	Value	Description
Maximum acceleration	$a_{max} = 10$ m/s^2	Physical acceleration limit
Comfort acceleration	$a_{c,max} = 2$ m/s^2	Passenger comfort limit
Maximum velocity	$v_{max} = 180$ km/h	Speed limit (50 m/s)
Minimum velocity	$v_{min} = 0$ m/s	Vehicle cannot reverse
Minimum safety distance	$\delta s_{min} = 2$ m	Absolute minimum gap
Lead acceleration bound	$w_{a,lead} = 3$ m/s^2	Maximum lead vehicle acceleration

Table A.3: Road Profile Parameters

Parameter	Value	Description
Friction Profile		
Friction segments	$\mu_{mean} = [0.8, 0.3, 0.8]$	Dry-wet-dry sequence
Transition positions	$s_{trans} = [500, 700]$ m	Friction change locations
Transition steepness	$k_{fric} = 0.1$	Sigmoid transition rate
Curvature Profile		
Curve start position	$s_{curve,start} = 800$ m	Beginning of curved section
Curve end position	$s_{curve,end} = 1000$ m	End of curved section
Curvature steepness	$k_{curve} = 0.1$	Sigmoid transition rate
Curvature uncertainty	$\kappa_{interval} = 0.05$	Fixed uncertainty interval

Table A.4: Simulation and Uncertainty Parameters

Parameter	**Value**	**Description**
Simulation duration	$T_{sim} = 120$ s	Total simulation time
Beta distribution shape	$\rho = 8$	Distribution concentration factor
Minimum friction bound	$\mu_{min} = 0.1$	Physical lower limit
Maximum friction bound	$\mu_{max} = 1.1$	Physical upper limit
Random seed	0	When reproducibility required

Table A.5: Derived Quantities from Fixed Parameters

Quantity	**Formula/Value**	**Description**
Prediction horizon time	$T = N \cdot dt = 5$ s	Total preview time
Simulation steps	$N_{sim} = T_{sim}/dt = 240$	Total discrete time steps
Curve radius	$R = 1/\kappa_{max} = 25$ m	Minimum curve radius
Optimization variables	$n_z = 7N + 3 = 73$	Total decision variables
Max safe curve speed (dry)	$v = \sqrt{0.8 \cdot 9.81/0.04} \approx 14$ m/s	At peak curvature with $\mu = 0.8$
Max safe curve speed (wet)	$v = \sqrt{0.3 \cdot 9.81/0.04} \approx 8.6$ m/s	At peak curvature with $\mu = 0.3$

B

Appendix - Use Case Parameter Specifications

B.1 Discretization Method Parameter Table

Table B.1: MPC Controller Parameters

Parameter	Value	Description
Sampling time	$dt = 0.5$ s	Discretization interval
Prediction horizon	$N = 10$	Number of prediction steps
State penalty matrix	$Q = \text{diag}(0, 0.1, 0.1)$	Position, velocity, acceleration weights
Control penalty	$R = 1$	Jerk penalty weight
Slack penalty matrix	$R_s = \text{diag}(10^3, 10^2, 10^2, 1)$	Safety, velocity, accel., comfort penalties
Reference speed	$v_{ref} = 100$ km/h	Desired cruise speed (27.78 m/s)

Table B.2: Vehicle Physical Constraints

Parameter	Value	Description
Maximum acceleration	$a_{max} = 10$ m/s^2	Physical acceleration limit
Comfort acceleration	$a_{c,max} = 2$ m/s^2	Passenger comfort limit
Maximum velocity	$v_{max} = 180$ km/h	Speed limit (50 m/s)
Minimum velocity	$v_{min} = 0$ m/s	Vehicle cannot reverse
Minimum safety distance	$\delta s_{min} = 2$ m	Absolute minimum gap
Lead acceleration bound	$w_{a,lead} = 3$ m/s^2	Maximum lead vehicle acceleration

Table B.3: Road Profile Parameters

Parameter	Value	Description
Friction Profile		
Friction segments	$\mu_{mean} = [0.8, 0.3, 0.8]$	Dry-wet-dry sequence
Transition positions	$s_{trans} = [700, 1100]$ m	Friction change locations
Transition steepness	$k_{fric} = 0.1$	Sigmoid transition rate
Curvature Profile		
Curve start position	$s_{curve,start} = 900$ m	Beginning of curved section
Curve end position	$s_{curve,end} = 1000$ m	End of curved section
Curvature steepness	$k_{curve} = 0.05$	Sigmoid transition rate
Curvature uncertainty	$\kappa_{interval} = 0.05$	Fixed uncertainty interval

Table B.4: Simulation and Uncertainty Parameters

Parameter	Value	Description
Simulation duration	$T_{sim} = 120$ s	Total simulation time
Beta distribution shape	$\rho = 8$	Distribution concentration factor
Minimum friction bound	$\mu_{min} = 0.1$	Physical lower limit
Maximum friction bound	$\mu_{max} = 1.1$	Physical upper limit
Random seed	0	When reproducibility required

Table B.5: Derived Quantities from Fixed Parameters

Quantity	Formula/Value	Description
Prediction horizon time	$T = N \cdot dt = 5$ s	Total preview time
Simulation steps	$N_{sim} = T_{sim}/dt = 240$	Total discrete time steps
Curve radius	$R = 1/\kappa_{max} = 25$ m	Minimum curve radius
Optimization variables	$n_z = 7N + 3 = 73$	Total decision variables
Max safe curve speed (dry)	$v = \sqrt{0.8 \cdot 9.81/0.04} \approx 14$ m/s	At peak curvature with $\mu = 0.8$
Max safe curve speed (wet)	$v = \sqrt{0.3 \cdot 9.81/0.04} \approx 8.6$ m/s	At peak curvature with $\mu = 0.3$

Table B.6: Configuration Details for Uncertainty-Aware ACC Simulation

Parameter	Value	Description
Initial Conditions		
Ego initial speed	50 km/h	$v_0 = 13.89$ m/s
Ego initial position	0 m	$s_0 = 0$
Lead initial separation	70 m	$ds_{L0} = 70$ m
Lead initial speed	70 km/h	$v_{L0} = 19.44$ m/s
Operational Settings		
Lead vehicle mode	Random acceleration	$a_L \in [-3, 3]$ m/s^2
Friction mode	Stochastic	Beta distribution
1st Discretization	Taylor expansion	Third-order accuracy
2nd Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	0.04 rad/m	κ_{max} (25m radius)
Curve start position	900 m	Curve entry point
Curve end position	1000 m	Curve exit point
Curve steepness	0.05	Transition parameter
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

B.2 Use Case Parameter Tables

Table B.7: UC1: Baseline ACC with Nominal Friction Uncertainty

Parameter	Value	Description
Initial Conditions		
Ego initial speed	50 km/h	$v_0 = 13.89$ m/s
Ego initial position	0 m	$s_0 = 0$
Lead initial separation	70 m	$ds_{L0} = 70$ m
Lead initial speed	70 km/h	$v_{L0} = 19.44$ m/s
Reference speed	$v_{ref} = 100$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	No acceleration
Friction mode	Deterministic	Mean profile only
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.8: UC2: Robustness Test with Random Lead Acceleration

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	100 km/h	$v_0 = 27.78$ m/s
Ego initial position	0 m	$s_0 = 0$
Lead initial separation	50 m	$ds_{L0} = 50$ m
Lead initial speed	100 km/h	$v_{L0} = 27.78$ m/s
Reference speed	$v_{ref} = 200$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Random acceleration	$a_L \in [-3, 3]$ m/s^2
Friction mode	Stochastic	Beta distribution
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.9: UC3: High Uncertainty Friction Transition

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	80 km/h	$v_0 = 22.22$ m/s
Ego initial position	400 m	$s_0 = 400$ m
Lead initial separation	60 m	$ds_{L0} = 60$ m
Lead initial speed	80 km/h	$v_{L0} = 22.22$ m/s
Reference speed	$v_{ref} = 100$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	No acceleration
Friction mode	Stochastic	Beta distribution
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.7	$\mu_{unc,far} = 0.7$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.10: UC4: Curve Navigation at High Speed

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	140 km/h	$v_0 = 38.89$ m/s
Ego initial position	600 m	$s_0 = 600$ m
Lead initial separation	150 m	$ds_{L0} = 150$ m
Lead initial speed	90 km/h	$v_{L0} = 25$ m/s
Reference speed	$v_{ref} = 200$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	No acceleration
Friction mode	Deterministic	Mean profile only
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	0.04 rad/m	κ_{max} (25m radius)
Curve start position	900 m	Curve entry point
Curve end position	1000 m	Curve exit point
Curve steepness	0.05	Transition parameter
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.11: UC5: Combined Following and Curve Challenge

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	90 km/h	$v_0 = 25$ m/s
Ego initial position	700 m	$s_0 = 700$ m
Lead initial separation	80 m	$ds_{L0} = 80$ m
Lead initial speed	85 km/h	$v_{L0} = 23.61$ m/s
Reference speed	$v_{ref} = 200$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Random acceleration	$a_L \in [-3, 3]$ m/s^2
Friction mode	Stochastic	Beta distribution
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	0.04 rad/m	κ_{max} (25m radius)
Curve start position	900 m	Curve entry point
Curve end position	1000 m	Curve exit point
Curve steepness	0.05	Transition parameter
Uncertainty Parameters		
Current uncertainty	0.15	$\mu_{unc,0} = 0.15$
Far uncertainty	0.4	$\mu_{unc,far} = 0.4$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.12: UC6: Low-Speed Convergence Test

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	30 km/h	$v_0 = 8.33$ m/s
Ego initial position	0 m	$s_0 = 0$
Lead initial separation	100 m	$ds_{L0} = 100$ m
Lead initial speed	70 km/h	$v_{L0} = 19.44$ m/s
Reference speed	$v_{ref} = 100$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	No acceleration
Friction mode	Stochastic	Beta distribution
Discretization	Taylor expansion	Third-order accuracy
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.13: UC7: Close-Following Robustness Scenario

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	80 km/h	$v_0 = 22.22$ m/s
Ego initial position	450 m	$s_0 = 450$ m
Lead initial separation	30 m	$ds_{L0} = 30$ m
Lead initial speed	80 km/h	$v_{L0} = 22.22$ m/s
Reference speed	$v_{ref} = 180$ km/h	Desired cruise speed (50 m/s)
Operational Settings		
Lead vehicle mode	Random acceleration	$a_L \in [-3, 3]$ m/s^2
Friction mode	Stochastic	Beta distribution
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.14: UC8: Stochastic vs Deterministic Comparison

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	70 km/h	$v_0 = 19.44$ m/s
Ego initial position	0 m	$s_0 = 0$
Lead initial separation	70 m	$ds_{L0} = 70$ m
Lead initial speed	70 km/h	$v_{L0} = 19.44$ m/s
Reference speed	$v_{ref} = 100$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	No acceleration
Friction mode	Both modes	Run twice
Discretization	RK4	Fourth-order accuracy
Random seed	0	Fixed for reproducibility
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.15: UC9: Extreme Initial Speed Differential

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	150 km/h	$v_0 = 41.67$ m/s
Ego initial position	0 m	$s_0 = 0$
Lead initial separation	120 m	$ds_{L0} = 120$ m
Lead initial speed	50 km/h	$v_{L0} = 13.89$ m/s
Reference speed	$v_{ref} = 100$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	No acceleration
Friction mode	Stochastic	Beta distribution
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.1	$\mu_{unc,0} = 0.1$
Far uncertainty	0.3	$\mu_{unc,far} = 0.3$
Preview distance	150 m	$s_{unc,max} = 150$ m

Table B.16: UC10: Minimal Uncertainty Validation

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial speed	90 km/h	$v_0 = 25$ m/s
Ego initial position	400 m	$s_0 = 400$ m
Lead initial separation	70 m	$ds_{L0} = 70$ m
Lead initial speed	90 km/h	$v_{L0} = 25$ m/s
Reference speed	$v_{ref} = 100$ km/h	Desired cruise speed (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	No acceleration
Friction mode	Stochastic	Beta distribution
Discretization	RK4	Fourth-order accuracy
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Curve location	N/A	Essentially straight
Uncertainty Parameters		
Current uncertainty	0.05	$\mu_{unc,0} = 0.05$
Far uncertainty	0.1	$\mu_{unc,far} = 0.1$
Preview distance	150 m	$s_{unc,max} = 150$ m

C

Appendix - Additional Scenarios Parameter Specifications

Table C.1: AS1: Driver-Imposed Speed Limitation

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial state	[0; 60/3.6; 0]	Position 0m, speed 16.67 m/s
Lead initial separation	80 m	$ds_{L0} = 80$ m
Lead initial speed	100/3.6 m/s	$v_{L0} = 27.78$ m/s
Reference speed	80/3.6 m/s	Driver preference (22.22 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	random_leading_vehicle_acceleration = false
Friction mode	Deterministic	random_friction = false
Discretization	RK4	use_exact_discretization = false
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Friction profile	[0.8, 0.3, 0.8]	Standard dry-wet-dry
Friction transitions	[500, 700] m	Standard positions
Curve location	[800, 1000] m	Standard curve section
Uncertainty Parameters		
Uncertainty bounds	[0.1, 0.3]	$\mu_{unc} = [0.1, 0.3]$
Preview distance	150 m	pos_mu_unc = 150 m

Table C.2: AS2: Progressive Friction Degradation

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial state	[0; 80/3.6; 0]	Position 0m, speed 22.22 m/s
Lead initial separation	70 m	$ds_{L0} = 70$ m
Lead initial speed	80/3.6 m/s	$v_{L0} = 22.22$ m/s
Reference speed	100/3.6 m/s	Standard reference (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	random_leading_vehicle_acceleration = false
Friction mode	Stochastic	random_friction = true
Discretization	RK4	use_exact_discretization = false
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Friction profile	[0.8, 0.5, 0.2]	Dry-wet-icy degradation
Friction transitions	[400, 700] m	Progressive worsening
Curve location	[800, 1000] m	Standard curve section
Uncertainty Parameters		
Uncertainty bounds	[0.1, 0.3]	$\mu_{unc} = [0.1, 0.3]$
Preview distance	150 m	pos_mu_unc = 150 m

Table C.3: AS3: Curve Coincident with Low Friction

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial state	[300; 100/3.6; 0]	Position 300m, speed 27.78 m/s
Lead initial separation	100 m	$ds_{L0} = 100$ m
Lead initial speed	90/3.6 m/s	$v_{L0} = 25$ m/s
Reference speed	100/3.6 m/s	Standard reference (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	random_leading_vehicle_acceleration = false
Friction mode	Stochastic	random_friction = true
Discretization	RK4	use_exact_discretization = false
Road Geometry		
Maximum curvature	0.05 rad/m	κ_{max} (33m radius)
Friction profile	[0.8, 0.3, 0.8]	Standard dry-wet-dry
Friction transitions	[500, 700] m	Coincides with curve
Curve location	[500, 700] m	Aligned with low friction
Uncertainty Parameters		
Uncertainty bounds	[0.15, 0.4]	$\mu_{unc} = [0.15, 0.4]$
Preview distance	150 m	pos_mu_unc = 150 m

Table C.4: AS4: Extended Low-Friction Region

Parameter	Value	Description
Initial Conditions		
Ego initial state	[200; 100/3.6; 0]	Position 200m, speed 27.78 m/s
Lead initial separation	90 m	$ds_{L0} = 90$ m
Lead initial speed	80/3.6 m/s	$v_{L0} = 22.22$ m/s
Reference speed	100/3.6 m/s	Standard reference (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	random_leading_vehicle_acceleration = false
Friction mode	Stochastic	random_friction = true
Discretization	RK4	use_exact_discretization = false
Road Geometry		
Maximum curvature	N/A	κ_{max} (straight road)
Friction profile	[0.8, 0.3, 0.8]	Standard profile
Friction transitions	[400, 800] m	Extended low region
Curve location	N/A	No curve on straight road
Uncertainty Parameters		
Uncertainty bounds	[0.1, 0.3]	$\mu_{unc} = [0.1, 0.3]$
Preview distance	150 m	pos_mu_unc = 150 m

Table C.5: AS5: High-Speed Curve Entry with Mid-Curve Friction Change

Parameter	Value	Description
Initial Conditions		
Ego initial state	[400; 120/3.6; 0]	Position 400m, speed 33.33 m/s
Lead initial separation	150 m	$ds_{L0} = 150$ m
Lead initial speed	90/3.6 m/s	$v_{L0} = 25$ m/s
Reference speed	100/3.6 m/s	Standard reference (27.78 m/s)
Operational Settings		
Lead vehicle mode	Constant speed	random_leading_vehicle_acceleration = false
Friction mode	Stochastic	random_friction = true
Discretization	RK4	use_exact_discretization = false
Road Geometry		
Maximum curvature	0.025 rad/m	κ_{max} (40m radius)
Friction profile	[0.8, 0.5, 0.8]	Dry-wet-dry
Friction transitions	[700, 750] m	Mid-curve transition
Curve location	[600, 800] m	Contains friction change
Uncertainty Parameters		
Uncertainty bounds	[0.1, 0.35]	$\mu_{unc} = [0.1, 0.35]$
Preview distance	150 m	pos_mu_unc = 150 m

Table C.6: AS6: Ice Patch on a Curve

Parameter	**Value**	**Description**
Initial Conditions		
Ego initial state	[350; 90/3.6; 0]	Pos. 350 m, 25 m/s (90 km/h)
Lead initial separation	75 m	$ds_{L0} = 75$ m
Lead initial speed	90/3.6 m/s	$v_{L0} = 25$ m/s
Reference speed	100/3.6 m/s	Same as ego start speed
Operational Settings		
Lead vehicle mode	Constant speed	`random_leading_vehicle_acceleration=false`
Friction mode	Stochastic	`random_friction=true`
Discretisation	RK4	`use_exact_discretization=false`
Road Geometry		
Maximum curvature	0.05 rad/m	Tight curve (20 m radius)
Friction profile	[0.8, 0.1, 0.8]	Dry–ice–dry
Friction transitions	[650, 700] m	Ice patch centred at apex
Curve location	[600, 800] m	Curve spans the ice section
Uncertainty Parameters		
Uncertainty bounds	[0.1, 0.35]	$\mu_{unc} = [0.1, 0.35]$
Preview distance	150 m	`pos_mu_unc=150 m`

Department of Electrical Engineering - *Division of Systems, Control and Mechatronics*
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY