CHALMERS
UNIVERSITY OF TECHNOLOGY

Adaptive Lateral Control of Autonomous Trucks

Development and evaluation of adaptive strategies for lateral control of autonomous heavy vehicles in haulage systems

Master's thesis in Systems, Control and Mechatronics

ANIRUDH VINOD NAIR

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis 2026

Adaptive Lateral Control of Autonomous Trucks

Development and evaluation of adaptive strategies for lateral control of autonomous heavy vehicles in haulage systems

ANIRUDH VINOD NAIR

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Division of Systems and Control
Chalmers University of Technology
Gothenburg, Sweden 2026

Adaptive Lateral Control of Autonomous Trucks
Development and evaluation of adaptive strategies for lateral control of autonomous heavy vehicles in haulage systems
ANIRUDH VINOD NAIR

Supervisor: Nikolce Murgovski, Department of Electrical Engineering
Supervisor: Robert Hult, Volvo Autonomous Solutions
Supervisor: Stefan Kojchev, Volvo Autonomous Solutions
Examiner: Nikolce Murgovski, Department of Electrical Engineering

Master's Thesis 2026
Department of Electrical Engineering
Division of Systems and Control
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Autonomous hauling truck by Volvo.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Adaptive Lateral Control of Autonomous Trucks
Development and evaluation of adaptive strategies for lateral control of autonomous heavy vehicles in haulage systems
ANIRUDH VINOD NAIR
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Autonomous heavy vehicle control systems depend on a suite of parameters, both vehicular and environmental, and their performance is directly related to these parameter values. These parameters often include variables such as the total mass of the loaded vehicle, lateral and longitudinal tire cornering coefficients, wheelbase geometry, load distribution, etc. In theory, these parameters are often considered constant and known. However, in practice, this might not always be the case, especially in dynamic operating environments of a quarry or mine. Therefore, to ensure robust and reliable performance, the controller must be adaptive, updating itself online to continue operating efficiently.

This thesis focuses on the investigation, development, and implementation of different types of adaptive lateral controllers for autonomous haulage trucks, whose parameters initially deviate from the true value and are subject to change during operation. The study focuses on evaluating both direct and indirect approaches to adaptive control. The direct adaptive controller uses a *Lyapunov*-based approach towards adaptation, meanwhile the indirect controller uses *Kalman filter*-based estimation methods to estimate unknown system parameters.

Overall, the findings suggest that the indirect adaptive control strategies are more suitable for autonomous haulage applications. The integration of Kalman filter-based parameter estimation with model based controllers like LQR resulted in consistent performance and robustness, and better handling of system uncertainties compared to direct adaptive approaches.

Keywords: Adaptive Lateral Control, Autonomous Trucks, Parameter Estimation, Extended Kalman Filter, Unscented Kalman Filter, PID Control, Linear Quadratic Regulator.

Acknowledgements

I would like to express my most sincere gratitude towards my supervisors at Volvo Autonomous Solutions, Robert Hult and Stefan Kojchev. Throughout the entirety of the thesis work, they provided great support and guidance which aided me in achieving the goals of my thesis topic. Their insight and experience in the field of heavy-vehicle dynamics and control helped me in my investigation of adaptive strategies and their implementation. Furthermore, their knowledge of industry standards helped me in keeping my ideas and implementations grounded in nature. I would also like to thank my supervisor and examiner at Chalmers, Nikolce Murgovski, for providing me with both theoretical and administrative support throughout the thesis work. Finally, I would like to thank Edo Drenth at Volvo Autonomous Solutions, with whom I had many insightful discussions with respect to certain topics in my thesis work.

Anirudh Vinod Nair, Gothenburg, June 2026

List of Acronyms

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

AHS	Autonomous Hauling System
CAE	Cumulative Absolute Error
COM	Center of Mass
EKF	Extended Kalman Filter
LQR	Linear Quadratic Regulator
LTI	Linear Time Invariant
MPC	Model Predictive Control
MRAC	Model Reference Adaptive Control
OEM	Original Equipment Manufacturer
PID	Proportional Integral Derivative
RLS	Recursive Least Square
UKF	Unscented Kalman Filter
VAS	Volvo Autonomous Solutions

Nomenclature

Below is the nomenclature of indices, parameters, and variables that have been used throughout this thesis.

Indices

k	Index for discrete time instance

Parameters

m	mass of the vehicle
I_z	moment of inertia of the vehicle around its z-axis
v_x	longitudinal velocity of the vehicle
L	wheelbase length of the vehicle
L_f	distance between the COM of the vehicle and the front axle
L_r	distance between the COM of the vehicle and the rear axle
$F_{z,f}$	normal force on the front wheels
$F_{z,r}$	normal frictional force on the rear wheels
C_f	lateral cornering stiffness coefficient on the front wheel
C_r	lateral cornering stiffness coefficient on the rear wheel
γ	adaptation constant
$\mathbf{\Gamma}$	adaptation constant vector
w_k	gaussian process noise
Q	Process noise covariance matrix
v_k	gaussian measurement noise
R	Measurement noise covariance matrix
P	State error covariance matrix
K_p	Proportional gain of the PID controller

K_i	Integral gain of the PID controller
K_d	Derivative gain of the PID controller
K_ψ	Proportional gain for the heading
K_{v_y}	Proportional gain for the lateral velocity
K_r	Proportional gain for the yaw rate
σ_0	leakage coefficient
ν	ratio for soft constraint

Variables

V	Velocity of the vehicle
v_y	lateral velocity of the vehicle
$v_{y,eq}$	lateral velocity of the vehicle at equilibrium condition
δ	total input steering angle of the vehicle
δ_{ff}	feedforward input steering angle of the vehicle
δ_{fb}	feedback input steering angle of the vehicle
α_f	slip angle of the front wheels of the vehicle
α_r	slip angle of the rear wheels of the vehicle
ψ	heading (or the yaw angle) of the vehicle
r	yaw rate of the vehicle
r_{eq}	yaw rate of the vehicle at equilibrium condition
$F_{l,f}$	longitudinal frictional force on the front wheel
$F_{l,r}$	longitudinal frictional force on the rear wheel
$F_{c,f}$	lateral frictional force on the front wheel
$F_{c,r}$	lateral frictional force on the rear wheel
$\mathbf{x}$	state vector of the model
s	distance covered by the vehicle on the reference path
d	lateral offset of the vehicle from the reference path
$\kappa(s)$	curvature of the path at s
$\varphi(s)$	path tangent at s
β	velocity slip angle
$\mathbf{e}$	error vector of the model
A, B, C, D	linear state space matrices
A_m, B_m	linear state space matrices of the reference model

K^*_{fb}	ideal feedback gain
K^*_{ff}	ideal feedforward gain
K_{fb}	feedback gain
K_{ff}	feedforward gain
$\tilde{K}$	error in gain
J	cost function
$\mathbf{\Phi}$	regressor vector

Contents

List of Figures

List of Tables

1

Introduction

The field of autonomous heavy vehicles has experienced a rapid growth in both research and industrial deployment during the last two decades, particularly in structured and controlled environments of open-pit mines, quarry sites, and remote hauling operations. Since the first commercial deployment of Autonomous Hauling Systems (AHS) in 2008, it has expanded significantly throughout Australia, North America, and South America [1].

Recent industry reports indicate that AHS has evolved from a pilot-scale experimentation into an industrial-scale deployment. Major OEMs and operators such as Komatsu, Caterpillar, and Rio Tinto have deployed large-scale autonomous fleets to improve safety, productivity, and enable continuous 24/7 operation in environments with minimal public interaction and highly structured traffic conditions [2, 3, 4]. These developments, supported by mining automation surveys and industrial reports, show that autonomous transport systems are now a mature and rapidly expanding solution for repetitive transportation tasks in mining and quarry operations [5].

Volvo Autonomous Solutions (V.A.S.) has also contributed to autonomous heavy vehicle development and its applications in open-quarries and hub-to-hub freight corridors. After early research and prototype demonstrations of autonomous truck platforms, Volvo has progressed toward commercial deployment of autonomous transport systems, including driverless operations in quarry and industrial logistics sites in Scandinavia [6].

1.1 Motivation

Autonomous heavy vehicle control systems rely on a range of vehicle and environmental parameters that, in conventional control design, are often assumed to be constant or slowly varying. These include, among others, total vehicle mass, tire-road friction coefficient, wheelbase geometry, and load distribution. However, in practice, these parameters can vary during operation due to changing payload conditions, uneven loading, tire wear, road surface variations, and weather effects. These can lead to a model mismatch between the controller model assumptions

and the actual vehicle dynamics, which may degrade the control performance and compromise system stability.

In light of these challenges, the idea of an adaptive control strategy has emerged as a promising approach for maintaining robust performance in the system. By constantly or periodically updating the control law in a direct or indirect fashion, these controllers can compensate for modeling errors and environmental changes, thereby improving tracking accuracy and overall system performance.

Within an autonomous vehicle control system, both longitudinal and lateral control play critical roles in ensuring safe and efficient vehicle operation. This thesis focuses specifically on the lateral control problem, which is responsible for trajectory tracking and path following. The objective is to investigate, develop, and evaluate different adaptive lateral control strategies for autonomous heavy vehicles, and to assess their performance under varying operating conditions and modeling uncertainties.

1.2 Literature Review

Adaptive lateral control strategy is an extremely broad topic, which derives knowledge primarily from the fields: vehicle dynamics modeling, lateral controllers, and adaptive controllers. A straightforward approach is to start with a conventional lateral controller and then augment it with an adaptive layer. Thus, we focus on a literature review of the above mentioned fields before combining them together.

1.2.1 Modeling of vehicle dynamics

Modeling of the non-linear vehicle dynamics of the autonomous trucks forms the foundation of the control design problem. A widely adopted reference in this domain is the textbook by Rajamani, which provides various comprehensive models of longitudinal and lateral vehicle dynamics, including the single-track bicycle models and tire force representations commonly used in lateral control design [7].

In addition to this, several recent studies have further supported and extended these modeling approaches. A technical review of lateral control for autonomous vehicles by Yassine Kebbati et al. that reinforces the applicability of such classical vehicle models in modern autonomous driving frameworks and highlights how they are commonly employed as a baseline for controller design [8].

Furthermore, the work presented in the master's thesis by Jawahar et al. demonstrates how these simplified vehicle models can be effectively used in the design and evaluation of lateral control systems for real-world heavy vehicle platforms [9].

Together, these sources establish a consistent vehicle modeling framework in which classical dynamics theory is used as a baseline, while contemporary research validates and adapts these models for autonomous heavy vehicle applications.

1.2.2 Lateral control strategies

With trajectory tracking and path following in mind, lateral control strategies in vehicles can be classified into geometric-based and model-based approaches. Geometric controllers, such as Pure Pursuit and Stanley, relying on the geometric relationship between the vehicle and a reference path, are often favored for their simplistic and intuitive design. In contrast, model-based approaches such as linear quadratic regulator (LQR) and model predictive control (MPC), are derived from an explicit vehicle dynamics model and typically offer improved performance in terms of stability and tracking accuracy.

The technical review by Yassine Kebbati et al. primarily focuses on model-based approaches, like LQR and MPC, and discusses their formulation, advantages and limitations in autonomous driving applications [8]. In contrast, the work by Jawahar et al. provides a broader comparison of both geometric and model-based controllers, including Pure Pursuit, Stanley and LQR, in the context of heavy vehicle applications, comparing their practical performance under realistic operating conditions [9].

1.2.3 Adaptive control

Adaptive control is a broad class of control strategies which are capable of handling systems with uncertain or time varying parameters, as comprehensively presented in the textbook Adaptive Control by Åström K. J. et al. [10]. These controllers are commonly divided into two categories: direct and indirect adaptive controllers.

A widely used framework for direct adaptive controllers is the Model Reference Adaptive Control (MRAC) approach, where the controller gains are directly updated so as to follows a desired reference model. This framework can be implemented using the MIT rule or the Lyapunov based approach, as illustrated in the various literature and tutorials [11, 12, 13, 14]. In this work, a pseudo-MRAC approach is adopted, inspired by these formulations, where the adaptation law is reformulated for the purposes of the thesis.

Indirect adaptive control, in contrast, first estimates the unknown system parameters and then accordingly updates the controller. Parameter estimation can be performed using methods such as Recursive Least Squares (RLS) and Kalman-filter-based approaches, including the Extended Kalman Filter (EKF) and Unscented Kalman Filter (UKF). The technical review by Amal Chebbi et al develops and illustrates how different variation of the aforementioned Kalman filters can be used for parameter estimation [16].

1.2.4 Research gap

The existing literature, that we have briefly discussed in the above sections, has extensively studied vehicle dynamics modeling, lateral control strategies and adaptive control methods; however, these areas are often treated as separate research

directions rather than being fully integrated into a unified framework.

In direct adaptive control, most studies rely on the standard Model Reference Adaptive Control (MRAC) formulation, with limited exploration of modified or application specific variations that are tailored to particular vehicle control problems. Similarly in indirect adaptive control, few works combine parameter estimation with feedback control in an integrated manner.

Furthermore, within different approaches to parameter estimation, Kalman filter based approaches have been explored to some extent, but their use remains limited compared to classical methods such as Recursive Least Squares. The application of adaptive variants of Kalman filtering for real-time parameter estimation in vehicular control systems is even less common.

1.3 Contributions from the Thesis

This thesis addresses the aforementioned research gaps by proposing and evaluating a set of adaptive lateral control strategies. The main contributions are summarized as follows:

- The thesis investigates, develops and evaluates a family of adaptive lateral control strategies specifically designed for heavy vehicle applications.
- An indirect adaptive lateral control framework is developed by combining Kalman filter based parameter estimation techniques with model-based lateral controllers, enabling real-time adaptation to changing vehicle dynamics.
- A novel parameter drift-rate-based adaptive tuning strategy for the process noise covariance (Q) in Kalman filter based estimators is developed to improve parameter estimation performance under varying operating conditions and sudden parameter variations.
- A Model Reference Adaptive Control (MRAC) based direct adaptive lateral PID controller is developed and tailored specifically for lateral path tracking of autonomous heavy vehicles.
- The performance and robustness of the proposed direct and indirect adaptive controllers are tested and evaluated under multiple driving scenarios, including varying vehicle parameters and environmental conditions.

1.4 Thesis Overview

The remainder of this report continues in the following structure. Before stipulating a problem statement, chapter 2 first goes through the necessary theoretical background, which includes understanding the vehicle dynamics, and exploring the

various types of lateral controllers and adaptive control strategies. Chapter 3 then states the various limitations and assumption taken, before presenting the problem formulation. Based on the theory discussed in chapter 2 and the assumptions from chapter 3, chapter 4 develops and presents the various types of adaptive lateral controllers, while using certain output plots to illustrate key development points. Finally, the various controllers are tested and their results are compared in chapter 5 and in chapter 6 the discussion of the final conclusions and future work is done.

2

Theoretical Background

As explained in section 1.2, in order for us to develop an adaptive lateral controller, we must discuss and understand the foundations upon which they are built, namely the vehicle dynamics, lateral control strategies and the derivation of an adaptive approach. This chapter goes through each of these core concepts, by first modeling the essential lateral dynamics of the vehicle in a suitable coordinate frame, to then exploring the various lateral controllers which are viable in the chosen frame and finally discussing the different types of indirect and direct adaptive control strategies.

2.1 Vehicle Modeling

Before considering the different types of adaptive control techniques available at our disposal, we must first establish an appropriate model to describe the dynamics of the vehicle system. Most vehicle models can be categorized into two types: kinematic model and kinetic model.

A kinematic model assumes that motion is purely geometric and ignores forces and masses entirely. They are usually valid at low speeds and no slip assumptions. Meanwhile, kinetic models account for forces, inertia and other factors such as slip. These are used when we need to model high speed systems with aggressive maneuvers and are much more complex than kinematic models.

The added complexity of the kinetic model allows for a more realistic representation of the vehicle behavior under real-world operating conditions, thereby providing a more reliable basis for controller design and simulation. Thus, in this thesis, a kinetic model was chosen to model the vehicle dynamics.

2.1.1 Kinetic bicycle model

A bicycle model is a common choice for a vehicle model, owing to its ability to represent the essential motion characteristics while maintaining a relatively simple mathematical structure. By reducing the elaborate multi wheel configuration to an equivalent two wheel representation, the model significantly decreases the complexity without sacrificing key vehicle dynamics required for analysis.

Considering the dynamics bicycle model in Figure 2.1, we can note that it accounts for the 2D planar motion of the vehicle along the X and Y axes, and rotation around the Z axis (yaw) of the reference frame. This yields a vehicle model with three degrees of freedom (3 DoF), along which, we can apply Newton's law of motion,

$$\sum F_x = ma_x, \tag{2.1}$$
$$\sum F_y = ma_y, \tag{2.2}$$
$$\sum M_z = I_z\ddot{\psi} \tag{2.3}$$

where F_x and F_y represent the longitudinal and lateral forces, while I_z and M_z represent the moment of inertia and moment around the Z axis respectively. Using Figure 2.1, we can further expand the Newtonian equations as,

$$F_x = m\dot{v_x} = (F_{lf} \cdot cos\delta) - (F_{cf} \cdot sin\delta) + F_{lr} \tag{2.4}$$
$$F_y = m\dot{v_y} = (F_{lf} \cdot sin\delta) + (F_{cf} \cdot cos\delta) + F_{cr} \tag{2.5}$$
$$I_Z\ddot{\psi} = L_f(F_{cf}\cos\delta + F_{lf}\sin\delta) - L_rF_{cr} \tag{2.6}$$

where m is the mass of the vehicle, F_{lf} and F_{lr} are the longitudinal frictional forces of the front and rear tires, and F_{cf} and F_{cr} are the lateral frictional forces. The figure also describes the velocity of the vehicle V as a combination of lateral velocity v_y and longitudinal velocity v_x. If we assume that V acts on a point on the wheelbase L, then L_f and L_r is the distance of this point from the front and rear tires. The angles α_f and α_r are the slip angles of the respective tires, δ_f is the steering angle of the front wheel, and ψ is the yaw angle that the vehicle makes with the coordinate axis.

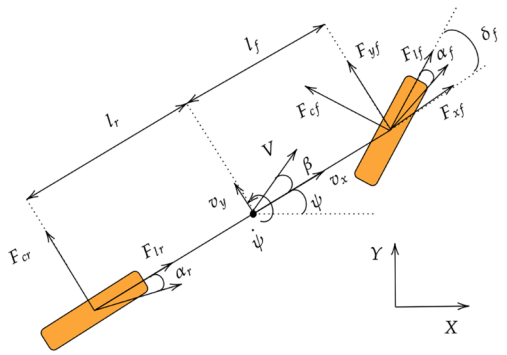

Figure 2.1: Dynamic bicycle model [8]. The various longitudinal and lateral forces acting on the vehicle can be observed here. The direction of vector V indicates the direction of motion of CoM of the vehicle.

Since we are focusing on only modeling the lateral dynamics, we do not need to model the longitudinal dynamics, as these can be considered as part of the set of parameters that affect the model dynamics.

2.1.1.1 Tire force modeling

The lateral frictional force, also known as the lateral tire force, acting on the vehicle wheels is dependent on how we model the tire forces [20]. For computational efficiency and assuming that the slip angles are sufficiently small, we can model the lateral tire forces linearly as,

$$F_{cf} = C_f F_{z,f} \alpha_f \quad and \quad F_{cr} = C_r F_{z,r} \alpha_r \tag{2.7}$$

where, C_f and C_r are the lateral cornering stiffness coefficients, and $F_{z,f}$ and $F_{z,r}$ are the normal forces acting on the tires. Furthermore, we can represent the slip angles by the equations,

$$\alpha_f = \arctan\left(\frac{v_y + L_f\dot{\psi}}{v_x}\right) - \delta \tag{2.8}$$

$$\alpha_r = \arctan\left(\frac{v_y - L_r\dot{\psi}}{v_x}\right) \tag{2.9}$$

2.1.1.2 Cartesian frame system

Based on the aforementioned force equations, we can write the dynamics of bicycle model as a non linear state space equation of the form $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$ in the cartesian frame as,

$$\dot{x} = v_x \cos\psi - v_y \sin\psi \tag{2.10}$$

$$\dot{y} = v_x \sin\psi + v_y \cos\psi \tag{2.11}$$

$$\dot{\psi} = \omega \tag{2.12}$$

$$\dot{v}_y = -\frac{1}{m}\left(C_f F_{z,f}\left(\arctan\frac{v_y + L_f\omega}{v_x} - \delta\right)\cos\delta + C_r F_{z,r}\arctan\frac{v_y - L_r\omega}{v_x}\right) - v_x\omega \tag{2.13}$$

$$\dot{\omega} = -\frac{1}{I_z}\left(L_f C_f F_{z,f}\left(\arctan\frac{v_y + L_f\omega}{v_x} - \delta\right)\cos\delta - L_r C_r F_{z,r}\arctan\frac{v_y - L_r\omega}{v_x}\right) \tag{2.14}$$

where the states vector is defined as:

$$\mathbf{x} = \begin{bmatrix} x & y & \psi & v_y & \omega \end{bmatrix}^T$$

with ω representing the yaw rate of the vehicle ($\dot{\psi}$) and δ as the steering input $\mathbf{u}$ for the system.

While the cartesian frame state equations perfectly encapsulate the dynamics of the vehicle, from a lateral control perspective, it would be beneficial if the equations were more nuanced, especially the equations describing the position of the vehicle.

2.1.1.3 Frenet-Serret frame system

The Frenet–Serret frame, also called the curvilinear or road-aligned frame, is a coordinate system frame that moves along a reference path, making it natural for describing a vehicle's deviation from that path.

As can be seen in Figure 2.2, the center line of the road can be chosen as the reference line for the frame. The direction along this reference line is known as the longitudinal axis or the s-axis, while the directions perpendicular to any point on the reference line are known as the lateral axes or the d-axes [21]. These axes can be related to the Cartesian frame by the property,

$$(x, y) = p(s) = (p_x(s), p_y(s)) \tag{2.15}$$

Based on the definition, it is clear that a vector along the d-axis is lateral offset of the vehicle, which should ideally converge to zero at steady state, and s is the total distance covered by the vehicle along the reference path. Thus, we can convert the earlier cartesian frame dynamic vehicle model (equations 2.10 - 2.14) into the Frenet-Serret frame as,

$$\dot{s} = \frac{v_x \cos(\psi - \varphi(s) + \beta) - v_y \sin(\psi - \varphi(s) + \beta)}{1 - \kappa(s)d} \tag{2.16}$$

$$\dot{d} = v_x \sin(\psi - \varphi(s) + \beta) + v_y \cos(\psi - \varphi(s) + \beta) \tag{2.17}$$

$$\dot{\psi}_{err} = \dot{\psi} - \kappa(s)\dot{s} \tag{2.18}$$

$$\dot{v}_y = -\frac{1}{m}\left(C_f F_{z,f}\left(\arctan\frac{v_y + L_f\omega}{v_x} - \delta\right)\cos\delta + C_r F_{z,r}\arctan\frac{v_y - L_r\omega}{v_x}\right) - v_x\omega \tag{2.19}$$

$$\dot{\omega} = -\frac{1}{I_z}\left(L_f C_f F_{z,f}\left(\arctan\frac{v_y + L_f\omega}{v_x} - \delta\right)\cos\delta - L_r C_r F_{z,r}\arctan\frac{v_y - L_r\omega}{v_x}\right) \tag{2.20}$$

where $\kappa(s)$ and $\varphi(s)$ are the path curvature and path tangent angle at the point s along the path, and $\beta = \arctan(v_y/v_x)$ is the velocity slip angle.

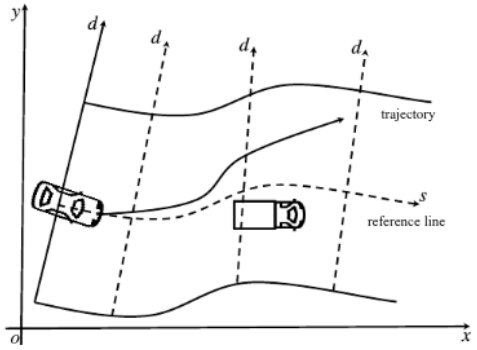

Figure 2.2: Frenet-Serret frame vs Cartesian frame [21]. The dotted line labeled s is the longitudinal axis, while the dotted lines labeled d perpendicular to the reference line are the lateral axes of the frame.

2.2 Lateral Control

Autonomous vehicle control can be broadly divided into two parts, longitudinal and lateral control. While longitudinal control is responsible for maintaining desired speed, application of acceleration and brakes, and ensuring safe following distances, lateral control governs the steering behavior and maintains stability while following a desired path.

As mentioned before, in this work, the focus lies on lateral control part, where the objective of the controller would be to minimize lateral displacement d and heading error ψ_{err} with respect to the reference path and road direction, by providing an optimal steering command input δ. To that effect, the Frenet-Serret frame dynamic model defined earlier would be a better choice, as compared to the cartesian frame model. As the steady state that the controller should achieve is much more intuitive and practical,

$$e = \begin{bmatrix} d & \psi_{err} & v_y & \omega \end{bmatrix}^T \tag{2.21}$$

On a straight path, the curvature κ reduces to zero, and thus the steady state is simply,

$$e = \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix}^T$$

However, on a curved path (where $\kappa \neq 0$) the vehicle must have a non zero yaw rate r and a non zero lateral velocity v_y. To be more precise, for a vehicle to track a curved path, the yaw rate must match the angular rate at which the path turns, that is

$$\omega_{eq} = \kappa \cdot v_x \tag{2.22}$$

similarly, at steady state, $\dot{v}_y$ and $\dot{r}$ must be zero. Thus using equations 2.19, 2.20 and 2.22, we can compute the equilibrium value of v_y as

$$v_{y,eq} = \kappa\, v_x \left(L_r - \frac{m\, v_x^2\, L_f}{C_r^*\, L} \right) \tag{2.23}$$

we can verify here that on a straight path, both ω_{eq} and $v_{y,eq}$ reduces to zero, recovering the intuitive all-zero equilibrium condition. Here, since the lateral controller is designed independently of longitudinal motion, there is no need to explicitly control the longitudinal position s of the vehicle.

2.2.1 Controller design

In practice, lateral steering controllers are often structured as a combination of two control actions: a feedforward and a feedback control action,

$$\delta = \delta_{ff} + \delta_{fb} \tag{2.24}$$

The feedforward term δ_{ff} provides the nominal steering input required to follow a curved path at steady state. From the kinematic bicycle model, the steering angle required to negotiate a curve of curvature κ is given by

$$\tan(\delta_{ff}) = \kappa \cdot L \tag{2.25}$$

$$\delta_{ff} = \arctan(\kappa \cdot L) \tag{2.26}$$

where $L = L_f + L_r$ is the wheelbase. This is the classic geometric Ackermann steering relationship.

While the feedforward provides the correct steady-state steering, it has a few key limitations. Since it is derived from steady-state kinematics, it provides no mechanism for correction of initial lateral offsets and heading errors, or regulate dynamic states such as v_y and ω. To address these limitations, a feedback controller is introduced, which not only resolves the earlier limitations, but also compensates for disturbances, model mismatches, etc.

As established earlier, on a curved path the vehicle naturally requires non-zero values of ω_{eq} and $v_{y,eq}$. The feedforward handles the nominal steering input for this, but the feedback controller must be aware of this equilibrium otherwise it would interpret the steady-state cornering values as errors and attempt to steer them back to zero, partially canceling the feedforward. Therefore, the feedback acts on the *deviation* from the curved-path equilibrium,

$$\tilde{e} = \begin{bmatrix} d \\ \psi_{err} \\ v_y - v_{y,eq} \\ \omega - \omega_{eq} \end{bmatrix} \rightarrow \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \tag{2.27}$$

such that the full control law becomes

$$\delta = \delta_{ff} - K\,\tilde{e} \tag{2.28}$$

where K is the feedback gain vector. Generally, a normal PID feedback controller is sufficient, but model based approaches like LQR provide a more optimal performance.

2.2.1.1 PID based feedback control

A Proportional Integral Derivative (PID) controller is a classical feedback control strategy that computes the control input based on the current, accumulation of past and predicted future errors $e(t) = r(t) - y(t)$,

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau)\,d\tau + K_d \frac{de(t)}{dt}$$

where, K_p, K_i and K_d are the proportional, integral and derivative gains respectively.

2.2.1.2 LQR based feedback control

A Linear Quadratic Regulator (LQR) is a model-based optimal control strategy that computes the feedback control input by minimizing a quadratic cost function of the system states and control effort. Unlike heuristic methods, LQR systematically balances tracking performance and control effort using a mathematical model of the system. For a LTI system,

$$\dot{x}(t) = Ax(t) + Bu(t)$$
$$u(t) = -Kx(t)$$

the LQR cost function is defined as,

$$J = \int_0^\infty \left(x(t)^T Q x(t) + u(t)^T R u(t) \right) dt$$

such that the control input $u(t)$ minimizes the cost function J. Here,

$$K = R^{-1} B^T P$$

and P is the solution of the algebraic riccati equation,

$$A^T P + PA - PBR^{-1}B^T P + Q = 0$$

Since LQR is a model-based strategy, it relies on the accuracy of the system model used. This makes it particularly well-suited for lateral vehicle control, where state feedback can be used to systematically minimize deviations in position, heading and yaw dynamics while maintaining the needed control input efficiently.

2.2.1.3 Miscellaneous controllers

Apart from the PID and LQR based lateral controller there are other geometric or model-based controllers. The Stanley controller and the pure pursuit controller are two such examples of path tracking controllers.

The pure pursuit controller is a geometric controller which chooses a "look-ahead" point on the reference path and then steers the vehicle towards it. Meanwhile, the Stanley controller is built upon the pure pursuit setup, where the aim is to decrease the heading error and the lateral offset.

The biggest limitation of these approaches is that since these controllers are purely geometric in nature, there is no dynamic optimization of the input involved. Furthermore, they are often based on kinematic models of the system and as such does not account for the various complex dynamics acting on the system. Thus, in this work, the two controllers were not considered for the development of the adaptive lateral controllers.

The model predictive controller (MPC) is a model-based optimization approach similar to LQR. While being more robust than LQR, the MPC is several times more

resource intensive and not suitable for systems where quick response is needed for dynamic models. However, despite its flaws, the adaptive control strategy developed using the MPC remains a valuable approach worth understanding. Thus, this idea briefly presented in the Appendix section A.1.3.

2.3 Adaptive Control

Intuitively, adaptive control could be categorized as a control scheme which modifies the system behavior, in response to changes in the dynamics or due to disturbances. However, conventional feedback controllers also fulfill the same role. To present a more precise distinction, let us define the system process model in such a way that it is composed of two parts: a set of states that evolve on a faster time scale, and a set of parameters which vary at a slower one. Conventional feedback controllers concern themselves only with the former. Thus, in a pragmatic sense, we can define adaptive control as a special type of non linear feedback control scheme in which the control input is dependent on both the rapidly changing states as well as the slowly varying parameters [10].

Before proceeding, let's establish what is meant by a slowly varying parameter. In the case of a linear system, we usually write the state space model in the form,

$$\dot{x}(t) = A(\theta)x(t) + B(\theta)u(t) \tag{2.29}$$
$$y(t) = Cx(t) + Du(t) \tag{2.30}$$

where θ is the set of parameters that define the system matrices. If the system were to be non-linear, then it would be of the form,

$$\dot{x}(t) = f(x, \theta, u) \tag{2.31}$$
$$y(t) = h(x, u) \tag{2.32}$$

where the functions f and h are non linear system functions that are dependent on the states, parameters and the input. As stated, these parameters (θ) may either be initially known but slowly varying over time, constant but initially unknown or incorrectly assumed, or a combination of both. The adaptive controller must either estimate their true values or compensate for their effects through suitable control input.

2.3.1 Direct and indirect adaptive controllers

Earlier we defined adaptive controllers as controllers which adjusts the control input based on both changes in states and plant parameters. Based on how the plant parameters affect the control law, adaptive controller can be broadly classified into two categories:

- Indirect adaptive control: where the controller must first estimate the plant parameters before computing the control parameters. As such an indirect adaptive controller is a combinatorial setup of a feedback controller and a parameter estimator, thus, a variety of controllers are possible based on the choice of estimator and feedback control strategy used.

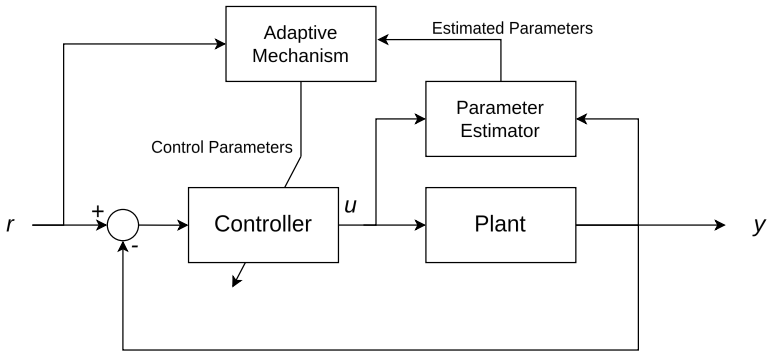

Figure 2.3: Block diagram for Indirect Adaptive Control. The parameter estimator block sends the estimation of the parameters to the adaptive mechanism block so that it can update the controller gains.

- Direct adaptive control: where the control strategy is to directly compute and update the control parameters. Since we are essentially skipping the estimation step, different types of direct controllers are possible based on how we relate control parameters to the controller performance.

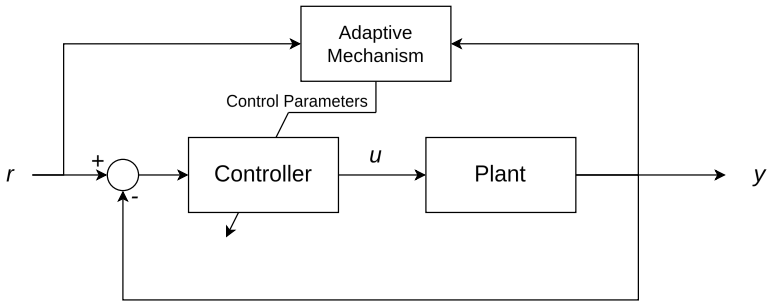

Figure 2.4: Block diagram for Direct Adaptive Control. The parameter estimator block is absent here, and the adaptive mechanism directly changes the controller gains based on output of the plant.

2.4 Model Reference Adaptive Control

The Model Reference Adaptive Control (MRAC) framework is one of the most well known approaches to direct adaptive control [10]. Figure 2.5 shows the block diagram for a typical MRAC controller, where K_{ff} and K_{fb} are the feed-forward and feedback gains respectively. The chief detail about this framework is that the desired performance of the plant is described using a reference model. An adaptation mechanism is then established, which uses the outputs of the reference model and the actual plant to directly update the feed-forward and feedback gains.

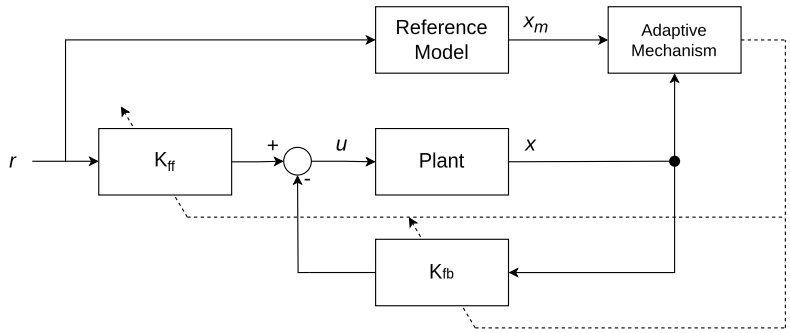

Figure 2.5: Block diagram for Model Reference Adaptive Control. The adaptive mechanism compares the outputs from the plant and the reference model to update the feed-forward and feedback gains.

There are essentially three basic approaches to designing an adaptation mechanism: the gradient approach, passivity theory and the Lyapunov approach [10]. For the scope of this thesis, we will be focusing on the Lyapunov based stability approach.

2.4.1 Adaptation based on Lyapunov's second method

Let us assume a simple plant process model and reference model of the form,

$$\dot{x}(t) = Ax(t) + Bu(t) \tag{2.33}$$

$$\dot{x}(t) = A_m x(t) + B_m r(t) \tag{2.34}$$

where, (A, B) and (A_m, B_m) are the system and input matrices for the plant and reference model respectively. From Figure 2.5, we can see that control law for the system can be described as,

$$u(t) = -K^*_{fb} x(t) + K^*_{ff} x_m(t) \tag{2.35}$$

Thus, the controlled system becomes,

$$\dot{x}(t) = (A - BK^*_{fb})x(t) + BK^*_{ff} r(t) \tag{2.36}$$

Ideally, this controlled system should match the reference system such that,

$$(A - BK^*_{fb})x(t) + BK^*_{ff} r(t) = A_m x(t) + B_m r(t) \tag{2.37}$$

This leads us to the implication that,

$$A - BK_{fb}^* = A_m \implies K_{fb}^* = \frac{A - A_m}{B} \tag{2.38}$$

$$BK_{ff}^* = B_m \implies K_{ff}^* = \frac{B_m}{B} \tag{2.39}$$

These values represent the ideal gains for the system so that it matches with the reference system. Thus, the control system that has to be designed must have gains that converge to these values at steady state. It can be observed that these ideal values are dependent upon the matrices A and B, both of which are matrices influenced by the system parameters.

Now, as mentioned before, the adaptation strategy is developed by comparing the outputs of the plant model against the reference model. Thus, let us define $e(t)$ as the error between the two outputs at time instant t,

$$e(t) = x(t) - x_m(t) \tag{2.40}$$

On differentiating this equation, we get

$$\begin{aligned}
\dot{e} &= \dot{x}(t) - \dot{x}_m(t) && (2.41)\\
&= (A - BK_{fb})x(t) - A_m x_m(t) + (BK_{ff} - B_m)r(t) && (2.42)\\
&= (A - BK_{fb})x(t) - A_m x_m(t) + (BK_{ff} - B_m)r(t) + A_m x(t) - A_m x(t) && (2.43)\\
&= A_m(x(t) - x_m(t)) + (A - A_m - BK_{fb})x(t) + (BK_{ff} - B_m)r(t) && (2.44)\\
&= A_m e + (A - A_m - BK_{fb})x(t) + (BK_{ff} - B_m)r(t) && (2.45)
\end{aligned}$$

By substituting the true values of the gains K_{fb}^* and K_{fb}^* here, we get

$$\begin{aligned}
\dot{e} &= A_m e + B(K_x^* - K_x)x(t) - B(K_r^* - K_x)r(t) && (2.46)\\
&= A_m e + B\tilde{K}_x x(t) - B\tilde{K}_r r(t) && (2.47)
\end{aligned}$$

Where, for simplification $K_x^* - K_x = \tilde{K}_x$ and $K_r^* - K_x = \tilde{K}_r$ represent the error in gains. Now, in order to develop an adjustment mechanism to push the gains to their ideal values, we must choose a suitable Lyapunov equation, which is dependent on not only $\tilde{K}_x$ and $\tilde{K}_r$ but also on e. For a single state plant, an apt equation would be,

$$V = \frac{1}{2}e^2 + \frac{1}{2\gamma}\tilde{K}_x^2 + \frac{1}{2\gamma}\tilde{K}_r^2 \tag{2.48}$$

where, γ is a constant scalar adaptation gain. We know that, for the system to be Lyapunov stable, $\dot{V} \leq 0$. Thus,

$$\begin{aligned}
\dot{V} &= e\dot{e} + \frac{1}{\gamma}\tilde{K}_x\dot{\tilde{K}}_x + \frac{1}{\gamma}\tilde{K}_r\dot{\tilde{K}}_r && (2.49)\\
&= e(A_m e + B\tilde{K}_x x(t) - B\tilde{K}_r r(t)) - \frac{1}{\gamma}\tilde{K}_x\dot{K}_x - \frac{1}{\gamma}\tilde{K}_r\dot{K}_r && (2.50)
\end{aligned}$$

Here, in order to guarantee the stability, the matrix the terms with $\tilde{K}_x$ and $\tilde{K}_r$ should cancel out. Thus, we get the adaptation equations,

$$\dot{K}_x = \gamma \cdot x(t) \cdot e \cdot sign(B) \quad and \quad \dot{K}_r = \gamma \cdot r(t) \cdot e \cdot sign(B) \tag{2.51}$$

Finally, the equations simplifies to,

$$\dot{V} = e^T A_M e \leq 0 \tag{2.52}$$

Hence, another condition, for a single state system to be stable is that A_m should be negative.

However, for a system with two or more states, the relevant equations and conditions change. As such, the Lyapunov function can be modified as,

$$V = \frac{1}{2} e^T P e + \frac{1}{2} \tilde{K}_x \Gamma_x^{-1} \tilde{K}_x^T + \frac{1}{2\gamma_r} \tilde{K}_r^2 \tag{2.53}$$

Where, Γ_x is the adaptation gain matrix for the state and γ_r is the scalar gain for the reference. Notice that $\tilde{K}_x$ is a matrix here since there are now more than one states. Moreover, P is a matrix which solves the Lyapunov equation $A_m^T P + P A_m = -Q$ for $Q > 0$. Thus, the derivative becomes:

$$\dot{V} = -\frac{1}{2} e^T Q e + e^T P B \tilde{K}_x x - e^T P B \tilde{K}_r r + \tilde{K}_x \, \Gamma_x^{-1} \, \dot{\tilde{K}}_x^T + \frac{1}{\gamma_r} \, \tilde{K}_r \, \dot{\tilde{K}}_r \tag{2.54}$$

Accordingly, the gain adaptation equations become,

$$\dot{K}_x = -\Gamma_x \, x \, e^T P B \tag{2.55}$$

$$\dot{K}_r = \gamma_r \, r \, e^T P B \tag{2.56}$$

Finally, for a unique positive definite solution P to exist, the matrix A_m must be Hurwitz. A matrix is Hurwitz if all its eigenvalues have strictly negative real parts. This ensures that the reference model is asymptotically stable.

The benefit of using the Lyapunov approach based MRAC is that the resulting controller not only directly converges the gains to their ideal values, but also guarantees stability while it does so. However, one chief requirement for this framework is that we must have a known linear or linearized reference model which satisfies the Hurwitz condition.

Another detail to note is that this approach works best when the reference model is linear in nature. Hence, for a non-linear reference model, we must first linearize the plant model around an operating point. Conveniently, since we are using the Frenet-Serret frame model, the operating point to linearize around would be (d, ψ) = (0, 0).

It should be noted that for the MRAC framework, only the reference model needs to be linear and not the actual plant model too. Since the goal is to directly update the gains for the controller by skipping the update step of the system matrices, the framework will still work for a non-linear actual plant model.

2.5 Parameter Estimation based Adaptive Controllers

Online determination of process parameters is a key element in adaptive control. Even in model reference adaptive control, we are implicitly performing parameter estimation by driving controller gains toward values that encode the unknown plant parameters. A typical indirect adaptive controller consists of two parts, the parameter estimator to compute the unknown parameters online, and a controller to use the updated parameters values and control the system.

To implement such estimators in practice, several recursive identification methods, such as deterministic least square or probabilistic estimation frameworks, can be employed. Among these, Kalman filter based estimators are particularly attractive due to their ability to explicitly account for uncertainties and provide optimal recursive updates. For this thesis, we shall explore different estimation frameworks based upon Extended Kalman Filters (EKFs) and Unscented Kalman Filters (UKFs), as discussed in the technical review paper by Amal Chebbi et al. [16].

2.5.1 EKF based estimators

Consider the discrete time non-linear state space model described by the process and measurement equations,

$$x_{k+1} = f(x_k, \theta_k, u_k) + w_k \tag{2.57}$$

$$y_k = h(x_k) + v_k \tag{2.58}$$

where x_k and y_k are the state and measurement vectors respectively. The process noise w_k and measurement noise v_k are zero mean gaussian noise with covariance matrices Q and R respectively. The vector θ represent the unknown parameters that the non-linear state evolution function f is dependent upon.

2.5.1.1 Simultaneous state and parameter estimation

In order to perform simultaneous estimation using EKF, we can define an augmented state vector: $z_k = \begin{bmatrix} x_k & \theta_k \end{bmatrix}^T$ such that the process model becomes,

$$z_{k+1} = \begin{bmatrix} f(x_k, u_k, \theta_k) \\ \theta_k \end{bmatrix} + w_k \tag{2.59}$$

Here we assume that the unknown parameters remains constant, aside from the addition of some noise: $\theta_{k+1} = \theta_k + w_{\theta,k}$, i.e., a random walk model.

Now, given that

$$F_k = \left.\frac{\partial f}{\partial z}\right|_{\hat{z}_{k|k}, \mathbf{u}_k} \quad and \quad H_k = \left.\frac{\partial h}{\partial z}\right|_{\hat{z}_{k|k}} \tag{2.60}$$

The prediction step of the Kalman Filter becomes,

$$\text{Estimation: } \hat{z}_{k+1|k} = f(\hat{z}_{k|k}, u_{k+1}) \tag{2.61}$$

$$\text{Covariance Propagation: } P_{k+1|k} = F_{k+1} \cdot P_{k|k}^T \cdot F_{k+1}^T + Q \tag{2.62}$$

and during the update step, we compute,

$$\text{Innovation: } r_{k+1} = y_{k+1} - h\left(\hat{z}_{k+1|k}\right) \tag{2.63}$$

$$\text{Kalman Gain: } K_{k+1} = P_{k+1|k} H_{k+1}^T \left(H_{k+1} P_{k+1|k} H_{k+1}^T + R\right)^{-1} \tag{2.64}$$

$$\text{State Update: } \hat{z}_{k+1|k+1} = \hat{z}_{k+1|k} + K_{k+1} r_{k+1} \tag{2.65}$$

$$\text{Covariance Update: } P_{k+1|k+1} = (I - K_{k+1} H_{k+1})\, P_{k+1|k} \tag{2.66}$$

Finally, the updated parameter estimates can be extracted from the extended state vector $\hat{z}_{k+1|k+1}$.

2.5.2 UKF based estimators

The Unscented Kalman Filter (UKF) is an alternative nonlinear state estimation method that is better suited for highly non linear systems. Instead of linearizing the system dynamics using jacobians, the UKF propagates a carefully chosen set of deterministic sample points (known as sigma points) through the nonlinear system, providing a more accurate approximation of the mean and covariance of the state distribution.

The scaling of sigma points in the UKF for n number of states is governed by the Merwe parameterization

$$\lambda = \alpha^2(n + \kappa) - n. \tag{2.67}$$

where, α controls the spread of the sigma points (smaller values result in a tighter distribution), κ is a secondary scaling parameter and β encodes the prior knowledge of distribution.

To begin the prediction, a set of $(2n+1)$ sigma points is generated from the posterior distribution:

$$\begin{aligned}
x_{k-1}^{(0)} &= \hat{x}_{k-1|k-1}, \\
x_{k-1}^{(i)} &= \hat{x}_{k-1|k-1} + \sqrt{n+\lambda}\left(P_{k-1|k-1}^{1/2}\right)_i, \quad i = 1, 2, \ldots, n, \\
x_{k-1}^{(i+n)} &= \hat{x}_{k-1|k-1} - \sqrt{n+\lambda}\left(P_{k-1|k-1}^{1/2}\right)_i, \quad i = 1, 2, \ldots, n.
\end{aligned}$$

The weights for the mean and covariance are computed as,

$$\begin{aligned}
W_0^m &= \frac{\lambda}{n+\lambda}, \\
W_0^c &= \frac{\lambda}{n+\lambda} + \left(1 - \alpha^2 + \beta\right), \\
W_i^m &= W_i^c = \frac{1}{2(n+\lambda)}, \quad i = 1, 2, \ldots, 2n.
\end{aligned}$$

Using these, the predicted values of mean and covariance becomes,

$$\hat{x}_{k|k-1} \approx \sum_{i=0}^{2n} W_i^m \, f\left(x_{k-1}^{(i)}\right),$$
$$P_{k|k-1} \approx Q_{k-1} + \sum_{i=0}^{2n} W_i^c \left(f\left(x_{k-1}^{(i)}\right) - \hat{x}_{k|k-1}\right)\left(f\left(x_{k-1}^{(i)}\right) - \hat{x}_{k|k-1}\right)^T.$$

To perform the update a new set of sigma points is generated from the predicted distribution,

$$\begin{aligned} x_k^{(0)} &= \hat{x}_{k|k-1}, \\ x_k^{(i)} &= \hat{x}_{k|k-1} + \sqrt{n+\lambda}\left(P_{k|k-1}^{1/2}\right)_i, \quad i = 1, 2, \ldots, n, \\ x_k^{(i+n)} &= \hat{x}_{k|k-1} - \sqrt{n+\lambda}\left(P_{k|k-1}^{1/2}\right)_i, \quad i = 1, 2, \ldots, n. \end{aligned}$$

Compute the desired moments:

$$\hat{y}_{k|k-1} \approx \sum_{i=0}^{2n} W_i^m \, h\left(x_k^{(i)}\right),$$
$$S_k \approx R_k + \sum_{i=0}^{2n} W_i^c \left(h\left(x_k^{(i)}\right) - \hat{y}_{k|k-1}\right)\left(h\left(x_k^{(i)}\right) - \hat{y}_{k|k-1}\right)^T,$$
$$P_{xy} \approx \sum_{i=0}^{2n} W_i^c \left(x_k^{(i)} - \hat{x}_{k|k-1}\right)\left(h\left(x_k^{(i)}\right) - \hat{y}_{k|k-1}\right)^T.$$

Finally, use the moments to calculate the Kalman gain and apply the correction:

$$\begin{aligned} K_k &= P_{xy} \, S_k^{-1}, \\ \hat{x}_{k|k} &= \hat{x}_{k|k-1} + K_k \left(y_k - \hat{y}_{k|k-1}\right), \\ P_{k|k} &= P_{k|k-1} - K_k \, S_k \, K_k^T. \end{aligned}$$

2.5.2.1 Multi-scaled UKF

The UKF uses the sigma points to capture the mean and the covariance of random variables propagated through a nonlinear system. The spreading of these sigma points is crucial for the filter to capture the actual behavior of the states. The general formulation of the UKF uses a uniform scaling set for all states. The paper by Amit Levy et al. introduces the concept of independently tunable scaling parameters for each state dimension [18].

Based on this, we can implement an extended scaling set that can be used to independently tune the spread of each state, maintaining theoretical accuracy with negligible additional computational cost.

$$\lambda_i = \alpha_i^2(n + \kappa_i) - n \tag{2.68}$$

Accordingly, the covariance compensation factor becomes,

$$\gamma = 1 - (\Pi\alpha_i)^{2/n} + \beta \tag{2.69}$$

2.5.2.2 Variations in Kalman filter based estimation

The estimation methods discussed till now are also known as joint estimation methods [16] where a single filter is used to estimate both states and parameters. Another framework for estimation methods is the dual estimation approach, where the parameters and state are estimated separately one after the other.

While this method has not been implemented in this work, the reader can find the theory behind it in the appendix section A.1.1.

2.5.3 Adaptive Kalman filter based estimation

In a standard augmented Kalman filter setup, the augmented state covariance matrix P converges over time based on the values of the matrices Q and R. However, if there is a change in the parameters of the system, like a change in the cornering stiffness or load mass, the state covariance does not grow again to account for them, as a result of which the re-convergence of the parameter estimates is very slow.

An adaptive Kalman filter is a special type of filter where the process and measurement noise covariance matrices Q and R are updated online based on incoming measured data [19].

2.5.4 Limitations to estimation methods

When we look at the cartesian frame dynamic vehicle model (equations 2.10 - 2.14), we can clearly observe that the system is currently *under-determinate*. This means that there are more unknown parameters than equations. Furthermore, most of the unknown parameters appear in the last two equations, pertaining to the evolution of $\dot{v}_y$ and $\dot{\psi}$, which makes accurately estimating each parameter much more difficult.

Finally, another issue is lack of *persistent excitation* in the system. In order to accurately estimate the unknown parameters in a system, the input signals should persistently excite or stimulate all relevant system dynamics. In a closed-loop system, there exists a trade-off between controller performance and parameter estimation. A controller will aim to regulate the system states to a desired steady-state value or towards zero. While this achieves stability and performance objectives, it also unintentionally reduces the level of persistent excitation in the system. Thus, even if a system is theoretically observable, closed-loop operation can make parameter estimation infeasible due to lack of excitation.

3

Problem Formulation

Earlier, in the section 2.3, we defined an adaptive controller as a controller that provides optimal feedback based on changes in states and parameters of the system. The section also broadly explored different direct and indirect strategies to the adaptive control problem. Subsequently, we also encountered the various caveats to each approach, which forces us to place certain limitation and restrictions on our system model. Thus, before we narrow down to a problem formulation, we must first clearly establish what these restrictions are.

3.1 Assumptions and Simplification

Since the direct and indirect approaches to adaptive control differ significantly, multiple separate assumptions and simplifications must be made for each approach. However, some of these remain the same for both,

- It is assumed that the single track dynamic bicycle model perfectly represents the dynamics of the vehicle, instead of using a more complex fully dynamic model,

- The vehicle dynamics are continuous in time and are represented by the equations 2.10 - 2.14 in the Cartesian frame, and by equations 2.16 - 2.20 in the Frenet-Serret frame,

- The unknown parameters θ that influence the plant dynamics are assumed to be constant with a poor initial estimation, unless specified to be otherwise,

- The velocity of the vehicle along the longitudinal axis v_x is constant, unless specified to be otherwise, and known. Consequently, the longitudinal frictional forces F_{lf} and F_{lr} are removed from the model as well, since their net effect is implicitly captured by the maintenance of a constant forward velocity.

3.1.1 Direct adaptive control

As discussed earlier, the main assumption we must take is that we have an accurate reference model that not only demonstrates the ideal behavior of the vehicle, but is also linear in nature and Hurwitz stable.

Furthermore, the performance of the controller is dependent upon the chosen adaptive gain matrix Γ and the initial values of the controller gain K. A degree of trial and error is necessary to raise the efficiency of this control approach.

3.1.2 Indirect adaptive control

From the (equations 2.10 - 2.14), we can observe that the vehicle dynamics is heavily dependent on multiple potentially unknown parameters. Since our system is underdeterminate and due to poor excitation in the system, it is not viable to estimate all the parameters from the aforementioned set of equations. Thus, we must settle on a smaller subset of unknown parameters to estimate.

The first simplification we can make is regarding the tire stiffness and normal force applied. As can be seen from (equations 2.10 - 2.14), the lateral cornering stiffness of the front and rear wheels (C_f and C_r) always appear together with the normal forces acting on them ($F_{z,f}$ and $F_{z,r}$). In practice, there is no real benefit in estimating the values of all these parameters separately, so we can clump these together such that,

$$C_f^* = C_f F_{z,f} \quad and \quad C_r^* = C_r F_{z,r}$$

Furthermore, we can also assume that the ratio of these clumped friction parameters (C_f^*/C_r^*) is always constant. The application of this assumption will be expanded upon in the later chapters.

Thus, for the scope of this thesis, we shall restrict ourselves to only estimating the clumped friction parameters (C_f^* and C_r^*) and the wheelbase parameters L_f and L_r.

3.2 Aim

The aim of this thesis is to use the concepts discussed in chapter 2 and assumptions in section 3.1, to develop and implement a family of adaptive lateral controllers, and to test them in varying conditions. The thesis shall also compare the performance of these controllers with one another and finally also list out the merits and drawback of each approach.

To guide this work, the following research questions are addressed:

- Which adaptive control methods are most suitable for lateral control of autonomous heavy vehicles?

- How effectively do the proposed controllers adapt to time-varying vehicle dynamics, and how robust are they across changes in parameter values?

- What are the computational requirements and onboard implementation constraints of the proposed adaptive controllers, and how generalisable are they across different truck configurations?

4

Methodology

Based on the discussions in the previous chapters, we can clearly see that an adaptive lateral controller is but a lateral controller with an adaptive sub-step in it. This chapter will discuss how a Lyapunov based direct and an Kalman filter based indirect adaptive lateral controller can be developed.

4.1 The Adaptive Lateral PID Controller

The adaptive lateral PID controller is a type of direct adaptive controller developed by combining a lateral PID controller with the MRAC framework. Using the Lyapunov's second method based adaptation law we shall be updating the gains of the PID controller online.

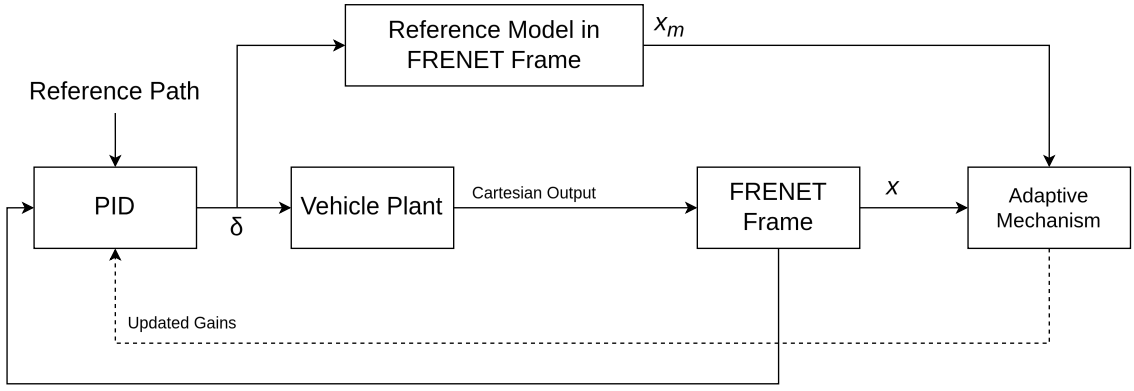

Figure 4.1: Block diagram for MRAC based Adaptive lateral PID control flow. The PID gains are updated by the adaptive mechanism block on comparing the outputs of the reference model and plant.

For the lateral control problem, the primary tracking objective is to drive the lateral offset d to zero. Thus, the complete PID action is applied to the state d. However, as discussed in section 2.2, the feedback must act on the full deviation vector $\tilde{e}$. Therefore, additional proportional feedback terms are introduced on the remaining

error components, giving us the control law

$$\delta = \delta_{ff} - \delta_{fb} \tag{4.1}$$

$$= \delta_{ff} - \Big(K_p\, d + K_i\, \xi + K_d\, \dot{d} + K_\psi\, \psi_{err} + K_{v_y}(v_y - v_{y,eq}) + K_\omega(\omega - \omega_{eq}) \Big) \tag{4.2}$$

where, d is the lateral offset, ξ is the integral of the lateral error ($\int d\ dt$), e_ψ is the heading error, v_y is the lateral velocity and $\dot{\psi}$ is the yaw rate. For a compact representation, the control law can be re-written as,

$$\delta = \delta_{\text{ff}} - K^T \phi \tag{4.3}$$

with

$$K = \begin{bmatrix} K_p & K_i & K_d & K_\psi & K_{v_y} & K_\omega \end{bmatrix}^T \tag{4.4}$$

and

$$\phi = \begin{bmatrix} d & \xi & \dot{d} & \psi_{err} & (v_y - v_{y,eq}) & (\omega - \omega_{eq}) \end{bmatrix}^T \tag{4.5}$$

Recalling equation 2.47, we can model the error dynamics of our system as,

$$\dot{e} = A_m e + B\tilde{K}_x x(t) - B\tilde{K}_{ref} r(t)$$

with A_m being the Hurwitz stable reference model matrix, and $K_x^* - K_x = \tilde{K}_x$ and $K_{ref}^* - K_x = \tilde{K}_{ref}$. Since, we are aiming to drive the equilibrium deviation $\tilde{e}$ to zero at steady state, we can rewrite the error dynamics as,

$$\dot{e} = A_m e + B\tilde{K}_x^T \phi$$

In order to derive the adaptive law, let us consider the updated form the Lyapunov equation 2.53, where we can ignore the term $\tilde{K}_r$ as $r(t) = 0$,

$$V = e^T P e + \frac{1}{2}\tilde{K}^T \Gamma^{-1} \tilde{K} \tag{4.6}$$

where P is the positive definite solution of the Lyapunov equation and Γ is a positive definite adaptation gain matrix. Taking the derivative of the Lyapunov equation yields,

$$\dot{V} = -e^T Q e + \tilde{K}^T \left(\Gamma^{-1} \dot{\tilde{K}} - \phi e^T P B \right) \tag{4.7}$$

In order to ensure $\dot{V} < 0$, the adaptive gain update law becomes,

$$\Gamma^{-1} \dot{\tilde{K}} = \phi e^T P B$$

$$\dot{\tilde{K}} = \Gamma \phi e^T P B$$

Substituting $\sigma = e^T P B$ and since K^* is constant, we have the gain update as:

$$\dot{K} = \Gamma \phi \sigma \tag{4.8}$$

However, if the product $\phi \cdot \sigma$ were to be small yet non zero for long operation durations, the gains could slowly drift away without bounds [15]. Thus, a damping factor is added such that,

$$\dot{K} = \Gamma \phi \sigma - \sigma_0 K \tag{4.9}$$

where $\sigma_0 > 0$ is a small leakage coefficient used to bound the adaptive gains and improve numerical robustness.

From the equations, we can clearly observe that the performance of the controller is dependent upon the tuning of the constant adaptive gain matrix Γ and the reference model chosen. The choosing of the reference model is also a main caveat to this approach. The reference model should not only be a linear function in structure, but also, the matrix A_m must be Hurwitz. One way of doing this would be to design a reference model using pole-placement. Subsequently, the adaptive lateral PID controller can be tuned by adjusting Γ and the poles of the reference model.

4.2 Adaptive Lateral LQR Controller

As discussed before, indirect adaptive controllers are always a combination of a parameters estimation and a feedback controller. In this work, we will be focusing on the implementation of an adaptive lateral LQR controller which uses a Kalman filter for parameter estimation, more precisely, we will be investigating the variations of extended Kalman filters (EKFs) and unscented Kalman filters (UKFs).

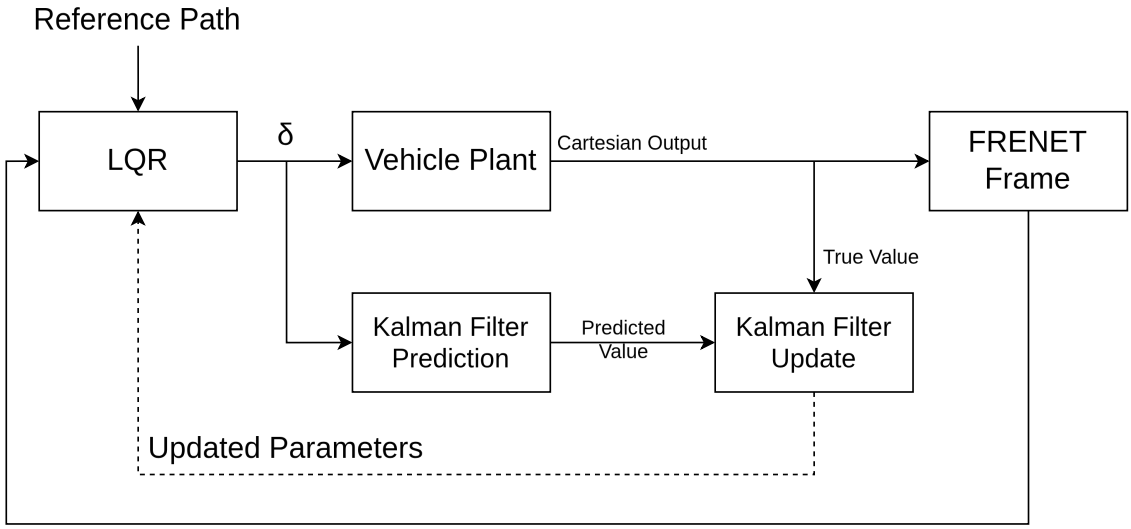

Figure 4.2: Block diagram for Adaptive lateral LQR control flow. The Kalman filter estimates the parameters online and updates the estimates when new measurements are received. The updated parameters are shared to the LQR controller where it updates its gains to adapt.

4.2.1 Implementation of Constraints

As discussed in section 2.5.4, a significant hurdle in the parameter estimation step of an indirect adaptive controller is the existence of an under-determinate system and lack of persistent excitation. While reducing the number of parameters to be estimated (section 3.1.2) helps with improving estimation efficacy, it is still not enough. This is why the implementation of constraints is important.

In the work by Sierra C, et al., a "beta-less method" is discussed where the cornering stiffness of the tires are connected by a ratio to determine each other [17]. In this work, we shall use a slightly different approach from the paper. The two constraints we will be adding to our estimation problem are, $C_f^*/C_r^* = \nu$ and $L_f + L_r = L$, such that

$$C_f^* - \nu C_r^* = 0 \tag{4.10}$$

$$L - (L_f + L_r) = 0 \tag{4.11}$$

We can also expand equation 4.10 as,

$$\nu = \frac{C_f^*}{C_r^*} = \frac{C_f F_{z,f}}{C_r F_{z,r}}$$

where, the normal forces $F_{z,f}$ and $F_{z,r}$ can be computed using torque balance equations such that,

$$\nu = \frac{C_f^*}{C_r^*} = \frac{\frac{mgL_r}{L_f + L_r}}{\frac{mgL_f}{L_f + L_r}} = \frac{L_r}{L_f} \tag{4.12}$$

Thus, the ratio between the clumped cornering stiffness parameters is equal to the ratio between the front and rear wheelbase lengths.

4.2.1.1 Soft constraints

Soft constraints are constraints that are only required to be approximately satisfied rather than be exactly satisfied. We might want to implement such constraints in cases where the constraints are heuristic rather than rigorous, or in cases where the constraint function has some uncertainty or fuzziness.

Thus, the equation 4.10 becomes,

$$C_f^* - \nu C_r^* \approx 0 \tag{4.13}$$

And we can use this to perform a scalar pseudo Kalman update after the estimation update step is done. For this, let us assume an observation vector $h_r = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & -\nu \end{bmatrix}$ where the parameters to be estimated are placed at the end of the vector. Then the update step would be of the form,

$$S_\nu = h_\nu P h_\nu^\top + R_\nu \tag{4.14}$$

$$K_\nu = \frac{P h_\nu^\top}{S_\nu} \tag{4.15}$$

$$e_\nu = -(C_f^* - \nu C_r^*) \tag{4.16}$$

$$z^+ = z + K_\nu \cdot e_\nu \tag{4.17}$$

$$P^+ = P - K_\nu K_\nu^\top S_\nu \tag{4.18}$$

where R_ν is the measurement noise variance of this pseudo measurement step.

Unlike the hard constraint approach (see appendix section A.1.2), which forces the parameters to take a certain relation, the soft constraint approach instead pushes the estimated values to a certain numerical relation, thereby significantly decreasing the state uncertainties.

Thus, as illustrated in Figure 4.3, this layer of soft constraint enforcement acts as a pseudo update step after the prediction step by the Kalman filter.

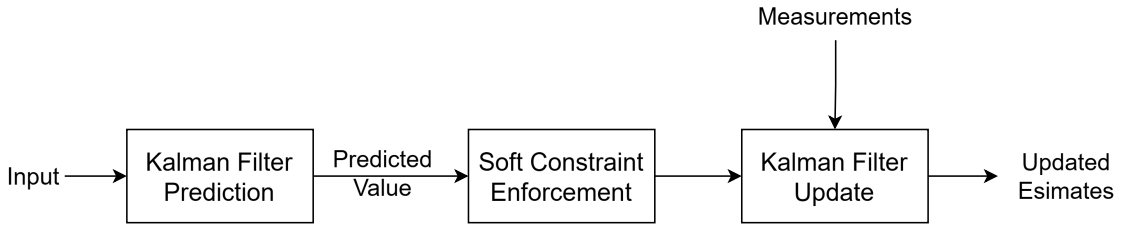

Figure 4.3: Block diagram showing the layer of soft constraint enforcement between prediction and update step in the Kalman filter. Effectively two updates happen when an estimation is performed.

4.2.2 Adaptive Q based Kalman filter

As described in Section 2.5.3, one established approach to developing a Kalman filter capable of adapting to changes in unknown parameters is to employ an adaptation law that dynamically scales the process noise covariance matrix Q. The underlying idea is that when the estimated value of a parameter θ_i begins to drift away from its previously converged value, for example ,due to a change in the vehicle's effective cornering stiffness caused by a shift in mass, or due to traversing a low-friction path, the corresponding entry $Q[\theta_i, \theta_i]$ should be temporarily inflated. This increases the associated value of the state uncertainty matrix P, allowing the filter to instead trust the new measurements more heavily and re-converge to the shifted parameter value quickly.

4.2.2.1 Parameter drift rate adaptive Q

In this thesis, we explore a strategy in which Q is updated based on the *drift rate* of each estimated parameter during a time interval. To detect a change in a given parameter θ_i, a buffer is used to store the last N estimated values of $\hat{\theta}_i$. Using this buffer, we compute the drift between the newest and oldest estimates over the capture window,

$$\Delta\hat{\theta}_i = \left|\hat{\theta}_{i,\text{newest}} - \hat{\theta}_{i,\text{oldest}}\right|$$

This drift is then normalized by the current uncertainty of the estimate,

$$\rho_i = \frac{\Delta\hat{\theta}_i}{\sigma_{\theta_i} + \epsilon}, \qquad \sigma_{\theta_i} = \sqrt{P[\theta_i, \theta_i]}$$

where ϵ is a small constant introduced to avoid division by zero. The normalized drift ρ_i has a clear interpretation: if $\rho_i \ll 1$, the estimate is stable within its own uncertainty band, whereas if $\rho_i > \rho_{\text{threshold}}$, the estimate is persistently moving away from its previous value, indicating an unresolved mismatch between the model and the true system.

Based on this metric, the adaptation law for the process noise covariance becomes

$$Q[\theta_i, \theta_i]_k = \begin{cases} \min\Big(Q[\theta_i, \theta_i]_{k-1} + \gamma \cdot q_\theta^0 \cdot \rho_i,\ Q_{\max}\Big) & \text{if } \rho_i > \rho_{\text{threshold}} \\ q_\theta^0 + \alpha \cdot \Big(Q[\theta_i, \theta_i]_{k-1} - q_\theta^0\Big) & \text{otherwise} \end{cases} \tag{4.19}$$

where q_θ^0 is the baseline process noise for parameter θ_i, γ is the inflation gain, and $\alpha \in [0, 1)$ is the decay rate. In words, when the normalized drift exceeds the threshold, the process noise Q is inflated in proportion to both the drift magnitude and the baseline noise level, up to a maximum of $Q_{\max}$; once the drift subsides, the process noise decays geometrically back toward its baseline value q_θ^0. The figure 4.4 neatly illustrates the work flow for this adaptive Q framework.

As the process noise covariance increases in response to a detected parameter change, the corresponding entry of the state uncertainty matrix P also increases, allowing the estimator to adapt quickly to the new parameter value. In Appendix A.1.4, we compare the re-convergence performance of an EKF with and without this adaptive layer.

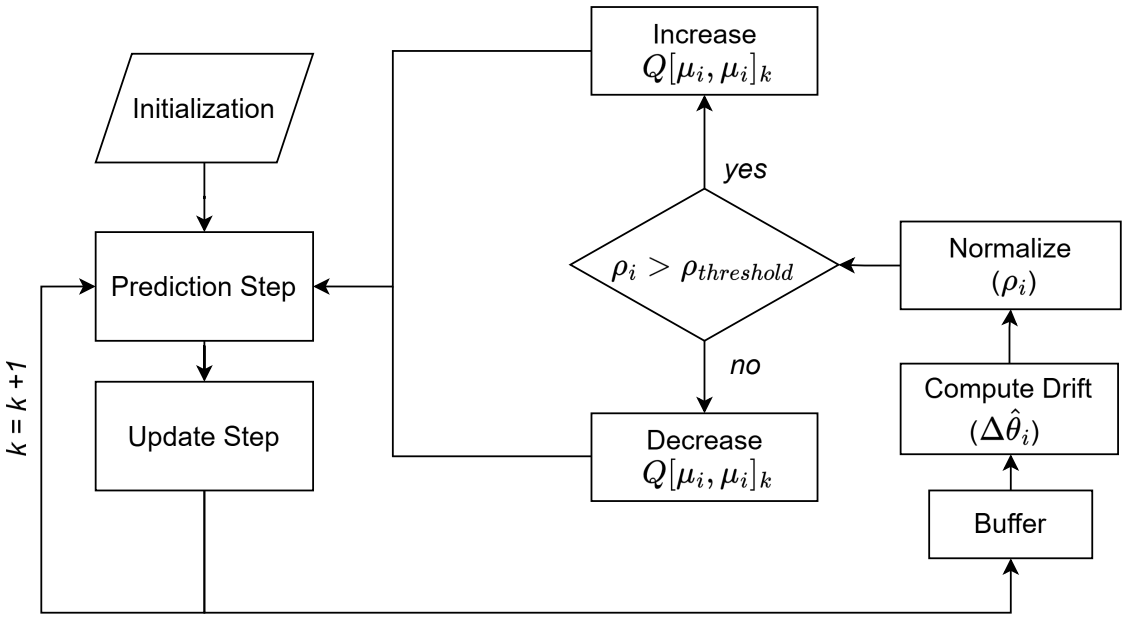

Figure 4.4: Block diagram showing the for the parameter drift rate adaptive Q framework. Effectively two updates happen when an estimation is performed.

4.2.3 LQR based feedback control

For the linearized lateral error dynamics derived from the Frenet-frame bicycle model,

$$\dot{\tilde{e}}(t) = A\,\tilde{e}(t) + B\,\delta_{fb}(t)$$

where $\tilde{e} = [d,\ \psi_{err},\ v_y - v_{y,eq},\ r - r_{eq}]^T$ is the deviation from the curved-path, and the matrices A and B encode the linearized vehicle dynamics, the LQR minimizes the infinite-horizon cost function,

$$J = \int_0^\infty \left(\tilde{e}(t)^T Q\,\tilde{e}(t) + \delta_{fb}(t)^T R\,\delta_{fb}(t)\right) dt$$

The weight matrices $Q \succeq 0$ and $R \succ 0$ are tunable parameters which decides how aggressively the state deviations and control input are penalized. The optimal

feedback gain is given by,

$$K = R^{-1}B^{T}P$$

where P is the unique positive-definite solution of the continuous algebraic Riccati equation (CARE),

$$A^{T}P + PA - PBR^{-1}B^{T}P + Q = 0$$

Since the matrices A and B depend on the vehicle parameters, the performance of the LQR controller is directly tied to the accuracy of the estimated values of the unknown parameters provided by the Kalman filter.

5

Results and Discussions

5.1 Introduction to scenarios

Before we look at the performance of the various controllers discussed in the previous chapters, let us first understand the background behind these tests, starting with the vehicle model used.

The vehicle modeled in this work is a *Volvo FH 8x4 rigid truck* operating at a gross vehicle weight (GVW) of approximately 40 tonnes. The 8×4 drivetrain layout comprises two front steering axle groups and two rear driven axle groups. As stated in the section 3.1, this multi-axle configuration is simplified into an equivalent dynamic bicycle model with a single front and a single rear axle. The relevant vehicle parameters for the test can be listed as,

Parameter	**Symbol**	**Value**	**Unit**
Vehicle mass	m	40 000	kg
Yaw moment of inertia	I_z	500 000	$\mathrm{kg\,m^2}$
Front axle distance (CoM to front)	L_f	2.8	m
Rear axle distance (CoM to rear)	L_r	1.8	m
Total effective wheelbase	$L = L_f + L_r$	4.6	m
Front axle normal load	$F_{zf} = mg\,L_r/L$	153 504	N
Rear axle normal load	$F_{zr} = mg\,L_f/L$	238 896	N
Front normalized cornering stiffness	C_f	≈ 5.21	1/rad
Rear normalized cornering stiffness	C_r	≈ 5.02	1/rad

Table 5.1: Vehicle parameters for Volvo FH 8×4 rigid truck (equivalent bicycle model)

The tests will be conducted for each controller individually. For the indirect adaptive lateral controllers, the system will be tested on two scenarios,

- **Scenario 1:** The truck would be simply made to take laps around a closed reference path, with the parameters C_f^*, C_r^*, L_f and L_r set to 80% of their true values. The goal of this scenario is to test the performance of base Kalman filter based estimator and LQR combination when provided with a poor initial

estimate of the necessary parameters.

- **Scenario 2:** Same as the previous scenario, however, at a certain point during the simulation the longitudinal velocity of the vehicle would be decreased significantly, after which the mass of the vehicle would be increased (to simulate loading of the vehicle) and then longitudinal velocity would return back to its original value. After some time, this sequence would be performed again, but the mass of the vehicle would be reduced (to simulate unloading). The goal of this scenario is to test the performance of adaptive Kalman filter based estimator.

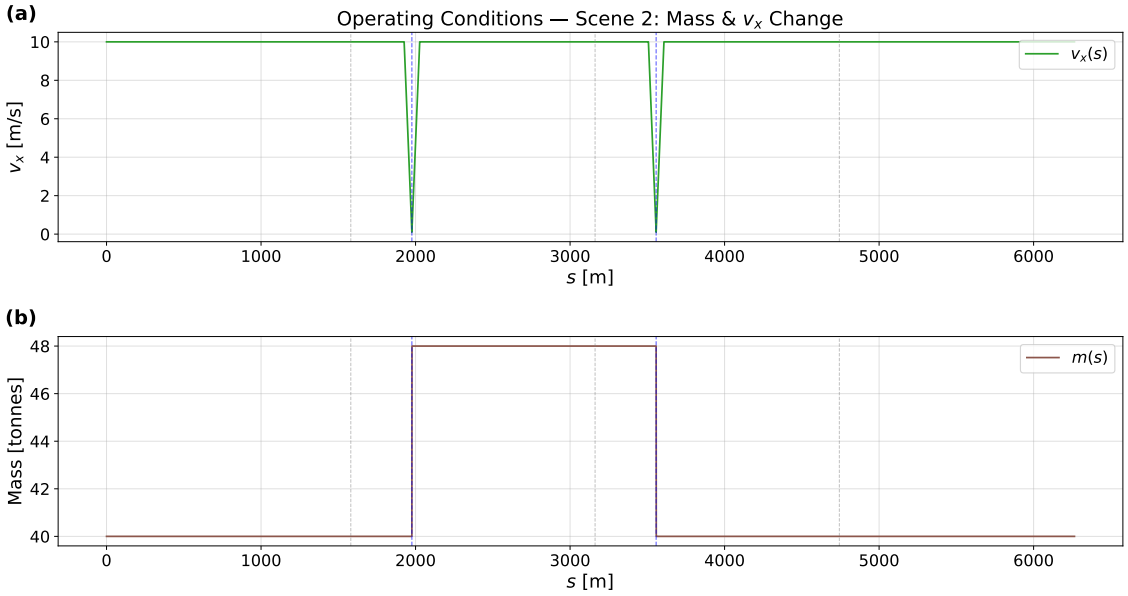

Figure 5.1: Plot showcasing the operating conditions for the second scenario. Subplot (a) shows the change in velocity whenever the vehicle is being loaded or unloaded, which is shown in subplot (b)

- **Scenario 3:** Same as scenario 1, but this time, a third of the path is considered to be muddy, where the true values of the lateral cornering stiffness (C_f and C_r) of the front and rear tire would be reduced. At the same time, the longitudinal velocity of the vehicle will also be reduced slightly. The velocity is returned to the normal value when the vehicle is out of the muddy portion.

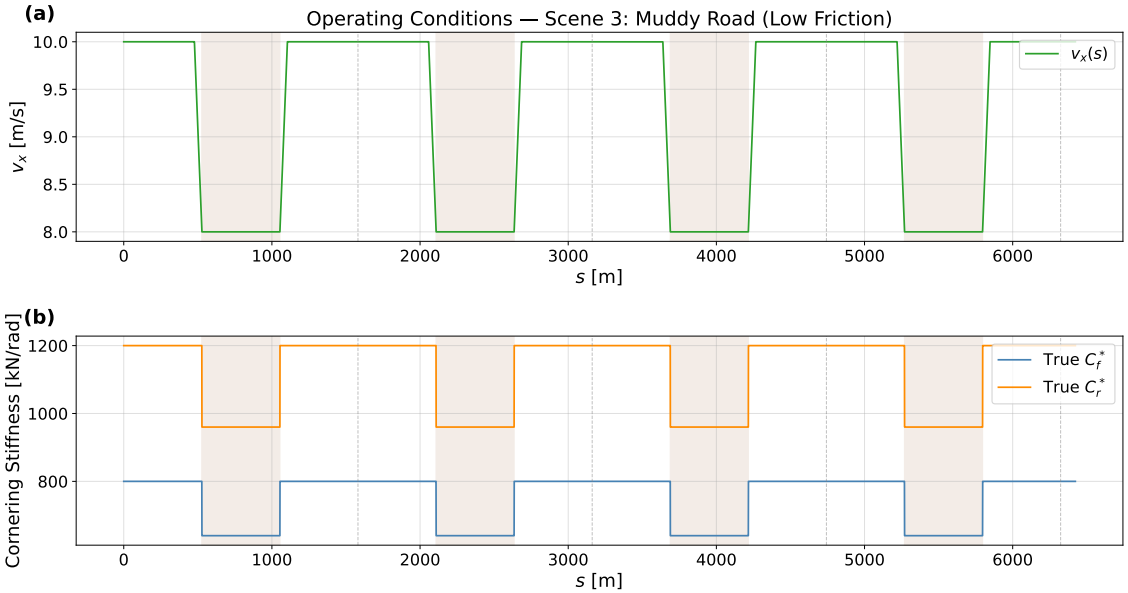

Figure 5.2: Plot showcasing the operating conditions for the third scenario. Subplot (a) shows the change in velocity whenever the vehicle when the truck enters a muddy zone, which is shown as a shaded area in the plots. The subplot (b) shows the true values of the cornering stiffness at different parts of the track

Meanwhile, for the direct adaptive lateral PID controller, the test performed would be similar to the first scenario for the indirect adaptive controller case, however, the initial parameter conditions would be different. The tuning parameters of the controller, i.e., the reference model poles and the adaptation constants vector Γ, will be chosen via trial and error. While it is understandable that finding an optimal set of these tuning parameters is crucial yet difficult, the aim is to choose a set of values which can showcase the performance of the controller properly.

5.2 Indirect Adaptive Lateral Control

For the indirect approach to adaptive lateral control, four variations of Kalman filters were investigated in total: normal EKF and UKF based estimators, and their respective adaptive versions. The simulation was performed for a total of 35,000 steps with a time step of 0.02 seconds. Within this duration, the vehicle was able to successfully complete three laps around the chosen track. Furthermore, the parameters (C_f^*, C_r^*, L_f and L_r) were given a prior estimate of 0.8 times their true values.

5.2.1 Scenario 1: Normal vehicle motion

For an indirect adaptive controller, the performance must be evaluated on two domains, that is, the accuracy of the estimated parameter and the speed at which it converges to the true values, and the performance of the LQR controller which uses these estimated parameters.

5.2.1.1 Parameter estimator performance

From the figures 5.3 - 5.6, it can be observed that both the extended and unscented Kalman filters estimate and converge the set of parameters to their true values early on in the simulation (see table 5.2). By the end of the simulation, the performance of the estimators can be noted in the table 5.3.

	EKF			**UKF**		
Param	**Step**	**Time (s)**	**s (m)**	**Step**	**Time (s)**	**s (m)**
C_f^*	24	0.48	4.80	8	0.16	1.69
C_r^*	24	0.48	4.80	8	0.16	1.69
L_f	1	0.02	0.25	1	0.02	0.17
L_r	2	0.04	0.48	2	0.04	0.43

Table 5.2: Convergence Speed: Steps to Reach $\leq 5\%$ Error. It can be seen that both the filters converge to the true value of the parameters within the first few time steps.

			EKF		**UKF**	
Param	**True**	**Initial**	**Final Est**	**Err (%)**	**Final Est**	**Err (%)**
C_f^*	800 000	640 000	810 348.63 ± 4 283.42	1.29	812 758.30 ± 5 008.58	1.59
C_r^*	1 200 000	960 000	1 215 522.78 ± 6 421.35	1.29	1 219 138.07 ± 7 509.69	1.59
L_f	2.80	2.24	2.82 ± 0.05	0.72	2.83 ± 0.06	0.91
L_r	1.80	1.44	1.78 ± 0.05	1.12	1.77 ± 0.06	1.42

Table 5.3: Comparison of EKF and UKF final Parameter Estimation Results. The sub 2% errors and small standard deviation proves how precise the estimator are

From the plots and the tables, it is clear that both the EKF and UKF based parameter estimator perform equally well and thus it is difficult to recommend one

estimator over the other. It should be noted that an EKF is computationally less expensive, requiring only a first order Jacobian evaluation, unlike computing $2n+1$ sigma point propagation per set in the UKF. However, for a highly non-linear system, it is difficult to accurately derive a Jacobian linearization. In such cases, it is much more efficient to use a UKF based estimator.

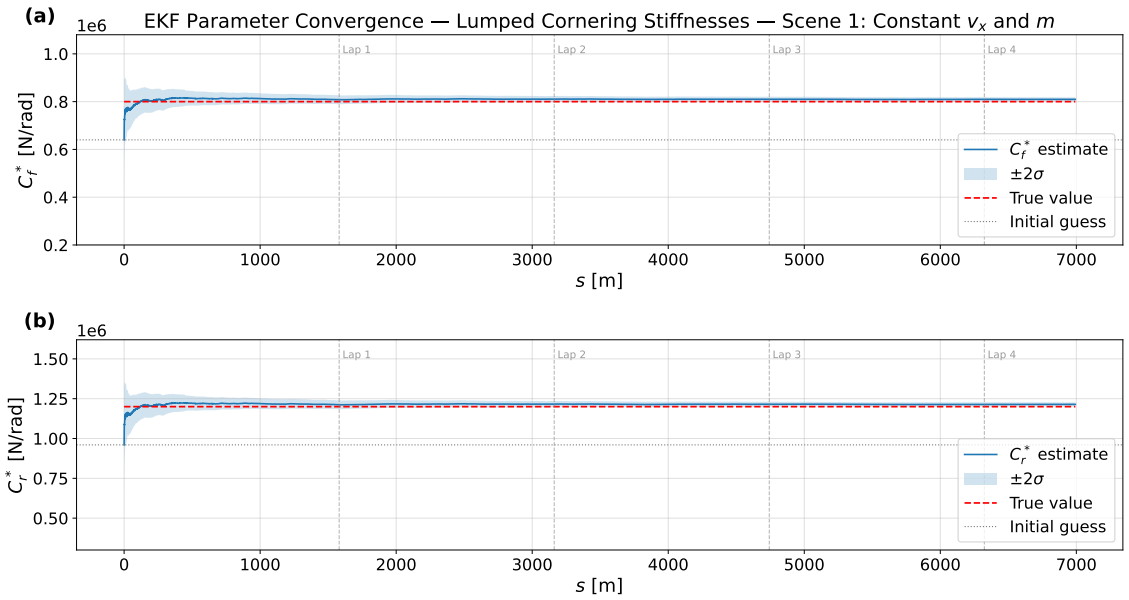

Figure 5.3: Plots showing estimation of C_f^* (a) and C_r^* (b) using EKF. The black dotted line indicates the original value and the red dashed line indicates the true value. The shaded blue region represents the 2σ standard deviation band.

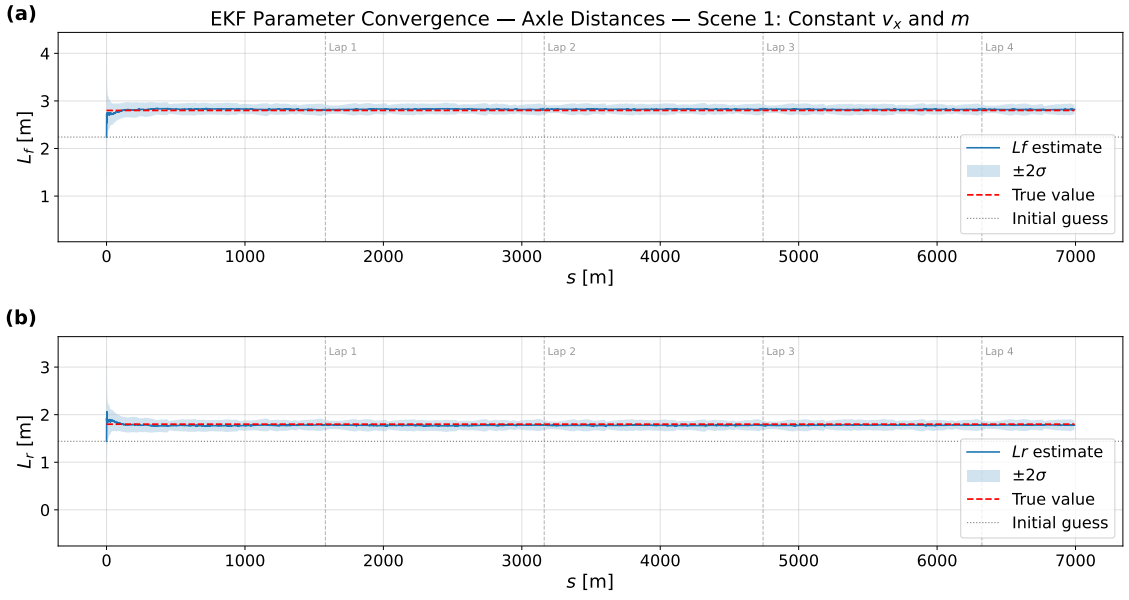

Figure 5.4: Plots showing estimation of L_f (a) and L_r (b) using EKF.

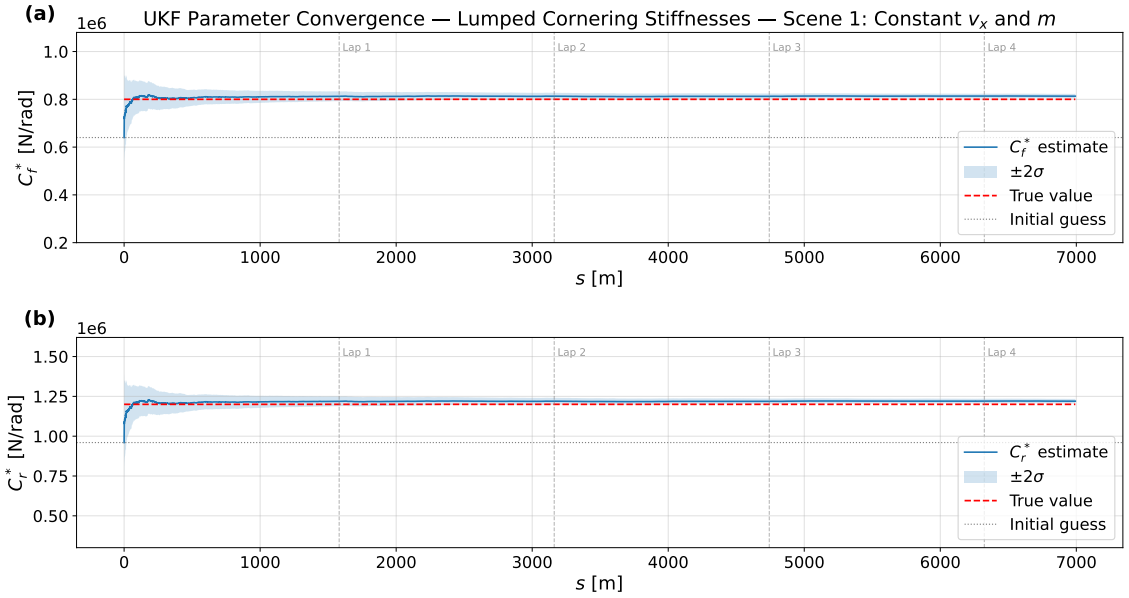

Figure 5.5: Plots showing estimation of C_f^* (a) and C_r^* (b) using UKF. The plot details are same as 5.3, furthermore the estimation performance and convergence speed is also similar to the EKF based approach.

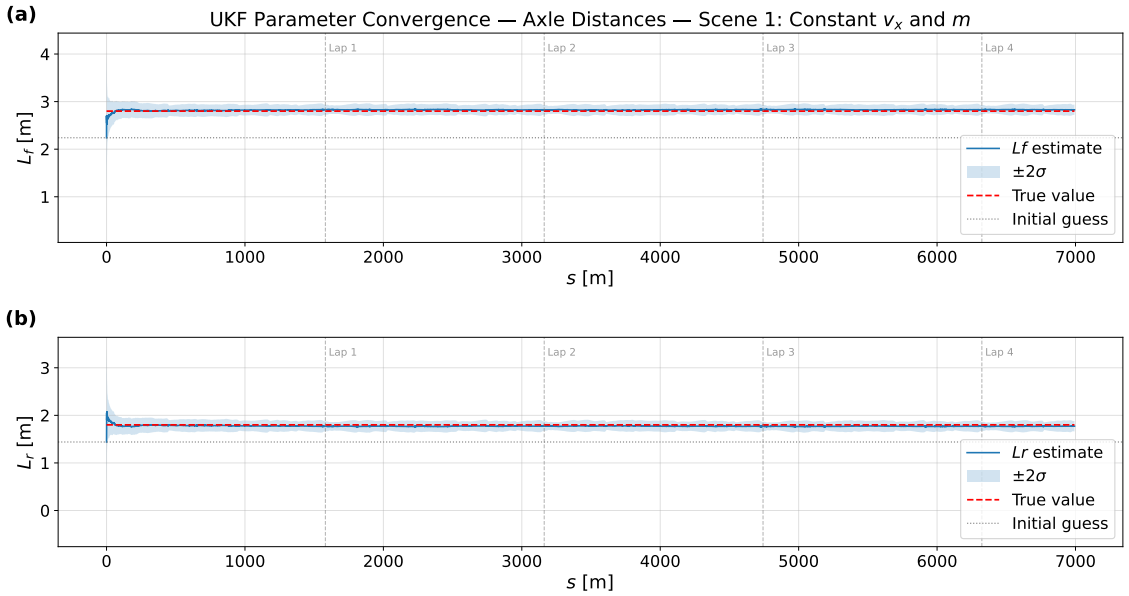

Figure 5.6: Plots showing estimation of L_f (a) and L_r (b) using UKF.

5.2.1.2 Lateral LQR performance

As explained in section 4.2.3, the estimated values of the parameters from the Kalman filter based estimators are fed into the Linear Quadratic Regulator (LQR), so as to update the gains of the controller. As the estimators converge towards the true values early enough, the gains of the LQR also reach a steady state value at the same time (figures 5.7 and 5.8). The figures also show the gains of a non adaptive

LQR controller and we can observe the difference between the true gains and initial gains.

	$K[d]$	$K[e_\psi]$	$K[v_y]$	$K[\dot{\psi}]$
EKF	0.5474	3.3842	0.1048	0.5138
UKF	0.5474	3.3848	0.1049	0.5141

Table 5.4: Final LQR Gain Matrix K: EKF vs UKF. Since the final estimates from both approaches are almost equal in value, the final gains are virtually the same as well

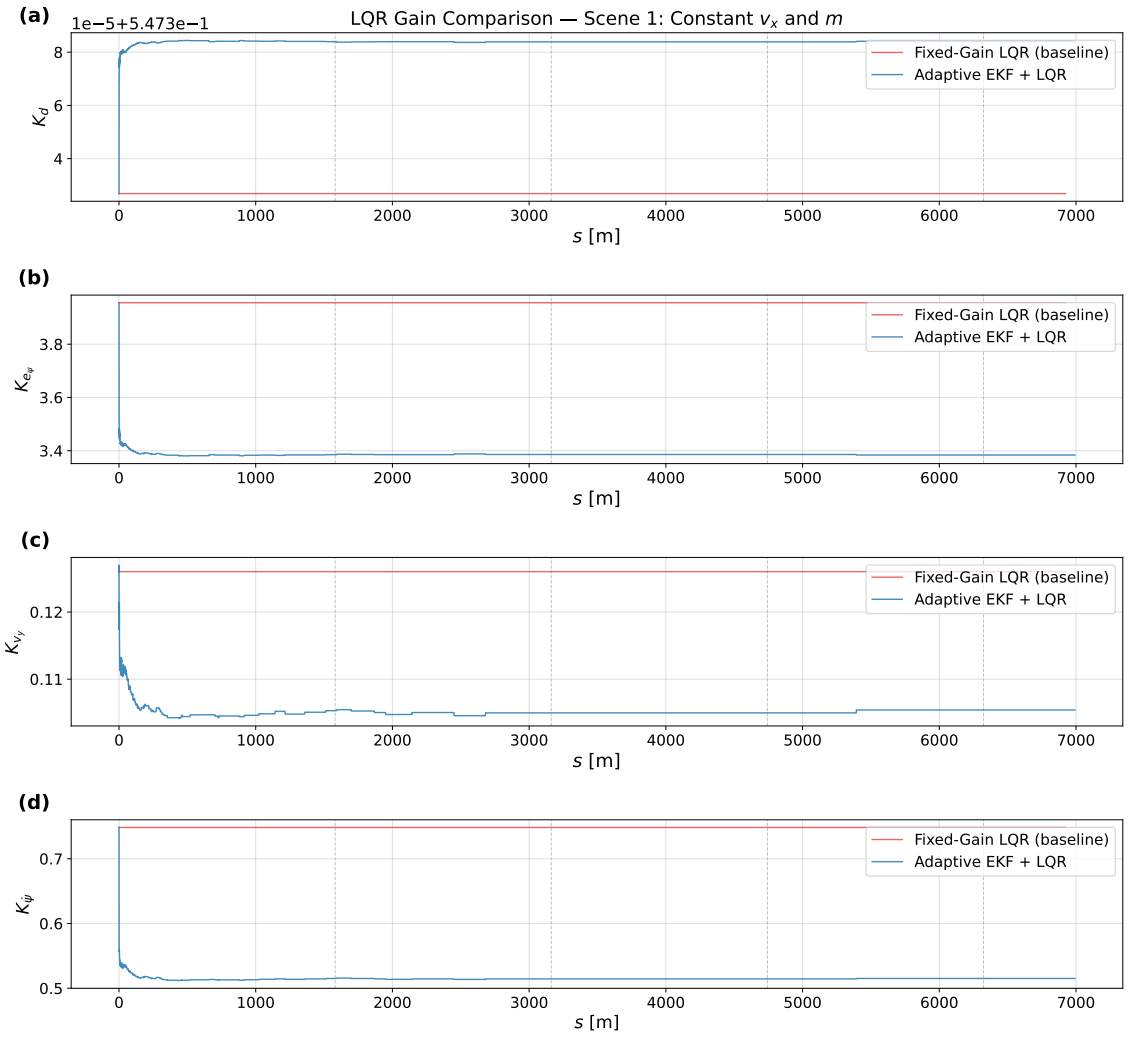

Figure 5.7: Plots showing evolution of LQR gains for an indirect adaptive lateral LQR controller using EKF. The solid red horizontal line shows the gain initially computed by the LQR. For a non adaptive controller this gain would be constant, but for an adaptive one, we can observe how quickly it evolves and reaches a steady state value.

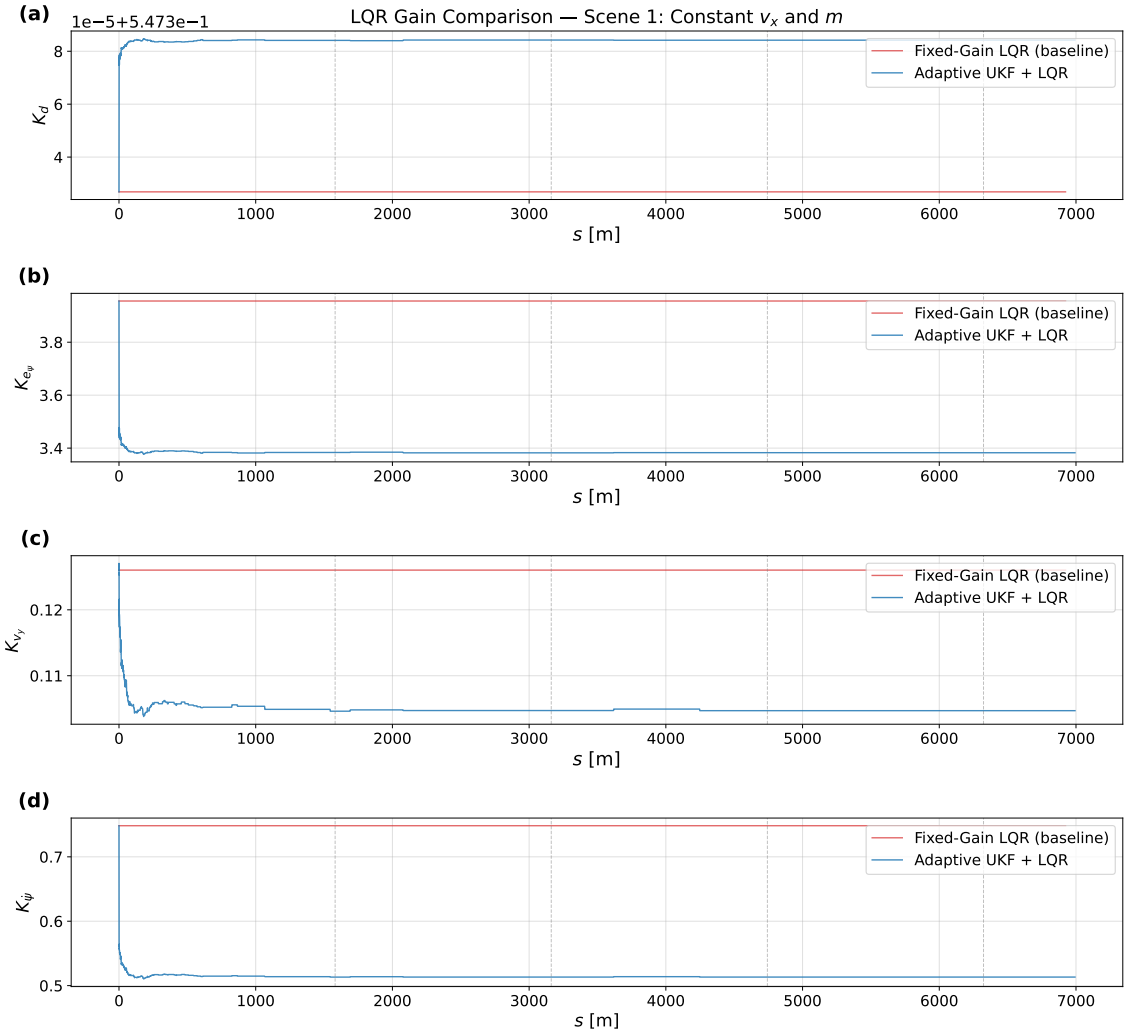

Figure 5.8: Plots showing evolution of LQR gains for an indirect adaptive lateral LQR controller using UKF. The features of the plot are similar to that of Figure 5.7 and showcases the benefits of using an adaptive mechanism.

In figures 5.9 and 5.10, we can see that the path taken by the vehicle completely overlaps the reference path. While these figures are a clear give away of the performance of the indirect adaptive controllers, additional figures shed a brighter light on the detailed performance of the controllers.

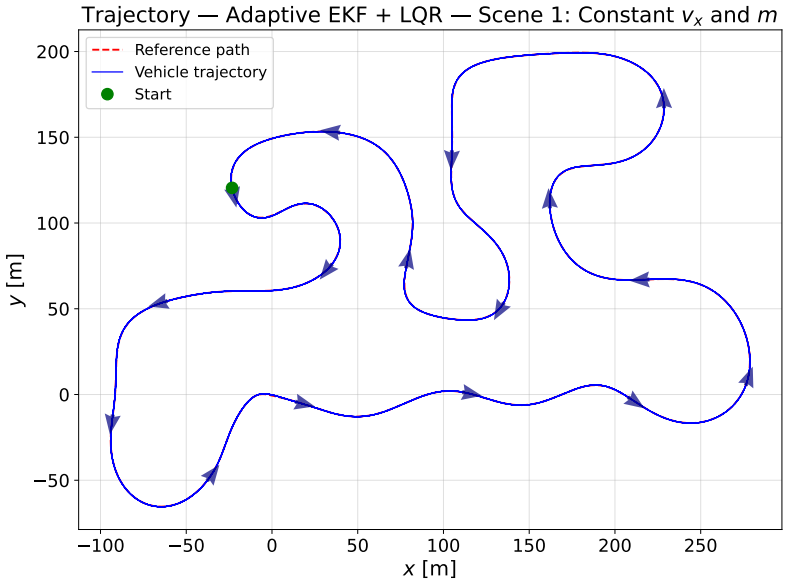

Figure 5.9: Plot shows the trajectory taken by a vehicle using an indirect adaptive EKF based lateral LQR controller.

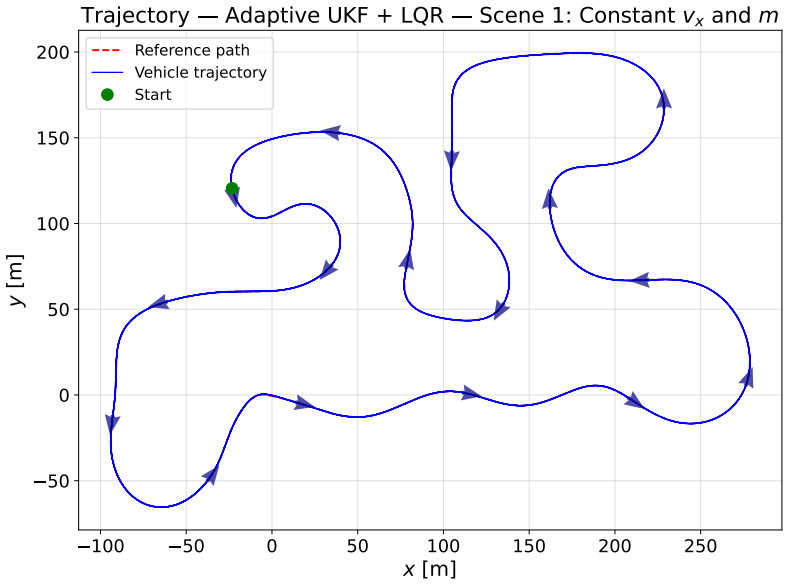

Figure 5.10: Plot shows the trajectory taken by a vehicle using an indirect adaptive UKF based lateral LQR controller. The trajectory taken is extremely similar to figure 5.9, giving an early indication that both approaches perform equally well.

The Figures 5.11 - 5.14, shows the lateral accuracy of the indirect adaptive lateral LQR controller. The figures also show the performance of a vehicle using a normal LQR controller as well. The comparison between the two controllers shows that the presence of adaptation in a lateral controller indeed increases the performance of the plant.

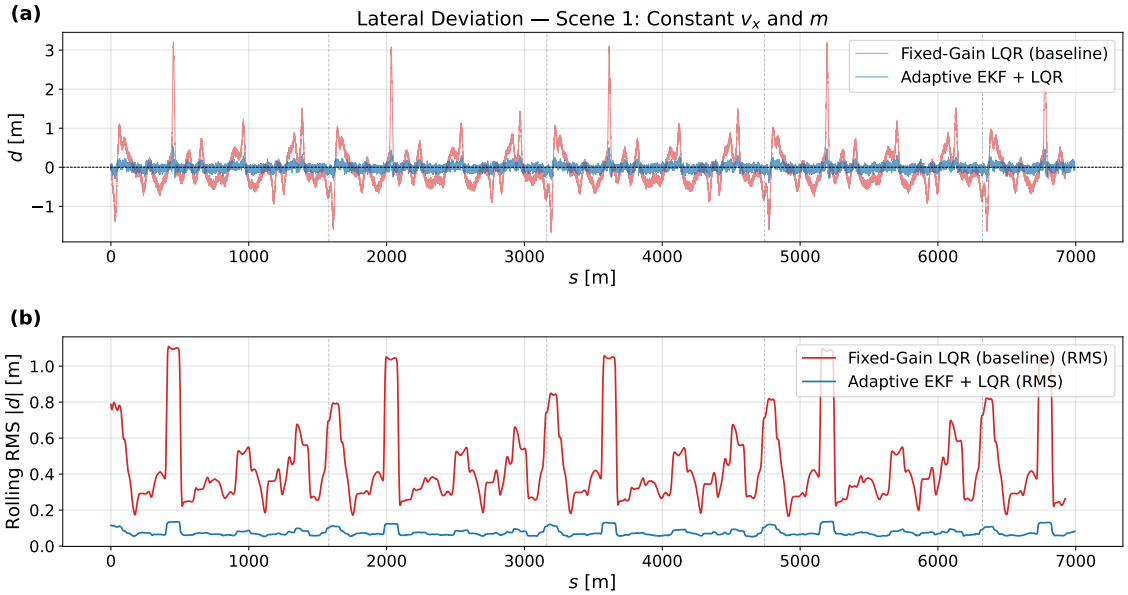

Figure 5.11: Plots showing lateral offset d of the vehicle with (blue line) and without (red line) an adaptive EKF based LQR controller. Subplot (a) shows the raw d value for the simulation, whereas subplot (b) shows the rolling RMS of $\mid d \mid$ within a window of 500 time steps.

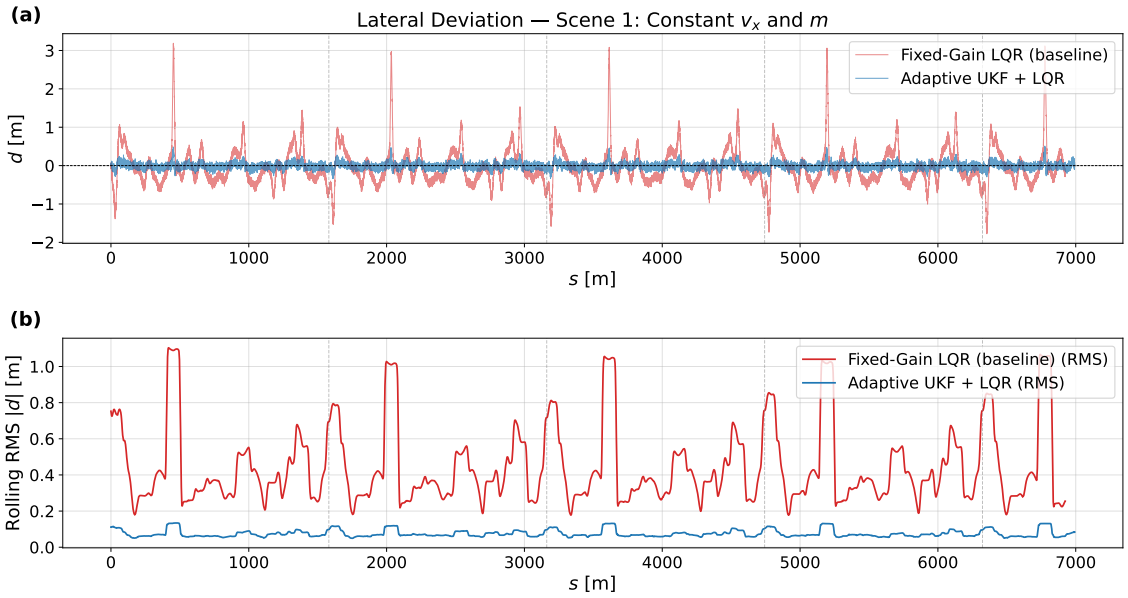

Figure 5.12: Plots showing lateral offset d of the vehicle with (blue line) and without (red line) an adaptive UKF based LQR controller. Subplot (a) shows the raw d value for the simulation, whereas subplot (b) shows the rolling RMS of $\mid d \mid$ within a window of 500 time steps.

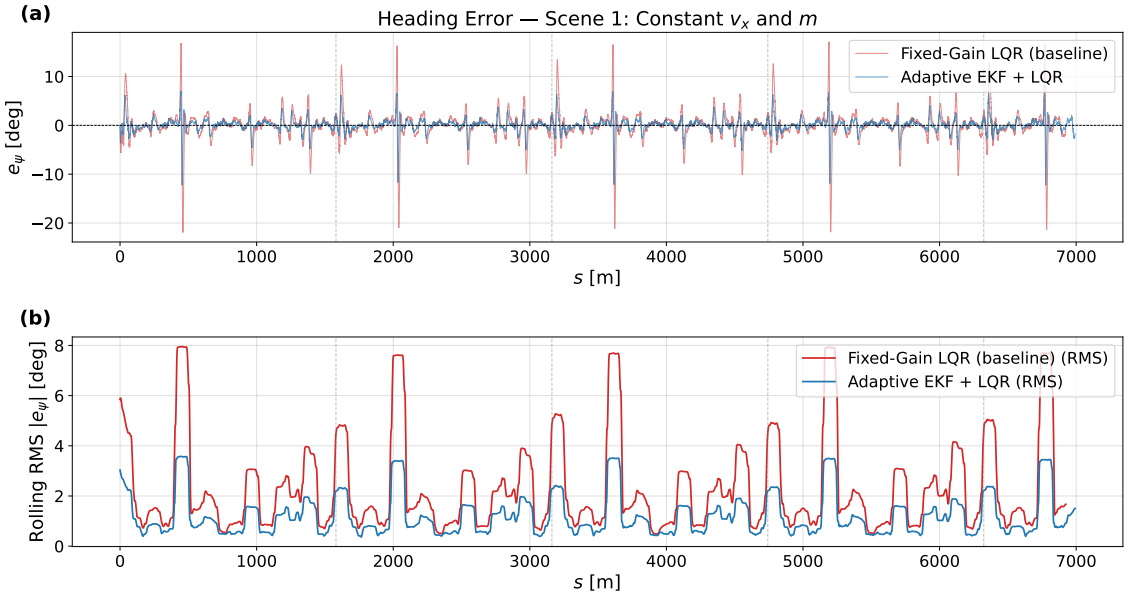

Figure 5.13: Plots showing heading error e_ψ of the vehicle with (blue line) and without (red line) an adaptive EKF based LQR controller. Subplot (a) shows the raw e_ψ value for the simulation, whereas subplot (b) shows the rolling RMS of $\mid e_\psi \mid$ within a window of 500 time steps.

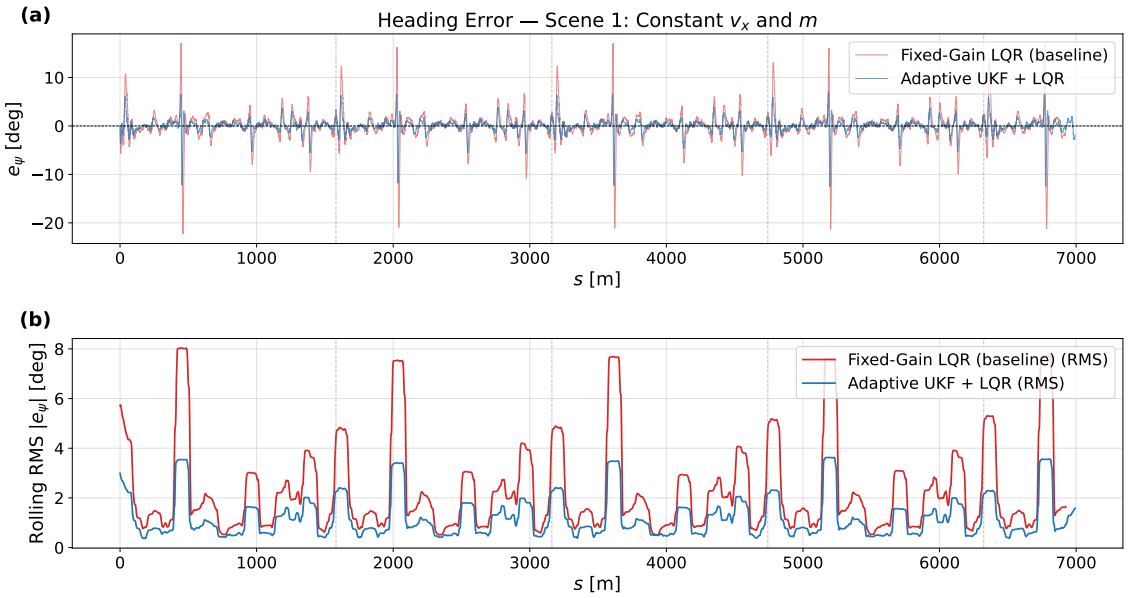

Figure 5.14: Plots showing heading error e_ψ of the vehicle with (blue line) and without (red line) an adaptive UKF based LQR controller. Subplot (a) shows the raw e_ψ value for the simulation, whereas subplot (b) shows the rolling RMS of $\mid e_\psi \mid$ within a window of 500 time steps.

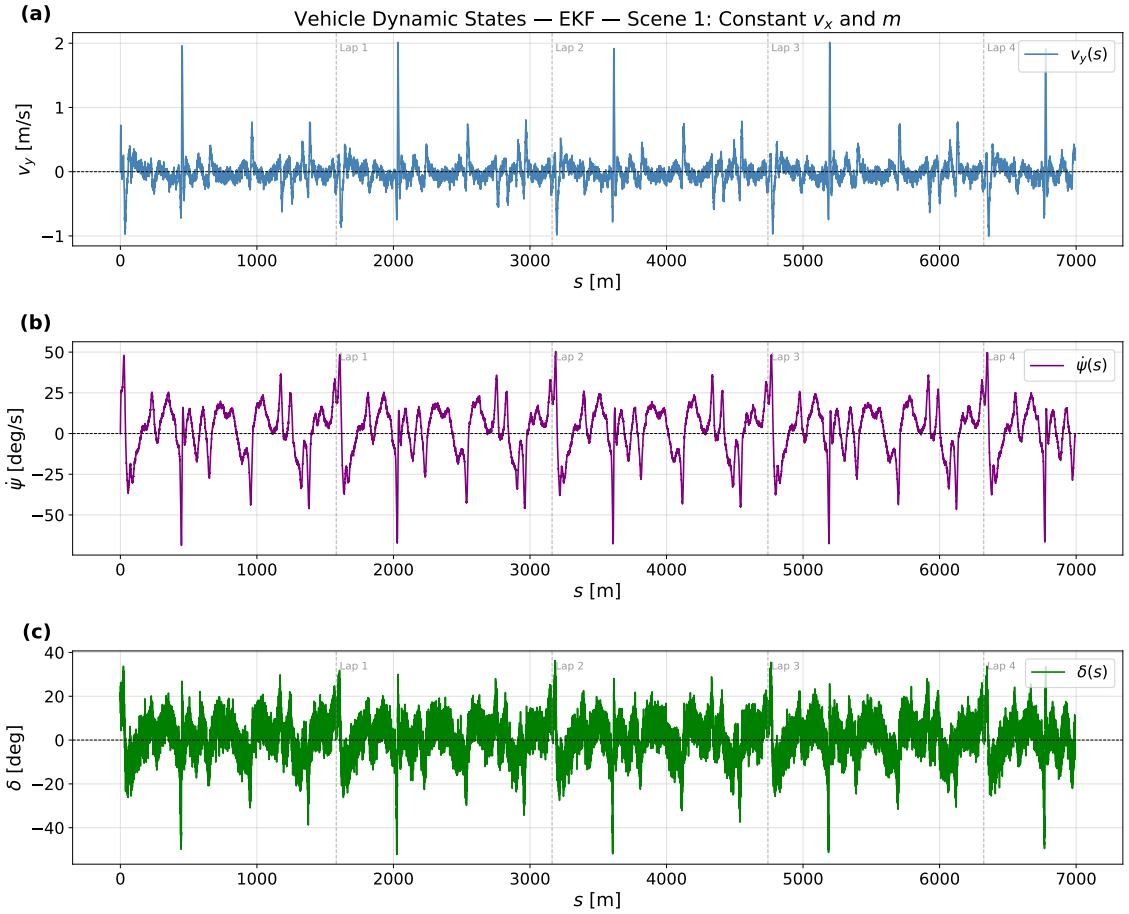

Figure 5.15: Plots showing other vehicle states when the adaptive lateral EKF based LQR controller is used. Subplot (a) shows the lateral velocity v_y of the vehicle at various points on the track, while subplot (b) shows the yaw rate $\dot{\psi}$. Subplot (c) shows the optimal control input δ provided by the adaptive lateral controller.

One important thing of note here is that in the current simulation method, the signal measurements that the LQR controller receives, includes a gaussian noise as well in order to make the scenario more realistic. It is because of this noise that the all the figures 5.11 - 5.16 appear to be noisy. The appendix section A.1.5 includes the same simulation but without the inclusion of sensor noise, and there we can observe that the plots are indeed smooth in nature.

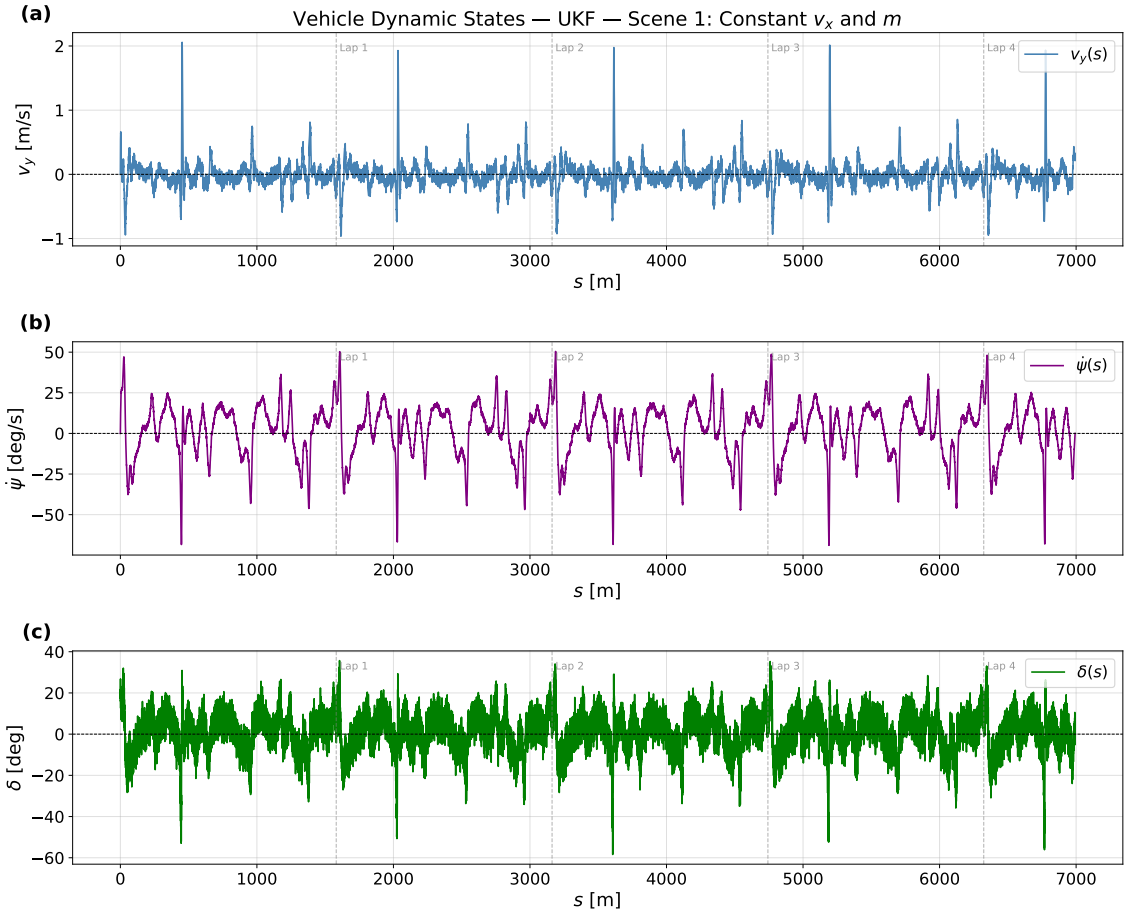

Figure 5.16: Plots showing other vehicle states when the adaptive lateral UKF based LQR controller is used. The subplots here convey the same message as those in figure 5.15.

5.2.2 Scenario 2: Vehicle is loaded and unloaded

While the test scenario is different from before, the evaluation criteria for the controllers remain the same, with the added criteria of the adaptation speed of the estimators. As specified before, in this scenario, the longitudinal velocity of the vehicle would decrease to a near zero value (0.1 m/s), after which it's mass would be increase by 20 %, before returning to it's normal velocity. After a while, the reverse of this would happen where the mass is decreased back to the original value.

5.2.2.1 Parameter estimation performance

The convergence speed and the accuracy of the final estimation results for this scenario is almost the same as those presented in tables 5.2 and 5.3. The main aim of this scenario is to observe how quick the Kalman filters can adapt (using the parameter drift rate adaptive Q framework) when there is change in the parameter to be estimated.

From the figures 5.17 - 5.20, we can observe that both the adaptive EKF and UKF are able to detect the change in the parameters and adapt themselves to it. On a closer look, we can further see that the UKF adapts to the new values much more efficiently compared to the EKF. This is because at velocities closer to zero, the Jacobian linearizations of the EKF starts to break, as a result of which the estimations drift away from the true values. Once the velocity returns to normal, the adaptive layer of the EKF aggressively corrects the estimates.

This issue is not encountered in the UKF case because the unscented Kalman filter does not perform Jacobian linearizations, and instead develops and propagates using $2n + 1$ sigma points through the non linear dynamics to estimate the parameters. This nature of the UKF also enables it to function at even lower velocities than 0.1 m/s, where the adaptive performance is further increased.

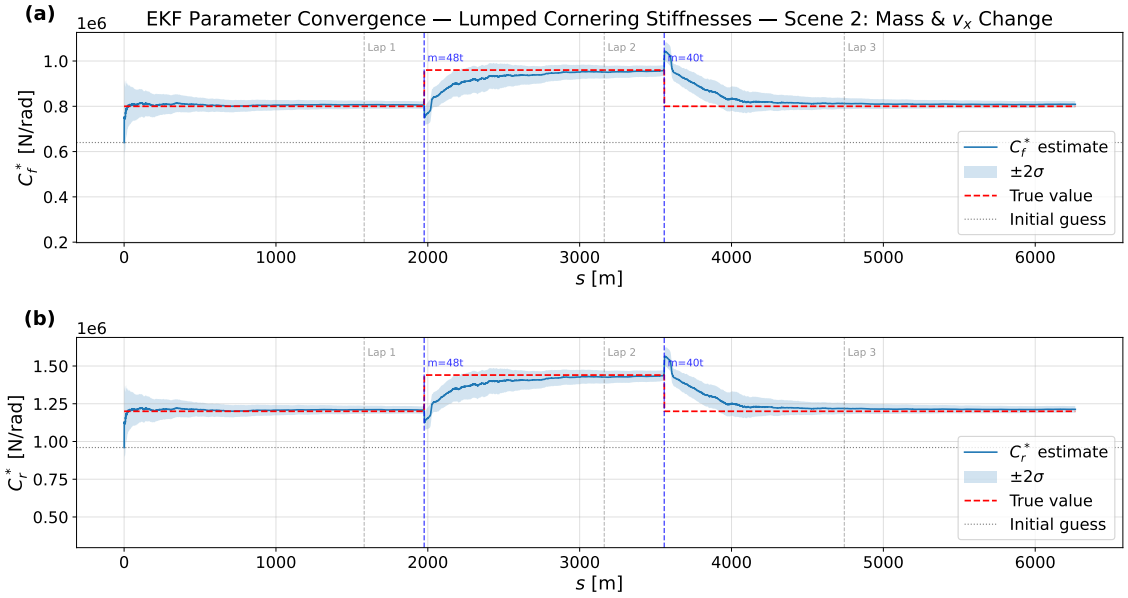

Figure 5.17: Plots showing estimation of C_f^* (a) and C_r^* (b) using adaptive EKF. We can observe there is a slight hiccup in the estimation when the reduction in velocity occurs, showing the effect of linearization errors at low speeds

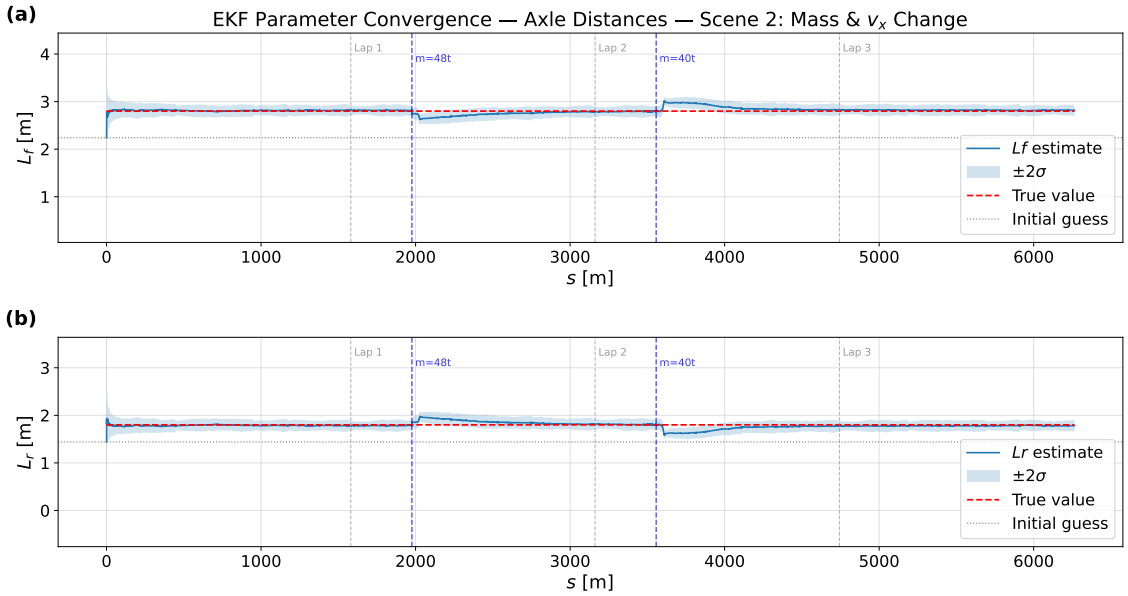

Figure 5.18: Plots showing estimation of L_f (a) and L_r (b) using adaptive EKF. We can observe that similar to figure 5.17 there is a slight hiccup in the estimation due to linearization errors at low speeds

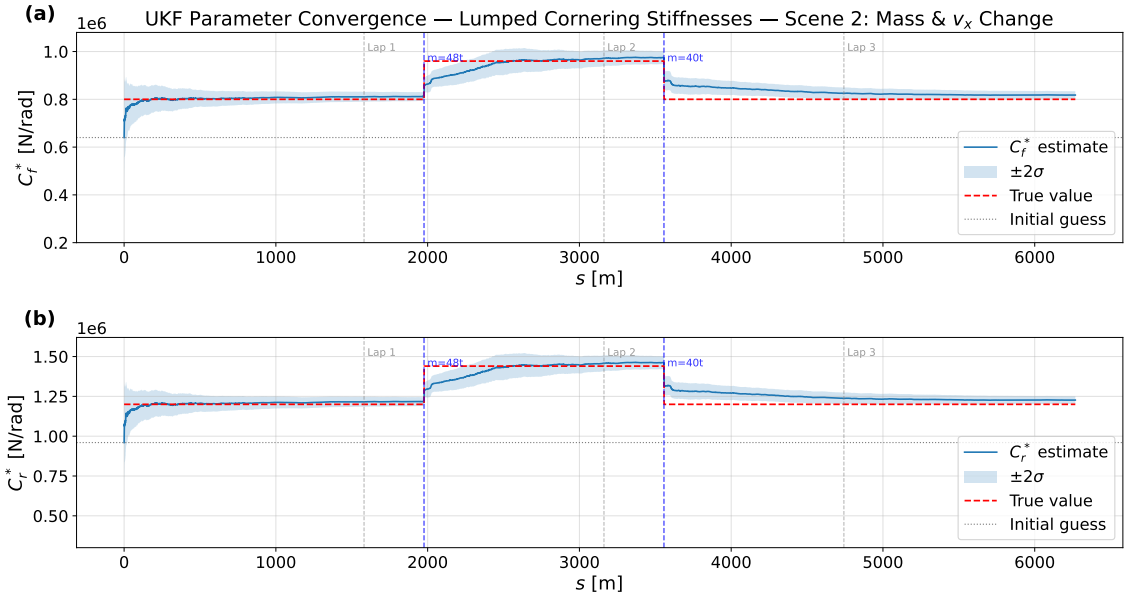

Figure 5.19: Plots showing estimation of C_f^* (a) and C_r^* (b) using adaptive UKF. We can observe that, unlike in the EKF case, UKF is unaffected by the low speed of the vehicle as it does not conduct any linearizations for estimation.

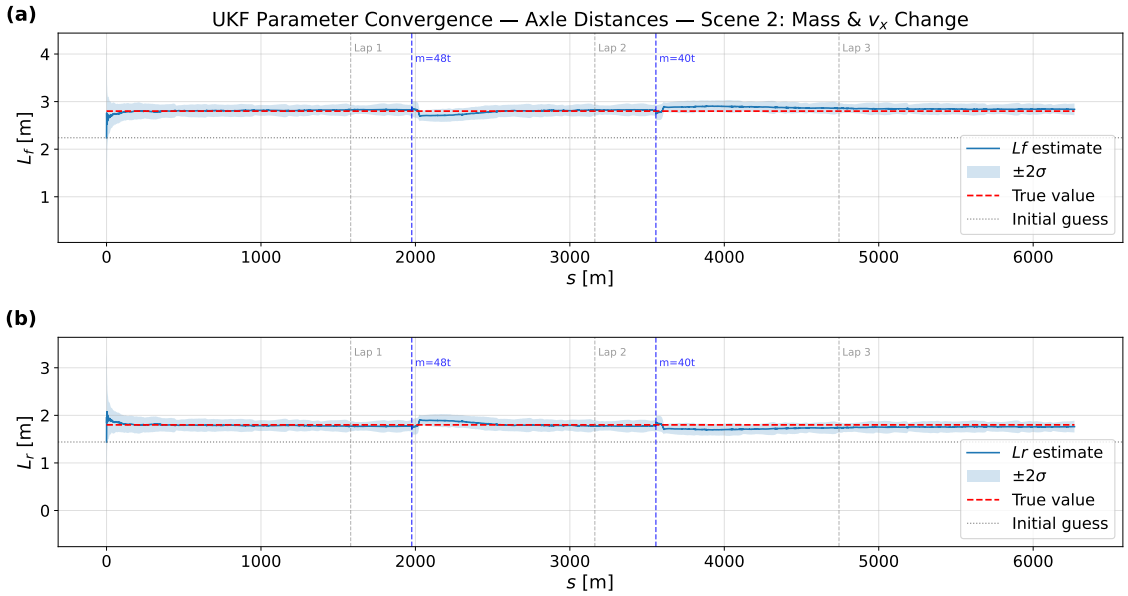

Figure 5.20: Plots showing estimation of C_f^* (a) and C_r^* (b) using adaptive UKF. Similar to figure 5.19, the filter is able to smoothly estimate the parameters even at low speeds.

We can observe the evolution of the process error covariance noise matrix Q in the figures 5.21 and 5.22. We can clearly observe how the adaptive layer in the Kalman filters inflates and deflates Q based on the estimated value of the parameters, especially when the values of the parameters change.

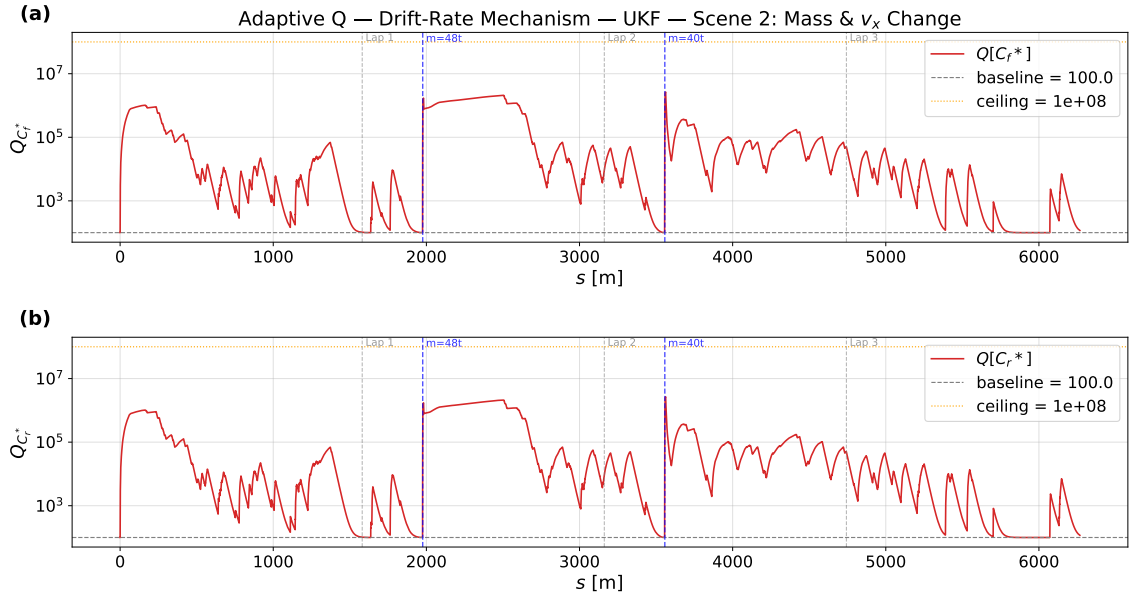

Figure 5.21: Plot showing evolution of process error covariance matrix Q using the parameter drift rate adaptive Q framework in EKF. It can be observed that Q inflates and deflates consistently with changes to parameter values.

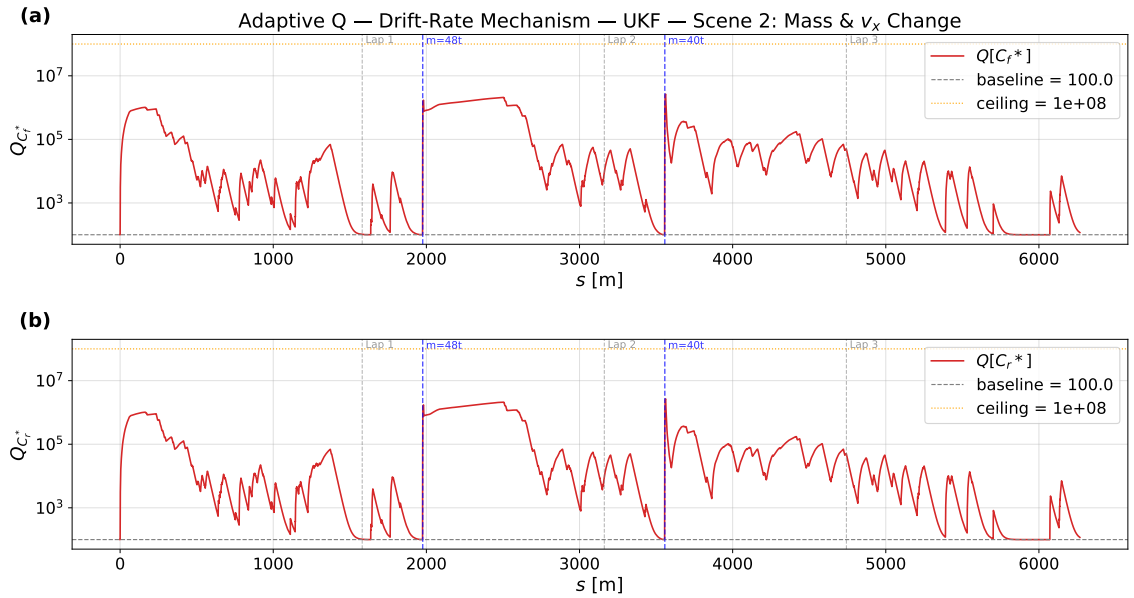

Figure 5.22: Plot showing evolution of process error covariance matrix Q using the parameter drift rate adaptive Q framework in UKF. Since the estimation performance of both filters were similar, the changes in Q are also similar.

5.2.2.2 Lateral LQR performance

Owing to the adaptation speed of the adaptive Kalman filters and their estimation accuracy, we can observe in figures 5.23 and 5.24 that even though the gains tank when the parameter change is encountered, they promptly return to their optimal

values. In the table 5.5 it can be observed that the optimal gains for the controller is the same as the previous scenario.

	$K[d]$	$K[e_\psi]$	$K[v_y]$	$K[\dot{\psi}]$
EKF	0.5474	3.3788	0.1041	0.5113
UKF	0.5474	3.3800	0.1043	0.5121

Table 5.5: Final LQR Gain Matrix K: EKF vs UKF. Even with the changes in true value of the parameters, the estimators still converged to nearly identical values, and consequently, the final gains are once again almost the same.

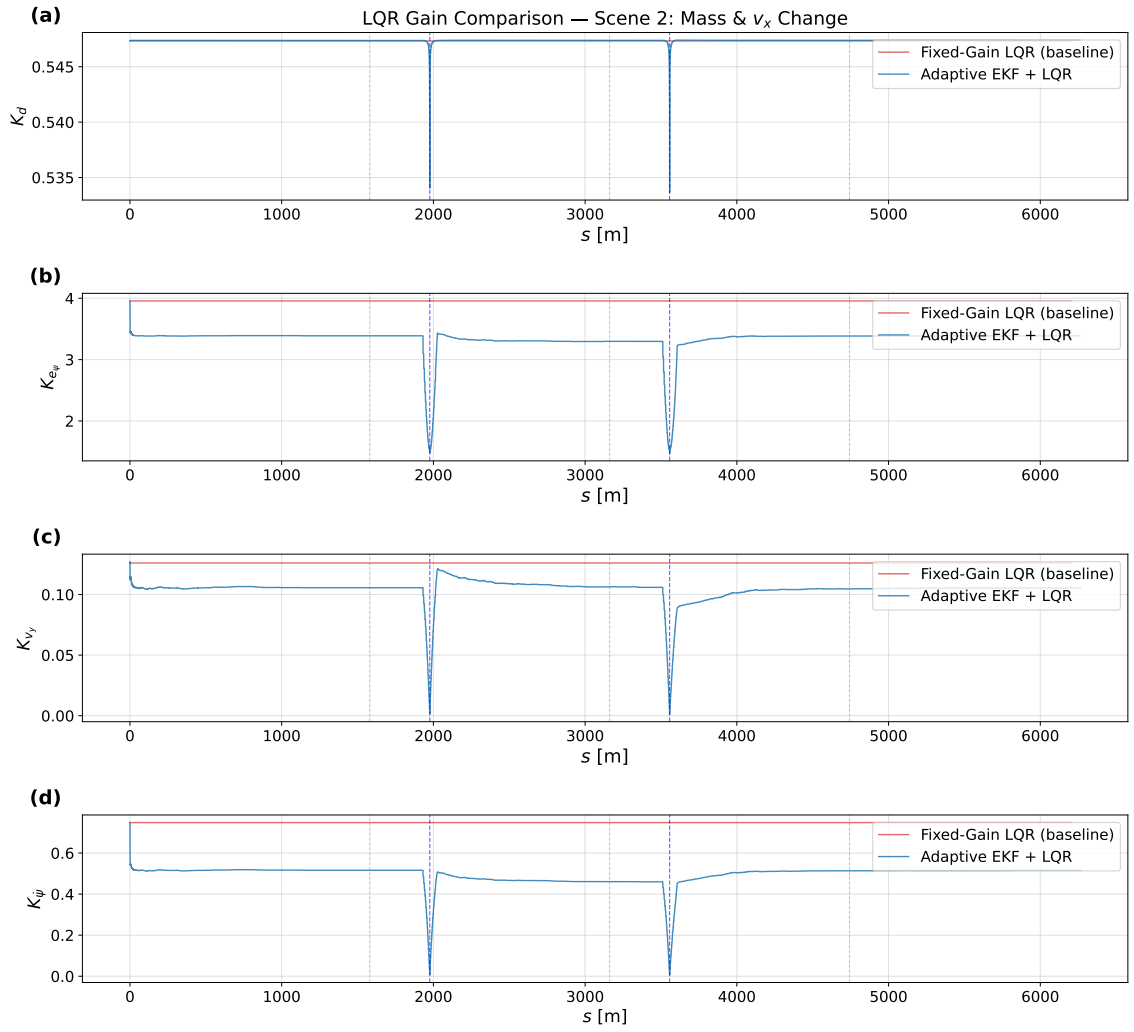

Figure 5.23: Plot shows the evolution of LQR gains for a lateral LQR controller with an adaptive EKF based estimator. We can also observe how the change in parameters affects the gains as well as how quickly the controller adapts to the changes and drives the gains to a steady state value.

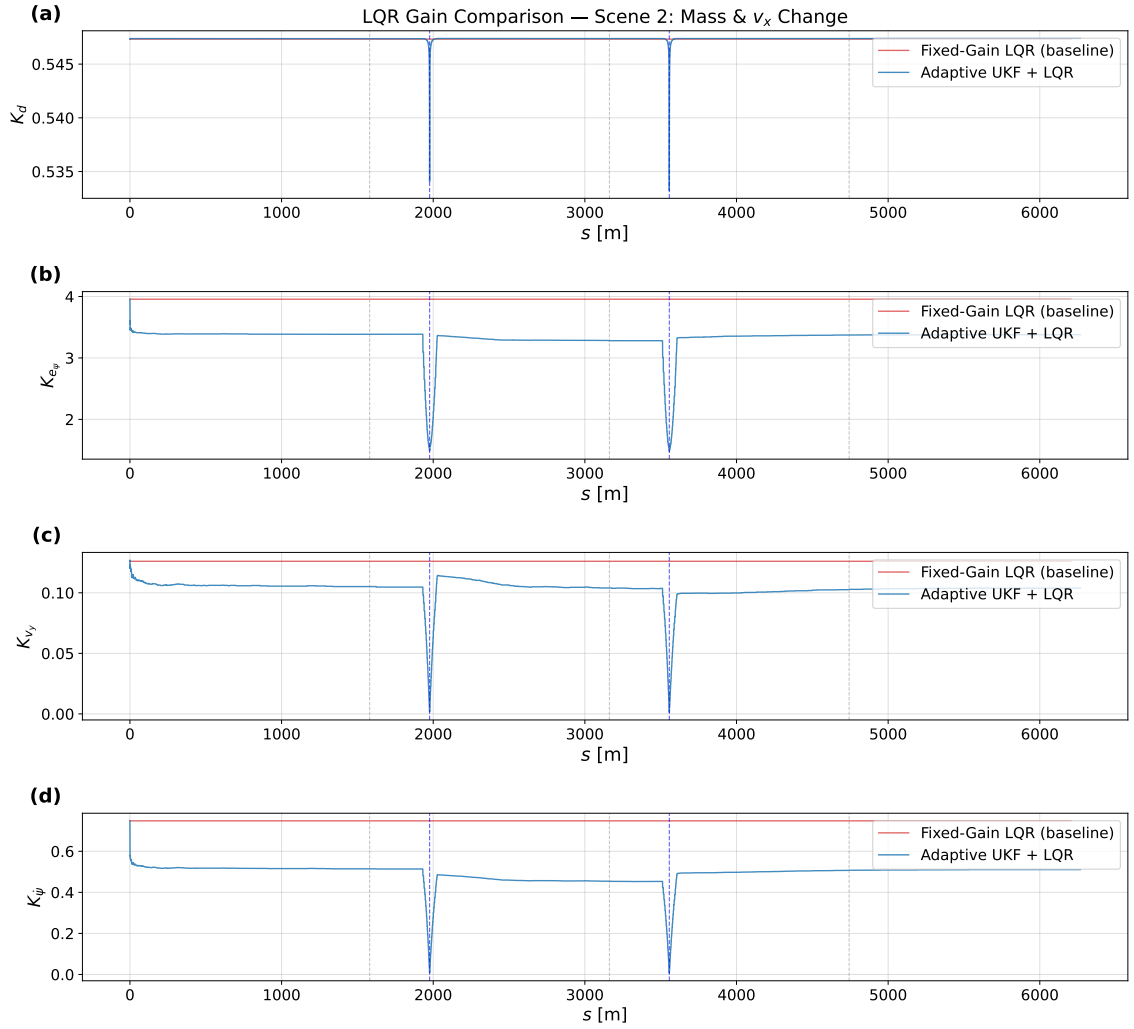

Figure 5.24: Plot shows the evolution of LQR gains for a lateral LQR controller with an adaptive UKF based estimator. Similar to figure 5.23, the controller deftly adapts the gains to the most optimal value.

The rest of the plots, showcasing the performance of the LQR controller, such as the path tracking ability, lateral and heading errors, etc., are similar to the figures 5.11 - 5.16, and are thus not shown here. These can be found in the appendix section A.1.6 for the readers reference.

5.2.3 Scenario 3: Vehicle drives through a muddy surface

This scenario builds upon the scenario in 5.2.2, but instead of varying the vehicle mass, the lateral tire stiffness coefficients C_f and C_r are subject to change. Specifically, one third of the track is designated as a muddy surface, over which the true values of C_f and C_r are reduced to 80% of their true values. For the same portion, the longitudinal velocity v_x is also reduced to 0.8 times its original value to reflect the more cautious driving behavior expected on such a surface.

The key distinction from Scenario 2 lies in the duration of the parameter change. Previously, the vehicle operated in a loaded state for a substantial fraction of the total runtime, giving the estimator prolonged exposure to the shifted dynamics. Here, the muddy segment is comparatively shorter, which places greater emphasis on the transient response of the parameter drift rate adaptive Q Kalman filters. Thus, this scenario primarily evaluates how rapidly each filter can track the abrupt changes in C_f^* and C_r^* upon entering and exiting the muddy region, rather than its steady state estimation accuracy.

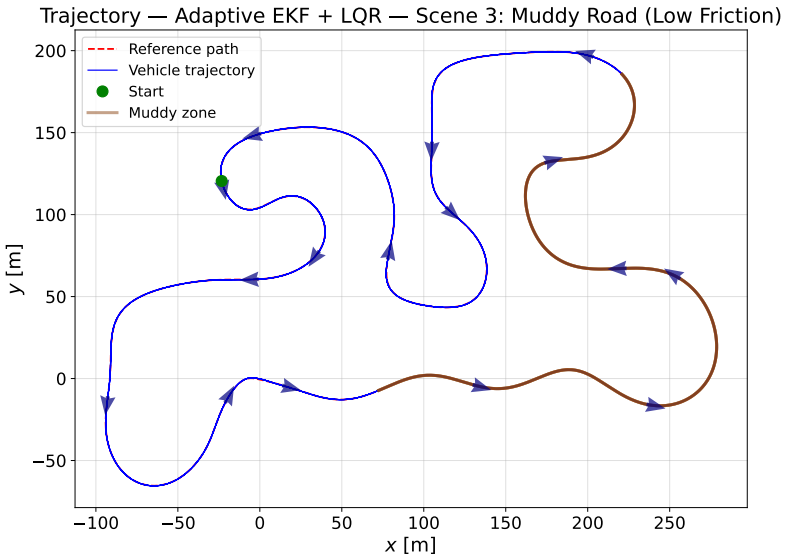

Figure 5.25: Plot shows the trajectory taken by the vehicle. The part of the track shaded brown is the muddy zone and it covers a third of the total track.

5.2.3.1 Parameter estimation performance

The figures 5.26 - 5.29 shows the performance of the adaptive Kalman filter based estimators. Unlike figures 5.17 and 5.19, the estimator is unable to converge swiftly enough to the updated values of the parameters. This is primarily because the dwell time of the vehicle in the new parameter regime in scenario 3 is much shorter than that in scenario 2, thereby not providing the estimator with sufficient time to converge.

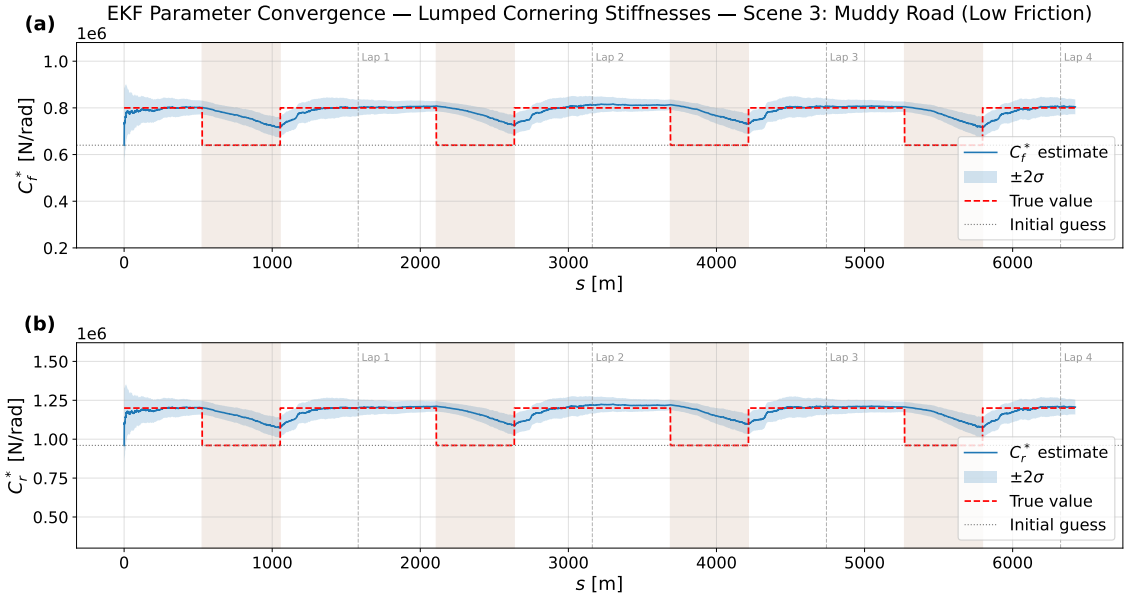

Figure 5.26: Plots showing estimation of C_f^* (a) and C_r^* (b) using adaptive EKF. It can be observed that the duration for which vehicle traverses the muddy path is quite short, as a result of which the estimator does not get enough time to converge to the update values, before the parameters shift again.

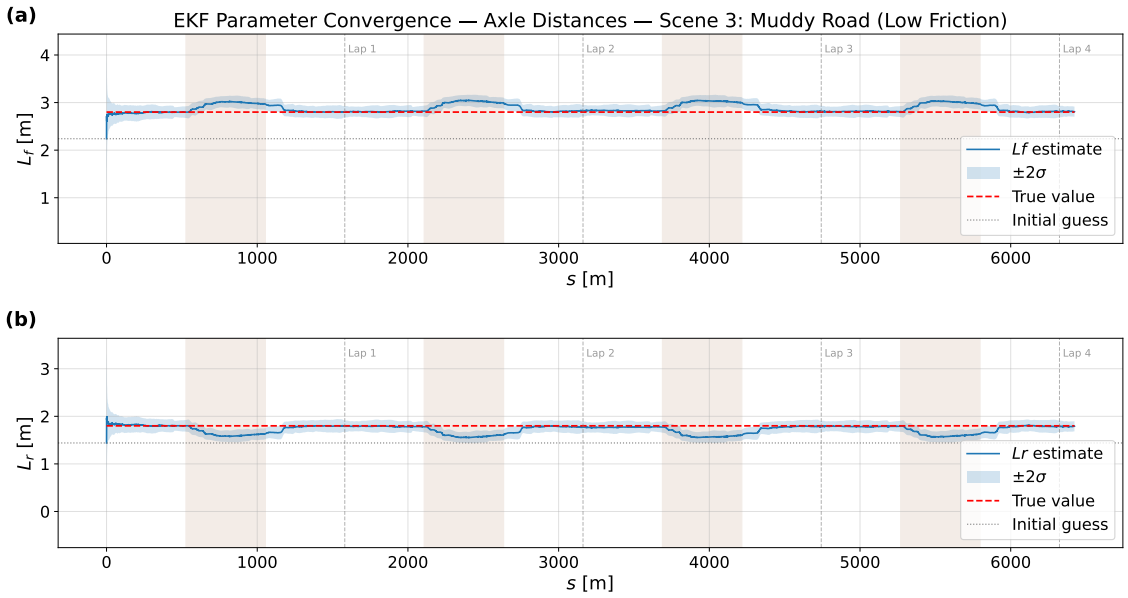

Figure 5.27: Plots showing estimation of L_f (a) and L_r (b) using adaptive EKF. We can observe that the overall performance of the estimator here is similar to figure 5.18.

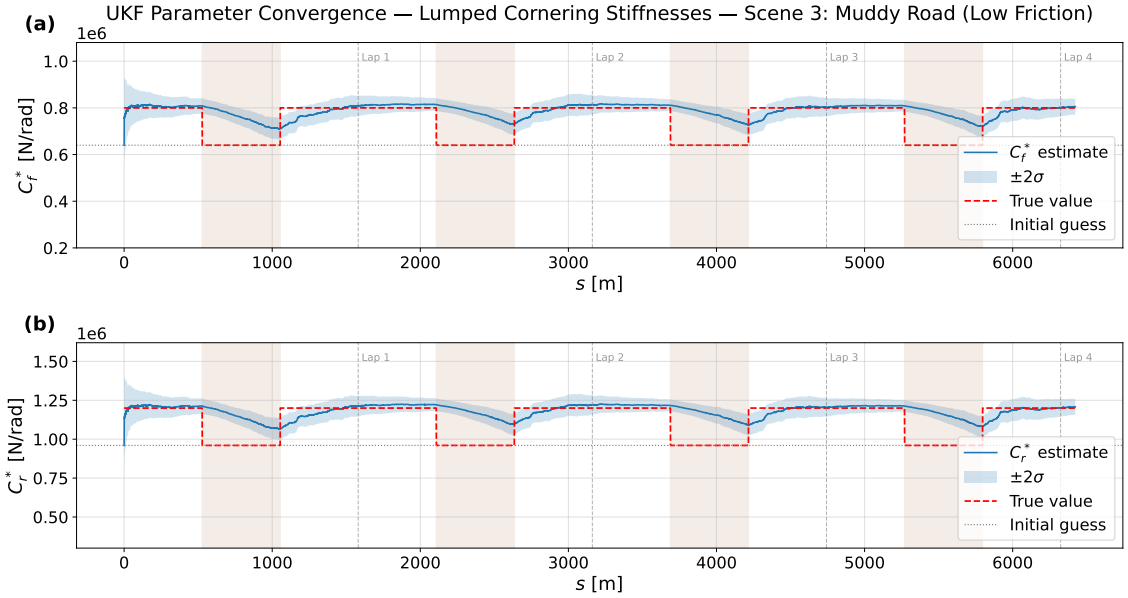

Figure 5.28: Plots showing estimation of C_f^* (a) and C_r^* (b) using adaptive UKF. Similar to 5.26, the estimator attempts to converge to the new value however the vehicle exits the part before that.

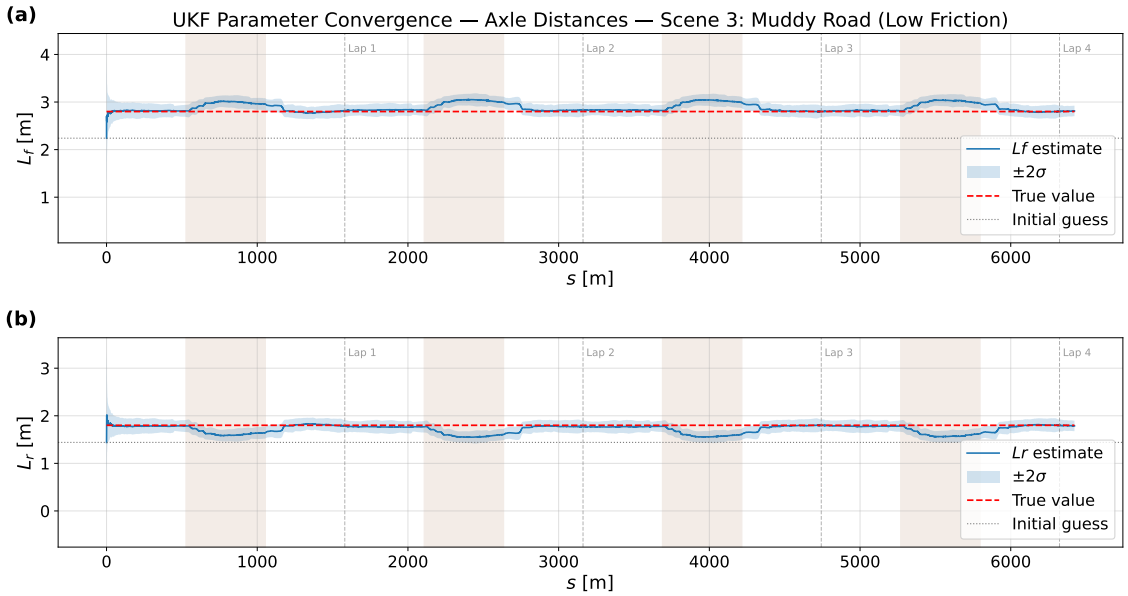

Figure 5.29: Plots showing estimation of C_f^* (a) and C_r^* (b) using adaptive UKF.

The estimation performance of the estimator can be improved by tuning the parameter drift rate adaptive Q to be more aggressive, however this would increase the steady state uncertainty of the parameters. Thus, this scenario reveals us the trade-off between faster tracking of the parameters versus increased estimation noise.

5.2.3.2 Lateral LQR performance

The figures 5.30 and 5.31 shows that the gains are evolving in a cyclic manner, which is consistent with the entry and exit of the vehicle in the muddy zone. While the estimator is unable to provide the controller with a steady state estimate for the transient period, the partial adaptation it does perform pushes the LQR gains towards the correct values.

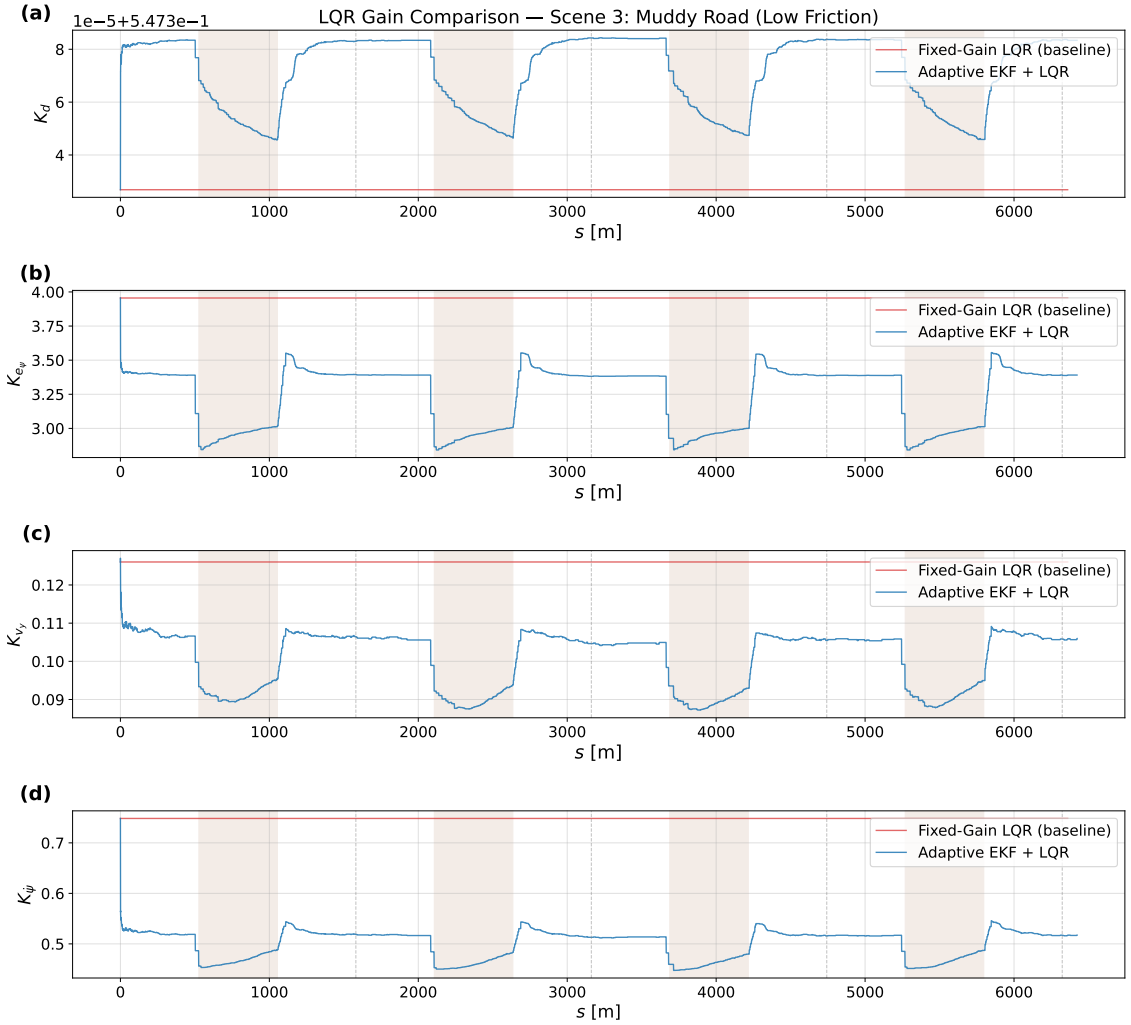

Figure 5.30: Plot shows the evolution of LQR gains for a lateral LQR controller with an adaptive EKF based estimator. We can also observe how quickly the optimal gains change whenever the vehicle enters the muddy region.

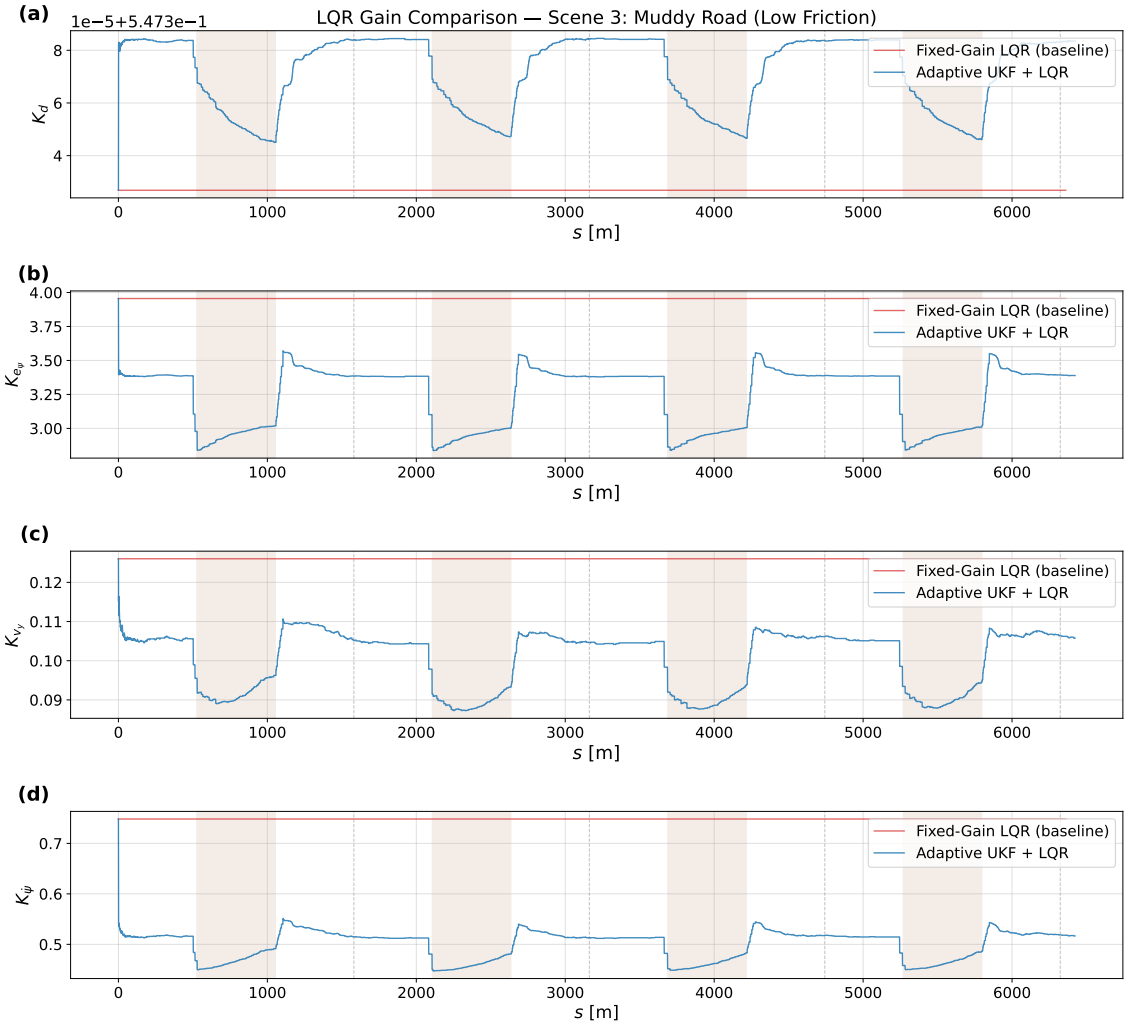

Figure 5.31: Plot shows the evolution of LQR gains for a lateral LQR controller with an adaptive UKF based estimator. Similar to figure 5.30, the controller deftly adapts the gains to the most optimal value.

The remaining plots, which illustrate other aspects of the lateral performance of the controller such as lateral offset, heading error, control input, etc. show similar behavior to those presented in Sections 5.2.1 and 5.2.2. As a result, since they do not provide additional insights, they are therefore omitted from this report.

5.2.4 Comparative lateral performance

To consolidate the per-scenario findings, the lateral performance of the three controllers, a normal LQR controller, an indirect adaptive lateral LQR controller with adaptive EKF and an indirect adaptive lateral LQR controller with adaptive UKF, are directly compared in the figures 5.32 - 5.37.

In the figures 5.32 - 5.34, the subplot (a) shows the maximum lateral deviation and the RMS value of the deviation for the full duration of the simulation and at the end during steady state. Subplot (b) on the other hand shows the various heading errors in a similar fashion.

Meanwhile, the figures 5.35 - 5.37 show the cumulative absolute error (CAE) for the driven distance for all three controllers. A steeper slope indicates a poorer tracking performance, while a flatter slope reflects tighter path following.

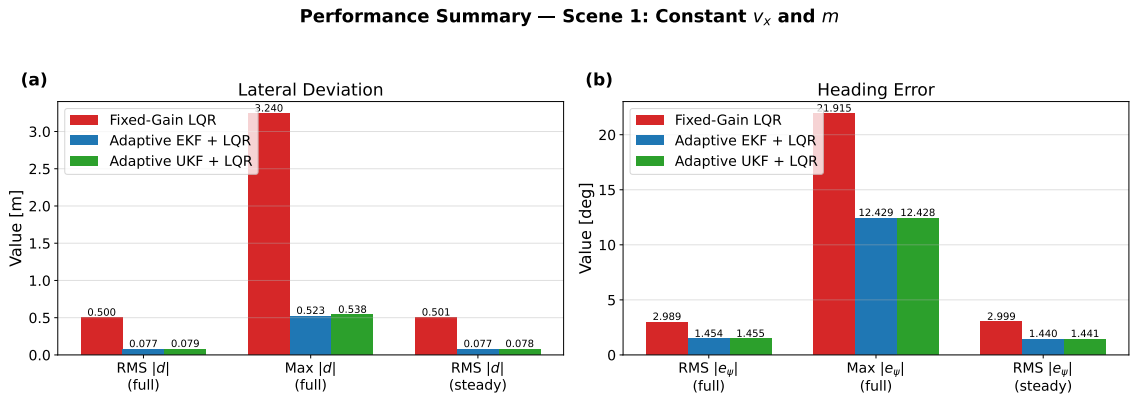

Figure 5.32: Plots for scenario 1, showing the overall lateral performance of the controllers. The red bar represents a normal lateral LQR controller, whereas the blue and green bar represent the lateral controllers using adaptive EKF and UKF respectively.

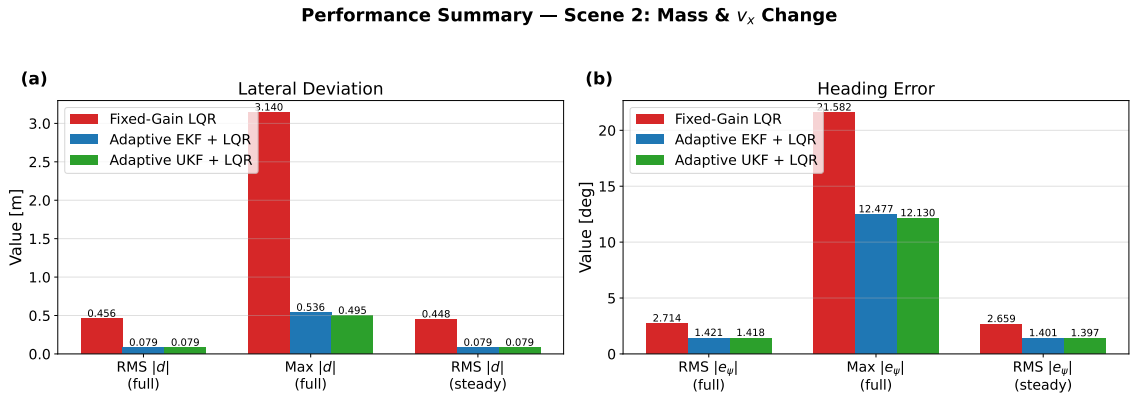

Figure 5.33: Plots showing the overall lateral performance of the controllers for scenario 2, where the loading and unloading of the vehicle is simulated.

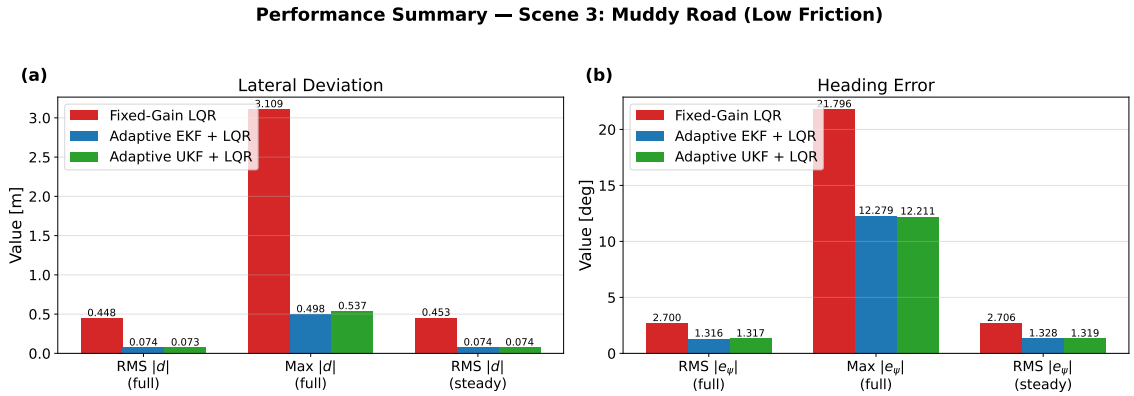

Figure 5.34: Plots showing the overall lateral performance of the controllers for scenario 3, where a portion of the track was considered as muddy (low friction).

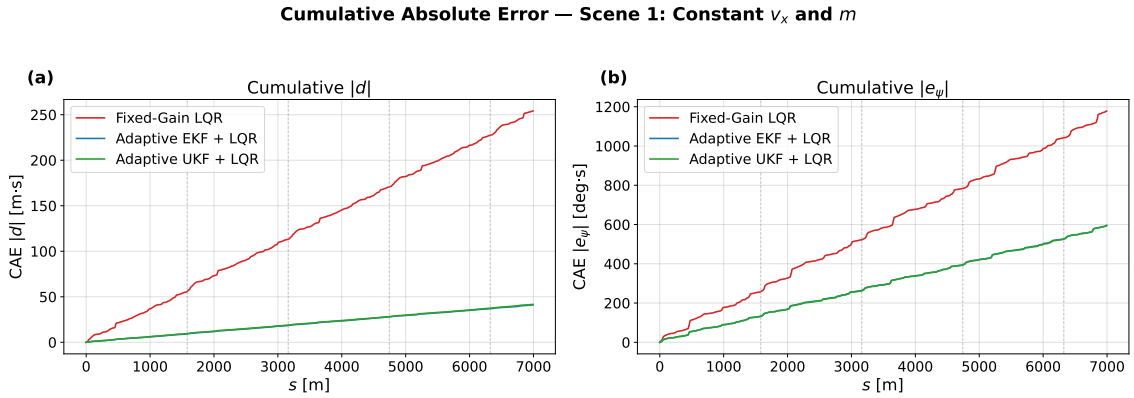

Figure 5.35: Plots showing the CAE of d (a) and e_ψ (b) for scenario 1. The plots for EKF and UKF overlap each other considerable. The plots of the adaptive controllers are significantly flatter than the non adaptive case.

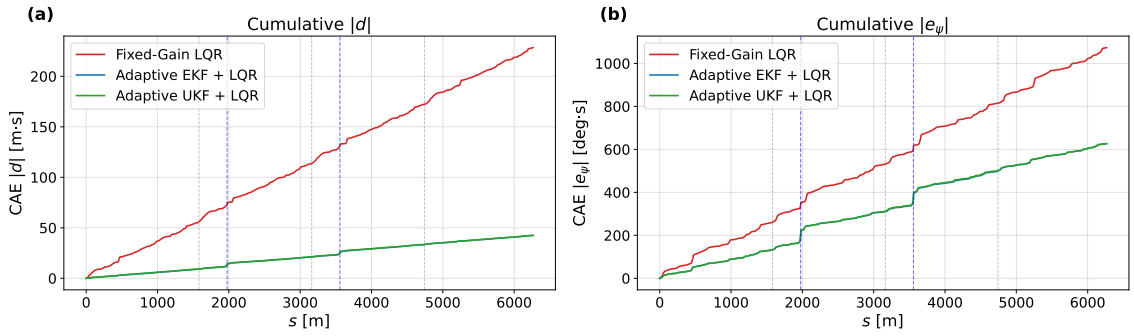

Figure 5.36: Plots showing the CAE of d (a) and e_ψ (b) for scenario 2. The vertical blue lines indicate the point at which the vehicle was loaded and unloaded. We can observe that there is a noticeable increase in CAE values at these points.

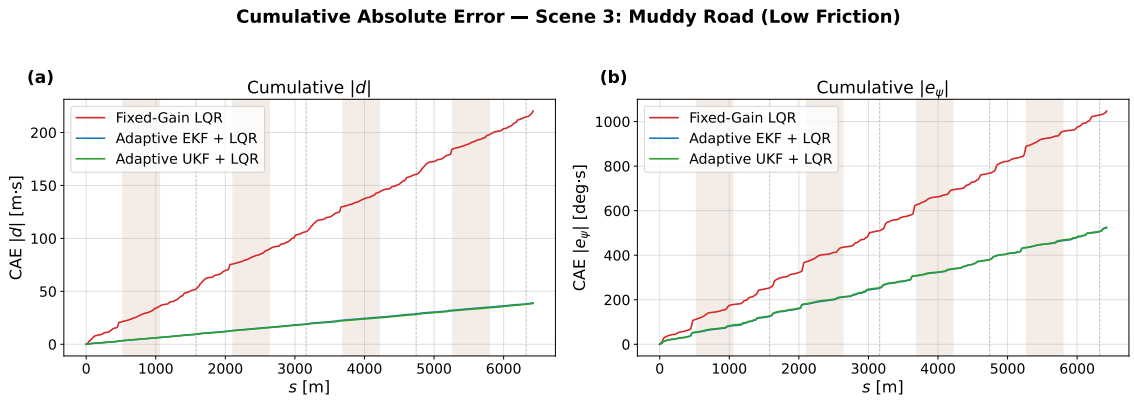

Figure 5.37: Plots showing the CAE of d (a) and e_ψ (b) for scenario 3. The shaded segments indicate the portions of the total path when the vehicle was in the muddy zone.

Across all three scenarios, the adaptive controllers reduce the steady state RMS lateral deviation by $83\% - 85\%$ and the steady state RMS heading error by $49\% - 53\%$, relative to the fixed gain LQR baseline. Furthermore, this improvement is consistent regardless of operating conditions of the three scenarios.

It can also be observed that the two adaptive controllers (based on EKF and UKF estimators) perform nearly identically in all three scenarios, their metrics differing at most by $3\% - 4\%$.

5.3 Direct Adaptive Lateral PID Control

As explained in the section 5.1, the testing scenario for the adaptive PID controllers is similar to the testing scenario 1 of the indirect adaptive controllers. The nominal (incorrect) values of the parameters has been specified in the table 5.6.

Parameter	Nominal	True	Err (%)
L_f [m]	2.8	2.8	0.0
L_r [m]	1.8	1.8	0.0
m [kg]	35 000	40 000	12.5
I_z [kg·m^2]	450 000	500 000	10.0
C_f^* [N/rad]	640 000	800 000	20.0
C_r^* [N/rad]	960 000	1 200 000	20.0

Table 5.6: Nominal vs true parameters of the vehicle. Since this is a direct approach, the nominal values won't be updated during runtime.

The A_m matrix of reference model has been constructed using pole-placement, and the adaptation constants vector Γ have been chosen using trial and error to be,

$$\text{Reference model poles} = \begin{bmatrix} -1.2 & -1.2 & -0.5 & -0.5 \end{bmatrix}^T \tag{5.1}$$

$$\Gamma = \begin{bmatrix} 1.5 & 0.8 & 1.5 & 0.2 & 0.2 & 0.2 \end{bmatrix}^T \tag{5.2}$$

Based on these specifications, the simulation was executed for a total of 30,000 steps with a time step of 0.02 seconds. Within this duration, the vehicle was once again able to successfully complete three laps around the track.

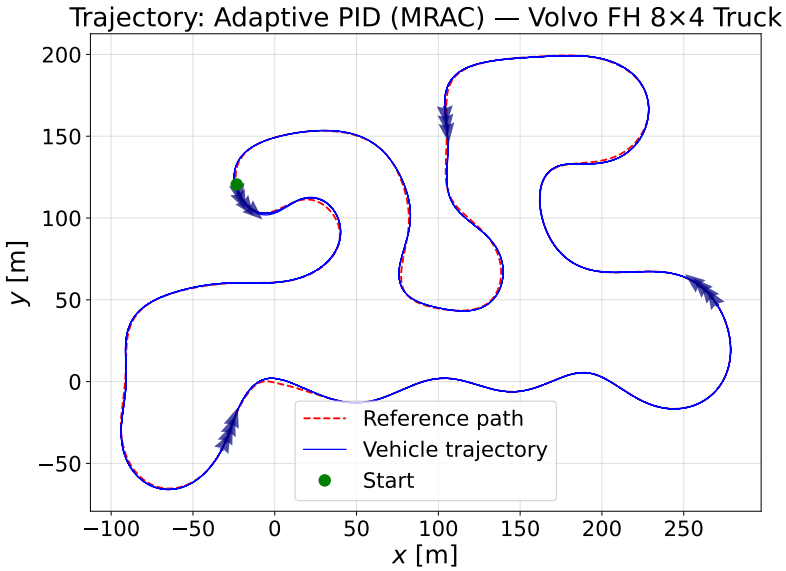

Figure 5.38: Plot shows the trajectory taken by a vehicle using an direct adaptive lateral PID controller.

From the plots in the figure 5.39, it is evident that the performance of the direct adaptive controller is significantly poorer than the indirect adaptive controllers. Especially during the turns, where the controller is unable to adapt the gains fast enough.

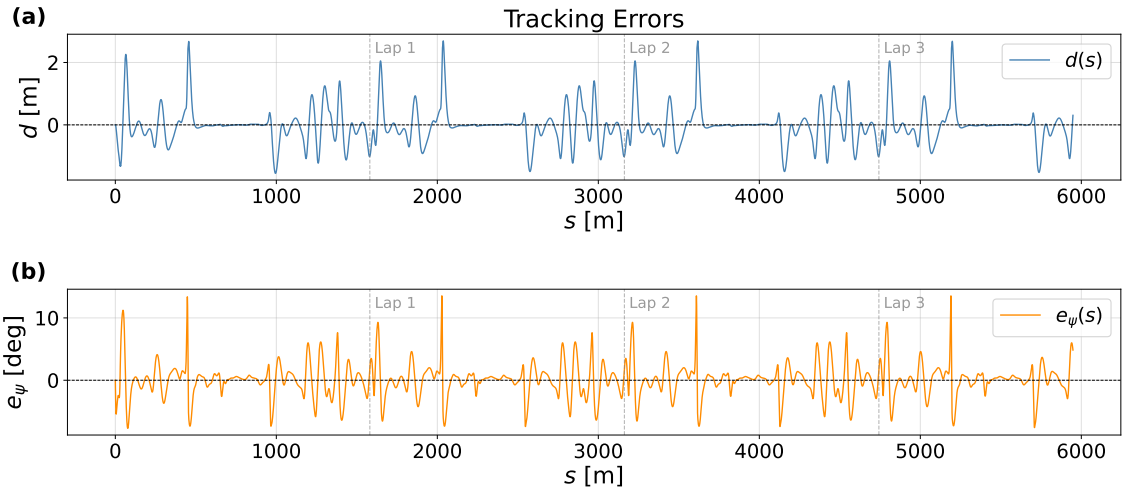

Figure 5.39: Plots showing lateral offset d (a) and heading error e_ψ (b) for the adaptive lateral PID controller. It is clearly observable that the lateral performance of the controller is not nearly as good as indirect controllers.

Furthermore, it can be observed in figure 5.40, that the control input experiences extreme changes whenever a turn is encountered. In figure 5.41, this is reflected in the form of sharp spikes in the gains.

The inadequacy of the adaptive PID controller can be tracked back to the tuning of the reference models and adaptation constants. The adaptive lateral PID controller that was developed in this work is a variation of the MRAC approach. As such, similar to the MRAC approach, the controller's performance is greatly dependent on how well the chosen reference model mirrors the actual plant model.

For a highly non-linear system like the single track dynamic vehicle model, one cannot simply choose a reference model via pole placement and expect it mirror the plant model perfectly. This one of the main disadvantages of the direct MRAC based adaptive lateral PID controller.

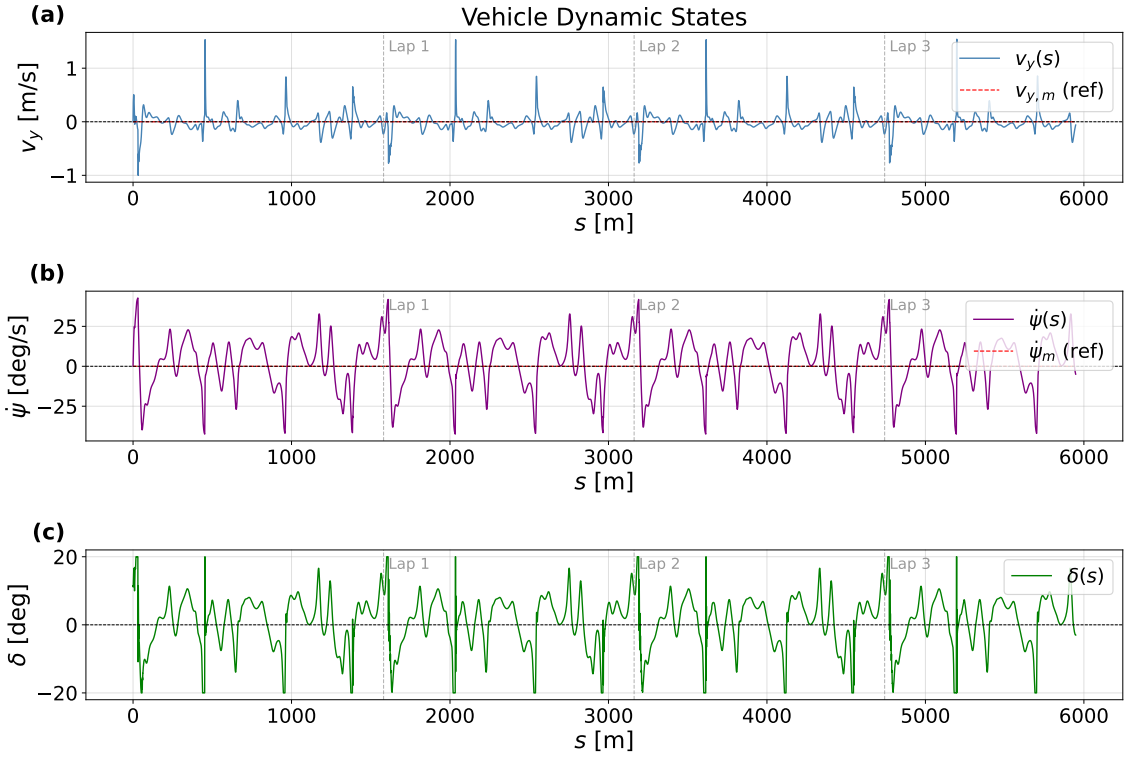

Figure 5.40: Plots showing other vehicle states when the adaptive lateral PID controller is used. Subplot (a) shows the lateral velocity v_y of the vehicle at various points on the track, while subplot (b) shows the yaw rate $\dot{\psi}$. Subplot (c) shows the optimal control input δ provided by the adaptive lateral controller.

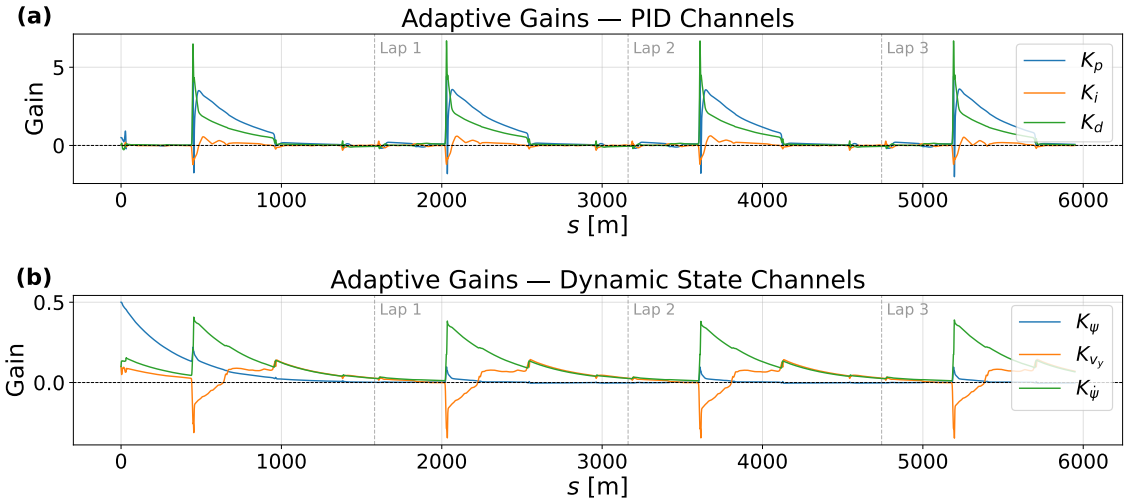

Figure 5.41: Plots showing the evolution of the gains $\begin{bmatrix} K_p & K_i & K_d & K_\psi & K_{v_y} & K_\omega \end{bmatrix}$ for the adaptive lateral PID controller. Since the steady-state value for the gains are near zero, it is clear that most of the control action is done by the feed-forward input.

5.4 Discussion of Simulation Findings

From the various scenarios and simulation tests, it is clear that the indirect approach to adaptive control is significantly more recommendable.

In figures 5.42 and 5.43, we can clearly observe the difference in lateral performance. The EKF and UKF based indirect adaptive lateral LQR controllers (represented in blue and green) substantially outperform the direct adaptive lateral PID controller (orange), with their results closely overlapping one another and remaining comparatively flatter.

This gap in performance is because of the two key obstacles encountered while implementing the direct adaptive approach,

- *Derivation of the adaptation law:* Depending on plant to be controlled, a unique adaptation law must be derived specifically for it. The complexity of the derivation and the Lyapunov equation to be used in once again dependent on the complexity of the plant model.

- *Accuracy of the reference model:* The performance of the controller is directly linked to the accuracy of the chosen reference model, which in turn has several restriction on it. The essential restrictions being it should be linear and be Hurwitz stable. For an extremely non-linear plant model, choosing a reference model satisfying these criteria is not simple.

Furthermore, since the indirect architecture decouples parameter estimation from controller design, the controller (LQR) gains are continuously recomputed from accurate parameter estimates regardless of the model's complexity. The adaptive PID controller, however, is constrained by the fixed structure and the stability requirements of its reference model, limiting how closely it can track the true plant dynamics as conditions change.

In contrast, continuing on the findings from section 5.2.4, for the development of the indirect adaptive lateral controllers, choosing a controller between the EKF and UKF variations is much more nuanced decision. The EKF based controller and its adaptive counterpart, is a computationally efficient approach which can be recommended for easy to linearize plant models and situations where Jacobian linearizations do not break.

On the other hand the UKF based controllers are an excellent choice for highly non linear plant models or for situations where Jacobian linearizations cannot be guaranteed.

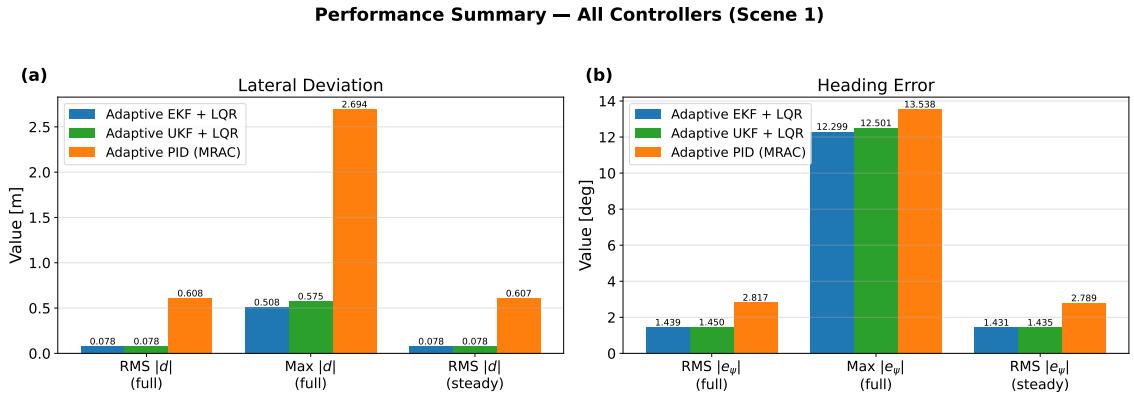

Figure 5.42: Plots showing the overall lateral performance of the three control approaches. The blue and green bars represent the adaptive EKF and UKF based indirect adaptive controllers respectively, and the orange bar represents the direct adaptive PID controller. Across the board, the indirect controllers perform significantly better.

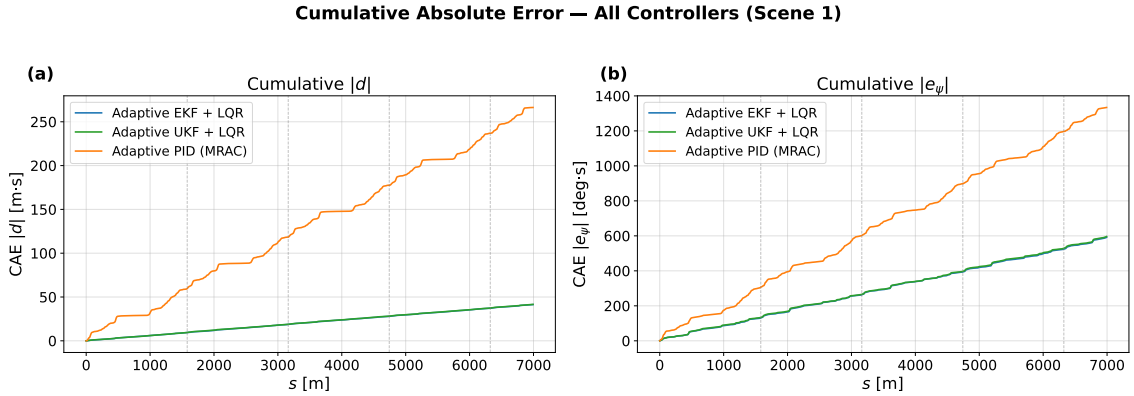

Figure 5.43: Plots showing the CAE of d (a) and e_ψ (b) of the three control approaches. The plots for the EKF and UKF based indirect controllers significantly overlap each other. Moreover these plots are much flatter than the plot for the direct adaptive PID controller, indicating significant difference in performance capabilities

6

Conclusion

The objective of this thesis was to investigate, develop and evaluate adaptive lateral control strategies for autonomous heavy vehicles operating in haulage systems. Particular emphasis was placed on testing the robustness, adaptability and generalizability of different adaptive control methodologies under varying operating conditions and dynamic testing scenarios.

6.1 Answers to Research Questions

RQ1: Which adaptive control methods are most suitable for lateral control of autonomous heavy vehicles?
In this work, two major classes of adaptive control approaches, direct and indirect adaptive control, were studied and integrated with established lateral control techniques such as PID and LQR controllers. The lateral PID controller used a direct adaptive control strategy developed using the MRAC framework, whereas the LQR used an indirect adaptive method where a Kalman filter based parameter estimator was implemented alongside the controller. Overall, it was found that the indirect adaptive approach yielded better lateral performance and is better suited for the autonomous control of heavy vehicles.

RQ2: How effectively do the proposed controllers adapt to time-varying vehicle dynamics, and how robust are they across changes in parameter values?
Among the direct and indirect adaptive strategies, the latter demonstrated strong adaptability across multiple scenarios. Both (EKF and UKF) estimators converged to a 5% error of true parameters within a few steps, enabling rapid LQR gain convergence, and successfully re-converged after a 20% mass change, with the UKF outperforming the EKF at low velocities. Under a short-duration disturbance, the estimators only partially adapted, exposing a re-convergence speed to noise trade-off, though this still outperformed the fixed-gain baseline. The direct adaptive PID, in contrast, exhibited sharp gain spikes during turns and required extensive manual tuning.

RQ3: What are the computational requirements and onboard implementation constraints of the proposed adaptive controllers, and how generalizable are they across different truck configurations?
In the indirect approach, the EKF estimator requires $\mathcal{O}(n^2)$ Jacobian evaluations per step, while the UKF one costs $\mathcal{O}(n^3)$ due to the Cholesky decomposition for its $2n + 1$ sigma points; for $n = 8$, both ran comfortably within the 0.02 s simulation time step, indicating real-time feasibility. Furthermore, the indirect architecture is highly generalizable, as the estimator operates independently and can pair with different controllers (LQR, MPC) and vehicle configurations by updating the model equations and estimated parameters. The direct adaptive PID controller is computationally cheaper, avoiding the matrix operations required by the EKF/UKF, but this efficiency comes at the cost of robustness: as noted in RQ2, it exhibited sharp gain spikes during turns and required extensive manual tuning of the reference model and adaptation gain matrix Γ, and its adaptation law must be re-derived for each vehicle configuration, limiting scalability.

6.2 Summary of Findings

Across all three test scenarios in the section 5.2, the indirect adaptive lateral LQR controllers reduced the steady-state RMS lateral deviation by $83\% - 85\%$ and the steady-state RMS heading error by $49\% - 53\%$ relative to the fixed-gain LQR baseline. This improvement was consistent regardless of whether the vehicle operated under constant conditions (5.2.1), experienced a 20% mass change (5.2.2), or traversed a low-friction surface (5.2.3).

Between the two Kalman filters and its adaptive variants, the EKF and UKF produced nearly identical steady-state metrics, differing by at most 3%–4% across all scenarios. The practical distinction lies in their transient behavior: the UKF maintained stable estimates at near-zero velocities during loading and unloading events (Figures 5.17 and 5.19), where the EKF exhibited estimation drift due to Jacobian singularities. More broadly, since the UKF propagates sigma points through the actual nonlinear dynamics of the system rather than a linearized model, it is better suited to highly nonlinear systems, where Jacobian based linearization errors can become significant. This makes the UKF based parameter estimator a suitable choice for applications involving frequent stops, low-speed maneuvers, or strongly nonlinear dynamics. Meanwhile, the EKF based estimator remains a viable and computationally lighter alternative for scenarios with milder nonlinearities.

A notable secondary benefit of the indirect adaptive architecture is that the estimated parameters are available in real time as standalone outputs. Beyond their use in controller adaptation, these estimates can serve as inputs to vehicle health monitoring, tire degradation tracking, or load verification systems, adding diagnostic value without additional sensor hardware.

The direct adaptive lateral PID controller, while functional, underperformed in terms of lateral performance and accuracy. Its peak lateral deviation during turns

was approximately five times larger than that of the indirect controllers (Figures 5.42), and the control input exhibited high-frequency oscillations at path curvature transitions (Figure 5.40 (c)). These limitations stem from the structural mismatch between the linear reference model and the nonlinear plant, rather than from the MRAC framework itself, suggesting that further work on nonlinear reference model design could improve its viability.

In conclusion, the findings of this thesis demonstrate that indirect adaptive control strategies, particularly those based on Kalman filter estimation combined with model based control strategies like LQR, offer a promising and practically viable solution for adaptive lateral control in AHS.

6.3 Future Work

While this thesis demonstrates the viability of adaptive lateral controllers for autonomous heavy vehicles, the work naturally opens multiple avenues for further research:

- **Higher-fidelity vehicle models:** This work uses a single-track dynamic bicycle model for the development and testing of the adaptive controllers. Extending the framework to a fully dynamic multi-axle vehicle model with nonlinear tire force representations (e.g., Pacejka Magic Formula) would improve the fidelity of both the simulation and the estimation.

- **Cubature Kalman Filter (CKF) based estimator:** The CKF offers similar nonlinear parameter estimation capabilities as the UKF but with fewer tuning parameters, as it uses a fixed set of cubature points rather than the user-defined α, β and κ scaling parameters. Investigating the CKF as an alternative estimator could simplify filter tuning while maintaining estimation accuracy.

- **Dual estimation framework:** As discussed in Appendix A.1.1, a dual estimation setup, where separate filters estimate states and parameters independently, may increase numerical stability, computational efficiency and offer better support for different update rates between state and parameter estimation.

- **Experimental validation on Functional Mock-up Units (FMU):** The controllers developed in this thesis have been evaluated solely in simulation. Validation on a real-world autonomous truck platform or FMU, with actual sensor noise, communication delays, and actuator dynamics, is a necessary step toward deployment.

- **Further investigation into direct adaptive control methods:** Although the direct adaptive MRAC-based lateral PID controller did not perform competitively in this study, as compared to the indirect controllers, the underlying

approach remains conceptually promising and warrants further investigation.

Bibliography

[1] Komatsu Ltd. (2018) *Komatsu celebrates 10th anniversary of commercial deployment of Autonomous Haulage System (AHS)*. Available at: https://www.komatsu.jp/en/newsroom/2018/20180129.

[2] Komatsu Ltd. (2018) *Autonomous Haulage System Overview and Deployment Milestones*. Available at: https://www.komatsu.jp/en/newsroom/2018/20180129.

[3] Rio Tinto (2024) *Autonomous Haulage Operations in Pilbara Mining*. Available at: https://www.riotinto.com.

[4] Caterpillar Inc. (2024) *Autonomous Mining and Haulage Systems Overview*. Available at: https://www.caterpillar.com.

[5] Teng, S., Li, X., Li, Y., Xuanyuan, Z., Ai, Y. and Chen, L. (2024) *Scenario Engineering for Autonomous Transportation: A New Stage in Open-Pit Mines*. arXiv:2405.15772. Available at: https://arxiv.org/abs/2405.15772.

[6] Volvo Autonomous Solutions (2023) *Autonomous Transport Solution (ATS) Quarry and Hub-to-Hub Demonstrations*. Available at: https://www.volvogroup.com/en/news-and-media.html.

[7] Rajamani, R. (2012) *Vehicle Dynamics and Control*. 2nd edn. New York: Springer.

[8] Kebbati, Y., Ait-Oufroukh, N., Ichalal, D. and Vigneron, V. (2023) *Lateral control for autonomous wheeled vehicles: A technical review*. Asian Journal of Control, 25(4), pp. 2539–2563. doi: https://doi.org/10.1002/asjc.2980.

[9] Jawahar, A. et al. (2023) *Lateral Control of Heavy Vehicles*. Master's thesis, KTH Royal Institute of Technology, in collaboration with Volvo GTT. Available at: https://kth.diva-portal.org/smash/get/diva2:1817230/FULLTEXT02.pdf.

[10] Åström, K. J. and Wittenmark, B. (2013) *Adaptive Control*. 2nd edn. Dover

Publications.

[11] Python Control Systems Library (n.d.) *MRAC SISO with MIT Rule Example*. Available at: https://python-control.readthedocs.io/en/0.10.1/mrac_siso_mit.html.

[12] Python Control Systems Library (n.d.) *MRAC SISO with Lyapunov Rule Example*. Available at: https://python-control.readthedocs.io/en/0.10.1/mrac_siso_lyapunov.html.

[13] MathWorks (n.d.) *Model Reference Adaptive Control (MRAC) Documentation*. Available at: https://se.mathworks.com/help/slcontrol/ug/model-reference-adaptive-control.html.

[14] Kersting, S. and Buss, M. (2017) *Direct and Indirect Model Reference Adaptive Control for Multivariable Piecewise Affine Systems*. IEEE Transactions on Automatic Control, 62(11), pp. 5634–5649. Available at: https://ieeexplore.ieee.org/abstract/document/7890417/.

[15] Ioannou, P. A. and Kokotovic, P. V. (1984) *Instability analysis and improvement of robustness of adaptive control*. Automatica, 20(5), pp. 583–594. doi: https://doi.org/10.1016/0005-1098(84)90009-8.

[16] Chebbi, A., Franchek, M. A. and Grigoriadis, K. (2025) *Simultaneous State and Parameter Estimation Methods Based on Kalman Filters and Luenberger Observers: A Tutorial and Review*. Sensors, 25(22), 7043. Available at: https://www.mdpi.com/1424-8220/25/22/7043.

[17] Sierra, C., Tseng, E., Jain, A. and Peng, H. (2006) *Cornering stiffness estimation based on vehicle lateral dynamics*. Vehicle System Dynamics, 44(sup1), pp. 24–38. doi: https://doi.org/10.1080/00423110600867259.

[18] Levy, A. and Klein, I. (2026) *Multi-Scaled Unscented Kalman Filter*. arXiv:2604.04792. Available at: https://arxiv.org/abs/2604.04792.

[19] Akhlaghi, S., Zhou, N. and Huang, Z. (2017) *Adaptive Adjustment of Noise Covariance in Kalman Filter for Dynamic State Estimation*. arXiv:1702.00884. Available at: https://arxiv.org/abs/1702.00884.

[20] Mendes, A. (n.d.) *OpenVD Tire Models Documentation*. Available at: https://andresmendes.github.io/openvd/build/html/tire.html.

[21] Huang, J., He, Z., Arakawa, Y. and Dawton, B. (2023) *Trajectory Planning in Frenet Frame via Multi-Objective Optimization*. IEEE Access, 11, pp. 1–1. doi: https://doi.org/10.1109/ACCESS.2023.3294713.

A

Appendix 1

A.1 Additional Theory and Plots

This section contains multiple additional background theory and plots relevant at various different points in the thesis. The aim of these sections is to help the reader gain a deeper understanding of the relevant topics.

A.1.1 Dual state and parameter estimation

While this concept has not been implemented in the current simulation setup, the theory is nevertheless important in the context of future work.

In a dual estimation setup, two independent filters, one for state estimation and parameter estimation each, operate in parallel. The state filter uses the current process and measurement equations to estimate the state at each time instant, whereas the parameter filter receives these state estimates to refine the predict and refine the parameter estimates.

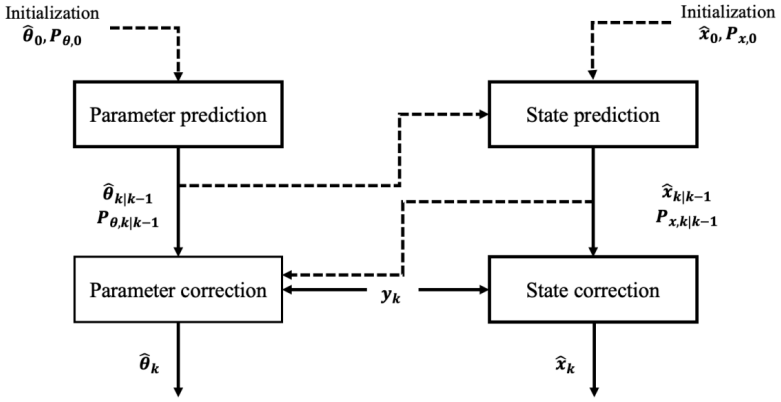

Figure A.1: Block diagram illustrating Dual Estimation Schematic [16]

While the traditional simultaneous estimation is a more optimal and theoretically more efficient approach, the dual estimation offers a more robust and tuning flexible approach, especially in highly non-linear and real-world scenario. In particu-

lar, it provides better numerical stability by avoiding the propagation of large and potentially ill-conditioned augmented covariance matrices, and reduces the risk of divergence caused by strong state–parameter coupling.

It also naturally supports different time scales. In a real-world implementation, it is more preferable to estimate the parameters at a slower time scale than the states, since they usually evolve much slowly. Furthermore, the dual estimation framework facilitates this by allowing the parameter filter to be updated with lower process noise, thereby preventing overreaction to measurement noise and reducing estimation variance.

A.1.2 Hard constraints

The constraint equations 4.10 and 4.11 are also known as hard constraints because it creates a fixed numerical link between two unknown constants. This reduces the effective number of unknown parameters to two (for example C_f^* and L_f). The other two parameters can be computed from the constraint equations.

While this will undoubtedly increase the performance of our Kalman estimator, there is one main issue. It should be known that the state uncertainty also propagates through these constraint equations. While this is not a concern for the computation of L_r, the same cannot be said for C_r^*, as the computed estimate and variance for it would be,

$$\hat{C}_r^* = \hat{C}_f^* \frac{\hat{L}_f}{\hat{L}_r}$$

$$\sigma_{C_r^*}^2 = \left(\frac{\partial C_r^*}{\partial C_f^*}\right)^2 \sigma_{C_f^*}^2 + \left(\frac{\partial C_r^*}{\partial L_f}\right)^2 \sigma_{L_f}^2 + \left(\frac{\partial C_r^*}{\partial L_r}\right)^2 \sigma_{L_r}^2$$

Meaning that the uncertainty for C_r^* will be much greater than the rest.

A.1.3 Indirect adaptive lateral MPC controller

Similar to the lateral adaptive control approach using a LQR controller, another alternative for a lateral controller is the model predictive controller (MPC). It is an indirect approach as well, where the unknown parameters are first computed using estimation methods such as RLS, EKF, UKF, etc., and then used to minimize an optimization problem over a finite horizon.

While the MPC approach is more powerful and accurate, as it accounts for future predictions too, than a LQR controller, it is significantly more computationally expensive and difficult to implement in a real time dynamic systems where execution speed matters.

Thus, this approach to an adaptive controller was not deeply investigated in this work. However, a brief introduction to the topic and it's implementation is provided here for reference.

A.1.3.1 The lateral MPC controller

Let us consider a Kinematic bicycle model described by the equations.

$$\begin{bmatrix} \dot{d} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} v_x \sin(\psi - \varphi(s)) \\ \frac{v_x}{L} \tan(\delta) \end{bmatrix} \tag{A.1}$$

where the model is dependent upon the parameters v_x and L. These parameters must be first computed via parameter estimation before computing an optimal input, thereby making the controller an indirect adaptive strategy. The optimal steering angle δ is then calculated by solving an optimization problem over a finite horizon,

$$\begin{aligned} \min\ J = \sum_{k=0}^{N-1} \left(Q_d d_k^2 + Q_\psi \left(\psi_k - \phi_k\right)^2 + R\delta_k^2 + R_\delta \left(\delta_k - \delta_{k-1}\right)^2 \right) \\ + k(Q_d d_N^2 + Q_\psi \left(\psi_N - \phi_N\right)^2) \end{aligned} \tag{A.2}$$

where,

- $Q_d d_k^2 + Q_\psi \left(\psi_k - \phi_k\right)^2 + R\delta_k^2$ is the running cost which encourages the controller to reduce error and not take large steering angles.
- $R_\delta \left(\delta_k - \delta_{k-1}\right)^2$ is the penalty on the control input.
- $k(Q_d d_N^2 + Q_\psi \left(\psi_N - \phi_N\right)^2)$ is the terminal cost. It encourages the controller to converge the state by the end of the horizon.

it should be taken care that while the solver is moving through the horizon, it take the curvature of the road $\kappa(s)$ into consideration. This is because, when the vehicle is traveling through a curve, there would be a non-zero curvature which subsequently would produce a change in the heading $\psi(s)$ of the vehicle.

From the optimization problem equation A.2, it is clear that the size and difficulty of the problem will increase with the increase in the complexity of the system model used. Furthermore, in order to increase the optimality of the input, the horizon length of the controller should also be sufficiently long. Both these factors significantly increase the computation time and effort required by the controller, making it a difficult to recommend approach to practical scenarios.

A.1.4 Estimation re-convergence for non adaptive Kalman filters

Below are a few plots that were obtained during the development of this work. The purpose of this plot is to show the difference in the re-convergence speed of an adaptive and non adaptive controller, and the conditions behind the simulation scenario are not the same as the ones discussed in section 5.1. It must also be noted that the parameters μ_f and μ_r represent the parameters C_f^* and C_r^* respectively.

The key detail to note in the figure A.3 is the broadening of the state uncertainty when the value of the unknown parameter changes. Since a normal Kalman filter is unable to do this, the estimation correction speed is significantly lower in figure A.2.

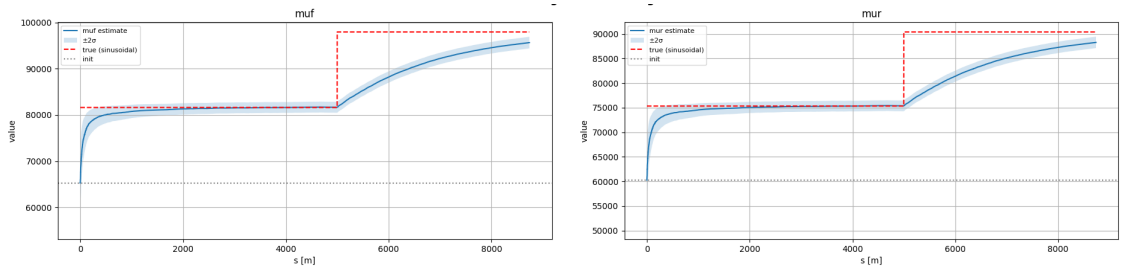

Figure A.2: Estimation plots for a non adaptive EKF estimator

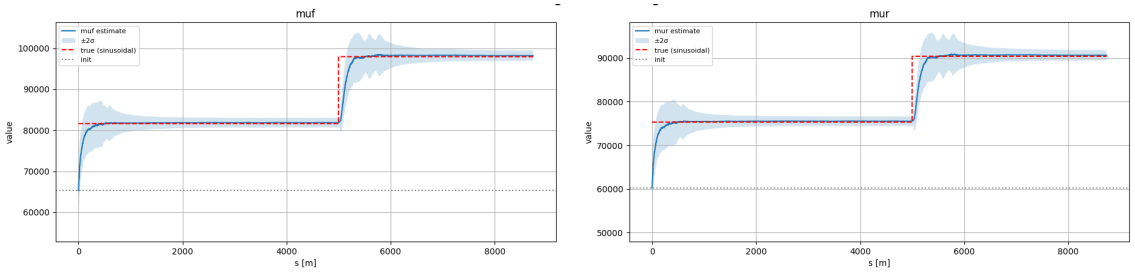

Figure A.3: Estimation plots for an adaptive EKF estimator

A.1.5 Noise-free EKF based adaptive lateral LQR control plots

In section 5.2.1.2, the noisy nature of the plots was attributed to the presence of a simulated sensor noise in the simulation. This section showcases a simulation setup similar to scenario 1 in section 5.2.1.2 with an EKF based parameters estimator, but without the use of the sensor noise.

From the figures A.4 and A.5, and the table A.1, we can see that the estimator performance slightly degrades overtime. This decrease in performance can be explained by the gradual loss in persistent excitation of the system. As the system approaches steady state operation, without the presence of the sensor noise to provide a form of excitation, the amount of information available for the estimator reduces since the controller actively pushes the states to their steady state values (see section 2.5.4).

Param	True	Initial	Final Est	Err (%)
C_f^*	800 000	640 000	834 762.9 ± 11 719.3	4.35
C_r^*	1 200 000	960 000	1 252 144.4 ± 17 577.6	4.35
L_f	2.8	2.2	2.9 ± 0.03	2.73
L_r	1.8	1.4	1.7 ± 0.03	4.25

Table A.1: EKF Parameter Estimation Results

While this test scenario does not mimic a real world situation, it helps in performing a sanity check on the system. It should be noted that model-based approaches such as LQR and MPC are best suited for systems where the states are accurately known. Thus, while these controllers are noise aware to a certain extent, they should always be used in tandem to a state-parameter estimator so as to minimize the effect of the noise in the system and to prevent the controller from generating a noisy and aggressive control input.

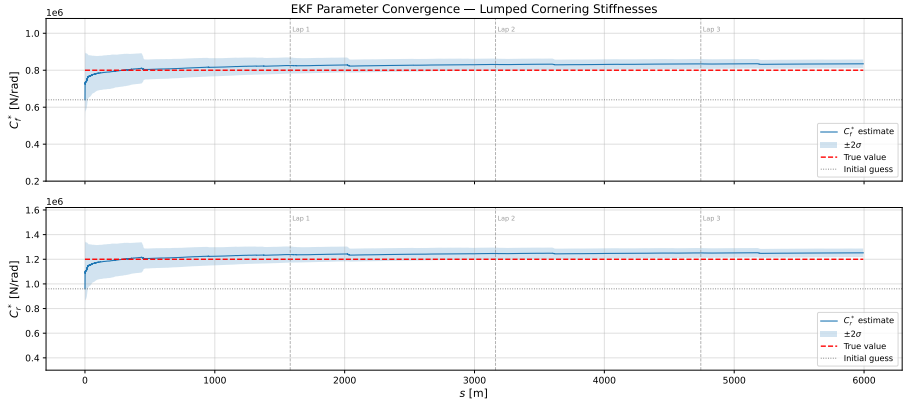

Figure A.4: Estimation of C_f^* and C_r^* using EKF (without sensor noise)

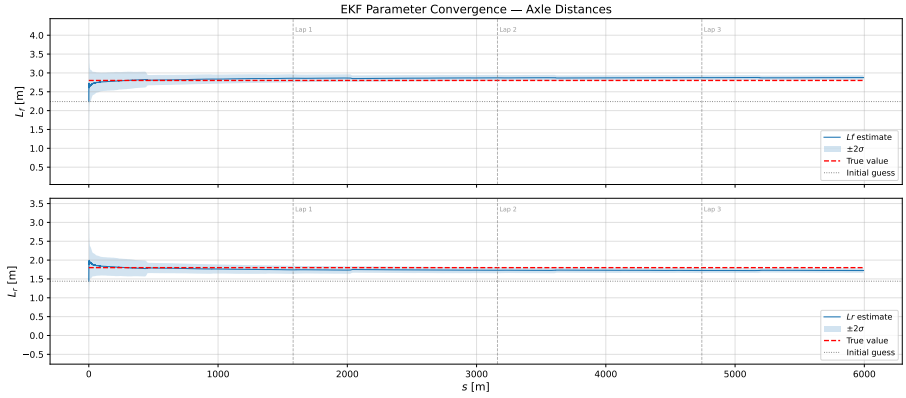

Figure A.5: Estimation of L_f and L_r using EKF (without sensor noise)

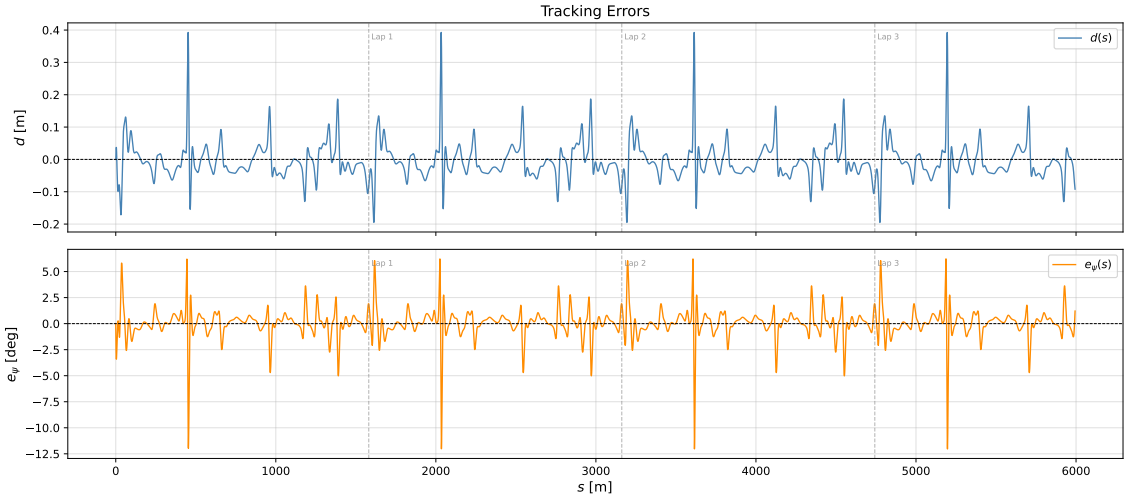

Figure A.6: Lateral tracking error using EKF (without sensor noise)

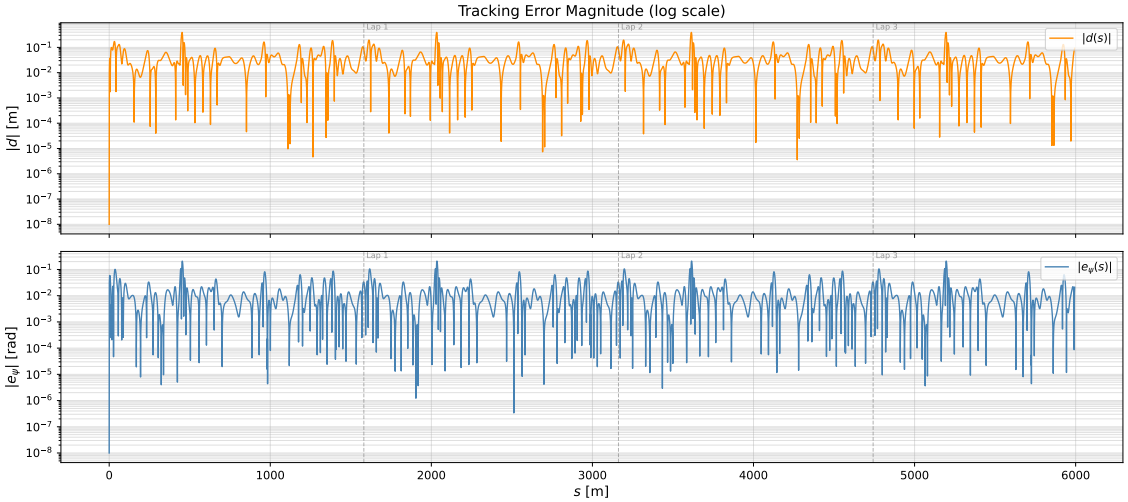

Figure A.7: Semilog lateral tracking error using EKF (without sensor noise)

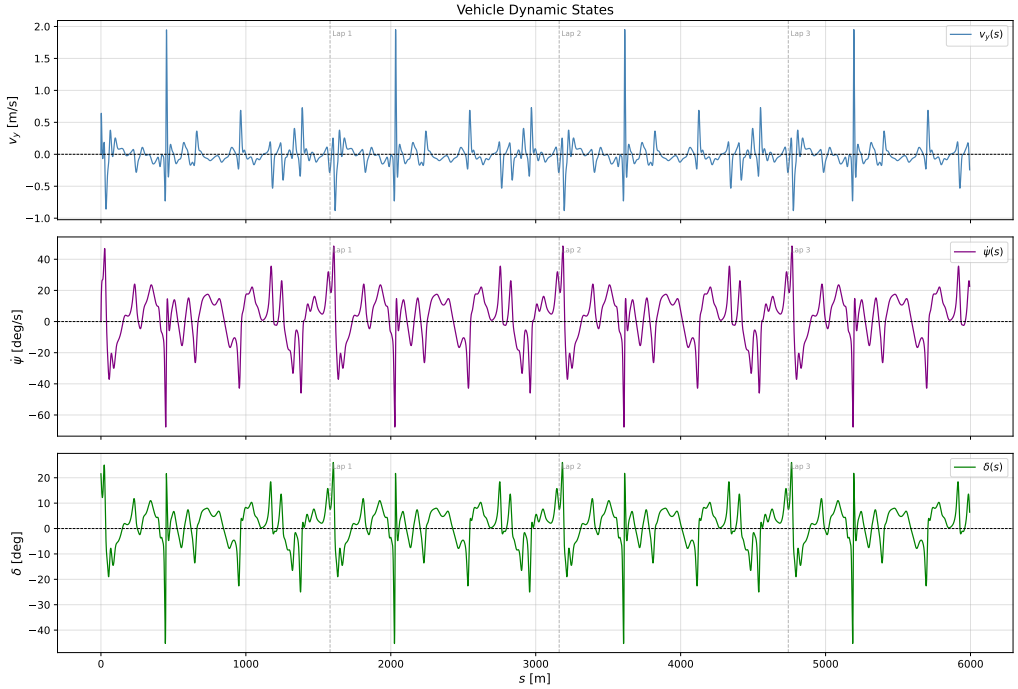

Figure A.8: Vehicle states and input using EKF (without sensor noise)

A.1.6 LQR performance plots for an adaptive Kalman filter

In the section 5.2.2.2, the performance of the lateral LQR controller was found to be extremely similar to that in scenario 1. As nothing particularly new could be gleaned from them, the plots relevant to them were not shown in chapter 5. Nevertheless, those plots have been presented here for the reader's reference.

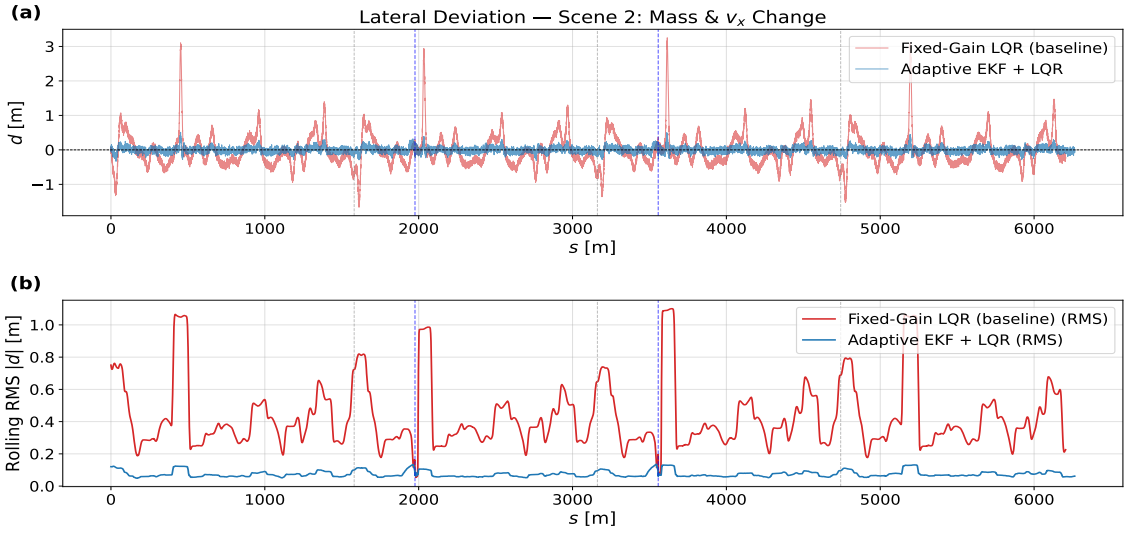

Figure A.9: Plots showing lateral offset d of the vehicle with (blue line) and without (red line) an adaptive LQR controller using a parameter drift rate adaptive EKF. Subplot (a) shows the raw d value for the simulation, whereas subplot (b) shows the rolling RMS of $|\, d \,|$ within a window of 500 time steps.

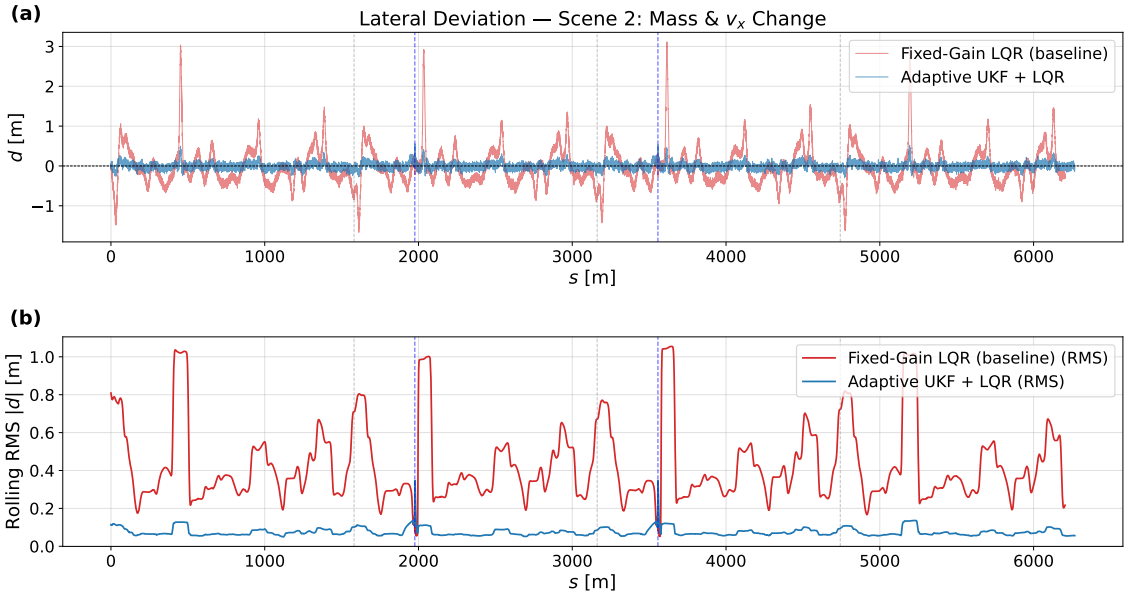

Figure A.10: Plots showing lateral offset d of the vehicle with (blue line) and without (red line) an adaptive LQR controller using a parameter drift rate adaptive UKF. Subplot (a) shows the raw d value for the simulation, whereas subplot (b) shows the rolling RMS of $\mid d \mid$ within a window of 500 time steps.

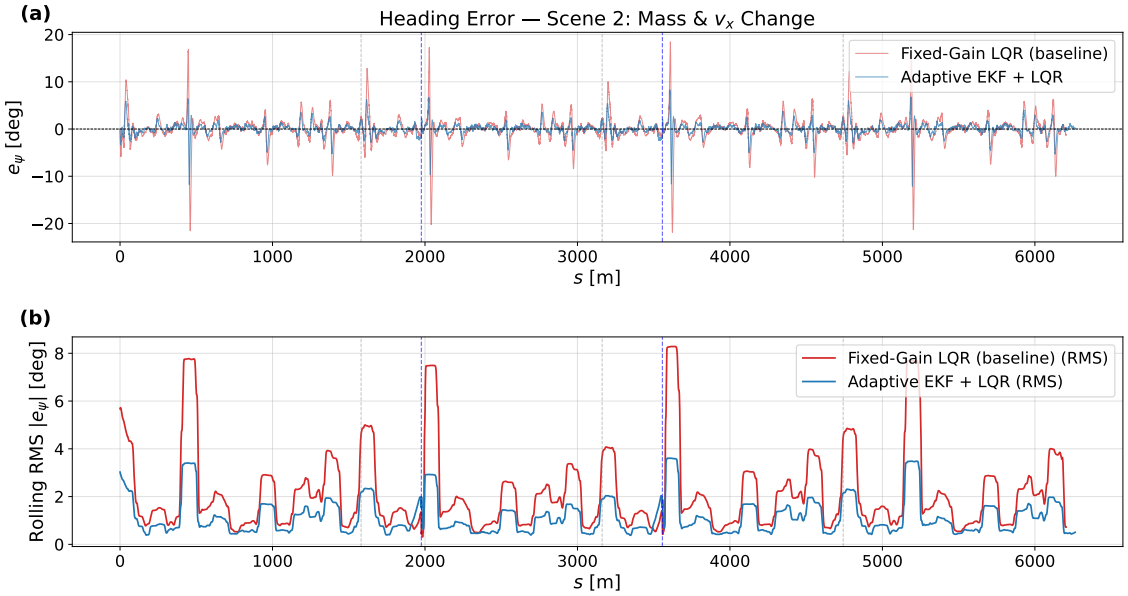

Figure A.11: Plots showing lateral offset e_ψ of the vehicle with (blue line) and without (red line) an adaptive LQR controller using a parameter drift rate adaptive EKF. Subplot (a) shows the raw e_ψ value for the simulation, whereas subplot (b) shows the rolling RMS of $\mid e_\psi \mid$ within a window of 500 time steps.

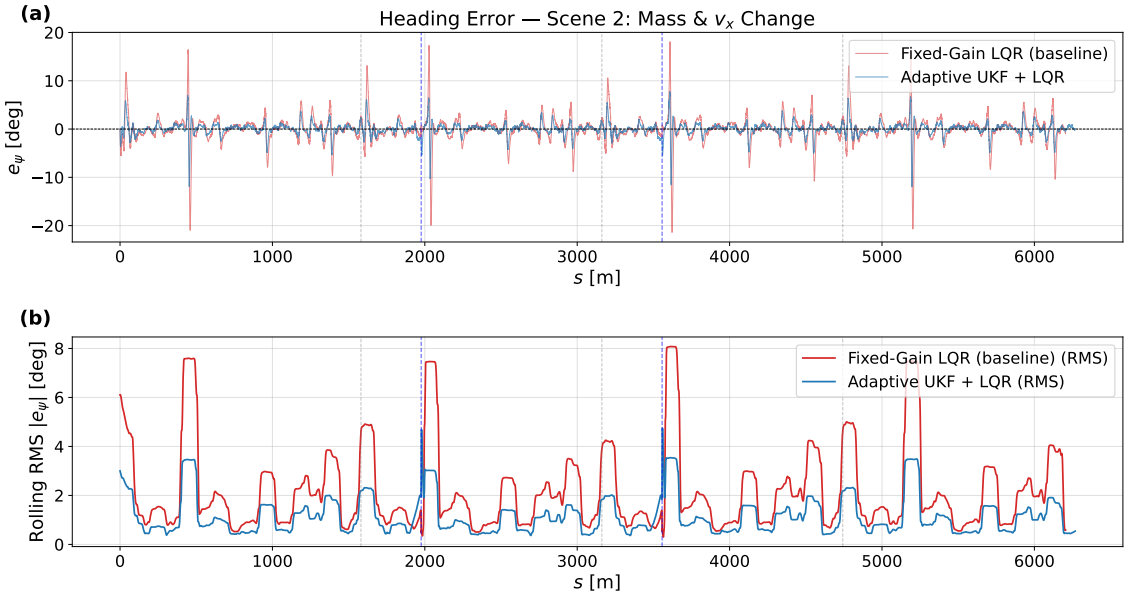

Figure A.12: Plots showing lateral offset e_ψ of the vehicle with (blue line) and without (red line) an adaptive LQR controller using a parameter drift rate adaptive UKF. Subplot (a) shows the raw e_ψ value for the simulation, whereas subplot (b) shows the rolling RMS of $\mid e_\psi \mid$ within a window of 500 time steps.

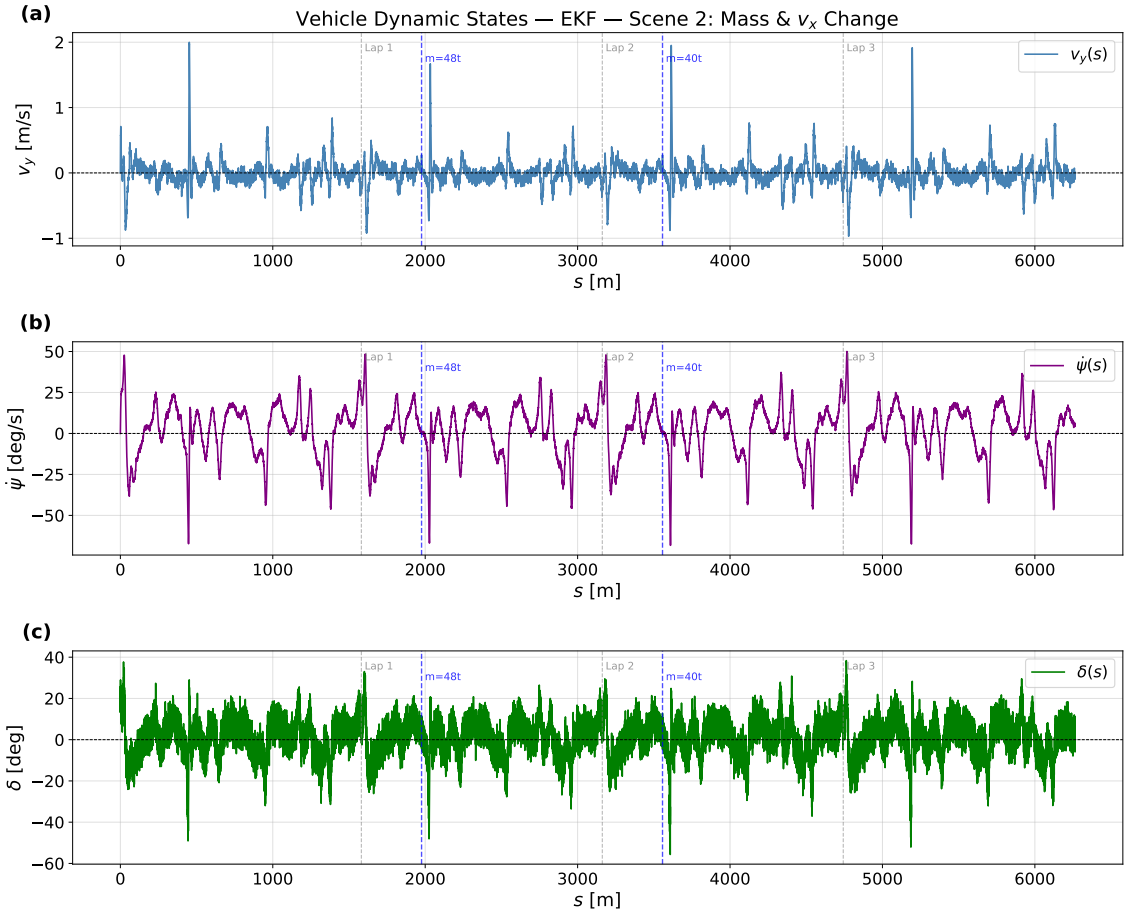

Figure A.13: Plots showing other vehicle states when the adaptive LQR controller using a parameter drift rate adaptive EKF is used. Subplot (a) shows the lateral velocity v_y, while subplot (b) shows the yaw rate $\dot{\psi}$. Subplot (c) shows the optimal control input δ provided by the adaptive lateral controller.

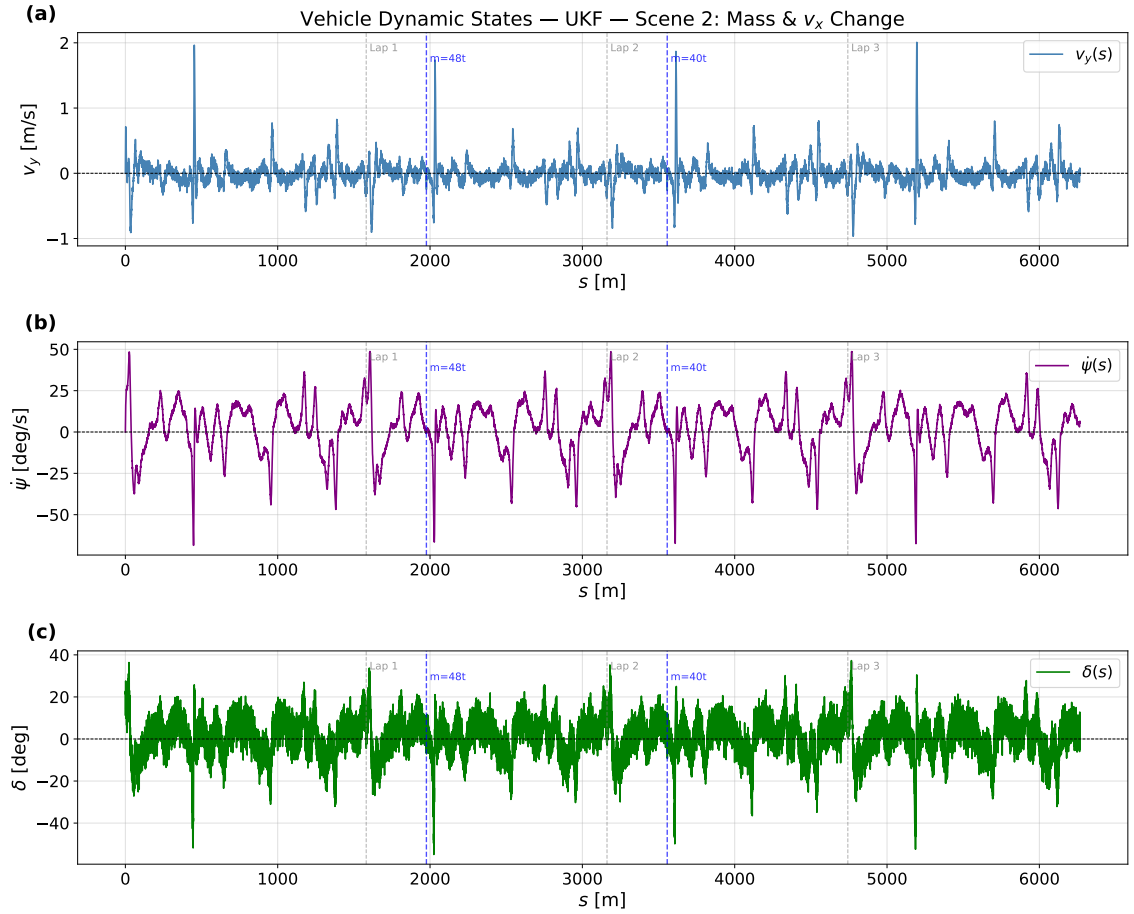

Figure A.14: Plots showing other vehicle states when the adaptive LQR controller using a parameter drift rate adaptive UKF is used. Subplot (a) shows the lateral velocity v_y, while subplot (b) shows the yaw rate $\dot{\psi}$. Subplot (c) shows the optimal control input δ provided by the adaptive lateral controller.

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

CHALMERS
UNIVERSITY OF TECHNOLOGY