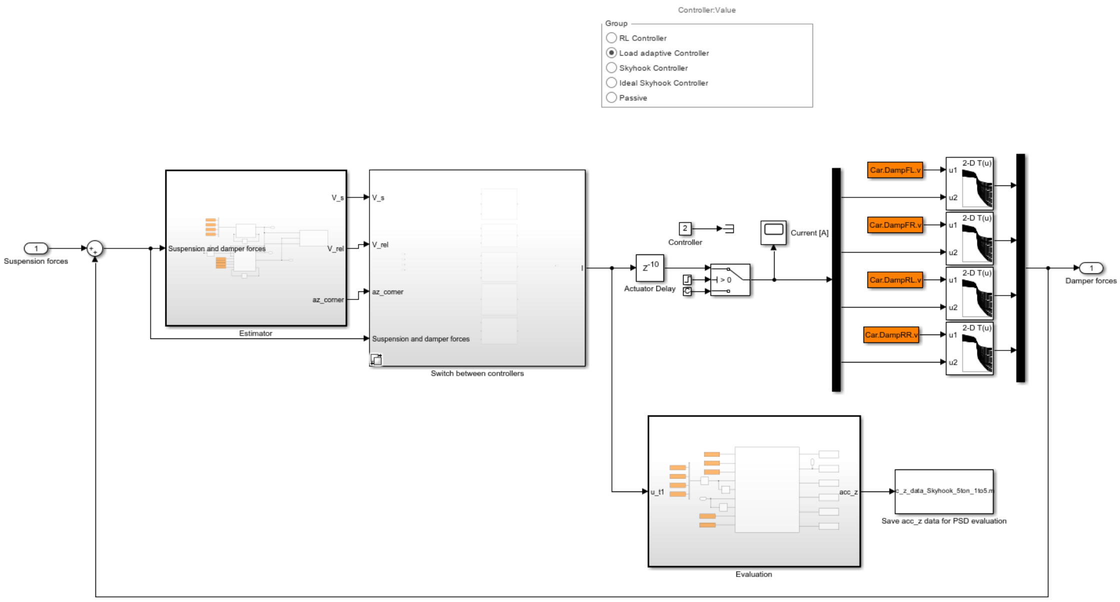

Controlling Semi-Active Damper for Heavy Duty Trucks

A Comparison of Control Algorithms for Increasing Ride Comfort with Semi-Active Dampers

Master's thesis in Systems, Control and Mechatronics

GUSTAV HALLER
GILLIS MAGNUSSON

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's Thesis in The Name of Systems, Control, and Mechatronics

Controlling Semi-Active Damper for Heavy Duty Trucks

A Comparison of Control Algorithms for Increasing Ride Comfort with Semi-Active Dampers

GUSTAV HALLER

GILLIS MAGNUSSON

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
Gothenburg, Sweden 2026

Controlling Semi-Active Dampers for Heavy Duty Trucks
A Comparison of Control Algorithms for Increasing Ride Comfort with Semi-Active Dampers

GUSTAV HALLER
GILLIS MAGNUSSON

Supervisor & Examiner: Mats Jonasson. Department of Mechanics and Maritime Sciences

Master's Thesis 2026
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Simulink model of the semi-active suspension control system, including state estimation, controller selection, actuator delay, and individual damper force mapping for each wheel.

Typeset in LaTeX
Gothenburg, Sweden 2026

Controlling Semi-Active Dampers for Heavy Duty Trucks
A Comparison of Control Algorithms for Increasing Ride Comfort with Semi-Active Dampers
GUSTAV HALLER
GILLIS MAGNUSSON
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology

Abstract

This thesis evaluates semi-active suspension control for a heavy-duty 4x2 truck under varying payload conditions, with the aim of improving ride comfort while maintaining implementation feasibility. A rigid full-vehicle model was implemented in IPG TruckMaker and integrated with MATLAB/Simulink. Five damper configurations were compared: Passive, Load-Adaptive, Skyhook with estimated states, Ideal Skyhook, and a Reinforcement Learning controller based on TD3.

The controllers were evaluated on a 1–5 Hz swept-frequency road and on stochastic road profiles generated according to ISO 8608 classes A–F, using unloaded, 5,000 kg, and 10,000 kg hitch load cases. Performance was assessed primarily through RMS of vertical acceleration, supported by frequency-domain analysis and power consumption.

The results show that the benefit of semi-active control depends strongly on road roughness. On smooth roads, the passive damper remained highly competitive. On rougher roads, active controllers consistently reduced vertical acceleration relative to passive damping. Reinforcement Learning achieved the lowest RMS vertical acceleration in many test cases, but required higher power and showed degraded performance outside its training distribution. Skyhook with estimated states provided the best overall balance between comfort improvement, low power consumption, and implementation feasibility, while Ideal Skyhook offered strong performance at the cost of additional sensing requirements.

The Load-Adaptive controller maintained the intended damping ratio across payloads and was the simplest active strategy to implement at the lowest computational cost. However, its comfort benefit was limited at high load while power consumption increased substantially, suggesting that disabling the load-adaptive logic in operating regions where passive damping already performs comparably would improve its energy efficiency.

Keywords: semi-active damper, heavy-duty truck, ride comfort, Skyhook control, reinforcement learning, TD3, load-adaptive damping, ISO 8608

Acknowledgements

We would like to express our sincere gratitude to everyone who supported this master's thesis. First, we would like to thank our supervisor and examiner at Chalmers University of Technology, Mats Jonasson, for his guidance, feedback, and support throughout the project.

We are also grateful to our supervisors at Volvo Group, Nrupatunga Ashok and Adam Kahlif, in the Suspension Control Systems group (BE48230), for their valuable insights, technical discussions, and continuous support during the thesis work. We also thank Carl-Johan Rundqvist, manager of the Suspension Control Systems group, for enabling and supporting this collaboration.

In addition, we would like to thank the Control System Simulations group (BE48210) at Volvo Group for their support, insights during the project and welcoming us into the group. We especially thank Robert Legeza Söderström, manager of the group, for supporting the work and making this collaboration possible.

Finally, we would like to thank Roots Food Market for all the amazing lunches and everyone else at Chalmers University of Technology and Volvo Group who contributed with knowledge, discussions, and practical support throughout the project.

Gustav Haller & Gillis Magnusson,
Gothenburg, May 2026

List of Acronyms

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

ADD	Acceleration-Driven Damping
COM	Center of Mass
DDPG	Deep Deterministic Policy Gradient
DOF	Degree of Freedom
DRL	Deep Reinforcement Learning
GH	Groundhook
LA	Load Adaptive
LUT	Lookup Table
MDP	Markov Decision Process
MPC	Model Predictive Control
PSD	Power Spectral Density
QCM	Quarter Car Model
RL	Reinforcement Learning
RMS	Root Mean Square
SH	Skyhook
TD3	Twin-Delayed Deep Deterministic Policy Gradient

Nomenclature

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

Indices

j	Index for wheel (e.g., $j \in \{\mathrm{FL}, \mathrm{FR}, \mathrm{RL}, \mathrm{RR}\}$)
ℓ	Index for axle (front, rear)
b	Index for body degree of freedom
k	Index for discrete time step or discrete vehicle state
i	Index for sampled data point, corner, or current grid point, depending on context

Parameters

m	Vehicle mass	[kg]
m_s	Sprung mass (vehicle body)	[kg]
m_u	Unsprung mass (wheel assembly)	[kg]
$m_{u,j}$	Unsprung mass at wheel j	[kg]
m_{corner}	Corner sprung mass	[kg]
I_x	Roll moment of inertia of the vehicle about the longitudinal axis	[kg m^2]
I_y	Pitch moment of inertia of the vehicle about the lateral axis	[kg m^2]
I_z	Yaw moment of inertia of the vehicle about the vertical axis	[kg m^2]
l_f	Distance from the vehicle center of gravity to the front axle	[m]
l_r	Distance from the vehicle center of gravity to the rear axle	[m]

t_f	Front track width	[m]
t_r	Rear track width	[m]
t	Track width	[m]
h_{CoM}	Height of the center of gravity above ground	[m]
k_s	Suspension stiffness	[N/m]
$k_{s,j}$	Suspension stiffness at wheel j	[N/m]
k_t	Tire vertical stiffness	[N/m]
$k_{t,j}$	Tire vertical stiffness at wheel j	[N/m]
c_s	Suspension damping coefficient	[N s/m]
$C_{\min}$	Minimum achievable damping coefficient	[N s/m]
$C_{\max}$	Maximum achievable damping coefficient	[N s/m]
C_{sky}	Skyhook damping coefficient	[N s/m]
C_{eq}	Equivalent damping coefficient	[N s/m]
C_{target}	Target equivalent damping coefficient	[N s/m]
ζ	Damping ratio	[–]
ζ_{ref}	Reference damping ratio	[–]
ω_n	Natural frequency	[rad/s]
ω_s	Sprung mass natural frequency	[rad/s]
ω_u	Unsprung mass natural frequency	[rad/s]
g	Gravitational acceleration	[m/s^2]
C_x	Longitudinal tire stiffness	[N/rad]
C_y	Lateral tire stiffness	[N/rad]
μ	Tire-road friction coefficient	[–]
ρ	Air density	[kg/m^3]
r_w	Effective wheel radius	[m]
Ω	Spatial angular frequency of the road profile	[rad/m]
Ω_0	Reference spatial angular frequency in ISO 8608	[rad/m]
G_{q0}	Road roughness coefficient in ISO 8608	[m^3/rad]
w	Waviness exponent in ISO 8608 road model	[–]
W_k	Frequency weighting factor for vertical vibration according to ISO 2631-1	[–]
T	Evaluation duration or time horizon	[s]

Variables

v_x	Longitudinal velocity of the vehicle	[m/s]
v_y	Lateral velocity of the vehicle	[m/s]
v_z	Vertical velocity of the vehicle	[m/s]
ϕ	Roll angle of the vehicle body	[rad]
θ	Pitch angle of the vehicle body	[rad]
ψ	Yaw angle of the vehicle	[rad]
z_s	Vertical displacement of the sprung mass	[m]
$z_{s,j}$	Vertical displacement of the sprung mass at wheel j	[m]
z_u	Vertical displacement of the unsprung mass	[m]
$z_{u,j}$	Vertical displacement of the unsprung mass at wheel j	[m]
z_r	Road vertical displacement	[m]
$z_{r,j}$	Road vertical displacement input at wheel j	[m]
z_{rel}	Relative vertical displacement across the suspension	[m]
$\dot{z}_s$	Vertical velocity of the sprung mass	[m/s]
$\dot{z}_{s,j}$	Vertical velocity of the sprung mass at wheel j	[m/s]
$\dot{z}_u$	Vertical velocity of the unsprung mass	[m/s]
$\dot{z}_{u,j}$	Vertical velocity of the unsprung mass at wheel j	[m/s]
$\dot{z}_{\text{rel}}$	Relative vertical velocity across the suspension	[m/s]
a_x	Longitudinal acceleration	[m/s^2]
a_y	Lateral acceleration	[m/s^2]
a_z	Vertical acceleration	[m/s^2]
$a_{z,\text{corner}}$	Estimated vertical acceleration at a vehicle corner	[m/s^2]
ω_w	Wheel angular speed	[rad/s]
F_d	Damper force	[N]
$F_{d,j}$	Damper force at wheel j	[N]
F_s	Suspension force	[N]
$F_{s,j}$	Suspension force at wheel j	[N]
$F_{x,j}$	Longitudinal tire force at wheel j	[N]
$F_{y,j}$	Lateral tire force at wheel j	[N]
$F_{z,j}$	Normal tire force at wheel j	[N]

F_x	Resultant longitudinal force on the vehicle	[N]
F_y	Resultant lateral force on the vehicle	[N]
$F_{\text{ideal},i}$	Ideal target damper force at corner i	[N]
$F_{\text{target},i}$	Target damper force at corner i	[N]
F_{min}	Minimum damper force	[N]
M_x	Roll moment about the vehicle center of gravity	[N m]
M_y	Pitch moment about the vehicle center of gravity	[N m]
M_z	Yaw moment about the vehicle center of gravity	[N m]
s	State in reinforcement learning	[–]
a_t	Action at time step t in reinforcement learning	[–]
$\mathcal{P}$	State transition probability in the Markov decision process	[–]
R	Reward function in the Markov decision process	[–]
γ	Discount factor in reinforcement learning	[–]
r_t	Reward at time step t	[–]
$Q(s, a)$	State-action value function	[–]
π	Policy function	[–]
I	Damper control current	[A]
I_i	Damper current at corner i	[A]
I_{opt}	Optimal damper current	[A]
I_{grid}	Discrete current search grid	[A]
u_t	Control input at time step t	[–]
P	Instantaneous power	[W]
$\bar{I}$	Mean current	[A]
a_w	Frequency-weighted RMS vertical acceleration	[m/s^2]
$G_q(\Omega)$	Power spectral density of the road surface profile	[m^3/rad]
x	Longitudinal position of the vehicle	[m]
y	Lateral position of the vehicle	[m]

Use of Generative AI

Generative AI tools, including Claude, ChatGPT, and Gemini, were used solely as writing aids during the preparation of this thesis. Specifically, these tools were used to assist with spelling and grammar correction, sentence restructuring, and improving the clarity of written expression. All technical content, analysis, results, and conclusions are the authors' own work. No AI-generated text has been used without review, editing, and verification by the authors.

Contents

1

Introduction

This chapter introduces the thesis topic and motivation, states the problem and research objectives, formulates the research questions, and defines the scope and report outline.

1.1 Background

Heavy-duty trucks operate across a wide range of duty cycles and payload conditions, from empty return trips to fully loaded transport missions. This variability places conflicting demands on the suspension system: the vehicle must provide acceptable ride comfort for the driver while maintaining stable road holding and controllable chassis motions under high load. Because conventional passive suspensions are tuned as a compromise between competing performance requirements, their performance can degrade significantly when operating far from the nominal design load. Semi-active damping offers a practical path to improve this trade-off by adjusting damper behaviour to current operating conditions with limited energy use and complexity compared to fully active systems [1].

1.2 Problem Statement

In passive suspensions systems, the damper characteristics are fixed, so the effective damping level changes with the vehicle's sprung mass. This can lead to excessive oscillations and reduced ride comfort when the vehicle is heavily loaded, while a lightly loaded vehicle may experience unnecessarily stiff damping and reduced vibration isolation. In heavy-duty applications, this load dependence makes it difficult to achieve consistently good comfort across operating conditions [2].

This thesis investigates semi-active suspension control algorithms, where damper resistance is adjusted through actuator current I, and compares their ability to maintain ride comfort across different controllers.

1.3 Aim and Objectives

The aim of this thesis is to evaluate and compare semi-active damper control strategies for a heavy-duty truck across varying payload conditions, with the purpose of identifying which approach offers the most comfort along with a practical path to real-world implementation. The following objectives are pursued:

- Compare the comfort performance of five control strategies: Passive, Load-Adaptive (LA), Skyhook (SH) with estimation, SH with measurements and Reinforcement Learning (RL), across a range of road surfaces and payload conditions.
- Evaluate the energy consumption of each active controller to assess operational cost relative to comfort gain.
- Assess the sensor and hardware requirements of each controller to determine implementation complexity on a production vehicle.
- Identify which control strategy offers the best balance of comfort, power efficiency, and implementation feasibility for deployment in a real truck.

1.4 Research Questions

To address the problem of load-variant comfort, this thesis seeks to answer the following questions:

1. **Load Sensitivity:** How does varying payload affect ride comfort, and can active control strategies maintain consistent comfort performance across load conditions compared to a passive baseline?
2. **Control Performance:** Is there a quantifiable comfort gain from semi-active control strategies relative to passive damping across varying road surfaces and payload conditions?
3. **Energy Consumption:** How much power do the evaluated control strategies consume, and how does this relate to the comfort improvement achieved?
4. **Implementation Feasibility:** Which control strategy offers the most viable balance between comfort performance, power consumption, sensor requirements, computational complexity, and practical feasibility for implementation in a production heavy-duty truck?

1.5 Scope and Limitations

This thesis is conducted as a simulation study using a rigid full-vehicle truck model in IPG TruckMaker. The controllers are developed and evaluated under a set of representative payload conditions and road inputs.

The vehicle model includes an air suspension system provided by the industry partner, Volvo TTI, and implemented in TruckMaker. Full pneumatic dynamics and levelling control are outside the scope of this work, and the air spring behaviour is represented through the parametrised model as given.

Physical testing on a real vehicle is outside the scope of this thesis. Results are therefore limited by the fidelity of the simulation model and its parametrisation, including assumptions in the vehicle, tyre, and road representations. On-vehicle validation is identified as future work.

1.6 Report Outline

The remainder of this report is structured as follows.

Chapter 2 presents the literature review on vehicle suspension and damper modelling with semi-active control, covering classical benchmark methods, learning-based approaches, and practical limitations related to delays and sensing.

Chapter 3 introduces the vehicle dynamics used in this thesis, including the full-car model, suspension dynamics, ride comfort metrics, and road surface representation.

Chapter 4 describes the evaluated control strategies: Reinforcement Learning, Skyhook, Load-Adaptive control and estimated states used for the control algorithms.

Chapter 5 outlines the methodology, including the simulation environment in IPG TruckMaker and MATLAB/Simulink, the semi-active damper model, controller implementation, estimation implementation, and the test framework used for evaluation.

Chapter 6 presents the results in terms of ride comfort, frequency-domain response, energy consumption across the different road classes and load cases, and the performance of the estimator.

Chapter 7 discusses the results with emphasis on performance, implementation feasibility, sensor and estimation validity, and energy consumption.

Chapter 8 summarises the main conclusions and outlines directions for future work.

2

Literature Review

This chapter summarizes the state-of-the-art studies and reports regarding vehicle suspension modelling and control.

2.1 Overview

This chapter surveys the state of the art in vehicle suspension modelling and control, tracing a path from classical physical models and heuristic strategies to modern data-driven approaches. The review is organised around three interconnected themes that motivate this thesis. First, the fidelity of the simulation environment determines whether a control strategy can transfer to a real vehicle. Second, the theoretical performance ceiling for comfort-oriented semi-active control without road preview is well-established, and any novel approach must be evaluated against it. Third, and most critically, the gap between simulation and reality is shaped by signal delays and sensor limitations, constraints that can render otherwise sound control strategies ineffective or even destabilising in practice.

2.2 Vehicle Dynamics and Modelling

The choice of vehicle model fundamentally determines the validity of any conclusions drawn from simulation. The Quarter Car Model (QCM) is the standard for preliminary component sizing and vertical dynamics analysis, yet it carries important limitations. [3] found that QCMs significantly overestimate body acceleration on random roads compared to higher-fidelity models. For accurately predicting passenger comfort or handling asymmetric road inputs such as single-wheel bumps, a 7 Degrees of Fredom (7-DOF) Full Vehicle Model is often indispensable [3].

Despite these limitations, QCMs remain prevalent in RL research due to their computational efficiency. [4] demonstrated that an RL agent trained on a QCM can transfer successfully to a real vehicle, provided that physical parameters such as friction and actuator delay are calibrated to match measured rig data prior to training. This finding underscores a broader principle: the model need not be fully faithful to reality, but the parameters governing the transfer-critical dynamics must.

Accurate modelling of non-linearities is equally important. [5] highlights that incorporating vertical degrees of freedom for the wheels and non-linear damper characteristics is essential for realistic steering feel and road irregularity response, effects that a linearised QCM cannot capture.

2.3 Classical and Optimal Control Benchmarks

Before evaluating control approaches, robust baselines must be established. The semi-active control literature distinguishes between heuristic methods, which are designed from physical intuition, and strategies derived from optimal control theory.

2.3.1 Skyhook and Groundhook

The field is anchored by two complementary heuristics. Skyhook (SH), introduced by [6], minimises sprung mass motion to maximise ride comfort. Groundhook (GH), developed by [7], minimises dynamic tyre loads to maximise road holding. These strategies represent the extreme ends of the fundamental comfort–handling trade-off and serve as primary benchmarks throughout the literature.

2.3.2 Mixed SH-ADD

A limitation of Skyhook is its poor performance near body resonance. [8] showed that minimising vertical body acceleration without road preview requires a frequency-dependent switching strategy, introducing Acceleration-Driven Damping (ADD), which is optimal at high frequencies but degrades near resonance. The *Mixed SH-ADD* algorithm [9] resolves this by using a frequency-range selector to switch between Skyhook at low frequencies and ADD at high frequencies. This strategy is mathematically proven to constitute the quasi-optimal lower bound for comfort-oriented semi-active control without road preview, meaning the best achievable performance under semi-active passivity constraints in the absence of preview information, making it the theoretical ceiling against which any non-preview strategy should be judged.

To solve this, the *Mixed SH-ADD* algorithm was developed [9]. It uses a frequency-range selector to switch between Skyhook (low frequency) and ADD (high frequency).

2.4 Modern Control: Predictive and Learning-Based Approaches

Recent work has shifted from fixed control laws towards methods capable of handling non-linearities, constraints, and uncertainty in a principled way.

2.4.1 Explicit Model Predictive Control

Model Predictive Control (MPC) optimises over a receding prediction horizon and handles physical constraints naturally, but the online computational cost has historically limited its use in real-time suspension control. [10] addressed this by approximating the MPC solution offline using a neural network, while [11] demonstrated real-time Explicit MPC with road preview on a hardware-in-the-loop test rig. Both approaches achieve near-MPC performance at execution speeds compatible with embedded hardware, suggesting that computational cost is no longer the fundamental barrier it once was.

2.4.2 Deep Reinforcement Learning

Deep Reinforcement Learning (DRL) has emerged as a promising framework for continuous suspension control, owing to its ability to learn non-linear policies directly from interaction without requiring an explicit system model [12]. Three findings from the literature are particularly relevant to this thesis.

Ultsch et al. [4] demonstrated real-world transfer of a Soft Actor-Critic agent trained in simulation, achieving approximately 2.5% better comfort than an offline-optimised SH/GH baseline-evidence that the sim-to-real gap is bridgeable under careful parameter matching. Savaia et al. [13] provided insight into *why* RL policies tend to perform well: learned policies frequently converge on an structure resembling augmented Skyhook, activating damping in quadrants where the classical law remains passive. This suggests that classical heuristics are close to locally optimal, and that RL adds value primarily by generalising beyond their fixed switching logic. Finally, [14] showed through Batch RL that competitive policies can be recovered from pre-collected data without any online interaction or physical model, demonstrating that RL is viable even when real-time simulation is impractical.

2.5 The Reality Gap: Time Delays and Sensor Constraints

A recurring finding across the literature is that theoretical performance is difficult to realise in practice. Two factors dominate: signal delays and limited sensor availability. These are not peripheral concerns but central obstacles to deployment, and they directly motivate the approach taken in this thesis.

2.5.1 Influence of Time Delays

Delays introduced by sensor processing and actuator dynamics are a primary cause of performance degradation. Prabhahar & Kamardeen [15] quantified that delays exceeding 40 ms can render semi-active Skyhook worse than a passive suspension, while active Linear Quadratic Regulator control can become unstable with delays beyond 60 ms. Critically, the uncertainty of the delay is often more damaging than its mean value [16]. By training a Twin-Delayed Deep Deterministic Policy Gradient (TD3) agent in an environment with stochastic delays, Wang et al. [16] maintained a comfort improvement exceeding 30% over passive suspension even under uncertain delay conditions, a result that underlines the importance of treating delay as a stochastic quantity during training rather than a fixed parameter.

2.5.2 Sensor Constraints and State Estimation

Standard strategies such as Skyhook and ADD rely on absolute sprung mass velocity or acceleration. Where accelerometers are unavailable or unreliable, these states must be estimated. Paul & Varadharajan [17] demonstrated that Unscented and Cubature Kalman Filters can fuse data from suspension level sensors and IMUs to accurately reconstruct chassis states. An alternative approach applies where pressure sensors are present, such as in air spring systems: by measuring spring force directly and combining it with relative height measurements, vertical dynamics can be reconstructed via Newton's second law, enabling advanced control laws including Mixed SH-ADD without direct acceleration measurements [18].

2.6 Conclusion

The literature establishes a clear foundation for this thesis. Classical semi-active strategies such as Skyhook and Mixed SH-ADD define strong performance benchmarks for comfort-oriented suspension control, with Mixed SH-ADD representing the quasi-optimal limit for non-preview semi-active systems. At the same time, modern learning-based methods show that Reinforcement Learning can approach or exceed classical benchmarks under carefully controlled conditions.

For the present work, however, the key lesson from the literature is not only which controller performs best in theory, but which assumptions determine whether that performance is achievable in practice. In a heavy-duty truck with large payload variation, the controller must remain effective across changing sprung mass, while operating under limited sensing, actuator delay, and semi-active passivity constraints. This makes implementation realism central to the problem formulation.

The review therefore motivates the controller set selected in this thesis. Skyhook is retained as the main comfort-oriented benchmark due to its strong literature support and practical relevance. Reinforcement Learning is included as a modern non-linear alternative with the potential to improve performance beyond fixed heuristics. LA control addresses the specific payload sensitivity problem directly, while offering a low-complexity and implementation-oriented baseline.

Overall, the literature suggests that the practical success of semi-active suspension control will depend less on theoretical peak performance alone, and more on how well the controller handles delays, sensing limitations, and variation in operating conditions.

3

Vehicle Dynamics

This chapter covers the fundamental principles of vehicle dynamics, establishing the mathematical and physical framework needed to understand how a vehicle responds to driver inputs and external forces [19][20]. Central to this are the interactions between the chassis, suspension, and tires, which together govern performance, safety, and ride comfort [21].

3.1 Coordinate System

A standardized coordinate system is used to define the interaction between the vehicle's mass and the external forces acting upon it. This thesis adopts the ISO 8855 convention, which utilizes a right-handed Cartesian coordinate system centered at the vehicle's Center of Mass [22].

3.1.1 Reference Frames

To accurately describe the motion and orientation of the 4x2 truck, two distinct coordinate systems are utilized in accordance with the ISO 8855 standard:

- **Inertial Frame (X, Y, Z):**
  - A fixed, earth-bound global coordinate system.
  - Primarily used to define the vehicle's absolute position and the road profile input z_r.
  - Acts as the stationary reference for calculating absolute displacements and velocities.
- **Body-Fixed Frame (x, y, z):**
  - A local frame that translates and rotates alongside the vehicle's sprung mass.
  - The origin is fixed at the vehicle's COM.
  - x **(Longitudinal):** Directed forward, parallel to the vehicle's longitudinal axis.
  - y **(Lateral):** Directed to the left, perpendicular to the direction of travel.
  - z **(Vertical):** Directed upwards, perpendicular to the xy-plane, tracking heave motion.

3.1.2 Degrees of Freedom (DOF)

A vehicle is a complex multibody system. To analyze the influence of damping across all dimensions, the model must account for several degrees of freedom [1]:

1. **Translational Degrees of Freedom:**
   - **Heave (z):** The vertical displacement of the entire vehicle body.
   - **Surge (x):** Longitudinal displacement, acceleration and braking.
   - **Sway (y):** Lateral displacement, turning the vehicle.
2. **Rotational Degrees of Freedom:**
   - **Roll (ϕ):** Rotation about the longitudinal x-axis. This occurs primarily during cornering due to lateral acceleration.
   - **Pitch (θ):** Rotation about the lateral y-axis. This occurs during braking or acceleration.
   - **Yaw (ψ):** Rotation about the vertical z-axis. This describes the vehicle's heading during steering.

3.2 The Full-Car Model Geometry

While simplified models exist, the comparison of damping controllers across all three domains requires a 7-DOF Full-Car Model [21]. This model treats the chassis as a rigid body with three degrees of freedom and each of the four wheels as independent unsprung masses with one vertical degree of freedom each.

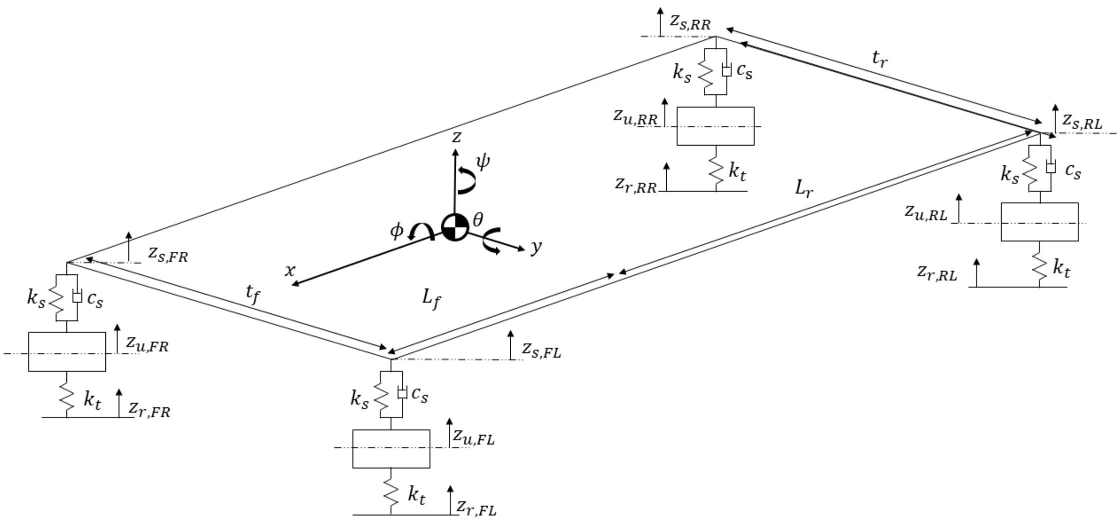

Figure 3.1: 7-DOF Vehicle model

3.2.1 Geometric Relationships

The vertical displacement at each corner of the vehicle, $z_{s,ij}$, is not independent; it is a kinematic function of the body's central heave, pitch, and roll [19]. For a vehicle with wheelbase L, divided into L_f and L_r from the Center of Mass (CoM), and track width t, the vertical position of the sprung mass at the four corners is defined by:

$$z_{s,FL} = z_s - L_f \sin\theta + \frac{t}{2}\sin\phi \tag{3.1}$$

$$z_{s,FR} = z_s - L_f \sin\theta - \frac{t}{2}\sin\phi \tag{3.2}$$

$$z_{s,RL} = z_s + L_r \sin\theta + \frac{t}{2}\sin\phi \tag{3.3}$$

$$z_{s,RR} = z_s + L_r \sin\theta - \frac{t}{2}\sin\phi \tag{3.4}$$

These equations bridge the gap between the global body motion and the local suspension displacement which the dampers control.

3.2.2 Coordinate Transformation

To calculate the velocities required for damping force calculations, we differentiate the corner displacements:

$$\dot{z}_{s,w} = \dot{z}_s - L_a\dot{\theta}\cos\theta + \frac{t}{2}\dot{\phi}\cos\phi \tag{3.5}$$

This kinematic foundation ensures that any damping force applied at a single wheel is correctly translated into a moment that affects the pitch and roll of the entire vehicle body.

3.2.3 Longitudinal

Longitudinal dynamics focus on the forces generated during acceleration and braking [20]. The central phenomenon here is load transfer. When a vehicle decelerates, the inertial force causes a weight shift from the rear axle to the front axle. This manifests as pitch, θ, where the front of the vehicle dives. The damping system must manage the rate of this pitch to maintain stability and prevent the front suspension from hitting its travel limits.

3.2.4 Lateral

Lateral dynamics involve the vehicle's response to steering commands and cornering forces [23]. As the vehicle turns, centrifugal force acts on the center of gravity, causing the body to roll toward the outside of the curve. This creates lateral load transfer between the inner and outer wheels. Damping is crucial during the transient phase as it determines how quickly the tires gain grip and how stable the vehicle feels to the driver.

3.2.5 Vertical

Vertical dynamics address the isolation of the vehicle body from road irregularities [24]. The primary goal is to minimize the energy transmitted from the road profile,

z_r, to the passenger cabin, z_s, to ensure ride comfort. The motion of a quarter-car model is governed by the following coupled differential equations:

$$m_s \ddot{z}_s + F_d + k_s(z_s - z_u) = 0 \tag{3.6}$$

$$m_u \ddot{z}_u - F_d - k_s(z_s - z_u) + k_t(z_u - z_r) = 0 \tag{3.7}$$

3.3 Equations of Motion

Applying Newton's second law to the rigid sprung mass and each of the four unsprung masses yields the seven coupled equations of motion that govern the system [19]. The sprung mass has three degrees of freedom, heave, pitch, and roll, while each wheel assembly contributes one vertical degree of freedom.

3.3.1 Sprung Mass

The net suspension force at each corner, $F_{s,w}$, acts as the coupling term between the body and the wheel assemblies:

$$F_{s,w} = F_{d,w} + k_{s,w}(z_{s,w} - z_{u,w}) \tag{3.8}$$

where $F_{d,w}$ is the damper force and $k_{s,w}$ is the spring stiffness at wheel $w \in \{\text{FL, FR, RL, RR}\}$. Summing these corner forces over the entire body gives the heave equation:

$$m_s \ddot{z}_s = -\sum_w F_{s,w} \tag{3.9}$$

The pitch equation balances the moment about the lateral (y) axis:

$$I_{yy}\ddot{\theta} = \sum_a \sum_{w \in a} F_{s,w} \cdot L_a \tag{3.10}$$

where $L_a = -L_f$ for the front axle ($a = f$) and $L_a = L_r$ for the rear axle ($a = r$). The roll equation balances the moment about the longitudinal (x) axis:

$$I_{xx}\ddot{\phi} = \sum_w F_{s,w} \cdot d_w \tag{3.11}$$

where $d_w = +t/2$ for left wheels $w \in \{\text{FL, RL}\}$ and $d_w = -t/2$ for right wheels $w \in \{\text{FR, RR}\}$.

3.3.2 Unsprung Masses

Each wheel assembly is treated as an independent mass $m_{u,w}$ suspended between the spring-damper unit above and the tyre stiffness $k_{t,w}$ below:

$$m_{u,w}\ddot{z}_{u,w} = F_{s,w} - k_{t,w}(z_{u,w} - z_{r,w}) \tag{3.12}$$

where $z_{r,w}$ is the road profile input at wheel w.

3.4 Damping

Damping is the mechanism of energy dissipation within the suspension system. Its primary role is to control the oscillations of the springs k_s and k_t, preventing the vehicle from bouncing indefinitely after a disturbance.

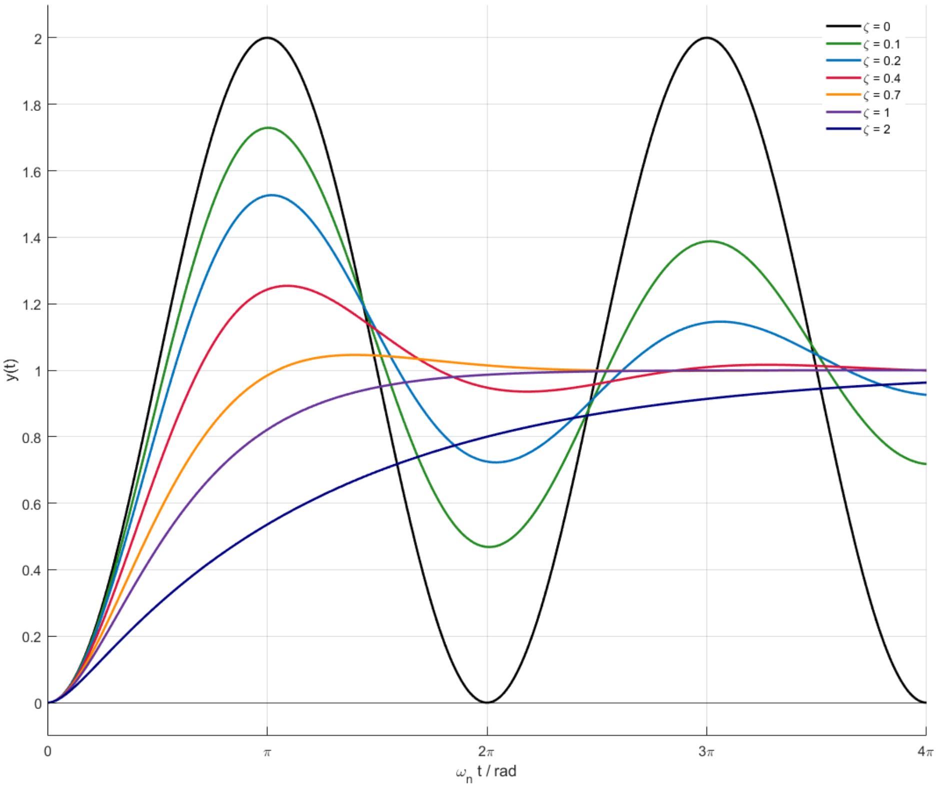

Figure 3.2: Step Response for Varying Damping Ratios ζ

3.4.1 Passive Damping

In a passive suspension, the damping coefficient c is fixed. This creates a fundamental design conflict: a soft damper provides excellent comfort but poor handling, while a stiff damper provides sharp handling but a harsh ride [21]. The damping ratio ζ is defined as:

$$\zeta = \frac{c}{2\sqrt{km}} \tag{3.13}$$

where k is the spring stiffness and m is the sprung mass. This ratio dictates the system's response to disturbances, as illustrated in Figure 3.2. An underdamped response ($\zeta < 1$) occurs when the damping force is insufficient to dissipate energy quickly, causing persistent oscillations. As the sprung mass m increases under load, ζ decreases, leading to degraded handling. Conversely, an overdamped response ($\zeta > 1$) prevents the suspension from absorbing impacts, resulting in poor driver comfort. Active suspensions can address these issues but are often impractical in heavy-duty applications due to energy consumption and complexity. Semi-active dampers offer a compromise by modulating the damper force characteristic via an input current I.

3.4.2 Semi-Active Damping

Semi-active damping allows the damping coefficient c to be adjusted in real-time. Semi-active dampers used for vehicle systems cannot add energy to the system, they can only vary the rate of dissipation. This is defined by the passivity constraint, which dictates that the damping force must always oppose the direction of the suspension's relative velocity.

$$F_d = c_{variable} \cdot (\dot{z}_s - \dot{z}_u) \tag{3.14}$$

$$P = F_d \cdot (\dot{z}_s - \dot{z}_u) \geq 0 \tag{3.15}$$

Passive dampers on heavy vehicles with high payload variances are inherently compromised [7]. Because passive damping characteristics are fixed, the suspension must be tuned to guarantee ride safety at maximum Gross Vehicle Mass [5]. Consequently, the system becomes over-damped during unladen states, significantly degrading vibration isolation and driver comfort [1].

3.5 Ride Comfort Metric

Human sensitivity to whole-body vibration is frequency-dependent and varies significantly across the vibration spectrum [25]. The ISO 2631-1 standard [26] defines methods for quantifying this exposure, specifying a frequency weighting filter W_k applied to vertical vibration that reflects human sensitivity across frequency. The standard identifies the 4–8 Hz range as the region of highest sensitivity, which overlaps directly with the sprung mass resonance of a heavy truck. Prolonged exposure in this range is associated with increased driver fatigue and long-term musculoskeletal health risks. The comfort metric used in this thesis is the frequency-weighted Root Mean Squared (RMS) acceleration:

$$a_w = \sqrt{\frac{1}{T}\int_0^T \left[W_k(f) \cdot a(t)\right]^2 dt} \tag{3.16}$$

where $a(t)$ is the measured vertical acceleration and $W_k(f)$ is the ISO 2631-1 frequency weighting filter for vertical vibration.

3.5.1 Ride Frequency

A vehicle suspension system exhibits two distinct natural frequencies. Using the QCM as an analytical basis, the sprung mass natural frequency is given by:

$$\omega_s = \sqrt{\frac{k_s}{m_s}} \tag{3.17}$$

For a heavy truck this typically falls in the range of 1–2 Hz, overlapping with the frequency band to which humans are most sensitive [26], making it the primary target for comfort-oriented control. The unsprung mass natural frequency is given by:

$$\omega_u = \sqrt{\frac{k_s + k_t}{m_u}} \tag{3.18}$$

This typically falls in the range of 10–15 Hz [19]. The condition $k_t \gg k_s$ ensures a clear separation between the two resonances. In this thesis, the primary focus is on attenuating sprung mass motion near ω_s, as this directly governs driver comfort.

3.6 Road Surface Profiles

Road surface profiles are classified by ISO 8608:2016 [27] into categories A through H based on their Power Spectral Density (PSD), defined as:

$$G_q(\Omega) = G_{q0} \left(\frac{\Omega}{\Omega_0} \right)^{-w} \tag{3.19}$$

where G_{q0} is the roughness coefficient, Ω is the spatial frequency, $\Omega_0 = 2\pi \cdot 0.1\,\text{rad/m}$ is the reference spatial frequency, and $w = 2$ is the waviness exponent. Each class represents a doubling of roughness relative to the previous, as summarised in Table 3.1.

Table 3.1: ISO 8608:2016 road surface classification

Class	G_{q0} (m^3)	IRI (m/km)
A	16×10^{-6}	0.75
B	64×10^{-6}	1.5
C	256×10^{-6}	3.5
D	1024×10^{-6}	7.0
E	4096×10^{-6}	14.0
F	16384×10^{-6}	28.0

Classes A and B are representative of well-maintained paved roads, while Classes C and D correspond to typical operating conditions for heavy-duty trucks [24]. Classes E and F represent severely damaged surfaces and serve as robustness benchmarks [27]. Since ISO 8608 classifies road roughness in terms of spatial frequency, the temporal frequency experienced by the vehicle depends on travel speed:

$$f = \frac{v}{\lambda} \tag{3.20}$$

3.6.1 Frequency Weighting

The W_k weighting factors defined by ISO 2631-1 [26] reflect the frequency dependence of human sensitivity to vertical whole-body vibration. As shown in Table 5.6, the weights increase from 1 Hz, peak in the 5–6.3 Hz range, and decrease towards

12.5 Hz [19]. This indicates that vibrations in the 4–8 Hz band are perceived as most uncomfortable, which directly overlaps with the sprung mass resonance of a heavy truck. A controller that attenuates energy in this band will therefore produce a disproportionately larger improvement in weighted RMS acceleration than one that targets higher or lower frequencies.

Table 3.2: ISO frequencies and W_k weighting factors [19]

Frequency (Hz)	W_k **weight**
1.0	482
2.0	513
2.5	631
3.15	804
4.0	976
5.0	1039
6.3	1054
8.0	1036
10.0	988
12.5	902

4

Control Algorithms

This chapter proposes different control algorithms that based on Chapter 2, yielded good results for increasing comfort in a heavy duty truck. Table 4.1 provides an overview of the explored control algorithms and their respective purposes.

Table 4.1: Overview of explored control algorithms

Control Algorithm	Purpose	Justification
Reinforcement Learning (TD3)	Data-driven control using neural networks	Demonstrated good performance in previous studies on semi-active damper control
Skyhook	Model-based control using available sensor measurements	Widely used benchmark for semi-active suspension comfort; proven performance in literature
Load-Adaptive	Cost-effective rule-based control	Simple implementation requiring minimal sensors; scales passive damper characteristics based on load

In the following sections, each control strategy is explained in detail.

4.1 Reinforcement Learning

Consider a self walking robot, learning to navigate a maze with the goal to reach a certain position. The robot gains no previous information about the maze or how to walk, thus it must learn to walk and navigate through trial and error. The robot can only gain sparse feedback through a series of positive signals when it successfully walks or navigates to the goal position. Furthermore, it will receive negative signals when it e.g. falls over, or walks into a dead end. This is the core idea of reinforcement learning, to learn sequential decisions based on previous interactions with the environment that maximizes the outcomes.

The reinforcement learning problem can be traced back to optimal control theory along with the Markov Decision Process framework (MDP). An MDP is formally defined as a tuple (s, a, P, R, γ), where s is the state space, a is the action space, $P : s \times a \times s \to [0, 1]$ is the state transition probability function, $R : s \times a \to \mathbb{R}$ is the reward function, and $\gamma \in [0, 1]$ is the discount faction for future rewards.

An MDP consists of an agent interacting with an environment at all time steps, making a decision at all possible moments. The MDP follows the workflow:

- Observes the state s_t at time t.
- An agent selects an action a_t at time t.
- Receives a reward r_{t+1}.
- Transitions to the next state s_{t+1}, based on the action.

This is the basic idea of Reinforcement Learning, an illustration of the RL framework can be seen in Figure 4.1.

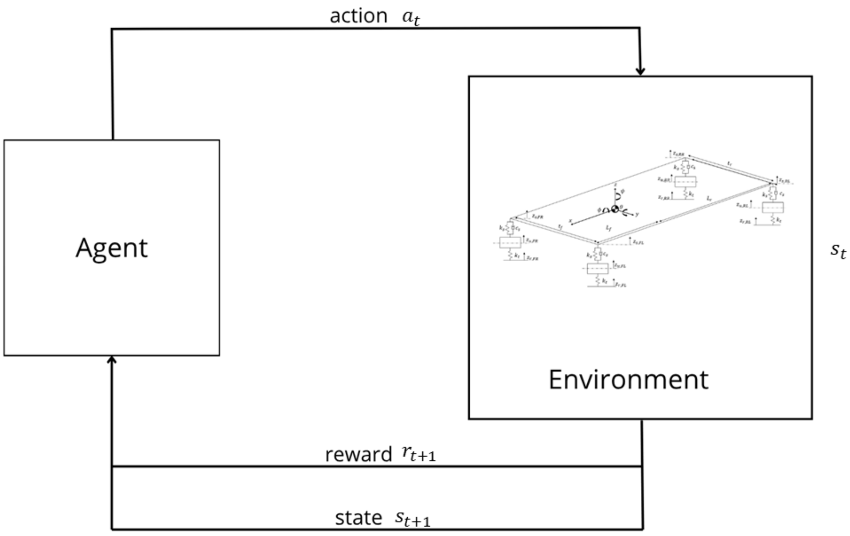

Figure 4.1: RL framework overview

The objective of the agent is to maximize the cumulative reward over an entire episode, thus the agent will learn to choose actions that create a certain behaviour given a certain policy. A policy $\pi : S \to A$ defines the agent's behavior by mapping states to actions. In deterministic policies, denoted $\mu(s)$, each state maps to a single action. The objective is to find the optimal policy π^* that maximizes the expected cumulative discounted reward over time. The agent in traditional RL schemes is a neural network [28] that consists of many connected processors. The processors or neurons are aligned in layers that form a network where layers are connected to each other. The neurons get activated through weighted connections to other active neurons in previous layers via backpropagation. The inputs are the sensor inputs from a system and the outputs are the actions that control the environment or system. For a comprehensive introduction to reinforcement learning theory and algorithms, readers are referred to [29].

4.1.1 Deep Deterministic Policy Gradient

The RL algorithm in general learns by estimating the Q-value that the model will take at a certain state, where the Q-value is the measure of how well an agent fulfils the reward function at a given state. However, learning and computing the maximum is not an easy task for a system that is complex in the continuous action space.

Deep Deterministic Policy Gradient (DDPG) [30] addresses the complexity of learning in a continuous environment by combining actor-critic methods with techniques from Deep Q-Networks. The actor-critic architecture consists of two components: an actor network that learns the policy $\mu(s)$ to directly output actions, and a critic network that learns a Q-function $Q(s, a)$ to evaluate the quality of state-action pairs. This separation allows the algorithm to handle continuous action spaces efficiently, where the actor proposes actions and the critic provides feedback on their quality.

The idea is to learn a deterministic policy to output actions directly, while simultaneously learning and using a Q-function that evaluates the performance of the actions. The agent will learn the Q-function by using the Bellman equation

$$V^*(s) = \max_a(R(s,a) + \gamma \sum_{s\acute{t}} P(s,a,s')V(s') \tag{4.1}$$

where where $V^*(s)$ is the optimal value function under a given policy, $R(s, a)$ is the reward given a state and an action, γ is the discounting factor for future rewards, $P(s, a, s')$ is the transition probability from state s to s' and $V(s')$ is the value function of state s' under the policy.

4.1.2 TD3 Algorithm

This project makes use of the TD3 algorithm [31]. The algorithm addresses overestimation problems that are consistent in the DDPG algorithm by introducing three key improvements. First, TD3 uses two separate critic networks (twin Q-functions) and takes the minimum of their estimates when computing the target value, which reduces overestimation bias through clipped double-Q learning. Second, the algorithm delays policy updates by updating the actor network less frequently than the critic networks, allowing the value estimates to stabilize before adjusting the policy. Third, TD3 applies target policy smoothing by adding noise to the target actions during training, which helps prevent the policy from exploiting errors in the value function.

The architecture of TD3 is illustrated in Figure 4.2, showing both the local and target networks. The local network consists of one actor network that outputs actions given the current state s, and two critic networks that independently estimate Q-values $Q1$ and $Q2$ for the state-action pair. The twin critics compute separate TD-errors which are used to update their respective networks. The target networks maintain slowly-updated copies of the actor and critics, providing stable target values during training. For detailed implementation and theoretical analysis of the

TD3 algorithm, readers are referred to the original paper [31].

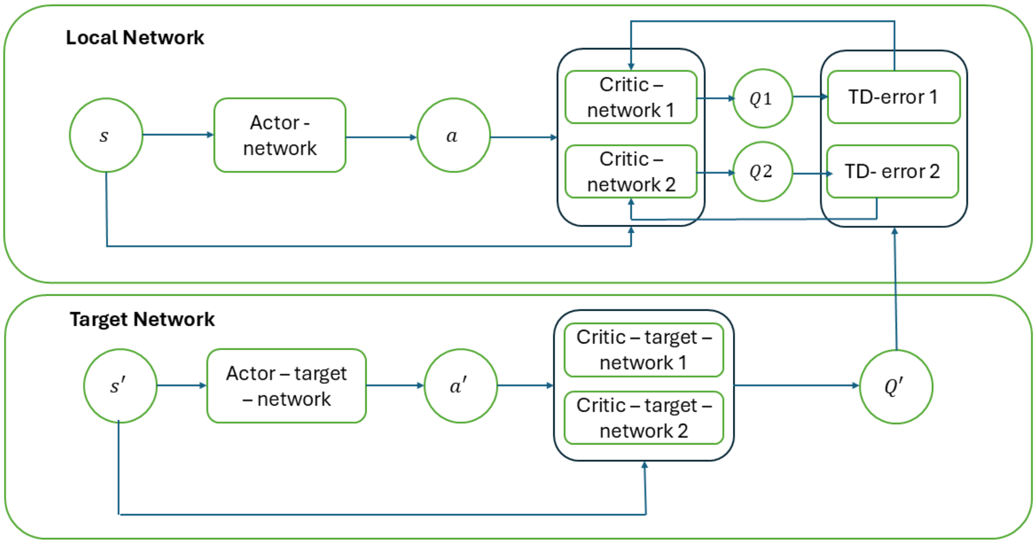

Figure 4.2: TD3 architecture overview showing local and target networks

4.2 Skyhook Control

The Skyhook damping control strategy, first proposed by Karnopp et al. [32], represents one of the most influential semi-active control algorithms in vibration isolation [33]. The fundamental concept behind Skyhook control is to emulate an ideal damper connected between the sprung mass and an inertial reference frame, which is physically impossible to implement with passive components but can be approximated using semi-active devices [34][6].

4.2.1 Control Principle

The Skyhook control strategy is based on the hypothetical system shown in Figure 4.3, where a damper c_{sky} is connected between the sprung mass and an inertial reference frame. This configuration would provide optimal damping for the sprung mass without being influenced by base excitation. While such a configuration cannot be physically realized, a semi-active damper can approximate this behaviour by modulating its damping coefficient based on the relative motion between the sprung mass velocity and the base [35].

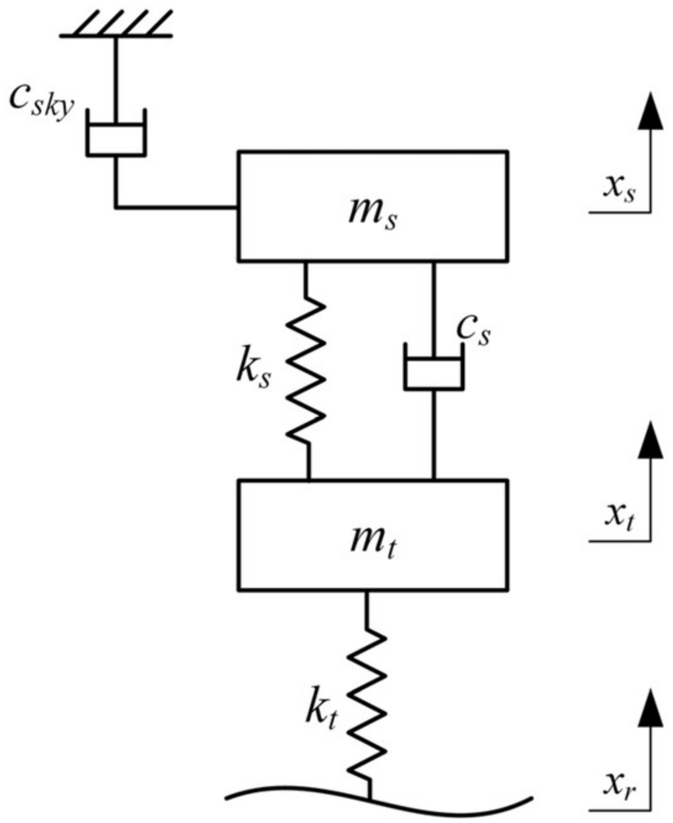

Figure 4.3: Conceptual representation of Skyhook damping

The ideal Skyhook damping force is expressed as:

$$F_{sky} = -c_{sky}\dot{z}_s \tag{4.2}$$

where c_{sky} is the desired Skyhook damping coefficient and $\dot{z}_s$ is the absolute velocity of the sprung mass. Since c_{sky} is load dependent it is hard to define for each load case. Thus the ideal Skyhook damping force can be redefined as:

$$F_{sky} = -m\ddot{z}_s \tag{4.3}$$

4.2.2 Control Law Implementation

Since the ideal Skyhook force cannot be directly implemented, the semi-active damper force is governed by:

$$F_d = -c(t)(\dot{z}_s - \dot{z}_u) \tag{4.4}$$

where $c(t)$ is the time-varying damping coefficient, $\dot{z}_s$ is the sprung mass velocity, and $\dot{z}_u$ is the unsprung mass velocity.
The Skyhook control law determines the damping coefficient based on the following logic [32]:

$$c(t) = \begin{cases} c_{sky} & \text{if } \dot{z}_s(\dot{z}_s - \dot{z}_u) > 0 \\ c_{min} & \text{otherwise} \end{cases} \tag{4.5}$$

This control law can be understood through the sign condition:

- When $(\dot{z}_s - \dot{z}_u) > 0$: The sprung mass velocity and relative velocity have the same sign, meaning the damper would dissipate energy from the sprung mass motion. The damping is set to maximum (c_{max}) to approximate the ideal Skyhook behaviour.
- When $(\dot{z}_s - \dot{z}_u) \leq 0$: The damper would add energy to the sprung mass or work against the desired damping effect. The damping is set to minimum (c_{min}) to minimise this adverse effect.

An alternative formulation expresses the control law in terms of the desired Skyhook damping coefficient:

$$c(t) = \begin{cases} \min\left(c_{sky}, c_{max}\right) & \text{if } \dot{z}_s(\dot{z}_s - \dot{z}_u) > 0 \\ c_{min} & \text{otherwise} \end{cases} \tag{4.6}$$

4.3 Load-Adaptive Controller

The Load-Adaptive controller adjusts damper current based on static payload, with the objective of maintaining a consistent effective damping ratio ζ regardless of vehicle load. Unlike dynamic controllers that respond to road inputs, the Load-Adaptive controller operates exclusively at standstill, computing a set-and-hold current command before the drive cycle begins.

4.3.1 Control Principle

The controller is based on the observation that a passive damper tuned for a specific load condition will be over- or under-damped at any other load, since the damping ratio depends directly on the sprung mass:

$$\zeta = \frac{c_{eq}}{2\, m_{corner}\, \omega_n} \tag{4.7}$$

By measuring the current static load and solving for the damper current that restores the target ζ, the controller normalises vehicle behaviour across the full payload range.

4.3.2 Equivalent Damping via Weighted Least Squares

Since damper force-velocity characteristics are non-linear, a single representative equivalent damping coefficient c_{eq} is derived by applying a Weighted Least Squares fit to the linear model $F = c \cdot v$ over a comfort-relevant velocity band:

$$c_{eq} = \frac{\sum_i w_i F_i v_i}{\sum_i w_i v_i^2} \tag{4.8}$$

where F_i and v_i are force-velocity pairs from the damper characterisation curve and w_i are weighting factors.

4.3.3 Target Damping and Current Selection

Given the measured corner sprung mass m_{corner} and a reference damping ratio ζ_{ref} established at a known load condition, the target equivalent damping coefficient is:

$$c_{target} = 2\,\zeta_{ref}\,m_{corner}\,\omega_n \tag{4.9}$$

Since the damper characteristic is non-linear, the current corresponding to c_{target} cannot be computed analytically. Instead, an inverse search is performed over a discrete current grid, selecting the current that minimises the error between the achieved and target damping:

$$I_{opt} = \arg\min_{I \in I_{grid}} |c_{eq}(I) - c_{target}| \tag{4.10}$$

4.4 State Estimation Theory

This section presents the state estimations needed for the aforementioned control algorithms.

4.4.1 The State Estimation Problem

State estimations are the process of estimating internal states from noisy and indirect partial measurements [36]. Vehicle suspension systems often do not contain direct measurements of vertical velocities and accelerations. For this thesis, the internal states must be estimated using measurements of:

- Suspension and damper forces
- Relative height

These measurements are further explained in Section 5.2. These types of measurements have complementary strengths and weaknesses:

- Force measurements are accurate at high frequencies, but require accurate mass estimation to estimate accelerations.
- Relative height measurements are accurate but require differentiation which amplifies noise.

The state estimation problem is used to optimally combine the measurements to compute a more accurate estimation compared to any single source [1].

4.4.2 Complementary Filtering

By using two different measurements with complimentary frequency characteristics, the state estimation problem can be solved [37]. Complementary filtering is an efficient approach to sensor fusion where:

- y_1 is accurate at low frequencies and noisy at high frequencies.
- y_2 is accurate at high frequencies and noisy at low frequencies.

A complementary filter combines the measurements as:

$$\hat{y} = H_{LP}(s) \cdot y_1 + H_{HP}(s) \cdot y_2 \tag{4.11}$$

where $H_{LP}(s)$ is a low-pass filter and $H_{HP}(s)$ is a high-pass filter satisfying

$$H_{LP}(s) + H_{HP}(s) = 1 \qquad \forall s \tag{4.12}$$

This complementary property, ensures that the frequency response is of the same magnitude across all frequencies [38]. The filters are defined for first orders as:

$$H_{LP}(s) = \frac{1}{1+\tau s}, \quad H_{HP}(s) = \frac{\tau s}{1+\tau s} \tag{4.13}$$

where τ is the filter time constant and the crossover frequency is defined as $\omega_c = 1/\tau$ [36].

4.4.3 Numerical Differentiation

Differentiating a measurement amplifies noise for high-frequencies. This can be shown for a continuous time sinusoidal component at frequency ω.

$$f(t) = Asin(\omega t) \Rightarrow \frac{df}{dt} = A\omega sin(\omega t) \tag{4.14}$$

The differentiation amplitudes the noise linearly with increasing frequency, making direct implementation of a differentiator impractical [39].

To combat the problem of amplified noise, the states are estimated using a Tustin differentiator with a filter. The Tustin approach uses the trapezoidal rule to approximate the derivative:

$$s \approx \frac{2}{T_s} \cdot \frac{z-1}{z+1} \tag{4.15}$$

This method provides a good frequency domain matching to the the continuous time derivative, particularly in phase response [40]. However, the method still provides unlimited high frequency gain, which makes it necessary to combine with a low-pass filter.

4.4.4 Digital Filter Design

Two similar filters that correctly filters the noises present in the measurements are presented.

Exact Discritization: The continuous time first order low pass filter:

$$\frac{1}{1+\tau s} \tag{4.16}$$

can be discretized by sampling the impulse response, which yields the discrete time difference equation:

$$y[k] = a \cdot y[k-1] + b \cdot x[k] \tag{4.17}$$

where:

$$a = e^{-T_s/\tau}, \quad b = 1 - e^{-T_s/\tau} \tag{4.18}$$

This method preserves all the continuous time frequency responses below the Nyquist frequency [41, 39]

Bilinear Transform Approximation: For small ratios T_s/τ, the filter coefficient can be approximated using the bilinear approach:

$$\alpha \approx \frac{T_s}{\tau + T_s} \tag{4.19}$$

The difference equation can be defined as:

$$y[k] = (1 - \alpha) \cdot y[k-1] + \alpha \cdot x[k] \tag{4.20}$$

For the parameter ranges used in this work ($T_s = 0.01$ s, $\tau = 0.032 - 0.64$ s), the bilinear approximation introduces less than 1 magnitude error at frequencies below f_c. The bilinear approximation is used in the sprung mass estimator for its computational simplicity [41].

The complementary filter pair for the bilinear approach is defined as:

$$\alpha_{LP} = \frac{T_s}{\tau + T_s}, \quad \alpha_{HP} = \frac{\tau}{\tau + T_s} \tag{4.21}$$

Note that $\alpha_{LP} + \alpha_{HP} = 1$ perseves the complimentary property [37].

5
Methodology

This chapter presents the description of the workflow and applied theory along with test metrics for evaluating proposed control algorithms.

5.1 Simulation Environment

This thesis makes use of IPG TruckMaker integrated with MATLAB/Simulink as the primary simulation platform.

5.1.1 IPG TruckMaker

IPG TruckMaker is an industry-standard tool for vehicle dynamics simulation, widely used for virtual prototyping in automotive development. It models the full multi-body dynamics of the vehicle, simulating the physical force interactions between every component including bushings, air springs, and tyre contact patches. A rigid 4x2 Volvo truck model is used as the plant, thus the entire sprung mass is treated as a solid block. Aerodynamics and powertrain are kept constant across all simulations to isolate the effect of the damping controllers.

5.1.2 Simulink Integration

TruckMaker connects to Simulink via its tm4sl interface, which allows custom control logic to override the damper force signals computed by TruckMaker's internal model. The Simulink model follows the workflow shown in Figure 5.1:

- Read sensor data from TruckMaker's signal dictionary.
- Estimate states needed for the active controllers.
- Compute an optimal current for the selected control scheme.
- Convert current and damper velocity to damper force via a 2D Lookup Tables (LUT).
- Return the computed force to the tm4sl force bus for TruckMaker's backend.

5.2 Modelling Assumptions

The chassis is modelled as a rigid body and frame torsion is not considered. Sensor signals are emulated by TruckMaker's virtual signal dictionary, providing ideal measurements at the defined sensor locations. Actuator dynamics are represented by a 10 ms delay on the damper current command, consistent with the response time of

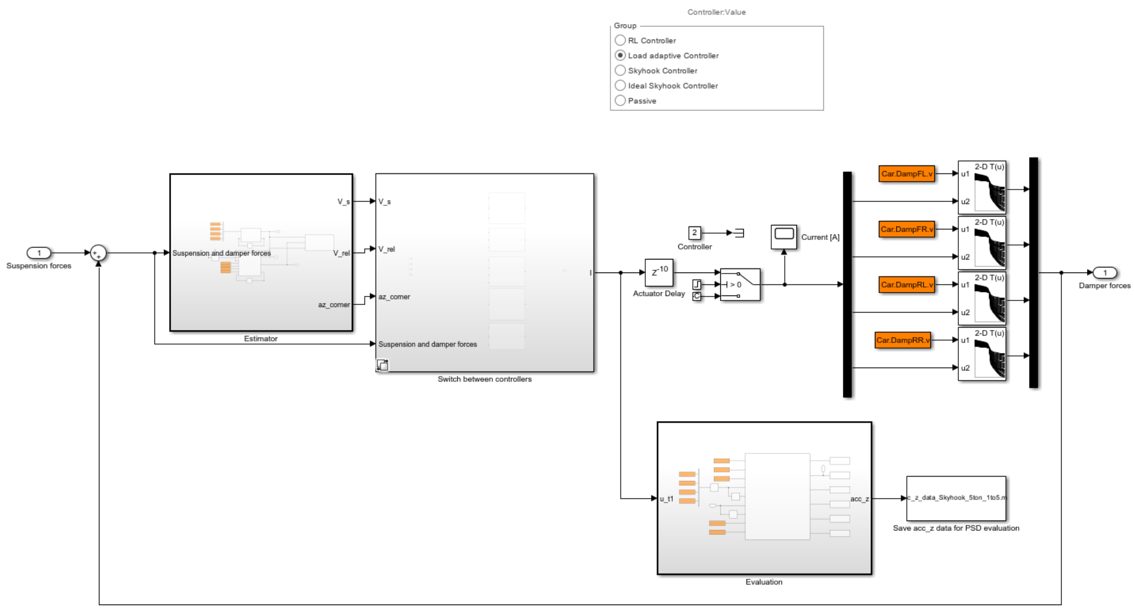

Figure 5.1: Simulink model

the physical hardware as summarised in Table 5.1. The geometric parameters of the vehicle used throughout this work are listed in Table 5.2. Note that t_f and t_r denote the half track widths, i.e. the lateral distance from the centreline to each wheel.

Table 5.1: Available System Signals

Signal	Symbol	Source	Delay
Bellow Pressure	P_{bellow}	Air spring pressure sensor	10 ms
Relative Height	z_{rel}	Linear height sensor	10 ms

Table 5.2: Vehicle geometry parameters

Parameter	Symbol	Value	Unit
Wheelbase	L	4.6	m
Front half-track	t_f	1.05	m
Rear half-track	t_r	1.05	m
Gravitational acc.	g	9.82	m/s^2
Sampling period	T_s	0.01	s

5.3 System Implementation

This section presents the implementation of physical components into the simulation environment.

5.3.1 Semi-Active Damper Model

The semi-active damper behaviour is captured using 2D LUTs indexed by damper velocity and control current. The effective control range is 0.6–1.9 A for both axles. The data sheets in their entirety can be seen in Appendix B.

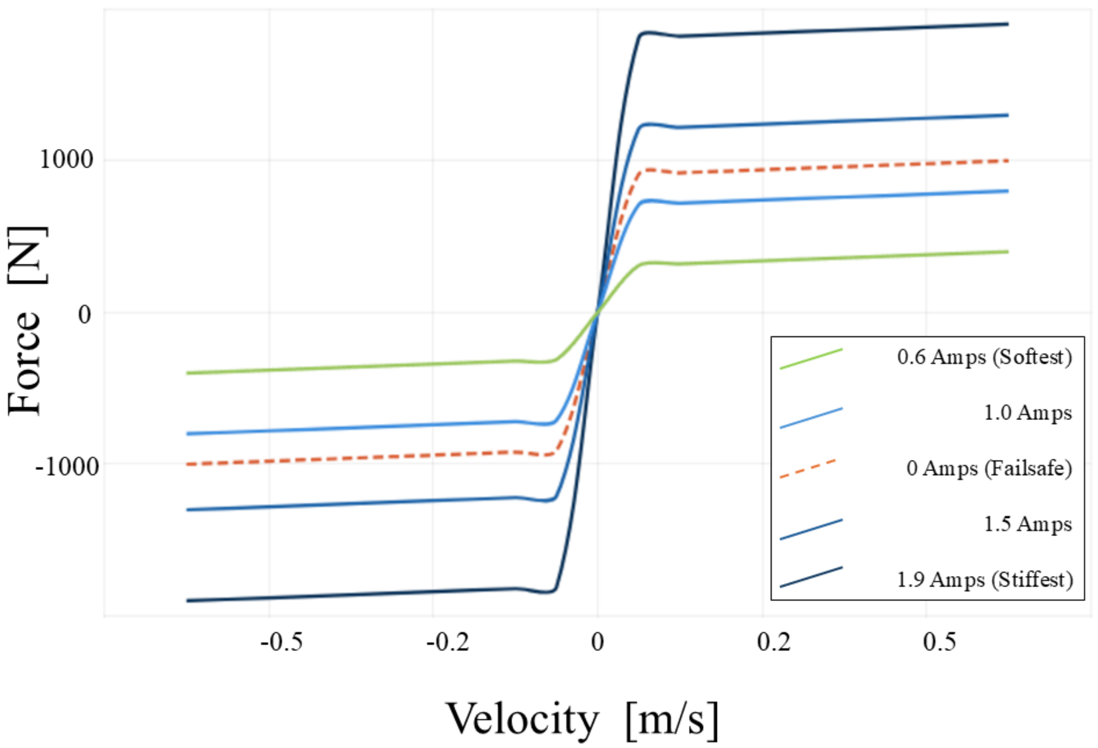

Figure 5.2: Illustrative damper force–velocity characteristic

FIGURES REDACTED DUE TO CONFIDENTIALITY

Figure 5.3: Front damper LUT

FIGURES REDACTED DUE TO CONFIDENTIALITY

Figure 5.4: Rear damper LUT

FIGURES REDACTED DUE TO CONFIDENTIALITY

Figure 5.5: Velocity-force plots with baselines, minimum, and maximum values

The exact lookup tables are redacted due to confidentiality, a generic illustration is shown in Figure 5.2. The damper force characteristics are captured in 2D LUT:s indexed by damper velocity and control current, shown in Figures 5.3 and 5.4. The front and rear dampers share the same hardware but differ in their failsafe behaviour, the front damper defaults to a soft baseline while the rear defaults to a stiffer baseline to support the higher rear axle loads. Figure 5.5 shows the achievable force envelope between the minimum and maximum curves for both axles. Figures 5.6 and 5.7 show the force-current relationship at a fixed velocity for compression and rebound respectively, illustrating the controllable range available to each controller within the 0.6–1.9 A operating band.

FIGURES REDACTED DUE TO CONFIDENTIALITY

Figure 5.6: Current-force in compression for baselines, minimum, and maximum values

FIGURES REDACTED DUE TO CONFIDENTIALITY

Figure 5.7: Current-force in rebound for baselines, minimum, and maximum values

5.3.2 Damper Model Verification

A Simulink assertion monitors the product of damper force and relative velocity throughout each simulation. A positive product indicates the damper is injecting energy into the system, violating the passivity constraint. No such violations were observed during the evaluation runs.

5.4 Control Architecture

This section describes the implementation of each control strategy evaluated in this thesis. The system overview is shown in Figure 5.8.

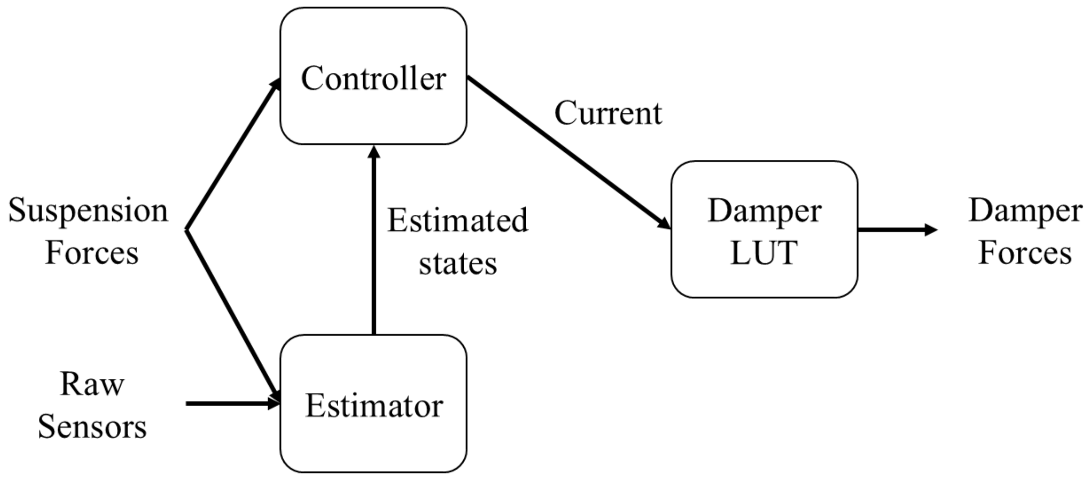

Figure 5.8: System overview

Table 5.3: Controllers used for evaluation

Controller	Purpose
Reinforcement Learning	Policy trained to minimize a reward function for optimal ride comfort. The purpose is to assess the potential of a model-free, data-driven approach.
Skyhook	SH control using sprung mass velocity estimations. Represents the practically deployable version of SH algorithm.
Load Adaptive	Adjusts damping based on estimated static corner load. Represents a simple, deployable benchmark that accounts for varying vehicle loading without dynamic sensing.
Ideal Skyhook	SH control with direct access to true sprung mass velocity. Serves as a theoretical performance ceiling, as it assumes perfect state information unavailable in practice.
Passive	Fixed damping coefficient with no active control. Acts as the baseline performance floor against which all other controllers are evaluated.

5.4.1 Reinforcement Learning Development

This thesis makes use of the reinforcement learning toolbox in MATLAB. The algorithm used is as mentioned in Chapter 4, the TD3 model. To establish the RL scheme, a number of parameters must be designed and integrated first:

- Observation space - which system states the agent can observe
- Action space - what the agent controls
- Reward function - how the agent evaluates performance
- Network architecture - size and structure of actor and critic networks
- Training configuration - hyperparameters for the learning process

The design of these parameters requires a thorough understanding of the suspension system dynamics and the control objectives. The following detail each design choice:

Observation space: The observation space defines what information is sent to the agent. It contains the necessary states for selecting an optimal action.

The selected observation for the TD3 agent are:

- Estimated sprung mass vertical acceleration at each corner $[a_{z,corner}]$. - To know which force the damper should apply, the vertical acceleration of the sprung mass must be observed since force is calculated using acceleration.
- Estimated relative vertical velocity between sprung and unsprung masses at each corner $[v_{rel,z,corner}]$. - It is necessary for the agent to see if the sprung and unsprung masses is moving towards or away from each other since the agent need to know how the unsprung masses move. Since the sensors needed to estimate the movement of the unsprung masses are not available, the relative vertical velocity will describe that motion to the agent.
- Longitudinal velocity $[v]$. - The longitudinal velocity is given to the agent to learn that hitting bumps or uneven road at higher speeds creates larger effects on the vertical acceleration.
- Static force at each corner measured when the truck has no movement $[F_{static,corner}]$. - This is to give the agent an observation of the current load at each corner.
- An observation buffer of 3 time steps prior. - To enable the agent to correlate previous observations to current observation.

The TD3 agent observes a 52-dimensional state vector containing the aforementioned states. The high-dimensional observation space allows trhe agent to capture the complex dynamics of the heavy-duty truck suspension system and make informed control decisions based ion the current driving conditions. Furthermore, the observations are normalized to match the interval of approximately [0, 1], which enables the agent to more clearly identify changes.

Action space: The action space defines what the agent controls. In this application, the action space defines a 2-dimensional continuous vector representing the current which is sent to the front and rear dampers $[0.6A - 1.9A]$. This is to simplify the training process and ensure the truck is stable laterally.

Reward function design: The reward function is critical as it defines the optimization objective for the agent. This thesis focuses on ride comfort, which means the primary goal is to minimize sprung mass vertical acceleration and velocity while maintaining acceptable suspension travel. Furthermore, switching the current in large amplitudes could potentially damage the actuators, thus this must optimally be kept low between time steps. The reward function is formulated as:

$$r = -(w_{acc}a_z^2 + w_{vel}v_z^2 + w_{heave}z_{def}^2 + w_{ctrl}(u_t - u_{t-1})^2) \tag{5.1}$$

where a_z is the sprung mass acceleration, v_z is the sprung mass velocity, z_{def} is the suspension deflection, and $w_{acc}, w_{vel}, w_{heave}, w_{ctrl}$ are weighting coefficients that balance the different objectives. The weight coefficients are defined as $w_{acc} = 0.1$, $w_{vel} = 0.01$,$w_{heave} = 0.005$ and $w_{ctrl} = 0.005$, which are tuned throughout the testing process. Note, the negative sign of the reward function ensure that lower values of these quantities result in higher rewards.

Network architecture: Both the actor and critic networks employ a deep architecture with three hidden layers, each containing 200 neurons, as illustrated in Table 5.4. This relatively large neural network was chosen to handle the high-dimensional observation space (52 inputs) and to capture the non-linear correlations between the observation spance and action space.

Table 5.4: TD3 network architecture

Network	Input Dimension	Hidden Layers	Output Dimension
Actor	52	3 × 200 (ReLU)	2 (tanh + scaling)
Critic (Q1, Q2)	52 (state) + 2 (action)	3 × 200 (ReLU)	1 (Q-value)

The actor network uses ReLU (Rectified Linear Unit) activation functions in the hidden layers to introduce non-linearity into the architecture. The output layer employs a tanh activation function that is followed by a scaling layer that maps the network output, which is defined as [-1, 1] to the action space [0.6, 1.9]. This is achieved through the transformation:

$$a = 0.65 \cdot \tanh(x) + 1.25 \tag{5.2}$$

where x is the pre-activation output and a is the final action.

The critic networks follow a dual-path architecture where state and action information are processed separately before being combined. The state path processes the 52 observations through two fully connected layers, while the action path processes the 2-dimensional action through a single fully connected layer. These paths are then merged using an addition layer, followed by additional processing layers before outputting a single Q-value estimate. This architecture, allows the critic to learn complex state-action value functions efficiently.

Training configuration: The TD3 agent was trained using the parameters listed in Table 5.5, which the hyperparameters were selected based on recommendations from the TD3 literature [31] and adjusted through experiments to suit the truck damper control problem.

Table 5.5: TD3 training hyperparameters

Parameter	Value	Description
Sample time	0.01 s	Control update frequency
Actor learning rate	1×10^{-6}	Step size for actor network updates
Critic learning rate	5×10^{-5}	Step size for critic network updates (higher than actor for stable value estimates)
Target smooth factor (τ)	1×10^{-3}	Soft update coefficient for target networks
Discount factor (γ)	Default (0.99)	Weight for future rewards
Mini-batch size	256	Number of experiences sampled per update
Experience buffer	1×10^{6}	Replay buffer capacity
Exploration noise (std)	0.3	Gaussian noise standard deviation for exploration
Noise decay rate	1×10^{-5}	Rate at which exploration noise decreases
Target policy smoothing	$\mathcal{N}(0, 0.1)$	Noise added to target actions, clipped to [-0.2, 0.2]

The learning rate for the critic networks is set significantly higher (5×10^{-5}) than the actor learning rate (1×10^{-6}). This ensures that value estimates converge before the policy is updated, following the delayed policy update principle of TD3. The target smooth factor of 1×10^{-3} controls how quickly the target networks track the main networks through the soft update rule:

$$\theta_{target} \leftarrow \tau\theta + (1 - \tau)\theta_{target} \tag{5.3}$$

Where θ is the current/online network and θ_{target} is the target network. Exploration during training is achieved through Gaussian noise with an initial standard deviation of 0.3, which gradually decays at a rate of 1×10^{-5} per step. This allows the agent to explore broadly early in training while progressively exploiting learned policies as training progresses. Additionally, target policy smoothing adds clipped Gaussian noise to target actions during critic updates, which helps prevent the policy from exploiting inaccuracies in the learned Q-function.

Training environment and process: The agent was trained in a Simulink environment integrated with IPG TruckMaker. The training process utilized 72 different test scenarios (TestRun1 through TestRun72), each representing different road conditions, speeds, and loading configurations typical for heavy-duty truck operations. For detailed documentation of test scenarios, see Appendix A. A random scenario is selected, at the start of each episode. This ensures the agent learns a robust policy across different driving conditions. The logic for selecting and loading training scenarios is defined as:

Listing 5.1: Episode reset function

```
function in = truckmaker_reset(in)
    idx = randi([1,72]);
    cmguicmd('SimStop');
    cmguicmd(['LoadTestRun␣TestRun' num2str(idx)]);
    pause(0.1);
    cmguicmd('Project.Restart');
end
```

The training was configured with the following episode parameters:

- Maximum episodes: 1000
- Maximum steps per episode: 10,000 (set to ensure a scenario is completed before moving on)
- Score averaging window: 100 episodes. Meaning, performance was evaluated using a rolling average over the most recent 100 episodes to smooth out fluctuations and assess learning trends.

5.4.2 Skyhook Implementation

The Skyhook controller estimates the ideal damping force for each corner using the corner sprung mass acceleration $a_{s,i}$ and the corner mass estimated from the bellow pressure:

$$F_{ideal,i} = m_i \cdot a_{s,i} \tag{5.4}$$

The passivity constraint is enforced by checking the sign of the relative velocity $\dot{z}_{rel,i}$ at each corner. When the relative velocity is positive, meaning the suspension is extending and the damper can usefully resist sprung mass motion, the target force is set to $F_{ideal,i}$. When the relative velocity is negative, the target force is set to the minimum damper force F_{min}, corresponding to the 0.6 A characterisation curve:

$$F_{target,i} = \begin{cases} F_{ideal,i} & \text{if } \dot{z}_s \dot{z}_{rel,i} > 0 \\ F_{min} & \text{otherwise} \end{cases} \tag{5.5}$$

The target force is clamped to the achievable envelope defined by the 0.6 A and 1.9 A characterisation curves evaluated at a reference velocity of the current relative velocity $\dot{z}_{rel,i}$. The corresponding current is then found by linear interpolation within this envelope:

$$I_i = 0.6 + (1.9 - 0.6) \cdot \frac{|F_{clamped,i}| - |F_{low}|}{|F_{high}| - |F_{low}|} \tag{5.6}$$

Front and rear corners use their respective characterisation curves, consistent with the different failsafe configurations of the front and rear dampers. If the force envelope collapses below a numerical threshold of 0.1 N, the current defaults to 0.6 A to avoid division by zero.

5.4.3 Load-Adaptive Implementation

The reference condition is defined as 70% of the maximum Gross Vehicle Mass of 19,000 kg, giving a total of 13,300 kg, with corner masses of 2,651 kg at each front wheel and 3,420 kg at each rear wheel. At this condition, the damping ratio naturally produced by the passive damper is computed from its rebound characteristic:

$$\zeta_{\text{ref}} = \frac{c_s}{2\sqrt{k_s\, m_{\text{corner}}}} \tag{5.7}$$

This ratio becomes the control target. For any other loading condition, the corner mass m_{corner} changes, and the controller adjusts the damper current I to maintain $\zeta = \zeta_{\text{ref}}$. By design, the load-adaptive and passive controllers are therefore identical at the reference load. The behaviour away from this point follows directly:

- **Below reference load:** the passive damper is over-damped; the controller reduces I to lower c_s and restore ζ_{ref}.
- **Above reference load:** the passive damper is under-damped; the controller increases I to raise c_s and restore ζ_{ref}.

The implementation proceeds in three steps at the start of each drive cycle:

1. **Characterisation.** The rebound characteristic is linearised via Weighted Least Squares over the velocity band $v \in [0.052,\ 0.52]$ m/s with uniform weighting, yielding an equivalent damping coefficient $c_{\text{eq}}(I)$ for each current level.
2. **Corner mass estimation.** The sprung mass at each corner is estimated from the measured air bellow forces at a standstill:
$$m_{\text{corner},i} = \frac{F_{\text{static},i}}{g} \tag{5.8}$$
3. **Current selection.** The target damping coefficient at each corner is computed as:
$$c_{\text{target},i} = 2\,\zeta_{\text{ref}}\sqrt{k_s\, m_{\text{corner},i}} \tag{5.9}$$
The optimal current is then found by searching a grid of 131 evenly spaced points over $I_{\text{grid}} \in [0.6,\ 1.9]$ A. The damping coefficient is interpolated between the 0.6 A and 1.9 A characterisation curves using piecewise cubic Hermite interpolation, and the current minimising the residual
$$I_{\text{opt},i} = argmin_{I\,\in\, I_{\text{grid}}} \left|c_{\text{eq}}(I) - c_{\text{target},i}\right| \tag{5.10}$$
is selected. The resulting currents are latched and held constant for the duration of the drive cycle.

5.5 Estimation and Filter Implementation

An overview of the estimation pipeline is shown in Figure 5.9. Since there are required states that cannot be measure by the available equipment, an estimation of these must be made. The state that are estimated are:

- Relative velocity between corners of the sprung mass and unsprung masses [$v_{rel,i}$]
- Vertical acceleration at the corners [$a_{z,i}$]
- Vertical velocity at the CoM [$v_{z,CoM}$]

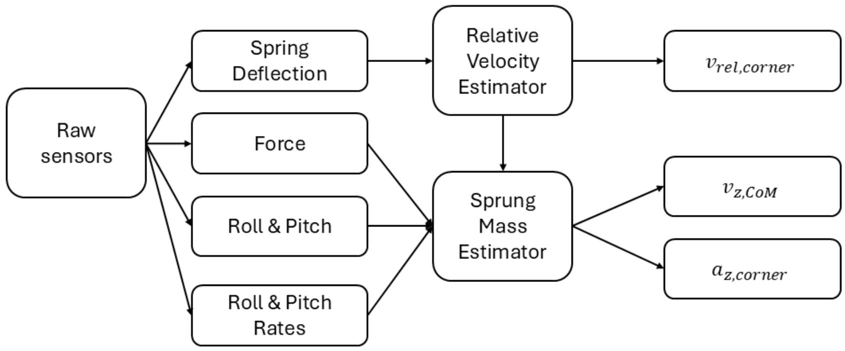

Figure 5.9: Estimation overview

The raw sensors for Roll, Pitch, Rollrate and Pitchrate, are assumed to be measured. This is a limitation that is made to simplify the estimation process and gain more accurate estimations used for evaluating control algorithms.

Relative Velocity Estimation: The relative corner velocities are estimated from the measured spring deflections via a cascade of first-order low-pass filtering and bilinear-transform differentiation. The deflection signal is first scaled to combat loss of amplitude in filtering,

$$z_{\text{rel,scaled}} = 50 \cdot z_{\text{rel}}, \tag{5.11}$$

and filtered using a first-order low-pass filter with time constant τ_{LP},

$$a_{\text{LP}} = e^{-T_s/\tau_{\text{LP}}}, \qquad b_{\text{LP}} = 1 - a_{\text{LP}}, \tag{5.12}$$

$$x_{\text{LP}}[k] = a_{\text{LP}}\, x_{\text{LP}}[k-1] + b_{\text{LP}}\, z_{\text{rel},i}[k]. \tag{5.13}$$

The filtered signal is then differentiated using the bilinear-transform differentiator with time constant τ_d,

$$\alpha = \frac{2\tau_d}{T_s}, \qquad k_1 = \frac{2}{T_s(1+\alpha)}, \qquad k_2 = \frac{1-\alpha}{1+\alpha}, \tag{5.14}$$

$$v_{\text{rel},i}[k] = k_1(x_{\text{LP}}[k] - x_{\text{LP}}[k-1]) + k_2\, v_{\text{rel},i}[k-1]. \tag{5.15}$$

The filter state vector for corner i is

$$\mathbf{s}_i[k] = \begin{bmatrix} x_{\text{LP}}[k] \\ v_{\text{rel},i}[k] \end{bmatrix}. \tag{5.16}$$

Vertical Acceleration and CoM Velocity: The sprung mass is estimated from static or spring forces depending on data availability,

$$m = \begin{cases} \dfrac{\sum F_{\text{static}}}{g} & \text{if } \sum F_{\text{static}} > 1, \\ \dfrac{\sum F_{\text{spring}}}{g} & \text{otherwise.} \end{cases} \tag{5.17}$$

The load-dependent axle distances are computed from the current spring force distribution,

$$F_{\text{total}} = \sum_{i=1}^{4} F_{\text{spring},i}, \qquad l_f = \frac{F_{\text{rear}}}{F_{\text{total}}} L, \qquad l_r = \frac{F_{\text{front}}}{F_{\text{total}}} L. \tag{5.18}$$

The raw vertical acceleration at the centre of mass is obtained by subtracting the gravity projection from the total spring force,

$$a_{z,\text{raw}} = \frac{F_{\text{total}} - mg\cos\phi\cos\theta}{m}. \tag{5.19}$$

This is smoothed with a low-pass filter of cutoff frequency $f_{c,\text{accel}}$ to suppress high-frequency noise,

$$\alpha_{\text{accel}} = \frac{T_s}{\tau_{\text{accel}} + T_s}, \qquad \tau_{\text{accel}} = \frac{1}{2\pi f_{c,\text{accel}}}, \tag{5.20}$$

$$a_{z,\text{CoM}}[k] = \alpha_{\text{accel}}\, a_{z,\text{raw}}[k] + (1 - \alpha_{\text{accel}})\, a_{z,\text{CoM}}[k-1]. \tag{5.21}$$

Corner accelerations are reconstructed by distributing the pitch rate $\dot{\theta}$ and roll rate $\dot{\phi}$ contributions,

$$a_{z,\text{FL}} = a_{z,\text{CoM}} + l_f\dot{\theta} - \tfrac{t_f}{2}\dot{\phi}, \qquad a_{z,\text{FR}} = a_{z,\text{CoM}} + l_f\dot{\theta} + \tfrac{t_f}{2}\dot{\phi}, \tag{5.22}$$

$$a_{z,\text{RL}} = a_{z,\text{CoM}} - l_r\dot{\theta} - \tfrac{t_r}{2}\dot{\phi}, \qquad a_{z,\text{RR}} = a_{z,\text{CoM}} - l_r\dot{\theta} + \tfrac{t_r}{2}\dot{\phi}. \tag{5.23}$$

The vertical velocity at the CoM is estimated using a complementary filter that fuses two sources. The high-frequency component is obtained by integrating $a_{z,\text{raw}}$,

$$v_{z,\text{int}}[k] = v_{z,\text{int}}[k-1] + a_{z,\text{raw}}[k]\, T_s, \tag{5.24}$$

$$v_{\text{HP}}[k] = \alpha_{\text{HP}}(v_{\text{HP}}[k-1] + v_{z,\text{int}}[k] - v_{z,\text{int}}[k-1]), \qquad \alpha_{\text{HP}} = \frac{\tau}{\tau + T_s}. \tag{5.25}$$

The low-frequency component is derived from the kinematic heave velocity at the CoM, computed by removing pitch and roll contributions from the corner relative velocities and projecting to the CoM,

$$v_{\text{z,FL}} = v_{\text{rel,1}} - \left(l_f\dot{\theta} - \tfrac{t_f}{2}\dot{\phi}\right), \qquad v_{\text{z,FR}} = v_{\text{rel,2}} - \left(l_f\dot{\theta} + \tfrac{t_f}{2}\dot{\phi}\right), \tag{5.26}$$

$$v_{\text{z,RL}} = v_{\text{rel,3}} - \left(-l_r\dot{\theta} - \tfrac{t_r}{2}\dot{\phi}\right), \qquad v_{\text{z,RR}} = v_{\text{rel,4}} - \left(-l_r\dot{\theta} + \tfrac{t_r}{2}\dot{\phi}\right), \tag{5.27}$$

$$v_{z,\mathrm{CoM,vrel}} = \frac{l_r}{L} \cdot \frac{v_{\mathrm{z,FL}} + v_{\mathrm{z,FR}}}{2} + \frac{l_f}{L} \cdot \frac{v_{\mathrm{z,RL}} + v_{\mathrm{z,RR}}}{2}, \tag{5.28}$$

$$v_{\mathrm{LP}}[k] = \alpha_{\mathrm{LP}}\, v_{z,\mathrm{CoM,vrel}}[k] + (1 - \alpha_{\mathrm{LP}})\, v_{\mathrm{LP}}[k-1], \qquad \alpha_{\mathrm{LP}} = \frac{T_s}{\tau + T_s}. \tag{5.29}$$

The two branches are finally summed,

$$v_{z,\mathrm{CoM}}[k] = v_{\mathrm{HP}}[k] + v_{\mathrm{LP}}[k], \tag{5.30}$$

combining the low-frequency accuracy of the suspension-derived estimate with the high-frequency accuracy of the acceleration integral.

5.6 Evaluation & Validation

Two road profiles are used for evaluation. The first is a harmonic frequency sweep from 1–5 Hz, targeting the sprung mass resonance range where human sensitivity to vertical vibration is greatest [25]. The second is a stochastic road input generated in accordance with ISO 8608:2016 [27], covering classes A through F to represent surfaces from smooth highway to severely damaged road. The complete test matrix is detailed in Appendix A.1–A.3.

5.6.1 1–5 Hz Swept Road

The swept frequency road profile covers 1–5 Hz in 0.1 Hz increments, targeting the sprung mass resonance range identified in Section 3.5.1. The corresponding spatial period lengths at a reference speed of 10 m/s are listed in Appendix A.4.

5.6.2 Weighted RMS Road Test Case

To evaluate ride comfort in accordance with ISO 2631-1 [26], a swept frequency road profile covering the ISO standardised one-third octave centre frequencies from 1 to 12.5 Hz is used [19]. The profile is constructed by concatenating sinusoidal road sections at the centre frequencies defined in Table 5.6, calculated at a reference speed of 10 m/s using $\lambda = v/f$. This frequency range spans the sprung mass resonance and extends into the unsprung mass region, covering the full band over which the W_k weighting filter of ISO 2631-1 assigns significant weight to vertical vibration. The resulting weighted RMS acceleration a_w is used as the primary comfort metric for this test case.

Table 5.6: ISO frequencies, period lengths at 10 m/s, and W_k weighting factors [19]

Frequency (Hz)	Period Length (m)	W_k weight
1.0	10.0000	482
2.0	5.0000	513
2.5	4.0000	631
3.15	3.1746	804
4.0	2.5000	976
5.0	2.0000	1039
6.3	1.5873	1054
8.0	1.2500	1036
10.0	1.0000	988
12.5	0.8000	902

5.6.3 Stochastic Road Generation

Stochastic road surfaces for ISO 8608 classes A through F were generated in MATLAB using a spectral synthesis approach. The PSD of each profile follows Equation 3.19, with generation parameters summarised in Table 5.7. Lateral correlation between wheel tracks was modelled using an exponential covariance kernel. The profiles were exported in OpenCRG format [42] and imported into TruckMaker via the Road Editor.

Table 5.7: Road surface generation parameters

Parameter	Symbol	Value
Road length	L	500 m
Road width	W	7.0 m
Longitudinal grid spacing	Δu	0.05 m
Lateral grid spacing	Δv	0.05 m
Waviness exponent	w	2.0
Reference spatial frequency	Ω_0	$2\pi \cdot 0.1$ rad/m
Lateral correlation length		1.5 m

Classes A and B represent smooth highway surfaces, C and D typical mixed-use and rural roads, and E and F severely damaged surfaces or constructions sites [27]. The speeds of the tests are shown in Table 5.8 and are chosen based on realistic operating conditions for each road class. Smoother surfaces permit higher speeds representative of highway driving, while rougher surfaces needs reduced speeds consistent with real-world operation on damaged roads. The generated profiles are shown in Figure 5.10. A detailed version of the CRG road profiles can be seen in Appendix C.

Table 5.8: Evaluation speeds per ISO 8608 road class

ISO 8608 Class	Speed (km/h)
A	80
B	80
C	80
D	60
E	40
F	20

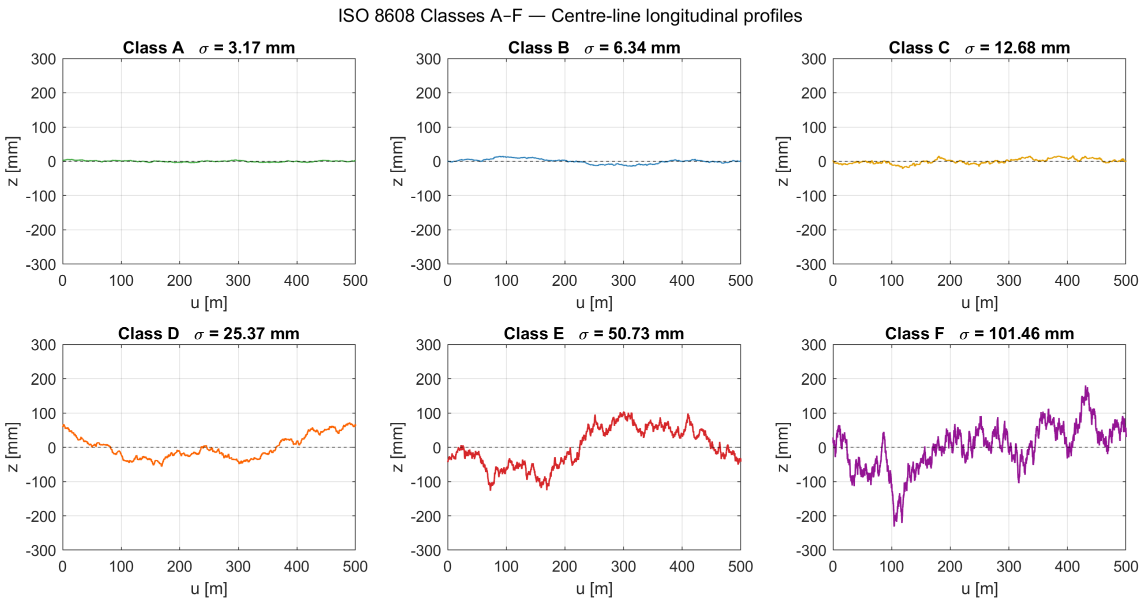

Figure 5.10: Generated road surface profiles for ISO 8608 classes A through F

5.6.4 Estimator Validation

The state estimators described in Section 5.5 are validated against ground truth signals available in TruckMaker. Accuracy is quantified by the RMS error and peak error of $v_{\text{rel},i}$, $a_{z,i}$, and $v_{z,\text{CG}}$ relative to the simulator's internal states, evaluated at a 2 Hz road and across the load cases defined in Section 5.6.5.

5.6.5 Load Cases

Three load cases are defined by the mass applied at the hitch, as summarised in Table 5.9.

Table 5.9: Axle loads for varying hitch loads

Hitch Load	Total Weight (kg)	Front Axle (kg)	Rear Axle (kg)
Unloaded	7,000	4,903	2,097
5,000 kg	12,000	5,673	6,321
10,000 kg	17,000	6,447	10,541

5.6.6 Controller Benchmarks

Five configurations are evaluated under identical conditions in TruckMaker:

- **Passive:** Fixed failsafe current, representing the baseline hardware response.
- **Load-Adaptive:** Rule-based damping adjustment based on static payload.
- **Skyhook:** Classical semi-active control using estimated sprung mass acceleration.
- **Ideal Skyhook:** Skyhook with perfect velocity information from corner accelerometers.
- **RL:** TD3 agent trained across the full scenario matrix.

5.6.7 Logged Metrics

Performance is quantified using the following metrics:

- **Comfort:** RMS vertical acceleration a_z, RMS roll ϕ, RMS pitch θ.
- **Energy:** Mean current $|\bar{I}|$, average power $\bar{P}$, total energy (Wh), where instantaneous power at each corner is computed as

$$P_i(t) = F_{d,i}(t) \cdot \dot{z}_{\mathrm{rel},i}(t) \tag{5.31}$$

6
Results

This chapter presents the result of the thesis including, RL training results, simulation results of different input frequencies and road classes, energy consumption and finally estimation of required states.

6.1 Training Results

The TD3 agent, trained across 72 scenarios converge to a lower average reward shown by the Dark blue line in Figure 6.1. It shows that the agent converge over time, however due to hardware limitations (memory being fully allocated), the agent could not converge completely.

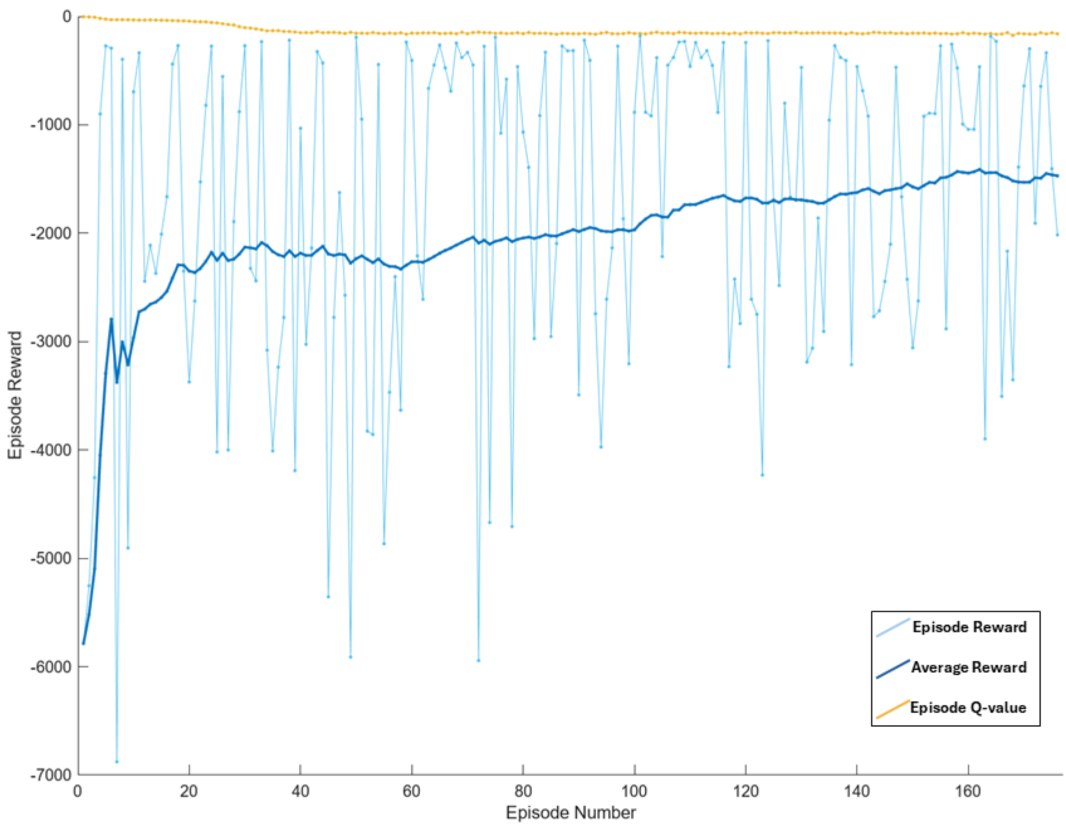

Figure 6.1: Training results for the TD3 agent

6.2 Performance Across Frequencies

This section presents the performance of the discussed control schemes across different road frequencies.

6.2.1 Swept Frequency Road (1–5 Hz)

Figure 6.2 compares RMS vertical acceleration across all controllers under the different load cases. In the unloaded condition, RL delivers the best performance, slightly outperforming both SH and Ideal SH. As the load increases, however, Ideal SH becomes the strongest performer, while RL shows a clear decline. LA performs poorly at high load, ending up nearly identical to the passive system despite consuming significantly more power which is expected.

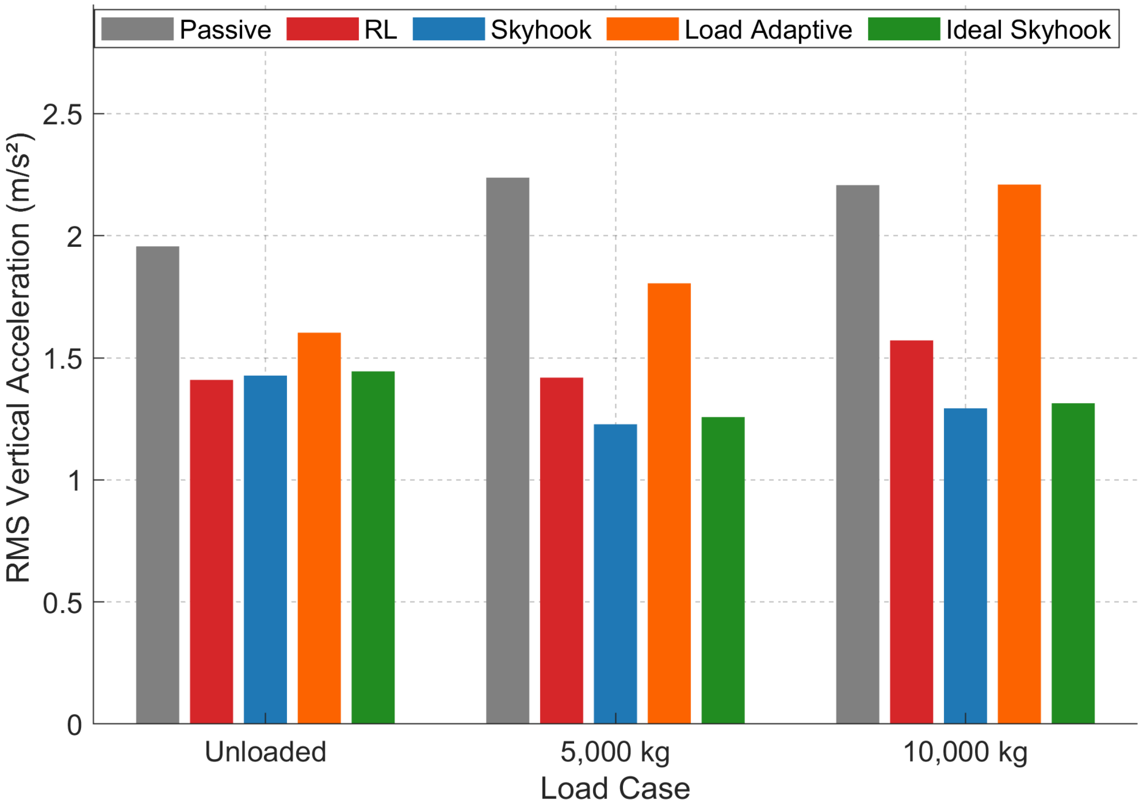

Figure 6.2: RMS Vertical Acceleration, swept frequency road 1–5 Hz

6.2.2 Frequency Domain Analysis

The PSD plots for all load cases clearly show two resonance peaks. As the hitch load increases, both the SH and Ideal SH controllers maintain a consistent gap from the passive baseline.

The LA controller grows closer to the passive baseline as the load increases, which is expected. However, the RL controller also moves closer to the passive baseline at frequencies below 3 Hz. This is not expected and indicates that the RL agent is not effectively adapting to the 10,000 kg load at the primary resonance. Notably, the RL controller maintains a clear separation from the passive baseline at frequencies above 3 Hz across all load cases, suggesting better performance in high-frequency vibration control.

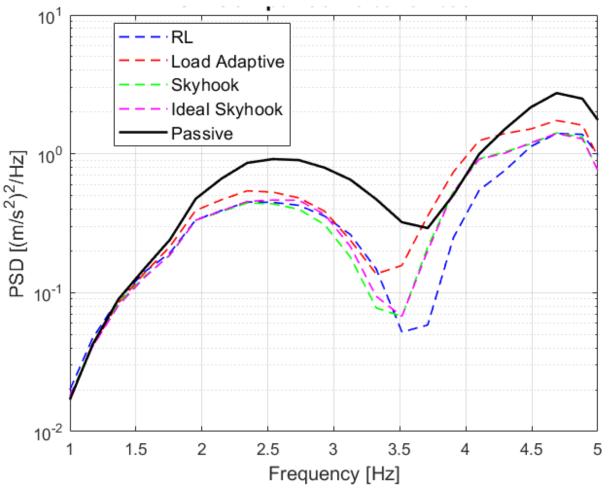

Figure 6.3: PSD comparison of all controllers, 1–5 Hz, unloaded.

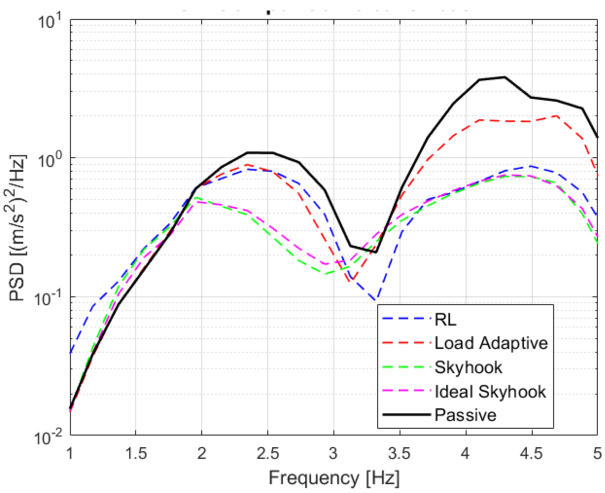

Figure 6.4: PSD comparison of all controllers, 1–5 Hz, 5,000 kg hitch load.

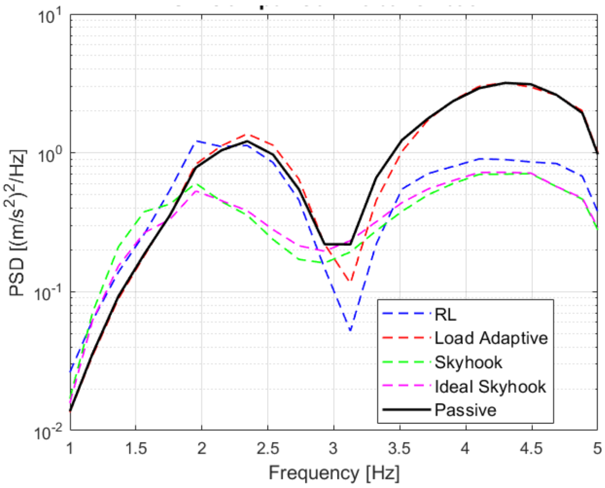

Figure 6.5: PSD comparison of all controllers, 1–5 Hz, 10,000 kg hitch load.

6.2.3 Weighted RMS 1-12.5 Hz

The control schemes was evaluated based on the weighted parameters presented in Section 5.6.2. It combines the RMS of the vertical accelerations from the different frequencies, along with the weights. In Figure 6.6 the results from the weighted RMS tests are presented.

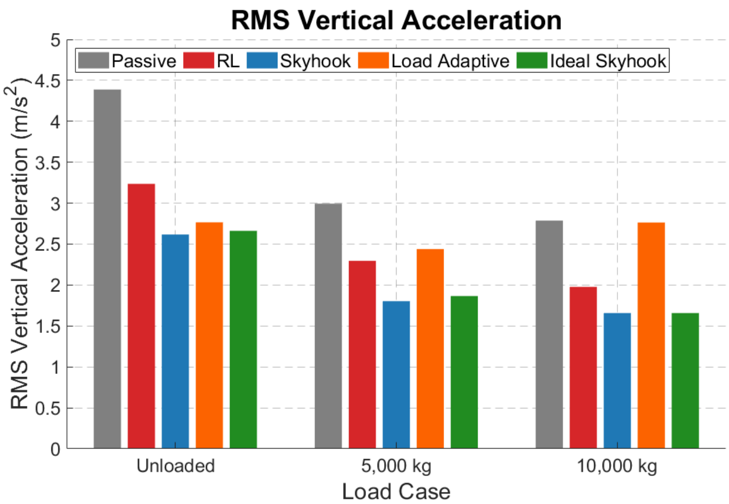

Figure 6.6: Weighted RMS according to ISO 2631

6.3 Performance Across Road Roughness Classes

Controller performance was evaluated across ISO 8608 road classes A through F at three load cases per class. The RMS vertical acceleration results are visualised in Figures 6.7 through 6.9. On smooth surfaces, such as Class A and Class B roads, the passive suspension remains highly competitive across all load cases. Specifically, it achieves the lowest RMS vertical acceleration at the 5,000 kg load case, see Figure 6.8, while active controllers show marginal or no meaningful improvement over passive, despite the higher power consumption inherent to the LA strategy.

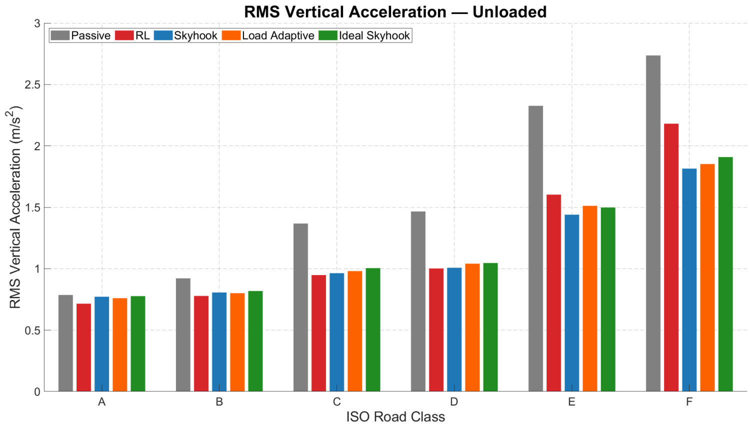

Figure 6.7: RMS vertical acceleration across road classes A–F: Unloaded case.

As road roughness increases to intermediate levels, Classes C and D, the active controllers begin to demonstrate a clear advantage over the passive system. For Class C roads, the RL controller achieves the best performance at 5,000 kg, while the Load Adaptive controller begins to outperform passive at the 10,000 kg condition, see Figure 6.9. This trend solidifies on Class D roads, where active controllers provide consistent improvements across all load cases. Notably, the Ideal SH controller reduces RMS acceleration by approximately 31% relative to the passive system in the unloaded condition (Figure 6.7).

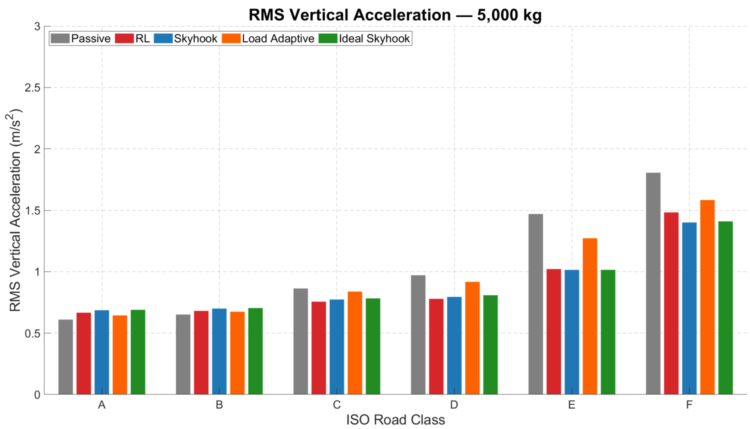

Figure 6.8: RMS vertical acceleration across road classes A–F: 5,000,kg case.

On severe road profiles, Classes E and F, active controllers deliver substantial reductions in RMS acceleration. On Class E roads, SH achieves the best results in the unloaded case, deviating from the typical patterns observed in Classes C and D, while Load Adaptive closely approaches passive performance at 10,000 kg. Finally, Class F serves as a critical robustness benchmark where the active controllers yield the largest absolute reductions relative to passive across the entire dataset. Interestingly, the RL controller underperforms both SH variants at all load cases on this profile.

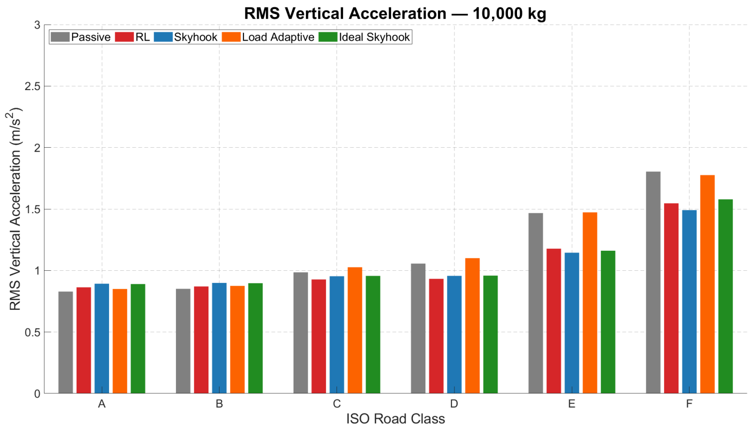

Figure 6.9: RMS vertical acceleration across road classes A–F: 10,000,kg case.

6.4 Energy Consumption

Figure 6.10 illustrates the power consumption trends across all load cases on 1–5 Hz swept frequency road. While SH and Ideal SH maintain near-constant power requirements, the other controllers scale with the mass. Ideal SH consumes more power which is not expected, still less than RL. The RL agent consumes moderately more power than SH, with its requirements increasing alongside the load. In contrast, the LA controller shows a drastic increase, reaching more than double the power consumption of SH at the 10,000,kg case. Despite this high energy cost, the LA system provides no measurable comfort benefit over the passive baseline at maximum load, which is expected due to the LA:s tuning follows the Passive damper.

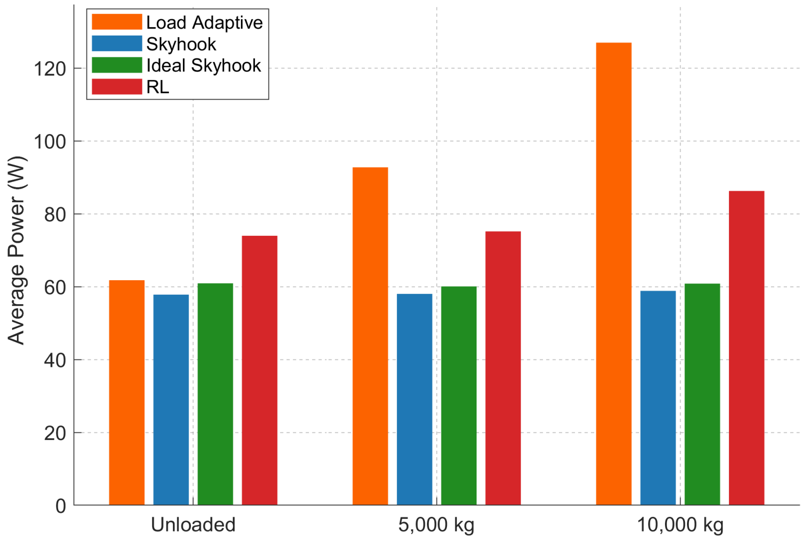

Figure 6.10: Average power consumption of controllers across hitch loads on the 1–5 Hz swept frequency road.

6.5 Estimation results

The quality of the estimated states are compared to measured states from TruckMaker for evaluation. The tests are made with the aforementioned test cases of; unloaded, 5,000 kg and 10,000 kg and driven on a straight road with a sinusoidal wave of 2 Hz. In Figures 6.11 through 6.13 a focused presentation of the results for the unloaded case is presented. For full results see Appendix D. Furthermore, these results were produced using the LA controller to maintain consistency across all results.

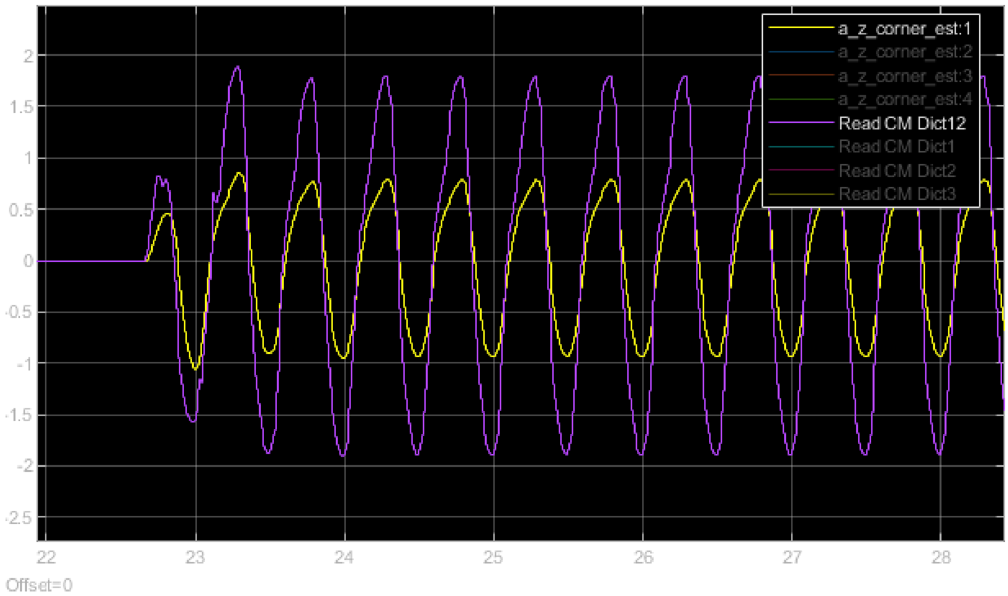

Figure 6.11: Estimated (yellow) compared against measured (purple) accelerations in corners.

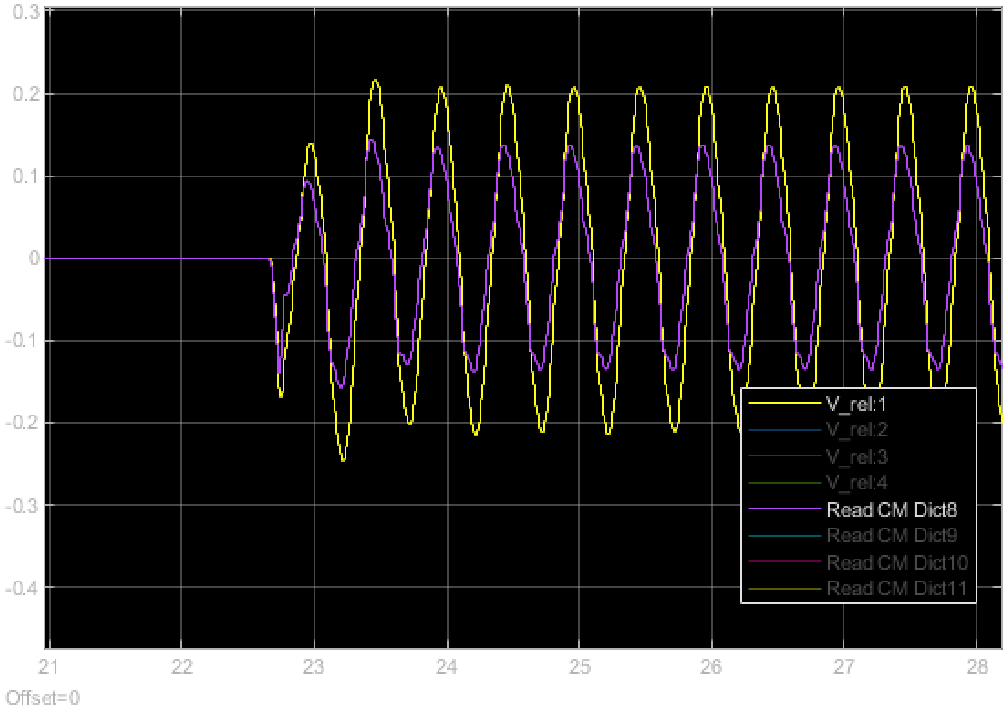

Figure 6.12: Estimated (yellow) compared against measured (purple) relative velocity between sprung and unsprung masses.

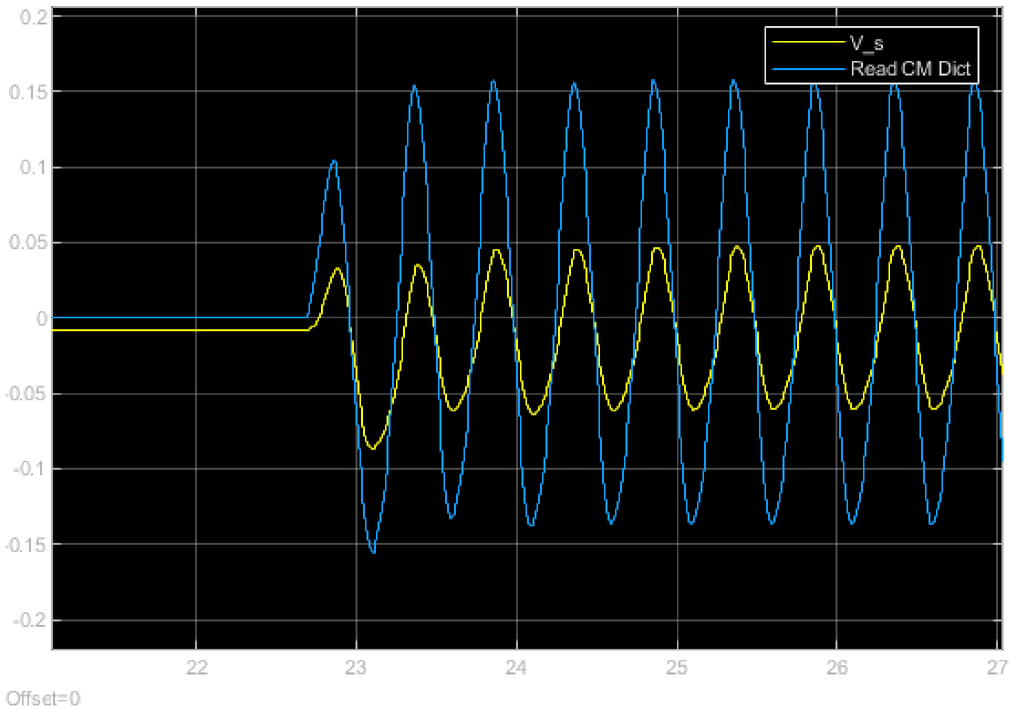

Figure 6.13: Estimated (yellow) compared against measured (blue) sprung mass velocity.

7

Discussion & Verification

In this chapter the results are discussed and validated.

7.1 Performance Analysis

Reinforcement learning has the potential to outperform a conventional Skyhook controller given sufficient training, since the agent learns situation-dependent behaviour directly from the simulation environment rather than following a fixed heuristic. This allows it to adapt to non-linearities, varying load conditions, and transient operating states in ways that a classical controller cannot. Furthermore, the RL agent uses the speed as an observation and has only been trained at lower speeds across a smaller set of frequencies and cases. This might be a big limitation to the performance of the controller. By introducing a more diverse set of training data with more training data in lower speeds, the performance of the RL controller might be improved. This is a phenomenon that is demonstrated in the road class tests, where at higher road classes and thus higher speeds the RL agent performs better than other controllers. However, with more training data the RL agent would need more episodes to train which could be a limiting factor in a development process.

The results indicate that RL achieved strong performance in many test cases, in some cases approaching Ideal Skyhook and consistently outperforming the passive and load-adaptive baselines. However, these gains come at a practical cost. RL requires substantial development time for reward design, hyperparameter tuning, and validation across scenarios. A Skyhook controller, by comparison, is fast to implement, easy to calibrate, and straightforward to reason about. If the performance improvement over estimated Skyhook is modest, the additional development burden is difficult to justify.

A further limitation is interpretability. Skyhook behaviour follows directly from its control logic and is easy to troubleshoot. An RL policy offers no such transparency, which complicates systematic tuning and validation. In an industrial setting, traceability and engineer confidence in the controller are often as important as performance, making a simpler and more interpretable controller preferable even if RL achieves marginally better results.

The RL agent was trained exclusively on the road profiles defined in the training matrix, which targets temporal frequencies between approximately 0.4 Hz and

30 Hz at the evaluated speeds. Smooth surfaces such as ISO 8608 Class A and B, which produce low-amplitude excitation at the tested speeds, were not explicitly represented as distinct training conditions. The degraded RL performance on these surfaces is therefore consistent with a distribution shift rather than a fundamental limitation of the algorithm, and could likely be addressed by including low-roughness scenarios in the training matrix. In particular, the agent appears to have learned a policy that assumes active intervention remains beneficial even when the road excitation is weak, which is less suitable for smooth highway-like conditions.

7.2 Implementation Feasibility

The results across ISO 8608 classes A and B demonstrate that on smooth road surfaces, representative of highway driving, active suspension control provides marginal or no comfort benefit over a well-tuned passive system. From a pure comfort perspective, the case for semi-active damping on vehicles operating primarily on highways is therefore weak. The passive damper, requiring no sensors, no control logic, and no power consumption, performs comparably or better under these conditions.

The argument for active control strengthens considerably on Class C and above, which represent typical mixed-use and poor road conditions encountered in heavy-duty transport. Here, Skyhook and Ideal Skyhook consistently reduce RMS vertical acceleration by 20–30% relative to passive, which is a meaningful improvement for driver fatigue and long-term occupational health.

The load-adaptive controller requires no dynamic sensing, no training, and no online optimisation beyond a standstill calibration. It represents the simplest implementation path among the evaluated active strategies and has the lowest computational cost, since the current is determined once and then held constant during the drive cycle. Its main limitation is that power consumption increases significantly with payload, while the comfort benefit becomes small at higher loads. A practical improvement would be to deactivate or bypass the load-adaptive logic when the vehicle operates near the reference load region of 70% GVW, where the passive damper is already close to the target damping ratio, and in higher-load operating regions where the comfort benefit is limited, thereby reducing unnecessary current draw. The controller could be used as a low-complexity comfort enhancement in favourable operating conditions while avoiding unnecessary power consumption when the benefit is small. Combined with a threshold-based handling controller, it represents the lowest-complexity path to a deployable semi-active system. Both controllers operate on signals already present in air-sprung trucks and introduce no additional hardware requirements. This combination could be introduced in production at minimal cost and serves as a viable intermediate step before committing to a more complex strategy such as Ideal Skyhook or Reinforcement Learning.

Among the active controllers evaluated, estimated Skyhook represents the most cost-effective option. It requires no additional sensors, consumes the least power, and delivers competitive comfort performance across all road classes. Ideal Skyhook offers marginally better comfort at the cost of four additional accelerometers. The RL agent offers no hardware advantage over estimated Skyhook, consumes more power, and requires significant development and validation effort, making it difficult to justify for production deployment in its current form.

7.3 Sensor and Estimation Validity

The estimations generally track the correct frequency content with only small phase errors; the dominant limitation is amplitude accuracy. Amplitude errors are significant because the Skyhook and RL controllers both scale their damping commands directly to the estimated velocity and acceleration magnitudes, an underestimated amplitude produces a weaker damping response than the road disturbance warrants, and an overestimated one risks over-damping. Each state carries its own amplitude uncertainty:

- **Relative corner velocity** [$v_{\mathrm{rel},i}$]: obtained by differentiating the height sensor signal, which inherently amplifies high-frequency noise. The low-pass filter time constant τ_{LP} partially recover amplitude, but the trade-off between noise suppression and amplitude fidelity means the peak velocity magnitude is systematically attenuated at higher excitation frequencies. Therefore, the measurements are preprocessed by a scalar constant to recover the amplitude.
- **Corner vertical acceleration** $a_{z,i}$: derived from total spring force, which depends on an accurate mass estimate m and assumes a known, stable spring characteristic. In practice, bellow stiffness varies with temperature, supply pressure, and wear, directly scaling the computed acceleration. Errors in m or the force-to-mass conversion appear as a constant amplitude gain error across all operating conditions. Additionally, the low-pass filter applied to $a_{z,\mathrm{raw}}$ at cutoff $f_{c,\mathrm{accel}}$ attenuates acceleration amplitudes at frequencies approaching the cutoff, meaning fast transient events, such as sharp bumps, are underreported in $a_{z,i}$
- **CoM vertical velocity** [$v_{z,\mathrm{CoM}}$]: the complementary filter fuses the acceleration integral (high-frequency branch) with the suspension-derived kinematic velocity (low-frequency branch). The crossover frequency τ determines how much each source contributes near the body resonance frequency (1-2 Hz), where control authority matters most. Drift in the integration branch and amplitude attenuation in the low-pass branch mean that near resonance, precisely where Skyhook is most active, the fused estimate can underreport peak velocity magnitude.

The Skyhook controller relies on the estimated sprung mass velocity derived from bellow pressure and relative height measurements. In simulation, the air spring model is exact by definition, meaning the estimation operates under ideal conditions. In a real deployment, the pressure-to-force conversion assumes a known and stable spring characteristic that varies with temperature, air supply pressure, and component wear. Velocity obtained by differentiating height sensor signals amplifies noise, and the 10 ms actuator delay introduces a small phase error that does not effect performance massively. The real-world performance of estimated Skyhook should therefore be expected to fall between the simulated results and passive, depending on operating conditions.

The 10 ms actuator delay and 50 ms CAN bus delay also had a measurable impact on controller performance, which is consistent with literature reporting that semi-active Skyhook begins to degrade beyond similar delay levels. The RL agent, trained in an environment that includes these delays, showed greater robustness to their effect than the classical controllers. This suggests that delay-inclusive training can improve robustness, even if the resulting policy is still limited by its training distribution.

7.4 Energy Consumption Analysis

The energy consumption results reveal a clear separation between the controllers. Skyhook and Ideal Skyhook maintain stable power consumption across all load cases, reflecting the nature of their control logic. Both respond only to the current dynamic state of the vehicle and are therefore largely insensitive to static payload.

The load-adaptive controller behaves differently. Its power consumption grows substantially with increasing hitch load, consistent with its objective of maintaining a target damping ratio under varying mass. However, the additional energy expenditure does not translate into a corresponding comfort improvement. At the highest load case, load-adaptive achieves virtually identical RMS acceleration to passive while consuming significantly more power, representing a poor energy-to-performance ratio in loaded operating conditions.

The RL agent consumes moderately more power than Skyhook and Ideal Skyhook across all load cases, with consumption increasing alongside load. This suggests the agent has learned a more active damping policy, intervening more frequently or at higher currents to achieve its marginal comfort gains. Whether this additional energy expenditure is justified depends on the weight placed on comfort versus operating cost in a given application.

The energy metrics reported here capture only the electrical power consumed by the damper actuators. A more complete picture would account for the indirect effect of damping on drivetrain losses: changes in damper force alter tyre normal loads, which influence longitudinal slip and therefore the torque and speed at the wheels. A controller that reduces normal force variation may thus yield drivetrain savings beyond what actuator power alone reflects. This coupling between suspension behaviour and drivetrain efficiency is not captured by the current metrics, and represents a worthwhile direction for future analysis.

7.4.1 Performance on Smooth Road Surfaces

The Class A results expose a limitation common to all active controllers evaluated in this thesis. On smooth road surfaces, the benefit of active damping is reduced and in some load cases reversed. This is consistent with the theoretical expectation that semi-active control adds value primarily when the road excitation is significant enough to warrant intervention. On a Class A surface the passive damper is already operating close to its optimal condition, and the switching behaviour of the active controllers introduces small perturbations that marginally worsen comfort.

For the RL agent specifically, the degraded performance at higher load cases on Class A suggests the agent has not learned to recognise low-excitation conditions as situations requiring minimal intervention. A practical mitigation would be to include a road roughness estimate in the agent state space, allowing it to modulate its intervention level based on current surface conditions.

7.5 Further Work

This thesis has focused primarily on ride comfort as the central performance metric. Several directions remain open for future investigation.

Handling performance. This thesis evaluates comfort in isolation. Handling represents the other half of the suspension performance problem, and a comprehensive evaluation of the semi-active dampers should address both. Future work could quantify the handling benefit of active damping under cornering, braking, and lane-change manoeuvres to provide a more complete picture of the value these systems bring.

Real-world implementation and validation. Simulation necessarily idealises sensor behaviour. Real sensors introduce noise, drift, and calibration errors that are difficult to replicate faithfully in a virtual environment. On-vehicle testing would reveal limitations not visible in simulation and constitutes the essential next step before any of the evaluated controllers could be considered for production deployment. This includes validation under conditions where vehicle parameters such as tyre pressures, suspension wear, and air spring characteristics deviate from the nominal model assumptions used in this thesis.

Subjective driver evaluation. A structured driver evaluation, such as a paired comparison protocol, would either validate or challenge the objective rankings presented here. Two controllers with similar RMS values can feel substantially different due to differences in transient sharpness or low-frequency body sway, and subjective assessment remains the most direct measure of the comfort this work ultimately targets.

RL trained with weighted frequencies. The results showed that the RL based controller did not perform as well as the skyhook based controllers on the weighted

RMS tests. By integrating the weights into the RL training, such that the frequencies with highest weight are penalized more by higher vertical accelerations, a lower overall weighted RMS might be achieved. However, the implementation of weighted RL training might negatively affect the performance of other frequencies.

RL-based Skyhook tuning. The traditional Skyhook controller relies on fixed parameters that are tuned manually. A reinforcement learning agent could be used to optimise these parameters online across varying load and road conditions, potentially closing the performance gap between estimated and Ideal Skyhook without the full complexity of an end-to-end RL policy.

Increased model fidelity. The simulation uses a rigid chassis model with simplified tyre dynamics. Higher fidelity models including flexible frame torsion, advanced tyre models, trailer dynamics, and alternative truck configurations would improve the transferability of the results to real deployment conditions.

Cornering and combined manoeuvres. All evaluations in this thesis were conducted on straight roads. Real-world operation involves combined lateral and vertical excitation, and future work should evaluate controller performance on curved roads and during combined braking and cornering events.

Drivetrain interaction. The energy analysis in this work considers only the electrical power consumed by the damper actuators. Damping forces influence tyre normal loads, which affect longitudinal slip and drivetrain losses. Future work could quantify this coupling by comparing drivetrain output power across controllers, as reductions in normal force variation may yield energy savings beyond what actuator power alone captures.

8 Conclusion

Load-Adaptive emerges as the most implementation-ready controller among the evaluated semi-active strategies. It is the simplest to deploy, requires no dynamic sensing, no training, and no online optimisation, and has the lowest computational cost. Although its comfort benefit is limited, especially at high load, its simplicity and compatibility with already available vehicle signals make it a realistic low-risk path toward production implementation. With a suitable enable/disable logic, it could be used only in operating regions where a comfort benefit is expected, thereby limiting unnecessary power consumption.

Skyhook with accurate state estimation remains a strong comfort-oriented candidate, offering good performance with relatively low power consumption. However, its practical reliability cannot be concluded from this study alone, since the estimation was evaluated using ideal virtual sensor signals and a rigid-body simulation model without sensor noise, drift, or model mismatch. A more realistic assessment would require higher-fidelity modelling and experimental validation on vehicle.

Reinforcement Learning demonstrated the ability to match or exceed the classical controllers in several test cases, but at higher power consumption, greater implementation complexity, and reduced robustness outside its training distribution. In its current form, it is therefore better viewed as a promising research direction than a near-term production candidate.

Bibliography

[1] E. Guglielmino, T. Sireteanu, C. W. Stammers, G. Gheorghe, and M. Giuclea, *Semi-Active Suspension Control: Improved Vehicle Ride and Road Friendliness.* Springer London, 2008.

[2] C. Lauwerys, J. Swevers, and P. Sas, “A Model Free Control Design Approach for a Semi-Active Suspension of a Passenger Car,” in *Proceedings of the International Symposium on Advanced Vehicle Control (AVEC).*

[3] R. Desai, A. Guha, and P. Seshu, “A Comparison of Quarter, Half and Full Car Models for Predicting Vibration Attenuation of an Occupant in a Vehicle,” *Journal of Vibration Engineering and Technologies*, vol. 9, no. 5, pp. 983–1001, 7 2021.

[4] J. Ultsch, A. Pfeiffer, J. Ruggaber, T. Kamp, J. Brembeck, and J. Tobolář, “Reinforcement Learning for Semi-Active Vertical Dynamics Control with Real-World Tests,” *Applied Sciences*, vol. 14, no. 16, 8 2024.

[5] E. Obialero, “A Refined Vehicle Dynamic Model for Driving Simulators,” Master’s thesis.

[6] D. Karnopp, “Active and Semi-Active Vibration Isolation,” in *Current Advances in Mechanical Design and Production, Sixth Cairo University International MDP Conference.*

[7] M. Valášek, M. Babíc, Z. Šika, and L. Magdolen, “Development of Semi-Active Truck Suspension,” in *Vehicle System Dynamics Supplement.*

[8] S. M. Savaresi, E. Silani, and S. Bittanti, “Acceleration-Driven-Damper (ADD): An Optimal Control Algorithm for Comfort-Oriented Semiactive Suspensions,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 127, no. 2, pp. 218–229, 6 2005.

[9] S. M. Savaresi and C. Spelta, “Mixed Sky-Hook and ADD: Approaching the Filtering Limits of a Semi-Active Suspension,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 4, pp. 382–392, 7 2007.

[10] R. Dessort and C. Chucholowski, “Explicit Model Predictive Control of Semi-Active Suspension Systems Using Artificial Neural Networks,” in *Proceedings in Automotive Engineering*, 2017, pp. 207–228.

[11] F. Kjellberg and S. Sundell, “Real-Time Nonlinear Model Predictive Control for Semi-Active Suspension with Road Preview,” Master’s thesis.

[12] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous Control with Deep Reinforcement Learning,” *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016. [Online]. Available: http://arxiv.org/abs/1509.02971

[13] G. Savaia, S. Formentin, G. Panzani, M. Corno, and S. M. Savaresi, "Enhancing Skyhook for Semi-Active Suspension Control via Machine Learning," *IFAC Journal of Systems and Control*, vol. 17, 9 2021.

[14] S. Tognetti, M. Restelli, S. M. Savaresi, and C. Spelta, "Batch Reinforcement Learning: An Application to a Controllable Semi-Active Suspension System," in *Proceedings of the 6th International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, 2009, pp. 228–233.

[15] M. Shaikh and E. Kamardeen, "Influence of Signal Delays and Accuracy on the Performance of Suspension Control System and Controller Design," Master's thesis.

[16] Y. Wang, C. Wang, S. Zhao, and K. Guo, "Research on Deep Reinforcement Learning Control Algorithm for Active Suspension Considering Uncertain Time Delay," *Sensors*, vol. 23, no. 18, 9 2023.

[17] E. Solomon, A. Paul, and S. Varadharajan, "Semi-Active Suspension Control Using Deep Reinforcement Learning," Master's Thesis, KTH Royal Institute of Technology.

[18] D. Kim and Y. Jeong, "Road-Adaptive Static Output Feedback Control of a Semi-Active Suspension System for Ride Comfort," *Actuators*, vol. 13, no. 10, 10 2024.

[19] B. Jacobson, "Vehicle Dynamics Compendium," Chalmers University of Technology, Tech. Rep. [Online]. Available: www.chalmers.se

[20] ——, "Vehicle Motion Engineering," Chalmers University of Technology, Tech. Rep. [Online]. Available: https://research.chalmers.se/en/person/bengtja

[21] Y. Huang, "Target-Driven Road Vehicle Suspension Design," Licentiate Thesis, Chalmers University of Technology.

[22] "Road Vehicles — Vehicle Dynamics and Road-Holding Ability — Vocabulary," International Organization for Standardization, Geneva, Switzerland, Tech. Rep. ISO 8855:2011, 2011.

[23] B. Schofield, "Vehicle Dynamics Control for Rollover Prevention," Licentiate Thesis, Department of Automatic Control, Lund Institute of Technology, Lund University, 2006.

[24] F. Skoglund, "Strategies for Road Profile in Adaptive Suspension Control," Master's thesis, 2020.

[25] C. Deubel, S. Ernst, and G. Prokop, "Objective Evaluation Methods of Vehicle Ride Comfort — A Literature Review," *Journal of Sound and Vibration*, vol. 548, 3 2023.

[26] "Mechanical Vibration and Shock — Evaluation of Human Exposure to Whole-Body Vibration — Part 1: General Requirements," International Organization for Standardization, Geneva, Switzerland, Tech. Rep. ISO 2631-1:1997, 1997.

[27] "Mechanical Vibration — Road Surface Profiles — Reporting of Measured Data," International Organization for Standardization, Geneva, Switzerland, Tech. Rep. ISO 8608:2016, 2016.

[28] H. Min, J. Lu, S. Xu, N. Duan, and W. Chen, "Neural Network-Based Output-Feedback Control for Stochastic High-Order Non-Linear Time-Delay Systems with Application to Robot System," *IET Control Theory and Applications*, vol. 11, no. 10, pp. 1578–1588, 6 2017.

[29] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA: MIT Press, 2018.

[30] N. Heess, G. Wayne, D. Silver, T. Lillicrap, T. Erez, and Y. Tassa, "Learning continuous control policies by stochastic value gradients," in *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds., vol. 28. Curran Associates, Inc., 2015.

[31] S. Fujimoto, H. van Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 2018, pp. 1587–1596. [Online]. Available: https://proceedings.mlr.press/v80/fujimoto18a.html

[32] D. Karnopp, M. J. Crosby, and R. A. Harwood, "Vibration Control Using Semi-Active Force Generators," *Journal of Engineering for Industry*, vol. 96, no. 2, pp. 619–626, 1974.

[33] S. M. Savaresi, C. Poussot-Vassal, C. Spelta, O. Sename, and L. Dugard, *Semi-Active Suspension Control Design for Vehicles*. Elsevier, 2010.

[34] D. Karnopp, "Active and Semi-Active Vibration Isolation," *Journal of Mechanical Design*, vol. 117, pp. 177–185, 1995.

[35] D. Zhang, J. Wang, Z. Huang, H. Hong, S. Yu, W. Shi, and X. Zhang, "Skyhook-Based Techniques for Vehicle Suspension Control: A Review of the State of the Art," *Machines*, 2025.

[36] R. G. Brown and P. Y. C. Hwang, *Introduction to Random Signals and Applied Kalman Filtering*, 4th ed. Wiley, 2012.

[37] W. T. Higgins, "A Comparison of Complementary and Kalman Filtering," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-11, no. 3, pp. 321–325, 5 1975.

[38] R. Mahony, T. Hamel, and J.-M. Pflimlin, "Nonlinear Complementary Filters on the Special Orthogonal Group," *IEEE Transactions on Automatic Control*, vol. 53, no. 5, pp. 1203–1218, 6 2008.

[39] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd ed. Prentice Hall, 2010.

[40] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*, 4th ed. Pearson, 2007.

[41] G. F. Franklin, J. D. Powell, and M. L. Workman, *Digital Control of Dynamic Systems*, 3rd ed. Addison-Wesley, 1998.

[42] ASAM e.V., "OpenCRG — Open File Format and Toolbox for Road Surface Description," https://www.asam.net/standards/detail/opencrg/, 2018.

A

Training Run and Road setup Matrices

This appendix presents the complete training run matrix used to train the Reinforcement Learning agent. The 72 scenarios are organised across three road profile categories, Fast Waves, Big Slow Waves, and Slow Bumps, each evaluated at four load conditions and multiple vehicle speeds. Together these scenarios cover a temporal excitation frequency range of approximately 0.4 Hz to 30 Hz, ensuring the agent is exposed to both sprung mass and unsprung mass resonance conditions during training.

Table A.1: Test Run Matrix: Fast Waves (Periods 0.8, 2, and 15)

Load / Velocity	**20 kph**	**40 kph**	**60 kph**	**85 kph**
Fast Waves 0.8 Period				
Frequency (approx.)	∼7 Hz	∼14 Hz	∼21 Hz	∼30 Hz
0 kg	Testrun1	Testrun2	Testrun3	Testrun4
3000 kg	Testrun5	Testrun6	Testrun7	Testrun8
7000 kg	Testrun9	Testrun10	Testrun11	Testrun12
12,000 kg	Testrun13	Testrun14	Testrun15	Testrun16
Fast Waves 2 Period				
Frequency (approx.)	∼3 Hz	∼5 Hz	∼8 Hz	∼12 Hz
0 kg	Testrun17	Testrun18	Testrun19	Testrun20
3000 kg	Testrun21	Testrun22	Testrun23	Testrun24
7000 kg	Testrun25	Testrun26	Testrun27	Testrun28
12,000 kg	Testrun29	Testrun30	Testrun31	Testrun32
Fast Waves 15 Period				
Frequency (approx.)	∼0.4 Hz	∼0.7 Hz	∼1.1 Hz	∼1.5 Hz
0 kg	Testrun33	Testrun34	Testrun35	Testrun36
3000 kg	Testrun37	Testrun38	Testrun39	Testrun40
7000 kg	Testrun41	Testrun42	Testrun43	Testrun44
12,000 kg	Testrun45	Testrun46	Testrun47	Testrun48

Table A.2: Test Run Matrix: Big Slow Waves

Load / Velocity	**50 kph**	**65 kph**	**80 kph**
0 kg	Testrun49	Testrun50	Testrun51
3000 kg	Testrun52	Testrun53	Testrun54
7000 kg	Testrun55	Testrun56	Testrun57
12,000 kg	Testrun58	Testrun59	Testrun60

Table A.3: Test Run Matrix: Slow Bumps

Load / Velocity	**10 kph**	**15 kph**	**20 kph**
0 kg	Testrun61	Testrun62	Testrun63
3000 kg	Testrun64	Testrun65	Testrun66
7000 kg	Testrun67	Testrun68	Testrun69
12,000 kg	Testrun70	Testrun71	Testrun72

Table A.4: 1–5 Hz road setup, temporal frequency is calculated at a speed of 10 m/s

Frequency (Hz)	Period Length (m)
1	10,00
1,1	9,09
1,2	8,33
1,3	7,69
1,4	7,14
1,5	6,67
1,6	6,25
1,7	5,88
1,8	5,56
1,9	5,26
2	5,00
2,1	4,76
2,2	4,55
2,3	4,35
2,4	4,17
2,5	4,00
2,6	3,85
2,7	3,70
2,8	3,57
2,9	3,45
3	3,33
3,1	3,23
3,2	3,13
3,3	3,03
3,4	2,94
3,5	2,86
3,6	2,78
3,7	2,70
3,8	2,63
3,9	2,56
4	2,50
4,1	2,44
4,2	2,38
4,3	2,33
4,4	2,27

B

Damper data sheets

Below are the front and rear damper data sheets. Only the base settings differ between the dampers. The physical damper is the same product.

FIGURES REDACTED DUE TO CONFIDENTIALITY

C

CRG Road Surface Profile

The following figures present the generated ISO 8608 road surface profiles. Individual profile details are shown for classes A through D.

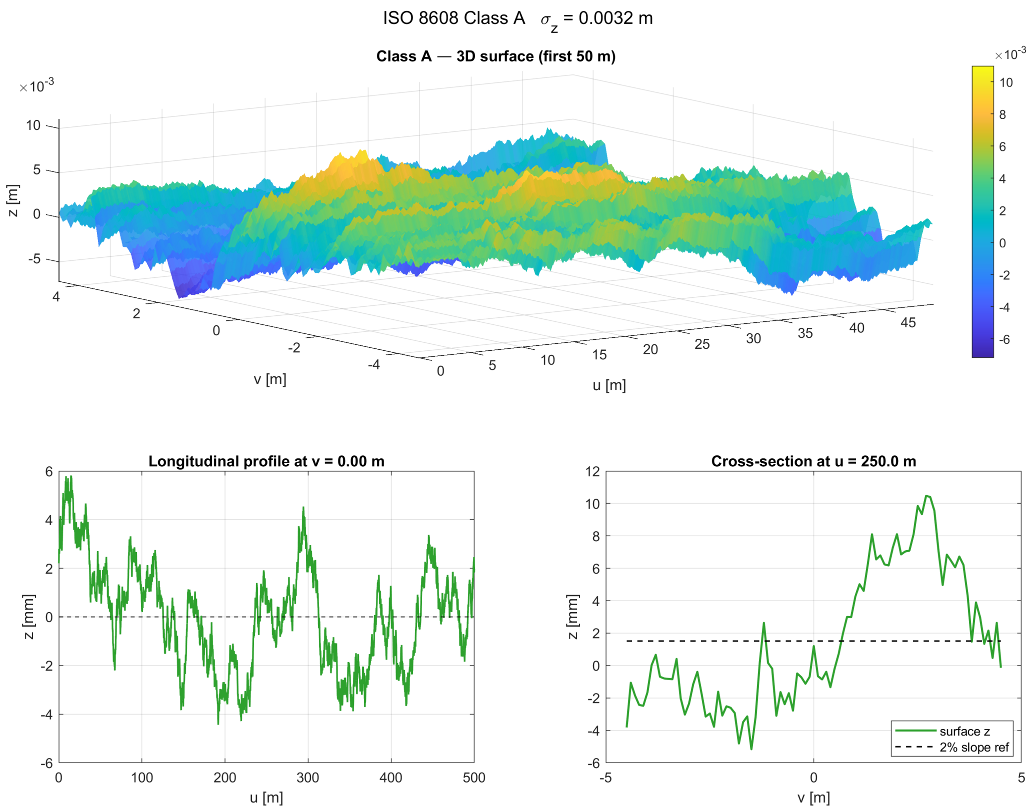

Figure C.1: ISO 8608 Class A road surface profile detail

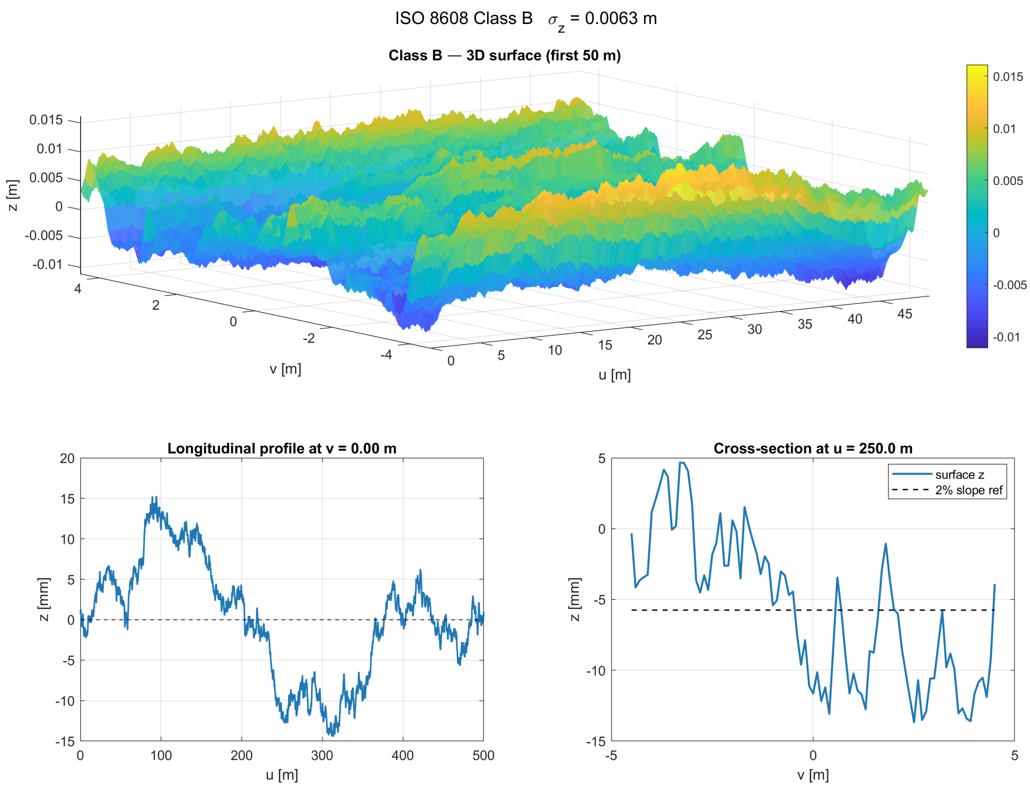

Figure C.2: ISO 8608 Class B road surface profile detail

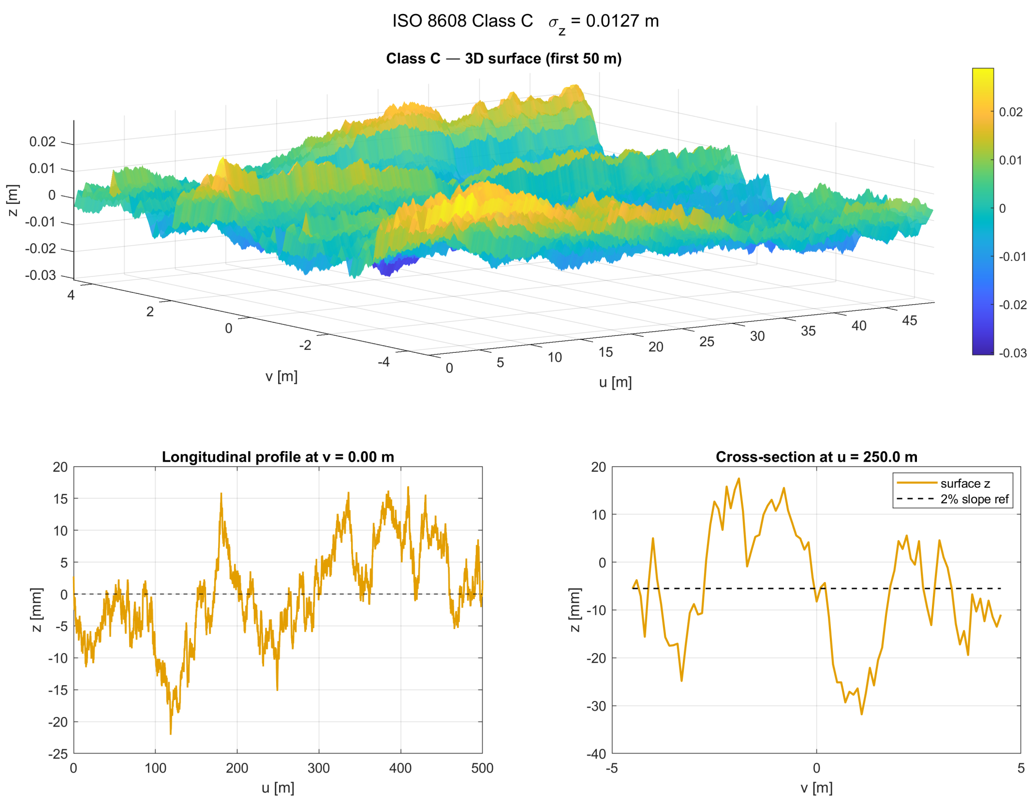

Figure C.3: ISO 8608 Class C road surface profile detail

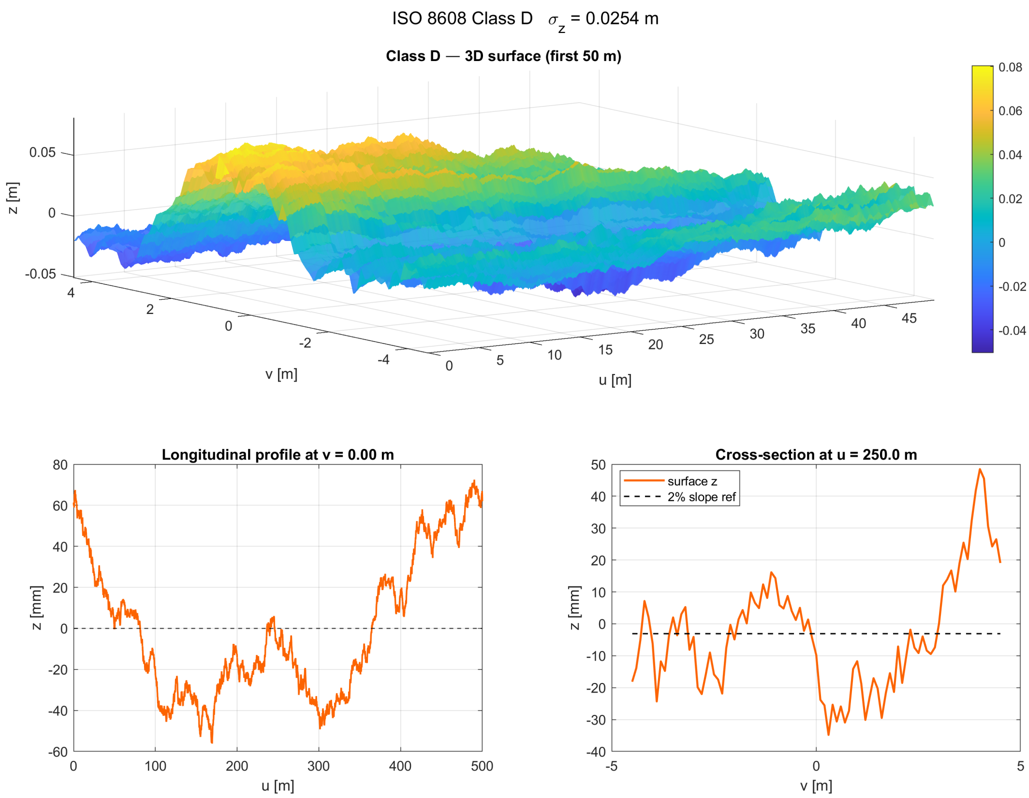

Figure C.4: ISO 8608 Class D road surface profile detail

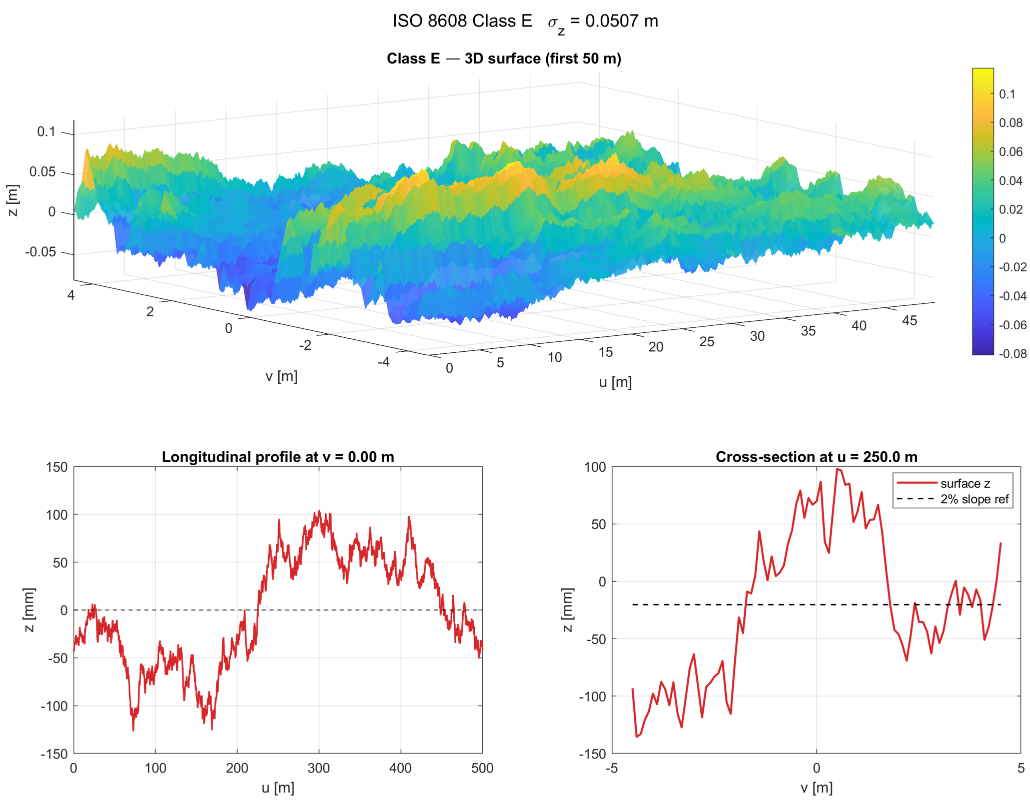

Figure C.5: ISO 8608 Class E road surface profile detail

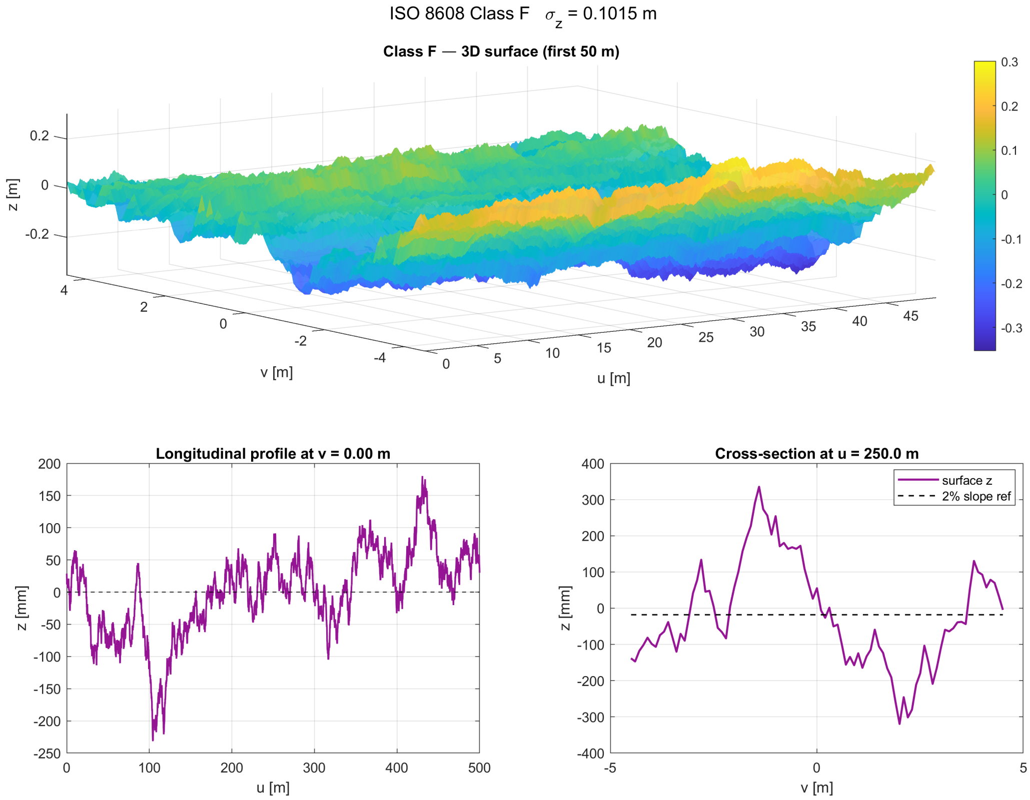

Figure C.6: ISO 8608 Class F road surface profile detail

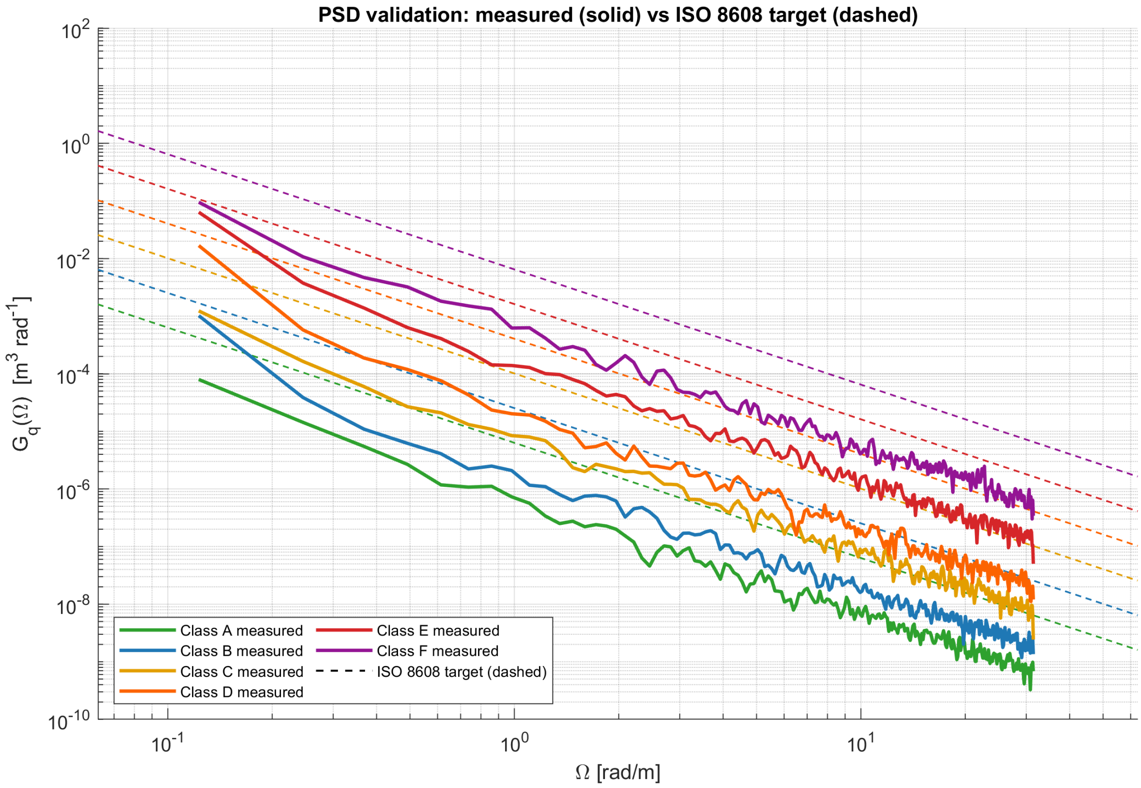

Figure C.7: PSD validation of generated road surfaces against ISO 8608 theoretical targets for classes A through F

D

Estimation Results

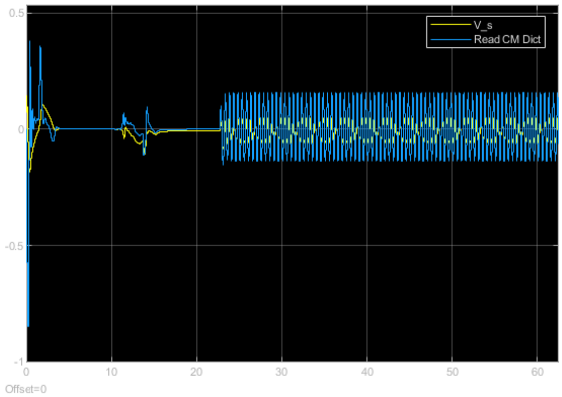

Figure D.1: Unloaded sprung mass velocity estimation (yellow) along with measured (blue)

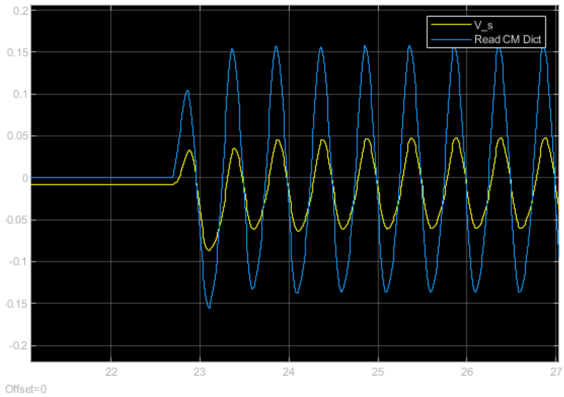

Figure D.2: Unloaded sprung mass velocity estimation (yellow) along with measured (blue)

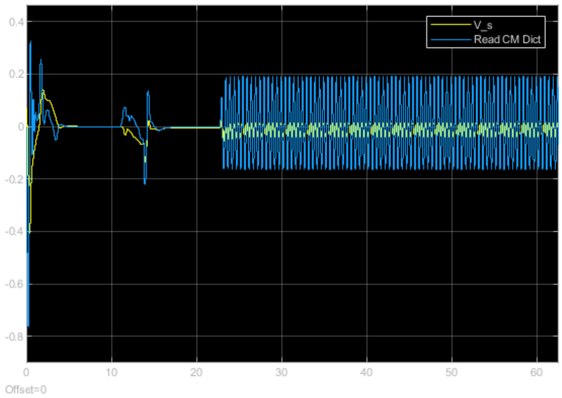

Figure D.3: 5,000 kg sprung mass velocity estimation (yellow) along with measured (blue)

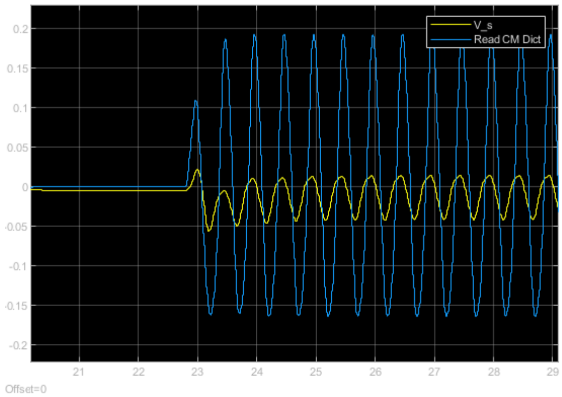

Figure D.4: 5,000 kg sprung mass velocity estimation (yellow) along with measured (blue)

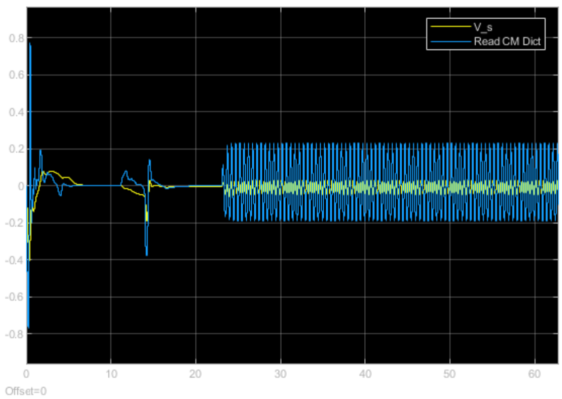

Figure D.5: 10,000 kg sprung mass velocity estimation (yellow) along with measured (blue)

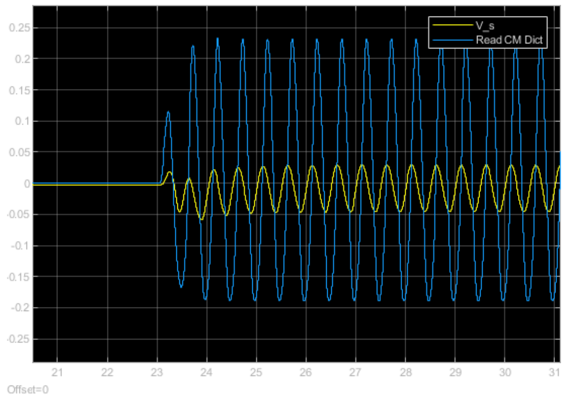

Figure D.6: 10,000 kg sprung mass velocity estimation (yellow) along with measured (blue)

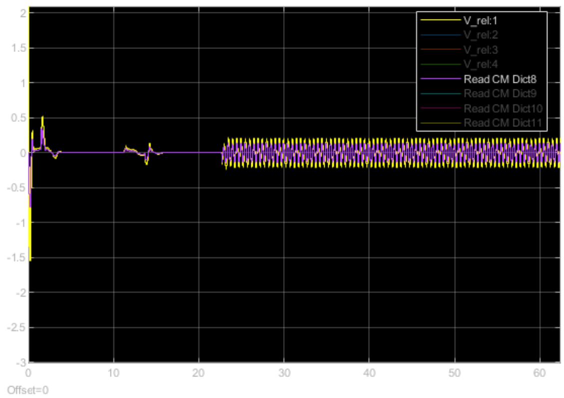

Figure D.7: Unloaded relative velocity between sprung mass and unsprung mass estimation (yellow) along with measured (purple)

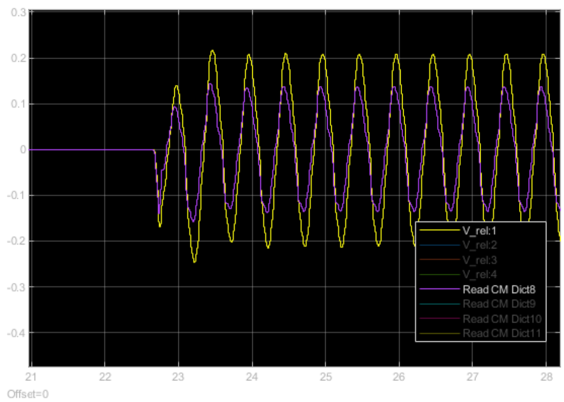

Figure D.8: Unloaded relative velocity between sprung mass and unsprung mass estimation (yellow) along with measured (purple)

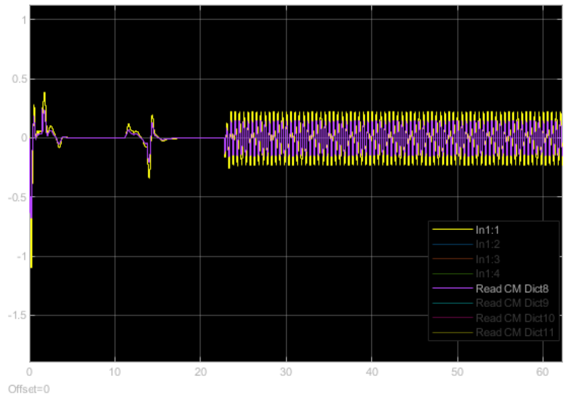

Figure D.9: 5,000 kg relative velocity between sprung mass and unsprung mass estimation (yellow) along with measured (purple)

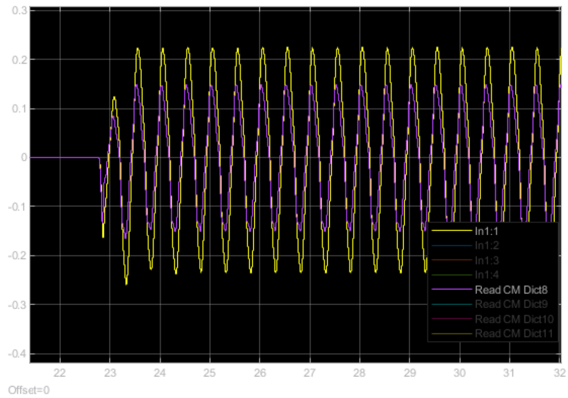

Figure D.10: 5,000 kg relative velocity between sprung mass and unsprung mass estimation (yellow) along with measured (purple)

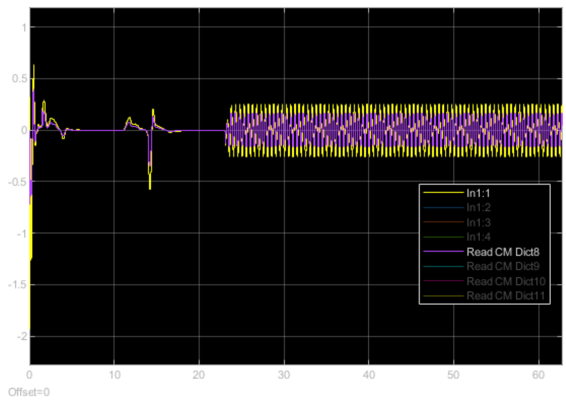

Figure D.11: 10,000 kg relative velocity between sprung mass and unsprung mass estimation (yellow) along with measured (purple)

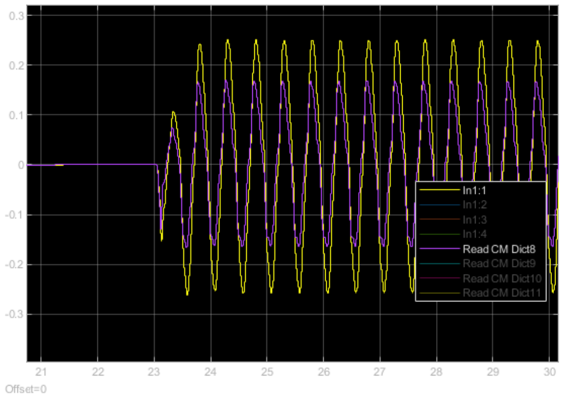

Figure D.12: 10,000 kg relative velocity between sprung mass and unsprung mass estimation (yellow) along with measured (purple)

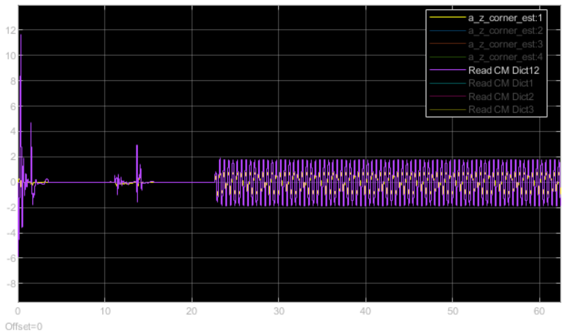

Figure D.13: Unloaded corner vertical acceleration estimation (yellow) along with measured (purple)

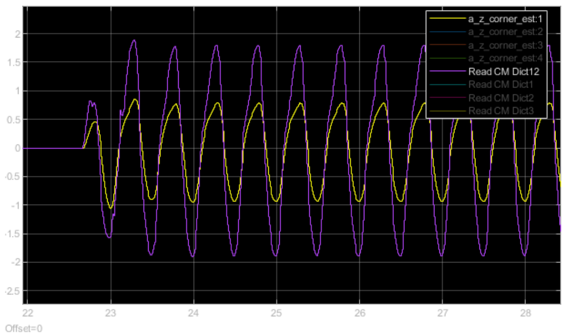

Figure D.14: Unloaded corner vertical acceleration estimation (yellow) along with measured (purple)

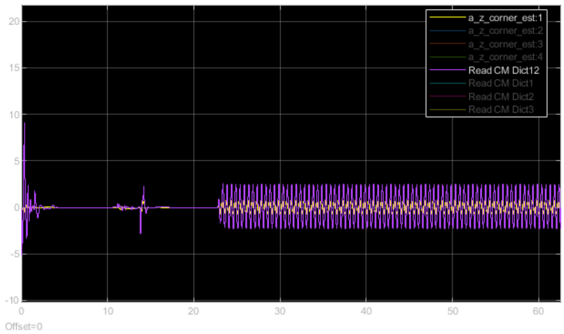

Figure D.15: 5,000 kg corner vertical acceleration estimation (yellow) along with measured (purple)

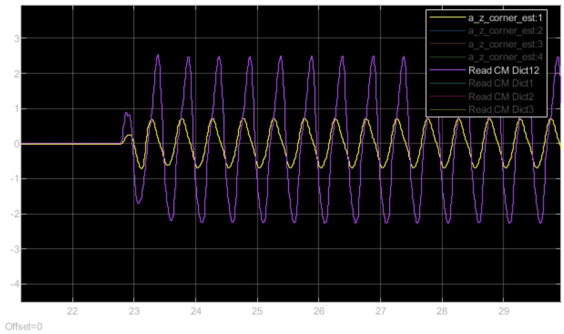

Figure D.16: 5,000 kg corner vertical acceleration estimation (yellow) along with measured (purple)

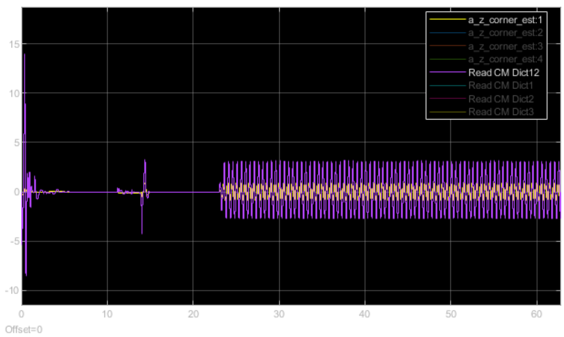

Figure D.17: 10,000 kg corner vertical acceleration estimation (yellow) along with measured (purple)

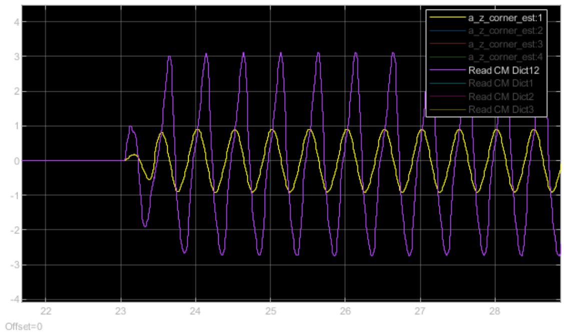

Figure D.18: 10,000 kg corner vertical acceleration estimation (yellow) along with measured (purple)

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY