




CHALMERS
UNIVERSITY OF TECHNOLOGY



Human Senses Mimicking on Virtual Public Roads

Master's Thesis Report

Master's thesis in Mobility Engineering

Himanshu Upendra Chaudhari

Prakash Raju Sridharraju

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden 2026

www.chalmers.se

MASTER'S THESIS IN MOBILITY ENGINEERING

Human Senses Mimicking on Virtual Public Roads

Himanshu Upendra Chaudhari
Prakash Raju Sridharraju



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mechanics and Maritime Sciences
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Human Senses Mimicking on Virtual Public Roads

© Himanshu Upendra Chaudhari, Prakash Raju Sridharraju, 2026.

Supervisor: Filip Lindström, Volvo Cars Corporation

Examiner: Petri Piiroinen, Department of Mechanics and Maritime Sciences

Master's Thesis 2026

Department of Mechanics and Maritime Sciences

Chalmers University of Technology

SE-412 96 Gothenburg

Sweden

Telephone +46 31 772 1000

Gothenburg, Sweden 2026

Human Senses Mimicking on Virtual Public Roads

Himanshu Upendra Chaudhari
Prakash Raju Sridharraju,
Department of Mechanics and Maritime Sciences
Chalmers University of Technology

Abstract

Automobile industry is continuously innovating and progressing towards autonomous solutions. An essential requirement for true vehicle autonomy is the capability to autonomously detect mechanical faults, as this function is critical for maintaining safety and ensuring that autonomous vehicles can operate independently, without an human intervention. As a part of a PhD research project aimed to develop a reliable diagnostic method concerning structural behavior of a vehicle and mitigate the malfunctions, this thesis is intended to investigate the influence of confounding variables on a set of pre-defined faults through multi-body simulations and signal processing technique. This project also incorporates the elements of machine learning framework to evaluate the performance metrics such as False Alarm Rate, Fault Detection Rate and Accuracy. A comparative study is done based on Score which is a harmonic mean of these performance metrics.

The two mechanical faults related to Anti-Roll Bar and Knuckle were induced. The confounding variables were shortlisted based on their influence and feasibility to replicate them in ADAMS Car. This study highlights the inherent trade-off between false alarm rate and fault detection capability in autonomous vehicle fault models, emphasizing the need to balance sensitivity and specificity for safe deployment. The analysis demonstrates that confounders affect fault detection performance differently depending on the component and fault type. Additionally, certain operating conditions can reduce the distinguishability between fault types, leading to similar performance outcomes. Importantly, incorporating confounders during model training significantly improves robustness by reducing false positives and enhancing sensitivity, enabling more reliable differentiation between true faults and external influences, an essential requirement for safe and effective autonomous vehicle operation.

Keywords: Local Rational Model, Multi-body Simulation, Frequency Response Function, Confounding Variables, Fault Detection, Transfer Function Estimation, Diagnostics, Logistic Regression

Acknowledgements

This master's thesis was carried out at the Department of Mechanics and Maritime Sciences, Chalmers University of Technology, in collaboration with Volvo Cars Corporation.

We express our sincere gratitude to our academic supervisor and examiner, Petri Piironen, for his support throughout this work. You encouraged us to work independently while always being approachable and available for discussion, providing valuable insights for entire course of the project.

We would also like to thank Volvo Cars and our industrial supervisors, Filip Lindström and Anders Nord, for providing us with the opportunity to work on this exciting project. Your continuous support, valuable feedback, and technical insights significantly contributed to the improvement of our work. We are also thankful to Mattias Robertson, Engineering Manager, Durability CV, RLD and Chassis Analysis for providing such great work culture. We also extend our gratitude to the members of the Road Load Durability (RLD) department at Volvo Cars — Kristian Carlsson, Suhail Jagwani, Sagar Parampalli Mahabaeshwar, Joakim Zakrisson, Rickard Petersson, and Anders Tengvall for their assistance and guidance in helping us become acquainted with the ADAMS Car software.

Finally, we are indebted to our family and friends for motivating us throughout our master's degree. Without them this wouldn't have been accomplished.

Himanshu Upendra Chaudhari, Prakash Raju Sridharraju, Gothenburg, June 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis.

ACC	Accuracy
ADAS	Advanced Driver Assistance Systems
ARB	Anti-Roll Bar
ARX	Auto-Regressive Exogenous Model
BL	Baseline condition
BP	Belgian Pave
C	Confounder
CC	Coarse Cobblestone
CF	Confounder with Fault
DDT	Dynamic Driving Task
DR	Detection Rate
DoF	Degree of Freedom
ETF _E	Empirical Transfer Function Estimate
F	Fault condition
FAR	False Alarm Rate
FRF	Frequency Response Function
F1-score	Harmonic mean of precision and recall
LRM	Local Rational Model
LUCA	Left Upper Control Arm
MA	Maude Ave
MIMO	Multiple Input Multiple Output
PA	Palo Alto
PSD	Power Spectral Density
PT	Pullinger Triangle
RC	Rural Concrete
RTL	Rig Test Load
SAE	Society of Automotive Engineers
SISO	Single Input Single Output

Nomenclature

Symbols	Description	Unit
Signals and System Representation		
$u(t)$	Input signal (road excitation)	–
$y(t)$	Output signal (vehicle response)	–
$u[n]$	Discrete-time input signal	–
$y[n]$	Discrete-time output signal	–
$g[k]$	Discrete-time impulse response of the system	–
n	Discrete-time index	–
k	Summation index	–
f	Frequency	Hz
ω	Angular frequency	rad/s
$G(\omega)$	Frequency Response Function (FRF)	–
$\hat{G}(f)$	Estimated frequency response	–
$U(f)$	Fourier transform of input signal	–
$Y(f)$	Fourier transform of output signal	–
Local Rational Model (LRM)		
$N(f)$	Numerator polynomial of local model	–
$D(f)$	Denominator polynomial of local model	–
L	Local frequency window length	Hz
T_{wl}	Time window length	s
$\hat{H}(f)$	Estimated MIMO FRF	–
$\hat{H}_{ij}(f)$	Estimated MIMO FRF from input j to output i	–
m, n	Model order of numerator and denominator	–
a_i, b_i	Complex coefficients of LRM	–

Symbols	Description	Unit
Feature Representation and Machine Learning		
$X \in \mathbb{R}^d$	Feature vector in d -dimensional space	–
d	Dimension of feature space	–
x_i	Individual feature component	–
y_i	Label (baseline/fault class)	–
β	Weight vector of logistic regression	–
β_0	Bias (intercept term)	–
z	Projected scalar feature ($z = \beta^T X + \beta_0$)	–
$p(x)$	Probability of fault condition	–
λ	Regularization parameter (L1 penalty)	–
Statistical Modeling		
Z	Projected random variable	–
Z_b	Distribution of baseline samples	–
Z_f	Distribution of faulty samples	–
μ	Mean of Gaussian distribution	–
σ	Standard deviation of Gaussian distribution	–
μ_b, μ_f	Mean of baseline and faulty distributions	–
σ_b, σ_f	Standard deviation of baseline and faulty distributions	–
$\mathcal{N}(\mu, \sigma^2)$	Normal (Gaussian) distribution	–
$\Phi(\cdot)$	Standard normal cumulative distribution function	–
Validation		
K	Number of folds in cross-validation	–
D_k	Dataset corresponding to fold k	–
D_{train}	Training dataset	–
D_{val}	Validation dataset	–
M_k	Performance metric for fold k	–
M_{final}	Averaged performance metric across folds	–

Contents

List of Acronyms	ix
Nomenclature	xi
List of Figures	xv
List of Tables	xix
1 Introduction	1
1.1 Background	1
1.2 Purpose	2
1.3 Research Questions	3
1.4 Limitations	4
2 Theory	5
2.1 Linear Time-Invariant model and frequency function	5
2.2 Estimating the frequency function	5
2.2.1 Local Rational Model	6
2.2.2 Multiple-Input Multiple-Output LRM	7
2.3 Feature Representation and Machine Learning	7
2.3.1 Clustering	7
2.3.2 Regression	9
2.3.2.1 Linear regression	9
2.3.2.2 Logistic Regression	9
2.3.2.3 L1 Regularized Logistic Regression	10
2.3.3 Geometric Interpretation	11
2.4 Validation	11
2.4.1 Statistical modeling: Gaussian distribution	11
2.4.2 Performance Metrics	12
2.4.2.1 False Alarm Rate and Detection Rate	13
2.4.2.2 Accuracy	13
2.4.2.3 F1-score	13
2.4.3 K-fold cross validation	14
3 Methods	16
3.1 Multi-body Simulation for Fault Detection and Diagnosis	16
3.1.1 Simulation Setup of XC90	16

3.1.2	Physical Testing	17
3.1.3	Correlation with physical testing	18
3.1.4	Public Roads	19
3.2	Fault Modeling	21
3.2.1	Bushings as joint replacements	22
3.2.2	ARB Fault Modelling in ADAMS	22
3.2.3	ARB Static Simulation of Fault Behaviour	23
3.2.4	Knuckle Fault Modelling in ADAMS	24
3.2.5	Knuckle Static Simulation of Fault Behaviour	25
3.3	Confounding Variables	26
3.3.1	Types of Confounders	27
3.3.2	Modeling of confounder	27
3.3.2.1	Trailer and Rim	27
3.3.2.2	Damper, RTL load and Speed	28
3.3.2.3	Tyre Pressure	28
3.3.2.4	Bushing Stiffness Due To Tolerance Effect	30
3.4	Sensor Configuration	31
3.5	Signal Processing Pipeline	31
3.6	Result Approach	32
4	Results	33
4.1	Fault Detection under Controlled Conditions	33
4.1.1	Effect of Confounding Variables on ARB Fault and Knuckle Fault Detections	35
4.1.1.1	Effect on ARB Fault Detection	35
4.1.1.2	Effect on Knuckle Fault Detection	40
4.2	Training and Improved Fault Detection under Confounding Conditions	44
4.2.1	ARB Training and Solution	44
4.2.2	Knuckle Training and Solution	48
5	Conclusion	53
	Bibliography	56
A	Appendix A	58
A.1	Road Excitations Different Speed	58
B	Appendix B	60
B.1	Training Results for ARB Fault	60
B.2	Training Results for Knuckle Fault	63
C	Appendix C	66
C.1	Validation Results for ARB Fault	66
C.2	Validation Results for Knuckle Fault	69
D	Appendix D	72
D.1	Solution Results for ARB Fault	72
D.2	Solution Results for Knuckle Fault	75

List of Figures

1.1	Mechanical health monitoring of car through frequency response. . . .	2
2.1	Illustration of LRM applied over multiple frequency windows to estimate the FRF.	6
2.2	Scatter split of baseline and fault.	8
2.3	Scatter overlap of baseline and fault.	8
2.4	Regression Classifier.	10
2.5	Gaussian Distribution for Separated Baseline and Fault.	12
2.6	Gaussian Distribution for overlapped Baseline and Fault.	12
3.1	Baseline: Volvo XC90 2024 ADAMS Model.	17
3.2	<i>Belgian Pave</i> road at Hallered Proving Ground.	17
3.3	Physical testing sensor locations (represented through ADAMS). . . .	18
3.4	Transfer Function comparison of baseline and faulty vehicle from ADAMS Car - <i>Belgian Pave</i> 40 Kmph.	18
3.5	<i>Belgian Pave</i> road in ADAMS Car.	19
3.6	Road Excitation(PSD) of different speed and direction Belgian pave Road on Left and Right, front Knuckle.	19
3.7	Public road surface types used in the study.	20
3.8	Road Excitation(PSD) of different Public Roads on Left and Right, front Knuckle.	20
3.9	Actual ARB and Knuckle faults induced during physical testing. . . .	21
3.10	Degrees of Freedom for ARB fault and Knuckle fault.	21
3.11	ARB joint replacement with linear bushing.	23
3.12	ARB joint replacement linear bushing stiffness curve.	23
3.13	ARB Static Simulation: Roll and Vertical force.	24
3.14	Knuckle joint replacement with linear bushing	25
3.15	Bushing translational and rotational stiffness characteristics.	25
3.16	ADAMS Static Simulation: Parallel Wheel Travel.	26
3.17	Knuckle fault validation: Time vs Relative displacement between Knuckle and Upper Control Arm.	26
3.18	XC90 with Trailer.	28
3.19	XC90 with 22in. Wheel rim (r22) configuration.	29
3.20	XC90 maximum load (RTL) variant.	29
3.21	Speed variation through Event Builder in ADAMS.	29
3.22	Soft and firm damper curves.	29
3.23	Bushing locations in front and rear suspension.	30

3.24	Example of bushing stiffness tolerance curves.	30
3.25	Signal processing workflow.	32
4.1	3-fold Validation Baseline Belgian pave and ARB Fault.	33
4.2	3-fold Validation Baseline Belgian pave and Knuckle Fault.	34
4.3	3-fold Validation Public Roads and ARB Fault.	34
4.4	3-fold Validation Public Roads and Knuckle Fault.	34
4.5	Belgian Pave against Public Roads as Confounder.	35
4.6	Public Roads 50 Kmph against Public Roads 30 Kmph.	36
4.7	Public Road against 270kpa tyre Pressure Confounder.	37
4.8	Validation: Public Road against Positive tolerance bushing.	37
4.9	Effect of All Confounding Variables on ARB fault validation.	38
4.10	Effect of Public Roads on Confounders with ARB fault Detection.	39
4.11	Public road against maximum load (RTL) confounder.	40
4.12	Public road against stiff damper confounder.	41
4.13	Public road against 270 kpa tyre pressure confounder.	41
4.14	Public road against negative bushing tolerance confounder.	42
4.15	Effect of all confounding variables on knuckle fault.	43
4.16	Effect of public roads on confounders with knuckle fault detection.	43
4.17	Training: 3-fold, public road against 270kpa tyre pressure with ARB fault.	45
4.18	Training: 3-fold, public road against positive tolerance bushing with ARB fault.	45
4.19	Effect of all confounding variables on ARB fault during training for robustness.	46
4.20	Solution: 3-fold, public road against 270kpa tyre pressure with ARB fault.	47
4.21	Solution:: 3-fold, public road against positive tolerance bushing with ARB fault.	47
4.22	Effect of all confounding variables on ARB fault after adding robustness.	48
4.23	Training: 3-fold, public road against RTL with knuckle fault.	49
4.24	Training: 3-fold, public road against stiff damper with knuckle fault.	49
4.25	Effect of all confounding variables on knuckle fault during training for robustness.	50
4.26	Solution: 3-fold, public road against RTL with knuckle fault.	51
4.27	Solution: 3-fold, public road against stiff damper with knuckle fault.	51
4.28	Effect of all confounding variables on knuckle fault after adding ro- bustness.	52
5.1	Solution F1-score comparison between ARB and Knuckle fault.	54
5.2	F1-score comparison between training, validation and solution for ARB and Knuckle fault.	54
A.1	PSD of road excitation at 30 kmph, 40 kmph, and 50 kmph for Maude Ave.	58
A.2	PSD of road excitation at 30 kmph, 40 kmph, and 50 kmph for Palo Alto.	58

A.3	PSD of road excitation at 30 kmph, 40 kmph, and 50 kmph for Pullinger Triangle.	58
A.4	PSD of road excitation at 30kmph, 40kmph and 50kmph for Rural Concrete.	59
A.5	PSD of road excitation at 40kmph and 50kmph for Coarse Cobblestone.	59
B.1	3-fold Training Public Road 30kmph Confounder and ARB fault.	60
B.2	3-fold Training Trailer confounder and ARB fault.	60
B.3	3-fold Training 22inch Wheel confounder and ARB fault.	61
B.4	3-fold Training 240 Kpa Tyre pressure confounder and ARB fault.	61
B.5	3-fold Training RTL confounder and ARB fault.	61
B.6	3-fold Training negative tolerance bushing confounder and ARB fault	62
B.7	3-fold Training soft damper confounder and ARB fault.	62
B.8	3-fold Training stiff damper confounder and ARB fault.	62
B.9	3-fold Training 270 Kpa tyre pressure confounder and Knuckle fault.	63
B.10	3-fold Training 22inch rim confounder and Knuckle fault.	63
B.11	3-fold Training Positive Bushing Tolerance Confounder and Knuckle Fault.	63
B.12	3-fold Training 240 kPa Tyre Pressure confounder and Knuckle fault.	64
B.13	3-fold Training soft damper and Knuckle fault.	64
B.14	3-fold Training 30 Kmph and Knuckle fault.	64
B.15	3-fold Training trailer and Knuckle fault.	65
B.16	Training: 3-fold, Public Road against negative bushing tolerance with knuckle fault.	65
C.1	Public Road training data.	66
C.2	Validation: Public Road against 240Kpa tyre pressure.	66
C.3	Validation: Public Road against negative tolerance bushing.	67
C.4	Validation: Public Road against maximum load (RTL).	67
C.5	Validation: Public Road against soft damper.	67
C.6	Validation: Public Road against stiff damper.	68
C.7	Validation: Public Road against trailer.	68
C.8	Validation: Public Road against 22inch wheels.	68
C.9	Validation: Public Roads against 22in. rim.	69
C.10	Validation: Public Roads against 30 kmph.	69
C.11	Validation: Public Roads against 240 kPa tyre pressure.	69
C.12	Validation: Public Roads against positive bushing tolerance.	70
C.13	Validation: Belgian Pave against Public Roads.	70
C.14	Validation: Public Roads against soft damper.	70
C.15	Validation: Public Roads against trailer.	71
D.1	3-fold Solution Public Roads and ARB Fault.	72
D.2	3-fold Solution Public Roads 30kmph and ARB Fault.	72
D.3	3-fold Solution Public Roads 22inch Wheel and ARB Fault.	73
D.4	3-fold Solution Public Roads Trailer and ARB Fault.	73
D.5	3-fold Solution Public Roads 240kpa Tyre Pressure and ARB Fault.	73

D.6 3-fold Solution Public Roads Negative Tolerance Bushing and ARB Fault.	74
D.7 3-fold Solution Public Roads RTL and ARB Fault.	74
D.8 3-fold Solution Public Roads Soft Damper and ARB Fault.	74
D.9 3-fold Solution Public Roads Stiff Damper and ARB Fault.	75
D.10 3-fold Solution Public Roads 270kpa Tyre Pressure and Knuckle Fault.	75
D.11 3-fold Solution Public Roads 22in. Rim and Knuckle Fault.	75
D.12 3-fold Solution Public Roads 30kph Speed and Knuckle Fault.	76
D.13 3-fold Solution Public Roads 240 kPa Tyre Pressure and Knuckle Fault.	76
D.14 3-fold Solution Public Roads and Knuckle Fault.	76
D.15 3-fold Solution Public Roads Trailer and Knuckle Fault.	77
D.16 3-fold Solution Public Roads Negative Bushing Tolerance and Knuckle Fault.	77

List of Tables

3.1	XC90 2024 baseline parameters at vehicle and subsystem level.	17
3.2	Detailed Characteristics of Public Road Profiles.	20
4.1	Effect of All Confounding Variables on ARB fault validation.	38
4.2	Effect of Public Roads on 270kpa Tyre Pressure and Stiff Damper Confounders with ARB fault Detection.	39
4.3	Effect of all confounding variables on knuckle fault validation.	42
4.4	Effect of public roads on RTL and stiff damper confounders with knuckle fault detection.	44
4.5	Effect of all confounding variables on ARB fault during training for robustness.	46
4.6	Effect of all confounding variables on ARB fault after adding robustness.	48
4.7	Effect of all confounding variables on knuckle fault during training for robustness.	50
4.8	Effect of all confounding variables on knuckle fault after adding ro- bustness.	52

1

Introduction

1.1 Background

Volvo Cars Corporation aims to create a future with zero collisions by continuously improving vehicle safety through advanced technologies. To support this, the company is investing on the development of Advanced Driver Assistance Systems (ADAS) and autonomous driving technologies [1].

Vehicles could encounter mechanical faults into their subsystems, such as suspension, steering, brakes, chassis etc. After certain period of time, which affect their ride comfort and overall handling, making them unsafe to drive. These faults can be in the form of but not limited to worn bushings, loose ball joints, degraded dampers, structural cracks etc. For manually driven vehicles, such faults can be intuitively identified by the driver by experiencing unpleasant noise or vibration inside the vehicle. However, for autonomous vehicles, they must mimic the perception of 'Human Senses' concerning mechanical integrity. In other words, the faults must be identified by the vehicle itself as it can be dangerous if remained undetected. According to SAE J3016, at higher levels of automation (Levels 4 and 5), the Automated Driving System (ADS) is responsible for performing both the Dynamic Driving Task (DDT) and the DDT fallback; consequently, the system is expected to bring the vehicle to a Minimal Risk Condition (MRC) when the trip cannot be continued [2]. This means, the vehicle must be able to protect itself and others by carrying out fallback and reaching a safe state (MRC) because there may not be a human ready or able to intervene in time. Hence, vehicle must monitor its mechanical health continuously. To achieve this, a fault detection system is incorporated into the autonomous vehicle to ensure safe and reliable operation.

One effective approach for monitoring mechanical health is to develop a fault-free baseline model that accurately represents the behavior of a healthy vehicle. For training, some fault data is also required to fix the detection thresholds. This sets a benchmark for a faulty vehicle to detect abnormalities without activating the false alarms. The one method to monitor a vehicle's mechanical health is through its frequency response, which shows how the vehicle reacts to vibrations at different frequencies. A healthy vehicle does not exhibit a single fixed frequency response pattern; instead, its response varies with operating conditions such as speed and road profile. A baseline model therefore captures this range of normal behavior whereas, mechanical faults introduce additional consistent deviations beyond this

normal variability, enabling reliable fault detection. To evaluate the frequency response, accelerometers are placed at strategic locations to capture input signals, such as road excitations from wheels (unsprung mass) and output signals, such as excitations of vehicle main frame (sprung mass). The frequency response can be estimated using various system identification methods. For instance, Auto-Regressive Exogenous Input (ARX) models are primarily time-domain techniques but can be used to derive the frequency response indirectly. In contrast, methods such as the Empirical Transfer Function Estimate (ETF) and the Local Rational Model (LRM) are more directly suited for frequency-domain analysis [3].

Sensor noise and leakage cause variance and systematic errors in the estimated Frequency Response Function (FRF) [4]. ETFE suffers from leakage due to DFT and non-periodic signals [5]. Additionally, the computational complexity of LRM is favorable as compared to the ARX [3]. Overall, LRM reduces leakage and improves robustness [6]. For this purpose, LRM method is used in our work. Further explanation regarding LRM method is mentioned in the Section 2.2.1.

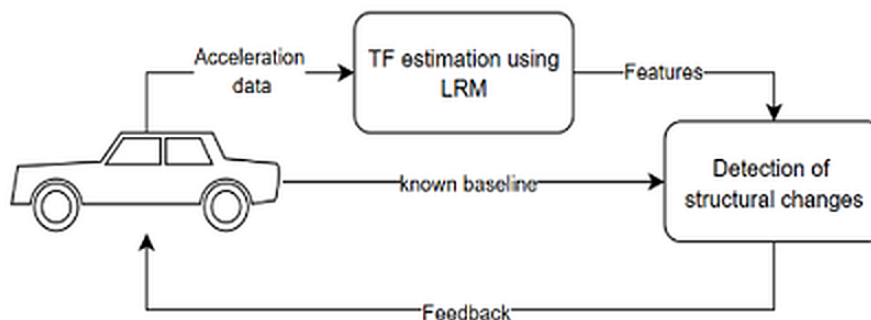


Figure 1.1: Mechanical health monitoring of car through frequency response.

The above Figure 1.1 illustrates the workflow of a vibration-based structural health monitoring system for a vehicle. Acceleration data is first collected from the vehicle and used to estimate its transfer function using the LRM, which captures the vehicle’s frequency response. From this, features are extracted and compared with a known baseline representing healthy behavior under normal operating conditions. Any significant deviation from this baseline is identified as a potential structural fault. A feedback mechanism then communicates these anomalies through alerts, warnings, and diagnostic signals to the driver or onboard vehicle system, enabling timely response and condition monitoring. Such a framework can be further extended with data-driven methods, such as logistic regression, to enhance fault classification and decision-making capabilities.

1.2 Purpose

With respect to fault detection, confounding variables are factors that introduce natural variations in the vehicle’s behavior within acceptable limits, influencing its

response without indicating a fault. However, these variations can affect sensor signals and lead to false alarms. Frequent false alarms reduce trust in the diagnostic system, causing users to potentially ignore or not respond to actual warnings. These confounding variables can be categorized under different levels - Environment Level, Vehicle Level, Subsystem Level and lastly at Component Level. For example, the environmental factor such as temperature can alter the mechanical properties of bushings and joints as they may become stiffer in cold conditions and softer in warm conditions. At vehicle level, load distribution between front and rear can deflect suspension travel and system natural frequencies. At subsystem level, tyre pressure variability will affect tyre stiffness and its natural frequency. At component level, we can regard possible sensor bias as one of the confounding variables. All these factors can cause variations in vibration behavior without any actual damage.

The purpose of this thesis is to investigate how confounding variables influence the detectability of predefined vehicle faults. In real driving conditions, vehicle behavior is influenced by many factors (confounding variables) besides the fault itself, such as road conditions, vehicle speed, loading, and variations between the vehicles etc. These factors can make fault detection more difficult by altering the signals that are normally used to identify abnormal behavior.

Since, naturalistic data of vehicles representing its dynamic condition is difficult to obtain and moreover, the physical testing with damaged components is limited by safety and repeatability, we rely on simulation-based methods to extract signals and build a decision-making logic (fault or no-fault) based on it. Through multi-body vehicle simulations in ADAMS Car on scanned public roads, this thesis aims to better understand how these external and operational variations influence fault detectability and diagnostic performance. By combining simulation, signal processing and statistical analysis, the study aims to contribute to the development of more robust and dependable vehicle diagnostic systems that can operate effectively in realistic driving environments.

1.3 Research Questions

This section tells about the key research questions that guide the direction and focus of the work, highlighting the main aspects to be investigated. These questions are made to address the core challenges and objectives identified in the analysis and they are follows:

- How different confounding variables affect the detectability and which confounding variable has the most influence?
- How can we mitigate confounding influence by improving diagnostic methods?
- Which Road type makes it easy/challenging to identify fault?

1.4 Limitations

This section outlines the key limitations of the study, highlighting the factors that may influence the results and the assumptions made during the analysis. Understanding these limitations is important for interpreting the findings and assessing the scope and applicability of the results and they are as follows:

- The results depend on how accurate and realistic the data is generated.
- The output can change depending on where the sensor is placed and will change the frequency response accordingly.
- The model assumes linear behavior, but real systems can be nonlinear.
- Supervised learning approach by manually labeling data into categories (e.g., baseline, faulty, different fault types), and incorrect labels can cause the model to learn wrong patterns and produce inaccurate results.

2

Theory

2.1 Linear Time-Invariant model and frequency function

The relationship between the input and output signals of a dynamic system can be modelled as a discrete-time Linear Time-Invariant (LTI) system. The system is assumed to be causal and linear, and the signals are sampled at uniform time intervals. The output signal $y[n]$ is expressed as [3]:

$$y[n] = \sum_{k=0}^{\infty} g[k] u[n - k] \quad (2.1)$$

where $u[n]$ is the input signal, while n and k are discrete-time index and summation index, respectively. Eq. 2.1 above shows that the output is a weighted sum of present input signal $u[n]$ and past input $u[n - k]$ values. The impulse response $g[k]$ fully characterizes the system in the time domain.

To analyse and understand the system behavior more conveniently, systems are often represented in the frequency domain. The frequency response function is obtained by applying the Discrete-Time Fourier Transform (DTFT) to the impulse response of a causal system and is given by Eq. 2.2 [3]:

$$G(\omega) = \sum_{k=0}^{\infty} g[k] e^{-j\omega k} \quad (2.2)$$

The function $G(\omega)$ is a complex-valued function that describes how the system responds to sinusoidal inputs at different frequencies.

2.2 Estimating the frequency function

In practical applications, the exact model of the system is generally unknown. Therefore, the frequency response function $G(\omega)$ must be estimated from measured input-output data, specifically by comparing the frequency content of the output to that of the input, which can be expressed as:

$$G(\omega) \approx \frac{Y(\omega)}{U(\omega)} \quad (2.3)$$

where, $Y(w)$ and $U(w)$ are Fourier Transform of output and input respectively.

Usually, the estimation is affected by noise and measurement errors. Therefore, techniques such as LRM are used to improve the accuracy of the estimated FRF.

2.2.1 Local Rational Model

The LRM is a frequency-domain system identification technique used to estimate the FRF. This method is used to study how a system behaves at different frequencies. Instead of trying to model the whole system at once, it looks at small sections of the frequency range one at a time. In each small section (called a frequency window), it fits a mathematical model (a low-order fraction of polynomials) to the nearby data points.

The general idea of LRM is shown below in Figure 2.1:

- Choose DFT lines for which frequencies we want to estimate the frequency function
- For each DFT line, approximate the frequency response function in a local frequency window width around that line using a low-order rational function given in the below Eq2.4

$$G(f) = \frac{N(f)}{D(f)} \quad (2.4)$$

where $N(f)$ and $D(f)$ are low-order polynomials of frequency f .

- For each DFT line, the coefficients of the numerator and denominator polynomials, defined by a selected model order, are estimated using input–output data within a local frequency window around the DFT index. Least Squares method is used.
- For each DFT line, compute the frequency function estimate.




Figure 2.1: Illustration of LRM applied over multiple frequency windows to estimate the FRF.

Based on the input parameters, also called as Hyper-Parameters, such as Frequency Window Width (L), Time Window Length (T_{wl}), and Local Model Order (n), the LRM outputs discrete FRF, a complex (magnitude and phase) estimate at each DFT frequency line.

2.2.2 Multiple-Input Multiple-Output LRM

The Local Rational Model, originally formulated as Single-Input Single-Output (SISO) system extends to Multiple-Input Multiple-Output (MIMO) system [7] by treating the system response as a matrix of transfer function. Hence, the MIMO FRF can be described as below:

$$\hat{H}(f) = \begin{bmatrix} \hat{H}_{11}(f) & \cdots & \hat{H}_{1j}(f) \\ \vdots & \ddots & \vdots \\ \hat{H}_{i1}(f) & \cdots & \hat{H}_{ij}(f) \end{bmatrix} \quad (2.5)$$

where each element $\hat{H}_{ij}(f)$ represents the estimated transfer function from input j to output i .

2.3 Feature Representation and Machine Learning

Traditional fault detection methods depend on predefined physical models, but in real-world automotive systems, the presence of confounding factors, such as varying driving conditions, road profiles, and other factors, makes it difficult to define consistent faults. Some systems behave differently at different speeds, and noise overlaps with fault effects, so the fault is not clearly separable in raw measurements. Therefore, instead of manually defining rules, machine learning is used to learn patterns directly from data, enabling the separation of faulty and healthy behaviour even under varying conditions, i.e., confounder effects.

The proposed method follows a supervised learning approach [8] where the data sample is labelled as baseline and faulty. This label enables the model to learn the separating normal system behaviour and faulty behaviour.

Before applying supervised learning, the features in space are analysed using scatter representation, which reveals that the data naturally form clusters corresponding to different system conditions. This observation motivates the use of classification models to separate these clusters.

2.3.1 Clustering

Analysing acceleration signals makes things complex, and in the time domain its noisy and not well interpretable, so to enable pattern recognition, the raw time-domain signals are transformed into a feature space using frequency domain representation obtained from the LRM. Each time a window of data is mapped to a

2. Theory

high-dimensional feature vector, representing the transfer function across multiple frequencies and input and output channels.

$$X_n \in \mathbb{R}^d \quad (2.6)$$

If the features perfectly separated the classes, simple clustering techniques such as K-means would be sufficient to distinguish between baseline and faulty behaviour. Figure 2.2 and 2.3 shows 3D plot of features i.e 3 features.




Figure 2.2: Scatter split of baseline and fault.

However, the clusters overlap and cannot be separated by simple unsupervised methods; a supervised classification approach is used. By utilising labelled data, the model can learn an optimal decision boundary that maximises separation between baseline and faulty conditions.




Figure 2.3: Scatter overlap of baseline and fault.

2.3.2 Regression

Regression concept in statistical learning aims to model the relationship between input and output variables. Given a set of observations, regression methods seek to learn a function that describes how the output depends on the inputs, which can be generally expressed as $y = f(X) + \varepsilon$, where $X \in \mathbb{R}^d$ represents the input feature vector (as defined in Eq. 2.6), y denotes the output variable, $f(X)$ is the underlying mapping, and ε represents the noise term.

In the context of fault detection, the aim is to establish a relationship between the extracted feature representation and the system condition (baseline or faulty). This enables prediction of the system state for unseen data.[9].

2.3.2.1 Linear regression

One of the simplest forms of regression is linear regression, where the output is modelled as a linear combination of input features.

$$Z = \beta^T X + \beta_0 \quad (2.7)$$

where $\beta \in \mathbb{R}^d$ is the coefficient vector, and β_0 is the intercept. Linear regression defines a hyperplane in the feature space. In two dimensions this corresponds to a line, in three dimensions to a plane, and in higher dimensions to a hyperplane. In the context of feature representations, the model attempts to find a linear relationship that best fits the data.

Although linear regression is simple and intuitive, it is not suitable for classification problems where the output variable is discrete. The predicted output is continuous in $(-\infty, +\infty)$ and no probabilistic interpretation, cannot define classes(fault vs normal). Therefore, a model that can map the linear relationship into probabilistic representation is required for classification tasks [9].

2.3.2.2 Logistic Regression

Logistic regression extends linear models to classification problems by transforming the linear output into a probability to address the limitations of linear regression, using sigmoid function and using Eq. 2.7.

$$P(y = 1 | X) = \frac{1}{1 + e^{-z}} \quad (2.8)$$

where Z is defined in Eq. 2.7. The sigmoid function maps the linear combination to a probability value between 0 and 1, enabling probabilistic interpretation of class members.

However, directly minimising a quadratic loss on this nonlinear function results in a nonconvex optimisation problem, which may result in unstable parameter estimation. To address this, a logistic regression is trained using the principle of maximum likelihood estimation.

The likelihood is formulated based on the probabilistic model, and its logarithm leads to the cross-entropy loss function, which is convex and ensures a unique global optimum [9].

Extending Eq. 2.7 to the case of multiple independent features and combining it with the sigmoid function (Eq. 2.8), the relationship can be expressed as:

$$\ln \left(\frac{P(y = 1 | X)}{1 - P(y = 1 | X)} \right) = \beta_0 + \beta^T X \quad (2.9)$$

where $\beta^T X = \sum_{i=1}^d \beta_i X_i$ representing the contribution of multiple independent features.

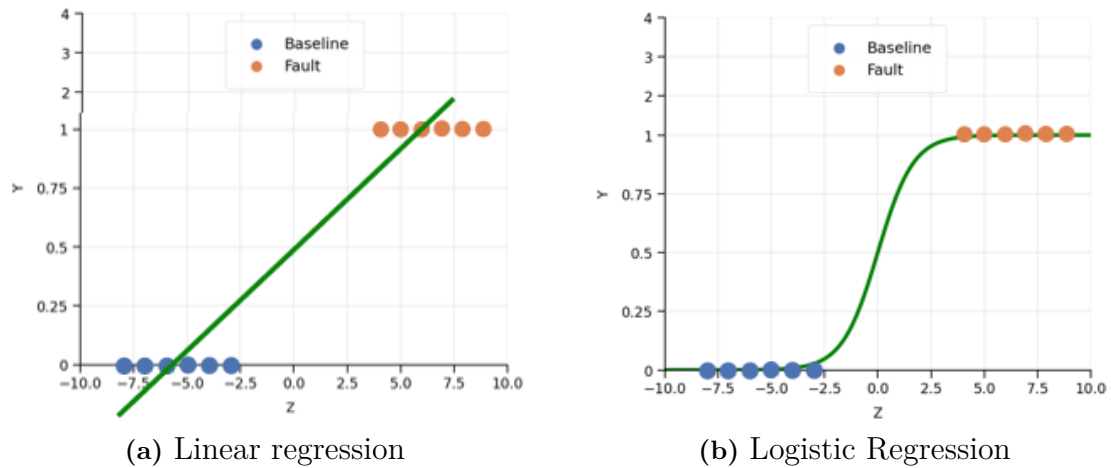


Figure 2.4: Regression Classifier.

For illustrative purposes, the difference between linear regression and logistic regression is shown in Figure 2.4. Figure 2.4b shows how the sigmoid function maps the linear combination enabling probabilistic interpretation of classes.

2.3.2.3 L1 Regularized Logistic Regression

In high-dimensional feature spaces, the model may become overly complex, leading to overfitting and reduced generalisation performance. To achieve better generalisation and interpretability, L1 regularisation is required, resulting in a sparse model that focuses on the most relevant features for fault detection.

This regularization promotes sparsity by shrinking coefficients toward zero [10], resulting in automatic feature selection by retaining only the most informative features while eliminating redundant inputs.

The resulting optimization problem is defined as:

$$\min_{\beta} \left[- \sum_{i=1}^N (y_i \log p_i + (1 - y_i) \log(1 - p_i)) + \lambda \sum_{j=1}^d |\beta_j| \right] \quad (2.10)$$

where $p_i = P(y_i = 1 | X_i)$ denotes the predicted probability for sample i , β is the coefficient vector defined in Eq. 2.7, and λ is the regularization parameter controlling the degree of sparsity [10].

2.3.3 Geometric Interpretation

The Regression model can be interpreted as defining a linear decision boundary in the high-dimensional feature space. The model computes a linear combination of features from Eq. 2.7. This formulation projects the high-dimensional features vector onto a one-dimensional space defined by β . As a result, each sample is represented by a scalar value of Z , simplifying the classification problem.

$$Z > 0 (\textit{Faulty}) \quad , \quad Z < 0 (\textit{Baseline})$$

The vector β defines the direction used to project the feature vectors onto a 1D score Z . The logistic regression model determines this direction by maximizing the likelihood of the observed class labels.. When L1-regularisation is applied, this direction is determined using only a subset of relevant features. The reduced one-dimensional representation plays an important role in the validation stage, when the statistical distribution of Z is analysed to evaluate detection performance.

Then the projected variable z enables the use of statistical modeling approaches, such as Gaussian distributions, to quantify the separation between baseline and faulty conditions (Figure 2.5 and 2.6).

2.4 Validation

Validation plays an important role in statistical and machine learning approaches when handling large datasets. It is used to assess how well the results of a model generalise to independent data sets, therefore helping to identify issues such as over-fitting and selection bias.

In predictive modelling, the primary goal is to evaluate how accurately a trained model will perform on unseen data. Validation techniques provide a systematic framework for estimating this predictive performance. They also allow assessment of the quality of a fitted model, across different datasets.

2.4.1 Statistical modeling: Gaussian distribution

To quantify the separability between baseline and faulty samples, the projected values from Eq (2.7) are modelled using a Gaussian distribution [9].

$$Z \sim N(\mu, \sigma^2) \quad (2.11)$$

$$Z_b \sim N(\mu_b, \sigma_b^2) \quad , \quad Z_f \sim N(\mu_f, \sigma_f^2) \quad (2.12)$$

The mean μ and standard deviation σ describe the location and spread of the data in the projected space for baseline and faulty conditions. A clear separation between these distributions indicates good fault detection, whereas overlap leads to misclassifications or is challenging to classify clearly.

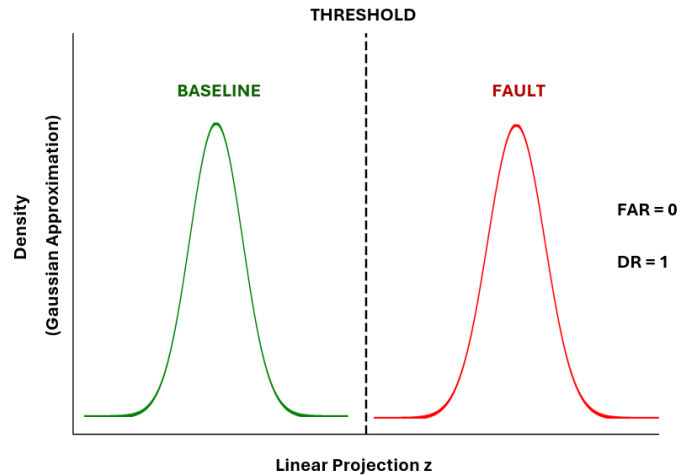


Figure 2.5: Gaussian Distribution for Separated Baseline and Fault.

In real applications, the decision boundary does not perfectly separate the data; instead, both baseline and faulty samples exhibit statistical variability around this boundary.

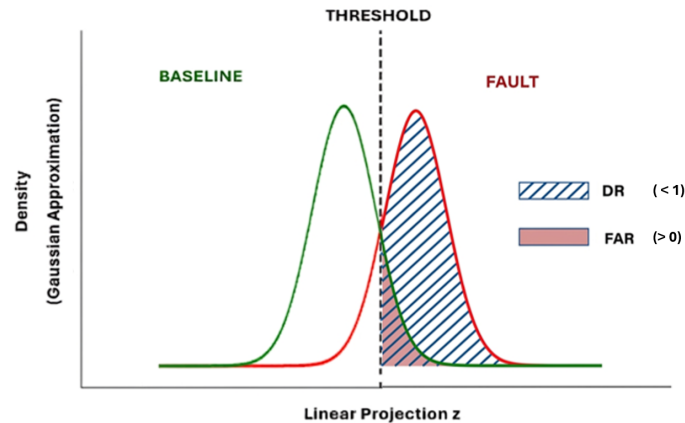


Figure 2.6: Gaussian Distribution for overlapped Baseline and Fault.

The extent of overlap between these distributions determines the probability of misclassification and directly impacts the performance of the detection.

2.4.2 Performance Metrics

The statistical distributions obtained from Gaussian modelling and KDE allows quantitative evaluation of the classification performance. In particular, the overlap

between baseline and faulty distributions directly determines the probability of misclassification.

The performance evaluation in this work is primarily based on detection-oriented metrics such as false alarm rate and detection rate, which are commonly used in fault detection systems.

2.4.2.1 False Alarm Rate and Detection Rate

The False Alarm Rate (FAR) represents the probability of incorrectly classifying a baseline sample as fault, and Detection Rate (DR) represents the probability of correctly identifying a fault condition. In terms of the projected variable Z , FAR can be defined as the probability that a baseline sample exceeds the decision threshold T , i.e., $\text{FAR} = P(z > T \mid \text{Baseline})$, whereas DR corresponds to the probability that a faulty sample exceeds the same threshold, i.e., $\text{DR} = P(z > T \mid \text{Fault})$, (Figure 2.5 and 2.6).

Under the Gaussian assumption:

$$\text{FAR} = 1 - \Phi\left(\frac{t - \mu_{\text{baseline}}}{\sigma_{\text{baseline}}}\right) \quad (2.13)$$

$$\text{DR} = 1 - \Phi\left(\frac{t - \mu_{\text{fault}}}{\sigma_{\text{fault}}}\right). \quad (2.14)$$

where $\Phi(\cdot)$ represents the cumulative distribution of the standard normal distribution.

2.4.2.2 Accuracy

Accuracy represents the proportion of correctly classified samples and is defined as:

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + TN + FN}, \quad (2.15)$$

where TP, FP, TN and FN denote true positives, false positives, true negatives and false negatives, respectively. In practice, accuracy is computed as the ratio of correctly predicted samples to the number of samples.

2.4.2.3 F1-score

The F1-score is a widely used performance metric in classification tasks that provides a balance between precision and recall. Precision represents the proportion of correctly identified fault samples among all predicted faults, while recall (or detection rate) measures the proportion of actual faults that are correctly identified [11]. The F1-score is defined as:

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}, \quad (2.16)$$

In the context of fault detection, recall corresponds to the DR, while precision is influenced by the FAR. Therefore, the F1-score provides a combined measure of detection capability and false alarm suppression.

Unlike accuracy, which may be dominated by class imbalance, the F1-score provides the trade-off between correctly detecting faults and avoiding false alarms. This makes it particularly suitable for evaluating performance under varying confounding conditions.

In this work, the F1-score is used to quantify the overall reduce in classification performance due to confounding variables. Lower F1-scores indicate increased overlap between baseline and faulty conditions, while higher in the F1-score reflect better separation and robustness. Consequently, the F1-score serves as an indicator for identifying critical confounders that significantly impact fault detection performance.

2.4.3 K-fold cross validation

To evaluate the robustness and generalisation capability of the proposed fault detection method, K-fold cross validation is used. Cross-validation is a widely used technique in statistical learning for assessing how well a model generalises to unseen data [8],[12].

In K-fold cross validation, the dataset is divided into K mutually exclusive subsets (folds). The model is trained on $K - 1$ folds and validated on the remaining fold. This process is repeated K times, such that each fold is used once as the validation set [8].

Let the dataset be partitioned as $D = \{D_1, D_2, \dots, D_K\}$. For each fold $k = 1, 2, \dots, K$, the subset D_k is used as the validation set, while the remaining subsets $\{D_1, \dots, D_{k-1}, D_{k+1}, \dots, D_K\}$ are used for training..

In this Thesis, a 3-fold cross-validation is used, where the dataset is divided into three segments. Each segment corresponds to data collected under different operating conditions, ensuring that the model is validated across varying conditions.

Aggregation of Performance Metrics

For each fold, the performance metrics such as FAR, DR, and accuracy are computed. The final performance is obtained by averaging across all folds:

$$M_{\text{final}} = \frac{1}{K} \sum_{k=1}^K M_k \quad (2.17)$$

where M_k represents the value of the performance metric (e.g., FAR, DR, and accuracy) computed on the K_{th} fold.

This approach provides a reliable estimate of the model's performance, reduces the likelihood of overfitting, and provides a reliable estimate of metrics. It ensures that the model is not biased toward a specific subset of the data and performs consistently across different operating conditions [12].

K-fold cross validation enables systematic evaluation of the metrics of the proposed fault detection method.

3

Methods

3.1 Multi-body Simulation for Fault Detection and Diagnosis

Multibody dynamics (MBD) models a system as multiple rigid and/or flexible bodies connected by joints and forces. MSC ADAMS is a widely used multibody dynamics simulation tool for analysing mechanical systems. ADAMS/Car is a specialised module within ADAMS focused on vehicle dynamics, enabling simulation of suspension systems, steering, drivetrain, full vehicle analysis, etc.

Fault detection and diagnosis in vehicle vertical dynamics relies on identifying deviations in system response from nominal behaviour [13]. This requires an accurate representation of the underlying physical system, as faults manifest through changes in dynamic characteristics.

Vehicle suspension systems exhibit strongly coupled, nonlinear behaviour involving multiple interacting components such as dampers, bushings, and structural elements. These interactions cannot be adequately captured using simplified or lumped models.

Therefore, MBD software is used to model the vehicle system, enabling realistic simulation of component interactions, nonlinearities, and load transfer mechanisms. This provides a physically consistent basis for generating data used in fault detection and classification.

3.1.1 Simulation Setup of XC90

Multibody simulation of baseline XC90 2024 using ADAMS/Car, as shown below in Figure 3.1 was performed to simulate its behavior under varying operating conditions, confounders and pre-defined mechanical faults - Anti Roll Bar fault and knuckle fault which are explained in Section 3.1.4, Section 3.3, and Section 3.2 respectively. The baseline model is defined as the final production configuration of the vehicle. Any mechanical fault or deviation from this configuration, such as changes in tyres, rims, drive modes, or suspension settings, categorizes the vehicle as a non-baseline model. Some of the baseline parameters at the vehicle and subsystem level, which are important for this thesis, are mentioned in Table 3.1 below.




Figure 3.1: Baseline: Volvo XC90 2024 ADAMS Model.

Table 3.1: XC90 2024 baseline parameters at vehicle and subsystem level.

Vehicle Baseline Parameters	Value
Wheel Rim Size	21 inches
Front Tyre Pressure	300 kpa
Rear Tyre Pressure	320 kpa
Loading	<i>curb2</i> (Driver and Passengers)
Damper Setting	Comfort

3.1.2 Physical Testing

The objective was to correlate the simulated results with physical testing results obtained from prior work conducted at Hallered Proving Ground before the commencement of this thesis as shown in Figure 3.2. Initially, simulations were performed on scanned *Belgian Pave* road with both baseline model and faulty models (ARB fault and knuckle fault) at three different speeds - 30, 40 and 50 Kmph.



Figure 3.2: *Belgian Pave* road at Hallered Proving Ground.

For testing, triaxial accelerometers were used for vibration measurements on the front suspension system. A total of four accelerometers were installed in symmetric left-right positions with input location at knuckle and output location at subframe, shown in Figure 3.3.




Figure 3.3: Physical testing sensor locations (represented through ADAMS).

The knuckle is the primary interface between the Wheel and the suspension system, and thus captures the excitation originating from road–tyre interaction, whereas subframe serves as the structural connection between the suspension and the vehicle body. These sensors measure the transmitted vibrations after they propagate through the suspension system.

3.1.3 Correlation with physical testing

For simulation, acceleration signals were obtained from the same input–output locations in the ADAMS Car model. In the simulation, these points correspond to the ideal knuckle center and subframe attachment point, which slightly differ from the exact physical sensor positions.

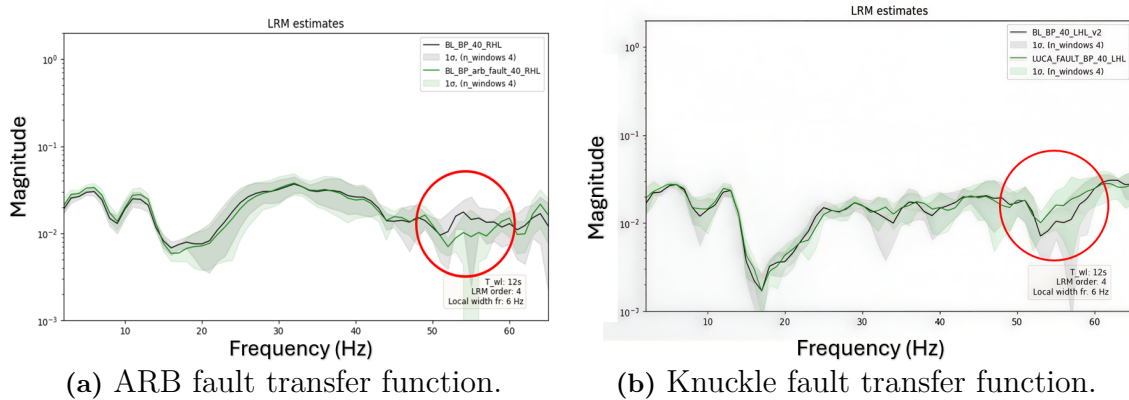


Figure 3.4: Transfer Function comparison of baseline and faulty vehicle from ADAMS Car - *Belgian Pave* 40 Km/h.

Figure 3.4a and Figure 3.4b show a comparison between the baseline (black curve) and faulty (green curve) transfer function magnitudes, estimated using LRM, for both ARB and knuckle faults respectively. The shaded regions represent the $\pm 1\sigma$ standard deviation of the estimates. Overall, the baseline and faulty responses follow a very similar trend across a wide frequency range (approximately 5–45 Hz), suggesting that the system behaviour remains largely unchanged in this region.

However, a difference appears in the higher frequency range, around 50–60 Hz, as highlighted in Figures 3.4a and 3.4b. In this region, the faulty response deviates from the baseline, showing a change in magnitude. The separation occurs mainly in that frequency range because the fault affects the system dynamics more strongly at those frequencies. This represents one of the 36 transfer functions with the highest detection. The total of 36 transfer functions arises from 6 input and 6 output channels, obtained by measuring each of the 2 inputs and 2 outputs in three directions (triaxial sensors). Subsequently, the simulations were extended to public roads, mentioned in Section 3.1.4.

3.1.4 Public Roads

The Belgian Pave road is used as the baseline reference to compare with physical testing as mentioned in Section 3.1.3 and also compare the public roads. Figure 3.6 shows the Power Spectral Density (PSD) of road excitation for different vehicle speeds (30 Kmph, 40 Kmph, and 50 Kmph). It can be observed that increasing speed leads to higher excitation levels and slight shifts in frequency content.

However, the excitation measured at the left-hand-side (LHL) and right-hand-side (RHL) locations remains nearly identical across all speeds, indicating symmetric excitation characteristics for the baseline case.




Figure 3.5: *Belgian Pave* road in ADAMS Car.




Figure 3.6: Road Excitation(PSD) of different speed and direction Belgian pave Road on Left and Right, front Knuckle.

3. Methods

Five public road profiles with varying surface characteristics are considered, as shown in Table 3.2 and Figure 3.7 [14]. Each represents different levels of roughness and excitation.

The public road profiles exhibit varying levels of excitation depending on surface roughness. As shown in Figure 3.8, different roads produce distinct PSD magnitudes, with rougher surfaces such as Coarse Cobblestone resulting in higher excitation levels compared to smoother roads such as Palo Alto. Additionally, the effect of speed further amplifies these differences, leading to increased variability in the system response. Detailed PSD plots for all public road profiles under different speeds are provided in Appendix A.

Table 3.2: Detailed Characteristics of Public Road Profiles.

Road Type	Description	Length (m)	Roughness
Palo Alto	Asphalt good condition	1000	2.2
Pullinger Triangle	Asphalt with tar repairs	1000	2.9
Maude Ave	Contains manhole covers	725	4.5
Rural Concrete	Concrete plates	850	7.7
Coarse Cobblestone	Belgian blocks with patches	850	35.2




Figure 3.7: Public road surface types used in the study.




Figure 3.8: Road Excitation(PSD) of different Public Roads on Left and Right, front Knuckle.

3.2 Fault Modeling

The two faults related to ARB and Knuckle were induced into the vehicle during physical testing as shown in Figure 3.9. The faults in ARB systems and related suspension components can lead to subtle changes in system dynamics. Such faults often produce responses that resemble parameter variations rather than distinct fault signatures, making their detection challenging [15].




Figure 3.9: Actual ARB and Knuckle faults induced during physical testing.



Figure 3.10: Degrees of Freedom for ARB fault and Knuckle fault.

The fault in the ARB drop link assembly is related to a loss of clamping force in the shoulder joint bolts on both the front left and front right sides. As shown in Figure 3.9a there is a visible gap (orange dashed line) of approximately 5mm between the joint faces, which corresponds to about 2–3 threads not being fully engaged. This insufficient tightening results in zero effective clamp force across the joint, meaning the connected components are no longer properly held together.

As a result, the joint is able to move freely, introducing four degrees of freedom, including translation in the Y-direction and rotation about the X, Y, and Z axes as shown in Figure 3.10a. This reduces the stiffness and constraint of the connection, which can affect load transfer through the ARB and lead to unwanted vibrations, noise, and reduced vehicle handling stability.

The fault in the knuckle is also related to reduction of clamping force between Knuckle and Upper Control Arm on the front left side. As shown in Figure 3.9b, a visible gap (orange dashed line) of approximately 2mm is present between the

mating components, corresponding to about 2–3 threads not being fully engaged. In this condition, the clamping force is reduced but not completely lost, meaning the fastener still holds the components together. However, the reduced preload allows relative translational motion along the Z-direction, introducing one degree of freedom as shown in Figure 3.10b. This free play compromises the stiffness of the joint and can affect load transfer and overall suspension performance.

Both of these faults are modelled in ADAMS for the simulation of a faulty vehicle. The process of fault modelling is explained in the Section 3.2.2 and 3.2.4.

3.2.1 Bushings as joint replacements

To replicate clearance effects observed in real joints, a dead-band nonlinearity is introduced in the bushing model. This represents the presence of mechanical play or small gaps within the joint, where no force is transmitted for small relative displacements.

Such behavior is commonly associated with wear, tolerances, or degradation in suspension joints and cannot be captured by purely linear stiffness models. The dead-band is implemented as a region around zero displacement within which the force is set to approximately zero, after which normal stiffness characteristics are applied. This modification introduces a strong nonlinearity in the system response and affects the transmission of vibration signals, thereby making fault detection more challenging.

This approach is consistent with modelling for mechanical joints with clearances, where nonlinear effects significantly influence system dynamics [16]. Additionally, rubber bushings exhibit sensitivity to environmental conditions such as temperature, leading to variations in transfer functions and dynamic response [17].

3.2.2 ARB Fault Modelling in ADAMS

As described in Section 3.2, the ARB fault is defined as a loss of clamping force at the drop link joint, resulting in relative motion between connected components. In the physical system, this is typically caused by loosening of the joint, leading to the introduction of additional degrees of freedom.

To replicate this behaviour in the ADAMS, the default revolute joint between the ARB torsion bar and the droplink is replaced with a linear bushing element, as shown in Figure 3.11. This modification allows controlled compliance in the joint to replicate the faulty condition.

A deadband is introduced in specific DoF to represent the loss of stiffness. In the translational Y axis, a deadband of ± 2.5 mm is provided with near-zero stiffness, allowing free motion within this range. Similarly, rotational stiffness in the X, Y and Z axes is reduced to near-zero within $\pm 7^\circ$, enabling free rotation Figure 3.12.




Figure 3.11: ARB joint replacement with linear bushing.

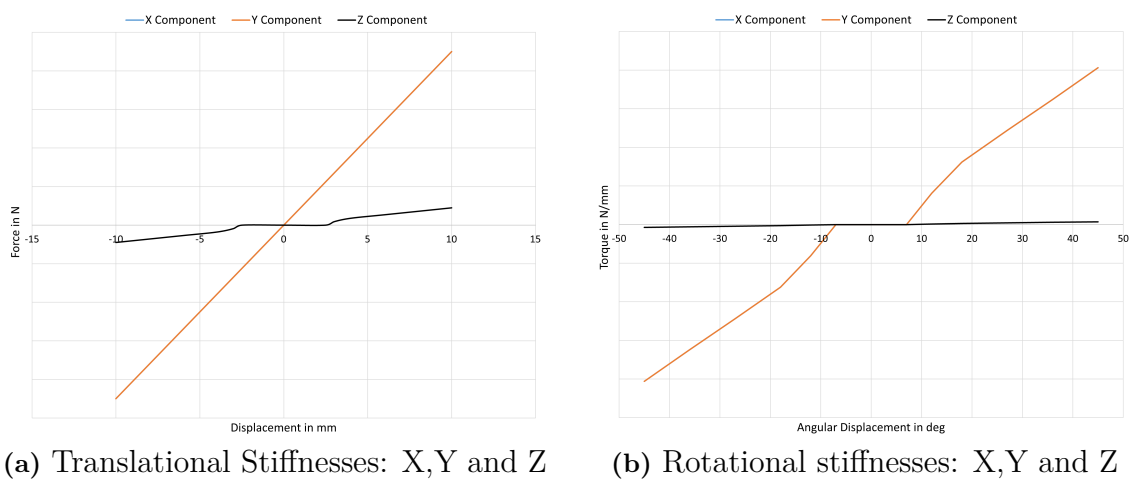


Figure 3.12: ARB joint replacement linear bushing stiffness curve.

This model effectively introduces additional DoFs in the joint, capturing the required behaviour of the fault while maintaining numerical stability in the simulation.

3.2.3 ARB Static Simulation of Fault Behaviour

To verify the implemented fault model and understand its behavior, a static simulation is performed under controlled loading conditions. The simulation considers suspension roll and vertical loading, as shown in Figure 3.13.


A roll angle of $\pm 4.25^\circ$ and a vertical force of approximately 904 N are applied to the system. These conditions are chosen to excite the joint and observe the response under typical loading scenarios like cornering at high lateral acceleration, Figure 3.13a.

The stiffness of the modified joint is evaluated in the dominant DOF, *i.e.* translational motion in the Y axis and rotational motion about the X axis, at the joint location. The response of the faulty joint (represented by the bushing with dead-band) is compared with the ideal joint.


As shown in Figures 3.13b and 3.13c, the faulty joint exhibits a distinct dead-

band region, where the stiffness remains near zero within a specified displacement or angular range. In contrast, the ideal joint shows a continuous increase in force without any deadband, indicating fully constrained behaviour.


Outside the deadband region, both the faulty and ideal joint responses follow a similar trend, indicating that the stiffness remains unchanged beyond the looseness range. This confirms that the introduced bushing model accurately captures the fault behaviour by allowing free motion within the deadband while preserving the original stiffness outside it.



(a) ADAMS Static Simulation: Roll and Vertical force.



(b) Translation Stiffness of joint Y-Axis



(c) Rotational Stiffness of joint X-Axis

Figure 3.13: ARB Static Simulation: Roll and Vertical force.

3.2.4 Knuckle Fault Modelling in ADAMS

In the ADAMS model, the default ideal spherical joint of the knuckle was replaced with a linear bushing at Left Upper Control Arm ball-joint location to represent the fault condition by introducing free play in the translational Z-direction, as the linear bushing permits compliant motion along the Z-axis as shown in Figures 3.14a and 3.14b. A deadband of $\pm 2.5\text{mm}$ was introduced in the translational Z-direction, with negligible stiffness to account for the presence of free play, resulting in a total free movement range of 5mm. Compared to the physical test setup, the free play in the model was intentionally kept slightly higher to better capture the fault dynamics. Furthermore, the stiffness in the translational X and Y directions (red and green curves respectively) was set to very high values to effectively restrict motion in those directions, ensuring that the modeled compliance is primarily limited to the Z-direction (blue curve) where the fault exists. Beyond the deadband region, the




Figure 3.14: Knuckle joint replacement with linear bushing

force–displacement curve follows the default stiffness characteristics of the bushing, allowing the joint to behave normally once the free play is taken up, as shown in Figure 3.15a.

Even though the fault in the knuckle joint is mainly in the form of free play along the Z-direction, it is still important to include rotational stiffness in the model. In reality, the ball joint does not become completely loose in all directions. It still offers some resistance to rotation due to its geometry and surface contact. So, while there is axial looseness, the joint does not behave like a freely rotating connection. The rotational stiffness for all three components was finalised through a trial-and-error process, ensuring that the selected values did not lead to unstable joint behavior or solver issues during simulation. Consequently, for X and Y components (red and green curves respectively), it was kept as the default; for the Z component (blue curve), it was increased comparatively, as shown in Figure 3.15b.

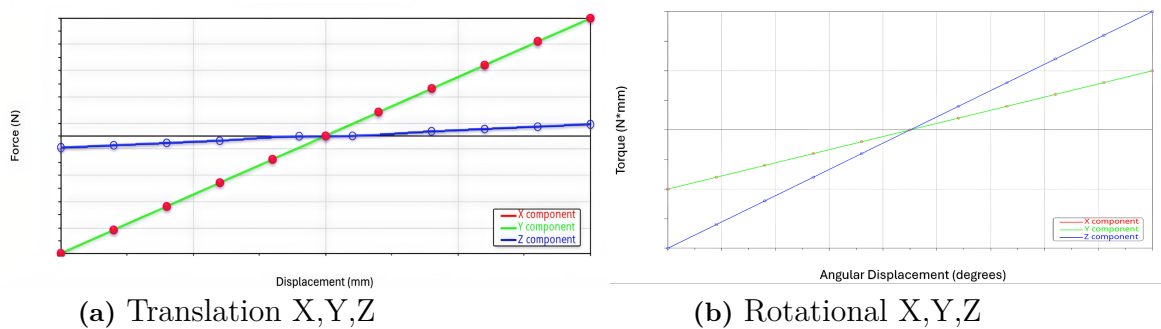


Figure 3.15: Bushing translational and rotational stiffness characteristics.

3.2.5 Knuckle Static Simulation of Fault Behaviour

To validate the knuckle joint fault, a quasi-static parallel Wheel travel test of the front suspension assembly was conducted, where both Wheels were moved together in bump and rebound (± 25 mm) as shown in Figure 3.16.




Figure 3.16: ADAMS Static Simulation: Parallel Wheel Travel.




Figure 3.17: Knuckle fault validation: Time vs Relative displacement between Knuckle and Upper Control Arm.

Figure 3.17 compares the relative displacement between the knuckle and the upper control arm versus time for both baseline and faulty cases. In the baseline (red curve), the displacement remains close to zero, indicating a well-constrained joint with no free play. In contrast, the faulty case (blue dashed curve) shows a clear variation from +2.5 mm to -2.5 mm. The flat regions at both ends represent the limits of the gap, while the transition in between corresponds to the deadband where the joint moves freely. This behaviour confirms the presence of approximately 5 mm total free play and shows that the model captures the effect of the loosened joint.

3.3 Confounding Variables

In practical fault detection scenarios, variations in system behaviour are caused not only by faults but also by changes in operating conditions and component characteristics. These variations, referred to as confounders, can alter the observed signal patterns and lead to ambiguity between normal and faulty conditions.

As discussed in the introduction, confounders play a critical role in evaluating the robustness of fault detection models. Therefore, they are explicitly considered and systematically incorporated into the analysis in this work.

3.3.1 Types of Confounders

The confounders considered in this work are grouped into 'Condition' and 'limit' confounders, based on their physical relevance and severity of impact on system behavior.

Condition confounders represent realistic variations encountered during normal vehicle operation. These include:

- Road excitation conditions (public roads)
- Driving conditions (speed variations: 30, 40 and 50 Kmph)
- Excitation character (Wheel sizing related inputs)
- Loading conditions (RTL condition)
- Combined system (Trailer)

These confounders primarily affect the input excitation and boundary conditions of the system. While they introduce variability in the observed signals, they do not fundamentally alter the intrinsic properties of suspension components.

Limit confounders correspond to extreme or boundary variations in component characteristics that significantly affect system dynamics. The variations beyond these limits are no longer considered normal and are classified as faults. These include:

- Bushing variations (pbush, nbush), capturing compliance and tolerance effects
- Tyre pressure variations (240 kpa, 270 kpa), affecting stiffness
- Damper variations (soft and stiff settings), influencing energy dissipation

Unlike operational confounders, these variations directly modify the physical parameters and nonlinear behavior of the system. As a result, they can produce responses that closely resemble fault signatures, making fault detection more challenging and raising the question of whether the variation should be considered a fault or remain within acceptable limits of normal behaviour.

3.3.2 Modeling of confounder

3.3.2.1 Trailer and Rim

The Trailer and Wheel models were implemented using a script-based approach by utilizing subsystem and template files. Each subsystem references its corresponding template, which defines its structure and parameterization. The Trailer was modeled using multiple subsystems, the Trailer body, the trailing arm suspension, and the trailer wheels. For the rim configuration, a Wheel subsystem was used, with a 22-inch rim size specified in the subsystem parameters.

All subsystem models were integrated into the full vehicle model using the assembly file, and the required vehicle assembly was created as shown in Figures 3.18, 3.19a and 3.19b.

3.3.2.2 Damper, RTL load and Speed

The Soft/Firm dampers, RTL load and Speed were modeled in ADAMS interface itself. RTL is the maximum weight limit of the vehicle including all occupants and cargo, which is approximately 2800 kg, where as *curb2* is approximately 2500 kg. It can be modeled by selecting the 'RTL' option in 'Switch Assembly Variant' dialogue box as shown in Figure 3.20 .

Similarly, we can modify the speed by manually entering the value in the 'Event Builder', after loading the Driver Control File of the road. In Adams Car, a Driver Control File is utilized to define vehicle maneuvers, specify target paths, and control longitudinal/lateral inputs. Overall, they dictate the sequence of driving commands over a given road or path. Hence, the speed was to set to 8.33 m/s which is 30 Kmph as shown in Figure 3.21.

The damper can be selected from 'Modify Damper' dialogue box in the interface provided that the required damper file is located in the directory. Figure 3.22a shows soft and firm dampers for front suspension system. Similarly, we can select for rear suspension. Front soft and firm dampers produce higher forces than rear dampers, even though their curve trends are similar as shown in Figures 3.22b and 3.22c.

3.3.2.3 Tyre Pressure

The tyre inflation pressure was modified through the tyre property (.tir) file by adjusting the operating condition parameter corresponding to inflation pressure.

Two configurations were used for the study: one with both front and rear tyres set to 240 kpa, and another with both front and rear tyres set to 270 kpa. In both configurations, the pressures were applied symmetrically across the front and rear axles, ensuring that no front-to-rear pressure variation was introduced.

Confounder Modeling Pictures


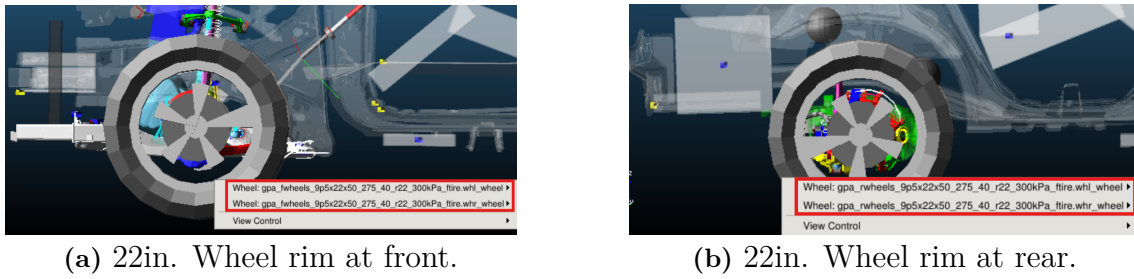


Figure 3.18: XC90 with Trailer.



(a) 22in. Wheel rim at front.

(b) 22in. Wheel rim at rear.

Figure 3.19: XC90 with 22in. Wheel rim (r22) configuration.

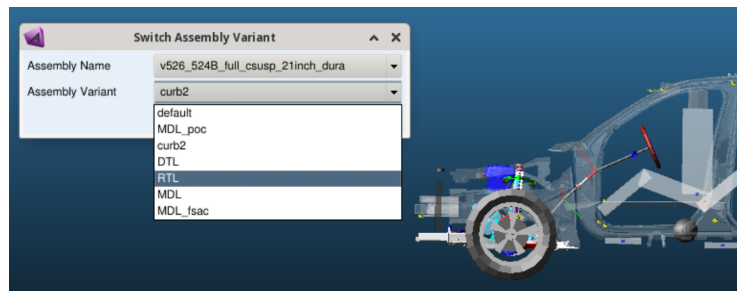


Figure 3.20: XC90 maximum load (RTL) variant.







Figure 3.21: Speed variation through Event Builder in ADAMS.



(a) Stiff and soft damper selection in ADAMS.



(b) Front suspension damper curves.



(c) Rear suspension damper curves.

Figure 3.22: Soft and firm damper curves.

3.3.2.4 Bushing Stiffness Due To Tolerance Effect

Bushing tolerance was introduced in the suspension model to account for variability arising from manufacturing deviations and operational factors such as temperature. Critical compliance joints in both the front and rear suspension systems were identified as shown in Figure 3.23a and 3.23b, and their corresponding bushing characteristics were modified. A bilateral tolerance of $\pm 15\%$ was applied to the nominal stiffness values to represent realistic component variability.

In this context, the positive tolerance ($+15\%$) corresponds to an increase in bushing stiffness, whereas the negative tolerance (-15%) corresponds to a decrease in stiffness. The $\pm 15\%$ deviation was applied in the linear region, altering the slope of the curve. In the nonlinear region, however, the characteristic curve was not rescaled but instead translated (offset) from the original curve while preserving its inherent shape and trend as shown in Figures 3.24a and 3.24b. Furthermore, these modifications were consistently applied across all six degrees of freedom, including translational stiffness in the X, Y, and Z directions, as well as rotational stiffness about these axes.

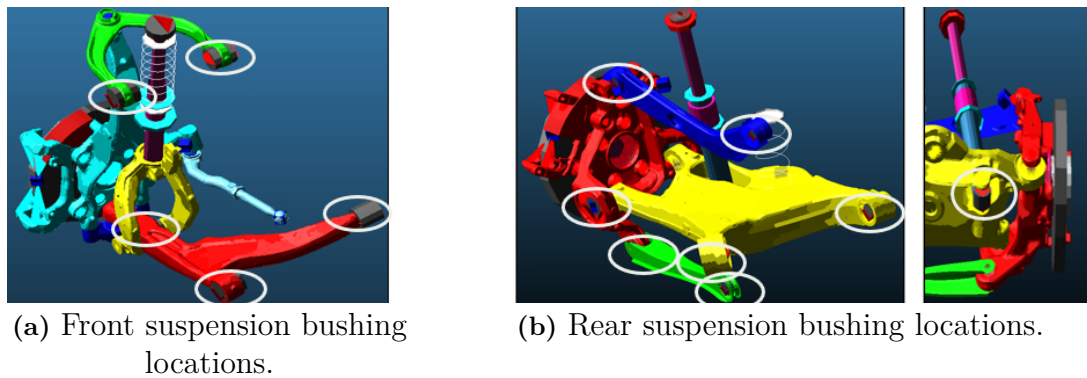


Figure 3.23: Bushing locations in front and rear suspension.

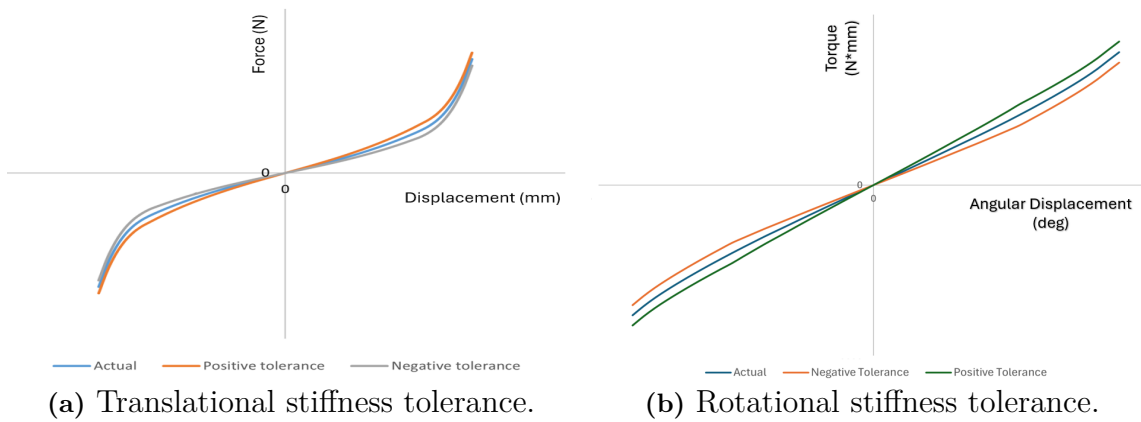


Figure 3.24: Example of bushing stiffness tolerance curves.

3.4 Sensor Configuration

A triaxial accelerometer is a type of inertial sensor that measures acceleration simultaneously along three mutually perpendicular axes, typically denoted as X, Y, and Z. This enables the sensor to capture the full three-dimensional motion of a body, including linear acceleration, vibration, and orientation changes. Accelerometers measure both static (gravity) and dynamic (motion-induced) acceleration, converting them into electrical signals for system analysis.

Overall, triaxial accelerometers provide a compact, reliable, and highly effective solution for measuring multi-axis motion and vibration, making them essential tools in modern automotive diagnostics, structural analysis, and dynamic system monitoring.

3.5 Signal Processing Pipeline

This section describes the data processing workflow used to transform simulated sensor signals into feature representations for fault detection. The overall pipeline is illustrated in Figure 3.25.

The individual processing and classification steps follow the theoretical framework discussed in Chapter 2 and are summarized here from an implementation perspective.

Data: The simulated acceleration signals are organized according to different operating conditions and prepared for further analysis.

Windowing: The signals are segmented into fixed-length time windows, and each segment is treated as an independent sample for subsequent processing.

Local Rational Model: The LRM is applied to each segmented window to estimate the frequency-domain transfer function between input and output signals

Feature Vector: The estimated transfer functions are used to construct feature vectors representing the system response.

The classification and validation stages follow the processing pipeline illustrated in Figure 3.25.

- Logistic regression with L1 regularization is used to classify baseline and faulty samples based on the extracted feature vectors.
- A K-fold cross-validation strategy is employed to evaluate the generalization performance of the model.
- The performance is evaluated using detection-oriented metrics such as FAR, DR, and accuracy.

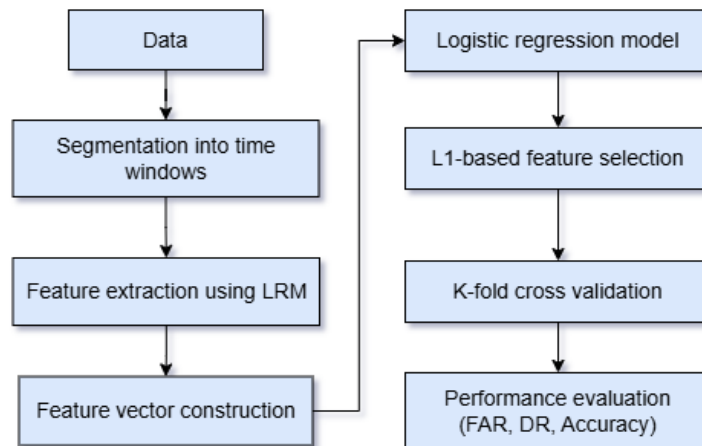


Figure 3.25: Signal processing workflow.

3.6 Result Approach

The overall idea behind result approach is to make the fault detection model reliable even when the confounders are present.

To do this, four types of models are considered: baseline, faulty, baseline with confounders, and faulty with confounders, representing different real driving situations where the system may or may not have faults.

During training, the model learns the difference between baseline and faulty systems and is further exposed to confounded versions (baseline with confounder vs faulty with confounder), helping it understand that not all variations are faults and improving robustness.

In the validation, the focus is on checking how the model behaves when confounders are present. Specifically, we check where the baseline with confounder and faulty with confounder samples are getting classified. If baseline with confounder is classified as faulty, it means the model is getting confused by external effects, which increases the FAR. If faulty with confounder is classified as baseline, it means the model is missing actual faults, which reduces the detection rate. By observing these behaviors, we understand whether the model is leaning too much toward sensitivity or robustness. Based on this, we tune the model accordingly to improve performance.

Finally, in the solution, a clear grouping strategy is applied. Both baseline and baseline with confounder are treated as the same (healthy class), and both faulty and faulty with confounder are treated as the same (fault class). This makes the model more practical, as it learns to ignore confounders and focus only on whether a fault is actually present. In simple terms, the goal is to build a model that does not get confused by normal variations, but still reliably detects real faults.

4

Results

4.1 Fault Detection under Controlled Conditions

The initial evaluation considers fault detection under controlled conditions, where the model is trained and validated using independent scenarios, including baseline vs faulty and individual confounding variables analyzed separately with their corresponding faulty cases.

For all 3-fold validation plots presented in this section, baseline and faulty distributions are represented by blue and red curves, respectively. The solid curves correspond to the trained distributions, while the dashed curves represent the validation distributions obtained at each fold in the 3-fold cross-validation.

The reported performance metrics, such as FAR and ACC, are computed as the average across all folds, providing a robust estimate of model performance.

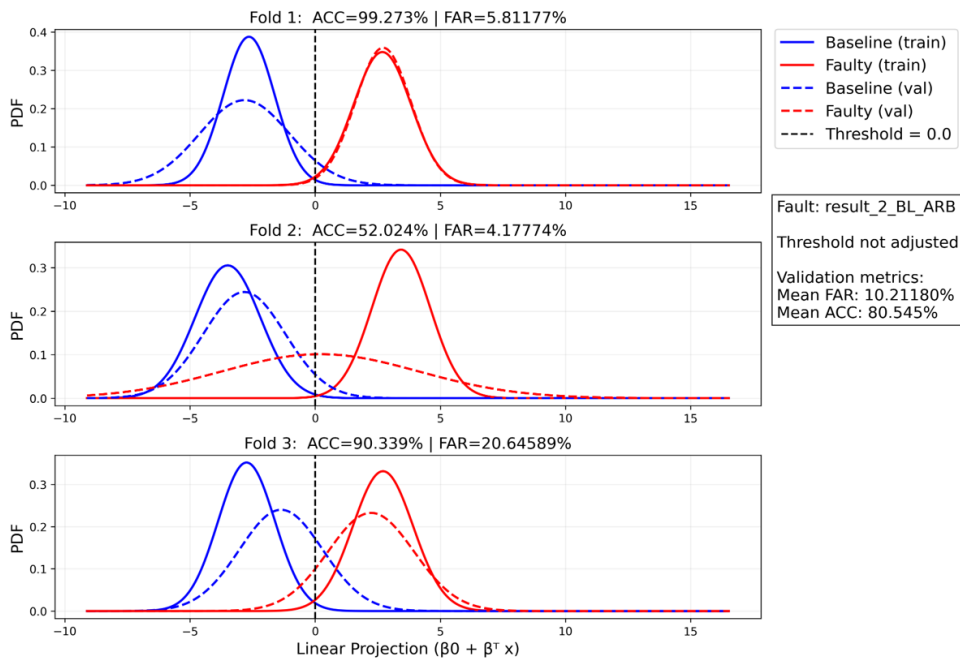


Figure 4.1: 3-fold Validation Baseline Belgian pave and ARB Fault.

4. Results

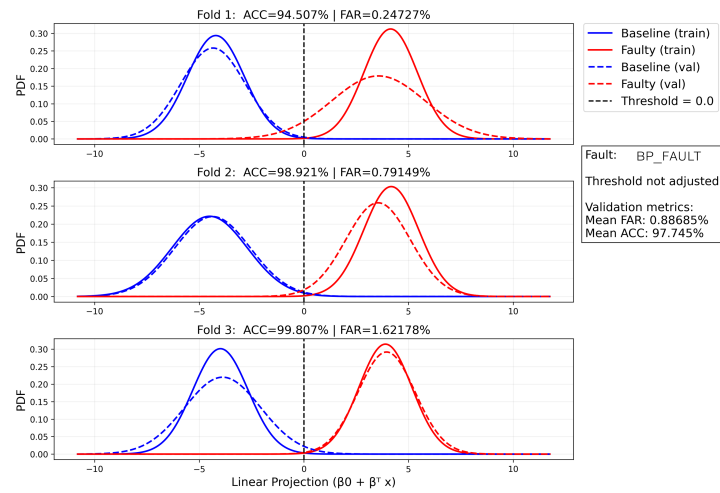


Figure 4.2: 3-fold Validation Baseline Belgian pave and Knuckle Fault.

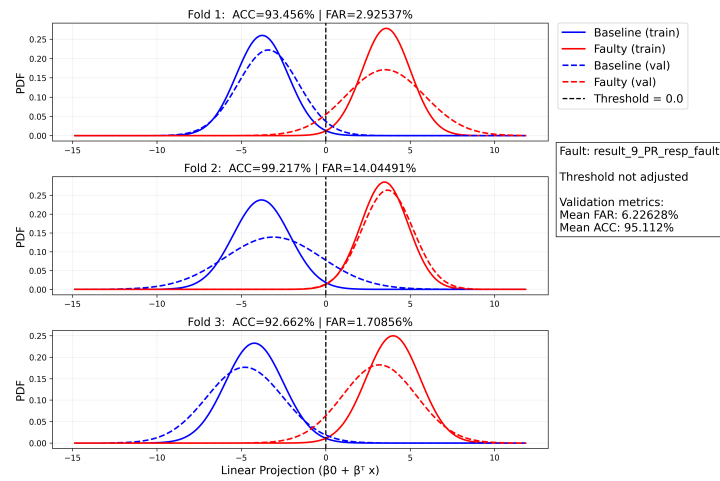


Figure 4.3: 3-fold Validation Public Roads and ARB Fault.

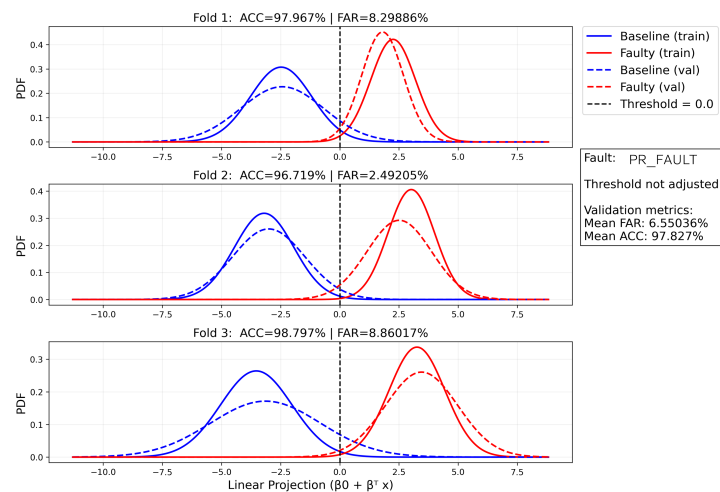


Figure 4.4: 3-fold Validation Public Roads and Knuckle Fault.

Under these conditions the model achieves good separation between baseline and faulty samples for Belgian Pave (Figure 4.1 and Figure 4.2) and Public Roads (Figure 4.3 and Figure 4.4), resulting in low FAR and good detection.

Although road variations are present, they are accounted for in the training set, confirming the effectiveness of the feature representation and classification framework as discussed in Chapter 2.

4.1.1 Effect of Confounding Variables on ARB Fault and Knuckle Fault Detections

This section evaluates the effect of different confounding variables on the detection performance of the ARB fault and knuckle fault. The analysis is based on the methods described in Chapter 3 and the performance metrics defined in Chapter 2.

The results are summarised in Table 4.1 and 4.3, where the impact of both external (Condition) and internal (limit) confounders is evaluated. Detailed visualization of the results is provided in Appendix C.

4.1.1.1 Effect on ARB Fault Detection

While the results in Section 4.1 show that the model performs well under controlled conditions, real world operation introduces variability that affects fault detection. In practice, operating conditions such as **road profiles** and **vehicle speed** are not fixed and may vary continuously, as discussed in Section 3.1.4.

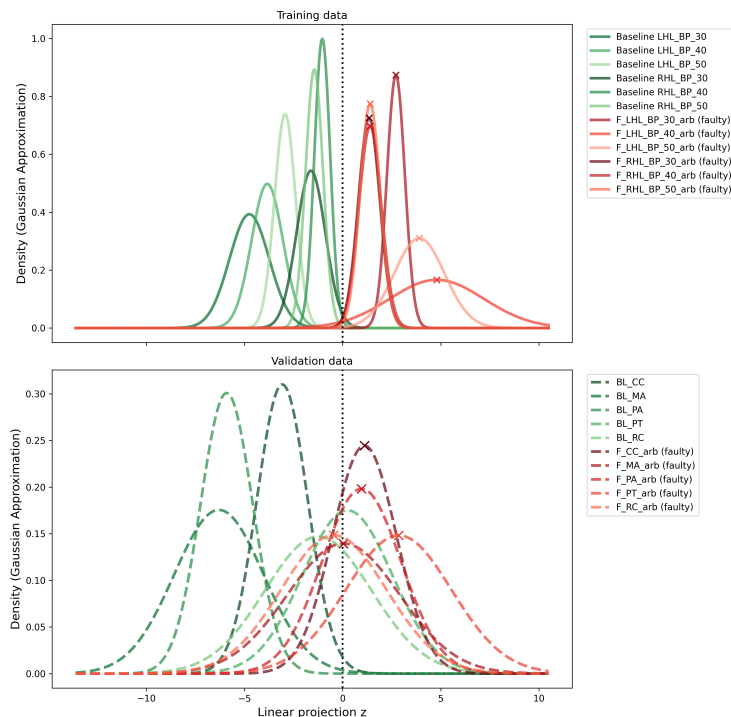


Figure 4.5: Belgian Pave against Public Roads as Confounder.

4. Results

The effect of road conditions is first analyzed by comparing simulations on Belgian pave and public road profiles under similar operating conditions. Although both cases individually show good separability between baseline and faulty conditions Figure 4.1 and 4.3, differences in road excitation result in shifts in the feature.

When analyzed together, this leads to overlap between baseline and faulty distributions, indicating that road conditions act as a confounding variable, Figure 4.5.

To isolate the influence of speed, simulations are performed on the same public road at 50 Km/h and 30 Km/h. Since the road profile is fixed, the observed differences in the system response are due to speed variation. As shown in Figure 4.6, speed changes the frequency content of the signals, increasing overlap between baseline and faulty conditions.

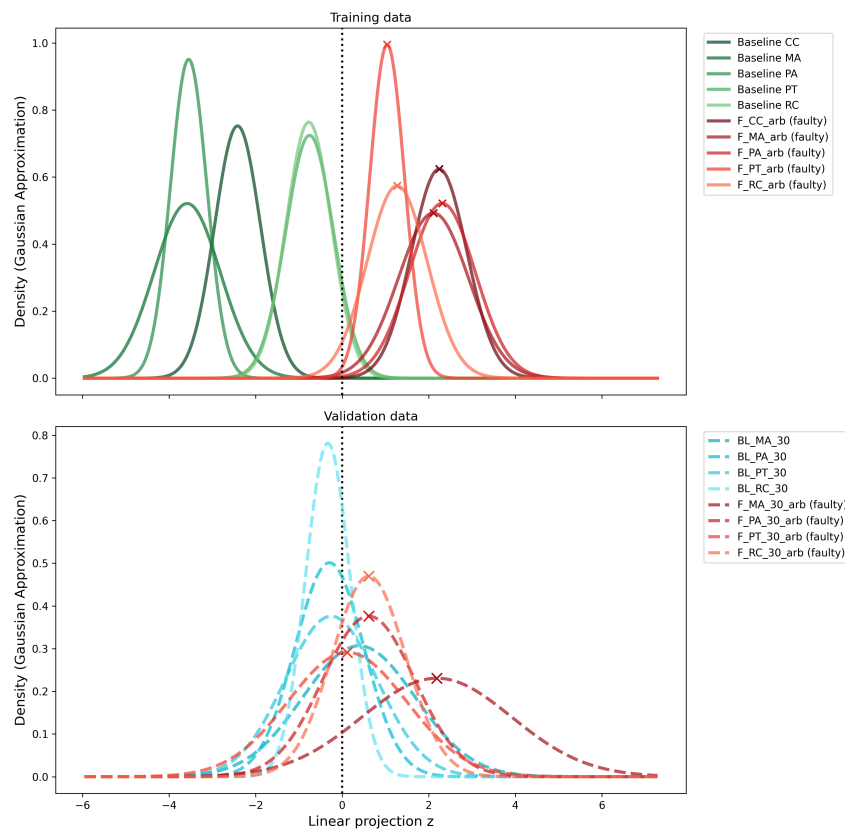


Figure 4.6: Public Roads 50 Km/h against Public Roads 30 Km/h.

These results show that both road conditions and vehicle speed act as confounding variables when not explicitly considered during training. As summarized in Table 4.1, such variations reduce detection performance, motivating the need for a more robust feature representation.

Following the analysis of operating conditions, the influence of additional confounding variables is studied based on the categories defined in Section 3.3.1. These vari-

ables include system-level variations that directly affect the dynamic response of the vehicle and, consequently, the separability between baseline and faulty conditions.




Figure 4.7: Public Road against 270kpa tyre Pressure Confounder.

Among the evaluated cases, certain confounders exhibit a more noticeable impact on fault detection performance. For instance, variations in system configuration, such as decrease in **Tyre pressure** (e.g., 270kpa configuration), introduce significant deviations in the system dynamics. As a result, a more overlap between baseline and faulty feature distributions is observed, leading to a high FAR in the validation case. This indicates that such parameters strongly influence the vibration characteristics of the system and can mask fault related issues.




Figure 4.8: Validation: Public Road against Positive tolerance bushing.

Similarly, variations in suspension related parameters, such as compliance **Bush tolerance**, affect the stiffness and damping properties of the system. These changes alter the dynamic behavior of the vehicle, resulting in less separability between baseline and faulty conditions. The corresponding results show that these confounders

4. Results

introduce non-trivial variations in the feature space, thereby degrading classification performance.

In contrast, other confounders exhibit relatively moderate influence on the detection performance. While they introduce variations in the system response, their effect on the overlap between baseline and faulty distributions is less significant compared to the dominant cases discussed above.

Table 4.1: Effect of All Confounding Variables on ARB fault validation.

Confounders	Mean FAR	Mean DR	Score
Public Roads	0.18	0.65	0.73
Speed 30kmph	0.41	0.73	0.65
Wheel 22inch	0.35	0.66	0.65
240kpa Tyre Pressure	0.57	0.86	0.57
270kpa Tyre Pressure	0.67	0.92	0.49
Positive Tolerance Bushing	0.39	0.86	0.71
Negative Tolerance Bushing	0.32	0.95	0.79
RTL	0.47	0.93	0.68
Trailer	0	0.60	0.75
Soft Damper	0	0.90	0.95
Stiff Damper	0	0.91	0.95

A comprehensive summary of the performance across all considered confounding variables is presented in Table 4.1. The results highlight the varying degree of influence of different confounders on the detection task. Detailed plots corresponding to each confounding condition are provided in Appendix C, which further illustrate the changes in feature space separability and distribution overlap.

Overall, the variation in performance across different confounding variables is summarized in Figure 4.9, where the relative impact of each condition is visualized.




Figure 4.9: Effect of All Confounding Variables on ARB fault validation.

Public Road Influence on ARB Fault Detection

To further investigate this effect, the influence of different road profiles is analyzed for selected confounders representing the most and least challenging cases. Table 4.2 shows that the performance drop caused by the 270 kpa tyre pressure confounder varies significantly across road types, with certain profiles resulting in extremely high false alarm rates.

Table 4.2: Effect of Public Roads on 270kpa Tyre Pressure and Stiff Damper Confounders with ARB fault Detection.

Road Type	270 kpa Tyre Pressure			Stiff Damper		
	FAR	DR	Score	FAR	DR	Score
Coarse cobblestone	1.00	1.00	0.00	0.00	1.00	1.00
Maude Ave	0.73	0.96	0.42	0.00	0.93	0.96
Palo Alto	0.85	1.00	0.26	0.00	0.99	0.99
Pullinger Triangle	0.19	0.70	0.75	0.00	0.86	0.92
Rural Concrete	0.56	0.96	0.60	0.00	0.78	0.88

In contrast, Table 4.2 shows that less sensitive confounders, such as stiff damping, remain robust across different road conditions, maintaining low FAR and consistent detection performance.

This trend is also reflected in Figure 4.10, where a wider spread (FAR) for the 270 kpa case indicates reduced performance, while the dot/compact shape (FAR) for stiff damping confirms stable behavior across roads.

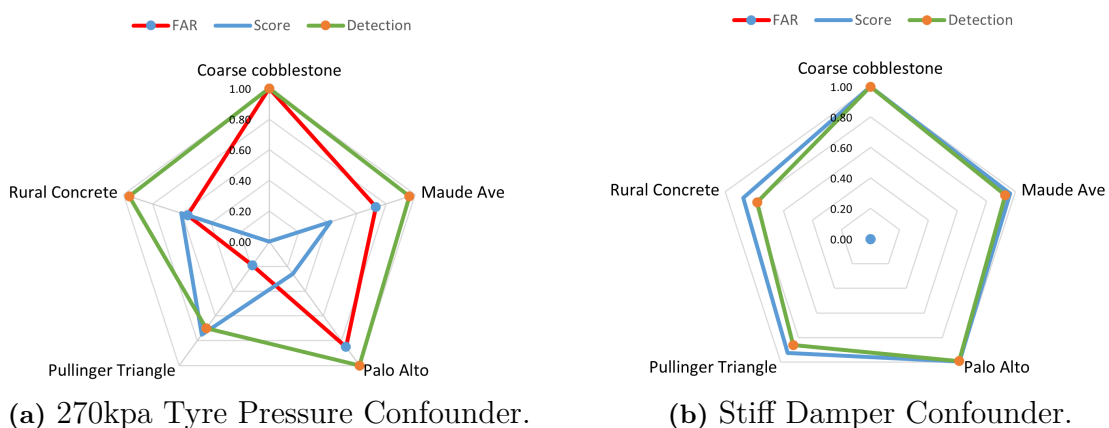


Figure 4.10: Effect of Public Roads on Confounders with ARB fault Detection.

These results show that confounding variables affect ARB fault detection to varying degrees, leading to reduced performance under realistic conditions. This motivates the need for a more robust training approach, discussed in the section 4.2.1.

4.1.1.2 Effect on Knuckle Fault Detection

For knuckle fault detection, RTL, Stiff Damper, 270 kpa Tyre Pressure, and Negative Bushing Tolerance are selected as representative cases to illustrate how performance metrics vary under different confounding conditions. The complete set of results is provided in Appendix C. The training dataset is kept constant, while only the validation dataset is varied according to different confounding conditions.

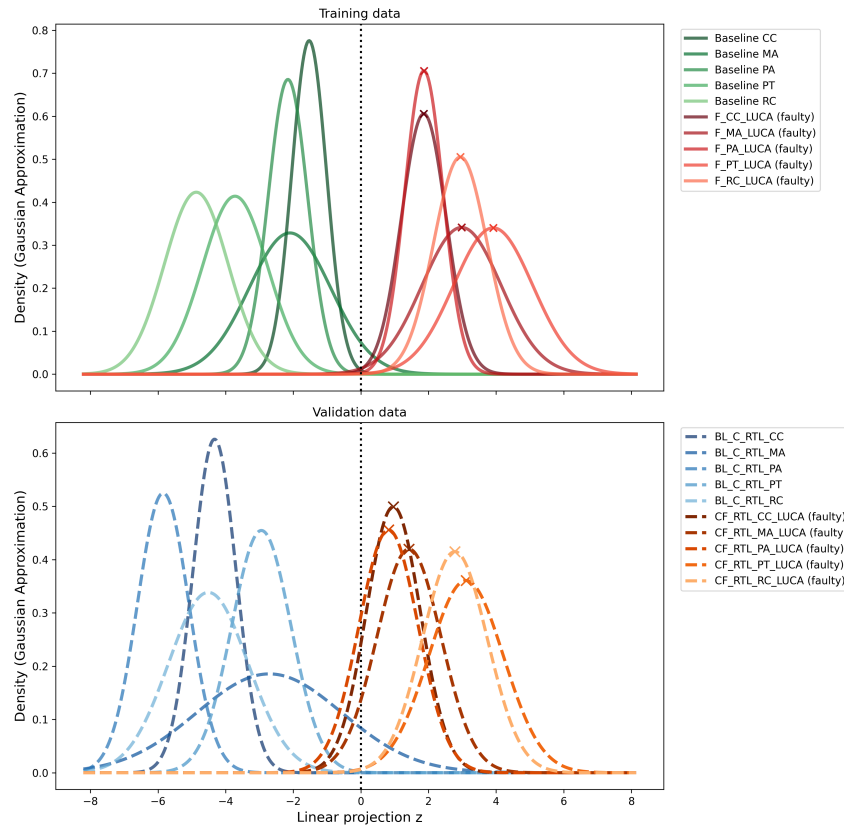


Figure 4.11: Public road against maximum load (RTL) confounder.

Figure 4.11 shows that there is strong separability between the effects of RTL confounder and the combined influence of RTL and fault condition. Specifically, the distributions corresponding to RTL (blue dashed curves) alone remain concentrated in the baseline region, while the addition of RTL to fault condition consistently shifts the distributions (orange dashed curves) towards the fault region. This systematic displacement demonstrates that the fault effect is sufficiently dominant to overcome the variability introduced by the RTL confounder. As a result, only negligible overlap is observed around the decision threshold at $Z = 0$ for both distributions showing there is low FAR and high DR. Figure 4.12 shows that there is relatively low separability between the effect of the stiff damper confounder and the combined influence of the stiff damper with the knuckle fault. The distributions corresponding to the stiff damper alone (blue dashed curves) extend across the decision threshold at $Z = 0$, either partially or completely depending on the public road condition. This indicates a high FAR, as a significant number of baseline samples affected




Figure 4.12: Public road against stiff damper confounder.

by the confounder are incorrectly classified as faults. At the same time, when the stiff damper is combined with the knuckle fault, the resulting distributions (orange dashed curves) remain consistently in the fault region. This leads to a high DR, since most of the fault-affected samples are correctly identified. Overall, this suggests that while the model remains highly sensitive to faults, the influence of the stiff damper confounder reduces its ability to reliably distinguish normal conditions. Figure 4.13 shows that there is relatively low separability between the effect of the




Figure 4.13: Public road against 270 kpa tyre pressure confounder.

270 kpa tyre pressure confounder and the combined influence of the 270 kpa with the knuckle fault. The distributions corresponding to the 270 kpa alone (blue dashed curves) remains mainly in the baseline region. This indicates a low FAR. At the same time, when the 270 kpa tyre pressure is combined with the knuckle fault, the

resulting distributions (orange dashed curves) tend to cross the decision threshold. This leads to a low DR. Figure 4.14 shows an unpredictable behavior of high FAR

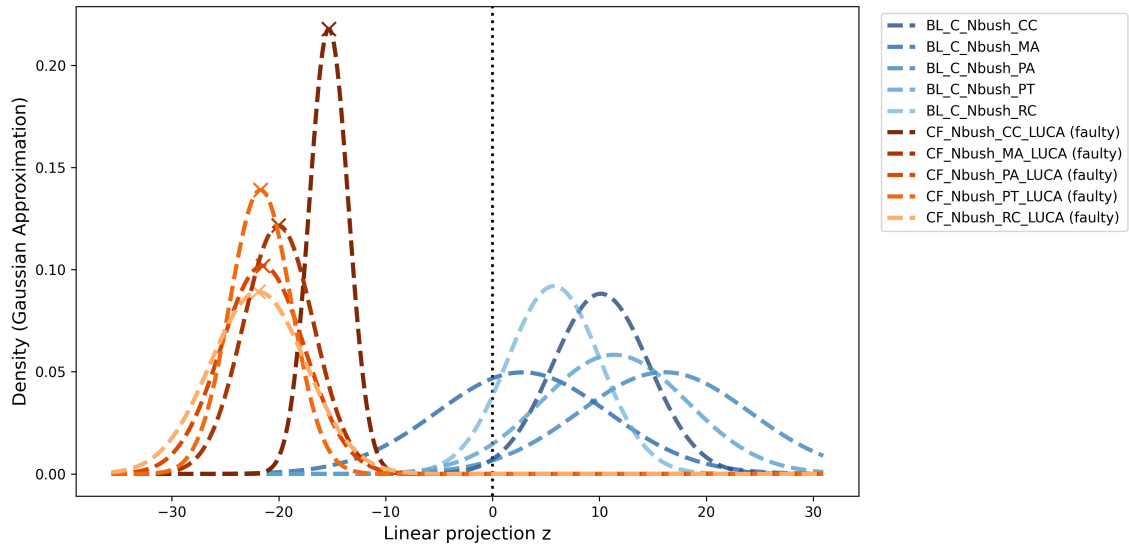


Figure 4.14: Public road against negative bushing tolerance confounder.

with zero fault detection.

A similar summary of performance across all confounding variables is presented in Table 4.3 and Figure 4.15, highlighting their varying influence on detection. Detailed plots for each condition are provided in Appendix C.

Table 4.3: Effect of all confounding variables on knuckle fault validation.

Confounders	Mean FAR	Mean DR	Score
Public Roads	0.43	0.83	0.68
Speed 30 kph	0.22	0.9	0.84
Wheel 22 inch	0.0	0.53	0.69
240 kpa Tyre Pressure	0.0	0.56	0.72
270 kpa Tyre Pressure	0.01	0.57	0.72
Positive Bushing Tolerance	0.0	0.8	0.89
Negative Bushing Tolerance	0.89	0.0	0.0
RTL	0.02	0.93	0.95
Trailer	0.4	0.91	0.72
Soft damper	0.44	0.65	0.6
Stiff damper	0.8	0.96	0.33

Public Road Influence on Knuckle Fault

The most and least challenging cases were considered for knuckle fault as well, with the least challenging being RTL, and the most challenging was considered as stiff damper instead of negative bushing tolerance due to its unpredictable and unreliable




Figure 4.15: Effect of all confounding variables on knuckle fault.

validation. The results in Table 4.4 and Figure 4.16b indicate that the influence of public road profiles on the RTL confounder with knuckle fault is substantial, with performance varying notably across surfaces. While most road types (Pullinger Triangle and Rural Concrete) achieve perfect detection and zero FAR, others such as Maude Ave exhibit a higher FAR (0.31), leading to reduced overall performance.

The stiff damper significantly degrades fault detection performance across all road profiles as shown in Table 4.4 and Figure 4.16a, with consistently high FARs and reduced scores. While detection rates remain relatively high, the system fails to maintain reliability, particularly on Pullinger Triangle and Rural Concrete where performance collapses entirely (score = 0). This indicates that the stiff damper confounder introduces substantial instability, making the model highly sensitive to road conditions.




Figure 4.16: Effect of public roads on confounders with knuckle fault detection.

Table 4.4: Effect of public roads on RTL and stiff damper confounders with knuckle fault detection.

Road Type	Stiff Damper			RTL		
	FAR	DR	Score	FAR	DR	Score
Coarse cobblestone	0.56	0.91	0.59	0.00	0.78	0.88
Maude Ave	0.45	0.85	0.67	0.31	0.99	0.81
Palo Alto	0.51	0.88	0.63	0.00	0.85	0.92
Pullinger Triangle	1.0	1.0	0.0	0.00	1.0	1.0
Rural Concrete	1.0	1.0	0.0	0.00	1.0	1.0

4.2 Training and Improved Fault Detection under Confounding Conditions

To address the issues related to false alarms and missed detections caused by confounding effects, the training process is designed to explicitly incorporate these variations into the model. Instead of training only on pure baseline and pure fault data, the dataset is expanded to include both baseline data with confounders and fault data with confounders. As a result, the baseline class consists of two types of samples: standard baseline data and baseline data affected by confounding conditions, while the fault class similarly includes both pure fault cases and fault cases combined with confounders.

This enables the model to distinguish between variations caused by normal operating conditions and those induced by faults, resulting in improved robustness and reduced false alarms under realistic scenarios.

4.2.1 ARB Training and Solution

Training

The results in Table 4.5 show a clear improvement in performance compared to the validation stage. The FAR is significantly reduced across most confounding variables, indicating better separation between baseline and faulty conditions.

In particular, confounders that previously shown strong decrease, such as tyre pressure variations and bush tolerance, show noticeable improvement after incorporating confounding variability during training. However, certain conditions, such as high tyre pressure (270 kpa) and positive tolerance bushing, still exhibit relatively higher FAR compared to other cases, indicating that these remain challenging even after training.

Overall, the training process improves robustness across a wide range of operat-

ing conditions, but does not completely eliminate the influence of all confounding variables.

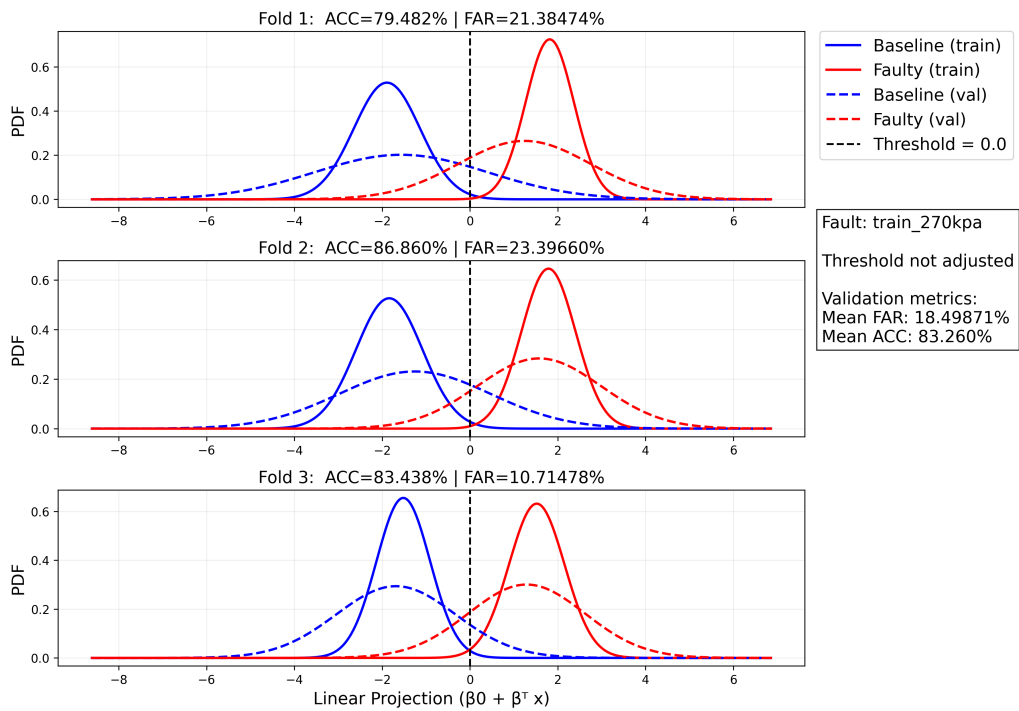


Figure 4.17: Training: 3-fold, public road against 270kpa tyre pressure with ARB fault.

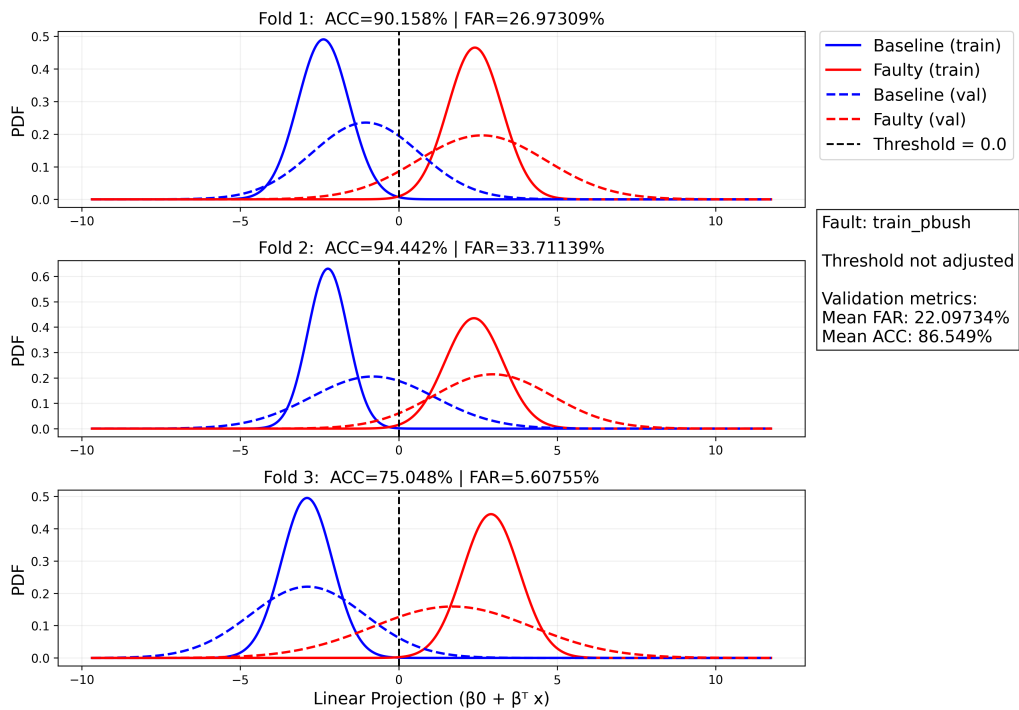


Figure 4.18: Training: 3-fold, public road against positive tolerance bushing with ARB fault.

Table 4.5: Effect of all confounding variables on ARB fault during training for robustness.

Confounders	Mean FAR	Mean ACC	Score
Baseline	0.10	0.80	0.85
Public Roads	0.06	0.95	0.94
Speed 30kmph	0.14	0.87	0.86
Wheel 22inch	0.10	0.95	0.92
240kpa Tyre Pressure	0.11	0.78	0.83
270kpa Tyre Pressure	0.18	0.83	0.82
Positive Tolerance Bushing	0.22	0.86	0.82
Negative Tolerance Bushing	0.00	1.00	1.00
RTL	0.10	0.86	0.88
Trailer	0.00	1.00	1.00
Soft Damper	0.00	1.00	1.00
Stiff Damper	0.00	1.00	1.00




Figure 4.19: Effect of all confounding variables on ARB fault during training for robustness.

Solution

The performance after applying the proposed solution is summarized in Table 4.6. Compared to the training stage, a further reduction in FAR is observed across all confounding variables, along with consistently high accuracy values.

Notably, previously challenging conditions such as tyre pressure variations show significant improvement, with FAR reduced to near-zero levels in some cases (e.g., 240 kpa). Similarly, system-level variations such as bush tolerance and damping

conditions demonstrate improved separability between baseline and faulty distributions.

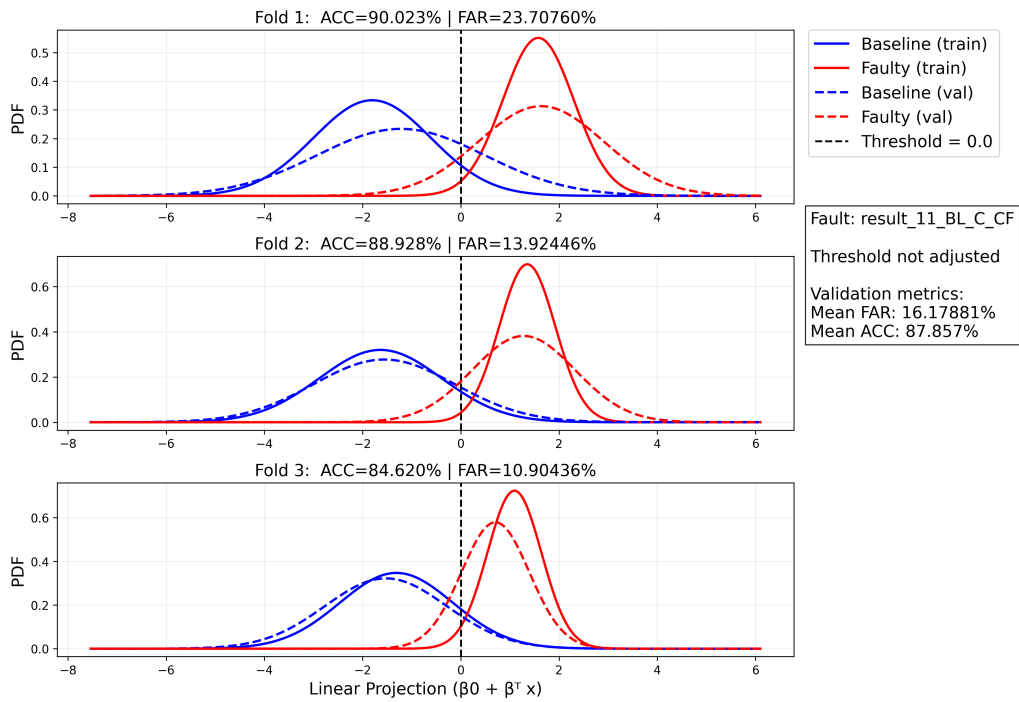


Figure 4.20: Solution: 3-fold, public road against 270kpa tyre pressure with ARB fault.

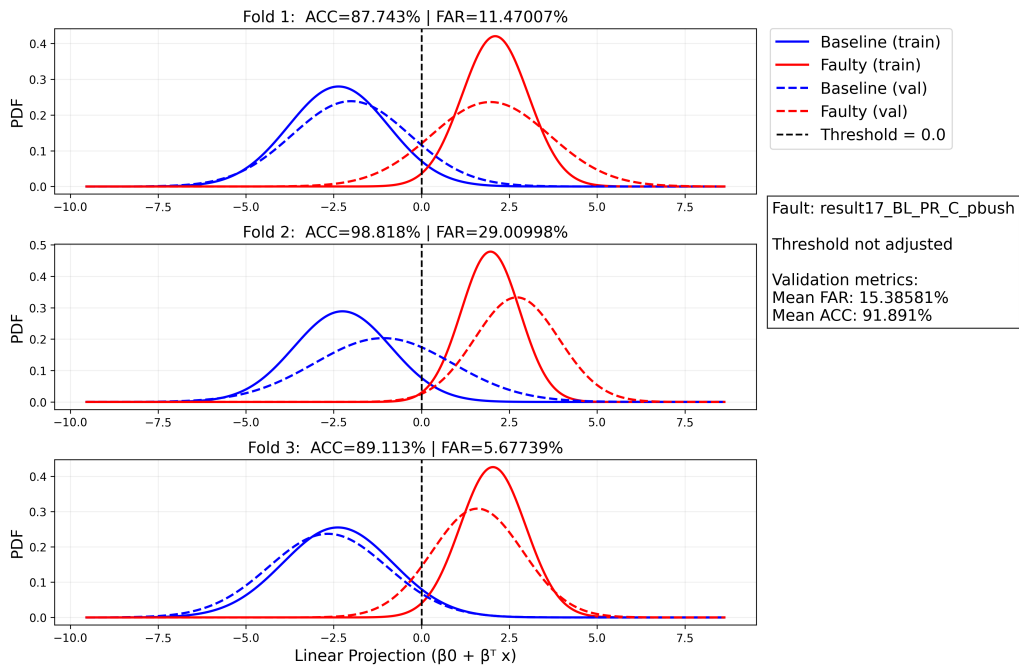


Figure 4.21: Solution: 3-fold, public road against positive tolerance bushing with ARB fault.

Although certain conditions, such as speed variation and high tyre pressure (270 kpa), still exhibit relatively higher FAR compared to other cases, the overall performance is substantially improved compared to both validation and training stages. This confirms that the proposed approach enhances the robustness of fault detection under varying operating conditions, additional results plots can be found in Appendix B and D.

Table 4.6: Effect of all confounding variables on ARB fault after adding robustness.

Confounders	Mean FAR	Mean ACC	Score
Public Roads	0.12	0.89	0.88
Speed 30kmph	0.21	0.85	0.82
Wheel 22inch	0.15	0.89	0.87
240kpa tyre pressure	0.01	0.98	0.98
270kpa tyre pressure	0.16	0.87	0.85
Positive Tolerance Bushing	0.15	0.91	0.88
Negative Tolerance Bushing	0.03	0.95	0.96
RTL	0.09	0.93	0.92
Trailer	0.03	0.96	0.96
Soft Damper	0.02	0.96	0.97
Stiff Damper	0.03	0.97	0.97




Figure 4.22: Effect of all confounding variables on ARB fault after adding robustness.

4.2.2 Knuckle Training and Solution

Training

The RTL and Stiff Damper are considered to illustrate the training as shown in Figure 4.23 and Figure 4.24 respectively. Accuracy and score metrics consistently re-

main close to 1, while the FAR is effectively negligible across all scenarios, including variations in road conditions, vehicle speed, tyre pressure, rim size, damping configurations as shown in Table 4.7. The training results for the knuckle fault–confounder model demonstrate high and uniform performance across all evaluated operating conditions, as depicted in Figure 4.25.

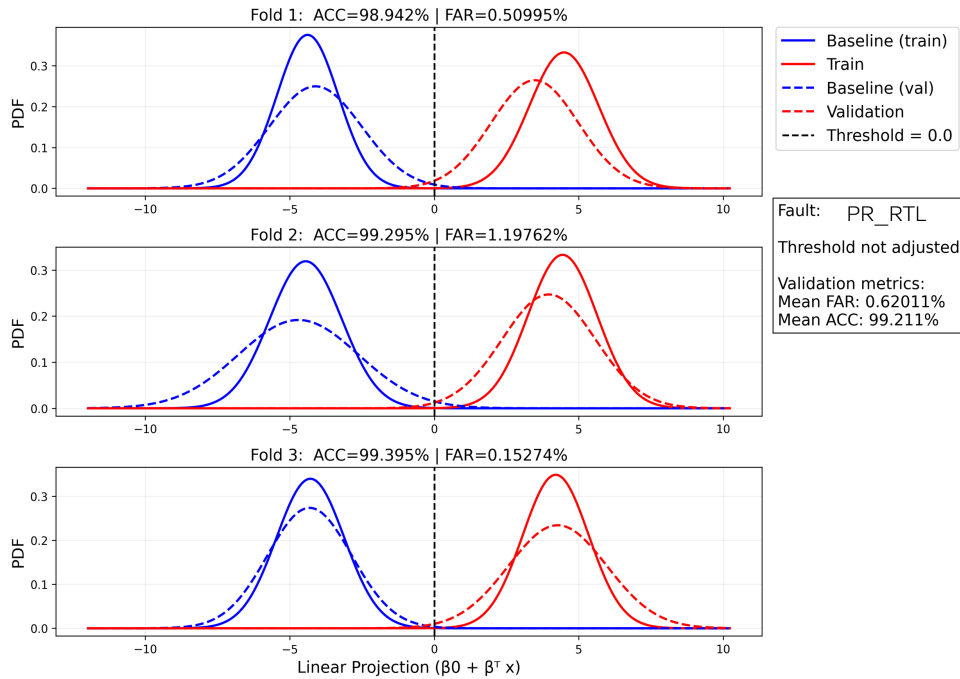


Figure 4.23: Training: 3-fold, public road against RTL with knuckle fault.

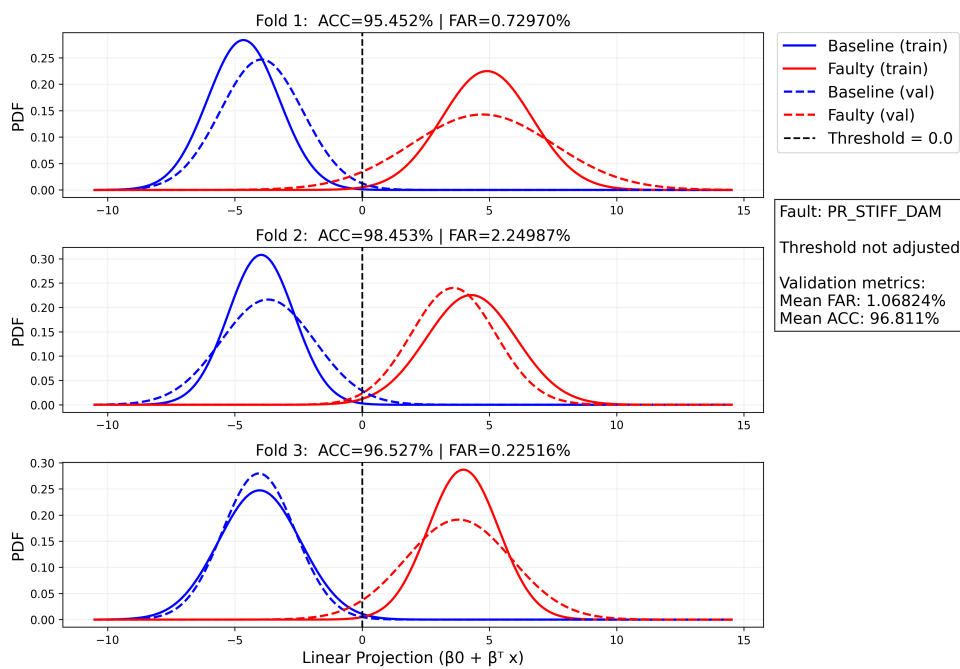



Figure 4.24: Training: 3-fold, public road against stiff damper with knuckle fault.

Table 4.7: Effect of all confounding variables on knuckle fault during training for robustness.

Confounders	Mean FAR	Mean ACC	Score
Baseline	0.01	0.97	0.98
Public Roads	0.07	0.97	0.95
Speed 30kmph	0.0	1	1
Wheel 22inch	0.0	1	1
240kpa Tyre Pressure	0.0	1	1
270kpa Tyre Pressure	0.0	1	1
Positive Tolerance Bushing	0.0	1	1
Negative Tolerance Bushing	0.00	1.00	1.00
RTL	0.0	1	1
Trailer	0.12	0.89	0.89
Soft Damper	0.01	0.97	0.98
Stiff Damper	0.14	0.89	0.87

**Figure 4.25:** Effect of all confounding variables on knuckle fault during training for robustness.

Solution

The validation results for the RTL confounder indicate that only minor tuning is needed to further improve model robustness. The model already demonstrates a low mean FAR of 0.02 and a DR of 0.93. As shown in Figure 4.26, although the mean FAR increases slightly by 3% to 0.05 as compared to validation, the model achieves an almost perfect mean DR of 0.99. For the stiff damper case, the validation results show that the model exhibits a very high mean FAR of 0.8 alongside a high mean DR of 0.96. However, as shown in Figure 4.27, the applied improvements lead to a substantial reduction in mean FAR by 59%, bringing it down to 0.21. This reduction comes with a decrease in mean DR by 12%, resulting in a value of 0.84, which still remains relatively high and acceptable.

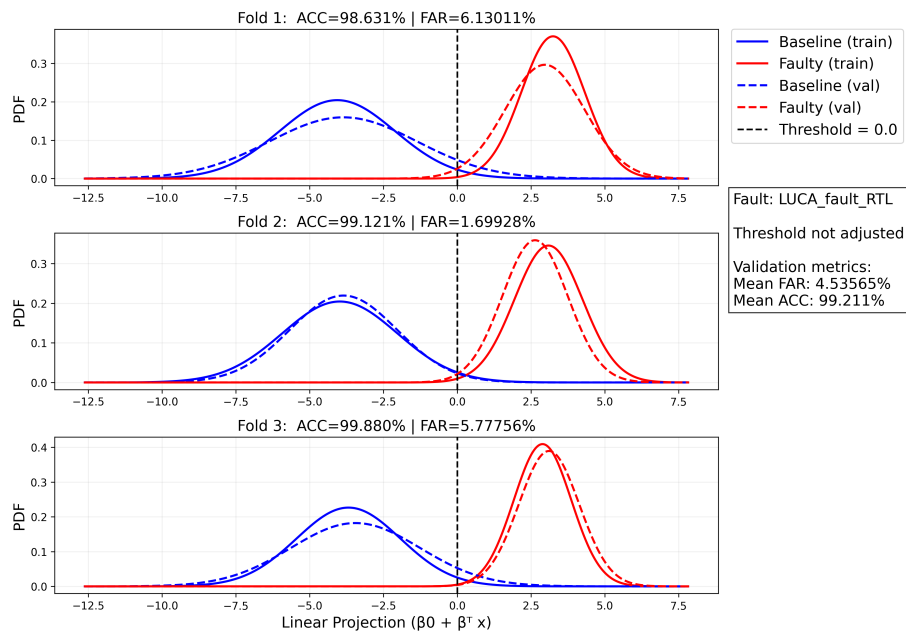


Figure 4.26: Solution: 3-fold, public road against RTL with knuckle fault.

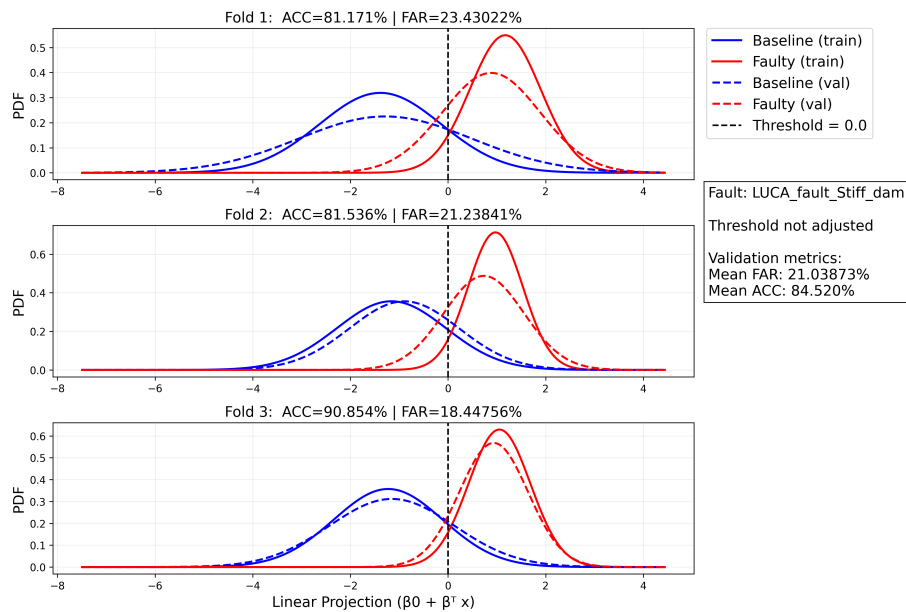


Figure 4.27: Solution: 3-fold, public road against stiff damper with knuckle fault.

The comprehensive summary of performance metrics for all confounders during training for robustness is shown in Table 4.8 and Figure 4.28. The comparative illustration of scores between training, validation and solution is shown in Figure 5.2b, which shows that the training performance remains consistently high across all confounders, indicating a stable model fit. In contrast, the validation performance varies significantly depending on the confounding condition, highlighting the model’s sensitivity to real-world variations. The solution curve generally lies between training and validation, reflecting a balanced performance.

4. Results

Table 4.8: Effect of all confounding variables on knuckle fault after adding robustness.

Confounders	Mean FAR	Mean ACC	Score
Public Roads	0.12	0.89	0.88
Speed 30 kph	0.05	0.93	0.94
22in. Wheel	0.03	0.97	0.97
240 kpa pressure	0.09	0.98	0.94
270 kpa	0.07	0.99	0.96
Positive Bushing Tolerance	0.15	0.91	0.88
Negative Bushing Tolerance	0.08	0.92	0.92
RTL	0.05	0.99	0.97
Trailer	0.12	0.89	0.88
Soft damper	0.2	0.92	0.86
Stiff damper	0.21	0.84	0.81




Figure 4.28: Effect of all confounding variables on knuckle fault after adding robustness.

5

Conclusion

The analysis across ARB and knuckle fault models reveals that achieving optimal performance is inherently a trade-off problem. Specifically, improvements in FAR consistently led to reductions in fault detection capability, highlighting the need to balance sensitivity and specificity depending on the application context.

A key finding is that confounders do not influence all fault models uniformly. Instead, their impact varies significantly depending on the component and fault type. In this context, tyre pressure emerged as the most critical confounder for ARB fault detection during validation, demonstrating a strong influence on model output and reliability. In contrast, the knuckle fault model is more sensitive to stiff damper conditions and negative bushing tolerance, indicating that mechanical stiffness related variations play a dominant role in affecting detection performance.

The comparative analysis after adding robustness showed significant improvement in the performance scores of both models as shown in Figure 5.1. The higher values (F1-scores close to 1) on the radar chart indicate better model performance, reflecting improved fault detection capability combined with lower false alarm rates under the given operating condition. However, under certain operating conditions, particularly public road scenarios and positive tolerance bushing variations, both ARB and knuckle fault models produced similar performance scores, suggesting that some environmental or operational conditions may reduce the distinguishability between different fault types. For the ARB model, the best performance is observed under 240 kpa tyre pressure (score = 0.98), while the lowest performance occurs in the 'Speed 30 vs 50' condition (score = 0.82). Similarly, for the knuckle model, the highest performance is achieved under 22 in. rim and RTL conditions (score = 0.97), whereas the lowest performance is observed under stiff damper conditions (score = 0.81). The radar chart visualization in Figure 5.1 supports these conclusions by illustrating performance across various operating conditions, where values approaching the outer boundary indicate strong and consistent model performance, while inward deviations represent reduced sensitivity and the influence of specific confounders.

Another important insight is the necessity of incorporating confounders during model training. Training model with confounders represented as normal operating variations significantly improves robustness by reducing false positives. Similarly, augmenting faulty models with confounder variations enhances sensitivity and enables the model to better differentiate between true faults and external influences. As shown in Figures 5.2a and 5.2b, both fault models show higher scores after adding

robustness to training (solution), outperforming the validation results.




Figure 5.1: Solution F1-score comparison between ARB and Knuckle fault.

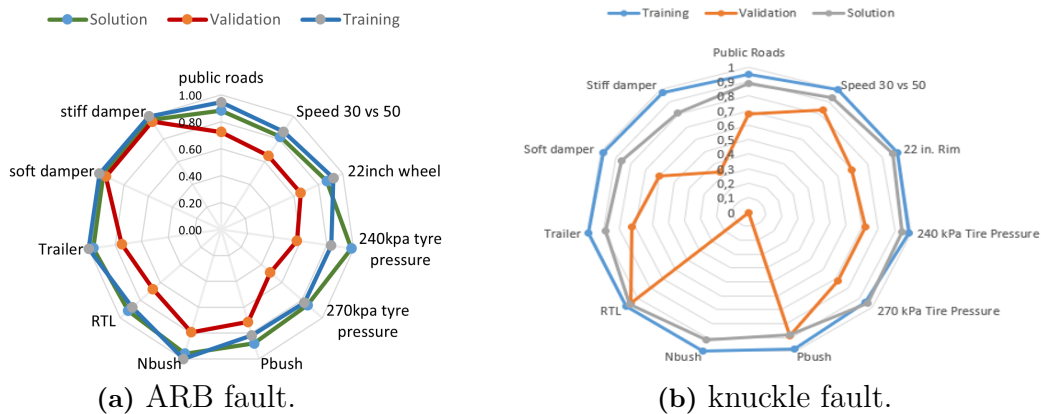


Figure 5.2: F1-score comparison between training, validation and solution for ARB and Knuckle fault.

Future Work

To begin with, the study can be expanded by including more confounding factors. In real driving conditions, many variables affect the system’s behaviour. Considering a wider range of these factors would make the models more realistic and better suited for practical applications.

Another important improvement would be to introduce nonlinear behavior in the bushing used for fault modeling. In the current work, the bushings are modeled in a simplified way, but in reality, their behavior is often nonlinear, especially under different loads and displacements. Capturing this complexity would help improve the accuracy of fault modeling and detection.

It is also essential to validate the results more thoroughly by testing the approach with real vehicle data or experimental setups would provide stronger confidence in the method. This step is important to ensure that the model performs well not just in simulations, but also in real-world scenarios.

Finally, the current approach is based on binary classification, where the system distinguishes between a healthy condition and faults. However, practical diagnostic systems need to identify multiple faults at once. Therefore, extending the method to a multiclass classification approach would be a valuable next step. This would allow the system to distinguish between different types of faults more effectively and provide more detailed diagnostic information.

Bibliography

- [1] Volvo Cars, “Volvo cars media document,” n.d., accessed: 2026-05-21.
- [2] SAE International, “J3016: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles,” SAE International, Tech. Rep., 2021, accessed: 2026-05-21.
- [3] D. McKelvey, “Detection and monitoring of mechanical faults in vehicles-using multivariate frequency function estimates,” 2023.
- [4] K. Coletti, R. Schultz, and S. Carter, “A practitioner’s guide to local frf estimation,” in *IMAC, A Conference and Exposition on Structural Dynamics*. Springer, 2024, pp. 67–80.
- [5] T. McKelvey and G. Guérin, “Non-parametric frequency response estimation using a local rational model,” *IFAC Proceedings Volumes*, vol. 45, no. 16, pp. 49–54, 2012.
- [6] X. Fang, Y. Xu, and T. Chen, “A local gaussian process regression approach to frequency response function estimation,” *IFAC-PapersOnLine*, vol. 58, no. 15, pp. 115–120, 2024.
- [7] R. Voorhoeve, A. van der Maas, and T. Oomen, “Non-parametric identification of multivariable systems: A local rational modeling approach with application to a vibration isolation benchmark,” *Mechanical Systems and Signal Processing*, vol. 105, pp. 129–152, 2018.
- [8] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer, 2009, accessed: 2026-05-21.
- [9] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006, accessed: 2026-05-21.
- [10] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B*, vol. 58, no. 1, pp. 267–288, 1996, accessed: 2026-05-21.
- [11] M. Sokolova and G. Lapalme, “A systematic analysis of performance measures for classification tasks,” *Information Processing & Management*, vol. 45, no. 4, pp. 427–437, 2009.
- [12] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, accessed: 2026-05-21.
- [13] T. Weispfenning, “Fault detection and diagnosis of components of the vehicle vertical dynamics,” *Journal of Dynamic Systems, Measurement, and Control*, 1997.
- [14] H. Klein and D. M. S. A. mbH, “2450 volvo roadselection track report,” 3D Mapping Solutions Auffanggesellschaft mbH, Tech. Rep., Mar. 2025.

- [15] X. Du, L. Mai, and H. Sadjadi, “Fault diagnostics and prognostics for vehicle springs and stabilizer bars,” in *General Motors Research and Development Report*, Warren, MI, USA.
- [16] J. Ambrósio and J. Pombo, “A unified formulation for mechanical joints with and without clearances/bushings and/or stops in the framework of multibody systems,” *Multibody System Dynamics*, 2018.
- [17] P. Koch, C. Angrick, D. Beitelschmidt, G. Prokop, and P. Knauer, “Influence of rubber temperature on transfer functions of bushings,” in *SAE Technical Paper 2015-01-9115*. SAE International, 2015.

A

Appendix A

This Appendix presents the PSD of road excitation for different public road profiles at varying vehicle speeds. These plots provide additional insight into how excitation characteristics vary with road type and speed.

A.1 Road Excitations Different Speed




Figure A.1: PSD of road excitation at 30 kmph, 40 kmph, and 50 kmph for Maude Ave.




Figure A.2: PSD of road excitation at 30 kmph, 40 kmph, and 50 kmph for Palo Alto.




Figure A.3: PSD of road excitation at 30 kmph, 40 kmph, and 50 kmph for Pullinger Triangle.




Figure A.4: PSD of road excitation at 30kmph, 40kmph and 50kmph for Rural Concrete.




Figure A.5: PSD of road excitation at 40kmph and 50kmph for Coarse Cobblestone.

B

Appendix B

This appendix provides detailed 3-fold training results for ARB and knuckle faults under different confounding conditions. These supplementary plots support the observations discussed in Section 4.2.

B.1 Training Results for ARB Fault

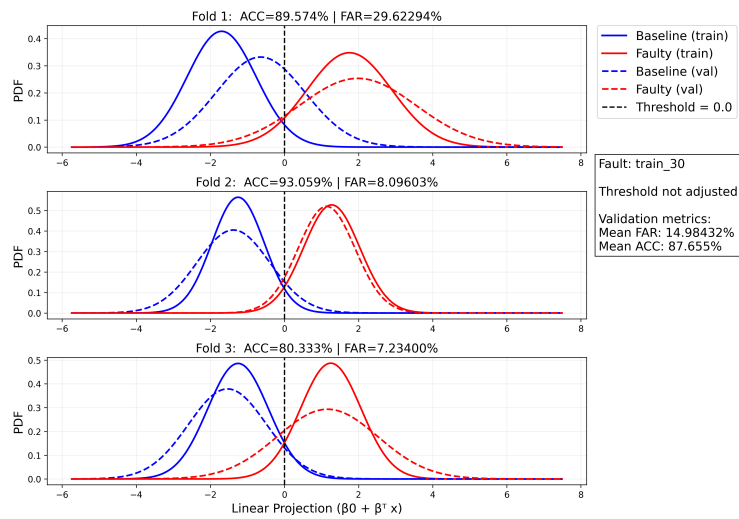


Figure B.1: 3-fold Training Public Road 30kmph Confounder and ARB fault.

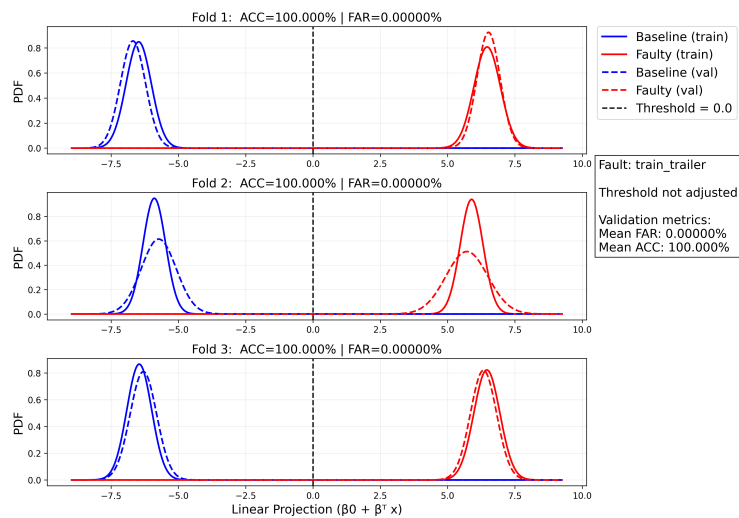


Figure B.2: 3-fold Training Trailer confounder and ARB fault.

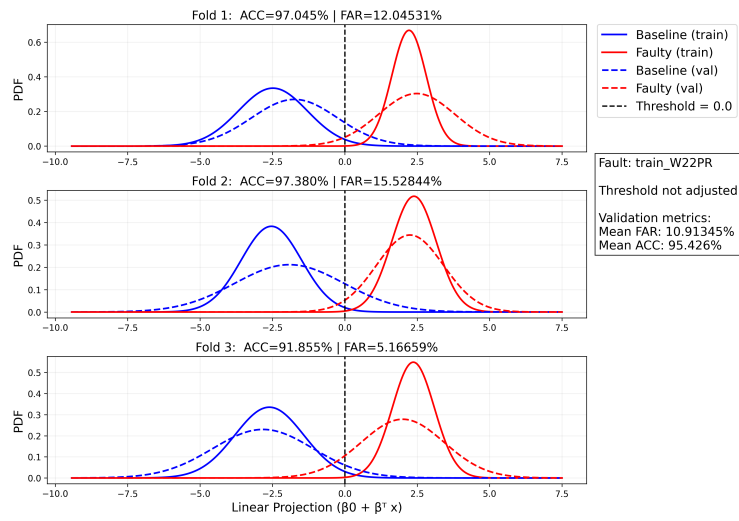


Figure B.3: 3-fold Training 22inch Wheel confounder and ARB fault.

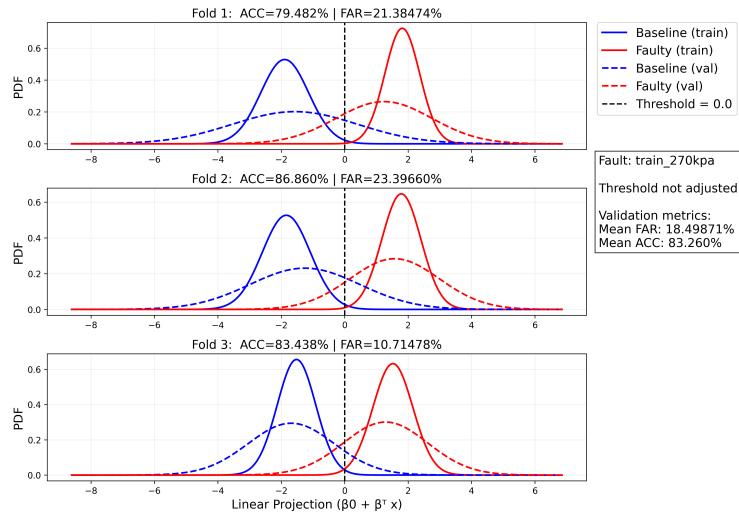


Figure B.4: 3-fold Training 240 Kpa Tyre pressure confounder and ARB fault.

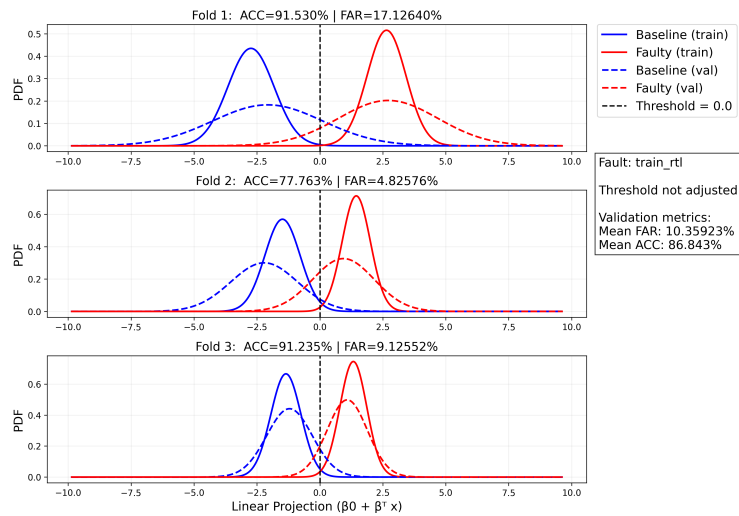


Figure B.5: 3-fold Training RTL confounder and ARB fault.

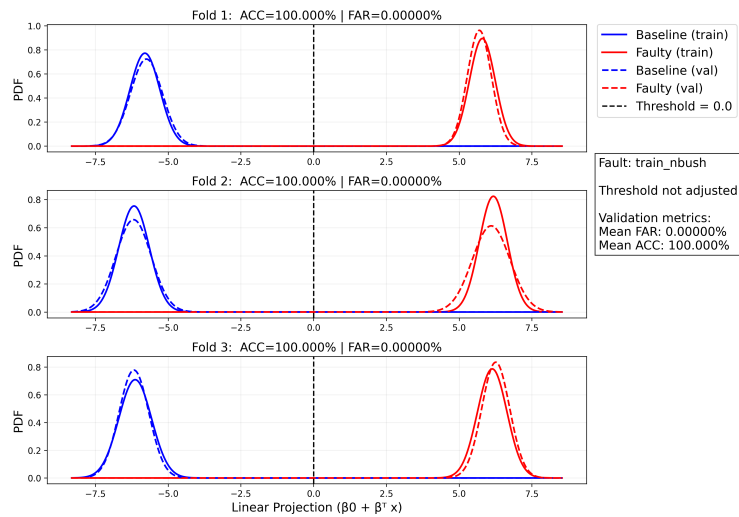


Figure B.6: 3-fold Training negative tolerance bushing confounder and ARB fault

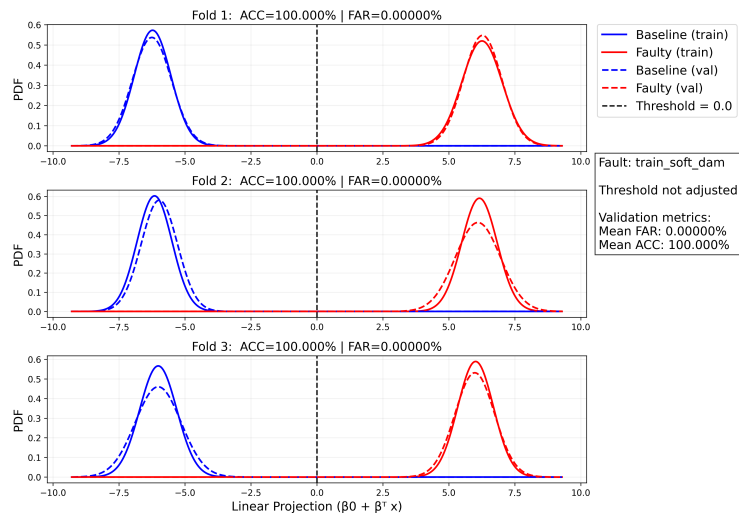


Figure B.7: 3-fold Training soft damper confounder and ARB fault.

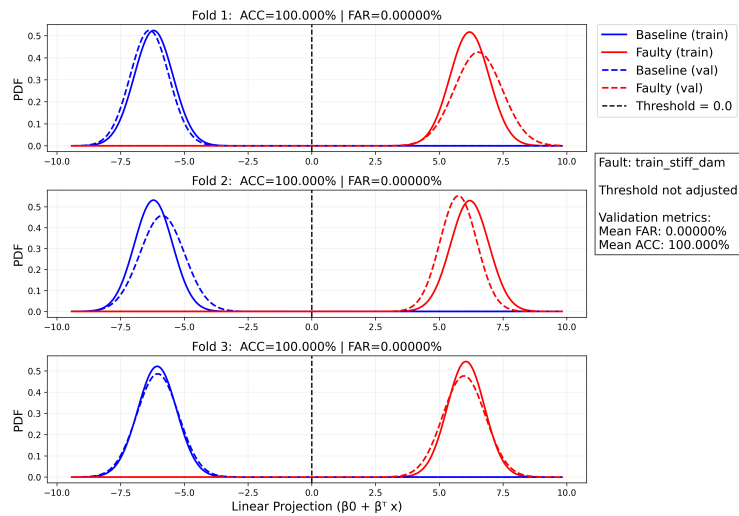


Figure B.8: 3-fold Training stiff damper confounder and ARB fault.

B.2 Training Results for Knuckle Fault

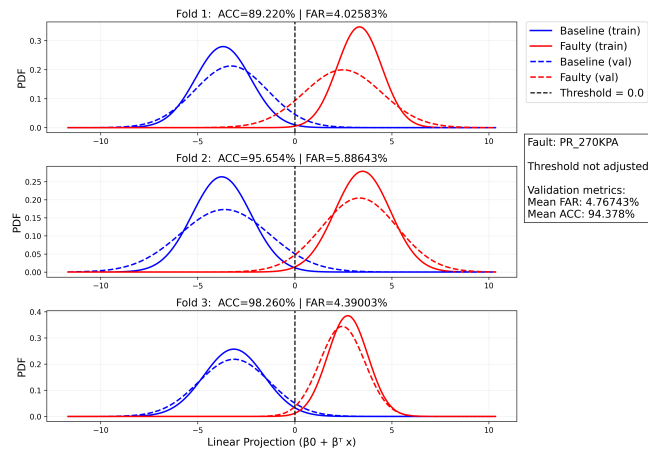


Figure B.9: 3-fold Training 270 Kpa tyre pressure confounder and Knuckle fault.

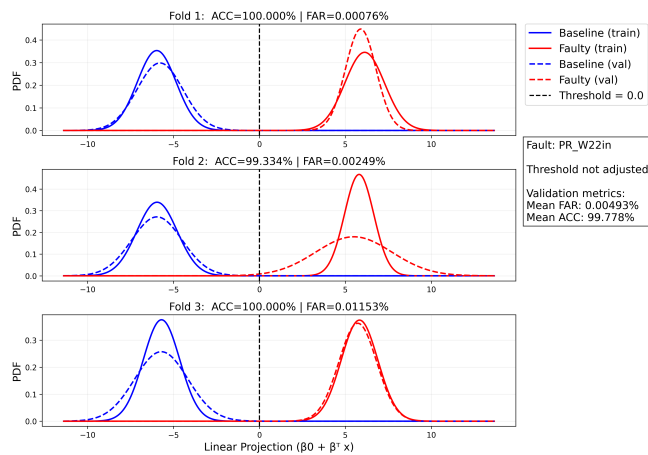


Figure B.10: 3-fold Training 22inch rim confounder and Knuckle fault.

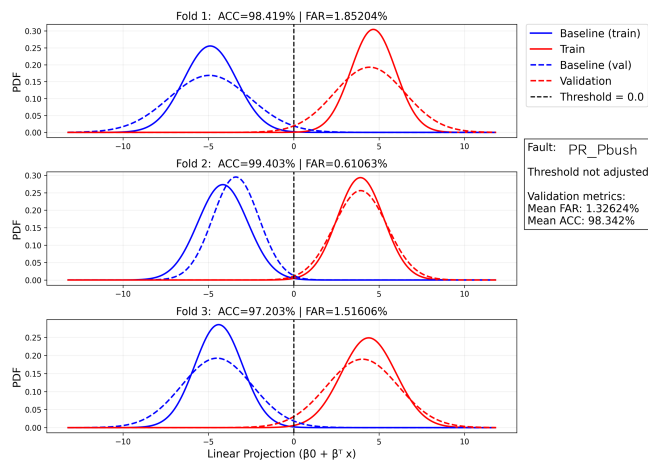


Figure B.11: 3-fold Training Positive Bushing Tolerance Confounder and Knuckle Fault.

B. Appendix B

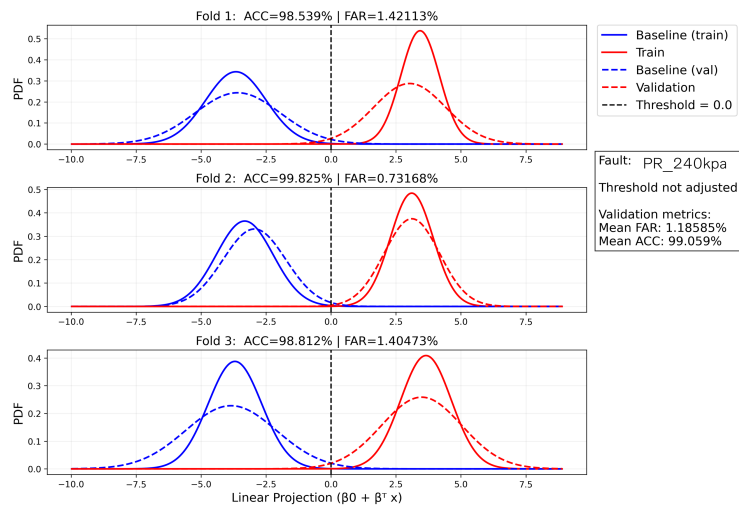


Figure B.12: 3-fold Training 240 kPa Tyre Pressure confounder and Knuckle fault.

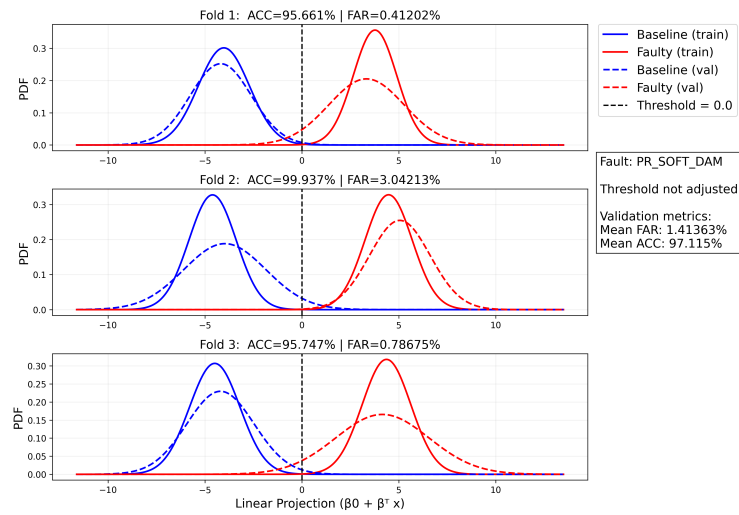


Figure B.13: 3-fold Training soft damper and Knuckle fault.

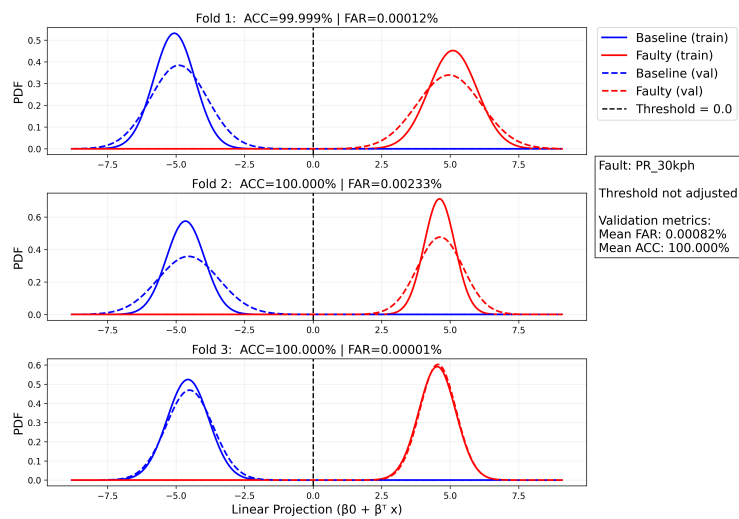


Figure B.14: 3-fold Training 30 Kmph and Knuckle fault.

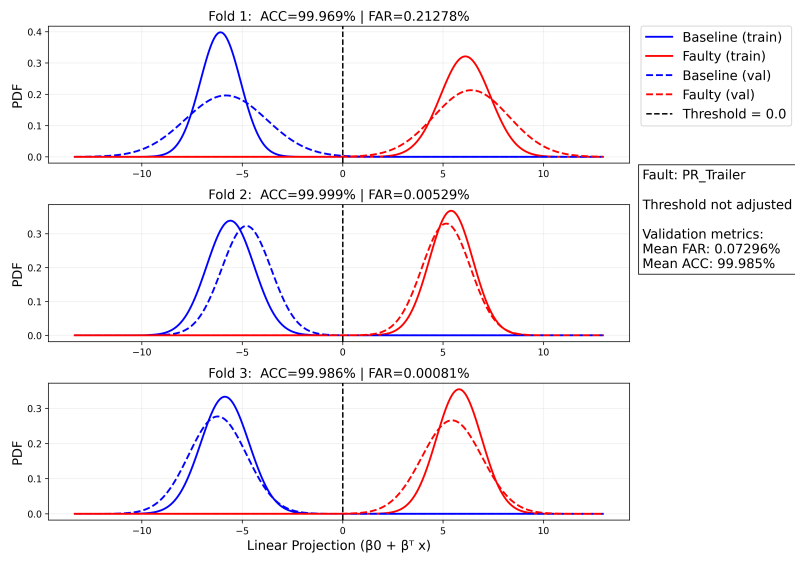


Figure B.15: 3-fold Training trailer and Knuckle fault.

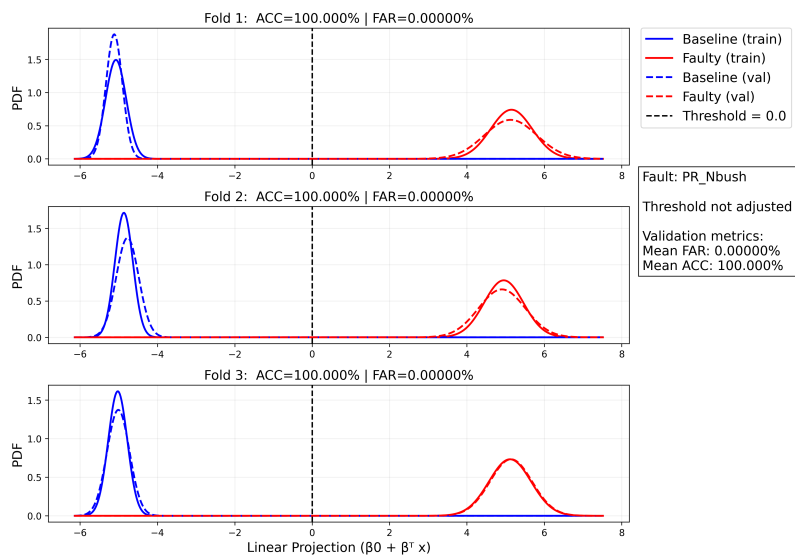


Figure B.16: Training: 3-fold, Public Road against negative bushing tolerance with knuckle fault.

C

Appendix C

This appendix presents additional validation results for ARB and knuckle faults across all confounding variables. These figures illustrate the distribution overlap and classification behavior under different conditions.

C.1 Validation Results for ARB Fault




Figure C.1: Public Road training data.




Figure C.2: Validation: Public Road against 240Kpa tyre pressure.




Figure C.3: Validation: Public Road against negative tolerance bushing.




Figure C.4: Validation: Public Road against maximum load (RTL).




Figure C.5: Validation: Public Road against soft damper.




Figure C.6: Validation: Public Road against stiff damper.




Figure C.7: Validation: Public Road against trailer.




Figure C.8: Validation: Public Road against 22inch wheels.

C.2 Validation Results for Knuckle Fault




Figure C.9: Validation: Public Roads against 22in. rim.




Figure C.10: Validation: Public Roads against 30 kmph.




Figure C.11: Validation: Public Roads against 240 kPa tyre pressure.




Figure C.12: Validation: Public Roads against positive bushing tolerance.




Figure C.13: Validation: Belgian Pave against Public Roads.




Figure C.14: Validation: Public Roads against soft damper.




Figure C.15: Validation: Public Roads against trailer.

D

Appendix D

This appendix contains detailed results obtained after applying the proposed robustness improvement method. These plots show the enhanced separability between baseline and faulty as discussed in Section 4.2.

D.1 Solution Results for ARB Fault

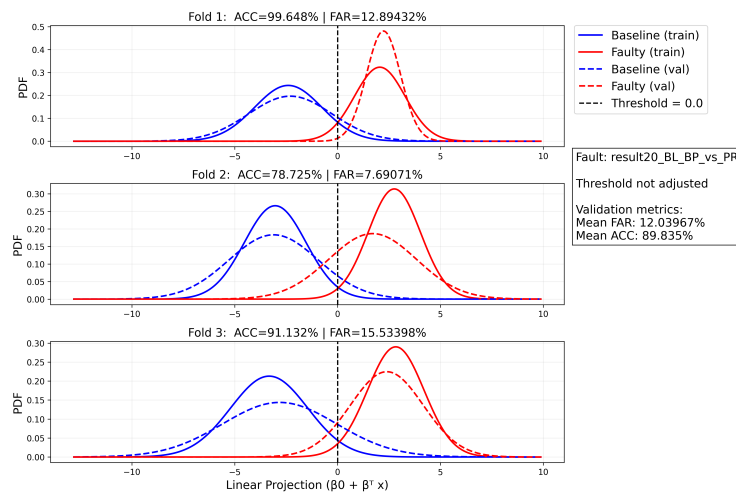


Figure D.1: 3-fold Solution Public Roads and ARB Fault.

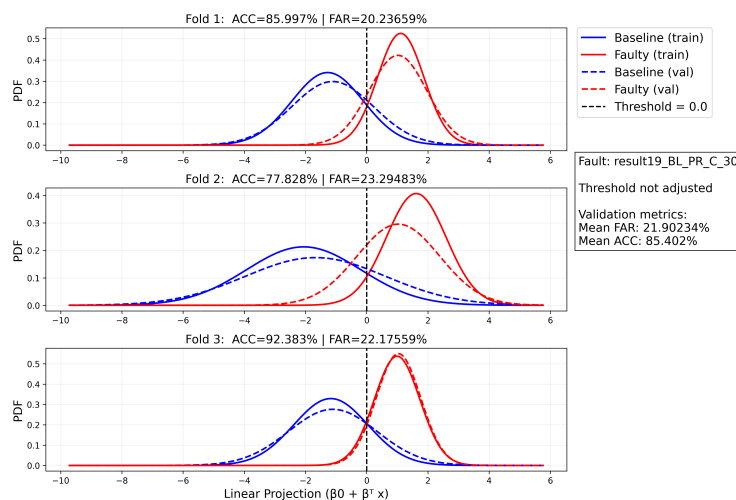


Figure D.2: 3-fold Solution Public Roads 30kmph and ARB Fault.

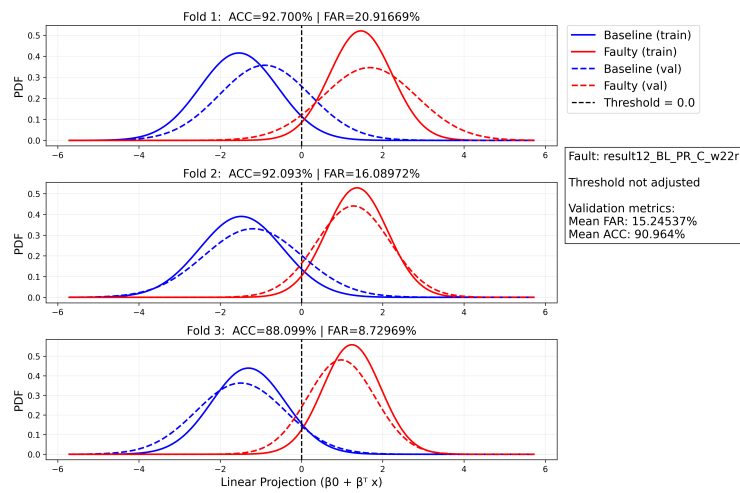


Figure D.3: 3-fold Solution Public Roads 22inch Wheel and ARB Fault.

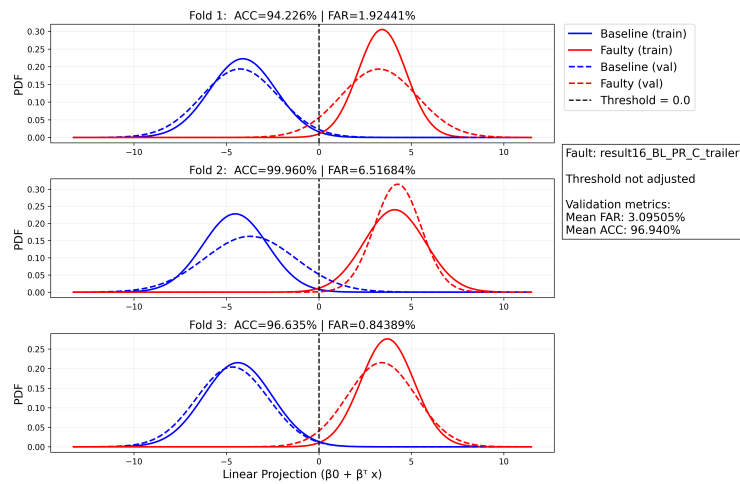


Figure D.4: 3-fold Solution Public Roads Trailer and ARB Fault.

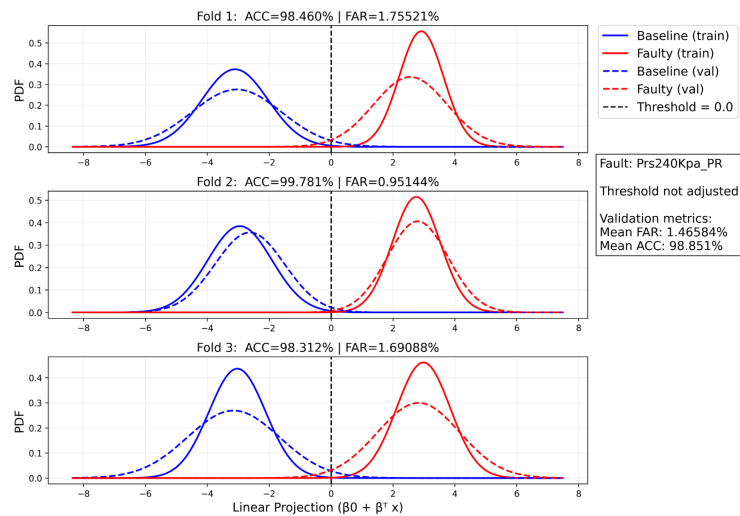


Figure D.5: 3-fold Solution Public Roads 240kpa Tyre Pressure and ARB Fault.

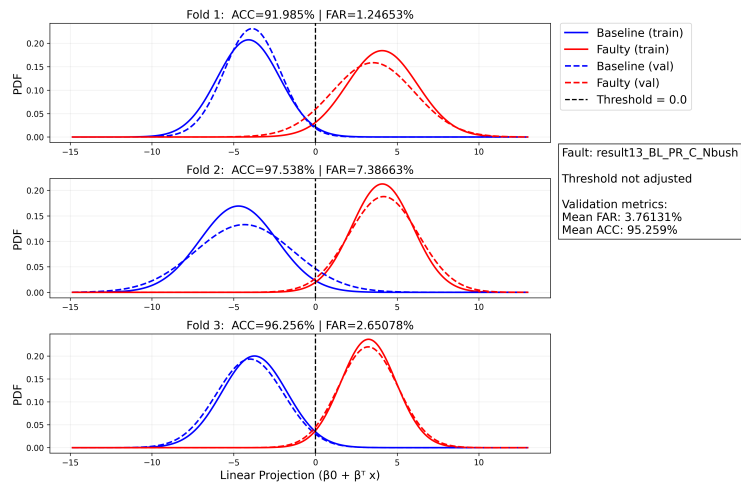


Figure D.6: 3-fold Solution Public Roads Negative Tolerance Bushing and ARB Fault.

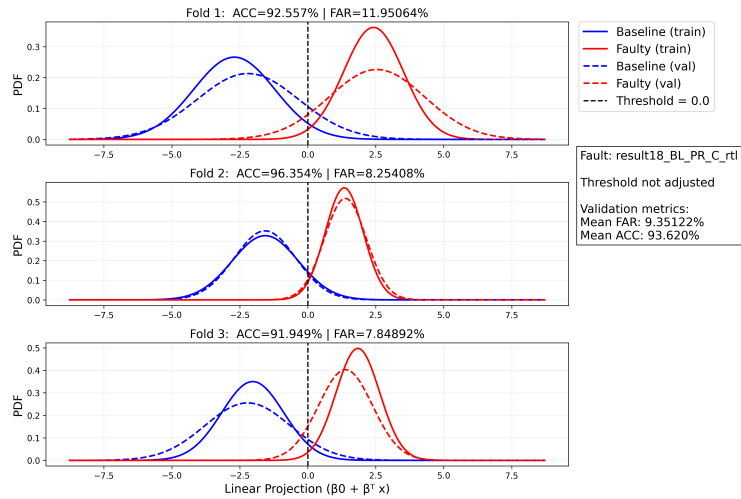


Figure D.7: 3-fold Solution Public Roads RTL and ARB Fault.

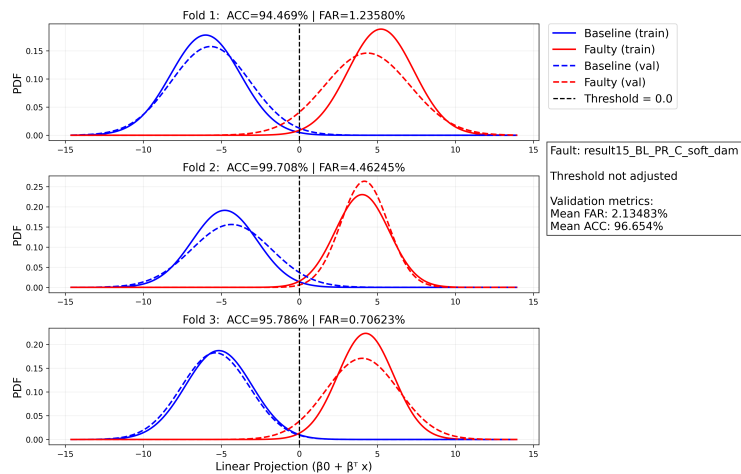


Figure D.8: 3-fold Solution Public Roads Soft Damper and ARB Fault.

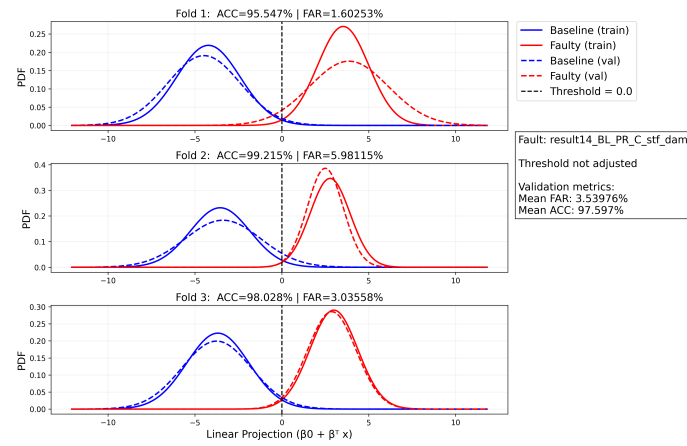


Figure D.9: 3-fold Solution Public Roads Stiff Damper and ARB Fault.

D.2 Solution Results for Knuckle Fault

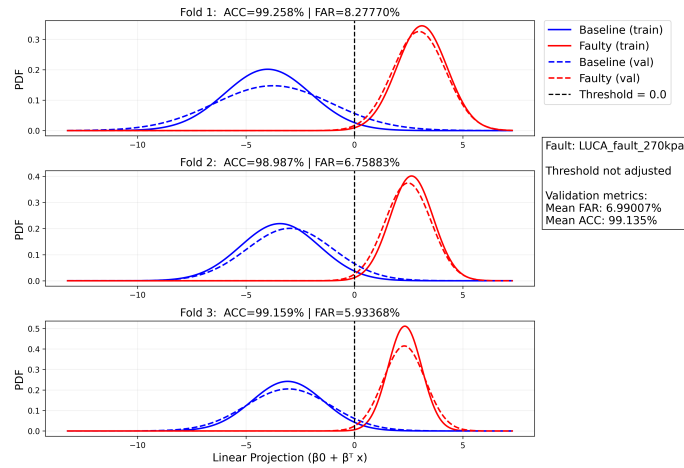


Figure D.10: 3-fold Solution Public Roads 270kpa Tyre Pressure and Knuckle Fault.

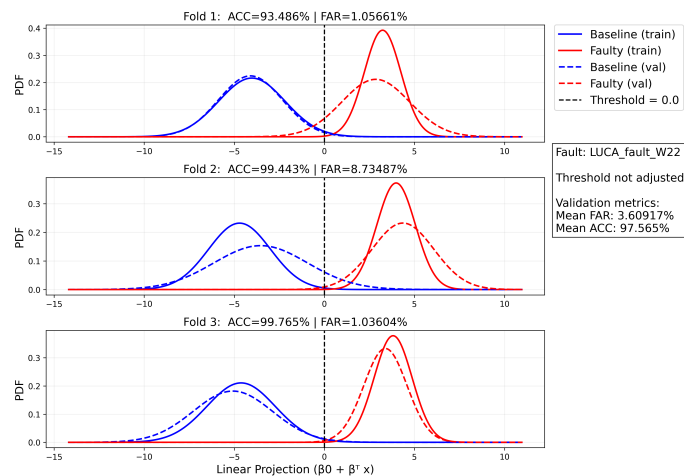


Figure D.11: 3-fold Solution Public Roads 22in. Rim and Knuckle Fault.

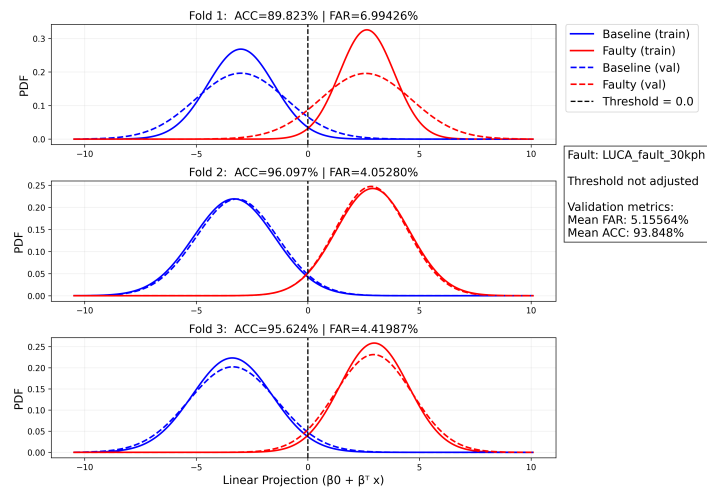


Figure D.12: 3-fold Solution Public Roads 30kph Speed and Knuckle Fault.

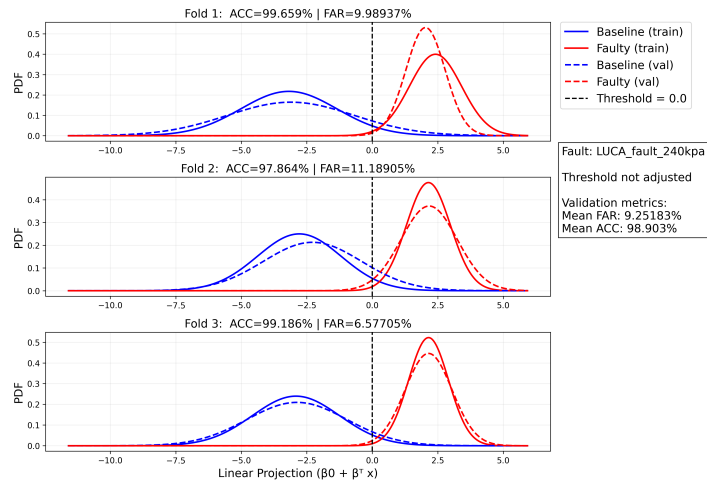


Figure D.13: 3-fold Solution Public Roads 240 kPa Tyre Pressure and Knuckle Fault.

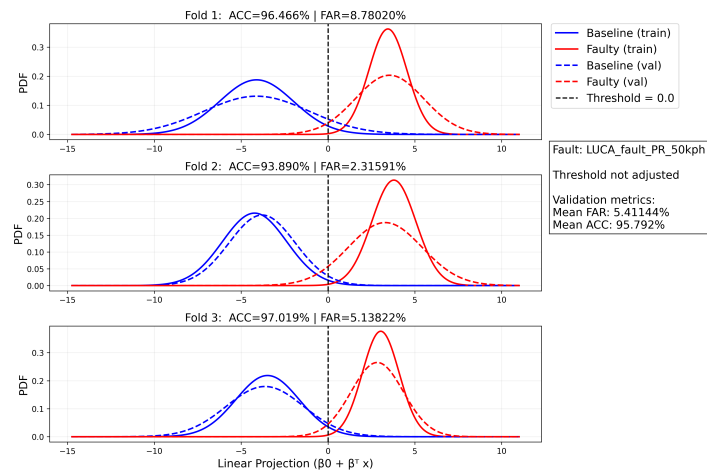


Figure D.14: 3-fold Solution Public Roads and Knuckle Fault.

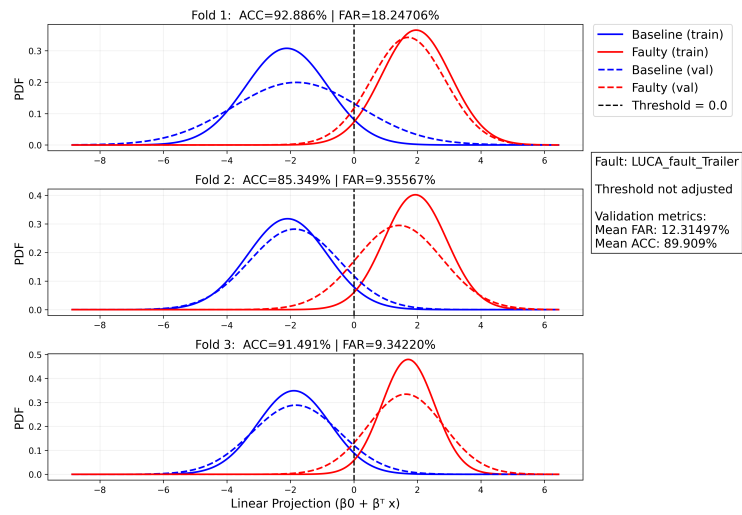


Figure D.15: 3-fold Solution Public Roads Trailer and Knuckle Fault.

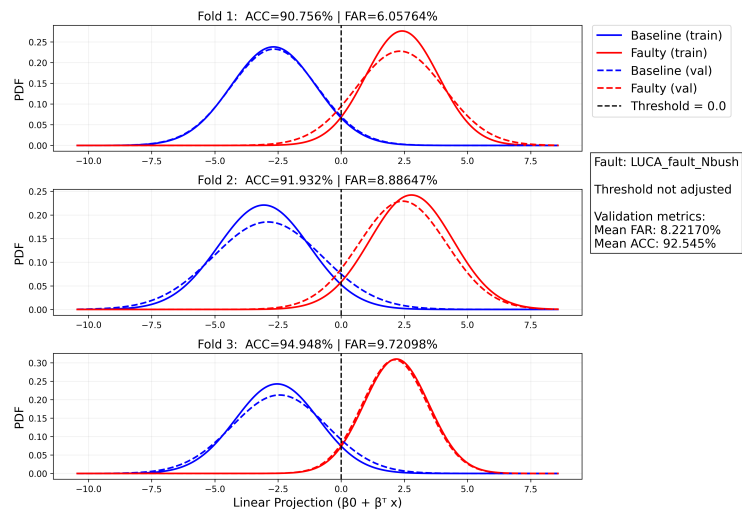


Figure D.16: 3-fold Solution Public Roads Negative Bushing Tolerance and Knuckle Fault.

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden 2026

www.chalmers.se



CHALMERS
UNIVERSITY OF TECHNOLOGY