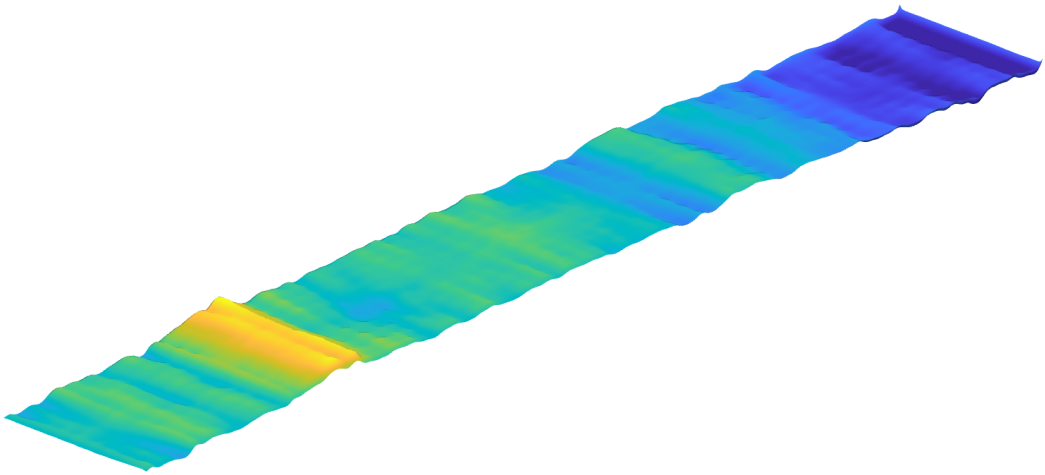

Minimal Statistical Equivalent for Virtual Public Road

Master's thesis

Houshi He Qiyu Tan

Department of Mechanics and Maritime Sciences, Division of Dynamics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis

Minimal Statistical Equivalent for Virtual Public Road

Houshi He Qiyu Tan

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mechanics and Maritime Sciences, Division of Dynamics
Chalmers University of Technology
Gothenburg, Sweden 2026

Minimal Statistical Equivalent for Virtual Public Road

Supervisor: Kristian Carlsson, Technical Expert, Volvo Cars
Examiner: Björn Pålsson, Associate Professor, Chalmers University of Technology

Master's thesis, 2026
Chalmers University of Technology
Department of Mechanics and Maritime Sciences, Division of Dynamics
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Synthetic road surface visualization generated in MATLAB.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Abstract

A minimal statistical equivalent for virtual public roads is developed for durability simulation. The road is a primary excitation source for vehicle loads and fatigue, but measured public roads are long, spatially uneven, and costly to use directly in repeated simulation workflows. The central question is therefore how far measured road data can be reduced while still remaining statistically equivalent to the original road environment for the intended vehicle response application.

Statistical equivalence is interpreted as more than agreement in visual road shape or in a single roughness descriptor. It requires preservation of the road statistics that govern excitation relevant to response, including spectral content, nonstationary roughness variation, localized transient events, and lateral surface structure. Two modelling routes are compared. The Nonstationary Laplace route, denoted NSL, represents nonstationarity implicitly through a compact stochastic description of road severity. The Wavelet Multi Component Reconstruction route, denoted WMCR, separates and reconstructs topography, background roughness, transients, and two dimensional surface content more explicitly.

The methods are evaluated through road statistics, reduced quarter car response, and nonlinear full vehicle simulation in Adams Car. The results show that compact statistical descriptors are useful for efficient virtual road generation, while localized transient structure must be treated carefully when fatigue relevant response is the target. A minimal statistical equivalent for virtual public roads should therefore be defined jointly by road statistics and vehicle response equivalence, rather than by compression, appearance, or average roughness alone.

Keywords: virtual road, road roughness, transient, wavelet, full vehicle simulation, durability.

Acknowledgements

We would like to express our sincere gratitude to everyone who contributed to this thesis project.

First and foremost, We would like to thank Kristian Carlsson for his continuous guidance, valuable discussions, constructive feedback, and support throughout this work. His expertise in durability simulation and road modelling has been essential to the development of this thesis.

We would also like to thank Anders Nord for his advice, technical discussions, and helpful insights during the project.

Special thanks are extended to Björn Pålsson for his support and guidance throughout the thesis work.

We would like to acknowledge Volvo Cars for providing the opportunity to conduct this research, as well as access to the road data, simulation tools, and computational resources required for the study. The work was carried out at the Durability, Safety Centre.

Finally, we are also grateful to each other for our productive discussions and collaboration throughout the project.

Gothenburg, Sweden

June 2026

Houshi He & Qiyu Tan

List of Acronyms

AR(1)	First-order autoregressive model
CRG	Curved Regular Grid road format
IQR	Interquartile Range
IRI	International Roughness Index
ISO	International Organization for Standardization
NSL	Nonstationary Laplace
PSD	Power Spectral Density
QC	Quarter Car
RMS	Root Mean Square
SWT	Stationary Wavelet Transform
WMCR	Wavelet Multi Component Reconstruction

Nomenclature

Road and Frequency Variables

u, v	Longitudinal and lateral road coordinates
$z_i(u)$	One dimensional profile for road i
$Z_i(u, v_k)$	Measured surface at lateral track v_k
$\hat{z}(u)$, $\hat{Z}(u, v)$	Reconstructed profile and surface
λ, n, n_0	Wavelength, spatial frequency, and reference frequency

Compact Road Models

$\mathcal{Q}_i$, $\mathcal{P}_i$	Compact NSL and WMCR parameter sets for road i
C_j, $\bar{C}$	Segment roughness coefficient and mean roughness level
w_j, $\bar{w}$	Segment waviness exponent and mean waviness
r_j, ν, a_r	Roughness intensity, variability, and persistence
L^{src}, L^{tar}, κ	Source length, target length, and compression ratio

WMCR Components

$g_i(u)$, $r_i(u)$, $t_i(u)$	Topography, roughness, and transient components
Θ_i^{g}, Θ_i^{r}, Θ_i^{t}	Geometry, roughness, and transient parameter sets
W_m, A_m, ψ_m	Width, amplitude, and waveform of event m
$C_{\mathrm{bg},m}$, $\hat{u}_m$	Event background roughness and target position

Response Quantities

D_{QC}	Quarter-car pseudo damage
D_{TOT}	Total counted damage measure
$f55$	Total axial load in the front suspension strut
$f6$	Longitudinal load at the lower ball joint
$f63$	Axial load in the front anti-roll bar drop link

Contents

List of Acronyms **vi**

Nomenclature **vii**

1 Introduction **1**

1.1 Background 1

1.2 The Problem of Road Nonstationarity 2

1.3 Implicit Modeling vs. Explicit Reconstruction 5

1.4 Scope and Contributions 6

1.5 Overall Road Compression Concept 8

1.6 Measured Roads and Data Preparation 9

1.7 Evaluation Framework 10

2 NSL Modeling **11**

2.1 NSL Pipeline for Virtual Road Synthesis 11

2.2 From Segment Statistics to NSL Parameters 13

2.3 IRI Based NSL 13

2.3.1 Roughness Coefficient Estimation 14

2.3.2 NSL Parameter Estimation 14

2.3.3 Persistence Calibration 15

2.4 Direct C NSL 15

2.4.1 Segmentwise PSD Fitting 15

2.4.2 Parameter Estimation 16

2.5 Tail Analysis and Calibration 16
2.5.1 Tail Detection 17
2.5.2 AR(1) Persistence Calibration 17
2.5.3 Hybrid Fitting Strategy 18
2.6 Two-Track Synthesis and Road Surface Generation 19
2.6.1 One-Dimensional Profile Synthesis 19
2.6.2 Lateral Coherence and Two-Dimensional Surface 19
2.6.3 Post-synthesis Excitation-Level Calibration 20
2.7 Summary 23

3 Wavelet Multi Component Reconstruction 24
3.1 Road Component Model 24
3.2 Topography 26
3.2.1 Topography Extraction 26
3.2.2 Spline Reconstruction 26
3.3 Wavelet Separation 27
3.3.1 SWT Decomposition 28
3.3.2 Event Detection and Split 29
3.4 Roughness 31
3.4.1 PSD Parameterization 31
3.4.2 Parameter Sampling and Reconstruction 32
3.5 Transient Events 33
3.5.1 Event Library 33
3.5.2 Event Placement and Assembly 34
3.6 2D Road Surface 35
3.6.1 Lateral Coherence 36
3.6.2 Surface Generation 37
3.6.3 Transient Injection and Export 38

4 Simulation and Response Results **41**

4.1 Preliminary Analysis 41

4.1.1 PSD 42

4.1.2 Quarter Car Pseudo Damage 43

4.2 Benchmark and Simulation Setup 45

4.3 Response Results 46

4.3.1 Monitored Response Channels 46

4.3.2 Channel-Based Adams Results 48

4.3.3 Overall Accuracy and Cross-Bin Robustness 51

4.4 Summary 52

5 Discussion and Conclusions **54**

5.1 Discussion 54

5.2 Limitations 55

5.3 Conclusions 56

References **59**

1

Introduction

1.1 Background

Virtual durability development increasingly relies on simulation rather than repeated physical testing. In such a workflow, the road cannot be treated as a secondary input or merely as a geometric boundary condition. It is one of the principal excitation sources governing suspension motion, transmission of wheel forces, body vibration, and the load histories that ultimately drive fatigue damage in structural components [1]. The credibility of virtual durability assessment therefore depends not only on the quality of the vehicle model, but also on how well the simulated road environment captures the excitation characteristics of real customer usage.

Real public road use is broad, heterogeneous, and spatially uneven. During its service life, a production vehicle may encounter highways, urban streets, repaired segments, coarse rural roads, joints, potholes, and other localized harsh events. These conditions vary across markets, customers, seasons, speeds, and route compositions. Even within a single road, excitation severity may change substantially with distance. Bogsjö shows that some road sections can dominate fatigue relevance despite occupying only a limited part of the travelled distance [1]. Charles' work on test severity supports the same engineering concern from the perspective of measured road environments [2]. The engineering challenge is therefore not to reproduce the visible shape of a road, but to represent the excitation mechanisms that matter for dynamics and damage.

From this perspective, the objective of virtual road modelling is not to generate a synthetic curve that merely looks like a road, but to construct a minimal statistical equivalent of the measured road environment. A *statistically equivalent virtual road* is defined here as a compact representation that is efficient enough for repeated simulation while still preserving the excitation characteristics of the measured environment that matter for vehicle response.

1.2 The Problem of Road Nonstationarity

Before introducing a compact road model, it is useful to clarify what is meant by the road profile considered here. Figure 1.1 presents a conceptual decomposition of a measured road centreline into topography, roughness, and texture. Topography describes the slowly varying road shape. Roughness or unevenness denotes the part of the longitudinal profile most directly connected to vehicle excitation and is formalised in ISO 8608, *Mechanical Vibration — Road Surface Profiles — Reporting of Measured Data* [3]. Texture represents finer pavement surface features. ISO 13473-1, *Characterization of Pavement Texture by Use of Surface Profiles — Part 1: Determination of Mean Profile Depth*, defines macrotexture measures such as mean profile depth [4]. ISO 13473-3, *Characterization of Pavement Texture by Use of Surface Profiles — Part 3: Specification and Classification of Profilometers*, specifies profiler requirements for pavement texture measurement [5]. The figure is therefore used as a conceptual guide to the road content considered here, rather than as a strict filtering procedure.

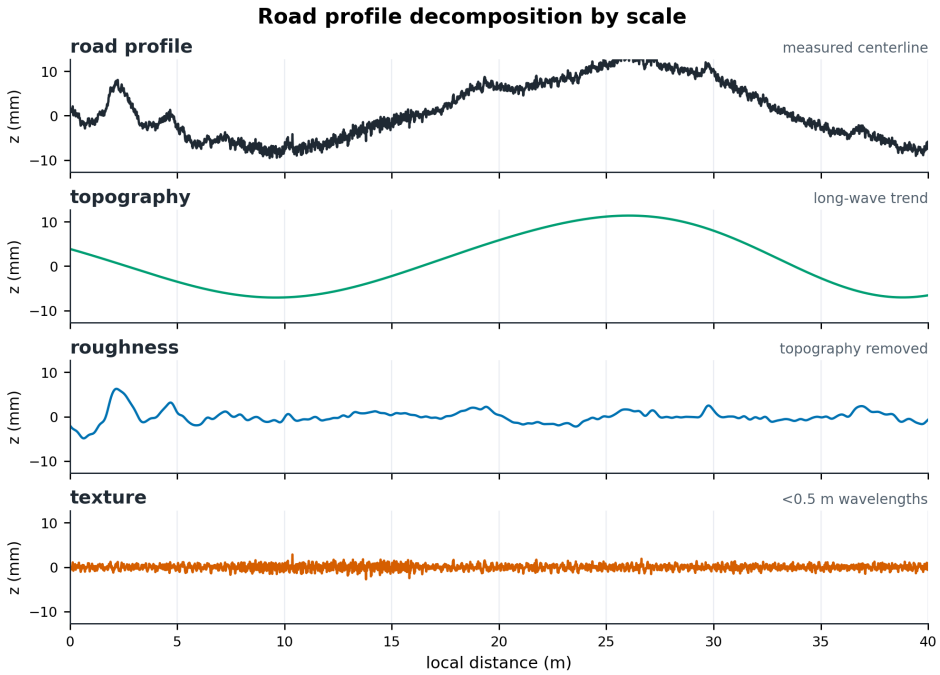

Figure 1.1: Conceptual road profile decomposition.

Within the roughness component of this decomposition, the longitudinal road profile is commonly described through a spatial power spectral density of the form

$$G_d(n) = G_d(n_0)\left(\frac{n}{n_0}\right)^{-w}, \tag{1.1}$$

where $G_d(n)$ is the displacement PSD at spatial frequency n, $G_d(n_0)$ is the roughness coefficient at the reference spatial frequency n_0, and w is the waviness exponent. Figure 1.1 defines the conceptual road components used in the discussion, while

Equation 1.1 and the schematic PSD in Figure 1.2 describe how the roughness component is represented spectrally.

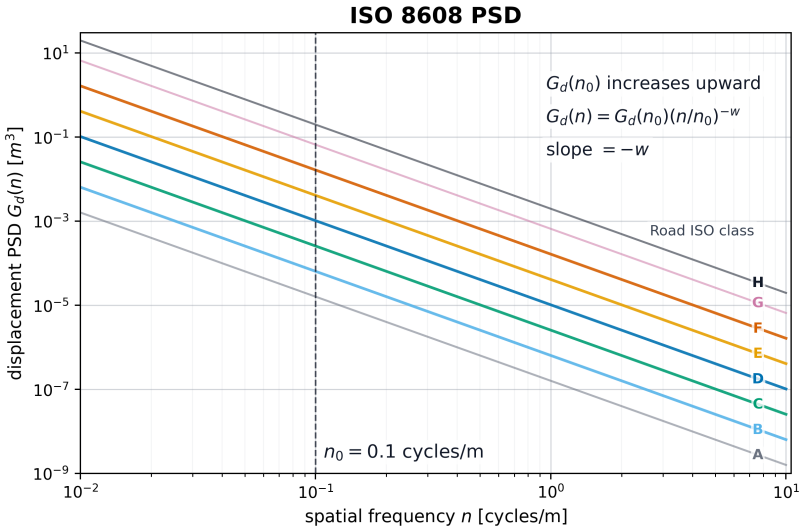

Figure 1.2: ISO 8608 PSD representation for road roughness.

For virtual roads intended for durability analysis, the main difficulty arises in the roughness range because real profiles are not spatially uniform. Measured roads contain localized events such as joints, repaired patches, potholes, and other short but severe departures from the background profile that are naturally interpreted as transients. Bruscella, Rouillard, and Sek show that road profiles contain components that should be separated before interpretation [6]. Johannesson, Podgórski, and Rychlik provide the probabilistic modelling background for such nonstationary road content [7]. For this reason, long measured roads are better viewed as nonstationary processes: the local roughness level changes along the road, and so does the response relevance of the excitation [1].

The literature also shows how road description has moved from stationary roughness models toward response oriented and damage oriented representations. Early random road studies established stationary Gaussian processes and PSD descriptions as the standard language for roughness representation [8]. That line of work was later consolidated in ISO 8608, which provided a common spectral framework for reporting and classifying measured road unevenness [3]. Andrén then showed how practical PSD approximations could support efficient engineering use of that framework for longitudinal profiles [9]. In parallel, response linked indices such as IRI became important because they reduce profile severity to compact measures connected to ride response [10]. Rouillard likewise argued for using predicted ride quality as a practical means of characterizing pavement roughness, reinforcing the engineering appeal of reduced scalar indicators [11].

As attention shifted from classification toward durability relevance, the limitations of purely stationary descriptions became clearer. Charles showed early on that descriptions of the road environment for testing cannot be reduced to average severity

alone when severe events govern loading exposure [2]. Bruscella, Rouillard and Sek emphasized that measured road profiles contain localized features whose analysis requires more than a single global roughness description [6]. Rouillard later proposed decomposition of pavement profiles into a Gaussian sequence, which can be understood as part of the broader effort to separate background roughness from localized or non Gaussian content [12]. In work focused on durability, Bogsjö made the same issue particularly concrete by showing that irregular sections may dominate fatigue damage out of proportion to their road length [1]. Johannesson, Podgórski and Rychlik then placed these observations in a unified probabilistic setting by comparing stationary Gaussian, nonstationary Gaussian, and nonstationary Laplace road models within one modelling family [7].

Taken together, these strands point to the main gap addressed here: compact roughness descriptors and nonstationary modelling theory are well established, but practical methods for generating compact virtual roads that preserve localized structure and are validated in a full vehicle durability setting remain less fully developed.

The usefulness and limitation of IRI can then be stated more precisely. As shown in Fig. 1.3, IRI is defined from a standardised linear quarter car reference model travelling at 80 km/h, so it is already a measure weighted by vehicle response rather than a purely geometric profile statistic [10]. This feature makes IRI attractive in pavement management and roughness screening. The vehicle model used in durability simulation, however, is not a linear quarter car. Even within quarter car analysis, nonlinear suspension and tyre elements improve response prediction relative to the linear approximation, particularly for sprung mass motion [13]. IRI is therefore informative, but it remains a heavily compressed scalar tied to one idealised reference system.

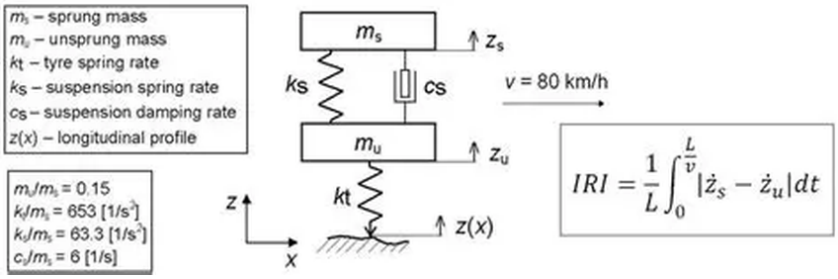

Figure 1.3: Linear quarter-car model used for IRI.

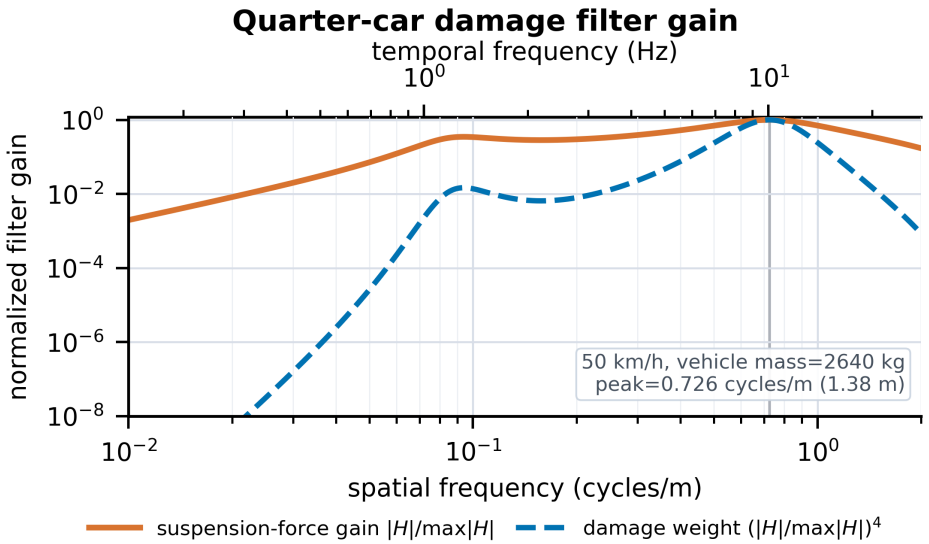

Figure 1.4: Quarter-car force and damage weighting at 50 km/h.

The example frequency weighting for suspension force and damage using a quarter car in Figure 1.4 explains why scalar roughness measures can be useful for screening: they do not treat all wavelengths equally, but emphasize the part of the road spectrum that excites the quarter car response and contributes most strongly to the damage proxy. At the same time, this filtering also shows why such measures cannot solve the synthesis problem by themselves. They compress the road into a response weighted indicator and therefore cannot describe where severe local events occur along the road or how that severity should be retained in a reduced representation.

ISO 8608 goes further than IRI by preserving the spectral description of the profile. This remains connected to the stationary random road tradition established by Dodds and Robson [8]. ISO 8608 makes that spectral description a reporting and classification standard [3]. However, a global PSD or roughness class does not state where isolated severe events occur along the road. The texture standards in ISO 13473 are not intended to solve that problem either. ISO 13473-1 focuses on macrotexture depth and is explicitly insensitive to microtexture and unevenness characteristics [4]. ISO 13473-3 addresses instruments for microtexture, macrotexture, and megatexture and excludes road unevenness from its scope [5]. Consequently, neither IRI, nor ISO 8608, nor texture descriptors by themselves ensure that transient content is retained in a reduced road representation.

1.3 Implicit Modeling vs. Explicit Reconstruction

Once the nonstationary nature of road excitation is acknowledged, two broad technical routes emerge. The first is an implicit stochastic route, in which nonstationarity is carried by the probability model itself rather than by explicitly listing every local

defect. This is the viewpoint emphasized in Johannesson, Podgórski and Rychlik's paper on Laplace distributions, where stationary Gaussian, nonstationary Gaussian, and nonstationary Laplace formulations are treated within one modelling family for road topography and roughness [7]. In this route, the key modelling decision concerns how the local variance or roughness level evolves as a random process. Behaviour that looks transient then appears through the stochastic structure rather than being inserted one event at a time.

The second is an explicit route, in which transient content is added by construction. In Bogsjö's work on fatigue, irregular sections are first identified from measured roads and then incorporated into new road models as explicitly inserted irregular sections at random locations, precisely because those sections dominate damage more than their road length would suggest [1]. Fauriat's framework for simulation differs in objective but is similar in spirit. Instead of relying only on a compact global roughness description, it builds loads induced by the road from mathematically constructed stochastic inputs for road topography and operating conditions, so that variability is generated through explicit scenario construction and then propagated through vehicle simulation [14].

These routes correspond to different modelling philosophies. The implicit route is attractive when the objective is a compact stochastic description with a small set of interpretable parameters. The explicit route is attractive when transient content, localization, or constructed load scenarios must remain visible in the reduced representation. Both aim at statistical equivalence, but they preserve different aspects of the measured road and therefore lead to different synthesis strategies.

1.4 Scope and Contributions

The title phrase "minimal statistical equivalent" is interpreted as an engineering target rather than as a purely mathematical identity. Minimal means that the measured road is reduced to the information needed for efficient virtual road generation. Statistical equivalent means that the reduced road must preserve the road statistics and response-relevant excitation needed for durability simulation, such that results from a shorter simulation can be scaled to represent the full measured driving distance. Two methodological directions are considered for this purpose.

Throughout this work, road slope and curvature are not treated as modelling targets. The measured roads are used after these large scale alignment components have been removed, so the analysis focuses on the roughness and localized excitation content relevant to the durability simulations.

The first contribution is an enhanced Nonstationary Laplace (NSL) framework. It builds on the Laplace road profile modelling line of Johannesson and Rychlik, whose work introduces Laplace processes for road profile description and connects this mod-

elling line to practical roughness-based road identification [15, 16]. Bogsjö, Podgórski and Rychlik further showed that Laplace excitation can represent road roughness with stronger transient content than the classical Gaussian description [17]. Here, NSL is developed into a complete synthesis pipeline with two baselines based on IRI and a direct spectral roughness coefficient route. The comparison asks whether compact descriptor statistics are sufficient when scanned road profiles are available, or whether local structure critical to response is lost when the road is represented mainly by segment statistics [3].

The second contribution is Wavelet Multi Component Reconstruction (WMCR). WMCR represents the explicit route and is developed as a dedicated reconstruction strategy for nonstationary roads. Öijer and Edlund identify transient road obstacle distributions and link them to durability and comfort, which supports treating localized obstacles explicitly [18]. Continuous wavelet analysis of pavement profiles shows how local pavement features can be studied across scale, which supports the use of wavelets for the separation step [19]. WMCR is formulated so that the generated profiles remain interpretable within the ISO 8608 roughness setting [3]. Its central idea is to keep the physical origin of the road content visible: long wavelength topography, background roughness, localized transients, and lateral surface structure are handled as separate components before being recombined for simulation. Rather than allowing localized severity to be absorbed mainly into the parameter evolution of a random model, WMCR seeks to preserve how such structure appears along the road, at different scales, and in a form that remains accessible for engineering interpretation.

The third contribution is the common validation framework used to compare the two routes. The question is not only whether a synthetic road can reproduce overall statistical levels, but also whether localized structure at different scales remains present when the end use is durability simulation. The workflow therefore combines road diagnostics, quarter car pseudo damage, and full vehicle response metrics. The final benchmark applies this framework to Adams Car simulations over the 30, 50, 70, and 100 km/h speed bins using monitored force channels that separate vertical strut loading, longitudinal lower ball joint loading, and axial loading in the stabilizer-bar drop link.

The work is therefore both methodological and comparative. It develops NSL as an implicit stochastic pipeline, develops WMCR as an explicit reconstruction pipeline, and evaluates both within the same identification, generation, and validation workflow. The intended conclusion is not a universal proof that one generated road is equivalent for all vehicles and operating conditions. It is a benchmarked assessment of which road features must remain explicit when compact virtual roads are used as durability inputs. Figure 1.5 summarizes this structure before the individual modelling chapters are introduced. The baseline branch keeps the measured road as the reference for comparison, the NSL branch compresses the road through stochastic roughness descriptors, and the WMCR branch keeps topography, background roughness, transients, and lateral structure as separate reconstruction components

before the common response evaluation.

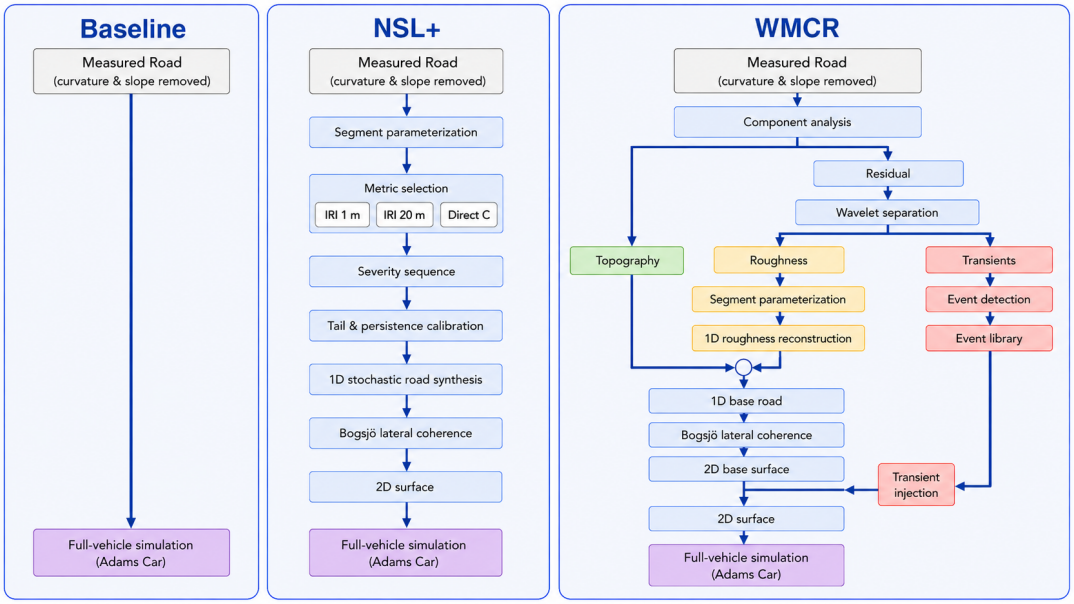

Figure 1.5: Overall comparison pipeline.

1.5 Overall Road Compression Concept

The compression problem addressed in this thesis starts from a practical simulation constraint. A measured public-road database may contain hundreds or thousands of kilometres of road data. Running all measured roads directly in full vehicle durability simulation would be too expensive and too slow. A compact virtual road is therefore needed, but the aim is not simply to make the road as short as possible. For durability, what matters is not total damage accumulated over the original distance, but the damage intensity per unit distance, for example damage per kilometre.

The compression target is therefore to transform a long measured source road of length L_{src} into a shorter virtual road of length L_{tar}, with a compression ratio $\kappa > 1$ such that

$$L_{\mathrm{tar}} = \frac{L_{\mathrm{src}}}{\kappa}. \tag{1.2}$$

This reduction should preserve the durability-relevant excitation rate of the source road. Low-importance sections that contribute little to response damage may be represented more compactly, while severe sections and transient events that dominate fatigue damage must remain present in the reduced road. In this sense, compression means preserving equivalent durability intensity, not preserving total road length or visual similarity.

This concept also explains why two complementary modelling routes are developed. NSL performs the reduction mainly through compact stochastic parameters that

describe the evolution of road severity. WMCR performs the reduction through an explicit component representation in which topography, background roughness, transient events, and lateral surface structure are separated before reconstruction. Both routes therefore follow the same compression objective: to produce the shortest practical virtual road that still preserves equivalent response-relevant severity, such as damage per kilometre, for the target simulation workflow.

1.6 Measured Roads and Data Preparation

The empirical basis is a set of high resolution scanned roads. In total, the data set contains 26 measured road sections with a combined length of approximately 26 km. The surface data are available at a spatial resolution of 5×5 mm, which makes it possible to retain both the broad roughness organization of the road and the localized features that are important for response analysis in durability work. Figure 1.6 shows one representative measured CRG surface in the Adams environment.

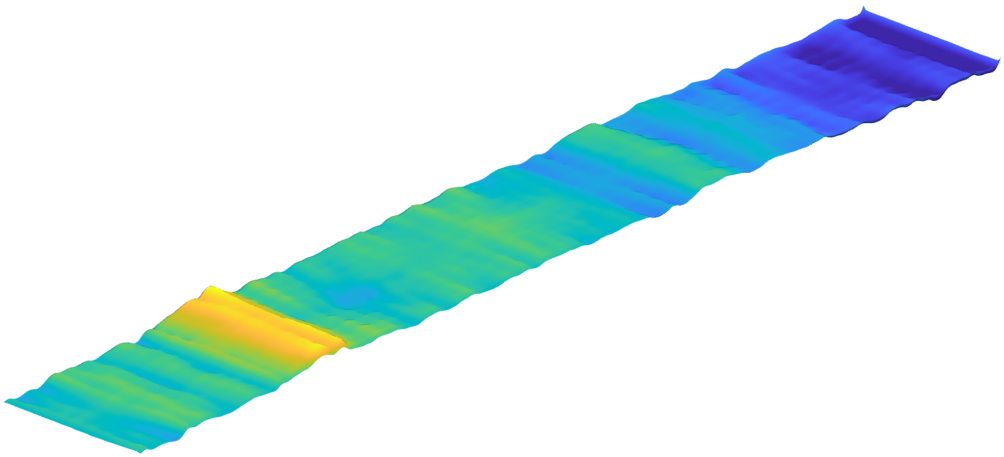

Figure 1.6: Measured CRG road visualised in Adams.

For identification and comparison, the measured roads are organized into four speed bins: 30, 50, 70, and 100 km/h. This grouping by speed is not introduced merely as a reporting convenience. Vehicle response is speed dependent, so the same road profile can lead to different excitation and damage relevance at different operating speeds. Higher-speed roads are generally smoother in terms of roughness level, although vehicle response characteristics may still differ due to speed-dependent dynamic effects. The speed bins should therefore be understood as operational categories that also reflect differences in road type, rather than as interchangeable samples from one pooled road population. The bins are kept separate so that road severity and response effects are compared within internally consistent operating conditions rather than pooled indiscriminately across all cases.

Before modelling, the measured roads are prepared for statistical identification, synthesis, and simulation through a common preprocessing workflow. That workflow

includes road selection within each speed bin, harmonization of longitudinal sampling, separation of left and right tracks where required, and conversion into formats suitable for both reduced road analysis and full vehicle simulation. The prepared data set is therefore used as the common basis for both NSL and WMCR. The speed bins support parameter identification and road comparison, while the final full vehicle response assessment is reported across the four speed-bin benchmark cases. This distinction is important methodologically, because it helps separate differences caused by the modelling route from differences caused by inconsistent road preparation or overly general validation claims.

1.7 Evaluation Framework

Proposing new modelling methods is not meaningful without a benchmark operating at the same level as the intended application. If comparison is limited to profile statistics or compact roughness indicators, important differences in load relevance may remain hidden. For road modelling aimed at durability analysis, the decisive question is whether the generated road preserves the excitation features that matter for vehicle response.

For that reason, the evaluation does not rely on profile descriptors alone. Road induced fatigue damage provides a direct response oriented validation perspective [20], and stochastic road descriptions can be combined with response relevant validation [17, 15, 16]. Reduced indicators remain useful as first screens, but their limitations must be kept in view. IRI is tied to a linear quarter car reference system [10], and even quarter car analysis can change when nonlinear suspension and tyre elements are included [13]. Fatigue evaluation therefore requires a further step from reduced indicators toward load histories and cycle counting, using the Palmgren–Miner framework as the common engineering basis for cumulative damage [21, 22]. Bogsjö's study on durability likewise shows that road models should ultimately be judged by their consequences for vehicle response relevant to fatigue rather than by profile appearance alone [1].

The evaluation framework adopted here is therefore deliberately broader. Model quality is assessed through both statistical agreement in the road description and equivalence in simulation metrics relevant to durability. The final assessment is not based only on a linear quarter car proxy, but on full vehicle simulation in Adams Car, where nonlinear subsystems and multiple response channels can be represented simultaneously. The final benchmark evaluates damage and range metrics for three force channels: $f55$ for vertical strut loading, $f6$ for gradient sensitive longitudinal lower ball joint loading, and $f63$ for left–right differential excitation through the front drop link. The evaluation framework therefore serves as the common basis on which the NSL and WMCR routes are assessed, while also making the scope of the final conclusion explicitly tied to the tested benchmark.

2

NSL Modeling

2.1 NSL Pipeline for Virtual Road Synthesis

The NSL road modelling strategy follows the Laplace road framework introduced by Johannesson and Rychlik [15] and [16], and its later extensions by Johannesson, Podgórski, and Rychlik and by Bogsjö, Podgórski, and Rychlik [7] and [17], where a Gaussian spectral road model is extended by random variance modulation between segments. Here, NSL is developed as a complete virtual road synthesis pipeline.

The pipeline operates on measured road profiles stored in the CRG format. For each measured road i, the left and right wheel-track profiles are extracted from the measured surface $Z_i(u, v)$, where u is the longitudinal road coordinate and v is the lateral coordinate. The wheel-track sampling positions are defined relative to a reference path centreline as

$$v_L^{\text{veh}} = -0.75 \text{ m}, \qquad v_R^{\text{veh}} = +0.75 \text{ m}, \tag{2.1}$$

corresponding to a representative track width of approximately 1.5 m.

To increase the number of available realizations from each measured road, additional paths are generated by laterally shifting the reference path relative to the CRG centreline. For a path offset v_o, the extracted profiles are

$$z_{i,L}^{(v_o)}(u) = Z_i(u,\, v_o + v_L^{\text{veh}}), \qquad z_{i,R}^{(v_o)}(u) = Z_i(u,\, v_o + v_R^{\text{veh}}). \tag{2.2}$$

Five offsets,

$$v_o \in \{-0.6, -0.3, 0, 0.3, 0.6\} \text{ m}, \tag{2.3}$$

are used throughout this work. The offsets are chosen to provide multiple lateral realizations while remaining within the valid lateral extent of the measured CRG data. After ramp trimming at both ends, each extracted left–right track pair forms an input realization for the NSL parameterisation stage.

Two parameterisation routes are considered. The first route starts from roughness characterisation based on IRI and then maps the local IRI values to NSL parameters. The second route bypasses IRI and starts directly from roughness coefficients estimated from road profile spectra in each segment. The first route is closer to a

response oriented proxy description of roughness, while the second route is more directly tied to the road itself. This distinction addresses a deeper methodological question: when scanned road profiles are available, is it still necessary to pass through IRI, or can one construct a statistically equivalent virtual road directly from spectral road statistics?

Figure 2.1 summarises the three pipeline variants. Pipeline A (IRI-1m) and Pipeline B (IRI-20m) follow the IRI route at two different segment resolutions. Pipeline C follows the Direct C route and carries the main methodological weight. All three pipelines share the same synthesis and road surface generation backend, but differ in how the input road statistics are estimated and compressed.

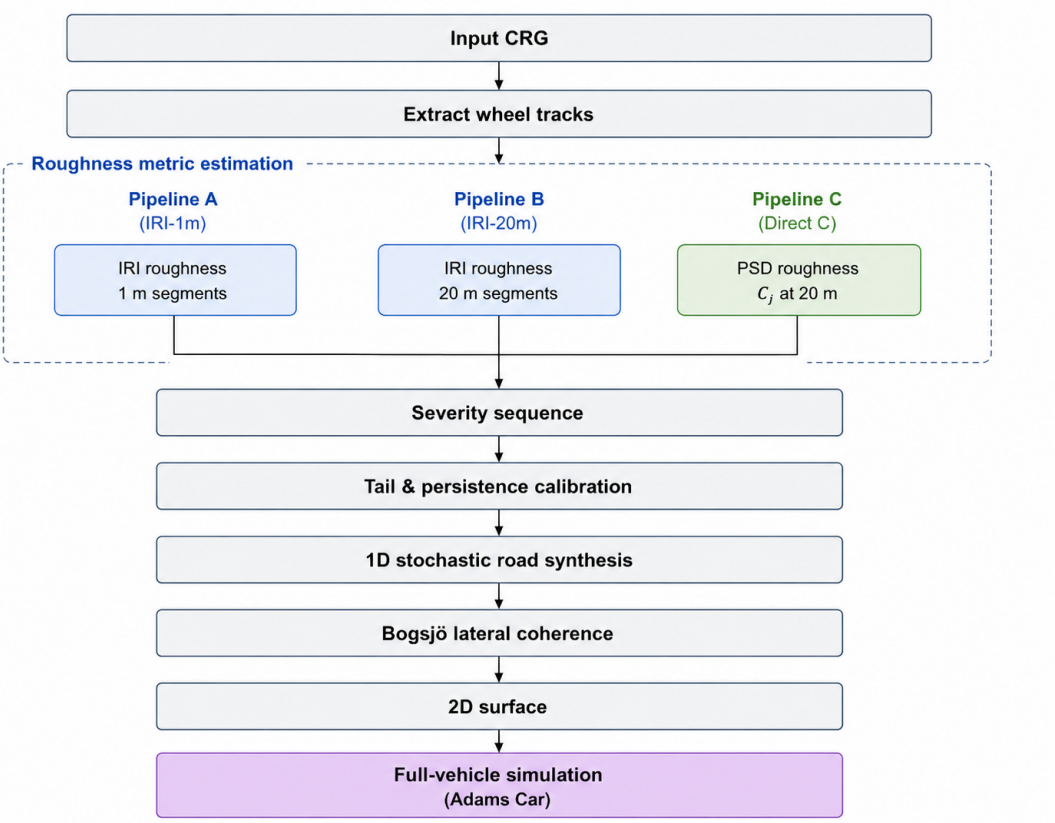

Figure 2.1: Parallel NSL workflows.

For each measured road i, the NSL pipeline produces a compact parameter set

$$\mathcal{Q}_i = \left(\bar{C}_i,\ \nu_i,\ a_{r,i},\ \bar{w}_i,\ C_{\mathrm{dif,hp},i},\ \mathrm{IRI}_{\mathrm{ref},i}\right), \tag{2.4}$$

where $\bar{C}$ is the mean roughness coefficient, ν is the nonstationary variability parameter, a_r is the AR(1) persistence coefficient, $\bar{w}$ is the mean spectral waviness, $C_{\mathrm{dif,hp}}$ is the high frequency differential roughness variance used for anti roll bar calibration, and $\mathrm{IRI}_{\mathrm{ref}}$ is the measured IRI vector by segment used as the low frequency calibration target. This parameter set is passed from the parameterisation stage to the synthesis stage. Unlike the WMCR parameter set, it does not retain

explicit topography or transient event templates; instead, it captures the statistical character of the road through a compact stochastic description.

2.2 From Segment Statistics to NSL Parameters

The NSL formulation is based on the idea that a road can be viewed as locally stationary but globally nonstationary. Within each segment, the road roughness is represented by a Gaussian spectral model. Across segments, the roughness intensity is allowed to vary through a random amplitude process. If the variance factors in the segments are gamma distributed, the resulting process becomes Laplace distributed rather than purely Gaussian. This enables the model to reproduce roughness fields with stronger intermittency and heavier tails than a stationary Gaussian model.

Let the local variance of segment j be written as

$$\sigma_j^2 = r_j \, \sigma^2, \tag{2.5}$$

where σ^2 is the variance implied by the mean road roughness level and r_j is a normalised local intensity factor. In the stochastic Laplace interpretation, r_j is replaced by a gamma random variable with mean one and variance ν. The parameter ν therefore governs how strongly roughness intensity fluctuates between segments. Large values of ν produce a more visibly nonstationary road with stronger alternation between harsh and smooth sections.

The sequence of segment intensities is not assumed to be independent in general. Real road quality often varies gradually, which means that rough segments tend to persist over multiple neighbouring blocks rather than appear as isolated single events. To account for this, the roughness evolution by segment is modelled using an AR(1) persistence description. The associated parameter a_r controls the spatial memory of roughness intensity. This addition is important because two roads with similar marginal roughness distributions may still differ strongly in durability relevance if one has isolated harsh segments while the other contains persistent rough regions.

The practical challenge is therefore not simply to estimate the mean roughness level, but to estimate the whole structure of roughness variation by segment in a way that remains useful for synthetic road generation. Two distinct but related parameterisation routes are used for this purpose.

2.3 IRI Based NSL

The route based on IRI uses IRI by segment as a compact descriptor of road severity. This route is convenient because IRI is widely available in road databases and has a

direct engineering interpretation. It is evaluated at two segment resolutions, referred to as Pipeline A (1 m segments) and Pipeline B (20 m segments).

2.3.1 Roughness Coefficient Estimation

Following Johannesson and Rychlik [16], with IRI itself based on the standard quarter-car roughness definition [23] and [10], the segment-wise roughness coefficient is estimated from the local IRI value by

$$\hat{C}_j = 10^{-6} \left(\frac{\hat{I}_j}{2.21} \right)^2, \tag{2.6}$$

for the commonly used waviness case $w = 2$, where $\hat{I}_j$ is the local IRI estimate in mm/m and $\hat{C}_j$ is the corresponding roughness coefficient.

The two IRI-based routes are retained for different methodological reasons. Pipeline B uses IRI on 20 m segments, which corresponds to a stable and widely available roughness descriptor in road databases and therefore provides an industry-standard benchmark for the NSL workflow. Its limitation is that local transients and short severe roughness variations may be smoothed by the segment averaging. Pipeline A uses IRI on 1 m segments to test whether a finer IRI entry point can better preserve severe roughness sections that may be important for durability. This route has higher local sensitivity, but it is also more sensitive to noise and less statistically stable. The comparison is therefore not intended to determine which IRI resolution is universally more accurate, but to evaluate whether shorter IRI segmentation improves retention of durability-relevant roughness content before the later tail protection and synthesis stages.

2.3.2 NSL Parameter Estimation

Once the local coefficients have been estimated, the global mean roughness and the Laplace variability parameter are obtained as

$$\hat{C} = \frac{1}{M} \sum_{j=1}^{M} \hat{C}_j, \tag{2.7}$$

$$\hat{\nu} = \frac{1}{M-1} \frac{\sum_{j=1}^{M} (\hat{C}_j - \hat{C})^2}{\hat{C}^2}. \tag{2.8}$$

The waviness exponent is fixed at $w = 2.0$ for the IRI-based route, consistent with the conversion formula. A normalised roughness intensity sequence $r_j = \hat{C}_j / \hat{C}$ is then formed as the input to the persistence calibration stage.

2.3.3 Persistence Calibration

The persistence coefficient a_r is calibrated from the normalised sequence $r_j = \hat{C}_j/\hat{C}$ using a multi-objective grid search. Let $\hat{\rho}_1$ denote the measured lag-one autocorrelation of r_j, $\hat{s}_{\log}$ the standard deviation of $\log r_j$, $\hat{\kappa}$ the kurtosis, and $\hat{Q}_p$ the p-th quantile. For each candidate value a_r, a long synthetic intensity sequence is generated and the following weighted absolute-error objective is minimised,

$$\mathcal{J}(a_r) = |\rho_1^{\text{syn}} - \hat{\rho}_1| + 0.5\,\frac{\left|s_{\log}^{\text{syn}} - \hat{s}_{\log}\right|}{|\hat{s}_{\log}|} + 0.5\,\frac{|\kappa^{\text{syn}} - \hat{\kappa}|}{|\hat{\kappa}|} + \sum_{p\,\in\,\{0.90,\,0.95\}} \frac{\left|Q_p^{\text{syn}} - \hat{Q}_p\right|}{\hat{Q}_p}, \quad (2.9)$$

where all statistics are evaluated from the respective normalised intensity sequences. The log-standard deviation $s_{\log}$ is used rather than the linear standard deviation, so that the objective is insensitive to scale and gives equal weight to proportional deviations across a wide dynamic range. The value of a_r that minimises $\mathcal{J}$ is retained as the fitted persistence coefficient. The search is performed first on a coarse grid and then refined on a fine grid centred on the coarse minimum. This route is therefore not a simple IRI fit. It is an NSL reconstruction pipeline in which IRI acts only as the entry point to a richer stochastic description.

2.4 Direct C NSL

The second route, which is central to the contribution of this thesis, bypasses IRI entirely and starts directly from the road spectrum. If full scanned road profiles are available, it is natural to ask whether one can estimate the segment-wise roughness level directly from the profile itself rather than first condensing it into IRI. The Direct C route answers this question positively.

2.4.1 Segmentwise PSD Fitting

In this approach, the left and right wheel track profiles are first extracted from the measured CRG road. Each track is then divided into fixed-length blocks, typically 20 m, corresponding to the NSL segment length used in the synthesis stage. For every segment, the local power spectral density is estimated using a Welch-type PSD procedure. The local spectrum is then fitted in log-log space by a power-law model of the form commonly used for road-profile PSD descriptions [8], [9], and [3]

$$S(\Omega) = C\left(\frac{\Omega}{\Omega_0}\right)^{-w}, \quad (2.10)$$

where $S(\Omega)$ is the spatial PSD, C is the roughness level, w is the waviness or spectral slope, and Ω_0 is the reference spatial frequency. The result of this fitting step is therefore a pair of segment-wise statistics (C_j, w_j) for each road block.

This is the key conceptual difference from the IRI-based route. Instead of estimating a vehicle-filtered roughness proxy and then converting it back into a road parameter, the Direct C route stays in the road-statistics domain from the beginning. As a consequence, it retains not only information about roughness intensity, but also information about spectral slope. This makes the method more closely tied to the physical road profile and less dependent on a specific intermediary roughness index.

2.4.2 Parameter Estimation

The segment-wise C_j values are then treated as the central roughness intensity sequence for NSL modelling. The mean roughness is estimated by

$$C_{\text{mean}} = \frac{1}{M} \sum_{j=1}^{M} C_j, \tag{2.11}$$

and the non-stationary variability parameter is estimated by

$$\nu = \frac{\text{var}(C_j)}{C_{\text{mean}}^2}. \tag{2.12}$$

A normalised roughness intensity sequence is formed as

$$r_j = \frac{C_j}{C_{\text{mean}}}. \tag{2.13}$$

This normalised sequence is then used in the same persistence calibration framework as the IRI-based method. In other words, the Direct C route preserves the NSL structure of the model but changes the statistical entry point from IRI space to PSD-C-space.

An additional consequence of this route is that the mean waviness can be estimated directly from the measured road. The w_j values are estimated together with C_j for each segment, and a representative mean waviness is used in the synthesis stage. This is a deliberate engineering simplification. It preserves the dominant spectral slope of the measured road while avoiding the complexity of allowing waviness to vary independently from segment to segment in the final synthetic generation.

2.5 Tail Analysis and Calibration

A major extension is that the NSL pipeline does not fit all segments as one homogeneous population. Instead, after the roughness descriptors have been extracted for each segment, the method performs an explicit tail analysis. The result is an NSL parameterisation with explicit persistence and tail treatment, which is a central engineering contribution.

2.5.1 Tail Detection

In the Direct C route, tail analysis is performed on the pooled sequence of left and right track roughness coefficients C_j. A Tukey IQR fence is applied to identify segments whose roughness level is unusually large relative to the bulk of the road [24]. A segment is labelled a tail segment if

$$C_j > Q_{0.75} + \delta_{\text{iqr}} \cdot \text{IQR}, \tag{2.14}$$

where $Q_{0.75}$ is the upper quartile, $\text{IQR} = Q_{0.75} - Q_{0.25}$ is the interquartile range, and δ_{iqr} is a tunable multiplier set to 1.5 here. Segments that exceed this threshold are labelled tail segments, while the rest form the normal set.

This step is important because road behaviour sensitive to fatigue is often dominated by relatively rare but severe segments [18] and [20]. If all segments were forced into a single fitted population, the tail would tend to be smoothed away. By treating normal and tail segments separately, the method preserves the asymmetry between common roughness and rare harsh roughness. This is not only a statistical improvement, but also a durability oriented one.

Figure 2.2 illustrates this tail-detection step by showing how the Tukey fence separates the normal roughness population from the high-severity tail segments.

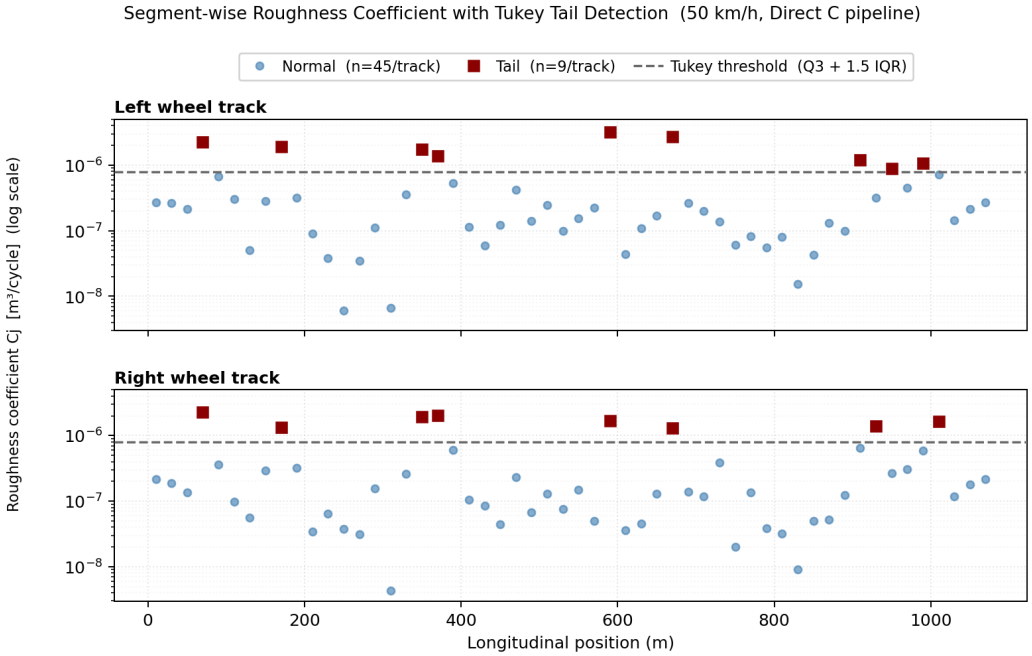

Figure 2.2: Direct C tail detection using the Tukey IQR fence.

2.5.2 AR(1) Persistence Calibration

After the tail split, persistence is calibrated separately for the normal and tail populations using the same multi-objective grid search described in Equation 2.9. In

the Direct C route, the target sequence is based on $r_j = C_j/C_{\text{mean}}$, and the AR(1) coefficient a_r is selected so that the simulated sequence reproduces multiple features of the measured one: lag-one correlation, standard deviation, kurtosis, and high quantiles at the 90th and 95th percentiles. This is a crucial part of the method because it ensures that the fitted NSL model preserves the way roughness severity evolves in space, not merely its overall level.

The underlying Laplace model from Johannesson and Rychlik [15] and [16] provides the stochastic mechanism, and the calibration procedure adapts it explicitly for road synthesis and later durability assessment.

2.5.3 Hybrid Fitting Strategy

Once the normal and tail populations have been defined and the persistence coefficient has been calibrated, the Direct C route fits NSL models separately for the two regimes. The fitting strategy is intentionally low-dimensional. The physically interpretable parameters are fixed directly from the data: the mean roughness level C, the variability parameter ν, and the persistence coefficient a_r. In addition, the mean spectral slope w is estimated from the measured segment spectra. The numerical search is then restricted to a small set of remaining shaping parameters, primarily an amplitude scale factor and a lower bound on the local roughness factor.

This design is important for both stability and interpretability. A high-dimensional black-box optimisation could in principle be used, but it would be much harder to interpret and more sensitive to local minima. By separating analytical estimation from low-dimensional search, the Direct C pipeline remains transparent. The final fitted model still reflects measured road statistics rather than hidden numerical artefacts.

The fitting objective is also defined in the same statistical space as the model input. For each candidate parameter set, synthetic left and right tracks are generated, and new segment-wise C_j values are extracted from the synthetic profiles using the same PSD fitting procedure applied to the measured road. The synthetic and measured C_j distributions are then compared in terms of logarithmic mean, logarithmic spread, and upper quantiles. In this way, the method forms a closed loop entirely in C-space. This is one of the most conceptually important differences from the IRI-based route: the Direct C route does not merely start from C_j; it also evaluates itself in the same space.

2.6 Two-Track Synthesis and Road Surface Generation

The NSL pipeline generates synthetic roads in two stages. The first stage produces a pair of one dimensional longitudinal profiles for the left and right wheel tracks. The second stage expands these profiles into a laterally coherent two dimensional surface suitable for export as an OpenCRG file and for import into Adams Car.

2.6.1 One-Dimensional Profile Synthesis

Synthesis begins by drawing a segment-wise intensity sequence $\{r_j\}$ from the fitted AR(1)-gamma NSL model. The normal regime and the tail regime each contribute blocks in proportion to their fitted occupancy fractions. The intensity r_j of each block scales the local spectral amplitude before the profile samples are generated.

Within each segment, a stationary Gaussian profile is synthesised from a power-law spectral shape governed by the fitted mean roughness coefficient $\bar{C}$ and mean waviness $\bar{w}$. The local spectral amplitude is then scaled by $\sqrt{r_j}$ so that the variance of each synthesised segment reflects the drawn intensity. Adjacent segments are stitched with a smooth height-continuity transition so that discontinuities in elevation and slope do not create artificial transients at block boundaries. Extreme local amplitudes are clipped to avoid unrealistic spikes that would be numerically problematic in vehicle simulation.

The synthesis is performed for left and right tracks using the same intensity sequence $\{r_j\}$, which ensures that both tracks experience the same pattern of rough and smooth regions along the road. The spectral noise draws within each segment are treated differently depending on spatial frequency, as described in the next subsection.

2.6.2 Lateral Coherence and Two-Dimensional Surface

The left and right track profiles are generated through a frequency-band decomposition. At a cutoff wavelength of $\lambda_c = 3.3$ m, each track profile is split into a low-frequency band ($\lambda > \lambda_c$) and a high-frequency band ($\lambda < \lambda_c$). The two bands are treated differently because measured roads show distinct lateral coherence behaviour at different spatial scales.

In the low-frequency band, the left and right profiles share a common roughness component. Following Bogsjö [1], and consistent with later two-track roughness modelling work [17], the lateral coherence at spatial frequency n and track separation

d is modelled as

$$\gamma^2(n,\, d) = \exp(-2\beta_\gamma\, d\, n)\,, \tag{2.15}$$

where β_γ is the coherence decay parameter identified from measured roads and $d = |v_L - v_R| = 1.5$ m is the lateral track separation. In this work, the Bogsjö coherence parameter is not treated as a single global constant; instead, it is estimated separately for each measured road and each CRG file before the corresponding synthetic road is generated. The left and right low-frequency profiles are written as

$$z_{L,\mathrm{lp}} = z_{\mathrm{com}} + z_{\mathrm{dif}}, \qquad z_{R,\mathrm{lp}} = z_{\mathrm{com}} - z_{\mathrm{dif}}, \tag{2.16}$$

where z_{com} is a common roughness component shared between the two tracks and z_{dif} is a differential component. The ratio of common to differential spectral amplitude is set to reproduce the Bogsjö coherence target at each frequency. In the high-frequency band, statistically independent noise draws are used for the left and right tracks, consistent with the rapid coherence decay observed in measured roads at short spatial wavelengths.

The one-dimensional left and right profiles serve as conditioning signals for the two-dimensional surface generation. The road surface is expanded over a lateral grid spanning from $v_L - \Delta v$ to $v_R + \Delta v$, with $\Delta v = 0.65$ m as the lateral margin beyond the wheel tracks. Five lateral anchor positions are defined: the two margins, the two wheel track positions, and the road centre. The elevation at the centre anchor is set as the arithmetic mean of the two wheel track profiles. Between anchors, the surface elevation is obtained by piecewise linear blending of the anchor values, supplemented by a small stochastic conditioning component that introduces spatially correlated lateral texture consistent with the measured differential roughness between adjacent anchors. The assembled two-dimensional surface is

$$\hat{Z}(u, v) = \hat{Z}_{\mathrm{lp}}(u, v) + \hat{Z}_{\mathrm{hp}}(u, v), \tag{2.17}$$

where $\hat{Z}_{\mathrm{lp}}$ and $\hat{Z}_{\mathrm{hp}}$ are the low-frequency surface and the high-frequency surface, respectively. Before export, ramps are added at both ends of the road to allow Adams Car to initialise the suspension. The final surface is converted to OpenCRG format for vehicle simulation.

Figure 2.3 shows the resulting synthesised two-dimensional road surface, including the longitudinal road texture and the lateral structure produced from the conditioned wheel-track profiles.

2.6.3 Post-synthesis Excitation-Level Calibration

After the NSL fitting and stochastic road synthesis steps, an optional post-synthesis gain calibration is applied to correct systematic deviations in road-input excitation level. This calibration is not based on Adams full-vehicle response results, nor is it intended to force agreement with the final validation channels. Instead, it uses road-based and quarter-car-based proxy quantities computed from the measured

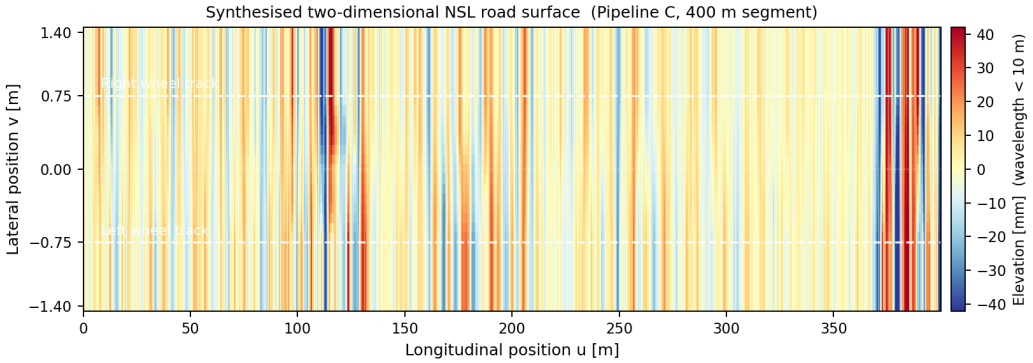

Figure 2.3: Synthesised two-dimensional NSL road surface.

and synthetic wheel-track profiles. The purpose is to preserve durability-relevant excitation levels after stochastic reconstruction while keeping the segment-level NSL structure fixed. The Adams Car force channels are not used as calibration targets in this step; they are reserved for the final independent validation.

The calibration is needed because the NSL parameters control the statistical structure of the road, whereas stochastic synthesis can still introduce systematic amplitude differences. The gain factors therefore quantify how much the synthesised road is modified after generation. They should be interpreted as road-input corrections rather than fitted vehicle-response parameters.

Low-frequency scalar calibration. The first calibration level corrects the overall low-frequency excitation level. Two complementary gain estimates are used. The first is based on a quarter-car damage proxy derived from the International Roughness Index. Both measured and synthetic profiles are passed through the Golden Car quarter-car model at the standard reference speed of 80 km/h [23, 10], and the segment-wise IRI values are accumulated using a fourth-power damage proxy. This gives

$$g_D = \left(\frac{\overline{\mathrm{IRI}_{\mathrm{ref}}^4}}{\overline{\mathrm{IRI}_{\mathrm{syn}}^4}} \right)^{1/4}, \tag{2.18}$$

where the overbar denotes the mean over all segments. The fourth-power weighting is used as a durability-oriented proxy consistent with Palmgren–Miner fatigue accumulation [21, 22]. The second estimate is based on the mean roughness coefficient,

$$g_C = \sqrt{\frac{\overline{C}_{\mathrm{ref}}}{\overline{C}_{\mathrm{syn}}}}. \tag{2.19}$$

The final low-frequency scalar gain is obtained as a weighted geometric average in the log domain,

$$\log g_{\mathrm{lf}} = \frac{w_{\mathrm{IRI}} \log g_D + w_C \log g_C}{w_{\mathrm{IRI}} + w_C}. \tag{2.20}$$

This gain compensates for systematic amplitude bias in the low-frequency band without changing the fitted NSL intensity sequence.

Low-frequency common–differential calibration. A single low-frequency scalar gain cannot independently control the excitation shared by the two wheel tracks and the left–right imbalance between them. Therefore, the low-pass filtered left and right road inputs are decomposed into common and differential modes,

$$z_{\mathrm{com,lp}} = \frac{z_{L,\mathrm{lp}} + z_{R,\mathrm{lp}}}{2}, \qquad z_{\mathrm{dif,lp}} = \frac{z_{L,\mathrm{lp}} - z_{R,\mathrm{lp}}}{2}. \tag{2.21}$$

The common component represents excitation shared by both tracks, while the differential component represents left–right road-input imbalance. These are not Adams response quantities; they are road-input modes used as physically interpretable calibration targets.

Separate low-frequency gains are then applied to the common and differential components,

$$g_{\mathrm{com,lp}} = \sqrt{\frac{C_{\mathrm{com,lp,ref}}}{C_{\mathrm{com,lp,syn}}}}, \qquad g_{\mathrm{dif,lp}} = \sqrt{\frac{C_{\mathrm{dif,lp,ref}}}{C_{\mathrm{dif,lp,syn}}}}, \tag{2.22}$$

where $C_{\mathrm{com,lp}}$ and $C_{\mathrm{dif,lp}}$ denote the low-frequency variances of the common and differential road-input modes, respectively. This step preserves the measured balance between shared low-frequency excitation and left–right low-frequency imbalance.

High-frequency common–differential calibration. The high-frequency calibration is introduced because left–right coherence is especially important for anti-roll-bar excitation. If the high-frequency left and right components are generated independently, the resulting differential variance is determined mainly by the random synthesis procedure rather than by the measured road. Therefore, the high-pass filtered left and right road inputs are also decomposed into common and differential modes,

$$z_{\mathrm{com,hp}} = \frac{z_{L,\mathrm{hp}} + z_{R,\mathrm{hp}}}{2}, \qquad z_{\mathrm{dif,hp}} = \frac{z_{L,\mathrm{hp}} - z_{R,\mathrm{hp}}}{2}. \tag{2.23}$$

The measured high-frequency differential variance is used as an explicit calibration target,

$$C_{\mathrm{dif,hp}} = \mathrm{var}\left(\frac{z_{L,\mathrm{hp}} - z_{R,\mathrm{hp}}}{2}\right). \tag{2.24}$$

The corresponding high-frequency gains are defined as

$$g_{\mathrm{hp,com}} = \sqrt{\frac{C_{\mathrm{com,hp,ref}}}{C_{\mathrm{com,hp,syn}}}}, \qquad g_{\mathrm{hp,dif}} = \sqrt{\frac{C_{\mathrm{dif,hp,ref}}}{C_{\mathrm{dif,hp,syn}}}}. \tag{2.25}$$

By calibrating these two quantities separately, the method preserves the total high-frequency energy while also controlling the left–right differential content produced by the Bogsjö coherence model and the stochastic high-frequency synthesis assumptions.

Reconstruction of calibrated wheel tracks. The final calibrated left and right wheel-track profiles are reconstructed by combining the calibrated low-frequency and high-frequency common–differential components,

$$z_{L,\mathrm{out}} = g_{\mathrm{com,lp}} z_{\mathrm{com,lp}} + g_{\mathrm{dif,lp}} z_{\mathrm{dif,lp}} + g_{\mathrm{hp,com}} z_{\mathrm{com,hp}} + g_{\mathrm{hp,dif}} z_{\mathrm{dif,hp}}, \tag{2.26}$$

$$z_{R,\mathrm{out}} = g_{\mathrm{com,lp}} z_{\mathrm{com,lp}} - g_{\mathrm{dif,lp}} z_{\mathrm{dif,lp}} + g_{\mathrm{hp,com}} z_{\mathrm{com,hp}} - g_{\mathrm{hp,dif}} z_{\mathrm{dif,hp}}. \tag{2.27}$$

In this formulation, the calibration changes only the amplitude of selected road-input modes. It does not re-estimate the NSL segment sequence, tail classification, or persistence structure.

Although the calibration does not use Adams response channels directly, the selected proxy quantities are physically related to the later validation responses: low-frequency common excitation is relevant for strut loading, longitudinal roughness content is relevant for lower ball joint loading, and differential excitation is relevant for anti-roll-bar link loading. The final Adams Car force channels are therefore used only to assess whether these road-input proxy calibrations lead to equivalent durability response, not to tune the generated roads.

2.7 Summary

The NSL virtual road synthesis framework consists of three pipeline variants. All three pipelines share the same tail detection, AR(1) persistence calibration, two track synthesis, lateral coherence model, and gain calibration backend. They differ only in the statistical entry point used to estimate the roughness intensity sequence. Table 2.1 summarises the key characteristics of each variant.

Table 2.1: Comparison of NSL pipeline variants.

	Pipeline A IRI-1m	**Pipeline B** IRI-20m	**Pipeline C** Direct C
Roughness descriptor	IRI at 1 m	IRI at 20 m	PSD C_j at 20 m
Waviness	Fixed ($w = 2$)	Fixed ($w = 2$)	Estimated from PSD
Domain	Vehicle filtered	Vehicle filtered	Road native
Transient sensitivity	High	Moderate	Moderate
Spectral fidelity	Moderate	Moderate	High
Closed loop C space eval.	No	No	Yes

The routes based on IRI (Pipelines A and B) provide a convenient link to established roughness databases and serve as engineering baselines. The Direct C route (Pipeline C) is the central methodological contribution: it estimates road roughness natively in the spectral domain, preserves information about spectral slope, and evaluates the fitted model through a closed loop in C space. Together, the three pipelines allow a direct experimental comparison between proxy based and road native parameterisation in the vehicle simulation results.

3

Wavelet Multi Component Reconstruction

3.1 Road Component Model

Wavelet Multi Component Reconstruction (WMCR) is the explicit reconstruction route for compact virtual road generation. It complements the NSL route by starting from the measured profile and keeping a structured record of the features that are expected to matter for vehicle response. The method is therefore less abstract than a purely statistical road model: it deliberately separates the road into components that have different physical roles in the simulation.

The motivation for this separation is that the dynamic effect of a road is not described well by a single roughness level. Long wavelength topography affects body motion and suspension stroke at low spatial frequencies. Distributed background roughness controls the continuous vibration environment. Localized events, such as joints, patches, pothole like depressions, or sharp bumps, can dominate short response intervals and fatigue relevant load cycles even if they occupy only a small part of the road length. If these mechanisms are compressed into one profile descriptor, the generated road may look statistically plausible but still fail to reproduce the content that drives the measured response.

WMCR addresses this problem by decomposing a measured road into physically distinct components and reconstructing them in a controlled sequence. The purpose is not to reproduce the exact measured profile point by point. Instead, the method preserves the road features that govern vehicle excitation while allowing a shorter road surface to represent the measured population. This makes the generated road compact enough for repeated durability oriented simulation, while still retaining explicit information about the origin of each part of the surface.

Figure 3.1 summarizes the component logic of the method. The measured profile is first separated into a low frequency topography baseline and a residual. The residual is then divided into background roughness and localized transient events. Roughness is described through local spectral parameters, while transient events are retained as explicit features with position, shape, amplitude, and context information. The reconstruction follows the same order: topography is generated first, roughness is

laid on top of it, transients are inserted into the one-dimensional profile, and the result is finally expanded into a laterally coherent two dimensional surface.

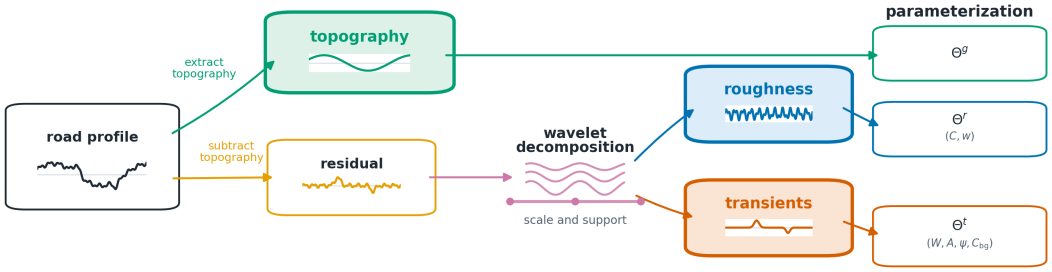

Figure 3.1: WMCR component decomposition workflow.

Let u denote longitudinal position and v denote lateral position. For a measured road i, the aligned surface is denoted by

$$Z_i(u, v_k), \qquad k = 1, \ldots, K_i, \tag{3.1}$$

where v_k are the available lateral track offsets. A representative track is selected for the one-dimensional component analysis,

$$z_i(u) = Z_i(u, v_0), \tag{3.2}$$

where v_0 is the selected lateral offset. On this track, WMCR uses the additive model

$$z_i(u) = g_i(u) + r_i(u) + t_i(u) + \epsilon_i(u). \tag{3.3}$$

Here g_i is the long wavelength topography, r_i is the background roughness, t_i is the transient content, and ϵ_i contains residual measurement noise and small modelling errors.

The component representation associated with a measured road can be written as

$$\mathcal{P}_i = \left(\Theta_i^{\mathrm{g}}, \Theta_i^{\mathrm{r}}, \Theta_i^{\mathrm{t}}, \Theta_i^{\mathrm{v}}\right), \tag{3.4}$$

where Θ_i^{g} stores topography controls, Θ_i^{r} stores roughness parameters, Θ_i^{t} stores transient events, and Θ_i^{v} stores lateral surface information. This parameter set is the compact representation passed from the measurement stage to the synthesis stage. It is intentionally modular: topography, roughness, transients, and lateral structure can be inspected independently, but they are recombined into one road surface before simulation. For a measured source length L^{src} and a compression ratio κ, the compact target length is

$$L^{\mathrm{tar}} = \frac{L^{\mathrm{src}}}{\kappa}, \qquad \kappa > 1. \tag{3.5}$$

The following sections describe this pipeline in the order in which it is used: topography extraction, wavelet-based residual separation, roughness parameterization, transient event reconstruction, and final two dimensional surface generation.

3.2 Topography

3.2.1 Topography Extraction

Road topography represents the long wavelength topographic baseline of the measured profile. It is separated before roughness and transient analysis so that grade, large scale waviness, and elevation drift are not incorrectly interpreted as road texture. For simplicity, this work removes macroscopic slope and curvature from the residual used for roughness and transient modelling. These terms are therefore not treated as separate roughness or event features; the topography component is the part of WMCR intended to handle such slowly varying road shape in the final road surface. With the profile expressed on the longitudinal grid, a low-pass separation operator $\mathcal{L}_{\lambda_g}$ extracts the topography baseline $g(u)$,

$$\begin{aligned} g(u) &= \mathcal{L}_{\lambda_g}\{z(u)\}, \\ x(u) &= z(u) - g(u). \end{aligned} \tag{3.6}$$

The cutoff is specified as a wavelength. Unless otherwise stated, $\lambda_g = 50$ m, so the cutoff spatial frequency is $n_g = 1/\lambda_g = 0.02$ m^{-1}. Components with $n < n_g$ form the topography, while the residual $x(u)$ is passed to roughness fitting and transient detection.

This split is a practical modelling choice rather than a claim that road content has a sharp physical boundary at exactly 50 m. The cutoff is chosen to remove grade and slow geometric drift from the texture analysis while leaving the wavelengths that excite suspension response in the residual. This is important for the later steps: if the baseline is too short, genuine roughness is absorbed into the topography layer; if it is too long, slow elevation drift contaminates the roughness spectrum and the transient detector. The topography extraction is therefore the first step that protects the physical meaning of the subsequent components.

3.2.2 Spline Reconstruction

The extracted topography is not stored as a dense signal. Instead, it is approximated by a smooth spline on a normalized longitudinal coordinate,

$$\Theta^{\mathrm{g}} = \{(\xi_\ell, h_\ell)\}_{\ell=1}^{N_g}, \tag{3.7}$$

where $\xi_\ell \in [0, 1]$ are normalized control locations and h_ℓ are the corresponding low frequency elevation values. The number of control points grows with road length, using approximately one support point per 15 m. This spacing is chosen as a compromise for the baseline representation. It gives several control points over the 50 m topography cutoff wavelength, so grade changes, curvature, and long wavelength undulations can be followed smoothly, but it is still coarse enough to avoid reintroducing roughness and transient content into the topography layer. Lower and upper limits keep the spline representation stable for very short or very long roads.

The normalized coordinate makes the stored topography independent of the original road length. In the synthesis stage, the same type of control set can be mapped to a shorter target interval without copying the measured elevation samples. This is the mechanism by which WMCR compresses the slowly varying road shape while still retaining realistic grade changes and long wavelength undulations. Figure 3.2 illustrates this process: the measured topography is reduced to spline control points, stored as a compact topography description, and then evaluated again to reconstruct the smooth baseline on the target road.

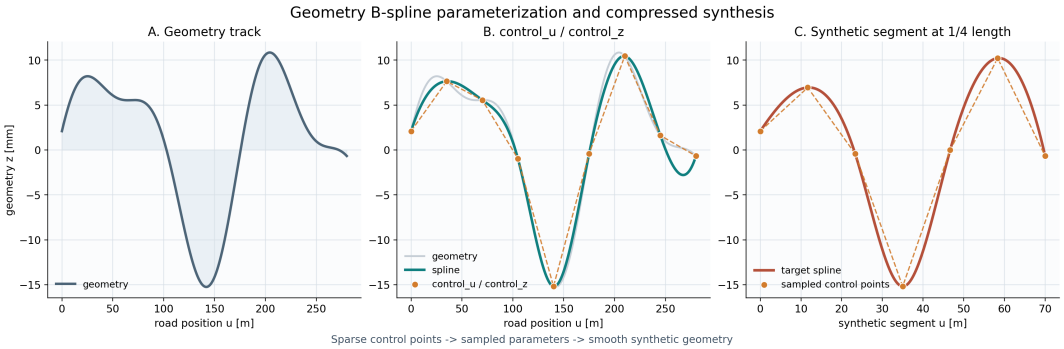

Figure 3.2: B-spline topography compression and reconstruction.

For the compact road, the spline controls are rescaled to the target interval and evaluated as a continuous topography baseline. Adjacent reconstructed intervals are joined with smooth transitions to avoid artificial discontinuities. If several measured lateral tracks are available, their topography traces act as anchors and are interpolated across the target width. Otherwise the centre topography is reused across offsets. The centreline of the reconstructed topography layer becomes the one-dimensional topography component $\hat{g}(u)$.

At this point the reconstruction contains only the smooth road shape. The next step operates on the residual $x(u)$, where the task is to separate continuous texture from localized events before either component is parameterized.

3.3 Wavelet Separation

Wavelet analysis is used here because the residual road profile is neither purely stationary nor purely impulsive. A PSD or Fourier description is effective for showing how variance is distributed over spatial frequency, but it averages over longitudinal position. It cannot say whether a short wavelength contribution is spread over the whole road or concentrated in a single joint, patch edge, pothole like depression, or bump. WMCR needs this positional information because a short localized feature can create a large vehicle response even when its contribution to the total profile energy is small.

The mathematical reason for using wavelets is that they analyse a signal in both position and scale. Mallat's multiresolution framework gives the standard signal

processing basis for separating a signal into coarse approximations and detail bands at successive scales [25]. The use of compact wavelet basis functions with finite filters and vanishing moments follows the same multiresolution setting and the construction of Daubechies type wavelets, which is important for local road profile operations [25, 26]. In a road context, continuous wavelet analysis has also been used to reveal localized features in pavement profiles that are not described by a single global PSD curve [19].

Conceptually, a wavelet coefficient is the local similarity between the residual profile and a shifted, scaled wavelet,

$$W_{j,k} = \langle x, \psi_{j,k} \rangle. \tag{3.8}$$

Here j denotes scale, k denotes longitudinal location on the sampled grid, and $\psi_{j,k}$ is the wavelet at that scale and location. A large coefficient means that the residual locally resembles the wavelet shape. Distributed roughness tends to produce moderate coefficients across many locations, whereas a localized event produces concentrated coefficients over a limited support. This is the property used by WMCR to separate background texture from transient content.

Different wavelet families emphasize different trade-offs. Haar wavelets are very compact, but their discontinuous step shape is too block-like for road events that are smoothed by pavement shape and tyre contact. Higher order Daubechies and Coiflet wavelets have more vanishing moments and reject smooth trends more strongly, but their longer filters can spread a short event over a wider interval. Continuous wavelets such as Morlet-type wavelets are useful for visual scale maps, but they are unnecessarily redundant for the compact component split used here. WMCR therefore uses a Symlet 4 wavelet. Symlet 4 is an orthogonal discrete wavelet basis function from the nearly symmetric Symlet family; the order four gives four vanishing moments and an eight coefficient compact filter. This makes it short enough to localize road events, while still smooth enough to avoid the block-like shape of the Haar basis. Its near symmetric filter shape also reduces phase distortion, which is important because the detected support must remain aligned with the position at which the vehicle encounters the road event.

3.3.1 SWT Decomposition

After the long wavelength topography has been extracted, the residual $x(u) = z(u) - g(u)$ still contains two different types of content. The background roughness is distributed along the road, whereas potholes, joints, and similar features appear as localized departures from that background. WMCR uses a stationary wavelet transform (SWT) with the Symlet 4 wavelet to identify these localized departures without shifting their longitudinal position.

The choice of an undecimated transform is deliberate. In a conventional decimated discrete wavelet transform, filtering is followed by downsampling. This is compact, but it makes the coefficients shift-sensitive: a small translation of the road profile

can move energy between neighbouring coefficients and between levels. The SWT removes the downsampling step. The filters are instead dilated from level to level, and every detail band remains on the original longitudinal grid. The residual is represented in the aligned multiresolution form

$$x(u) \approx a_J(u) + \sum_{j=1}^{J} d_j(u), \tag{3.9}$$

where a_J is the coarse approximation and d_j is the detail band at level j. Since all d_j are evaluated at the same road positions, a large detail coefficient can be interpreted as evidence for a local event at the same position where the vehicle will later encounter it.

The decomposition depth is chosen so that the coarsest analysed scale reaches the same wavelength range as the topography cutoff. For a longitudinal sample spacing Δu, level j represents wavelengths on the order of

$$\lambda_j \simeq 2^j \Delta u. \tag{3.10}$$

This relation is approximate because each wavelet detail band covers an octave rather than one exact wavelength. Its role is to connect the discrete wavelet levels to the physical road wavelengths used elsewhere in the work. At least six levels are used, and the residual is padded only to satisfy the SWT length requirement. After the decomposition, all signals are trimmed back to the measured road grid.

The wavelet stage is used only to guide the component split. It is not used to replace the measured residual by a smoothed wavelet reconstruction. This keeps the method conservative: the transform identifies candidate event supports, while the final roughness and transient signals are still formed in the original road profile domain.

3.3.2 Event Detection and Split

The detection rule treats each wavelet detail band as a scale specific residual signal. The working null hypothesis is that, after topography removal and away from rare local events, the background contribution within a band is centred, symmetric, and approximately Gaussian. The reason is local rather than global: a detail coefficient in one band combines many small background texture contributions after filtering, whereas a transient event appears as a sparse departure from that local background. The Gaussian approximation is used only to set a robust background scale for event detection in each band. Under the background-only model,

$$d_j(u) \approx \sigma_j \eta_j(u), \qquad \eta_j(u) \sim \mathcal{N}(0,1), \tag{3.11}$$

except at event locations. The centring is handled by the robust median estimate used below.

The scale σ_j must be estimated robustly because the same detail band is used to detect events. A standard deviation would be inflated by the large coefficients that

should be detected. Tukey's robust exploratory data analysis motivates using the median absolute deviation as a scale estimate when outliers are present [24]. Donoho and Johnstone's wavelet shrinkage work supports the related idea that sparse departures can be identified by thresholding wavelet coefficients against a background scale [27]. WMCR therefore uses

$$\begin{aligned} \mathrm{MAD}_j &= \mathrm{median}_u \, |d_j(u) - \mathrm{median}(d_j)|, \\ \sigma_j &= \frac{\mathrm{MAD}_j}{0.6745}, \qquad T_j = \alpha \sigma_j. \end{aligned} \tag{3.12}$$

The constant 0.6745 is not tuned. If η is standard normal, then $\mathrm{median}(|\eta|) = \Phi^{-1}(0.75) \approx 0.6745$, where Φ is the standard normal cumulative distribution. Dividing by this value makes σ_j consistent with the Gaussian standard deviation when the band contains only background texture. The multiplier $\alpha = 0.72$ is the calibrated detection sensitivity used in the present WMCR configuration. It is used as an event support detector, not as a final denoising threshold. This is important because a candidate support can still be rejected or adjusted by later context checks, whereas a missed severe event cannot be recovered after the split.

The thresholding is performed in two passes. The first pass marks unusually large coefficients and forms a preliminary transient estimate. This estimate is then removed before the background scale is recomputed, so that strong local events do not inflate the threshold used to detect themselves. This two-pass structure is a practical implementation of the background-only assumption in Equation 3.11: the background scale should be estimated from the part of the signal that is not already dominated by sparse localized departures.

The thresholded coefficients locate event supports rather than define the final event waveform directly. A local robust envelope, evaluated over a moving local window and bounded by a global noise floor, identifies intervals where the residual is dominated by localized content. The window is chosen to compare an event with its surrounding texture without treating slowly varying road severity as one stationary population over the whole route. Short inactive gaps are closed, supports are expanded to follow the event envelope, and compound supports can be split at valleys between dominant positive or negative peaks. The result is a set of event intervals

$$I_m = [u_m^-, u_m^+], \tag{3.13}$$

rather than isolated peak samples. This support-based definition is important because vehicle response depends on event width and local slope, not only on the largest sample value.

The component split is then made on the original residual rather than on the thresholded wavelet coefficients. Inside the detected supports, the transient component is the local departure from a baseline estimated from neighbouring clean context. Outside those supports, the transient component is set to zero. The roughness component is the complementary background signal, with event gaps bridged by the same local baseline before spectral fitting. At the profile level,

$$x(u) = r(u) + t(u), \tag{3.14}$$

where $t(u)$ contains compact localized departures and $r(u)$ is a background roughness signal with transient gaps filled. The event shape is therefore preserved in the road domain, while the roughness estimator is protected from the large local departures.

This construction gives the roughness and transient components different roles. The roughness signal is made suitable for local PSD fitting by removing the strongest localized departures, whereas the transient signal preserves those departures explicitly for later placement. Without this separation, severe events would either distort the fitted background roughness level or be lost when the roughness is regenerated stochastically.

3.4 Roughness

3.4.1 PSD Parameterization

The roughness component $r(u)$ represents the background texture after topography removal and transient support removal. It is parameterized locally, rather than stored as a measured waveform. The road is divided into fixed length segments of 20 m. For each segment q, WMCR estimates a one sided spatial power spectral density using Welch's method with constant detrending. The spectrum is then fit in logarithmic space with the power law model

$$S_q(n) = C_q \left(\frac{n}{n_0} \right)^{-w_q}, \tag{3.15}$$

where n is spatial frequency, $n_0 = 0.1$ cycle/m is the reference frequency, C_q is the spectral level, and w_q is the waviness exponent. The fit starts at the long wavelength floor, currently 1/50 cycle/m, and extends to the resolvable high frequency range of the sampled signal.

The 20 m segment length is chosen as a compromise between local adaptation and statistical stability. Shorter windows would follow rapid changes more closely, but the PSD estimates would become noisy and sensitive to individual samples. Longer windows would produce smoother parameter histories, but they would hide changes in roughness severity that are relevant for vehicle loading. The local pair (C_q, w_q) therefore acts as a compact descriptor of both amplitude level and spectral shape over a distance scale that remains meaningful for simulation.

The local roughness descriptor is therefore

$$\Theta_q^{\mathrm{r}} = (C_q, w_q). \tag{3.16}$$

Unstable edge estimates are filled from the interior parameter statistics, and w_q is constrained to a physically reasonable range. Additional scale measures, such as the local RMS, can be retained for calibration, but the core roughness parameterization remains the pair (C, w).

3.4.2 Parameter Sampling and Reconstruction

The roughness reconstruction is driven by the distribution of the fitted parameters rather than by copied profile samples. Figure 3.3 shows the parameter comparison for the bin 50 case. The two panels compare the marginal distributions of $\log_{10} C$ and waviness w between the measured parameter pool and the generated segments. These are the two core parameters used to control the local roughness spectrum: C sets the roughness level, while w sets the spectral slope. The generated curves follow the main peaks of the measured distributions, which indicates that the compact road keeps both the dominant roughness level and the dominant spectral slope. Small differences remain in the tails, especially for high roughness levels, where the measured pool contains fewer but more severe segments. The comparison is therefore a marginal distribution check rather than a full joint-support diagnostic.

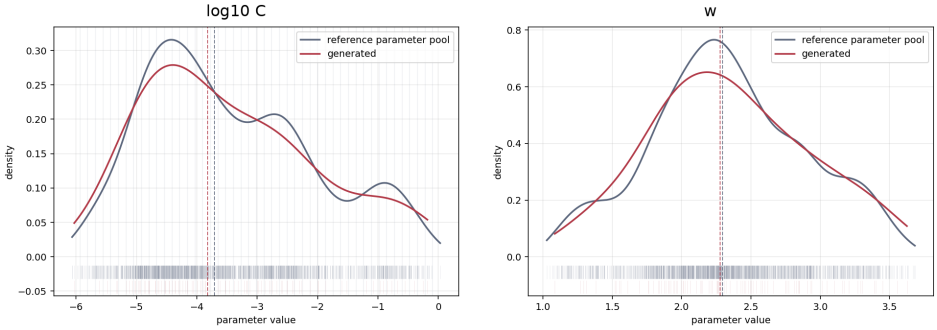

Figure 3.3: Bin 50 distributions of $\log_{10} C$ and waviness w.

During reconstruction, the selected segment parameters are converted into continuous fields of roughness level $C(u)$, waviness $w(u)$, and local amplitude scale. Rather than reusing measured roughness samples directly, these fields define a local target spectrum,

$$S(n; u) = C(u) \left(\frac{n}{n_0} \right)^{-w(u)} . \tag{3.17}$$

A new stochastic signal $\hat{r}(u)$ is synthesized over logarithmic frequency bands. The power law fit controls the spectral shape, while the local amplitude scale controls the absolute level. A soft calibration then aligns RMS, slope, and damage related targets, but this calibration only adjusts the generated realization; the main mechanism remains parameter-driven synthesis.

Figure 3.4 visualizes this reconstruction step. The sampled roughness level C and waviness w define the local target PSD rather than a fixed waveform. A random phase realization is then drawn from that target spectrum and transformed back into the spatial domain, giving a roughness signal whose local amplitude and spectral slope remain traceable to the compact parameter pair.

This design is deliberately different from copying roughness samples from the measured road. Copying would preserve local details but would also limit the compact road to a rearrangement of measured snippets. Parameter-driven synthesis allows a new realization to be generated while retaining the measured roughness statistics.

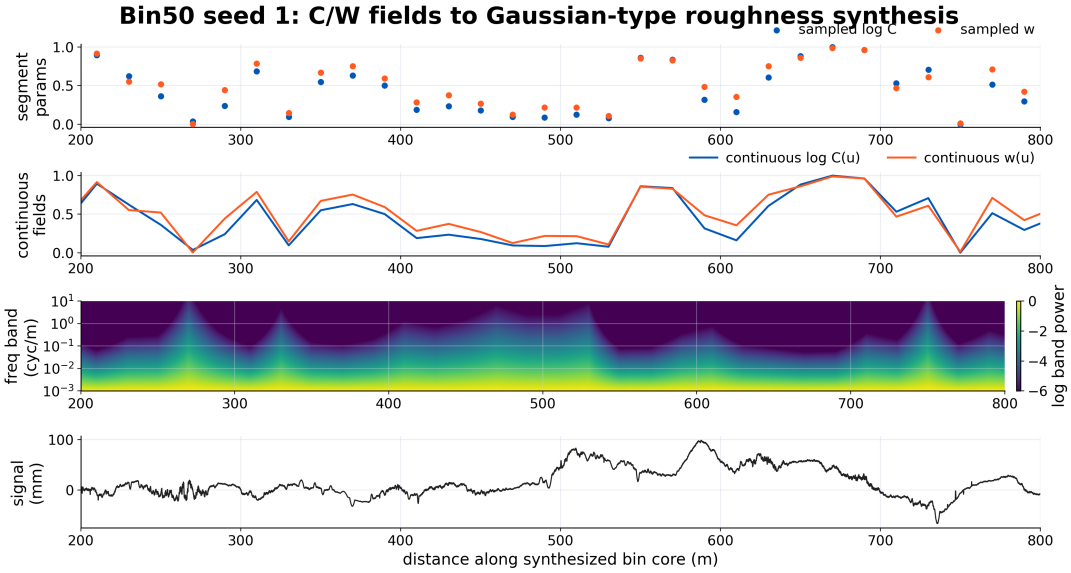

Figure 3.4: WMCR roughness reconstruction from sampled C and w fields.

The calibration step is kept secondary so that the reconstructed roughness remains traceable to the fitted PSD parameters rather than becoming a black-box response fit.

3.5 Transient Events

3.5.1 Event Library

The transient component $t(u)$ is represented as a library of compact localized events. For each detected support interval I_m, WMCR extracts the event waveform from the transient signal and records its peak, width, polarity, damage contribution, and local background condition. The compact descriptor used during synthesis is

$$\Theta_m^{\mathrm{t}} = (W_m, A_m, \psi_m, C_{\mathrm{bg},m}), \tag{3.18}$$

where W_m is the event support width, A_m is the signed peak amplitude, ψ_m is the normalized longitudinal event shape, and $C_{\mathrm{bg},m}$ is the roughness level of the segment in which the source event was found. When a measured waveform is available, ψ_m is a peak preserving event template. A half sine family is used as a fallback template representation.

The background term is not only auxiliary metadata. It records that a transient is defined relative to its local road context. The same waveform may stand out as a localized event on a smooth road section, but be absorbed into ordinary texture on a rough section. The value $C_{\mathrm{bg},m}$ therefore prevents the event library from treating amplitude and shape as absolute properties independent of the background on which the event was observed.

The event library is intentionally more explicit than the roughness parameterization. Localized events are rare, and their influence is often carried by their shape, polarity, and surrounding context rather than by their average spectral contribution. Storing them as templates avoids diluting them into the background PSD. At the same time, using normalized templates and separate amplitude and width parameters makes it possible to place events in a compact road without requiring the original event samples to be copied verbatim.

Template reconstruction is therefore sufficient only in a response-oriented sense, not as an exact reproduction of every measured defect. It is adequate when the template preserves the event properties that dominate the vehicle input: peak amplitude, support width, polarity, edge shape, and the background roughness context in which the event was observed. A measured event template is preferred because it retains asymmetry and multi-peak details that a simple analytic shape may miss. The half-sine fallback should be read as a regularized approximation for cases where only compact event descriptors are available. It can preserve the main amplitude and width effect, but it is less reliable for channels that are sensitive to local slopes and edge sharpness, especially the longitudinal $f6$ load path. The sufficiency of the template representation is therefore checked by the subsequent response validation rather than assumed from the road profile alone.

3.5.2 Event Placement and Assembly

Transient reconstruction is conditioned on the generated roughness context. Candidate events are drawn from the event library, but they are not placed arbitrarily. The sampler prefers events whose source background roughness is compatible with the target segment, preserves measured event density through a conservative density floor, and avoids unrealistic clustering through spacing constraints. During placement, each selected event receives a new target position $\hat{u}_m$, its width is scaled to the target segment, and its waveform is blended with the surrounding roughness so that it behaves as a localized road feature rather than an isolated numerical spike,

$$\hat{t}_m(u) = A_m \, \psi_m\left(\frac{u - \hat{u}_m}{W_m}\right). \tag{3.19}$$

The complete transient signal is the sum of all inserted events,

$$\hat{t}(u) = \sum_{m=1}^{M_t} \hat{t}_m(u), \tag{3.20}$$

where M_t is the number of events in the compact road.

The compression step changes the available road length, but it should not change what an event represents. WMCR therefore does not simply copy every event from the measured road into the shorter target profile. Doing so would increase the event density and create an unrealistically severe compact road. Conversely, discarding events only because the target road is shorter would remove rare features that may

dominate durability loads. The event sampler therefore treats compression as a density and severity preservation problem: the compact road keeps a representative number of localized features, keeps severe events as explicit inserts rather than averaging them into the roughness, and leaves enough spacing for each event to remain distinguishable from its neighbours.

The context condition prevents a common failure mode in compact reconstruction: placing severe events on unrealistically smooth background sections, or grouping many events so closely that the generated road contains an artificial obstacle cluster. WMCR therefore treats event placement as a constrained sampling problem. The event density controls how many severe features remain in the compact road, while the context and spacing controls determine where they can be inserted without breaking the statistical character of the source road bin.

This also makes transient insertion one of the most sensitive parts of WMCR. The roughness field is generated first, the event candidates are then matched to that realized local background, and only after that are the event waveforms blended into the profile. Changing this order, the local blend, or the spacing between events can change local slopes and pulse shapes, and therefore change the vehicle response even when the stored event template is unchanged.

The final one dimensional profile is assembled as the sum of reconstructed topography, generated roughness, and placed transients,

$$\hat{z}(u) = \hat{g}(u) + \hat{r}(u) + \hat{t}(u). \tag{3.21}$$

This signal is the conditioning profile for the 2D surface stage.

The one-dimensional assembly is a useful intermediate result because it can be checked before lateral expansion. It contains the three longitudinal components that dominate the road input: smooth topography, stochastic roughness, and explicit events. The remaining task is to convert this profile into a road surface that can excite left and right wheel paths with realistic lateral correlation.

3.6 2D Road Surface

The one dimensional reconstruction in Equation 3.21 is used as a conditioning signal for the 2D stage. Here, "one dimensional" means a univariate elevation signal over longitudinal distance, $\hat{z}(u)$. Such a signal is often plotted as a two-axis road profile, with distance on one axis and elevation on the other, but its spatial domain contains only one coordinate. In contrast, the exported WMCR road is a height field $\hat{Z}(u, v)$ over longitudinal and lateral position. It may be visualized as a three-dimensional surface plot because elevation is drawn vertically, but the modelled road domain is two dimensional: (u, v).

This terminology matters for the vehicle simulation. The vehicle model does not travel on a single centreline signal. The left and right tyres sample different lateral

positions, and suspension loads depend on both common vertical excitation and differential input between the wheel paths. The 2D stage therefore preserves the component logic of the 1D reconstruction while adding a lateral coherence model and explicit transient footprints.

3.6.1 Lateral Coherence

The continuous base surface is the sum of the topography and roughness layers,

$$\hat{Z}_{\text{base}}(u, v) = \hat{Z}_{\text{topography}}(u, v) + \hat{Z}_{\text{roughness}}(u, v). \tag{3.22}$$

It is conditioned by the non-transient part of the one dimensional reconstruction, so that the representative lateral position satisfies $\hat{Z}_{\text{base}}(u, v_0) \approx \hat{g}(u) + \hat{r}(u)$. The transient term $\hat{t}(u)$ is deliberately excluded from this background expansion.

Measured roads with multiple lateral tracks show correlation across lateral offsets. Nearby tracks share much of the long wavelength content, while shorter wavelengths decorrelate more quickly as the lateral distance increases. The functional form follows Bogsjö's treatment of road-profile statistics for laterally separated tracks [1]. Its use here is also consistent with the later two-track roughness modelling in Bogsjö, Podgórski, and Rychlik [17]. The numerical decay parameter, however, is not imported as a literature constant. WMCR estimates the Bogsjö-type coherence constant from the measured lateral track pairs in the road data used for each speed bin,

$$\gamma^2(n, d) = \exp[-2\beta_\gamma(d)n], \qquad d = |v_a - v_b|. \tag{3.23}$$

Here $\gamma^2(n, d)$ is the target magnitude-squared coherence at spatial frequency n and lateral separation d, while $\beta_\gamma(d)$ is the effective coherence decay parameter for that separation. The distance dependence is therefore carried by the fitted parameter $\beta_\gamma(d)$: larger values represent faster lateral decorrelation.

For each available pair of detrended tracks, the measured magnitude-squared coherence is estimated from Welch spectra,

$$\hat{\gamma}^2_m(n, d) = \frac{|\hat{S}_{ab}(n)|^2}{\hat{S}_{aa}(n)\hat{S}_{bb}(n)}. \tag{3.24}$$

Numerically unreliable points, including the zero-frequency value and coherence estimates close to the limits zero and one, are excluded before fitting. Taking the logarithm of Equation 3.23 gives the through-origin linear relation

$$\log \hat{\gamma}^2_m(n_i, d) \approx -2\beta_\gamma(d)n_i, \qquad \hat{\beta}_\gamma(d) = -\frac{\sum_i n_i \log \hat{\gamma}^2_m(n_i, d)}{2\sum_i n_i^2}. \tag{3.25}$$

Estimates from all valid track pairs are aggregated robustly by lateral separation to form a speed-bin-specific, distance-dependent decay table. During surface synthesis, the decay parameter for an anchor separation is obtained by interpolating within this measured table. Thus the reported WMCR surfaces use a coherence constant derived

from the measured road data in the corresponding bin, not a generic default value. The magnitude-squared target is then converted to an amplitude coherence before conditional anchor sampling. This calibration keeps long wavelength content shared across the road width while allowing shorter wavelength texture and differential excitation to vary between wheel paths.

3.6.2 Surface Generation

The 2D road is generated by expanding the non-transient 1D reconstruction to a small set of lateral anchors, interpolating these anchors to the dense surface grid, and finally adding transient increments. The final surface is written as

$$\hat{Z}(u, v) = \hat{Z}_{\text{topography}}(u, v) + \hat{Z}_{\text{roughness}}(u, v) + \hat{T}(u, v), \tag{3.26}$$

where $\hat{Z}_{\text{topography}}$ is the long wavelength surface topography, $\hat{Z}_{\text{roughness}}$ is the laterally coherent background roughness, and $\hat{T}$ is the 2D transient increment field.

For a standard wheel track layout, the anchor set is

$$\mathcal{A} = \{-1.4,\ -0.8,\ 0,\ 0.8,\ 1.4\}\,. \tag{3.27}$$

The coordinates are expressed in metres, giving a road width of 2.8 m. The centre anchor is conditioned on $\hat{g}(u) + \hat{r}(u)$, while the remaining anchors are sampled so that their cross spectra follow the coherence target in Equation 3.23. Topography anchors provide the low frequency lateral shape, and roughness anchors provide correlated stochastic texture.

The dense lateral grid is then obtained by piecewise linear interpolation between neighbouring anchors at each longitudinal sample. If $v_a \leq v \leq v_b$ are the two anchors surrounding a grid point v, the background surface value is

$$\hat{Z}_{\text{base}}(u, v) = (1 - \alpha)\hat{Z}_{\text{base}}(u, v_a) + \alpha\hat{Z}_{\text{base}}(u, v_b), \qquad \alpha = \frac{v - v_a}{v_b - v_a}. \tag{3.28}$$

This interpolation is only the geometric expansion from sparse anchors to the CRG surface grid. The lateral coherence is already imposed when the anchor signals are sampled; it is not created afterwards by smoothing the dense grid. The background grid is then combined with the transient field and exported as the 2D road surface. The coherence target is used during conditional anchor sampling, not as a post-processing filter; the magnitude-squared target is converted to an amplitude correlation before sampling.

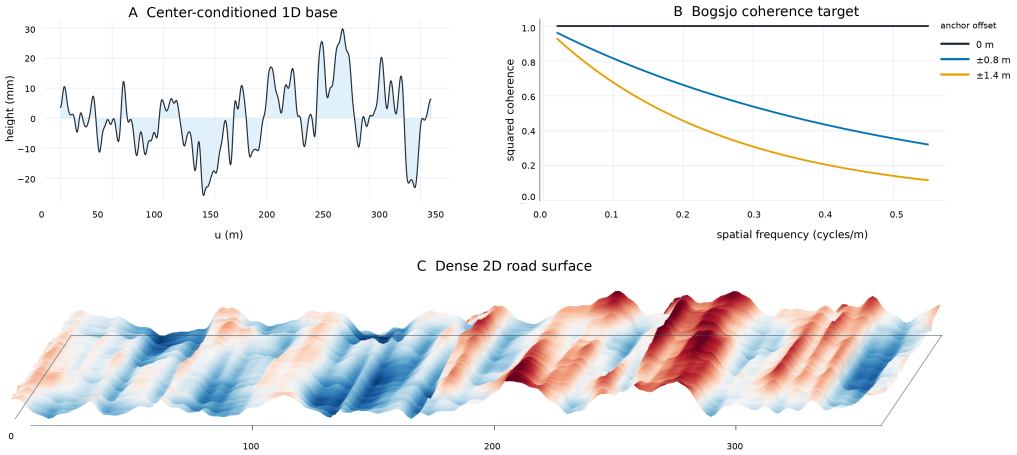

Figure 3.5: Staged generation of the 2D WMCR surface.

The decomposition in Equation 3.26 is maintained during surface generation. The topography layer is smooth in u and v and stores the long wavelength road shape. The roughness layer is generated around the topography layer as a zero mean background texture whose spectral level and waviness are inherited from the WMCR roughness parameters. Their sum forms the base surface in Equation 3.22. This surface is laterally coherent and continuous, but it contains no explicit transient events.

Keeping the layers separate avoids hiding localized events inside the background coherence model. It also makes the topography, roughness, and transient contributions easy to inspect before the final surface is exported.

Figure 3.5 illustrates this staged construction. The point is not only to obtain a dense grid, but to keep the origin of the grid values interpretable. A later response mismatch can then be traced back to topography, background roughness, transient placement, or lateral coherence rather than to an undifferentiated road surface.

3.6.3 Transient Injection and Export

Before injection, the compact-road compression has already determined how many transient events should be retained. The selected transient density is scaled to the target length while preserving the measured event rate and retaining severe localized features as explicit inserts. This avoids two opposite errors: filling the shorter road with too many copied events, or smoothing rare high-severity events into the background roughness because the road has been compressed.

Transient events are injected only after the base surface has been generated. They do not participate in the lateral coherence expansion of the background surface. The longitudinal part is the placed event in Equation 3.19; the lateral part is an envelope

$\beta_m(v)$. The 2D event is therefore

$$\hat{T}_m(u, v) = A_m \, \psi_m\left(\frac{u - \hat{u}_m}{W_m}\right) \beta_m(v). \tag{3.29}$$

When measurements from multiple lateral tracks provide a lateral footprint, $\beta_m(v)$ is estimated from those tracks. Otherwise, it is sampled from a bounded envelope family. The full transient surface is obtained by summing the inserted 2D events,

$$\hat{T}(u, v) = \sum_{m=1}^{M_t} \hat{T}_m(u, v). \tag{3.30}$$

The exported object then contains the dense surface grid, axis definitions, grid spacing, and minimal generation metadata. Before export, the grid is checked for finite values, regular axes, reasonable spacing, and plausible lateral coherence. The same surface can be saved as a MATLAB `.mat` payload or converted to OpenCRG for vehicle simulation.

The validation checks at export are practical but important. They ensure that the road surface is numerically well formed before it is passed to Adams Car, and they separate file format errors from modelling errors. After export, the WMCR road can be evaluated in the same way as the NSL road and the measured reference: first through road diagnostics and then through vehicle response metrics.

Figure 3.6 shows one generated WMCR surface example before export. It provides a visual check that the reconstructed road field is well formed before the road is used in the vehicle simulation.

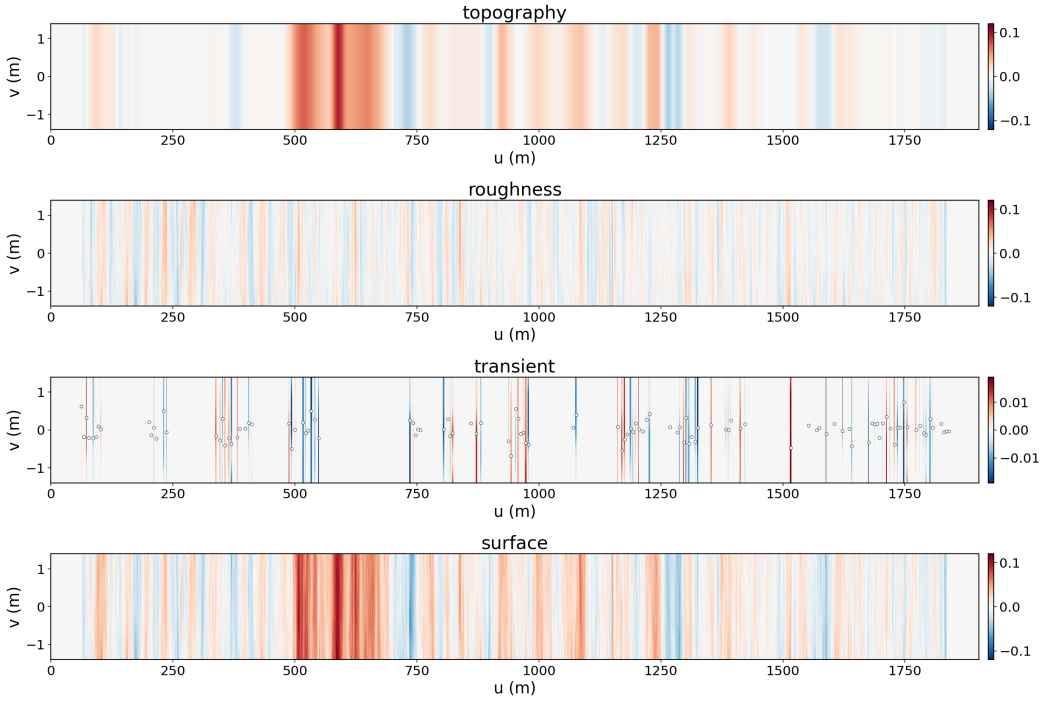

Figure 3.6: Generated 2D WMCR surface before export.

In summary, WMCR stores topography, roughness, transient events, and lateral surface information as separate component parameters. These parameters are sampled within a measured road bin, reconstructed as a one dimensional road, and expanded into a laterally coherent 2D surface for simulation. The method is therefore a route for component preserving compression. It reduces the amount of road data that must be simulated, but it retains the distinctions needed to interpret why a generated road does or does not reproduce the measured vehicle response.

4

Simulation and Response Results

The validation moves from road profile comparison to vehicle response simulation. A compact road is suitable for durability work only if the simplified representation preserves the excitation mechanisms that drive component forces, not only the statistical appearance of the measured profile. The resulting test is whether the generated roads remain equivalent after they are used as inputs to the vehicle model.

The validation chain is intentionally progressive and is applied to the 30, 50, 70, and 100 km/h speed bins. The preliminary analysis first checks spectral content and quarter car pseudo damage, which gives a reduced vertical response screen. The non-linear Adams Car benchmark then evaluates the same roads in a full multibody vehicle model, where tyre contact, suspension kinematics, load transfer, and left–right coupling can amplify differences that are partly hidden in profile descriptors.

This focus follows the durability road motivation used throughout the work: measured local irregularities can be more important for fatigue than their share of average roughness suggests. Bogsjö's road-profile statistics emphasize the fatigue relevance of local unevenness [1], and Charles' test-severity work supports deriving durability input from measured road events rather than from average roughness alone [2].

4.1 Preliminary Analysis

The preliminary analysis provides an intermediate check between road synthesis and vehicle response simulation. It evaluates whether the generated roads preserve the measured road spectrum and whether the remaining differences affect a reduced vehicle response measure before the Adams Car benchmark is interpreted.

The comparison is organized in two steps and is applied consistently to the 30, 50, 70, and 100 km/h speed bins. The PSD check first compares how the measured road, the external NSL roads, and the WMCR road distribute elevation energy over spatial frequency. Quarter car pseudo damage and suspension force RMS then convert the same road inputs into compact response severity measures. The two checks are interpreted together: the PSD shows where the road descriptions differ,

while the quarter car result indicates whether those differences are likely to influence durability simulation. The figures shown in this section are representative Bin50 validation views from the screening workflow rather than a complete set of all speed bin diagnostics.

4.1.1 PSD

Power spectral density describes how profile variance is distributed over spatial frequency. Dodds and Robson give the classical spectral description of road roughness [8]. ISO 8608 then standardizes how road-profile PSD content is reported [3]. Andrén discusses practical PSD approximations for longitudinal road profiles [9]. Figure 4.1 shows the Bin50 mean surface PSD comparison. It compares the measured reference with WMCR, the C route NSL surface, and the IRI 1m and IRI 20m NSL surfaces. The PSD is evaluated across the surface and reported as a mean curve, so the figure should be read as a Bin50 surface comparison rather than a single track example. The shaded band marks ±3 dB around the measured reference. In this discussion, low spatial frequency or long wavelength means approximately $n \lesssim 0.1$ cycle/m, corresponding to wavelengths $\lambda \gtrsim 10$ m.

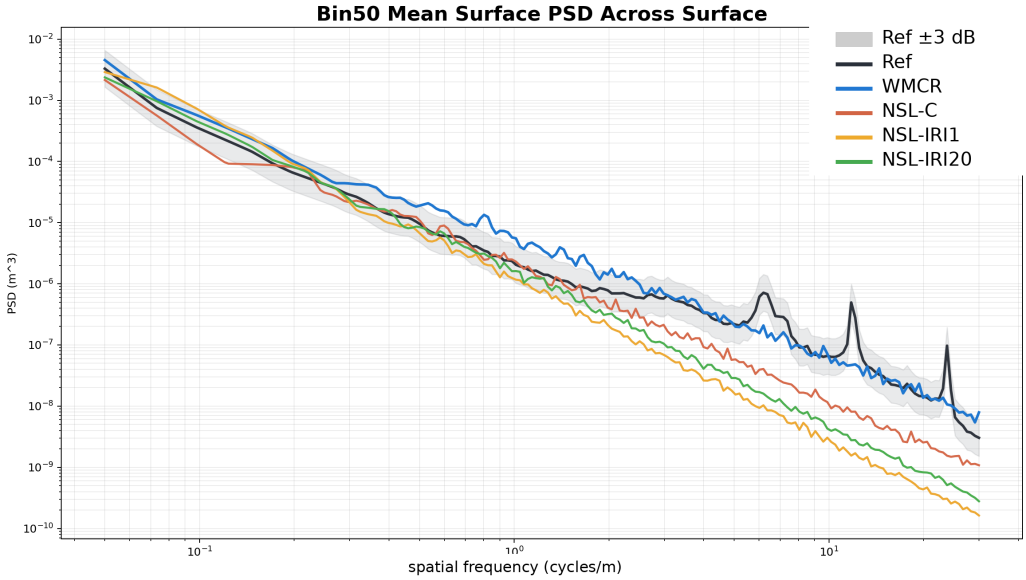

Figure 4.1: Bin 50 mean surface PSD comparison.

The displayed curves follow the same overall decay as spatial frequency increases, which indicates that the generated roads reproduce the main roughness trend of the measured surface to different degrees. This agreement is important, but it should be read as a spectral check rather than as evidence that nonstationary transient events are preserved. Bruscella and co-workers show that road profiles can require separation beyond a single global spectrum [6]. Johannesson, Podgórski, and Rychlik motivate nonstationary variance descriptions for road roughness [7]. Wavelet-based pavement analysis also demonstrates the need for spatially sensitive profile diagnostics [19].

The main differences are in spectral level and in the high frequency roll off. In the low and middle spatial frequency range, WMCR stays close to the reference envelope and generally remains inside or near the ±3 dB band. At higher spatial frequencies, especially above about 0.1 cycle/m, WMCR still falls below the reference and does not reproduce every narrow peak. The three NSL external roads lose high frequency energy more clearly. The C route NSL surface retains more of this content than the IRI based alternatives, while the IRI 1m and IRI 20m routes show the strongest downward separation from the reference.

The PSD result is therefore useful as a first check, but it is not sufficient on its own. A generated road can match the global spectral trend while still missing where transient energy occurs along the road, how sharp a local event is, and how that event is placed across the vehicle track width. These details affect wheel travel, suspension velocity, and fatigue relevant cycle amplitudes. The PSD comparison is therefore followed by a reduced response measure.

4.1.2 Quarter Car Pseudo Damage

Quarter car pseudo damage and suspension force RMS provide the response continuation of the PSD check. They recall the same reduced vehicle idea used by IRI. The International Road Roughness Experiment established the standardized quarter-car basis for IRI interpretation [10]. Sayers' work explains how IRI is computed from longitudinal road profiles [23]. Each road profile is evaluated with the same quarter car model, while the speed is selected from the corresponding road bin. The 30, 50, 70, and 100 km/h road bins are therefore simulated at 30, 50, 70, and 100 km/h, respectively. The spatial road profile is converted to a temporal input by

$$z_t(t) = z_x(vt), \qquad t = \frac{x}{v}, \tag{4.1}$$

where z_x is the spatial road elevation, z_t is the temporal road input, x is longitudinal distance, and v is the bin speed.

The quarter car response history is reduced to relative damage and RMS measures. After cycle counting [28], each response cycle is represented by a cycle count n_i and an amplitude or range S_i. Pseudo damage is accumulated as

$$D_{\mathrm{QC}} = \sum_{i=1}^{N_c} n_i \, S_i^m, \tag{4.2}$$

where N_c is the number of counted cycle levels and m is a fixed fatigue exponent. Material constants are omitted because only relative comparison is required. The same suspension force history is also reduced to RMS, so the comparison can separate fatigue sensitive cycle content from average force level.

Figure 4.2 applies this check to the Bin50 road set. Both panels are normalized by the measured reference, so a value of one indicates agreement with the reference road.

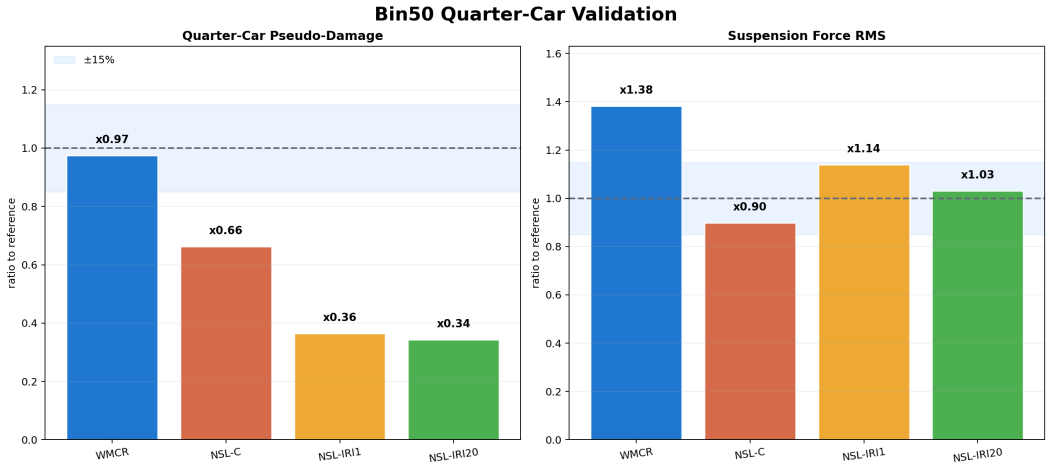

Figure 4.2: Bin 50 quarter-car validation.

WMCR gives a pseudo damage ratio close to the reference, about 0.97, while its suspension force RMS is higher, about 1.38. The three NSL external roads recover less pseudo damage than the reference. The C route NSL surface is the strongest of them at about 0.66, while the IRI 1m and IRI 20m routes are lower at about 0.36 and 0.34. The RMS panel gives a different view: the C route NSL surface is slightly below the reference, the IRI 1m route is moderately above it, and the IRI 20m route is close to one. This reduced model is valuable for ranking vertical severity, but it remains a simplified response path; nonlinear suspension and tyre effects can change the accuracy of quarter car predictions [13].

Across the speed bins, the quarter car trend should be interpreted through the speed mapping rather than through roughness alone. Increasing speed shifts a fixed spatial wavelength to a higher temporal excitation frequency, so pseudo damage and RMS can change when road content moves closer to or farther from the quarter car response bands. The Bin50 result shows why both metrics are needed. A road can have a reasonable RMS level while still losing fatigue relevant cycles, and a road can retain the pseudo damage level while remaining high in RMS. The 50 km/h plot is therefore a representative detailed view of the mechanism, not the only basis for the preliminary conclusion.

The measured component split gives the reduced-response explanation for this trend. The transient statistics are computed at profile level rather than by pooling all road samples. In the Chapter 4 screening set, one centreline WMCR profile is analysed for each speed bin, giving $N = 4$ profile-level observations for the 30, 50, 70, and 100 km/h bins. For each profile, the transient RMS share is computed as the RMS of the transient component divided by the RMS of the full profile. The response contribution is computed from the same profile before and after transient injection using quarter car pseudo damage. Figure 4.3 compares these two shares directly.

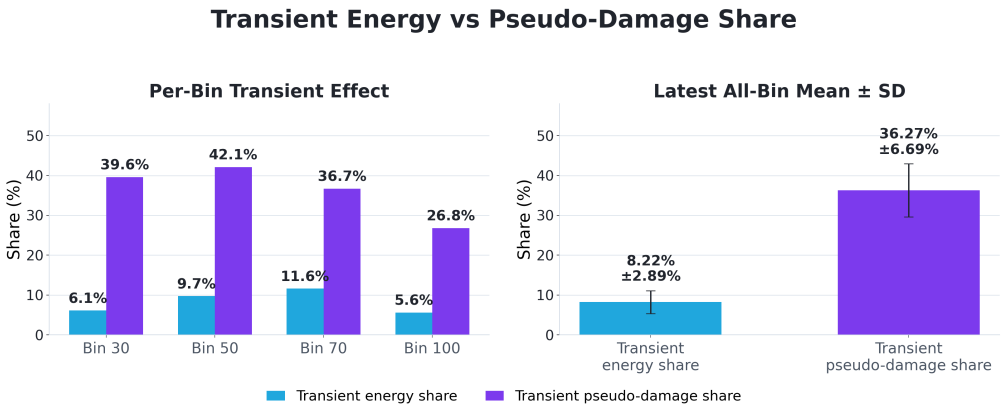

Figure 4.3: Transient energy and pseudo-damage shares.

The transient RMS share is $8.22\% \pm 2.89\%$, while the before–after quarter car comparison attributes $36.27\% \pm 6.69\%$ of the pseudo-damage increment to the same component. A small energy share can therefore carry a much larger fatigue share, which supports using the full Adams benchmark after the reduced vertical screen.

4.2 Benchmark and Simulation Setup

The quarter car model provides a reduced screen because it connects road input to vertical response through a compact reference system. It does not, however, represent suspension geometry, nonlinear component behaviour, multi axle coupling, or load transfer between the two sides of the vehicle. For that reason, the generated roads are also checked in a full vehicle benchmark before they are judged suitable for durability simulation.

The benchmark is run as a non-linear Adams Car full-vehicle simulation with the same multibody vehicle model for all road inputs. The manoeuvre is a straight, constant-speed road run over the imported CRG surface. The speed is assigned by road bin, so the 30, 50, 70, and 100 km/h bins are simulated at 30, 50, 70, and 100 km/h, respectively. The reference response is obtained from the original scanned road, and the NSL and WMCR roads are prepared as compact CRG road files and applied through the same road input format; CRG is a dedicated road surface format for vehicle simulation workflows [29].

The tyre and road-contact model are kept fixed between the measured, NSL, and WMCR simulations. No method-specific tyre parameters, contact settings, solver settings, or output sampling settings are retuned for one road representation. Where solver tolerances, integration options, and contact parameters enter the Adams model, they are treated as controlled benchmark settings rather than as calibration variables. The road surface is therefore the intended variable in the comparison, so differences in the monitored responses can be interpreted as consequences of the

road representation rather than changes in the simulation setup.

The post-processing workflow is also fixed. For each simulation, the left and right response histories for the monitored channels are extracted over the common travel interval, processed with the same filtering and cycle-counting procedure, and reduced to Damage, response range, maximum, and minimum diagnostics. The figures use Damage and range as the primary pair. Each bar is normalized by the corresponding measured-road value for the same speed bin, side, and response channel, so a value of one means agreement with the measured reference.

4.3 Response Results

The Adams Car response comparison is where road similarity is converted into response equivalence. The preliminary PSD and quarter car checks show whether the compact roads contain plausible spectral content and reduced vertical severity. The vehicle simulation then tests whether that content produces comparable force responses to the measured road. The result is therefore not judged from one scalar roughness descriptor. A compact road must recover fatigue relevant damage while also keeping the response amplitude realistic in several monitored force channels.

The response results are therefore organised by channel rather than by modelling method. Figures 4.5–4.7 show $f55$, $f6$, and $f63$ separately. Each figure contains the four speed bins and all four compact-road methods: NSL IRI-1m, NSL IRI-20m, Direct C, and WMCR. This layout makes the comparison speed-bin aware while keeping the methods in the same plot. The analysis can still distinguish the behaviour of the NSL routes and WMCR, but the evidence is read through the physical load path first.

4.3.1 Monitored Response Channels

The notation $f55$, $f6$, and $f63$ refers to three output channels from the Adams Car post processing. Here, a channel means a force response at a specific suspension or body connection. The channels are selected because they react to different physical aspects of the same road input. Damage, labelled as Fatigue Damage or DTOT $k = 4$ in the figures, is the primary durability quantity because large response cycles receive a high weight in cycle based fatigue evaluation. DTOT is the total rainflow-based pseudo damage accumulated over the processed response history. ASTM E1049 defines the cycle-counting practice used to reduce a response history to fatigue relevant cycles [28]. Palmgren's bearing fatigue work provides the early cumulative-damage basis [21]. Miner's later rule gives the linear damage accumulation principle used in many fatigue calculations [22]. Figure 4.4 identifies the three monitored channels before the channel-specific response plots are interpreted.

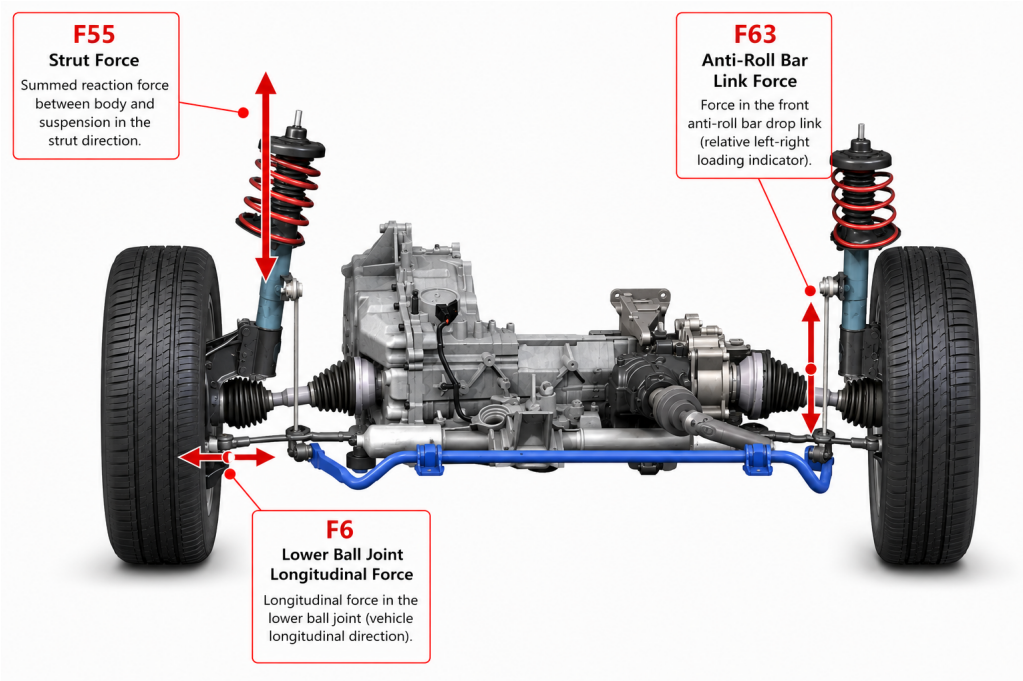

Figure 4.4: Monitored Adams force channels.

The exponent $k = 4$ is used as a representative fatigue-damage slope rather than as a material-specific calibration. In a Basquin-type relation $N = CS^{-k}$, the damage contribution of a counted response cycle is proportional to S^k after constants common to all compared roads are removed. Choosing $k = 4$ therefore emphasizes large response cycles while keeping the comparison as a relative road-severity metric. Range is used beside damage as an amplitude check, so a method does not appear successful only because it creates a few unrealistic peaks. In all response figures, a ratio of one indicates agreement with the measured-road reference for the same speed bin and channel.

Strut force sum $f55$. This channel represents vertical loading through the front strut assembly. In the post processing, $f55$ is the force sum along the strut direction, including the spring, damper, jounce bumper, and rebound stop force contributions. Physically, it is closely related to wheel vertical motion, damper velocity, spring compression, and the transfer of suspension force into the vehicle body. It therefore continues the quarter car screen at component level: if a compact road loses the roughness or localized events that drive suspension motion, the $f55$ damage usually decreases. If the road is made globally too severe, the $f55$ damage and range increase together.

Lower ball joint longitudinal force $f6$. This channel represents the longitudinal force at the lower ball joint. It is not controlled only by the vertical roughness level. Local slopes, event edges, tyre contact filtering, and suspension kinematics convert road elevation into longitudinal force. A short edge, bump, or depression can therefore produce a large $f6$ pulse even if it is not dominant in the road PSD.

This makes $f6$ a diagnostic check of whether the reconstructed road has realistic local event shapes.

Front drop link force $f63$. This channel represents the front anti roll bar link force. It is driven by the difference between the left and right front wheel excitations, so it tests the lateral meaning of the road surface. A road can match the roughness of each individual wheel track and still fail in $f63$ if the two tracks are too coherent, too independent, or if localized events have an unrealistic lateral footprint. The $f63$ channel therefore checks whether the compact road remains plausible as a two dimensional surface rather than as two separate longitudinal profiles.

4.3.2 Channel-Based Adams Results

Figure 4.5 compares the vertical strut-loading channel. At 30 km/h the descriptor-based routes under-recover $f55$ damage, while WMCR is conservative in damage but remains below the measured range. At 50 km/h, WMCR is close to the measured damage level, whereas Direct C strongly over-predicts $f55$ fatigue damage. At 70 km/h the NSL IRI variants become highly severe in $f55$, and Direct C and WMCR also exceed the reference damage level, although WMCR keeps the range lower. At 100 km/h, WMCR returns close to unity in damage and remains controlled in range, while the IRI variants remain too severe. The $f55$ evidence therefore supports the main WMCR mechanism for vertical loading, but it also shows that compression can still redistribute local peak content between damage and range.

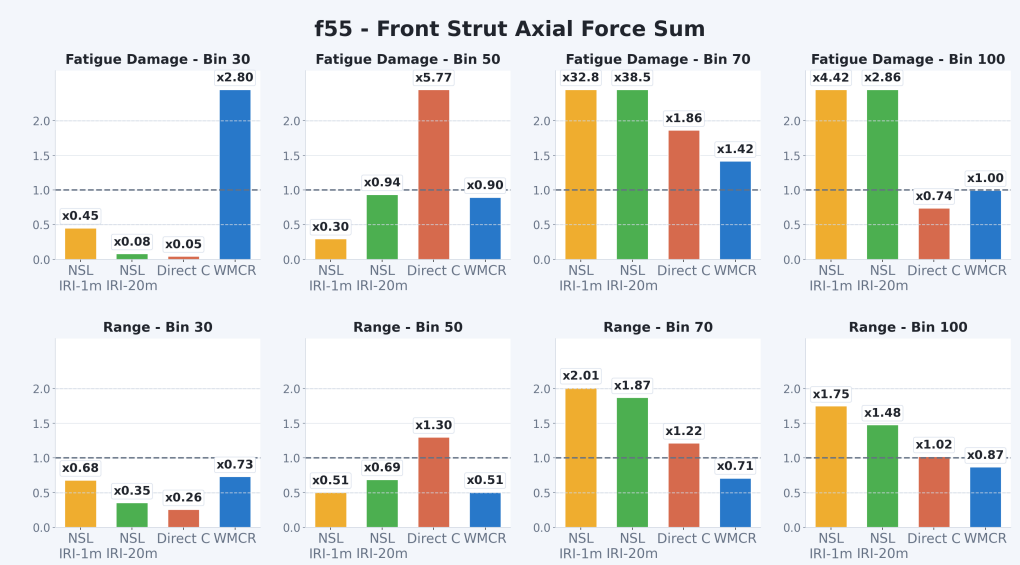

Figure 4.5: Adams $f55$ response over the four speed bins.

The large error of the IRI-based NSL routes in the 70 km/h bin can be explained by the conversion from IRI to the roughness coefficient C. For both the IRI-1m and

IRI-20m routes, the segment roughness coefficient is estimated as

$$\hat{C}_j = 10^{-6} \left(\frac{\hat{I}_j}{2.21} \right)^2 , \tag{4.3}$$

which relies on the empirical IRI–C relation under the assumption of a fixed waviness exponent $w = 2.0$. This assumption is reasonable for a broad road class, but it becomes problematic when the roads inside one speed bin are not homogeneous.

This effect is especially visible in the 70 km/h bin. Although roads in this speed class are normally expected to be relatively smooth, one road in this bin contains substantially rougher sections and also accounts for more than 70% of the total concatenated road length in the bin. As a result, this road dominates the fitted IRI-based NSL statistics. Because the IRI routes still impose a fixed $w = 2.0$, the conversion from IRI to C cannot represent the actual spectral shape of this dominant road. The generated road can therefore receive an incorrect roughness-energy distribution, which is then amplified in the full-vehicle response and leads to excessive fatigue damage.

The Direct C route does not show the same level of amplification because it estimates the PSD-based parameters directly from the measured profiles. In particular, the waviness exponent is evaluated segment by segment rather than being fixed to $w = 2.0$ for the whole road class. This allows the Direct C route to adapt to the unusual spectral characteristics of the dominant 70 km/h road and explains why its Adams response does not exhibit the same damage explosion as the IRI-based routes.

Figure 4.6 shows the longitudinal lower ball joint channel. This is the most gradient-sensitive response path. At 30 km/h all methods under-recover $f6$, indicating that small local slopes and event edges are difficult to preserve in the lowest-speed bin. At 50 km/h, Direct C and WMCR recover the damage level more clearly than the IRI variants, but WMCR still has a lower range ratio, which means that some peak amplitude is not fully recovered. At 70 and 100 km/h the damage ratios become method dependent: Direct C and WMCR can over-predict $f6$ damage, while the range values remain closer to the reference than the damage values. This confirms that matching the longitudinal load path is not only a roughness-level problem. Event width, edge shape, event placement, and tyre-contact filtering all affect the counted damage cycles.

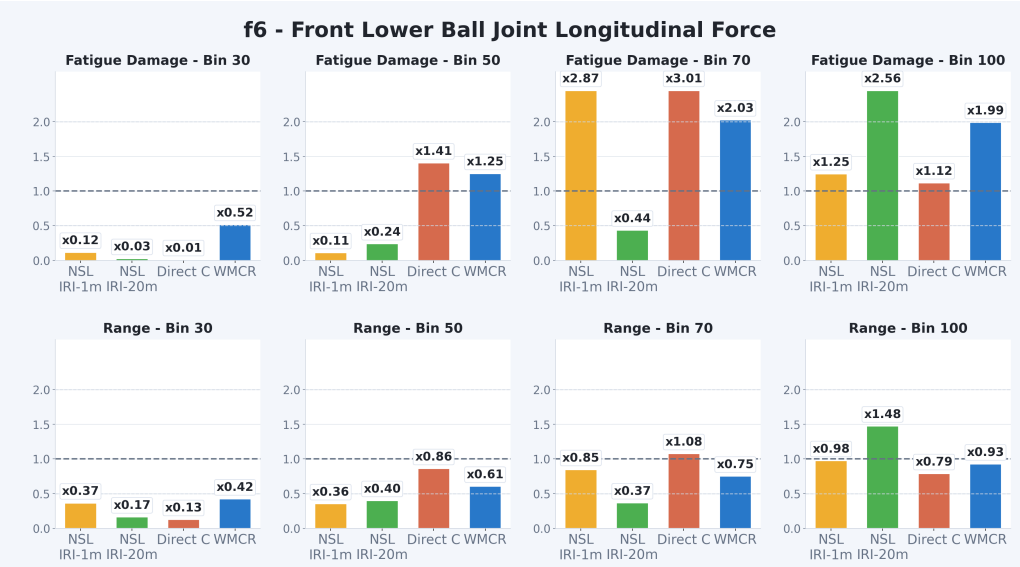

Figure 4.6: Adams $f6$ response over the four speed bins.

Figure 4.7 presents the front anti-roll-bar drop-link channel. The $f63$ result separates lateral surface plausibility from single-track roughness matching. The IRI-based NSL routes can become very severe in this channel, especially in the 50, 70, and 100 km/h bins, which indicates excessive or poorly balanced left–right differential excitation. Direct C is generally less extreme than the IRI variants, but it still over-predicts $f63$ damage in some bins and under-predicts it in the 100 km/h bin. WMCR is more controlled: its $f63$ damage and range are generally conservative rather than amplified. This is not a perfect match, especially at 100 km/h, but it supports the interpretation that the WMCR lateral reconstruction avoids broad artificial anti-roll-bar excitation.

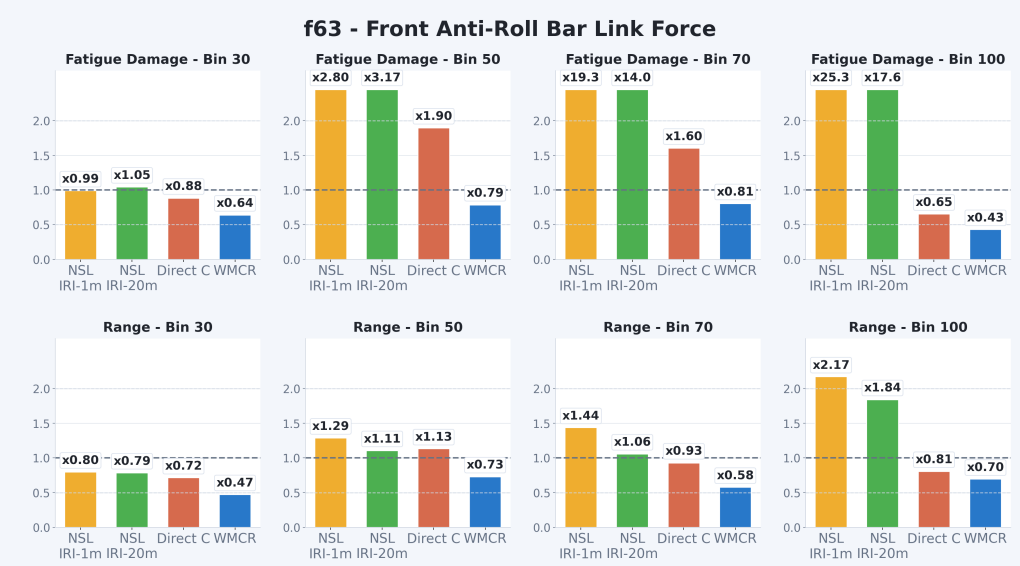

Figure 4.7: Adams $f63$ response over the four speed bins.

Taken together, the channel plots show why the methods should not be judged by

a single 50 km/h comparison or by a single scalar road descriptor. The IRI variants can be too smooth in some low-speed cases and too severe in high-speed response channels. Direct C is the strongest NSL baseline in several range comparisons and in some damage channels, but it can also recover damage through channel-specific over-prediction. WMCR does not reproduce every channel exactly, but it gives a more balanced response pattern: it tends to keep the vertical strut path and the left–right differential path controlled, while the remaining difficulty is concentrated in the gradient-sensitive $f6$ path.

4.3.3 Overall Accuracy and Cross-Bin Robustness

The channel plots provide the physical interpretation, but an aggregate view is needed to compare methods over all bins and channels. Here, the full Adams response set does not mean every output signal available from Adams Car. It means the six selected monitored force histories used in this validation: left and right $f55$ strut force, left and right $f6$ lower ball joint longitudinal force, and left and right $f63$ anti-roll-bar drop-link force. These six channels are chosen because they cover the main response mechanisms targeted by the compact-road validation: vertical suspension loading, gradient-sensitive longitudinal loading, and left–right differential excitation. Each channel is evaluated in the four speed bins using both fatigue damage and range. Figure 4.8 condenses this selected response set into a median absolute error relative to the measured reference.

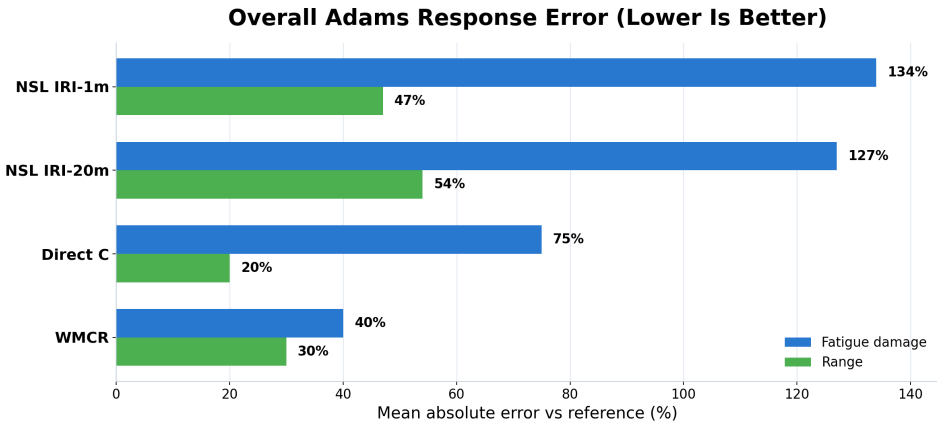

Figure 4.8: Overall Adams response error across methods.

WMCR gives the lowest fatigue-damage error among the compared methods and remains competitive in range error. Direct C has the lowest overall range error, but its fatigue-damage error is substantially larger than WMCR's. Figure 4.9 keeps the same error definition but separates the result by speed bin, so it shows whether the same method remains robust when road input is mapped into different temporal excitation ranges by vehicle speed.

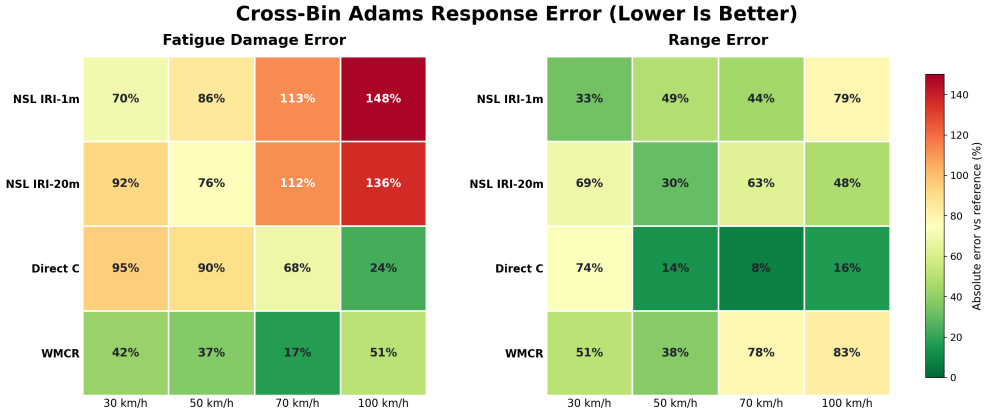

Figure 4.9: Cross-bin Adams response error.

Taken together, the aggregate figures show that the IRI variants have the largest overall damage errors and the strongest speed-bin sensitivity. This confirms that a compact scalar roughness route can miss or over-amplify fatigue relevant cycles when all speed bins and channels are considered together.

4.4 Summary

The compact roads are validated through a progressive chain applied to the four speed bins: PSD checks the spectral roughness content, quarter car pseudo damage adds a reduced vertical response screen, and the non-linear Adams Car benchmark tests component load paths in the full vehicle model. This progression is necessary because profile similarity and reduced vertical severity do not guarantee durability response equivalence.

The final response evidence is channel based. In $f55$, WMCR generally preserves the vertical strut-loading regime more consistently than the descriptor-based NSL routes, although some bins still show conservative or elevated damage. In $f6$, all methods remain sensitive to local event geometry, confirming that longitudinal force is the most difficult path for compact road reconstruction. In $f63$, WMCR is generally controlled rather than over-amplified, which supports the left–right coherence treatment but also shows some conservative under-recovery in high-speed bins.

The aggregate evidence supports the same interpretation. WMCR gives the lowest overall Adams fatigue-damage error and more stable cross-bin damage behaviour, while Direct C remains a strong NSL baseline for range but is less reliable in fatigue damage. The preliminary transient and quarter-car evidence explains why explicit localized content matters: transients account for a small share of surface energy but a much larger share of pseudo damage, and WMCR remains close to the measured quarter-car damage level before full vehicle validation. The main conclusion is therefore that durability validation must be channel-based and speed-bin aware,

and that preserving localized transient structure explicitly improves the robustness of compact-road response equivalence.

5

Discussion and Conclusions

5.1 Discussion

The main implication is that compact road models need to be evaluated by the excitation they preserve, not only by the amount of compression they achieve. In the validation cases reported here, the compact roads are generated with a compression ratio of $\kappa = 4$, so the target road length is $L^{\mathrm{tar}} = L^{\mathrm{src}}/4$. This corresponds to a four-to-one length compression, or a 75% reduction in simulated road length relative to the measured source road. The transient content introduced in the road nonstationarity discussion remains central to assessing whether this compression is acceptable. Long road sections can be relatively mild, while short local events can dominate fatigue response. The practical difference between a road that matches an average roughness level and a road that can replace measured input in durability simulation is whether the severe cycles produced by localized structure remain present after compression.

NSL and WMCR therefore represent two complementary responses to road nonstationarity. NSL makes nonstationarity compact by describing the evolution of road severity statistically. This is suitable when large road populations need to be summarized or when overall roughness variability is the main target. Its limitation is that transient location, shape, and lateral footprint are retained only indirectly. WMCR lowers the level of abstraction by carrying more of this information explicitly. This provides direct engineering control over the road features that are most relevant to response: local severity, event shape, and left–right coherence.

The results also refine the meaning of compression. In this context, compression is not only a reduction in road length, storage size, or number of parameters. It is a modelling decision about which features remain explicit and which features are represented statistically. For relatively homogeneous roads, a compact stochastic description may be sufficient. For roads whose durability effect is governed by localized severe events, the reduced representation needs to preserve enough transient information for cycles critical to response to survive the modelling process. This is illustrated by the Chapter 4 component split, where $N = 4$ speed-bin centreline WMCR profiles are analysed for the 30, 50, 70, and 100 km/h bins. In that screening calculation, the transient RMS share is computed from the centreline elevation

profile, and the response contribution is computed from quarter car pseudo damage before and after transient injection. Transients contribute only $8.22\% \pm 2.89\%$ of RMS elevation energy, but the before–after quarter car comparison attributes $36.27\% \pm 6.69\%$ of the damage increment to the same component. A small energy share can therefore carry a much larger fatigue share.

The channel-based interpretation provides a practical diagnostic link between road representation errors and vehicle response mechanisms. A mismatch in $f55$ primarily indicates an error in vertical roughness content or in the suspension loading produced by localized vertical excitation. A mismatch in $f6$ indicates insufficient preservation of local event shape, edge sharpness, slope, event width, or tyre-contact-relevant geometry. A mismatch in $f63$ indicates an error in left–right coherence or in the lateral construction of the surface, including differential roughness and the lateral footprint of localized events. The Adams Car benchmark is therefore not only a final comparison plot; it is also a tool for identifying which part of the compact road representation needs refinement.

This interpretation supports an iterative road design workflow. PSD and quarter car pseudo damage provide efficient early filters for spectral and reduced response errors. The full vehicle benchmark then determines whether those errors affect component loads. If the generated road agrees in PSD and quarter car severity but deviates in $f55$, the vertical roughness content or vertical transient severity requires attention. If the deviation is concentrated in $f6$, the local event geometry, edge shape, and slope require refinement. If the deviation appears in $f63$, the two dimensional surface construction and left–right coherence require further checking. The goal of the workflow is not to reproduce every visible detail of the measured profile, but to preserve the road excitation mechanisms that govern vehicle dynamics and fatigue.

5.2 Limitations

Several limitations define the scope of the conclusions. First, the final response benchmark is based on one vehicle model, one simulation environment, one manoeuvre setting, and a limited set of monitored force channels. The selected channels cover important but incomplete load paths: $f55$ is mainly a single-wheel vertical strut-loading channel, $f6$ is a longitudinal lower-ball-joint loading channel, and $f63$ is a left–right differential loading channel through the front drop link. They do not provide a dedicated pure lateral response channel. The conclusion therefore applies to the tested benchmark rather than to all vehicle architectures, components, payload conditions, speeds, manoeuvres, or durability targets.

The channel scope also means that additional response outputs could change or refine the method ranking. Relevant extensions include steering or lateral tyre force, yaw and roll response, rear axle channels, body acceleration, and wheel load variation. These channels could reveal response mechanisms that are not fully represented by

the present $f55$, $f6$, and $f63$ set, especially for manoeuvres or vehicles where lateral dynamics and rear-axle loading are more important.

Second, the measured road database defines the domain in which the compression methods can be calibrated and evaluated. This is not a shortcoming of the compression objective itself: NSL and WMCR are intended to preserve response-relevant information from the available measured roads, not to invent road states that are absent from the source population. The practical consequence is a scope boundary. Rare events, seasonal effects, regional pavement types, maintenance states, and very long route compositions should be represented in the source database, or treated as additional calibration bins, before the generated compact roads are used as population-general durability inputs.

Third, compression necessarily removes information. The objective is to remove redundant or response-irrelevant detail, but some combinations of roughness, transient shape, topography, and lateral structure can still be lost. This is visible in the WMCR range and peak behaviour, where some local response amplitudes remain lower than the measured reference. The $f6$ channel is the clearest example: even when the overall road representation is response relevant, longitudinal suspension forces remain sensitive to local event geometry that is difficult to preserve in a compact reconstruction.

Fourth, the transient definition is context dependent. A waveform that is clearly transient on a smooth background can be ordinary texture on a rough background, so transient selection depends on the local roughness condition represented in WMCR by $C_{\mathrm{bg},m}$. During synthesis, the local background, event ordering, spacing, and blending determine how inserted events couple to the generated road surface. These choices can change the vehicle response even when the stored event shapes are unchanged.

Finally, response equivalence is interpreted through selected metrics rather than through formal acceptance limits. Damage, range, maximum, and minimum values make the comparison physically meaningful, but they do not define universal tolerance bands for accepting or rejecting a generated road. The Adams Car benchmark is more representative than a reduced quarter car model, but it remains a simulation benchmark. Tyre–road contact assumptions, model parameters, channel selection, and post-processing choices all affect the response comparison. The thesis conclusion is therefore best interpreted as a benchmarked engineering assessment of NSL and WMCR, not as a complete proof of equivalence for all durability applications.

5.3 Conclusions

The investigation addressed how long measured public roads can be represented by compact virtual roads without losing the excitation mechanisms that govern vehicle

loads and fatigue. The durability motivation follows Charles' work on deriving test severity from measured road data [2]. It is also consistent with Bogsjö's road-profile statistics, which emphasize the fatigue relevance of local road irregularities [1], and with the fatigue study by Bogsjö and Rychlik, where local road content is shown to affect vehicle damage [20]. The results support the same principle: road equivalence cannot be defined by visual similarity, global PSD agreement, or a single roughness descriptor alone. A compact road can be statistically plausible as a profile while still being incomplete as a durability input if local content critical to response is removed.

Two modelling strategies were compared. NSL represents an implicit stochastic route, where local severity is absorbed into a compact parameter process and the road is regenerated from that description. This gives an efficient and interpretable representation of population-level roughness variation. WMCR represents an explicit reconstruction route, where topography, background roughness, localized transients, and lateral surface structure are separated and then recombined. This keeps severe local content visible in the reduced road representation. The comparison therefore addresses the main modelling question: whether localized road structure can be represented adequately through compact severity statistics, or whether it needs to be retained explicitly when the end use is durability simulation.

The validation chain shows why both road-based and response-based checks are needed. The preliminary PSD comparison confirms that the compact roads reproduce the broad spectral decay of the measured road, but it also shows that local spectral peaks and short-wavelength content are not preserved equally. The quarter car pseudo damage comparison converts these differences into a reduced vertical response measure. NSL generally under-recovers severity where localized content is important, while WMCR retains more severe response content and can be conservative at some lateral positions. These checks are informative early filters, but they do not include suspension geometry, tyre contact filtering, multibody load transfer, or left–right coupling. The full vehicle benchmark is therefore required to assess durability response equivalence.

The Adams Car benchmark provides the component-level response assessment. The final response screening is organised over the 30, 50, 70, and 100 km/h speed bins, and the response chapter presents the results directly by monitored channel rather than as separate NSL and WMCR case studies. The monitored force channels separate three physical load paths: $f55$ represents vertical strut loading, $f6$ represents longitudinal lower ball joint loading driven by local gradients and event shape, and $f63$ represents left–right differential excitation through the front drop link. This channel interpretation makes it possible to distinguish a general roughness error from a local event-shape error or a lateral surface-construction error.

The NSL baselines reveal two limitations of descriptor-based compact roads. The IRI variants can miss fatigue relevant cycles in some bins and become too severe in others because a scalar roughness descriptor does not preserve enough localized transient

structure. Direct C is the strongest NSL baseline in several range comparisons and reduces some of the IRI-route errors, but the all-bin response figures show that it can still recover damage through channel-specific over-prediction. The range response therefore shows that recovered damage is not automatically equivalent to recovered load behaviour. A damage match is meaningful only when the response envelope and the channel balance remain physically plausible.

WMCR gives the more balanced response reconstruction within the investigated benchmark and monitored channels. The overall Adams error summary shows the lowest fatigue-damage error for WMCR, while the range checks indicate that this agreement is not obtained through broad artificial amplification. The channel results explain the mechanism. In $f55$, WMCR generally retains the main vertical strut loading through the combination of background roughness and explicit transient content. In $f63$, the response remains controlled, supporting the two dimensional surface reconstruction and left–right coherence treatment. The clearest remaining limitation is $f6$, where local gradients, event width, event placement, and tyre contact filtering make the longitudinal force sensitive to small errors in transient shape.

With these scope boundaries in mind, the central conclusion is limited to the investigated four-bin Adams Car response screening and the monitored force channels. Within this scope, WMCR better satisfies the modelling objective: it provides a compact road representation that preserves fatigue relevant response more consistently across the selected vehicle force channels while maintaining controlled amplitude behaviour. NSL remains a valuable stochastic route when the target is compact representation of population-level roughness variation. However, when localized events, transient geometry, and lateral surface structure are important for the tested durability responses, the WMCR component representation provides the stronger engineering basis.

References

[1] K. Bogsjö, *Road Profile Statistics Relevant for Vehicle Fatigue*. Ph.D. dissertation, Dept. of Mathematical Statistics, Centre for Mathematical Sciences, Lund Univ., Lund, Sweden, 2007.

[2] D. Charles, "Derivation of environment descriptions and test severities from measured road transportation data," *Journal of the IES*, vol. 36, pp. 37–42, 1993.

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

[4] International Organization for Standardization, *ISO 13473-1:2019 Characterization of Pavement Texture by Use of Surface Profiles — Part 1: Determination of Mean Profile Depth*, Geneva, Switzerland, 2019.

[5] International Organization for Standardization, *ISO 13473-3:2002 Characterization of Pavement Texture by Use of Surface Profiles — Part 3: Specification and Classification of Profilometers*, Geneva, Switzerland, 2002.

[6] B. Bruscella, V. Rouillard, and M. Sek, "Analysis of road surface profiles," *Journal of Transportation Engineering*, vol. 125, no. 1, pp. 55–59, 1999, doi: 10.1061/(ASCE)0733-947X(1999)125:1(55).

[7] P. Johannesson, K. Podgórski, and I. Rychlik, "Laplace distribution models for road topography and roughness," *International Journal of Vehicle Performance*, vol. 3, no. 3, pp. 224–258, 2017, doi: 10.1504/IJVP.2017.085032.

[8] C. J. Dodds and J. D. Robson, "The description of road surface roughness," *Journal of Sound and Vibration*, vol. 31, no. 2, pp. 175–183, 1973, doi: 10.1016/S0022-460X(73)80373-6.

[9] P. Andrén, "Power spectral density approximations of longitudinal road profiles," *International Journal of Vehicle Design*, vol. 40, no. 1–3, pp. 2–14, 2006, doi: 10.1504/IJVD.2006.008467.

[10] M. W. Sayers, T. D. Gillespie, and C. A. V. Queiroz, *The International Road Roughness Experiment: Establishing Correlation and a Calibration Standard for Measurements*. Washington, DC, USA: World Bank, 1986.

[11] V. Rouillard, "Using predicted ride quality to characterise pavement roughness," *International Journal of Vehicle Design*, vol. 36, no. 2/3, pp. 116–131, 2004, doi: 10.1504/IJVD.2004.005352.

[12] V. Rouillard, "Decomposing pavement surface profiles into a Gaussian sequence," *International Journal of Vehicle Systems Modelling and Testing*, vol. 4, no. 4, pp. 288–305, 2009, doi: 10.1504/IJVSMT.2009.032021.

[13] D. Maher and P. Young, "An insight into linear quarter car model accuracy," *Vehicle System Dynamics*, vol. 49, no. 3, pp. 463–480, 2011, doi: 10.1080/00423111003631946.

[14] W. Fauriat, *Stochastic Modeling of Road-Induced Loads for Reliability Assessment of Chassis and Vehicle Components Through Simulation*. Ph.D. dissertation, Université Blaise Pascal – Clermont-Ferrand II, Clermont-Ferrand, France, 2016.

[15] P. Johannesson and I. Rychlik, "Laplace processes for describing road profiles," *Procedia Engineering*, vol. 66, pp. 464–473, 2013, doi: 10.1016/j.proeng.2013.12.099.

[16] P. Johannesson and I. Rychlik, "Modelling of road profiles using roughness indicators," *International Journal of Vehicle Design*, vol. 66, no. 4, pp. 317–346, 2014, doi: 10.1504/IJVD.2014.066068.

[17] K. Bogsjö, K. Podgórski, and I. Rychlik, "Models for road surface roughness," *Vehicle System Dynamics*, vol. 50, no. 5, pp. 725–747, 2012, doi: 10.1080/00423114.2011.637566.

[18] F. Öijer and S. Edlund, "Identification of transient road obstacle distributions and their impact on vehicle durability and driver comfort," in *The Dynamics of Vehicles on Roads and on Tracks*, pp. 744–753, 2004.

[19] F. Wang, I. M. Al-Qadi, S. H. Yang, and L. Zhang, "Continuous wavelet analysis of pavement profiles," *Automation in Construction*, vol. 63, pp. 134–143, 2016, doi: 10.1016/j.autcon.2015.12.013.

[20] K. Bogsjö and I. Rychlik, "Vehicle fatigue damage caused by road irregularities," *Fatigue & Fracture of Engineering Materials & Structures*, vol. 32, no. 5, pp. 391–402, 2009, doi: 10.1111/j.1460-2695.2009.01340.x.

[21] A. Palmgren, "Die Lebensdauer von Kugellagern," *Zeitschrift des Vereins Deutscher Ingenieure*, vol. 68, no. 14, pp. 339–341, 1924.

[22] M. A. Miner, "Cumulative damage in fatigue," *Journal of Applied Mechanics*, vol. 12, no. 3, pp. A159–A164, 1945.

[23] M. W. Sayers, "On the calculation of international roughness index from longitudinal road profile," *Transportation Research Record*, no. 1501, pp. 1–12, 1995.

[24] J. W. Tukey, *Exploratory Data Analysis*. Reading, MA, USA: Addison-Wesley, 1977.

[25] S. G. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, 1989, doi: 10.1109/34.192463.

[26] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA, USA: SIAM, 1992.

[27] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994, doi: 10.1093/biomet/81.3.425.

[28] ASTM International, *ASTM E1049-85(2023) Standard Practices for Cycle Counting in Fatigue Analysis*. West Conshohocken, PA, USA, 2023, doi: 10.1520/E1049-85R23.

[29] ASAM e.V., *ASAM OpenCRG Specification*, version 1.2.0. Höhenkirchen, Germany, 2020.

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY