

Applying Signal Processing Techniques to Characterize Rail Corrugation, Noise, and Vibration

**by
Julian Carneiro**

A Thesis submitted to the Faculty of Graduate Studies of
The University of Manitoba
in partial fulfilment of the requirements of the degree of
MASTER OF SCIENCE

Department of Electrical and Computer Engineering

University of Manitoba

Winnipeg, Manitoba, Canada

Copyright © 2023 by Julian Carneiro

Abstract

This thesis presents the application of signal processing techniques to rail transit data for the purposes of rail corrugation maintenance. There are three primary forms of data in this work to which signal processing techniques were applied: (1) the analysis of rail corrugation as a spatial signal, (2) the analysis of wayside noise and noise spectra, and (3) balancing time and frequency resolution for inspection of on-board vibration signals.

Grinding maintenance is common for addressing deteriorating rail surface condition, as surface defects like corrugation are removed. However, measuring the rail surface condition and then making the actionable decision to grind can be complex. This work develops a summary of corrugation roughness $RCI_x(\lambda)$, an extension of RCI_{2018} which uses a representative block root-mean square (RMS) measurement to summarize corrugation across a segment.

The act of measuring surface roughness is a time-consuming task for property maintainers, and the act of measuring rail that does not require maintenance is wasteful. The measurement and analysis of wayside noise can give insight into the state of the rail to help property maintainers determine whether rail in a particular location is believed to require direct measurement and intervention, thus minimizing wasted effort. $Lv_{eq}(\lambda, v)$ is developed in this thesis as a noise weighting scheme meant to isolate corrugation-related noise from a wayside noise sample spectrum.

The usage of vehicle-mounted accelerometers can provide massive amounts of up-to-date information on rail condition through regular revenue service. This scale of data collection can streamline measurement and maintenance decisions throughout a property, however, the information encoded in this data requires post-processing to identify the characteristics of interest. This work utilizes the wavelet transform as a time-frequency analysis tool to spatially locate a particular frequency attribute in an acceleration sample.

The results presented herein showcase the effectiveness of the signal processing techniques applied to these data streams. The growth of corrugation and its response to grinding were demonstrated in both corrugation and wayside noise datasets, while the wavelet transform has demonstrated a sensitivity to characteristics of interest in on-board vibration data.

Contributions

Journal Papers

A. Lasisi, **J. Carneiro**, J. D. Regehr, I. Jeffrey, E. Magel, S. Chénier, M. Reimer, “A rail corrugation index to characterize noise impacts and grinding effectiveness on rail transit systems,” Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit., vol. 0, no. 0, 2023.

Personal contributions: Data preparation, data analysis, visualization, manuscript preparation.

J. Carneiro, A. Lasisi, J. D. Regehr, I. Jeffrey, E. Magel, S. Chénier, M. Reimer, “Impacts of Noise Filtration and Train Velocity on Measuring Grind Effectiveness when Targeting Rail Corrugation,” Submitted to Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit.

Personal contributions: Data preparation, data analysis, visualization, manuscript composition, manuscript preparation, manuscript submission.

Conference Papers

A. Lasisi, J. D. Regehr, **J. Carneiro**, I. Jeffrey, E. Magel, S. Chénier, M. Reimer, “Characterizing rail corrugation through a data-driven assessment of impacts on noise and vibration,” Paper accepted to CCRC 2021.

Personal contributions: Data preparation, data analysis, visualization, manuscript preparation, led conference presentation.

Acknowledgements

Firstly, I want to thank my advisors, Dr. Ian Jeffrey and Dr. Jonathan Regehr, for taking the chance to work with and support me through my M.Sc. I couldn't have done any of this without your guidance, support, and continued friendship, and I will always be thankful for that.

I would like to thank Dr. Babak Mehran and Dr. Sherif Sherif for being members of my M.Sc. reviewing committee, and taking the time to evaluate my research.

I would like to thank Dr. Ahmed Lasisi, for his work on the NRC project that started this research. His expertise and dedication at every turn in that project made it a joy to be a part of.

Thank you to the rest of the people involved in the NRC project, specifically, Eric Magel, Mark Reimer, and Sylvie Chénier. Thank you for sharing your knowledge with me, and helping me find an appropriate space for my work in this sector. Thank you also to Dr. Puteri Paramita, Sanjay Abraham, and Trevor Nickel for their collaborative contributions on the NRC project.

Thank you to those at ARM and Mitacs for their financial support which made this research possible. Thank you to ARM, specifically Mark Reimer, for providing all of the data used for this research, and supporting the exploration of it at every turn.

Mom. Dad. Claire.

And the cats too.

Contents

List of Figures	viii
List of Tables	xiii
1 Introduction	1
1.1 Purpose	1
1.2 Background and Motivation	1
1.3 Objectives and Scope	5
1.4 Approach	6
1.5 Thesis Outline	7
2 Data Background and Description	8
2.1 Corrugation Data	8
2.2 Wayside Noise Data	12
2.3 Vibration Data	14
2.4 Chapter Summary	19
3 Corrugation Analysis	20
3.1 RCI_x for Monitoring Corrugation	20
3.1.1 Property A	21
3.1.2 Property B	24
3.2 Wavelength-Adapted $RCI_x (RCI_x(\lambda))$	26
3.2.1 Preliminary Investigation	26
3.2.2 Quantifying Corrugation Improvement using $RCI_x(\lambda)$	26
3.2.3 $RCI_x(\lambda)$'s Sensitivity to Post-grind Measurement Delay	29
3.2.4 Performance of $\% \Delta RCI_x$ for Low Pre-grind $RCI_x(\lambda)$	30
3.3 Chapter Summary and Next Steps	31

4	Analysis of Wayside Noise Generated by Trains	33
4.1	LA_{eq} and Spectrum Analysis	33
4.1.1	Property A Analysis	33
4.1.2	Property B Analysis	36
4.2	Velocity- and Wavelength-Adapted Boxcar Filtering ($Lv_{eq}(\lambda, v)$)	40
4.2.1	Property B Analysis	40
4.2.2	Property A Analysis	42
4.3	Chapter Summary and Next Steps	44
5	Vibration Analysis	46
5.1	ABAs and ABA Data	46
5.1.1	Data Characteristics	46
5.1.2	Towards Practical Use of ABA Data	49
5.1.3	Fourier Transform and Short Time Fourier Transform (STFT, Spectrogram)	52
5.2	ABA Data Analysis with the Wavelet Transform	53
5.2.1	Motivation For Using the Wavelet Transform	53
5.2.2	Continuous Wavelet Transform (CWT)	56
5.2.3	Discrete Wavelet Transform (DWT)	57
5.2.4	Stationary Wavelet Transform (SWT)	58
5.2.5	Applying the Wavelet Transform to ABA Data	60
5.3	Chapter Summary and Next Steps	75
6	Conclusion and Future Work	77
6.1	Summary of Key Findings and Contributions	77
6.2	Suggestions for Future Work	78
6.2.1	Corrugation and Rail Roughness	78
6.2.2	Noise	79
6.2.3	Vibration	80
	Appendix A Wavelet Transform Theory	81
A.1	Continuous Domain	81
A.1.1	Constructing a Family of Wavelets	81
A.1.2	Defining the Continuous Wavelet Transform	82
A.1.3	Analytic Wavelets, and Time-Frequency Resolution	83
A.1.4	The Scaling Function	85
A.2	Discrete Domain	86
A.2.1	Discretizing Wavelet Theory	86

A.2.2	Subband Encoding	88
A.2.3	Defining the Discrete Wavelet Transform	90
A.2.4	Shift-invariance (The Stationary Wavelet Transform)	91
	Bibliography	95

List of Figures

1.1	Block diagram illustrating a general corrugation mechanism [4].	2
1.2	Grinding car correcting rail corrugation. The sparks flying off are from the stones grinding on the metal rails.	3
1.3	Pre-grind and post-grind rail. Note that the wide flat contact point has been removed, and artificial roughness known as the grind signature has been left as a byproduct of maintenance. This will wear out into smooth rail through normal operation.	4
1.4	Illustration of “magic wear rate”, adapted from [9]. Removal of surface metal extends rail life by reducing the impact of RCF. However, it also reduces rail life as the rail surface is a finite resource. The optimal metal removal rate to maximize rail life is the intersect of these curves.	5
1.5	Workflow organization of this thesis. Corrugation and wayside noise research was done in tandem due to the paired nature of the datasets, while vibration research was completed in isolation.	6
2.1	Rail illustrating the running bands of interest at a particular property. The markings on the rail indicate that this is the right rail (RR), and the gauge side (G) of the rail is at the bottom of the image. The two ticks on the rail face denote 25 and 40 mm from gauge side, and the majority of the contact between the wheel and rail falls in the band between them, evidenced by the surface wearing.	9
2.2	Operator logging vertical displacement on the far rail face using the original CAT. The blue box on which the laptop is resting contains the sensor hardware.	9
2.3	Screenshot of RailMeasurement CAT software Version 14.26.21606 visualizing an unfiltered corrugation measurement and the calculated $1/3$ octave spectrum of the data.	10
2.4	Sample corrugation $1/3$ octave spectrum compared with the ISO 3095:2005 and EN 13231-3:2012 standards. A valid interpretation of this chart is that this surface roughness passes the EN standard, but fails the ISO standard.	11
2.5	L_{eq} noise for a given fluctuating noise sample. Note that if the dB signal were just integrated and averaged, the resulting calculation would slightly underrepresent the true L_{eq}	13

2.6	Regions of the A-weighting curve where sound is attenuated or amplified. The intercepts with 0 dB (no gain), denoted with a black \times , are located at 1000 Hz and \sim 6300 Hz.	14
2.7	ABA mounted on a vehicle, with the IMU's frame of reference.	15
2.8	Free body diagrams of acceleration captured by ABA at rest (a), when slid across the table (b), and when friction slows it down (c).	16
3.1	Corrugation growth captured by RCI_{65} , RCI_{85} , and RCI_{95} separated by rail and track geometry at Property A.	23
3.2	Pre- and post-grind $1/3$ octave spectra from one location in Property B. The large peak centred at approximately 100 mm appears to have been removed from grinding, while low wavelength roughness was leftover as the grind signature.	25
3.3	Box plot distribution of $\Delta RCI_x(\lambda)$ for the different tested wavelengths and percentiles ($n = 34$).	27
3.4	Box plot distribution of $\% \Delta RCI_x(\lambda)$ for the different tested wavelengths and percentiles ($n = 34$).	28
3.5	Distributions of $\Delta RCI_x(\lambda)$ based on how long after grinding the data was sampled.	29
3.6	4×3 grid of plots visualizing each percentile and wavelength band combination.	30
4.1	Correlation matrices of corrugation and wayside noise metrics for tangent and curve rail sections. The black boxes identify where wayside noise metrics and corrugation metrics are correlated. Note that the colourbar scale spans 0.4-1.0 for visual emphasis.	35
4.2	Corrugation growth captured by RCI_{65} , RCI_{85} , and RCI_{95} , and wayside noise growth captured by LA_{eq65} , and LA_{eq95} , separated by rail and track geometry at Property A.	37
4.3	Unweighted spectra of samples at a particular curve at Property B. Pre-Grind LA_{eq} was between 72.7 and 74.1 dB, while Post-Grind LA_{eq} was between 83.1 and 86.5 dB, as presented in the legend. Note the post-grind increase in noise where A-weighting is especially sensitive, denoted by black dotted lines, and how this is reflected in the dBA of each sample.	38
4.4	A-Weighting Curve showing the overlap between its passband and the usual noise frequencies of Top of Rail Squeal. The vertical black dotted lines represent the locations where A-Weighting has a gain of 0 dB or greater	40
4.5	An unweighted noise sample, with two different $Lv_{eq}(\lambda, v)$ weights applied to the spectrum.	42
4.6	A-weighting and v -weighting comparison at two curves. Figure 4.6a illustrates v -weighting disagreeing with A-weighting, while Figure 4.6b showcases agreement.	43
4.7	Box and whisker distribution plots for and $Lv_{eq}(30-100 \text{ mm}, 80-90 \text{ km/h})$, as a function of wear.	44

5.1	Path of travel for two trips , one up a line and one down a line of stations A through D, with the ABA’s frame of reference which is constant throughout both trips. Trip one is eastbound along the north track, made up of segments A→B, B→C, and C→D. Trip two is the reverse order (westbound on the south track) from station D to A. An interlock east of Station D (not pictured) allows a train to switch to the westbound track without turning the vehicle, keeping the positive x direction constant.	48
5.2	x , y , and z acceleration data for a station-to-station segment.	49
5.3	An abstract “positional uncertainty” as a function of position between two stations at the start and end of a segment. Static track characteristics like switches, which can anchor position in a segment, are marked with a blue \times in the top graph, and the positional uncertainty is 0 in the bottom graph at these locations. At locations between these key points, dead reckoning error adds uncertainty in the positional calculation.	51
5.4	x , y , and z dimensions FFTs calculated from acceleration during one station-to-station segment.	53
5.5	z acceleration, spectrogram with 50% overlap and window length of 1 second, spectrogram with 50% overlap and window length of 10 seconds.	54
5.6	A base mother wavelet $\psi_{s=1,u=0}(t)$ (top), shifted $\psi_{s=1,u=-100}(t)$ (middle), $\psi_{s=2,u=125}(t)$ (bottom). Note that this mother wavelet does not satisfy all the conditions to be a wavelet by definition as it does not have zero mean, however it illustrates the effects of changing s and u	55
5.7	Complex Morlet Wavelet in time, and its frequency response in radians at various scales s . Note the shift of centre frequency, as well as reduction of bandwidth as the scale increases. It is a one-sided frequency response because it is a complex function.	56
5.8	Continuous wavelet transform and spectrogram representation of two frequency-modulated Gaussian waves, called Gabor Pulses. Notice the non-linearity of the scale axis of the CWT.	57
5.9	Block diagram of DWT cascading filter algorithm. The \downarrow denotes the downsampling operator. $h[n]$ and $g[n]$ are the wavelet and scaling filters. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.	58
5.10	Maximum level DWT decomposition of Gabor Waves using the Haar wavelet. Upsampling each level by repeating 2^{level} times allows the transform to be plotted in this way.	59
5.11	Maximum level unnormalized SWT decomposition of Gabor Waves using the Haar wavelet. Unlike the DWT, this transform is time-invariant, and does not need to be upsampled at each level.	60
5.12	Block diagram of SWT cascading filter algorithm. $h[n]$ and $g[n]$ are the wavelet and scaling filters, and the j dots above the filter denote $2^j - 1$ holes inserted into the filter. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.	60
5.13	High frequency energy distributions of cD1 coefficients of wavelet families from 10 different ABA signals.	62

5.14	Low frequency energy distributions of cA1 coefficients from wavelet families from 10 different ABA signals.	63
5.15	DWT and SWT energies colour-coded according to which of the 10 ABA segments it is from. Note the terracing in some families, suggesting that all wavelets in that family are equally sensitive to the cA1 coefficients.	65
5.16	Wavelet and Scaling decomposition functions of Reverse Biorthogonal wavelet family with 3 vanishing moments.	66
5.17	Raw z acceleration signal, the reconstruction from only the cD1 coefficients (DWT), and the cD1 coefficient magnitudes.	69
5.18	Raw z acceleration signal, the reconstruction from only the cD1 coefficients (SWT), and the cD1 coefficient magnitudes.	70
5.19	Raw z acceleration signal, the untouched as well as smoothed (100 sample/0.03 second window) cD1 and cA1 coefficients (DWT), and the index-by-index product of the cD1 and cA1 coefficient magnitudes. Peaks in this plot indicate points where both the low and high frequency information shows a high amplitude, suggesting the location of switches or interlocks.	71
5.20	Quasi-periodic peaks in cA1 coefficients.	72
5.21	$ cD2 $, $ cD3 $, and $ cD4 $ coefficient zoom from $t = 254$ seconds to $t = 266$ seconds.	73
5.22	$ cD2 $, $ cD3 $, and $ cD4 $ coefficient zoom from $t = 116$ seconds to $t = 130$ seconds.	74
A.1	A sample Ricker wavelet with its frequency response. The effect of changing s can be seen in both domains. Note that as $s \rightarrow \infty$, the peak of the frequency response approaches 0. By utilizing all $0 < s$, the entire frequency spectrum $(-\infty, \infty)$ rad/sec of a signal can be represented.	84
A.2	The Ricker wavelet $\psi(t)$, $\hat{\psi}(\omega)$, and its $\hat{\phi}(\omega)$. Note that the scaling function's time domain representation $\phi(t)$ exists but is not visualized due to having to be generated from the inverse Fourier transform of $\hat{\phi}(\omega)$. The $\phi(t)$ associated with the Ricker wavelet does not have a closed form expression to my knowledge.	86
A.3	Passbands for a 3 level discretized wavelet decomposition. f_N denotes the Nyquist sampling rate of the original signal.	87
A.4	Figure A.3 under the coefficient umbrellas of cA_j and cD_j bandwidths for a level j	89
A.5	Two filters that are quadrature mirror to eachother. Note these filters also satisfy the unity power gain condition.	90
A.6	Figure 5.9 repeated for convenience. Block diagram of DWT cascading filter algorithm. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.	90
A.7	Figure 5.12 repeated for convenience. Block diagram of SWT cascading filter algorithm. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.	92
A.8	The effect of inserting zeros on the frequency response of the db10 wavelet.	93

A.9 Obtaining cD2 from the second level decomposition using the db10 wavelet. These coefficients are obtained by first filtering with $g[n]$ then by $\hat{h}[n]$, which has the effective frequency response shown in the bottom plot. 94

List of Tables

1.1	Types of corrugation and their characteristics [4].	3
2.1	ABA IMU Specifications.	18
3.1	Proposed scale and colour scheme for RCI_{2018} [12].	21
3.2	Improvement of RCI_x for different wavelength bands, and their rationale for testing.	26
4.1	Property A Data Summary.	34
4.2	Property B Data Summary. The data in Table 4.1 is repeated for reference.	38
4.3	Common rail noise emissions and their frequency presences, making note of Corrugation and Top of Rail Squeal [35].	39
5.1	Wavelets and wavelet families provided by pywt [46].	61

Chapter 1

Introduction

1.1 Purpose

The purpose of this research is to demonstrate the application of signal processing techniques to characterize rail corrugation and its noise and vibration covariates. Rail corrugation is the manifestation of roughness on the rail face, as a result of the wheel-rail dynamic loads. When trains operate on corrugated rail, they generate unwanted noise and vibration. To mitigate these problems, rail property owners monitor rail corrugation using manually-operated instruments such as a Corrugation Analysis Trolley (CAT). When corrugation or one of its noise or vibration covariates reach unacceptable levels, the most common mitigation measure is to remove the worn material from the surface of the rail through rail grinding to reveal unworn rail.

This thesis focuses on the application of signal processing techniques to better characterize rail corrugation, noise, and vibration. These techniques offer efficient ways to convert large data sets into summary statistics to support practical decision-making by those in industry. The nature of the signal varies for these three measurements, leading to the three main forms of data analysis in this work: (1) the analysis of corrugation as a spatial signal, (2) the analysis of noise and noise spectra, and (3) balancing time and frequency resolution for inspection of vibration signals. Ultimately, these tools offer opportunities to distill large data sets into key indicators that support rail maintenance and asset management decisions.

1.2 Background and Motivation

The passenger rail industry is central in the lives of many people, as it continues to be one of the most popular modes of public transportation in the world. In Canada, most major urban centres (Vancouver, Edmonton, Calgary, Ottawa, Waterloo, Brampton, Mississauga, Toronto, Hamilton, Montreal, and Québec City) have operational or planned public rail transportation systems [1, 2]. The American Public Transit Association (APTA) estimates 2.8 billion passengers were transported by rail in the United States in the year 2021, however this num-

ber has not recovered to the pre-COVID-19 pandemic estimate of 4.8 billion passengers in 2019 [3]. It is therefore crucial for transit agencies to maintain proper upkeep of these rail systems to best serve their riderships.

A primary concern for this industry is the phenomenon known as corrugation. Corrugation is defined in [4] as “a quasi-sinusoidal irregularity with a wavelength of less than approximately one metre”. Corrugation can happen either on the rail or on the wheel, however this thesis focused on investigating corrugation as it manifests on the rail surface.

Corrugation manifests due to feedback mechanisms in the wheel-rail interface over time [4], a process that is illustrated in Figure 1.1. The initial longitudinal rail profile elicits the ‘wavelength-fixing mechanism’, representing the dynamics of the wheel-rail interface. This system interacts with other parameters such as the traction between the wheel and rail, and stick-slip forces like friction. This mechanism fuels the ‘damage mechanism’, which is the physical manifestation of the forces acting upon the wheel and the rail. Most commonly, this manifests as wear at the point of contact, although plastic flow due to shear is another possible surface damage mechanism. This newly damaged rail surface now has a slightly new longitudinal profile, which, in turn, fuels a new wavelength-fixing mechanism. This continues cyclically until the corrugation is addressed by property maintainers.

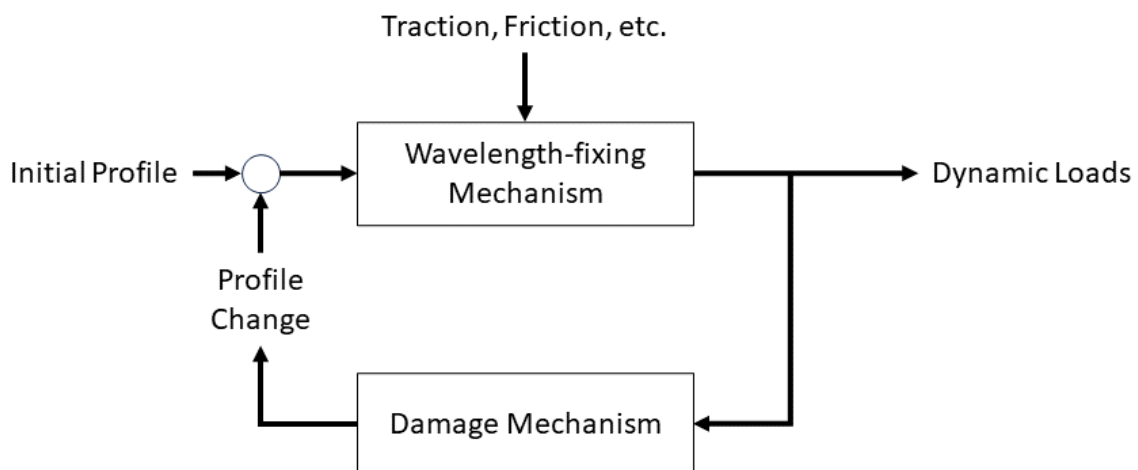


Figure 1.1: Block diagram illustrating a general corrugation mechanism [4].

Various wheel-rail and track dynamics precipitate different types of corrugation, and these different types of corrugation have defining characteristics such as the damage mechanism, or track geometry of common incidence. Some common types of corrugation and their characteristics are presented in Table 1.1.

Corrugation is a focus of property maintainers, as the formation and development of corrugation cause unwanted byproducts. Severe corrugation leads to increased noise and vibration experienced by both the riders and the surrounding public, as well as a reduction in lifespan to components of both the wheel and rail [5].

Additionally, rolling contact fatigue (RCF) is believed to be exacerbated and accelerated by corrugation. RCF is another insidious rail phenomenon and is a major safety concern that drives many system maintenance and replacement decisions. RCF manifests as a variety of rail defects such as cracking of the rail gauge corner and

Type	Wavelength-fixing mechanism	Location	Damage Mechanism
Pinned-pinned resonance	Pinned-pinned resonance	Straight track, high rail of curves	Wear
Rutting	Second torsional resonance of driven axles	Low rail of curves	Wear
Heavy rail	$P2$ resonance	Straight track or curves	Plastic flow in troughs
Light rail	$P2$ resonance	Straight track or curves	Plastic bending
Other $P2$ resonance	$P2$ resonance	Straight track or curves	Wear
Trackform-specific	Trackform-specific	Straight track or curves	Wear

Table 1.1: Types of corrugation and their characteristics [4].

wheel tread, shelling of the wheel and rail, and rail squats [6].



Figure 1.2: Grinding car correcting rail corrugation. The sparks flying off are from the stones grinding on the metal rails.

Management of corrugation is primarily handled by grinding, a process wherein artificial wear (i.e. purposeful removal of material) is imposed onto the rail to remove surface damage. This is done by a grinding machine, seen in Figure 1.2, which uses rough wheels called stones to chip away surface damage. Not unlike sandpaper, the grit of stone dictates the volume of rail removed. In addition to smoothing out roughness and spot defects longitudinally, grinding is also meant to reshape the lateral profile of the rail as needed. By resetting the state of the rail, ideally this improves the wheel-rail contact dynamics and extends the life of both the rail and the wheel. Figure 1.3 illustrates the condition of pre-grind rail, as well as the effect of maintenance on the rail face.

Rail milling is a similar technique to rail grinding, where a machine, instead of grinding metal off the rail, will cut away the rail using a cutting head. This head creates milling chips which are collected by the machine and can be recycled, where grinding disperses metal dust around the rail. Rail milling often necessitates polishing afterwards. Both milling and grinding are ways to achieve the goal of repairing the rail by removing surface defects [7].



(a) Pre-grind rail face.



(b) Post-grind rail face.

Figure 1.3: Pre-grind and post-grind rail. Note that the wide flat contact point has been removed, and artificial roughness known as the grind signature has been left as a byproduct of maintenance. This will wear out into smooth rail through normal operation.

The rail is a finite asset, and grinding necessarily discards some of the asset. This implies that grinding should not be overdone, be that grinding too deep or too often. As such, there is industry desire to move away from a reactive grinding philosophy towards a preventive grinding philosophy. In industry, properties are often grinding once corrugation becomes severe. Due to the feedback nature of corrugation's growth mechanism, the rate of corrugation growth is proportional to corrugation severity, meaning reactive grinding tends to remove more rail than may have been necessary at an earlier point. Preventive grinding aims to remove the minimum amount of metal to treat corrugation, prevent the onset of RCF, and maintain rail profile by grinding before the rail condition deteriorates significantly [8, 6]. There is a desire to optimize the life span of rail by finding what is called the "magic wear rate", which is balancing the removal of rail due to maintenance and the aging of rail due to usage. The optimal grinding rate minimizes the aging process from grinding, while maximizing the intervention with respect to RCF [9].

Taking this preventive grinding philosophy to its logical extreme means that a minuscule amount of metal would be ground every day to keep it in optimal condition. This is obviously impractical and infeasible, so the challenge of preventive grinding is finding a balance between maximizing system performance while minimizing cost, both in terms of dollars spent performing maintenance and in the volume of metal being removed from the rail.




Figure 1.4: Illustration of “magic wear rate”, adapted from [9]. Removal of surface metal extends rail life by reducing the impact of RCF. However, it also reduces rail life as the rail surface is a finite resource. The optimal metal removal rate to maximize rail life is the intersect of these curves.

Theorizing preventive grinding plans is out of the scope of this work. Instead, this work aims to clarify some corrugation assessment techniques so that, hopefully, these techniques can help property maintainers to diagnose rail quality in their preventive grinding plans.

Currently, the rail industry faces challenges in how rail management may be addressed. Rail inspection is time-consuming, costly, localized, and often done reactively as opposed to preemptively. Maintenance takes place in system off-hours (i.e. 2:00 am to 6:00 am), and the surveying range is limited to how much rail maintenance workers can cover in a night. By the time the rail is analyzed, the condition of the rail may have already deteriorated past the point of needing treatment. Additionally, quantifying the severity of corrugation on the rail face is not explicitly defined and agreed upon. Some standards attempt to address grinding quality indices [10, 11], but literature attempting to summarize corrugation severity with the intent to drive grinding decision-making is sparse. There is a place in this sector for an improvement to system monitoring and subsequent decision-making, whereby the quality of rail in a location can be indirectly measured properly through secondary rail characteristics to determine the level of care required.

1.3 Objectives and Scope

This thesis applies techniques from the field of signal processing to address the problem of rail corrugation and the unwanted noise emissions and vibration it causes. Specifically, the objectives of the thesis are:

1. to propose and evaluate a new rail corrugation index that converts raw rail corrugation data (i.e., a spatial signal) into a summary statistic that enables an evaluation of the effectiveness of grinding interventions.
2. to relate wayside noise measurements to rail corrugation and demonstrate, through spectral weighting, the potential use of noise as a predictive surrogate for rail corrugation.

3. to demonstrate the application of signal processing transforms to detect rail corrugation from accelerometer data.

This work uses signal processing techniques to aid those in the rail industry on their decision-making, and does not attempt to define decision-making thresholds. Instead, this work shows the validity and usefulness of these techniques, without having the necessary industry expertise to set actionable limits. Ideally, continuation of this work by those in industry might be to define standards and limits based around these techniques to drive maintenance decisions.

1.4 Approach

This research evolved from a collaborative research contract between the University of Manitoba, Advanced Rail Management (ARM), and the National Research Council of Canada (NRC). The data used is an extension of the data made available to us by ARM through this contract, and it largely came from two Canadian transit properties, referred to as Property A and Property B in this work.

Data from Property A was sampled at two locations in the property; one was at a curved segment, and the other was at a tangent segment (i.e. straight rail). The corrugation and noise data from this property was sampled approximately monthly over the course of two years, with grinding intervention as a part of standard operation. In comparison, the data from Property B was sampled at 15 unique curves throughout the property, and was only sampled immediately pre- and post-grinding intervention.

ARM also provided us with on-board vibration data that had no temporal or spatial connection to any corrugation or noise data, leaving that analysis completely separate. An organizational schematic outlining the research presented in this document can be found in Figure 1.5.

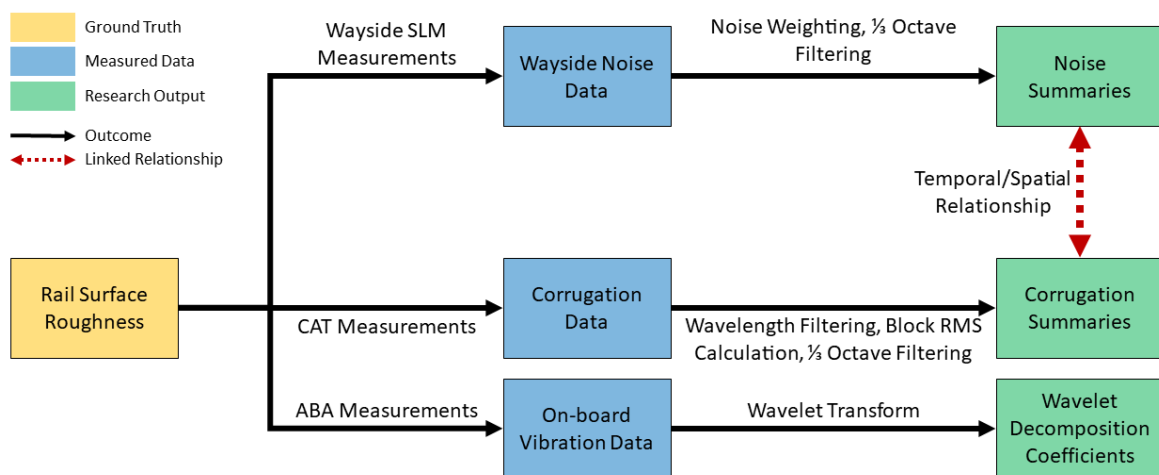


Figure 1.5: Workflow organization of this thesis. Corrugation and wayside noise research was done in tandem due to the paired nature of the datasets, while vibration research was completed in isolation.

In general, the analysis of corrugation involves qualitative investigation of a long spatial signal, which ben-

efits from appropriate data reduction to be practically analyzed. A block RMS approach is proposed by [12], meant to condense a longer corrugation signal into fewer values and score each block on a coarse scale. This research proposes an extension to the block RMS methodology of [12], whereby a section of rail is summarized using a single representative block instead of applying a several scores to the segment.

Audible noise is a known outcome of corrugation, where the dynamics of the wheel-rail interface generate louder sounds when corrugation is more severe. Sound measurement is a well-understood field, but not necessarily focused on detecting corrugation. This work looks to expand on corrugation detection through wayside noise emission by focusing on the growth as a function of usage, as well as its response to grinding, while attempting to isolate the frequencies of noise explicitly due to corrugation.

Lastly, even though corrugation is a phenomenon on the scale of microns, it has major impacts on the ride quality due to its known propensity to increase vibration experienced by the vehicle. This vibration can be measured and quantified using on-board accelerometer sensors. However, this data is often noisy and cannot be analyzed as a time series. The wavelet transform is identified as a possible signal processing technique meant to improve the corrugation detection by proxy from vibration data.

1.5 Thesis Outline

Chapter 2 summarizes the present-day goals and state-of-the-art assessment techniques for corrugation, noise, and vibration in industry, as well as the raw data format, pre-processing, and preparation of the data used in this analysis. Chapter 3 explains the work and outcomes focused on corrugation analysis at two properties. Chapter 4 investigates the wayside noise-related work at these same two properties. Chapter 5 looks at on-board vibration data, and provides a high level overview of the wavelet transform and its application to this vibration data. Chapter 6 concludes the thesis and offers recommendations for future work.

Chapter 2

Data Background and Description

Data needs for this research were serviced by ARM and NRC from two Canadian passenger rail sites. Some data was a result of a structured sampling routine for the sake of research, while some was only taken as part of a standard maintenance schedule. This chapter outlines the three ways this project's data was collected, the format in which it was received and processed for analysis, and some ways industry is using these forms of data for system maintenance.

Data collection was conducted by ARM in conjunction with the transit agencies from which the data was sampled. Because data collection was not in the scope of this work, it is assumed that the instrumentation was properly calibrated, either by the operators or at factory, and properly installed for data collection prior to this research.

2.1 Corrugation Data

In order to understand roughness on the rail surface, instrumentation is required to sample it. Corrugation data is sampled using the Corrugation Analysis Trolley (CAT), a tool proprietary to RailMeasurement, a division under RailTechnology GmbH. This is a manually operated trolley that is walked along the track by an operator to sample the vertical displacement of the rail at uniform intervals (periods) of either 1 or 2 mm, and collected by a user walking at a pace of 3-4 km/h. The CAT accomplishes this through a contact measurement system using an inertial sensor that follows the vertical displacement of the rail surface via physical contact. The vertical displacement of the sensor is then recorded at each sampling interval [13]. This sensor can be adjusted to measure the vertical displacement of a particular lateral position on the rail allowing inspection of the running band (shown in Figure 2.1) across different track geometries.

There are three distinct models of CAT device manufactured by RailMeasurement, which all accomplish the same task of corrugation assessment: the original CAT, the BiCAT, and the CAT-V. The original CAT only has one sensor module, meaning it can only measure vertical rail displacement on one rail at a time. The BiCAT



Figure 2.1: Rail illustrating the running bands of interest at a particular property. The markings on the rail indicate that this is the right rail (RR), and the gauge side (G) of the rail is at the bottom of the image. The two ticks on the rail face denote 25 and 40 mm from gauge side, and the majority of the contact between the wheel and rail falls in the band between them, evidenced by the surface wearing.

adds a second sensor module to the opposite side of the device allowing simultaneous left and right rail vertical displacement measurements. The CAT-V has only one sensor module like the original CAT, but also features a camera system for taking pictures of the rail face at either 0.2, 0.33, or 0.5 m intervals. Figure 2.2 shows the original CAT in use.



Figure 2.2: Operator logging vertical displacement on the far rail face using the original CAT. The blue box on which the laptop is resting contains the sensor hardware.

The CAT saves data in a raw proprietary format that can only be viewed through the proprietary CAT software. This is a tool that visualizes and quantifies characteristics of the sampled corrugation data. The two most important visualizations are the vertical displacement (raw or filtered), and the $\frac{1}{3}$ octave spectrum information. An anonymized visualization of corrugation data viewed in the CAT software can be found in Figure 2.3.

Once sampled, corrugation is a spatial discrete signal. While we may be more familiar with temporally



Figure 2.3: Screenshot of RailMeasurement CAT software Version 14.26.21606 visualizing an unfiltered corrugation measurement and the calculated $1/3$ octave spectrum of the data.

sampled signals like audio signals or voltage signals sampled several times per second, corrugation is sampled as a function of position in mm intervals. Corrugation sampled at a period of 1 mm, can be interpreted as a digital signal with a spatial sampling frequency (f_s) of 1 sample per mm (1 mm^{-1}) or 1000 samples per m (1000 m^{-1}). By extension, a spatial Nyquist frequency defined as $f_N = f_s/2 = 500 \text{ m}^{-1}$ quantifies the maximum spatial frequency detectable for a given sampling frequency f_s . More intuitively, this is the minimum wavelength observable given the sampling period $\lambda_s = f_s^{-1}$. In the case of $f_s = 1 \text{ mm}^{-1}$, we can detect information of wavelengths longer than $\lambda_N = f_N^{-1} = 2 \text{ mm}$.

Corrugation measurements, as spatial signals, can be analyzed using the same techniques and principles typically applied to temporal signals. Corrugation data can inform track maintainers about the smoothness of the rail surface by spatial observation, but specific wavelength information can help diagnose particular corrugation types like those found in Table 1.1. The CAT's proprietary software provides the ability to filter raw corrugation data for a specified wavelength range using a single forward pass Butterworth filter. The work in this thesis was done using the default filter order of $N = 3$, but this is a customizable parameter in the software.

Another way that frequency information can be used to assess corrugation severity is through a spectrum's $1/3$ octave decomposition. The wavelength spectrum is portioned into several wavelength bands each with the bandwidth of $1/3$ of an octave, where an octave is the doubling of a frequency. This means that each octave band

is proportionally the same size relative to its centre frequency, but they are not the same bandwidth, due to the logarithmic nature of the definition. For a frequency band defined with the centre frequency f_c , and lower and upper corner frequencies f_l and f_u , the bandwidth of this $1/3$ octave band is defined by $f_u/f_l = 2^{1/3}$. Since they are equally proportioned about f_c , they can also be defined as

$$2^{1/6} f_l = f_c = f_u / 2^{1/6}. \quad (2.1)$$

The spectral content of these bands can help property maintainers diagnose and make grinding decisions according to relevant roughness standards, driven by different corrugation outputs. Relevant standards include ISO 3095:2005 [14], a thresholding motivated by external noise generated by the train passing over the rail, and EN 13231-3:2012 [10], a standard developed with longitudinal and lateral track profiling in mind. By comparing the $1/3$ octave spectrum to spectrum guidelines in these standards, property maintainers have a roughness assessment tool for making grinding decisions. Figure 2.4 showcases a sample $1/3$ octave spectrum relative to the previously mentioned standards.

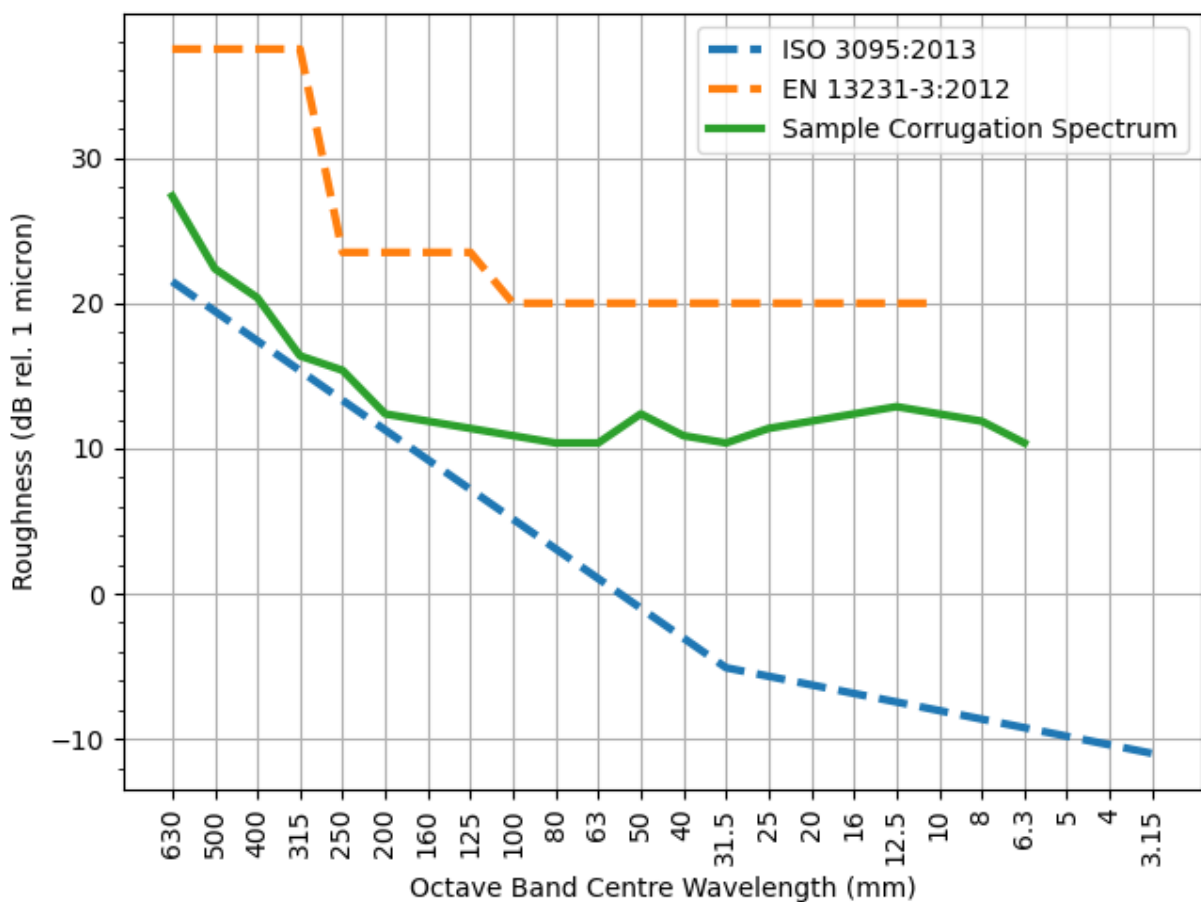


Figure 2.4: Sample corrugation $1/3$ octave spectrum compared with the ISO 3095:2005 and EN 13231-3:2012 standards. A valid interpretation of this chart is that this surface roughness passes the EN standard, but fails the ISO standard.

The CAT software allows for exporting of corrugation data into ASCII-encoded plain-text files. These ex-

ported files are formatted as a two column table where each row is a location measurement (distance from sample starting location) paired with a vertical displacement relative to the first sample. An exported file includes a header denoting how the data was filtered (if at all), and other sampling metadata like date of sample, time of sample, and CAT operator. Only the CAT software can read the raw corrugation data, meaning the data used in this research had to be exported into this plain-text tabulated format.

2.2 Wayside Noise Data

Different forms of audible noise have been measured to investigate the effects rail quality has on sound generated by the trains passing over corrugated rail. One way this noise can be captured is via wayside sound level meters, set up at fixed locations near the track, or mounting sound level meters on-board the train in order to monitor operational noise throughout the property.

A common summary of an audible sound is equivalent noise. This is a one-number summary meant to summarize some decibel level that, if generated for the same duration as a variable sound pressure sample, has the same total sound energy. The audible decibel scale is defined as [15]:

$$L_{\text{eq}} = 10\log_{10}\left(\frac{p^2}{p_{\text{ref}}^2}\right) = 20\log_{10}\left(\frac{p}{p_{\text{ref}}}\right), \quad (2.2)$$

where p denotes an instantaneous sound pressure, and L_{eq} is the equivalent sound level in dB relative to $p_{\text{ref}} = 20 \mu\text{Pa}$ for sound travelling through air per ISO 1683:2015 [16]. When p is a single value, it is trivial to calculate L_{eq} , but if p is a function of time (i.e. $p = p(t)$), Equation (2.2) does not produce a single value. L_{eq} for a varying noise sample is calculated by integrating over time as [15]:

$$L_{\text{eq}} = 10\log_{10}\left[\frac{1}{T}\int_T \frac{p^2(t)}{p_{\text{ref}}^2} dt\right] \quad (2.3)$$

$$\begin{aligned} &= 20\log_{10}\left[\frac{\sqrt{\frac{1}{T}\int_T p^2(t)dt}}{p_{\text{ref}}}\right] \\ &= 20\log_{10}\left[\frac{p_{\text{RMS}}}{p_{\text{ref}}}\right]. \end{aligned} \quad (2.4)$$

Here, T denotes the period over which the noise is sampled, and $p(t)$ is the fluctuating noise signal. It can be shown that Equation (2.2) is the degenerative case of Equation (2.3) when $p(t)$ is a constant sound pressure for any time duration T . The numerator, $\sqrt{\frac{1}{T}\int_T p^2(t)dt}$, is the root mean square calculation of $p(t)$, and is denoted as p_{RMS} as a simplification. An example of this calculation can be found in Figure 2.5.

The equivalent noise value can also be calculated using the spectral content of a signal. Given the frequency content of a signal $p(t)$, $\hat{p}(f)$, the equivalent sound level is defined as:

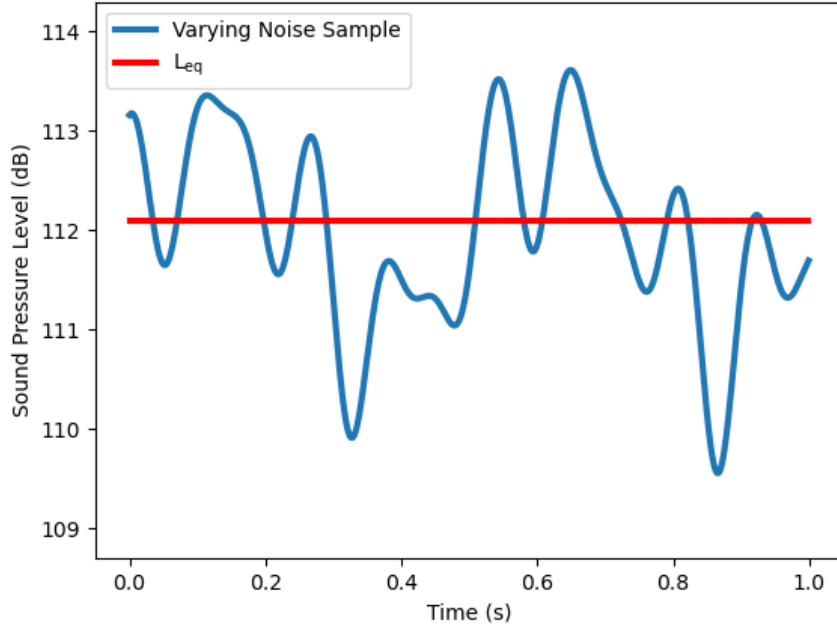


Figure 2.5: L_{eq} noise for a given fluctuating noise sample. Note that if the dB signal were just integrated and averaged, the resulting calculation would slightly underrepresent the true L_{eq} .

$$L_{eq} = 10\log_{10}\left(\int_F 10^{\frac{\hat{p}(f)}{10}} df\right) [15], \quad (2.5)$$

where F is the range of frequencies. This is useful for calculating the $1/3$ octave equivalent spectrum of a wayside noise sample, a spectrum of decibel levels summarizing equivalent sound energy in a $1/3$ octave frequency band, as F does not need to be the entire range of frequencies. For example, the equivalent sound pressure level in the 100 Hz $1/3$ octave spectrum uses Equation (2.5) where $F = [\sim 89, \sim 112]$ Hz. The bounds of F are calculated using Equation (2.1), in the same manner that they were for corrugation.

Lastly, different phenomena will cause different frequencies of noise, some of which will be audible to humans. In situations where sound levels as experienced by humans is a focus, it is appropriate to weight certain noise frequencies more heavily than others due to the fact that the human hearing profile is not uniform across all frequencies.

A standard way to weight sound for human hearing is called A-weighting. A-weighting's definition can be found in IEC 61672-2:2013 [17], a standard outlining sensitivity testing for sound level meters to ensure they conform to other sections of the IEC 61672 standard. A-weighting is designed to mimic human hearing sensitivities in order to better quantify loudness as perceived by the public by attenuating frequencies outside of the range of 1000-6300 Hz. The gain curve for A-weighting is shown in Figure 2.6.

A-weighting a noise sample to assess its audible impact is performed by first applying the gain curve to a noise spectrum's sample, and then performing L_{eq} calculations or similar calculations as necessary. The units of

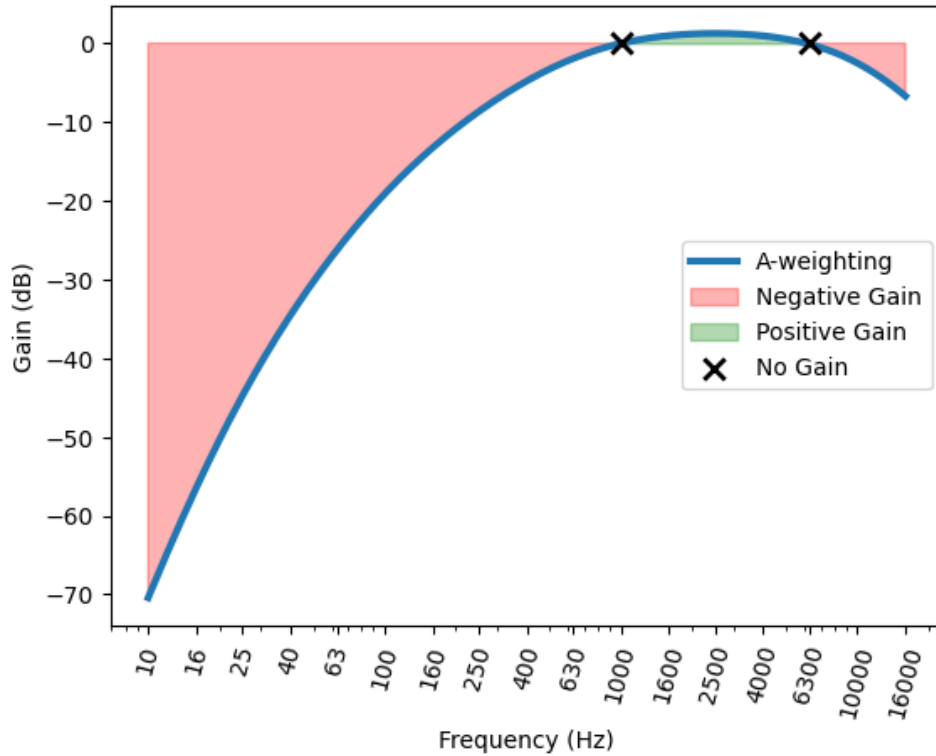


Figure 2.6: Regions of the A-weighting curve where sound is attenuated or amplified. The intercepts with 0 dB (no gain), denoted with a black \times , are located at 1000 Hz and \sim 6300 Hz.

measure for the sample reflect this through a change in notation from dB to dBA, and the equivalent noise measurement L_{eq} becomes LA_{eq} to inform those using the data. These conventions extend to other noise weighting standards like C-weighting.

Wayside noise data for this project was sampled with a Larson Davis Sound Level Meter (SLM), which measures the sound pressure experienced as a function of time. SLMs use a microphone, in which a diaphragm will respond to changes in sound pressure around it. This diaphragm is a transducer which will affect a measured electrical signal that can be converted into sound pressure experienced by the device.

The signal is saved in raw proprietary format to be viewed using Larson Davis G4 Utility software, which can then be exported into an Excel spreadsheet with information and statistics summarizing the audible noise signal in time. The Excel files were provided to us organized according to location and date of sample. These were scraped for the useful information listed above, specifically LA_{eq} , and A-weighted $\frac{1}{3}$ octave equivalent noise spectra.

2.3 Vibration Data

The vibration data in this project is fundamentally different from the corrugation and noise data. Where corrugation and noise are spatially restricted in this work, vibration data was collected throughout a property through the use of mounted (i.e. on-board) axle-box accelerometers (ABAs). Additionally, whereas corrugation and noise



Figure 2.7: ABA mounted on a vehicle, with the IMU's frame of reference.

were intentionally sampled to be paired with each other, the vibration data was sampled in isolation.

An ABA is a device which uses an inertial measurement unit (IMU) to measure proper acceleration, i.e. acceleration excluding gravity [18]. The ABA reports this acceleration in three orthogonal dimensions x , y , and z relative to the device using an accelerometer built into the IMU. IMUs generally also house a gyroscope which measures the rotation of the device, and, by extension, the rotation of the acceleration axes. In this work, the device is oriented such that x refers to longitudinal acceleration (acceleration in the direction of travel), y refers to lateral acceleration (left-to-right acceleration), and z is vertical acceleration (up-and-down acceleration), as shown in Figure 2.7.

An ABA at rest on a table experiences the force of gravity pointing down towards the floor, but is only reporting the normal force of the table counteracting gravity, i.e. proper acceleration. Assuming the gyroscope indicates that the device is oriented properly (i.e. with the z axis pointing up), the ABA would only report the acceleration due to the normal force of 9.81 m/s^2 in the z axis up from the table. If it were slid across the table in the positive direction of the x axis, it would report that initial positive x acceleration due to an applied force (a_F), and then the counteracting negative x acceleration due to kinetic friction (a_k) until it reached a stop, while continuing to report the normal acceleration from the table. The free body diagrams corresponding to these three stages of a hypothetical accelerometer sample can be found in Figure 2.8.

ABA data provides a potential opportunity to passively monitor system operation. ABAs are a non-invasive measurement tool that can be mounted and used during revenue service. If they are capable of monitoring rail condition accurately, they provide the potential to assess an entire property's rail quality in a single day, whereas maintenance workers are limited to the length of track they can physically inspect during a night with a device such as a CAT. Because rail wear is a slowly developing phenomenon arising over the course of months,




Figure 2.8: Free body diagrams of acceleration captured by ABA at rest (a), when slid across the table (b), and when friction slows it down (c).

ABA data can give up-to-date information on the state of operation with minimal time-consuming manual data collection.

Vibration has been an important aspect of rail maintenance for a long time, and several efforts have attempted to understand the effects of corrugation on vibration. Modelling has been used to understand the contact effect of the wheel and rail to simulate short-pitch rail corrugation more accurately [19]. Additionally, work has been done to understand the effect of various rail properties, constructions, and materials on vibration and corrugation [20]. There are several standards in place to assess vibration and its effects on both the rail car and the surrounding area. These include standardizing the vibrations experienced by the foundations of neighbouring buildings [21], the vibration of tunnels due to underground rail operation [22], the vibration and shock testing requirements for mechanical, pneumatic, electrical and electronic equipment/components on vehicles [23], and the vibration experienced by passengers and crew when experiencing whole-body vibration in the rail car [24].

ABAs have been used to investigate on-board acceleration previously. Some attempts experienced problems with suspension impacting the measurements [25], as well as issues with identifying track irregularities of long wavelength [26]. However, successful investigations have been completed using Kalman filters and mixed filtering to successfully identify various irregularities resulting in excessive vibration and strain [27].

One drawback with ABAs used in this study is a lack of position and velocity data sampling. Part of this comes from a desire for the system to be stand-alone and non-invasive, as well as the non-uniformity of vehicle interfacing across different makes of vehicle at a property. One possible solution that maintains the non-invasiveness of the unit would be the inclusion of a GNSS (global navigation satellite system) module in the ABA, however, GNSS signals suffer from dilution of precision, where error propagates due to error in measurement. [28] Additionally, GNSS suffers from multi-stage communication potentially affecting the reliability and consistency of the data. Clock synchronization between satellites and the unit, as well as propagation issues through the ionosphere and troposphere present problems that need to be addressed in the data handling process [29]. In the case of passenger rail applications, there is the added issue of urban noise interfering with the GNSS communications, either due to density of electromagnetic noise or physical interference between the satellite and receiver in the forms of tunnels, and tall structures reflecting or diffracting the signals [30]. Due to these problems, no position or velocity data is being explicitly collected by the unit responsible for collecting the data used herein.

IMUs house an accelerometer to measure proper three-axis acceleration as well as a gyroscope to measure rotation of the IMU. The ABAs used in this project do not use the gyroscope in the IMU and instead only sample the accelerometer. As a result, the rotation of the axes cannot be measured, and cannot be corrected. If the device is mounted at an angle, then the measured normal force due to gravity will have components in not only the z axis, and this cannot be unbiased confidently. Additionally, if the device rotates during operation due to curves or slopes, the normal force components will distort the measurements irreversibly without the gyroscope information.

The ABA units used in this research were made in-house by ARM using off-the-shelf (OTS) components.

Accelerometer Specifications		
Parameter	Performance/Range	Units
Full-Scale Range	± 16	g
ADC Word Length (two's complement)	16	bits
Sensitivity	2048	bits/g
Cross-axis Sensitivity	± 2	%
Power Spectral Density (low noise mode)	0.3	$\text{g}/\sqrt{\text{Hz}}$
Sampling rate	3225	Hz
Gyroscope Specifications		
Parameter	Performance/Range	Units
Full-Scale Range	± 2000	$^{\circ}/\text{s}$
ADC Word Length (two's complement)	16	bits
Sensitivity	16.4	$\text{bits}/(^{\circ}/\text{s})$
Cross-axis Sensitivity	± 2	%
Rate Noise Spectral Density	0.01	$^{\circ}/\text{s}/\sqrt{\text{Hz}}$
Sampling rate	3225	Hz

Table 2.1: ABA IMU Specifications.

The IMU inside the unit is anonymized in this document, but the specifications can be found in Table 2.1. Note that this IMU has independent variable full scale ranges for the accelerometer and gyroscope. The maximum full scale range was used for the accelerometer, but as mentioned the gyroscope was not enabled. The parameters listed for the gyroscope are given according to the maximum full scale range.

One potential shortcoming with OTS IMUs is their application design, which tends to be focused on human-scale acceleration measurements. For instance, wearable sensors, and motion-based controllers fit the design parameters of most OTS IMUs, while higher grade IMUs produced for aerospace applications will come with higher capabilities at a higher price point, and are likely built to order.

In the case of the ABAs to measure on-board acceleration information, it is possible that the application environment is not ideal for an OTS component, due to the violent dynamics of the vehicle. Some general issues facing IMUs are repeatability (consistent output for consistent input), stability (constant output for constant input), and drift (change of output unrelated to change in input) [31]. These can change systemically due to power supply problems, or randomly as stochastic processes. It is possible that, in the case of trains, these random processes can be exacerbated.

Each ubattery-powered unit is mounted to the train at the start of the day, and is recovered once the vehicle has returned to the yard regardless of when the device has run out of power. The data is passively collected throughout the day at a sample rate f_s of 3225 Hz, and is logged to an SD card timestamped according to time since power on. Each acceleration measurement is 16 bits signed, meaning without timestamps, the data is logged at a rate of 19.35 KB/s. As a result, the size of the logged data grows rapidly.

It is both infeasible and impractical to attempt to read the entire accelerometer sample into memory at once.

The solution is to post-process the raw data into more manageable chunks. The raw log file is decoded second by second (i.e. in chunks of 3225 timestamps), and the data is discarded if it is determined from the data that the train is not moving. This repeats until the train is determined to be in motion, at which point the data starts being tabulated and saved in memory. The acceleration data is converted from a 16 bit signed word to an acceleration measurement in g's, where one g is equal to 9.81 m/s^2 . Throughout daily revenue service, the vehicle will stop at stations to pick up and drop off passengers. Once the vehicle is determined to have stopped, the table is written to an ASCII-encoded plain-text file as a comma-separated value (CSV) table. Then, the process repeats where data is discarded second-by-second until the train moves again, and a new plain-text file is started. Our analysis is only done on data after it has been segmented into multiple plain-text files.

2.4 Chapter Summary

Corrugation as a maintenance problem is at the forefront of the rail industry, and has been for several decades. Its pervasiveness as well as inconvenience to diagnose have presented various maintenance problems that demand attention.

Corrugation is not a consistent problem across properties since wheel and rail design is not consistent across properties, meaning each properties will likely have its own unique corrugation problems. Similarly, its propensity to precede RCF brings other maintenance issues.

Monitoring rail condition like corrugation severity via a proxy like wayside noise or vibration is desirable for property maintainers due to the ease of data collection compared with direct corrugation measurement. Measuring corrugation directly is time-intensive, and inefficient in terms of length of track investigated.

Meanwhile, investigation of noise or vibration is non-invasive so it can be done during normal day-time operations. Workers can sample noise at several locations throughout the system during the day, travelling by car or the trains themselves and cover more ground than by foot. In addition to this, if noise or vibration can be measured passively on-board the vehicle, it allows the entire state of the system to be measured at once through normal transit service.

Chapter 3

Corrugation Analysis

This chapter describes the process of exploring corrugation growth and measurement. An initial study took place at Property A, a Canadian rail property cooperating with ARM and NRC to look into corrugation growth at a particular section of track with tangent followed by curve geometry. Property B, a second Canadian rail property, provided pre- and post-grind corrugation data from several curves. The data at this new location was used to assess the robustness and generality of the methodology applied at Property A. Sampling corrugation from the rail surface generates a discretely sampled spatial signal, which lends itself to various signal processing techniques like filtering. An interdisciplinary cooperation to bring signal processing expertise to the rail transit sector is a valuable partnership, as those in the rail industry do not have the opportunity to learn the signal processing knowledge those in the computer engineering space may take for granted. Thus, this work focuses on the treatment of corrugation as a spatial signal.

3.1 RCI_x for Monitoring Corrugation

The rail industry has developed and evaluated rail condition indices to characterize rail quality and assess grinding effectiveness. Specific characteristics include, but are not limited to, profile quality, surface damage, and rail corrugation. Magel and Oldknow [12] proposed a Rail Corrugation Index, denoted as RCI_{2018} , to measure the presence of corrugation on the rail surface. The work conducted in this thesis extends that work.

The presentation in [12] lays out the foundation of RCI_{2018} , a metric meant to assess the the severity of corrugation across many 5 m segments. Corrugation does not manifest across all wavelengths, and some spectral content of a corrugation signal will not be reflective of corrugation. For example, rail ties, supporting structures built to stabilize the rails and keep them in gauge, are spaced along the rail every 19-24 inches (480-600 mm) depending on material. As such, spectral content with wavelengths in that range is more likely to be due to rail ties and not corrugation. With intention to focus solely on corrugation, the methodology in [12] filters raw corrugation data for information in the 30-100 mm range.

For RCI_{2018} , a segment of rail is divided into contiguous 5 m blocks which are individually scored according to the corrugation in the block. In the case of an oscillating spatial signal, the arithmetic mean will not capture the roughness of a block since positive and negative vertical displacements in a block will cancel each other out in the calculation. Instead, root mean square (RMS) or the quadratic mean is used. Once the RMS of a block (RMS_{block}) is calculated, it is compared to a tolerance level ($TOL=4 \mu m$) using the following equation:

$$RCI_{2018} = \frac{RMS_{block} - TOL}{6 \times TOL}, \quad (3.1)$$

for $4 \mu m < RMS_{block} < 28 \mu m$. The results are then binned and a colour scheme proposed by [12] for visualizing these bins, as presented in Table 3.1.

$RMS_{block} (\mu m)$	RCI_{2018}	Colour
< 4	100.0	Blue
< 8	83.3	Green
< 12	66.7	Yellow
< 16	50.0	Orange
< 28	0.0	Red

Table 3.1: Proposed scale and colour scheme for RCI_{2018} [12].

While this is effective in providing some coarse grading of rail surface roughness, it has a major shortcoming in that it is inefficient and impractical to grind several short sections of rail. As a result, it is unclear what actions to take when, across some segment of rail, some blocks are red (indicating higher corrugation) amid a large number of blue blocks (indicating smooth rail). As a result, there is a need to summarize corrugation over longer stretches of rail to make actionable decisions practical.

3.1.1 Property A

Property A featured a robust corrugation sampling regimen. Data was sampled along a 0.2 mile (roughly 300 m) segment, at 5 different sensor positions along the rail - 25 mm, 30 mm, 35 mm, 40 mm, and 45 mm from the gauge side of the rail. This rail segment featured two different track geometries. It started in a tangent section of rail, and extended to a curve portion, with transition spirals before and after the curve. Corrugation was sampled monthly at this Property A track segment to investigate the corrugation growth as a function of accumulated tonnage experienced by the track (measured as million gross tonnes, or MGT).

To overcome the noted limitation of RCI_{2018} not providing actionable outcomes for longer track segments, the decision was made to select a single block, based on a percentile of the blocks in the segment, as a representation of the corrugation over that segment of corrugation. This replaces the RCI_{2018} grading scale in Table 3.1 with a


single number grading an entire segment of track. This new metric, RCI_x , improves upon RCI_{2018} , due to its ability to score a practical length of track for maintenance. Here, x denotes the block percentile selected, so the 65th percentile would be denoted as RCI_{65} . Formally, the algorithm for calculating RCI_x is as follows:

1. Filter a corrugation sample for the 30-100 mm wavelength. This uses the default filter in the CAT proprietary software which is a Butterworth filter of order $N=3$.
2. Segment the filtered corrugation sample into 5 m blocks.
3. For each block, calculate the RMS vertical displacement.
4. Select the block corresponding to the x^{th} percentile of the sequence of block RMS values, rounding up. That is, the smallest block value that is greater than or equal to $x\%$ of the block RMS values in the sample.

The data was divided into tangent and curve sections to identify the effect of track geometry on RCI_x growth for various percentiles. The tangent section was approximately 100 m, while the curve section was approximately 200 m as the spiral segments were concatenated with the curve segment. The 65th, 85th, and 95th percentiles of the RCI_x distribution on a segment were selected to investigate the percentile that most appropriately represented the distribution. The 95th percentile is more sensitive to the worst corrugation in a segment, the 85th percentile tends toward the worst corrugation while being a more outlier-resistant selection, and the 65th percentile is more of a measure of central tendency closer to the median corrugation block. RCI_{65} , RCI_{85} , and RCI_{95} were plotted as a function of MGT across two grind intervals to reveal whether the RCI_x metrics appropriately captured corrugation growth and the effectiveness of grinding.

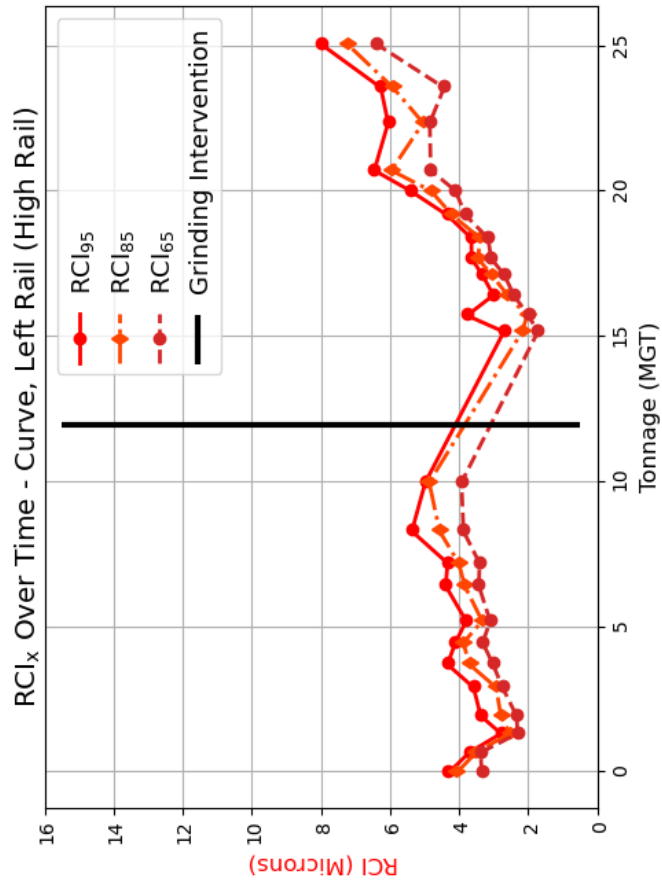
Investigation into these plots was separated according to geometry with each rail/geometry pairing visualized in Figure 3.1. In the case of tangent rail, the dynamics are theoretically even along the right and left rails. However, for curve rail, literature suggests that the low rail of the turn may experience more of the forces from the train than the high rail depending on the curve radius [32]. The data was captured at a right hand turn, meaning the right rail is the low rail. As visualized in the tangent data, there is a strong upward slope until grinding intervention leading to a sharp decline of each of the RCI_x metrics. In the curve data, there is a similar shape and response to grinding in the low rail, however the high rail sees a much less pronounced growth and grind response. This is explainable in that the low rail is likely to bear the brunt of the weight due to gravity when going around a turn for smaller radii turns. These trends were repeated during the second grind interval, suggesting that there was a systemic, repeatable relationship between MGT experienced post-grind, and RCI_x in the 65th, 85th, and 95th percentile values.

Research was conducted into the efficacy of the filter band selected in [12], as well as the effect of track geometry on RCI_x by Dr. Ahmed Lasisi and Sanjay Abraham, two members of our research team. The research is relevant to the work in this thesis, and so is presented here. The outcomes of that research are as follows:




(a) Left rail corrugation growth at tangent track location.

(b) Right rail corrugation growth at tangent track location.



(c) Left rail corrugation growth at curve track location.



(d) Right rail corrugation growth at curve track location.

Figure 3.1: Corrugation growth captured by RCI_{65} , RCI_{85} , and RCI_{95} separated by rail and track geometry at Property A.

- Investigation into the impact of wavelength band at this property suggested that 30-100 mm, while not only a recognized wavelength band in industry, demonstrated a stronger response in RCI_{2018} as a function of wear. While there may be desire to incorporate other wavelength bands in an “average-based” condition assessment, industry-focused condition reporting seems to have more use for the 30-100 mm band.
- Segment length plays a major role in the RCI_x calculation, as it will not only affect the granularity of the reporting, but will also make RCI_x more susceptible to outliers. For a fixed 5 m block length, shorter segment lengths will result in fewer blocks. The 95th percentile block will trend closer to the 100th percentile block as fewer blocks are present in a given segment.

These findings showed that RCI_x as a metric was sensitive to corrugation growth, and responded clearly to grinding intervention. While benchmarks for requiring intervention were not theorized, it was evident that RCI_x reflected rail surface roughness effectively and could be used to establish decision-making thresholds. Repetition at a new property was of interest to determine if any features unique to Property A were influencing outcomes and conclusions.

3.1.2 Property B

The research completed at Property A successfully identified corrugation growth, and a clear response to grinding in the 30-100 mm wavelength band. This methodology laid the groundwork for partial repetition of the study at a second Canadian rail property, Property B.

It is important to note that what follows is not a true repetition of the study due to a fundamentally different sampling regimen at Property B. Where Property A had two locations of unique track geometries (tangent and curve) with consistent corrugation sampling approximately monthly, Property B had 20 sampling locations, all of which were curves, with only pre- and post-grind corrugation data on the left and right rails. Of these 40 unique inspections (20 curves, 2 rails per curve), three sites were missing pre-grind data on at least one rail. These were removed, leaving 17 curves, or 34 unique rails. Thus, the validation at this property has a scope that is limited in terms of time and track geometry, but still allows verification of the grinding response of the metric.

The response to grinding was defined as $\Delta RCI_x = RCI_x^{\text{Pre-grind}} - RCI_x^{\text{Post-grind}}$, where x denotes the common percentile being compared. Rail maintenance is done with the intention of reducing corrugation roughness post-grind, meaning that $\Delta RCI_x > 0 \mu\text{m}$ indicates an improvement in rail condition.

ΔRCI_x was calculated for the 95th, 85th, and 65th percentiles on each of the 34 unique rails. Of these corrugation measurements, 24 (70.58%) of the ΔRCI_{95} values decreased with grinding as expected. Similarly, 24 (70.58%) and 22 (64.71%) of the ΔRCI_{85} and ΔRCI_{65} decreased with grinding respectively. The remaining samples for the three ΔRCI_x metrics suggested that corrugation *worsened* post-grind, contrary to field observations indicating improvement in rail condition. These findings suggested that RCI_x as defined and evaluated to be successful for Property A, was not consistently delivering the same kind of meaningful characterization at Property B.

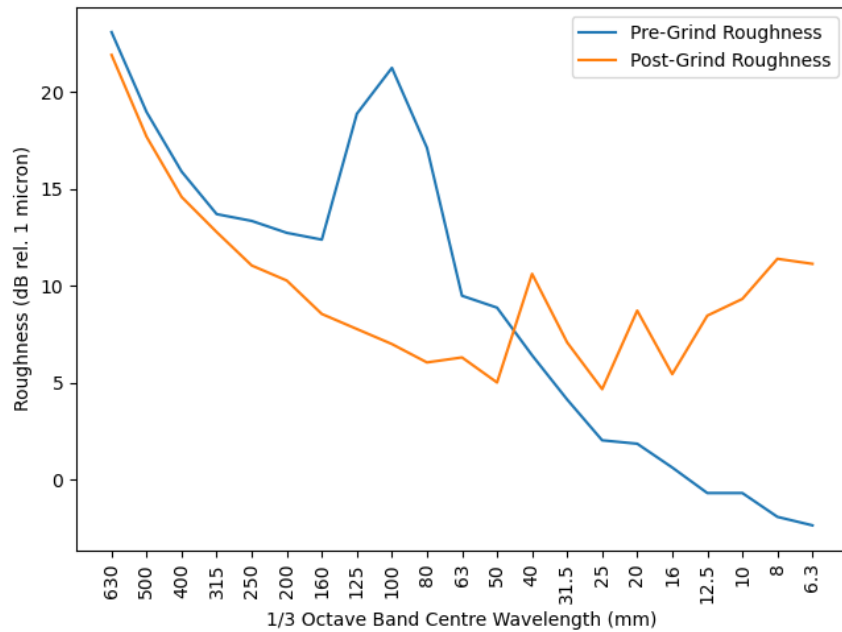


Figure 3.2: Pre- and post-grind $\frac{1}{3}$ octave spectra from one location in Property B. The large peak centred at approximately 100 mm appears to have been removed from grinding, while low wavelength roughness was leftover as the grind signature.

In an attempt to uncover potential issues with the metric at Property B, an investigation of the $\frac{1}{3}$ octave spectra of the pre- and post-grind data was performed to determine whether or not the 30-100 mm band was showing improvement after grinding. An example of this data is shown in Figure 3.2. Inspection of this figure and others like it raised two observations. First, there is an insertion of low-frequency roughness in the post-grind data at < 50 mm wavelengths, which may be artificially inflating the post-grind RCI_x measurement. When rail is ground or milled, the machinery leaves behind a characteristic roughness known as the grind signature; this eventually wears away into smooth rail through normal day-to-day operation. It is possible that RCI_x is responding to this leftover grind signature, thus obfuscating the removal of rail corrugation by grinding. The case can be made that while the grind signature left on the rail is a form of rail roughness, it is not rail roughness related to corrugation which is what RCI_x is primarily trying to quantify. Assessing the roughness of the grind signature was not considered during the work at Property A, and so it was deemed to be out of scope at Property B as well. Second, the major peaks in much of the pre-grind data removed in post-grind data was centred at 80 mm or higher, which suggests that the 30-100 mm wavelength band was not adequately capturing the target corrugation at Property B in the same way that it was at Property A.

These two observations motivated revisions to RCI_x to make it more flexible to accommodate the particular maintenance needs of different properties. While the original 30-100 mm filtration scheme was appropriate for Property A, it seems that Property B has slightly different corrugation maintenance characteristics. The solution was to add a parameterization to RCI_x , denoted as $RCI_x(\lambda)$, to allow property maintainers to hone the metric for their specific needs.

3.2 Wavelength-Adapted RCI_x ($RCI_x(\lambda)$)

3.2.1 Preliminary Investigation

As a first step to understanding RCI_x 's change in performance between Properties A and B, an investigation of the performance of various wavelength bands was undertaken, to better understand how the precision of cutoff frequency selection affects performance. The CAT tool's filtration scheme for a given wavelength band is a one-pass Butterworth filter with order $N = 3$, which was used as the filter design in the testing and the final $RCI_x(\lambda)$ implementation. Four wavelength bands were tested: 30-100 mm, 50-100 mm, 50-200 mm, and 100-300 mm. Table 3.2 presents the rationale for testing these wavelength bands, as well as their ΔRCI_{65} , ΔRCI_{85} , and ΔRCI_{95} improvement performance. The performance metric, denoted as ΔRCI_x Improvement, quantifies the percentage of pre- and post-grind samples a given filter band and percentile value that showed $\Delta RCI_x > 0$.

Wavelength	Testing Rationale	ΔRCI_{65} Improvement	ΔRCI_{85} Improvement	ΔRCI_{95} Improvement
30-100 mm	Baseline methodology from Property A	24/34 (70.6%)	24/34 (70.6%)	22/34 (64.7%)
50-100 mm	Baseline with removal of grind signature	24/34 (70.6%)	23/34 (67.7%)	23/34 (67.7%)
50-200 mm	Short/medium pitch with removal of grind signature	29/34 (85.3%)	29/34 (85.3%)	29/34 (85.3%)
100-300 mm	Exclusively long pitch	33/34 (97.1%)	34/34 (100%)	33/34 (97.1%)

Table 3.2: Improvement of RCI_x for different wavelength bands, and their rationale for testing.

The performance summaries in Table 3.2 answer the questions about the grind signature and filter band accuracy. Firstly, there is not a significant difference between the performance of the baseline methodology from Property A (30-100 mm), and the filter band designed to exclude the grind signature (50-100 mm). This suggests that the inclusion or exclusion of the grind signature has insignificant bearing on assessing the quality of corrugation removal post-grind. Secondly, shifting the wavelength band up beyond 100 mm improved the performance of ΔRCI_x in all percentiles. This indicates that parameterizing the filter passband more appropriately for a property will improve ΔRCI_x 's ability to detect the removal of corrugation. Looking at Figure 3.2, the large peak centred at approximately 100 mm suggests that properly enveloping a peak like that is necessary for diagnosing the presence and subsequent removal of corrugation.

3.2.2 Quantifying Corrugation Improvement using $RCI_x(\lambda)$

To contextualize improvements captured by ΔRCI_x , a relative measurement defined as $\% \Delta RCI_x = \Delta RCI_x / RCI_x^{\text{Pre-grind}} \times 100\%$ represents the corrugation removed by grinding as a percentage of the pre-grind corrugation level. In-




Figure 3.3: Box plot distribution of $\Delta RCI_x(\lambda)$ for the different tested wavelengths and percentiles ($n = 34$).

stances where ΔRCI_x is negative, $\% \Delta RCI_x$ indicates how much roughness has been added relative to the pre-grind amount.

As previously stated, we introduce the notation $RCI_x(\lambda)$ to highlight where the RCI_x methodology was used with a variable wavelength range, and this notation extends to $\Delta RCI_x(\lambda)$ and $\% \Delta RCI_x(\lambda)$.

To investigate the distribution of the datapoints that make up Table 3.2, the $\Delta RCI_x(\lambda)$ data is presented in Figure 3.3 as a series of box plots for each of the wavelength ranges. The red line indicates where $\Delta RCI_x = 0 \mu\text{m}$, and the orange line indicates the median of each distribution. The upper and lower edges of the boxes denote the first and third quartiles of the data respectively ($Q1$ and $Q3$), with the difference between them indicating the interquartile range ($IQR = Q3 - Q1$) of that distribution. The whiskers extend down to the first datapoint greater than $Q1 - IQR$, and extend up to the first datapoint less than $Q3 + IQR$. Lastly, the points denoted by circles indicate mathematical outliers by the convention of box plots.

Looking at Figure 3.3, the $\lambda=30-100 \text{ mm}$ and $\lambda=50-100 \text{ mm}$ distributions have whiskers extending further below the $0 \mu\text{m}$ line. Meanwhile, the $\lambda=50-200 \text{ mm}$ and $\lambda=100-300 \text{ mm}$ distributions extend less in that direction. For $n = 34$, each of the distributions have $n/4 = 8.5$ values in each of the quartile groups. Those are less than $Q1$, between $Q1$ and the median, between the median and $Q3$, and greater than $Q3$. The fact that the lower wavelength bands have such long whiskers and several outliers on the lower quartile section indicates that the spread is higher. Meanwhile, the higher wavelength bands have much shorter lower whiskers, indicating a clustering of the lowest quartile. This signifies that the higher wavelength range not only is more likely to have a

positive ΔRCI_x , but that its negative ΔRCI_x values are more likely to be less negative. All of this reinforces the idea that when RCI_x is applied using a more appropriate wavelength band for the data, the change in spectral content is echoed in the block RMS methodology, suggesting its robustness and flexibility.

Looking at the $\% \Delta RCI_x(\lambda)$ box plots of Figure 3.4, similar trends are evident. The lower wavelength bands have much wider spreads in the data on the low end of the distributions, and in some instances the metric demonstrates perceived post-grind corrugation growth up to 5 times the $RCI_x^{\text{Pre-grind}}$ value. The fact that the distributions of the low wavelength $\Delta RCI_x(\lambda)$ calculations are fairly symmetric about the median suggests that those particular wavelength bands are not strongly related to the corrugation removed from grinding. Meanwhile, the higher wavelength band $\Delta RCI_x(\lambda)$ calculations are positively skewed about the median, suggesting that there is a higher chance for those wavelength bands to properly detect corrugation removal. This trend is not only present in the lowest quartile section of data.

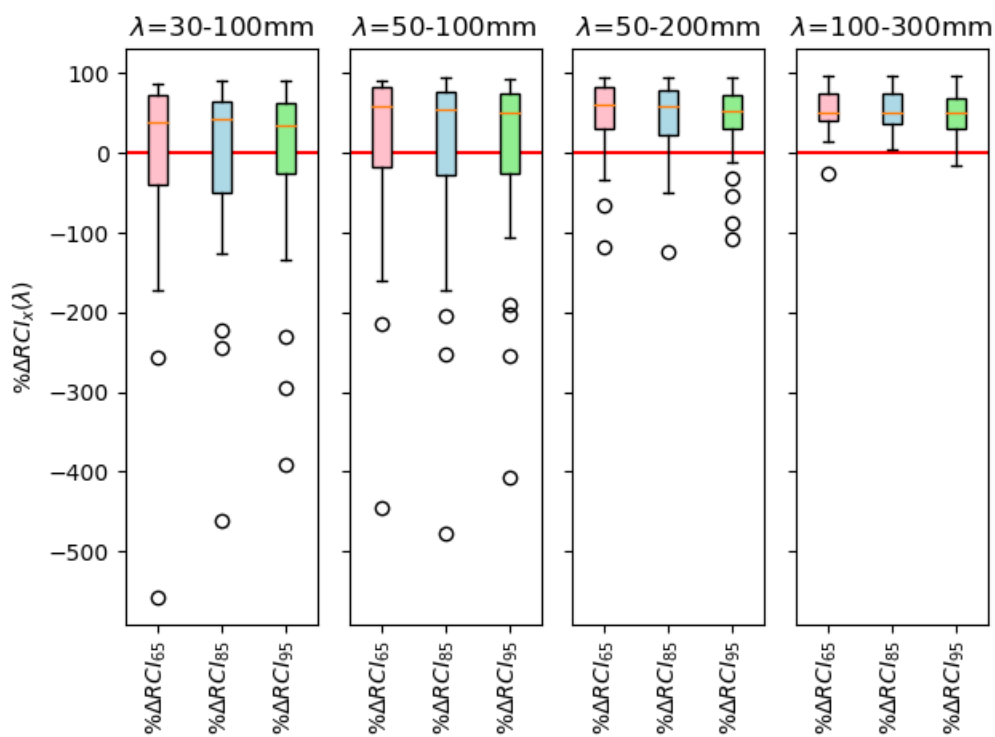


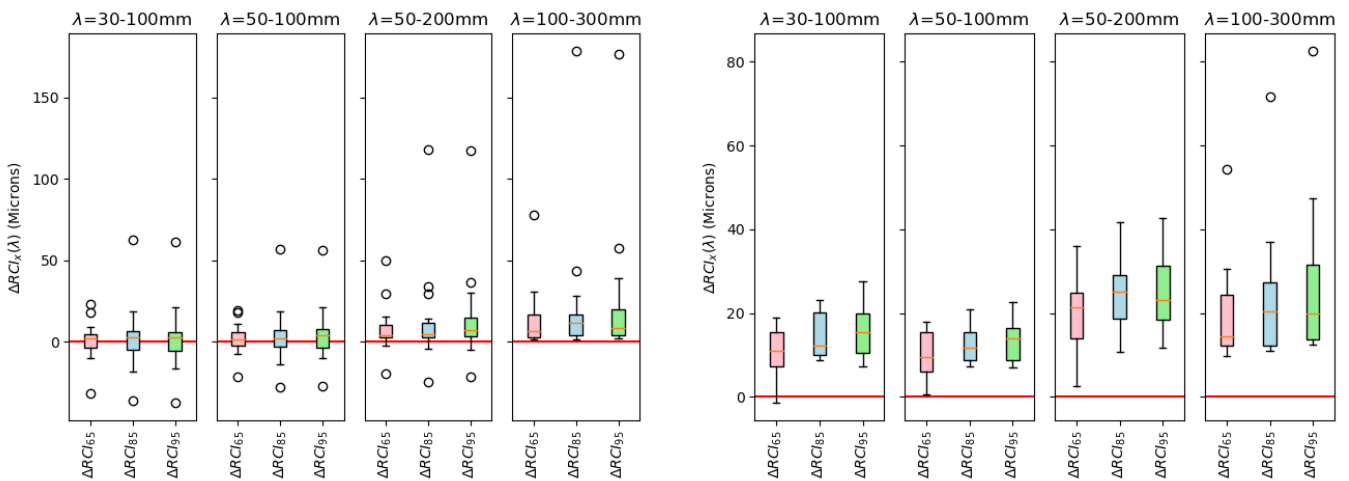
Figure 3.4: Box plot distribution of $\% \Delta RCI_x(\lambda)$ for the different tested wavelengths and percentiles ($n = 34$).

The performance of $RCI_x(\lambda)$ at Property B suggested that the block RMS technique foundational to RCI_x is capable of capturing the response to grinding as demonstrated at Property A only when the wavelength band is appropriately selected. This increases confidence that the metric is flexible enough to be moved from property to property to address their specific corrugation assessment needs.

3.2.3 $RCI_x(\lambda)$'s Sensitivity to Post-grind Measurement Delay

The grind signature left on the rail is not just an artifact of grinding. The roughness left on the rail is due to profiling the rail by concentrating metal removal where it is most-needed to address surface condition without wasting metal. This artificial roughness wears in through regular use, ideally leaving a smooth rail while minimizing the metal removal and rail aging during the maintenance process [33]. Corrugation and noise analysis at Property A indicated as much with the first post-grind samples not necessarily being the lowest results reported by the metrics. Property B's data sampling regimen featured post-grind measurements between 1 and 5 days inclusive post-grind, allowing us to see if the time between grinding and measuring allows $RCI_x(\lambda)$ to more accurately assess the quality of rail for the foreseeable future.

Figure 3.5a illustrates the distribution of data sampled 1 day post-grind, while Figure 3.5b shows 5 day post-grind data. When looking at Figure 3.5a, the low wavelength ranges tend to cluster around the 0 μm line, while the higher wavelength ranges tend to cluster just above this line. The low wavelength data still seems fairly symmetrical about the 0 μm line, whereas the higher wavelength ones skew positive. In contrast, when looking at Figure 3.5b, it is evident that across all bands the median has shifted up. The $\lambda=50\text{-}200$ mm and $\lambda=100\text{-}300$ mm bands still outperform the $\lambda=30\text{-}100$ mm and $\lambda=50\text{-}100$ mm bands in terms of the distribution being shifted up and having longer upper whiskers, continuing the trend of positive skew distributions.



(a) Distribution of $\Delta RCI_x(\lambda)$ for post-grind measurement samples taken 1 day post-grind (n=16). (b) Distribution of $\Delta RCI_x(\lambda)$ for post-grind measurement samples taken 5 days post-grind (n=8).

Figure 3.5: Distributions of $\Delta RCI_x(\lambda)$ based on how long after grinding the data was sampled.

These plots indicate that $RCI_x(\lambda)$ has a sensitivity to corrugation data sampled immediately post-grind. Best practice seems to indicate that assessing the response to grinding with $\Delta RCI_x(\lambda)$ is best served by waiting several days to allow the rail to wear in so that the working condition of the rail can be properly observed.

3.2.4 Performance of $\% \Delta RCI_x$ for Low Pre-grind $RCI_x(\lambda)$

The distributions found in Figure 3.3 and Figure 3.4 raised two related questions: is there a common trend in the severely negative $\% \Delta RCI_x(\lambda)$ datapoints, and is it an issue systemic to the metric? We sought to investigate how $\% \Delta RCI_x(\lambda)$ is related to the initial condition of the rail $RCI_x(\lambda)^{\text{Pre-grind}}$, seeing as it is a function of the metric's initial assessment of corrugation quality.

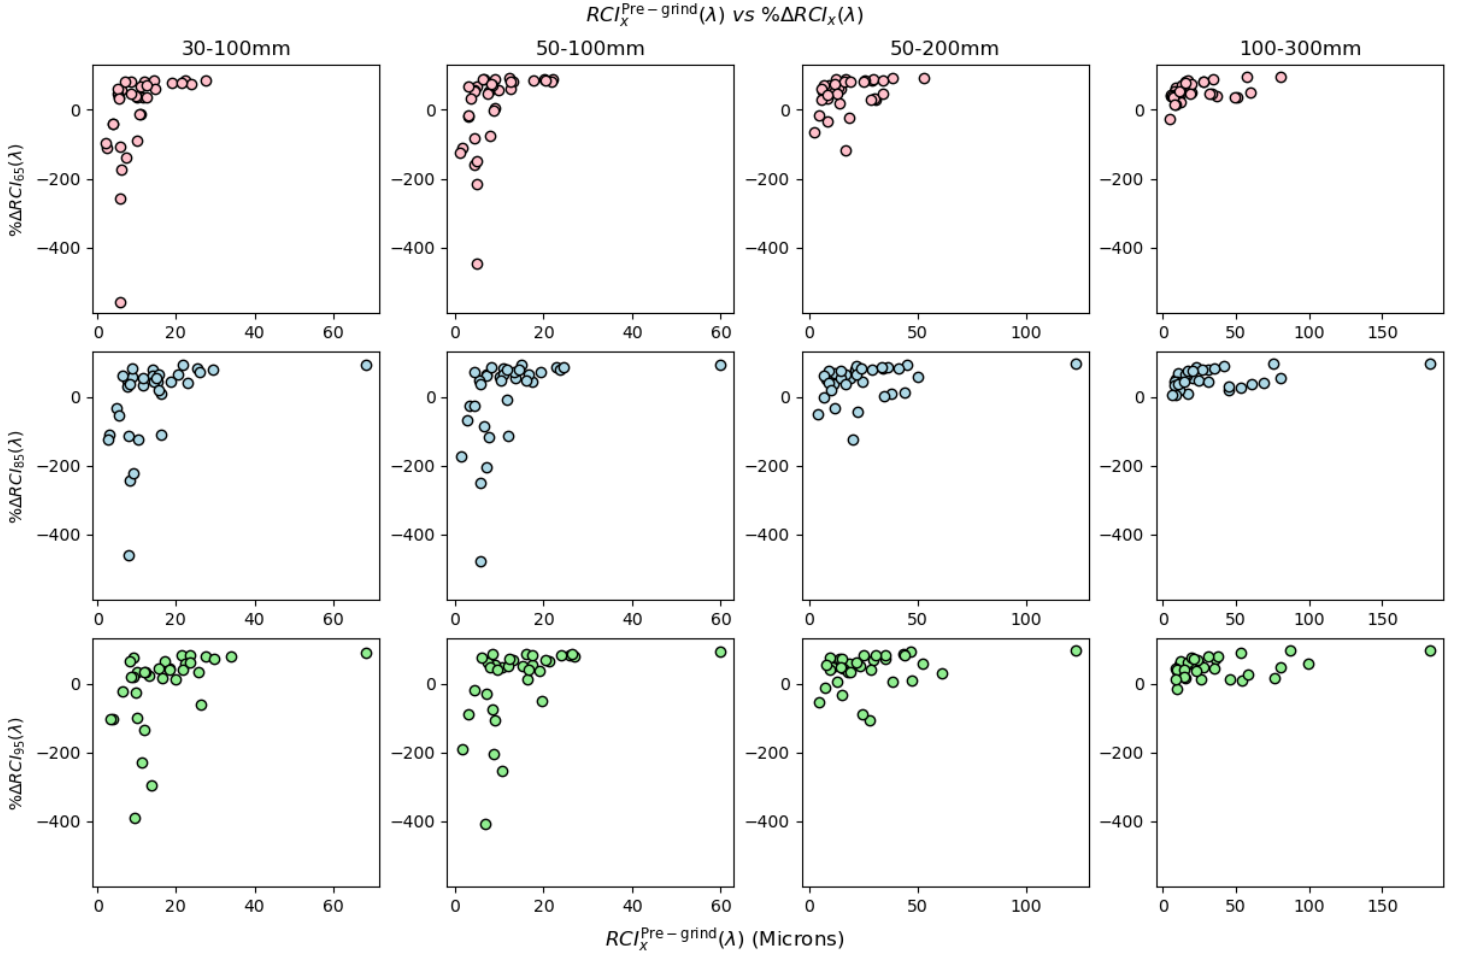


Figure 3.6: 4×3 grid of plots visualizing each percentile and wavelength band combination.

It is intuitive that any ΔRCI_x or $\Delta RCI_x(\lambda)$ value requires some context to determine if grinding was successful. If corrugation was improved by $10 \mu\text{m}$, that can be viewed as significant or not depending on whether the pre-grind rail quality was characterized as $12 \mu\text{m}$ or $120 \mu\text{m}$. As such, $\Delta RCI_x(\lambda)^{\text{Pre-grind}}$ was plotted against $\% \Delta RCI_x(\lambda)$ for each combination of percentile and wavelength band to investigate any potential trends that could be observed. This grid of plots can be found in Figure 3.6.

As was the case in Figure 3.4, the $\lambda=30-100 \text{ mm}$ and $\lambda=50-100 \text{ mm}$ wavelength bands have negative $\% \Delta RCI_x(\lambda)$ values that extend much further down than the $\lambda=50-200 \text{ mm}$ and $\lambda=100-300 \text{ mm}$ bandwidths. However, despite the magnitude of these values, all wavelength bands demonstrate similar trending in that the lower the $RCI_x(\lambda)^{\text{Pre-grind}}$, the wider and less precise the spread of the $\% \Delta RCI_x(\lambda)$ is shown to be. As the pre-grind

$RCI_x(\lambda)$ increases, the distribution tends to taper higher towards 100%. This suggests that while $RCI_x(\lambda)$ seems fairly adept at detecting corrugation removal with proper wavelength bracketing, it is best applied when the initial $RCI_x(\lambda)$ metric is high, indicating poor rail quality. When applied to low $RCI_x(\lambda)^{\text{Pre-grind}}$, small changes in $\Delta RCI_x(\lambda)$ are proportionally larger and if the metric shows corrugation has worsened slightly, this will indicate a large increase relative to the good quality rail pre-grind. It similarly suggests that when addressing severe pre-grind corrugation, proper grinding will be indicated very strongly using $RCI_x(\lambda)$.

This all suggests that while research into $RCI_x(\lambda)$ and its offshoots has shown promise when diagnosing and assessing corrugation correction, the metric is not infallible and is context-dependent in its usefulness. For example, using only $\% \Delta RCI_x(\lambda)$ or only $\Delta RCI_x(\lambda)$ to measure the quality of grinding may be misleading depending on the severity of $RCI_x(\lambda)^{\text{Pre-grind}}$. When rail is initially very rough, applying a rough grind will yield very high $\Delta RCI_x(\lambda)$ and $\% \Delta RCI_x(\lambda)$ values, but the resulting rail may still be rougher than the property desires. Similarly, polishing rail in good condition will likely return poor $\Delta RCI_x(\lambda)$ and $\% \Delta RCI_x(\lambda)$ performance, despite the finished rail being smooth. The work in this thesis indicates that RCI_x and $RCI_x(\lambda)$ demonstrate sensitivity to roughness, and that ΔRCI_x , $\Delta RCI_x(\lambda)$ and $\% \Delta RCI_x(\lambda)$ demonstrate an appropriate response to grinding. However, the grind response of any one sample may be less reliable than when applied to many samples as an aggregate due to the variance in initial rail condition.

3.3 Chapter Summary and Next Steps

The research into RCI_x , and its extension $RCI_x(\lambda)$, are relevant to modern day corrugation assessment and mitigation. The desire to quantify corrugation using a single number makes for straightforward rail management decisions, and allows transit properties to better serve their riders.

Initial RCI_x work focused on refining a block RMS-based corrugation assessment technique, and subsequent usage of that technique in modeling the growth and remediation of the phenomenon. Property A provided the location for a property-based long-term study of an RCI_x application to allow us to understand how the metric reflects corrugation.

Subsequent work into $RCI_x(\lambda)$ demonstrated how flexible this methodology needs to be in order to be usable at a variety of properties. Property B revealed that Property A is a location with its own unique corrugation issues, and that some of the rigidity in the metric needed to be relaxed. This allowed better understanding of the wavelength band selection and how crucial it is in using $RCI_x(\lambda)$ as a corrugation metric. Moreover, ensuring that adequate time has elapsed after grinding before measuring has significant impact on the metric's ability to assess the roughness on the rail. Lastly, $RCI_x(\lambda)^{\text{Pre-grind}}$'s impact on the effectiveness of the metric cannot be understated. The reliability of a single sample is heavily dependent on the state of rail pre-grind.

This work can be continued through repetition at a third property that can combine the best parts of Properties A and B. A property willing to measure corrugation at multiple locations throughout the property at regular

time intervals would allow testing of $RCI_x(\lambda)$ as a growth modeling metric, as well as giving the metric another test site for its flexibility. In addition, if RCI_x and $RCI_x(\lambda)$ are to be adopted, the derivation of decision-making thresholds would be a logical continuation to aid grinding decisions with the metric.

Chapter 4

Analysis of Wayside Noise Generated by Trains

This chapter describes the analysis and resulting conclusions drawn from work with wayside collection of audible noise data sampled from the same two Canadian rail transit properties studied in relation to corrugation in the previous chapter. Measuring corrugation is time-consuming for property maintainers, and spending this time by measuring corrugation on smooth rail is costly. A train passing over a section of corrugated rail generates sound that can be measured by a nearby microphone. Therefore, there is a desire to monitor this sound and understand if corrugation can be monitored by proxy using wayside vehicle noise measurements to avoid wastefully measuring smooth rail directly.

Throughout this work, uses of the term “noise” is common industry nomenclature when referring to audible sound, generated as a byproduct of vehicles passing over corrugated rail. This is not referring to noise from a statistical variability perspective, nor does it refer to noise in the context of unwanted data captured in a signal (i.e. the signal-to-noise ratio) unless explicitly indicated.

4.1 LA_{eq} and Spectrum Analysis

4.1.1 Property A Analysis

As previously discussed, Property A’s dataset provided us with a robust sampling regimen at a location with tangent track geometry, as well as one with curve track geometry. Wayside noise measurements and corrugation samples from both the left and right rail were taken monthly, starting the month after the rail was ground. This gave a clear picture of the growth in generated auditory noise as well as the growth in corrugation as functions of use (MGT) and the response to grinding in both noise and corrugation.

For a given date and track geometry (tangent or curve), between 14 and 24 noise measurements were sampled

using a sound level meter, where each measurement was the audible sound of one train passing by the location of interest. This was exported into an Excel-readable spreadsheet with A-weighted equivalent noise (LA_{eq}), and the A-weighted octave band analysis for both $1/3$ octaves and whole octaves. These samples were consistently taken a set distance from the guideway, at track height. This information is tabulated in Table 4.1.

Track Geometry	Tangent Curve
Noise Measurements	LA_{eq} LA_{eq} Spectrum
Data Collection Period	Monthly
Microphone Distance from Guideway	8 m (Tangent) 13 m (Curve)
Microphone Height	Track Height
Noise Samples Per Day	14-24


Table 4.1: Property A Data Summary.

For a given day, these samples were aggregated and summary information was extracted to be used as a representative sample. The mean LA_{eq} , as well as the 65th, and 95th percentiles were chosen as possible representatives for the overall A-weighted noise on the day, denoted by $LA_{eq_{mean}}$, $LA_{eq_{65}}$, and $LA_{eq_{95}}$ respectively. The standard deviation of the LA_{eq} values for the day, $LA_{eq_{SD}}$, was also calculated. Each LA_{eq} spectrum was averaged in each frequency bin to represent the noise sampled on that day. That is, the LA_{eq} value in the 63 Hz band was taken from each sample and averaged, then the 80 Hz band, and so on. The resulting spectrum was the LA_{eq} spectrum used in analysis.

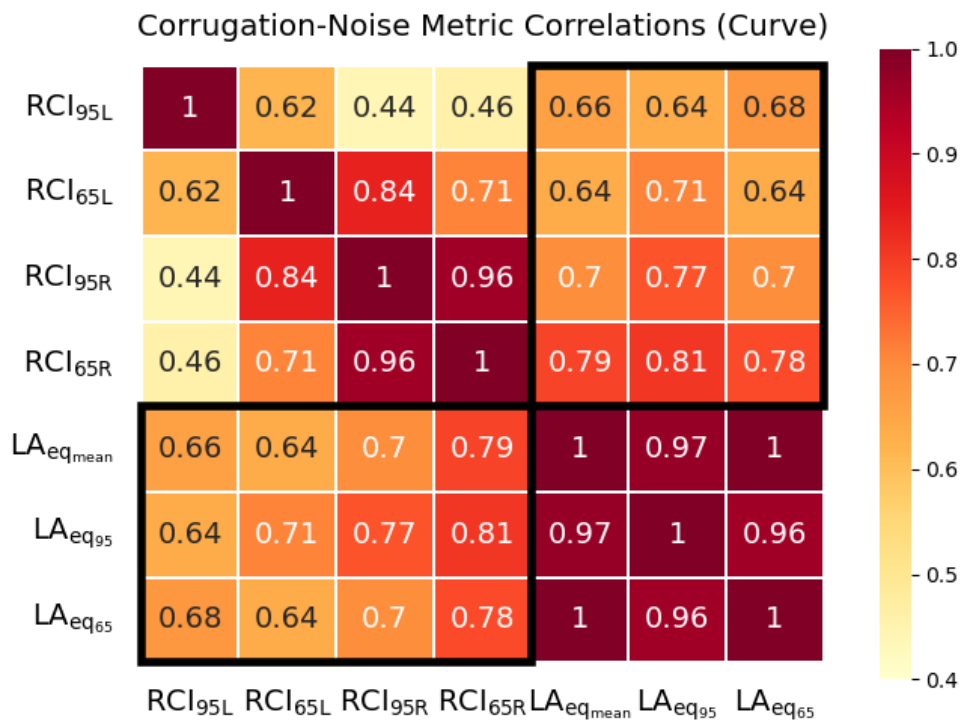
This wayside noise information was paired with corrugation information for the corresponding day. The corrugation data was sampled along the rail face 25 mm to 45 mm from gauge side in 5mm intervals, and the running band was selected based on the track geometry. For the tangent section of the rail, the 35mm sample was selected as the running band on both rails, while for the curve section, 25 mm and 45 mm were initially chosen for the high and low rail respectively. Throughout the study, the 25 mm and 45 mm bands stopped being regularly sampled, and so the running band for curve geometry was shifted to the 30 mm and 40 mm bands instead. The 5 m block RMS for each block in the sample was calculated, and the RCI_{65} , RCI_{85} , and RCI_{95} were selected appropriately on each rail, and denoted with a corresponding second subscript L or R .

Similar to the corrugation study of Chapter 3, cumulative MGT experienced by the rail due to usage was adopted as a proxy for time to generalize across properties with varying degrees of rail traffic.

The data was tabulated, and two correlation matrices were constructed to compare the relationships between each of the corrugation and wayside noise metrics, separated by track geometry. These figures are shown in



(a) Tangent rail data.



(b) Curve rail data.

Figure 4.1: Correlation matrices of corrugation and wayside noise metrics for tangent and curve rail sections. The black boxes identify where wayside noise metrics and corrugation metrics are correlated. Note that the colourbar scale spans 0.4-1.0 for visual emphasis.

Figure 4.1. The relationship between LA_{eq} and RCI_x measurements showed R^2 values between 0.64 and 0.97 depending on the track geometry and selected measurement (i.e. $LA_{eq_{mean}}$ and RCI_{65L} on tangent geometry), which was deemed sufficiently correlated to justify investigating the growth relationship between the two.

The plots in Figure 4.2 demonstrate a growth pattern in the corrugation that is reflected in the LA_{eq} information. Both the corrugation and vehicle noise information show a strong response to grinding (the occurrence of which is denoted by the vertical black line) by dropping back down to the initial levels. Figure 4.2c is the plot where the relationship visually appears the weakest, which is the left rail in the curve section. The curve in question is a right-hand curve, meaning the left rail is the high (outer) rail. Recall that for a train moving through a curve, literature suggests that the dynamic load may be experienced more substantially by the low (inner) rail [34]. Figure 4.2d shows a more pronounced relationship between noise and the right rail corrugation growth, which agrees more with this conventional knowledge of rail dynamics in curves.

The results of this analysis demonstrated a relationship between A-Weighted equivalent sound and various RCI_x measurements, and naturally led to validating these findings at another property.


4.1.2 Property B Analysis

Property B provided a different structure of dataset from Property A, and the work at this location showcased shortcomings with LA_{eq} when taken out of a controlled test site like the one at Property A. This property provided paired pre-grind and post-grind noise and corrugation measurements. Consequently, the scope of the validation was limited to an examination of the grind response rather than long-term growth and response trends.


The noise and corrugation measurements were taken from 10 unique curve track segments throughout the property, for both eastbound and westbound track. This yields 20 unique locations.

One shortcoming of this dataset's sampling regimen is the number of samples per day. At Property A, there were between 14 and 24 noise samples for a given location and date, whereas Property B has between 0 and 4 samples for a given location and date, where date refers to only either immediately pre- or post-grind. Due to the comparative lack of collected samples, it is unrepresentative to try to summarize an ensemble of points like at Property A. In comparison to Property A, the microphone distance from the track was not recorded, meaning the amplitude of noise recorded shouldn't be directly compared across properties. The pre- and post-grind noise and corrugation data was separated according to the curve at which it was recorded, and however many noise samples taken at that curve were viewed in aggregate. This data-handling regimen is tabulated in Table 4.2, and the reader can compare this information to that in Table 4.1.


An initial test was conducted on the data where the pre-grind and post-grind LA_{eq} were compared in order to view any decibel drop from pre-grind to post-grind. This test yielded inconsistent results, where the post-grind measurement was not reliably showing a drop in dB as it did at Property A. 5 of the 15 locations showed post-grind noise *rises*, between 3.3 and 13.7 dB. More detail was required, so the unweighted spectra of the noise




(a) Wayside noise and left rail corrugation growth at tangent track location.



(b) Wayside noise and right rail corrugation growth at tangent track location.



(c) Wayside noise and left rail corrugation growth at curve track location.



(d) Wayside noise and right rail corrugation growth at curve track location.

Figure 4.2: Corrugation growth captured by RCI_{65} , RCI_{85} , and RCI_{95} , and wayside noise captured by LA_{eq65} , and LA_{eq95} , separated by rail and track geometry at Property A.

	Property A	Property B
Track Geometry	Tangent Curve	15 Curves Throughout Property
Noise Measurements	LA_{eq} LA_{eq} Spectrum	LA_{eq} LA_{eq} Spectrum
Data Collection Period	Monthly	1-5 days Pre-Grind and Post-Grind
Microphone Distance from Guideway	8 m (Tangent) 13 m (Curve)	Unspecified
Microphone Height	Track Height	Unspecified
Noise Samples Per Day	14-24	0-4

Table 4.2: Property B Data Summary. The data in Table 4.1 is repeated for reference.

at these curve were examined.

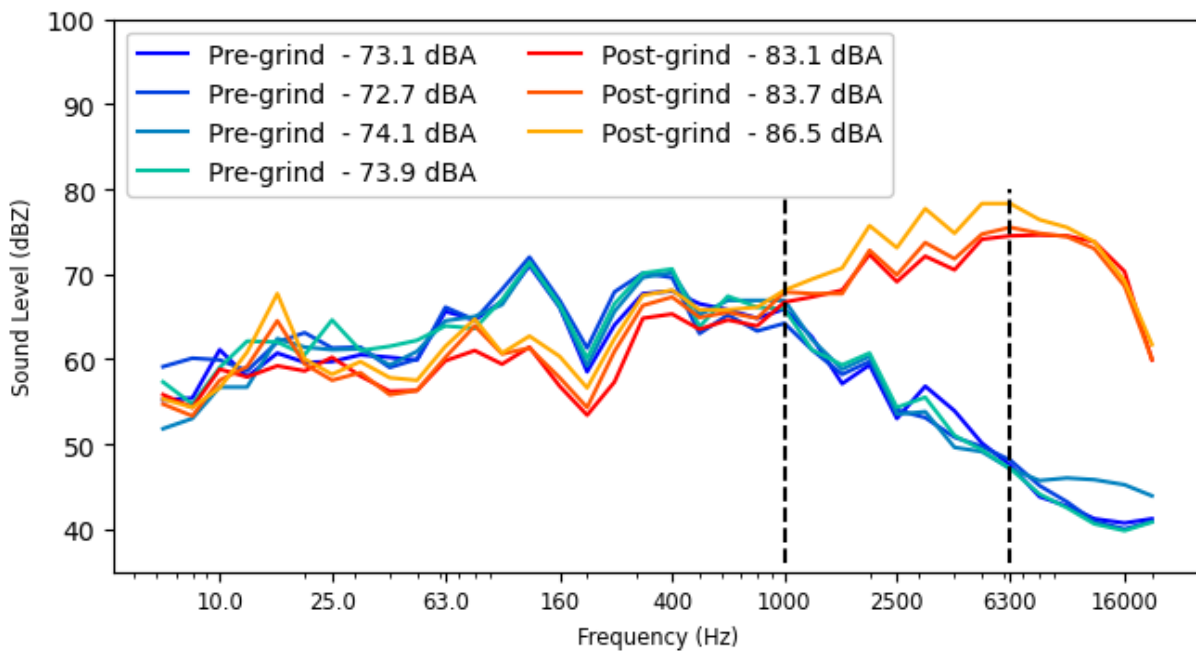


Figure 4.3: Unweighted spectra of samples at a particular curve at Property B. Pre-Grind LA_{eq} was between 72.7 and 74.1 dB, while Post-Grind LA_{eq} was between 83.1 and 86.5 dB, as presented in the legend. Note the post-grind increase in noise where A-weighting is especially sensitive, denoted by black dotted lines, and how this is reflected in the dBA of each sample.

The spectra shown in Figure 4.3 demonstrated unusual post-grind increases in volume at frequencies above 1000 Hz. Upon further research, it was hypothesized that this increase is due to rail effects separate from corrugation as summarized in Table 4.3.

<i>Noise Type</i>	<i>Frequency Range (Hz)</i>
Rolling	30-2500
Rumble (Corrugation)	200-1000
Flat Spots	50-250 (standard operating speed)
Ground-borne Vibrations	30-200
Top of Rail Squeal	1000-5000
Flanging	5000-10000

Table 4.3: Common rail noise emissions and their frequency presences, making note of Corrugation and Top of Rail Squeal [35].

The unusual presence of >1000 Hz sound is likely not due to corrugation, but rather due to friction between the wheelset and the top of the rail. Under normal operating conditions, a transit agency may apply a top of rail friction modifier to improve running conditions of the train [36]. This substance is removed for grinding maintenance, and then reapplied. At Property B, it is believed that a friction modifier was present during the pre-grind measurements, removed for the grinding, and then not reapplied before the post-grind noise measurements were taken at some of the locations. As a result, there is an anomalous increase in wayside noise level for the frequency bands associated with top of rail squeal.

By coincidence, the frequencies of noise associated with top of rail squeal align almost perfectly with where the A-Weighting curve is not attenuating, as demonstrated in Figure 4.4. Consequently, A-Weighting is more sensitive to top of rail squeal than it is to corrugation in situations where both are present. It is speculated that a lack of top of rail squeal at Property A, either due to the presence of a top of rail friction modifier or the dynamics of the wheelset, was obfuscating one of the potential shortcomings of the LA_{eq} measurement.

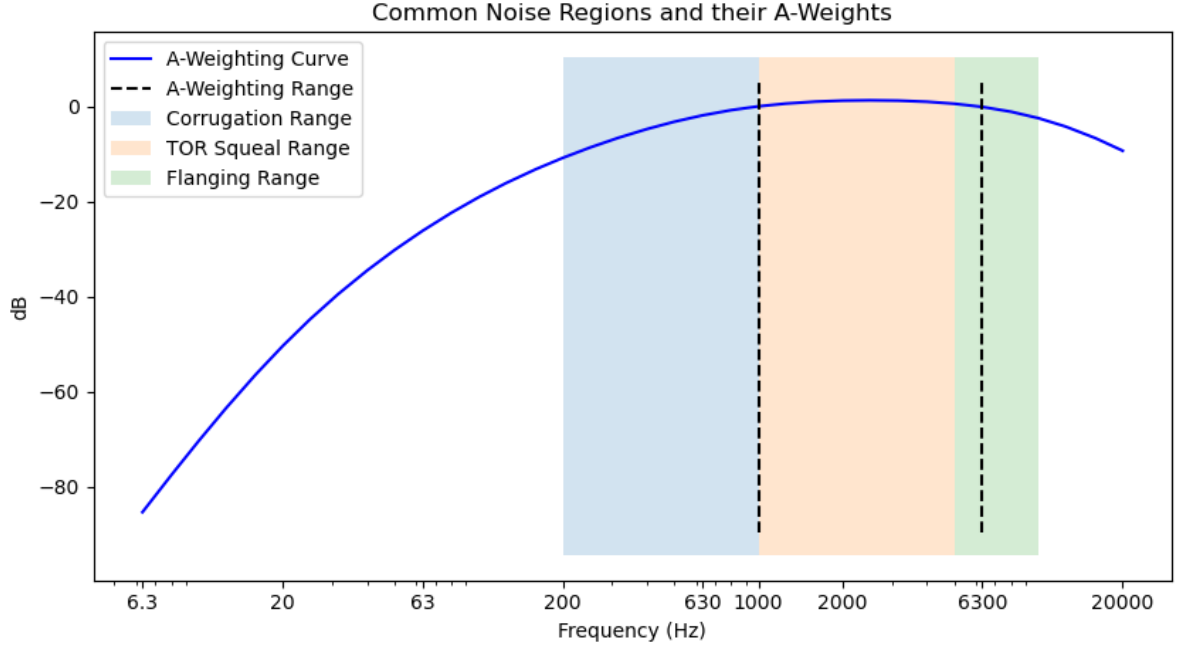


Figure 4.4: A-Weighting Curve showing the overlap between its passband and the usual noise frequencies of Top of Rail Squeal. The vertical black dotted lines represent the locations where A-Weighting has a gain of 0 dB or greater .

4.2 Velocity- and Wavelength-Adapted Boxcar Filtering ($Lv_{\text{eq}}(\lambda, v)$)

The use of LA_{eq} as a measure of wayside noise produced by corrugation shows promise (as demonstrated at Property A), yet the attempt to validate the potential relationship between the two at Property B revealed shortcomings when applied to less controlled environments. This fueled the definition of $Lv_{\text{eq}}(\lambda, v)$, a novel noise-weighting scheme parameterized by incorporating wavelength and velocity information to isolate frequencies of noise explicitly due to corrugation.

4.2.1 Property B Analysis

Corrugation is a spatial wave that, from the perspective of the vehicle passing over it at a velocity of $+v$, has a velocity of $-v$. The frequency of noise generated by this interaction f , is a function of the spatial wavelength λ , and the velocity of the train v passing over it. This is derived from the wave velocity equation,

$$f = \frac{v}{\lambda}.$$

In the case of corrugation, there may be a range of frequencies of interest, meaning the wavelength parameter λ actually denotes a wavelength range. Similarly, a given wayside noise measurement may not come with accompanying train velocity information meaning the v parameter must also tolerate some leniency in its precision. For example, the dataset at Property B was taken from locations where speed limits are between 60-90

km/h. As such, these parameters represent ranges, explicitly $\lambda = [\lambda_{\min}, \lambda_{\max}]$, and $v = [v_{\min}, v_{\max}]$.

The range of f is dependent on the ranges of λ and v , meaning $[f_{\min}, f_{\max}]$ can be calculated to include the entire range of possible noise from the tolerance ranges in the parameters. Explicitly,

$$f_{\min} = \min(f) = \frac{\min(v)}{\max(\lambda)} \quad (4.1)$$

$$f_{\max} = \max(f) = \frac{\max(v)}{\min(\lambda)}. \quad (4.2)$$

It is important to note that the metric remains valid if any or both of λ and v are single specific values, which gives $Lv_{\text{eq}}(\lambda, v)$ flexibility to be used broadly and precisely.

Once the frequency range is calculated, the transfer function $h(f)$ for $Lv_{\text{eq}}(\lambda, v)$ is a shifted and scaled boxcar function that defines the gain of a particular frequency. Explicitly,

$$h(f) = \begin{cases} -30 \text{ dB}, & \text{if } f_{\min} \leq f \leq f_{\max} \\ 0 \text{ dB}, & \text{otherwise,} \end{cases} \quad (4.3)$$

which is illustrated in Figure 4.5. For completeness, different smooth transfer functions were tested due to the signal processing problems that come from the inverse Fourier transform of a rectangular transfer function. Namely, these functions were overlapping sigmoids and a super-Gaussian. These were constructed to attenuate by a factor of -3 dB at f_{\min} and f_{\max} . It was determined that a smooth transfer function was unnecessary due to two reasons. First, there was a difference of at most 1.6% in the resulting equivalent noise calculations from spectra when compared to the boxcar transfer function. Second, there is no desire at present to filter the noise signal in time. In the future, some adaptations to $Lv_{\text{eq}}(\lambda, v)$ will need to be implemented to handle the discrete cut-off of the current transfer function, but the boxcar transfer function satisfies the need to window a specific part of the noise spectrum.

As illustrated in the previous chapter, the corrugation at Property B exists at longer wavelength than that of Property A, with both frequency ranges of 50-200 mm and 100-300 mm showing reliable performance. As such, λ was chosen to be the overlap of these ranges, specifically 100-200 mm. As previously mentioned, v was chosen to be 60-90 km/h (16.7-25 m/s), a wide range of possible velocities. This gives $Lv_{\text{eq}}(100\text{-}200 \text{ mm}, 60\text{-}90 \text{ km/h})$ cutoff frequencies of

$$f_{\min} = \min(f) = \frac{16.7 \text{ m/s}}{0.2 \text{ m}} = 83.3 \text{ Hz}$$

$$f_{\max} = \max(f) = \frac{25 \text{ m/s}}{0.1 \text{ m}} = 250 \text{ Hz.}$$

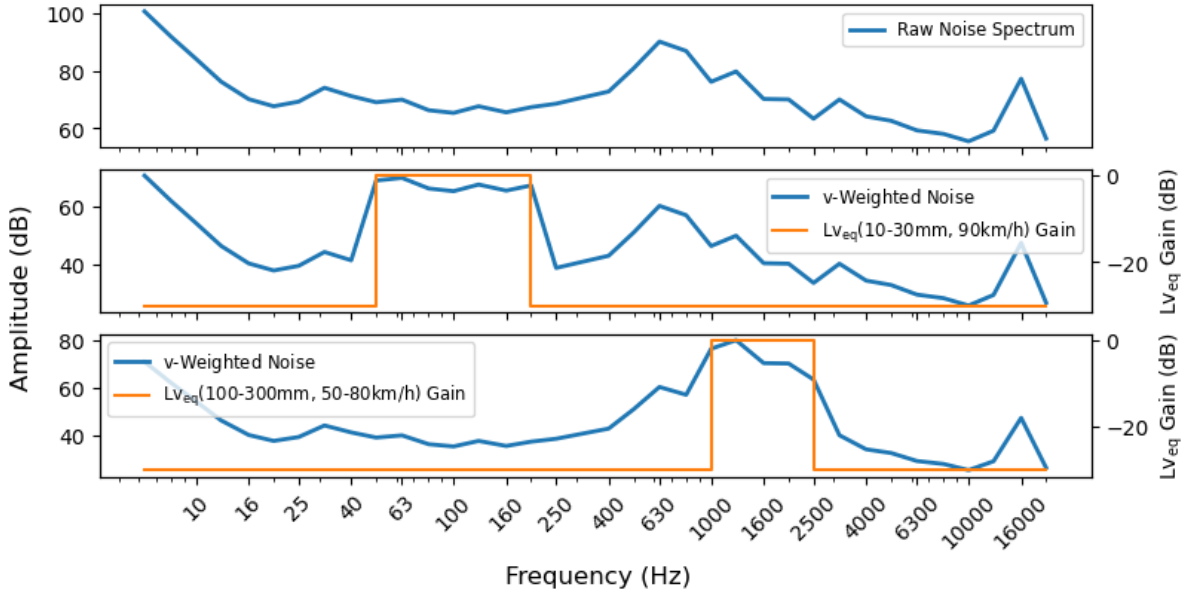


Figure 4.5: An unweighted noise sample, with two different $Lv_{eq}(\lambda, v)$ weights applied to the spectrum.

The results of applying $Lv_{eq}(100\text{-}200\text{ mm}, 60\text{-}90\text{ km/h})$ can be seen in Figure 4.6.

This methodology demonstrated an appropriate dB drop in the 5 previously anomalous locations where LA_{eq} rose post-grind, while retaining an appropriate dB drop in the 10 locations where LA_{eq} correctly assessed a reduction in corrugation. These results demonstrate that appropriately parametrizing velocity and corrugation leads to a robust mechanism for capturing wayside noise linked to corrugation, even when there is an introduction of sound frequencies unrelated to corrugation.


The final step in this research was to return to Property A to confirm $Lv_{eq}(\lambda, v)$'s flexibility.

4.2.2 Property A Analysis

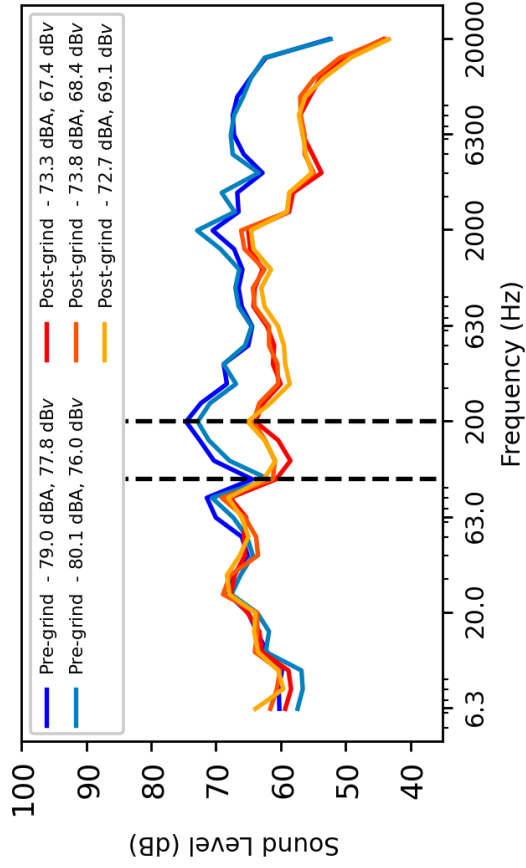
In applying $Lv_{eq}(\lambda, v)$ to Property A, some added information is required. The tangent/curve measurement site is known to have trains passing after accelerating to a posted 90 km/h speed limit, so $Lv_{eq}(\lambda, v)$ was applied using v of 80-90 km/h (22.2-25 m/s). The RCI_x research performed previously at Property A indicated that corrugation in the 30-100 mm range was reflective of the wear present at this system, making that the λ parameter.

Where the initial LA_{eq} research was conducted by averaging each frequency bin from the spectra for a given date, $Lv_{eq}(\lambda, v)$ did not have the dataset necessary to support this kind of methodology. As such, the validation of $Lv_{eq}(\lambda, v)$ at Property A used an overall analysis of spread and growth of spread for the data at each date, and verified these results qualitatively.

Figure 4.7 shows that the overall shape of the data growth remains similar to that of Figure 4.2, when the operating velocity and rail corrugation are known. The overall magnitude of the decibels has shifted down due slightly due to eliminating more frequencies with the $Lv_{eq}(\lambda, v)$ measurement than the LA_{eq} measurement, but




(a) LA_{eq} and Lv_{eq} (30-100 mm, 60-90 km/h) comparison at a specific curve, highlighting v -weighting's cutoff frequencies. Note that where A-weighting does not detect the noise reduction, v -weighting captures the low frequency noise difference.




(b) LA_{eq} and Lv_{eq} (30-100 mm, 60-90 km/h) comparison at a specific curve, highlighting v -weighting's cutoff frequencies. Note that even here where A-weighting observes a sound level reduction, v -weighting does as well, suggesting v -weighting does not overfit to anomalous A-weighting measurements.

Figure 4.6: A-weighting and v -weighting comparison at two curves. Figure 4.6a illustrates v -weighting disagreeing with A-weighting, while Figure 4.6b showcases agreement.



(a) $Lv_{eq}(30-100 \text{ mm}, 80-90 \text{ km/h})$ distribution as a function of wear, tangent rail.



(b) $Lv_{eq}(30-100 \text{ mm}, 80-90 \text{ km/h})$ distribution as a function of wear, curve rail.

Figure 4.7: Box and whisker distribution plots for and $Lv_{eq}(30-100 \text{ mm}, 80-90 \text{ km/h})$, as a function of wear.

the trend is consistent. This reinforces the idea that $Lv_{eq}(\lambda, v)$ works as a corrugation-derived sound measurement method for wayside noise.

4.3 Chapter Summary and Next Steps

This research demonstrated the relationship between corrugation growth and associated wayside noise. Initial research at a controlled location showed that standard LA_{eq} worked as a noise summary metric for wayside vehicle noise measurements, and it effectively echoed the growth in corrugation detected by RCI_x . It was also demonstrated that LA_{eq} has some shortcomings when applied at a property with a less sanitized collection regimen as it has the potential for other rail phenomena to impact the ability to measure noise. In situations where rail phenomena like top of rail squeal or flanging occur, the frequency of noise generated by corrugation

is obfuscated from the LA_{eq} measurement. This shortcoming of LA_{eq} was mitigated by $Lv_{eq}(\lambda, v)$, where the weighting scheme demands slightly more information with regards to velocity and corrugation wavelength, but more reliably reflects corrugation-related noise present in a sample.

The industry has a strong desire to minimize the time spent with workers on the rail. If property maintainers can glean insight into the state of the rail through a covariate like wayside noise, it can improve efficiency and allow a property to dedicate more effort to locations through the property that demand more attention rather than wasting time assessing rail that is in good condition. Using wayside noise measurements to passively observe rail quality is a powerful tool, but this research has demonstrated that it can be misleading when the noise due to corrugation is not weighted appropriately.

Logical next steps may be to investigate the effectiveness of this methodology with on-board noise measurement. By using measurement equipment affixed to the vehicle, the amount of track that can now be assessed increases. There are likely to be issues with vibration, ambient train noise, and wind noise affecting the measurements, but the principle of corrugation measurement by proxy through noise is still sound when applied to on-board noise.

A different analysis philosophy that can be applied to wayside noise analysis is the more general stochastic analysis of the sampled wayside noise. This likely would require more data than was available for this project, but a stochastic analysis of wayside noise as a random process is perhaps a stronger overall methodology than the deterministic scalar summary used in this work.

Chapter 5

Vibration Analysis

This chapter provides an overview of analysis done on vibration data collected from axle-box accelerometers (ABAs), and outlines potential use cases for a variety of signal processing techniques in rail maintenance. Recall that the purpose of using ABA data is to use another type of data for corrugation analysis by proxy. The size of dataset collected in this way is massive, and the time signal representation of the accelerometer data does not provide enough information by inspection. Instead, we apply time–frequency analysis of the data, namely the wavelet transform, to gain insight into the spectral content of the data at different points in time.

Material presented in this chapter assumes a basic understanding of the Fourier Transform and frequency decomposition.

5.1 ABAs and ABA Data

5.1.1 Data Characteristics

As mentioned in Section 2.3, the ABA logs data into a raw binary file, and in-house code extracts the acceleration data in the x , y , and z axes, as well as timestamps. It does this in five second batches as the raw file is too large to be read into memory in its entirety.

Based on how the ABA is mounted to the train (see Figure 2.7), the x acceleration is the dimension of longitudinal travel. As the data batches are being decoded from the raw file, this dimension is used to determine whether the vehicle is stationary. The x acceleration has the least noise when the vehicle is at rest as the dynamics of travel add white noise to the signal, so the x acceleration standard deviation is calculated for each batch of data to determine if the vehicle is stopped at a station.

The code divides the raw data into station-to-station **segments**, according to where it detects the train being stopped. It is assumed that during revenue service, the train will travel forward to the end of the line, switch to the track running the opposite direction, and travel backward to the start of the line. This means that the x acceleration will be positive in one direction of travel, and negative in the other. The code detects when the x

acceleration switches sign, suggesting the train has reached the end of the line. Travel of multiple segments in a given direction is called a **trip**. This is illustrated in Figure 5.1.

This partitioning of data is crucial because the only time the location of the train is precisely known is when the train is stopped at a station. The end of this processing yields several files of acceleration data like Figure 5.2, where the position is precisely known at the start and end of the files, but everything in the middle is uncertain.

As previously mentioned, the ABAs being used are limited in terms of their data collection. GPS is not mounted with them so there is no true position data. Similarly, the trade-off of being easy and non-invasive to mount means there is no connection to the vehicle's speedometer or odometer. Lastly, the gyroscope in the IMU is not logging data, meaning any rotation due to curves or slopes is not measurable, nor can the unit be unbiased if the ABA is not mounted such that the z axis is parallel to gravitational acceleration.

As such, a reasonable but limited solution for locating the device given the data being collected is the use of dead reckoning [37], where the accumulation of forward acceleration is used to calculate the velocity and longitudinal displacement of the sensor at each timestamp. Dead reckoning is a well-known technique in the navigational space, with uses in automotive [38], marine [39], and aeronautical transportation [40], so the limitations of the technique are known as well [41], namely the accumulation of error stemming from accelerometer measurement noise, errors, and biases. Including additional sensor information such as that from speedometer, odometer, and gyroscope data can help to improve certainty by estimating system properties [42]. One model of uncertainty reduction commonly used with dead reckoning calculations is a Kalman filter, where the measured system state (e.g. acceleration, velocity, and position at a given timestamp) and an estimated system state from measured information are combined to minimize uncertainty [43].

In the case of the collected accelerometer data, the forward acceleration, velocity, and position of the vehicle is only precisely known at the start (0 m/s^2 , 0 m/s , and 0 m) and end (0 m/s^2 , 0 m/s , and $d_{\text{seg}} \text{ m}$) of each segment. In between those points, there is some level of uncertainty.

This is exemplified when calculating the x displacement d_x throughout a segment by double integrating the x acceleration data a_x . Performing this integration on a variety of different segments, and comparing the final d_x value to the true d_{seg} of the station-to-station segment led to seemingly random results, as the final displacement could be over-reported, under-reported, or approximately correctly-reported (within 5% of the true displacement). Moreover, some measurements indicated that the vehicle travelled forward, reversed, and came to a stop behind where it started (i.e. the final d_x value was negative), while never realizing the condition for stop detection previously discussed. This is obviously an error as there is no way the vehicle could decelerate, stop, and reverse, all while the time signals were consistently noisy and erratic, suggesting the ABA was in a dynamic state throughout.

Instead, this points to measurement error and possible instability in the ABA's hardware. One possible explanation is that the error arises as a result of dead reckoning. The measurement error may be such that each sample is approximately correct, but the repeated accumulation of small errors at each sample eventually yields




Figure 5.2: x , y , and z acceleration data for a station-to-station segment.

these inconsistent results. An alternative explanation is that some level of error is due to rotational dynamics of the vehicle. One possibility is that due to slopes in the track or rotations through curves, the normal force counteracting gravity is coupling into the measured a_x , affecting the dead reckoning calculation. Due to a lack of gyroscope data, this cannot be rectified. Yet another explanation is that the ABA may be mounted improperly leaving the signal biased, however attempts to unbias it by detrending and zero-meaning the data did not fix the issue. Finally, it may be the case that the vehicle dynamics are so violent that the accelerometer in the IMU is drifting throughout the segment. This is causing essentially random behaviour where it cannot be fixed in post-processing. Some segments seem to suggest this may be possible because the resting value at the end of the segment was sometimes up to 0.4 g (3.92 m/s^2) different from the resting value at the start.

These shortcomings in the data precede the desire to investigate corrugation with the need to spatially locate the ABA through a segment: without knowing position at least approximately, there is no way to investigate the state of the rail at a particular point in the accelerometer signal.

5.1.2 Towards Practical Use of ABA Data

Despite the dataset's noted shortcomings, there was still a desire to learn something from it with respect to rail condition. The understanding was that the effects of corrugation would be overwhelmingly present in the z acceleration, meaning that irrespective of the inability to accurately locate the data in a segment, there may still be information to be gleaned from the vertical component of acceleration. As such, the questions to be answered

were twofold:

Question 1. Is there any way to improve the ability to spatially locate the ABA during segment-to-segment travel?

Question 2. How detectable is rail corrugation in the ABA data, specifically the z acceleration a_z ?

While dead reckoning was shown to have major flaws, it was the only method to spatially locate the ABA from the data available. The location of the vehicle is known at the beginning and end of each segment, meaning regardless of the final reported displacement, the true final displacement is guaranteed to be d_{seg} . Therefore, the displacement at each time step can be scaled to ensure that the final displacement measurement is d_{seg} . By extension, it may be the case that even though the forward acceleration is incorrectly reported, it may be linearly proportional to the true acceleration, meaning it would similarly be scaled appropriately. For instance, if the ABA passed over a patch of corrugation 50 m into a 250 m long segment, that would appear in the accelerometer signal 20% of the distance through the segment. If dead reckoning indicates that the segment is 400 m long, it could be the case that the ABA passing over that corrugation was recorded properly, but dead reckoning indicates it occurred 80 m into the segment, which is still 20% of the segment length. Based on this assumption, the position was calculated using dead reckoning from a_x , but then scaled according to the known distance between stations. Explicitly, for a particular segment of length d_{seg} the position of the ABA for each timestamp was calculated in Source Code 1.

```
1 import numpy as np
2 def dead_reckoning(ax, ts, d_seg):
3     """
4     Dead reckoning scaling function
5
6     Parameters
7     -----
8     ax : np.ndarray
9         Longitudinal acceleration of a segment in m/s**2
10    ts : float
11        Sampling period in s
12    d_seg : float
13        Known distance between stations in this segment
14
15    Returns
16    -----
17    ax : np.ndarray
18        Passed through longitudinal acceleration
19    ax_scaled : np.ndarray
20        Longitudinal acceleration scaled according to d_seg parameter
21    vx : np.ndarray
22        Accumulated velocity at every time step
23    dx : np.ndarray
24        Accumulated displacement at every time step
25    """
26    vx = np.cumsum(ax)*ts
27    dx = np.cumsum(vx)*ts
28
29    ratio = d_seg/dx[-1]
30
31    dx *= ratio #dx is now scaled so the end of the run is d_seg
32    vx *= ratio #vx is now scaled in the same way
33    ax_scaled = ax*ratio
34
35    return ax, ax_scaled, vx, dx
```

Source Code 1: Scaling dead reckoning according to known displacement data.

This technique assumes that as long as the displacement of a segment is monotonic and increasing (i.e. always away from the starting station), the start and end point of a segment will always be absolutely correct. This technique cannot fix the issue with accelerometer data indicating a polytonic displacement (i.e. travelling away from the starting station, then reversing).

Knowing position absolutely at the start and stop of a segment begged the question of whether there are ways to mark the location absolutely between segment endpoints: are there other constant stationary track characteristics that could be identified in a segment's a_x acceleration data? These anchor points could spatially locate the accelerometer at points throughout a segment, not just at the endpoints, allowing greater precision throughout more of a segment. These points would reset the dead reckoning calculation at those locations, thus providing better spatial certainty to property maintainers interested in the state of the rail at a location dictated by ABA data. A visualization of an abstract "positional uncertainty" is presented in Figure 5.3.

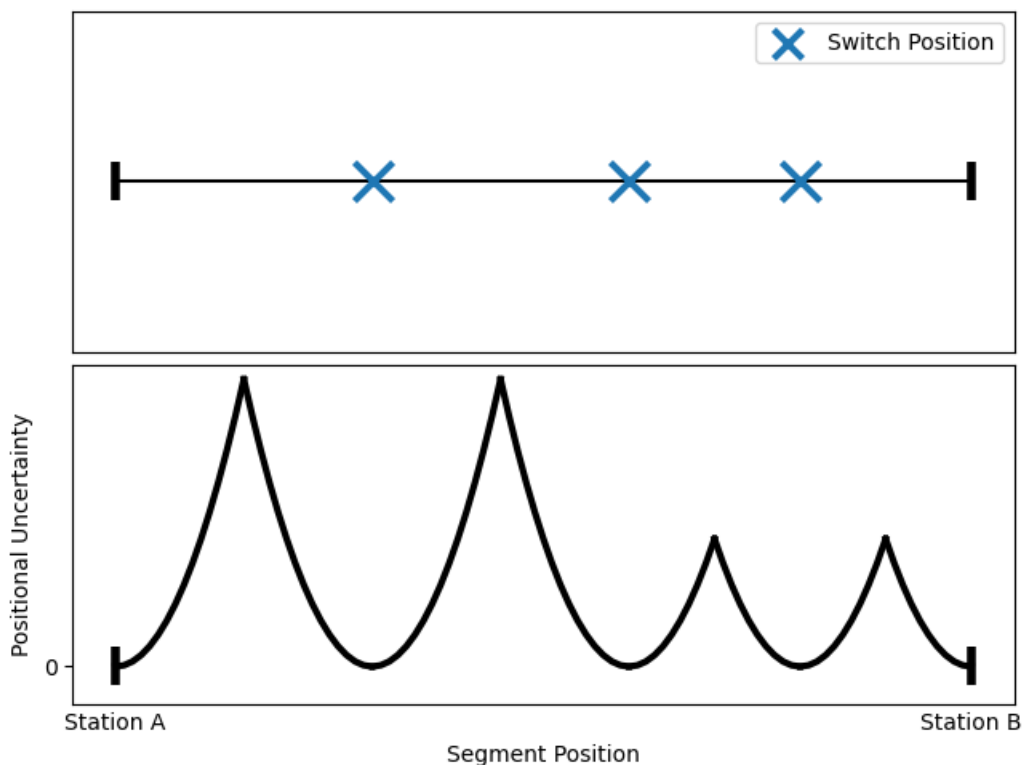


Figure 5.3: An abstract "positional uncertainty" as a function of position between two stations at the start and end of a segment. Static track characteristics like switches, which can anchor position in a segment, are marked with a blue \times in the top graph, and the positional uncertainty is 0 in the bottom graph at these locations. At locations between these key points, dead reckoning error adds uncertainty in the positional calculation.

Instead of integrating and scaling the full segment length, the integration and scaling would be done on all data between the starting station and the first switch, then on the data between the first and second switch, and so on until the data between the last switch and the end station. This technique attempts to minimize the error accumulation due to dead reckoning, by giving points at which the error accumulation is "reset", which seeks to address the first question to be answered.

It is expected that the target track characteristics will manifest as violent transients in the accelerometer

signals. It is understood that a transient is a coalescence in time of all frequencies, which is why, when taken to its logical extreme, the Dirac delta function has a Fourier transform which is constant for all frequencies $-\infty < f < \infty$. Therefore, locating these important points in a segment will involve finding points where the signal has high presence across the entire frequency spectrum.

The second question is less problematic to address. If corrugation can be identified in ABA data, it will manifest according to the wavelength of the corrugation and the velocity of the vehicle passing over it. As was the case for noise, this will be dictated by the relationship $f = v/\lambda$, where v is vehicle velocity in m/s, and λ is corrugation wavelength in m. For 30-100 mm corrugation, with velocities in the range of 60-90 km/h, this corresponds to frequency information in the range of roughly 170-830 Hz. For our ABAs with sampling frequency of $f_s = 3225$ Hz, this is very low frequency information.

Therefore, Question 1 can be addressed by finding short time, high frequency spikes in the frequency spectra of the data, while Question 2 will be addressed using long time low frequency rumbles in the frequency spectra. However, these outcomes can only be addressed qualitatively in this preliminary work as the uncertainty in hardware limits any conclusions that can be drawn. Without additional sensor data from either the IMU's gyroscope, or position/velocity data from the vehicle's motor control, the positional certainty in the accelerometer data from the ABA will be suspect. Additionally, with no temporally paired corrugation data, any estimations about the rail condition cannot be verified. If this work demonstrates that the a_z signal does contain information on track position and rail condition encoded in it, the future steps will necessarily involve improving the hardware shortcomings before this work can be expanded upon.

5.1.3 Fourier Transform and Short Time Fourier Transform (STFT, Spectrogram)

As shown in Figure 5.2, a typical ABA signal is sampled at such a high rate in such a noisy environment that visual inspection of the time domain signal is not instructive. In an effort to separate low frequency rail condition information from full spectrum (which will be present in high frequencies) track characteristic information, a frequency decomposition of the signal is warranted.

Frequency inspection proves to be useless when using just the standard Fast Fourier Transform (FFT), as demonstrated in Figure 5.4. Where the time domain representation in Figure 5.2 only contains time information, this representation only contains frequency information. If the lobe between ~ 300 and 400 Hz in Figure 5.4 is indicative of corrugation, it is unlikely that this exists throughout the entire length of the segment. Identifying and locating corrugation necessitates the inspection in terms of both time and frequency simultaneously.

Spectrograms are a common tool for simultaneous investigation of frequency and time. They are constructed by windowing a signal at different points in time, and calculating the FFT on this windowed subsection of the signal [44]. Alternatively, this can be done through a convolution of filter banks according to the frequency resolution of interest [45]. The resolution of this investigation in both time and frequency is directly related to the length of the window used to obtain the signal subsection. A longer window in time has more precise

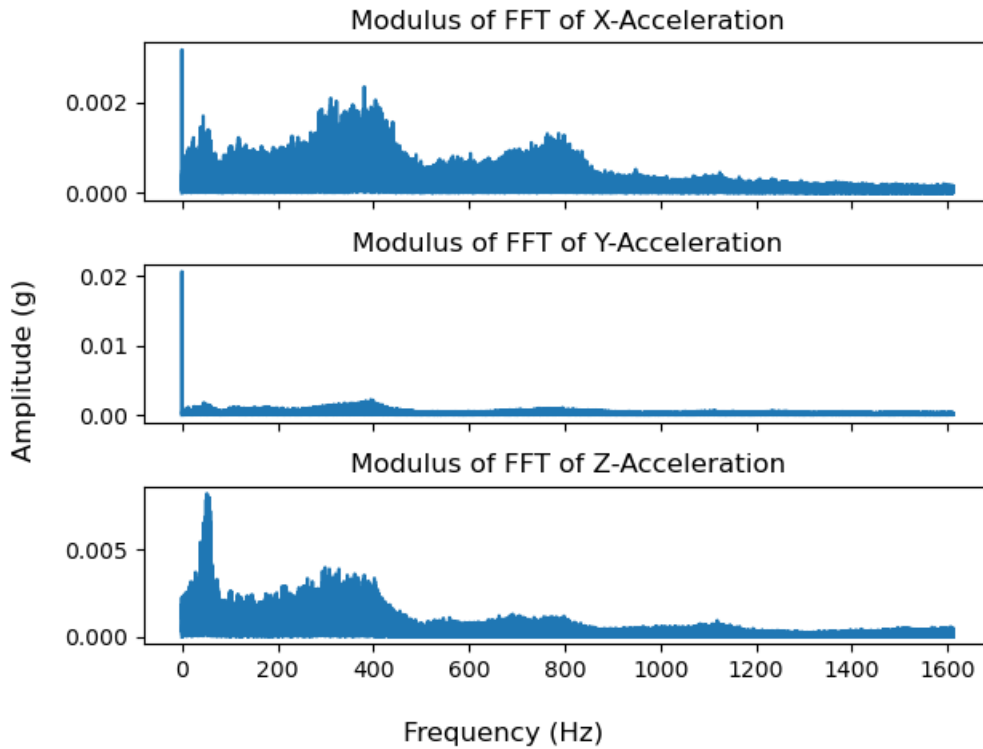


Figure 5.4: x , y , and z dimensions FFTs calculated from acceleration during one station-to-station segment.

frequency resolution, but less precise resolution in time. Similarly, very short time windows have high time resolution with poor frequency resolution. Looking at the spectrograms in Figure 5.5, the first spectrogram (with a shorter time window) highlights potential impulses, while the second spectrogram (with a longer time window) shows clear continuous horizontal lines. Both spectrograms have a shortcoming that is the other's strength that results naturally from the tradeoff between frequency and time resolution, a byproduct of the fixed windowing resolution. Therefore, if there is a desire to investigate signal characteristics of variable resolution, as is the case with this accelerometer data, a different time-frequency signal representation may be beneficial.

5.2 ABA Data Analysis with the Wavelet Transform

This section illustrates the need for, and the application of, the wavelet transform for investigating ABA data. An understanding of the Fourier transform is assumed, and a more rigorous mathematical explanation can be found in Appendix A.

The PyWavelets software package, also known as pywt was used to aid in implementations of various wavelet transform techniques [46].

5.2.1 Motivation For Using the Wavelet Transform

The wavelet transform is a time-frequency representation that lacks the fixed resolution trade-off of the spectrogram. It instead provides a variable resolution representation of a signal in both frequency and time, the

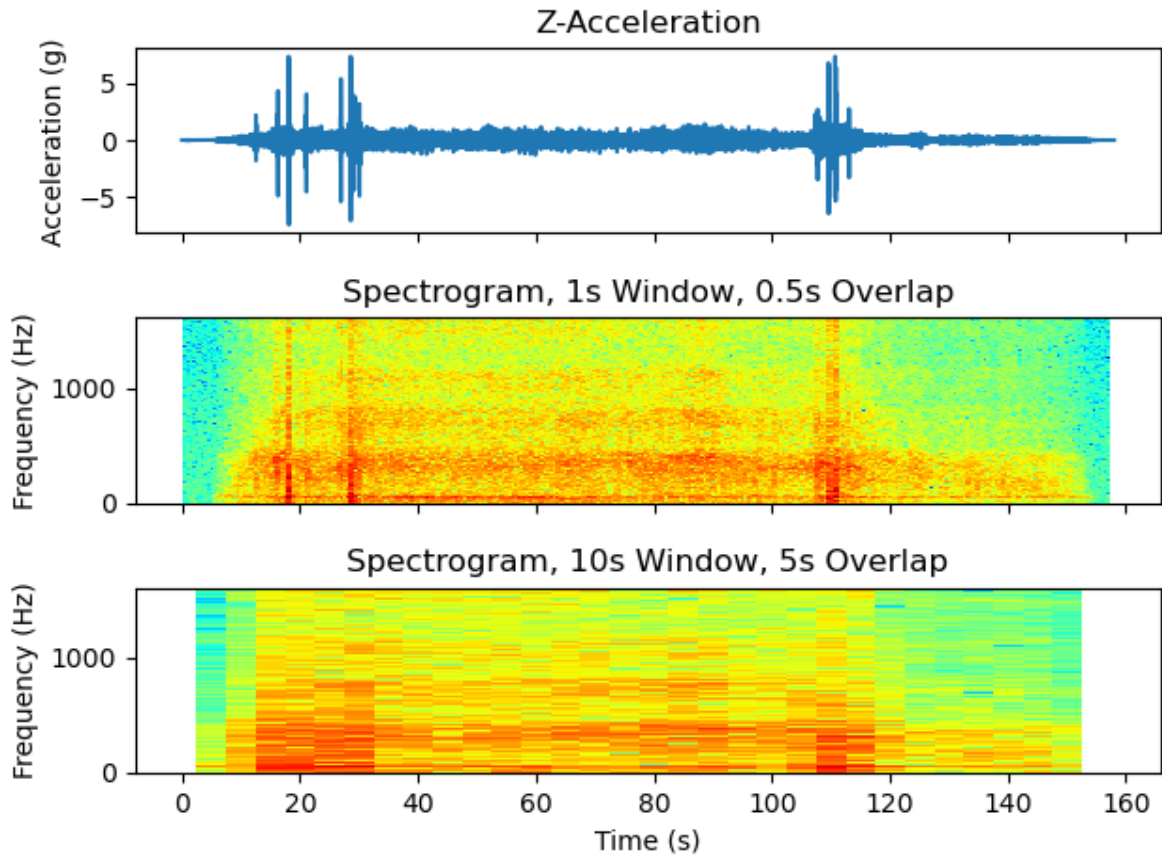


Figure 5.5: z acceleration, spectrogram with 50% overlap and window length of 1 second, spectrogram with 50% overlap and window length of 10 seconds.

precision of which varies as a function of frequency.

A spectrogram calculates frequency components of a signal at different points in time by projecting the signal onto a basis of windowed sinusoids of different frequencies, that are shifted through the function to localize the frequency information [47, 48]. This provides a grid of values with constant time and frequency resolution, that correspond to a particular frequency at a particular time in the signal.

The wavelet transform achieves a similar goal in a unique way. Instead of explicitly using windowed sinusoids, the wavelet transform can be computed on any windowed signal that abides by a set of rules in terms of amplitude, oscillating above and below the t axis, and requiring a mean of zero [44]. Similarly, because there is no necessity for a function explicitly defined by a frequency parameter like a sinusoid (thereby not having a frequency parameter to change), the wavelet is said to instead possess pseudofrequency information associated with the rate of oscillation in the wavelet [46]. The pseudofrequency is controlled by a scale parameter that dilates the wavelet, thereby changing the wavelet's pseudofrequency. This wavelet, similarly to the windowed sinusoid in a spectrogram, is shifted throughout the signal to obtain the pseudofrequency information at a particular time in the signal.

Explicitly, a wavelet basis is constructed from a mother wavelet $\psi(t)$. It is parametrized by a scale parameter

s , and a shift parameter u . Changing these parameters forms a basis in time and frequency according to the following general transformation:

$$\psi_{s,u}(t) = \frac{1}{\sqrt{s}} \psi\left(\frac{t-u}{s}\right). \quad (5.1)$$

Figure 5.6 illustrates the effects of changing these parameters on a sample wavelet. As s increases, the wavelet is stretched across a longer time, which can be thought of as a lower pseudofrequency, as the function oscillates more slowly. As u changes, the centre of the wavelet is moved accordingly.




Figure 5.6: A base mother wavelet $\psi_{s=1,u=0}(t)$ (top), shifted $\psi_{s=1,u=-100}(t)$ (middle), $\psi_{s=2,u=125}(t)$ (bottom). Note that this mother wavelet does not satisfy all the conditions to be a wavelet by definition as it does not have zero mean, however it illustrates the effects of changing s and u .

By manipulating the mother wavelet appropriately, this family of wavelets will span a signal in both frequency and time. Different mother wavelets have different frequency responses, which dictate how frequency information is pulled from a signal. An example of the frequency response of the complex Morlet wavelet, as well as the effect changing the scaling parameter s has on this response, is illustrated in Figure 5.7.

This property means the use of a particular wavelet can be more or less appropriate than another wavelet when applied to particular data. For instance, the Ricker wavelet has been shown to perform well in applications involving seismic data [49], while work has been done with Daubechies wavelets being used to identify characteristics of audible noise data [50]. This gives an added layer of control and flexibility in their application.

Theoretically, wavelets can reach infinite precision in either the time or frequency domains as s approaches ∞ and 0 respectively. Because of this, it is impossible to capture the entire signal without using an infinite




Figure 5.7: Complex Morlet Wavelet in time, and its frequency response in radians at various scales s . Note the shift of centre frequency, as well as reduction of bandwidth as the scale increases. It is a one-sided frequency response because it is a complex function.

number of scales. The solution is to define a scaling function $\phi_{s_0}(t)$ that captures all of the wavelet information for scales greater than a given s_0 .

5.2.2 Continuous Wavelet Transform (CWT)

Comparing the wavelet transform with a spectrogram is the clearest way to identify the differences. Figure 5.8 shows a signal of two complex sinusoidal chirps of different frequencies called Gabor pulses, and shows the difference between a spectrogram representation of this signal, and the complex Morlet wavelet transformation representation. Visualizing the wavelet transform like this across various scales and points in time is called a scalogram. This is a numerical implementation of the theoretical continuous wavelet transform (CWT), meaning s is discretely sampled, and doesn't cover the infinite range of $0 \leq s \leq \infty$.

Figure 5.8 illustrates the effect of the wavelet transform's variable resolution. The first lower frequency chirp has a fairly thin frequency resolution in the scalogram, just like the spectrogram, with a higher peak in the centre of the chirp than the spectrogram does at $t = 0.5$ s. The second higher frequency chirp has less bleed in time in the scalogram when compared with the spectrogram, but it has more bleed in pseudofrequency/scale, illustrating the variable time-resolution trade-off for different s values.

The principle of the wavelet transform provides variable resolution at different frequencies. The reduction in bleed in time at higher frequencies is going to be successful at removing any bleed in time from high frequency

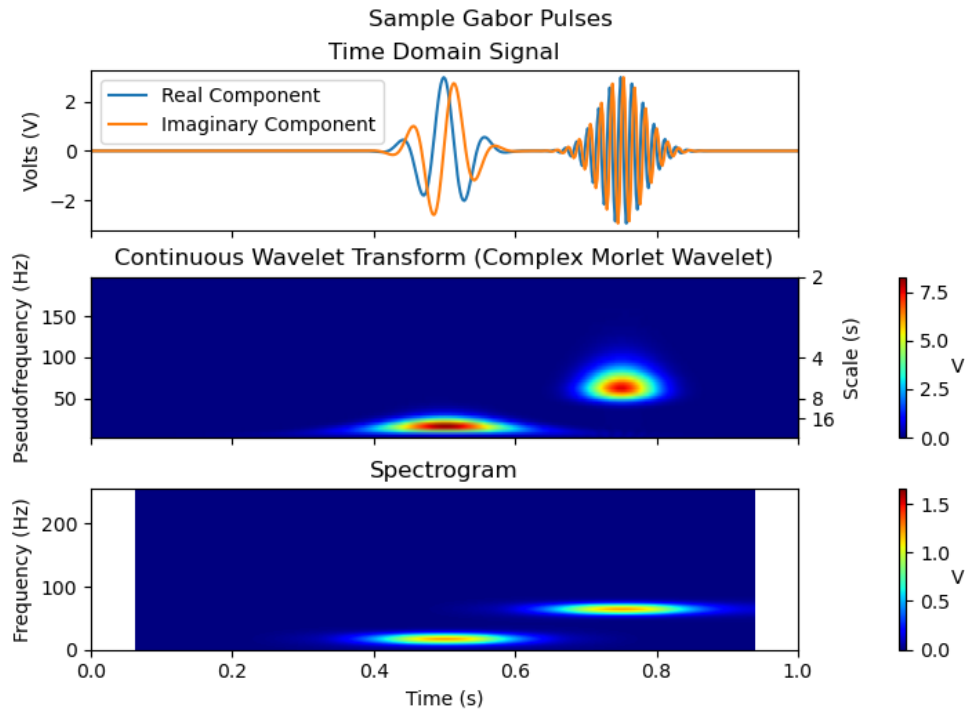


Figure 5.8: Continuous wavelet transform and spectrogram representation of two frequency-modulated Gaussian waves, called Gabor Pulses. Notice the non-linearity of the scale axis of the CWT.

impulses, while the high resolution at low frequency will be adept at pulling out the long, low rumbles like corrugation.

5.2.3 Discrete Wavelet Transform (DWT)

The previous section demonstrates the usefulness of the continuous wavelet transform, however it is an impractical implementation. For example, the y-axes of Figure 5.8 shows that between $s = 2$ and $s = 4$, there is a large range of frequency information encoded in a small range of s values. Meanwhile, a comparatively small range of frequency information exists at scales $s > 16$. Computing one convolution of the continuous wavelet transform on a signal of length N with a filter of length M is $O(MN)$, but for $M \ll N$ it is conventionally bound as a $O(N)$ operation. By extension, computing this on K scale parameters is $O(KN)$.

If s can be selected to minimize redundancy, the complexity of generating the data in a scalogram can be minimized as well. The derivation of the discrete wavelet transform (DWT) addresses the matter of minimizing s selection, while maximizing frequency support of the signal representation.

The DWT algorithm accomplishes this with a cascading filter bank that divides the signal into a high frequency content signal, and a low frequency content signal using two filters $g[n]$ and $h[n]$, followed by downsampling these component signals by a factor of two. The downsampled upper half is denoted as the first level detail coefficients (cD1), while the downsampled lower half is denoted as the first level approximation information (cA1). Then, cA1 is divided into upper and lower frequency content signal (cD2 and cA2) with the same pair of filters and then downsampled, cA2 is then divided into upper and lower frequency content signals (cD3 and




Figure 5.9: Block diagram of DWT cascading filter algorithm. The \downarrow denotes the downsampling operator. $h[n]$ and $g[n]$ are the wavelet and scaling filters. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.

$cA3$) in the same way, and so on. This corresponds to s values that are powers of 2, and the maximum s value is dictated by the length of the mother wavelet M , and the length of the signal N . This algorithm is visualized in Figure 5.9.

The downsampling in the DWT algorithm means each level of filter bank halves the sampling frequency of the previous level, meaning it halves the Nyquist frequency of the coefficients as well. The same filter designed according to the Nyquist frequency can therefore be reused with this downsampled signal to filter the same part of the spectrum relative to the *downsampled* Nyquist frequency, but a different part of the spectrum relative to the *original* Nyquist frequency. For instance, $x[n]$ has a sampling frequency of f_s , and a Nyquist frequency of $f_N = f_s/2$. $g[n]$ filters the data in the range of $[0, f_N/2]$, and then downsamples by a factor of two to produce $cA1$. $cA1$ has a sampling frequency of $f'_s = f_s/2$, and Nyquist frequency of $f'_N = f_N/4$. Applying $g[n]$ to $cA1$ now filters for the frequency range of $[0, f'_N/2 = f_N/4]$. Similarly, $cD1$ is obtained by downsampling the output of filtering $x[n]$ with $h[n]$, which filters in the range of $[f_N/2, f_N]$. Filtering $cA1$ with $h[n]$ filters in the range of $[f'_N/2, f'_N] = [f_N/4, f_N]$ because $cA1$ has the Nyquist frequency of $f'_N = f_N/2$ after downsampling. This recursive filtering and downsampling scheme allows full spectrum coverage, without redundantly sampling s . The result of the DWT with the Haar wavelet when applied to the same Gabor chirp signal is visualized as a scalogram in Figure 5.10.

The scaled wavelet length cannot exceed the length of the signal, meaning the maximum depth of the algorithm is $\log_2(\frac{N}{M})$. This results in a complexity bounded by $O(N \log(\frac{N}{M}))$, as there are $O(\log(\frac{N}{M}))$ convolutions of order $O(N)$. The shifting parameter u does not shift to every point in the signal, but rather is effectively shifted by the 2 at each step after downsampling. This still ensures that every scale covers the entire length of the signal, but localizes features to a section of points proportional to the 2^s .

5.2.4 Stationary Wavelet Transform (SWT)

The DWT improves on the CWT from an implementation feasibility standpoint. However, it suffers from a lack of time-invariance. For a given s value, the wavelet captures frequency information in fixed steps in time proportional to 2^s starting from $t = 0$. This means that the point in time during a signal that a particular waveform occurs dictates how a wavelet will overlap with it. It may happen on the edge of two wavelets, it may




Figure 5.10: Maximum level DWT decomposition of Gabor Waves using the Haar wavelet. Upsampling each level by repeating 2^{level} times allows the transform to be plotted in this way.

be centered in one. This is a potentially undesirable quality to have.

One final improvement to the wavelet transform is the addition of time-invariance at the expense of added redundancy and time complexity. Time-invariance means that generating a scalogram from a signal and shifting the scalogram by k samples will be the same as generating a scalogram from a signal that was shifted by k samples before generating the scalogram, which is not the case for all shifts k with the discrete wavelet transform. This time-invariant version of the transform is known as the stationary wavelet transform (SWT), an example of which using the Haar wavelet when applied to the Gabor pulse signal is shown in Figure 5.11.

Time-invariance in the SWT is accomplished with the same cascading filter methodology as the DWT, however there is no downsampling at each level. Instead, the SWT is implemented with the *algorithme à trous* or “algorithm with holes” [44]. Where the DWT downsamples each level, the SWT inserts holes (zeros) in the filter banks to upsample the filter, denoted with a dot over the filter. For a filter $h[n]$ with coefficients $[c_0, c_1, c_2, \dots]$, $\dot{h}[n]$ has coefficients $[c_0, 0, c_1, 0, \dots]$. Similarly, $\ddot{h}[n]$ has coefficients $[c_0, 0, 0, 0, c_1, 0, 0, 0, \dots]$. The j dots above the filter, correspond to the insertion of $2^j - 1$ zeros. These filters can be convolved at every step of the signal which preserves time invariance, and maintains the sampling frequency of the original signal, while altering the filters to maintain the same frequencies from the DWT downsampling scheme. The block diagram of the algorithm is found in Figure 5.12.

It is possible that the DWT will be sufficient for this application, and the added redundancy of the SWT will not justify the increased cost, especially with signals of length $N > 10^6$ like the ABA data.




Figure 5.11: Maximum level unnormalized SWT decomposition of Gabor Waves using the Haar wavelet. Unlike the DWT, this transform is time-invariant, and does not need to be upsampled at each level.




Figure 5.12: Block diagram of SWT cascading filter algorithm. $h[n]$ and $g[n]$ are the wavelet and scaling filters, and the j dots above the filter denote $2^j - 1$ holes inserted into the filter. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.

5.2.5 Applying the Wavelet Transform to ABA Data

Due to the high sampling frequency and long signal length, visualization of the DWT and SWT when applied to ABA data is infeasible in scalogram form. Throughout this chapter, the level of coefficient decomposition being visualized is explicitly stated, and the coefficient magnitude is displayed as a line.

Firstly, it is essential to select an appropriate wavelet for the application. `pywt` provides 14 wavelet families, for a total of 127 unique wavelets. Of these, 21 can only be used in a CWT, while the remaining 106 are reserved for the DWT. Table 5.1 illustrates the families of wavelets, and some of their properties.

It is not possible to thoroughly investigate every aspect of these wavelets to determine the most appropriate for all accelerometer data sampled. Instead, the investigation focused on understanding sensitivity to high frequency information, sensitivity to low frequency information, and which (if any) demonstrate sensitivity to both.

Family (Size)	Wavelet Members	Parameters	Transform
Haar (1)	haar	None	Discrete/Stationary
Daubechies (38)	db1, ... db38	Number of vanishing moments	Discrete/Stationary
Symlets (19)	sym2, ... sym20	Number of vanishing moments	Discrete/Stationary
Coiflets (17)	coif1, ... coif17	Number of vanishing moments	Discrete/Stationary
Biorthogonal* (15)	bior1.1, ... bior6.8	Number of vanishing moments (synthesis and decomposition)	Discrete/Stationary
Reverse Biorthogonal* (15)	rbio1.1, ... rbio6.8	Number of vanishing moments (synthesis and decomposition)	Discrete/Stationary
Discrete Meyer (1)	dmey	None	Discrete/Stationary
Gaussian (8)	gaus1, ... gaus8	Order of Gaussian derivative	Continuous
Ricker [†] (1)	mexh	None	Continuous
Morlet (1)	mor1	None	Continuous
Complex Gaussian (8)	cgau1, ... cgau8	Order of Complex Gaussian derivative	Continuous
Shannon (1)	shanB-C	Bandwidth (B) and centre frequency (C)	Continuous
B-Spline (1)	fbspM-B-C	Spline order (M), bandwidth (B) and centre frequency (C)	Continuous
Complex Morlet (1)	cmorB-C	Bandwidth (B) and centre frequency (C)	Continuous

* Biorthogonal and reverse biorthogonal families do not provide support for all parameters between 1.1-6.8.

† Ricker wavelet uses an outdated title in the package.


Table 5.1: Wavelets and wavelet families provided by pywt [46].

Due to signal length, it was decided early that the CWT was not feasible across a wide s range leaving the DWT and SWT as possible techniques. In order to understand the sensitivity of a wavelet to various high frequency accelerometer information, 10 random z axis ABA signals were selected, and the highest frequency DWT and SWT was performed with each applicable wavelet. This outputs 20 sets (10 signals, 2 transforms per signal) of level 1 detail coefficients (cD1), which correspond to information roughly in the range of $[f_N/2, f_N]$. The energy per coefficient for a given set of cD1 coefficients of length N ($cD1=[cD1_0, cD1_1, \dots, cD1_{N-1}]$) was calculated as


$$\text{Coefficient Energy} = \frac{\|cD1\|_2}{N} = \frac{1}{N} \sqrt{\sum_{i=0}^{N-1} (cD1_i)^2}, \quad (5.2)$$

which informed the high frequency sensitivity of each wavelet. This not only informed the performance of individual wavelets, but also wavelet families in general. Note that the length of the coefficient array is a function of the signal length, so this energy was normalized by the length of the coefficient array. These energies for every signal and every family were plotted in Figure 5.13 to understand the distribution, as well as the 85th percentile energy for each wavelet. That is, for the Daubechies family, the 85th percentile energy value for db1, db2, and so on are plotted. This process was repeated for the DWT and SWT coefficients.


Note that in these figures, a higher energy does not automatically mean a wavelet or wavelet family is a “better” signal representation. If the wavelet just increases the noise floor, then the increase in signal energy per coefficient is meaningless. The purpose of this technique is to observe any trends in sensitivity among wavelet families to justify further investigation.




(a) Distribution of DWT cD1 energies for each wavelet family.



(c) Distribution of 85th percentile DWT cD1 energies of each wavelet in the family.




(b) Distribution of SWT cD1 energies for each wavelet family.




(d) Distribution of 85th percentile SWT cD1 energies of each wavelet in the family.

Figure 5.13: High frequency energy distributions of cD1 coefficients of wavelet families from 10 different ABA signals.




(a) Distribution of DWT cA1 energies for each wavelet family.


93



(c) Distribution of 85th percentile DWT cA1 energies for each wavelet family.



(b) Distribution of SWT cA1 energies for each wavelet family.



(d) Distribution of 85th percentile SWT cA1 energies for each wavelet family.

Figure 5.14: Low frequency energy distributions of cA1 coefficients from wavelet families from 10 different ABA signals.

From Figure 5.13, the Reverse Biorthogonal family of wavelets showed the greatest range of sensitivities to high frequency information. The Daubechies, Symlet, and Coiflet families, semi-related wavelets designed by Ingrid Daubechies, showed similar clustering throughout the distributions.

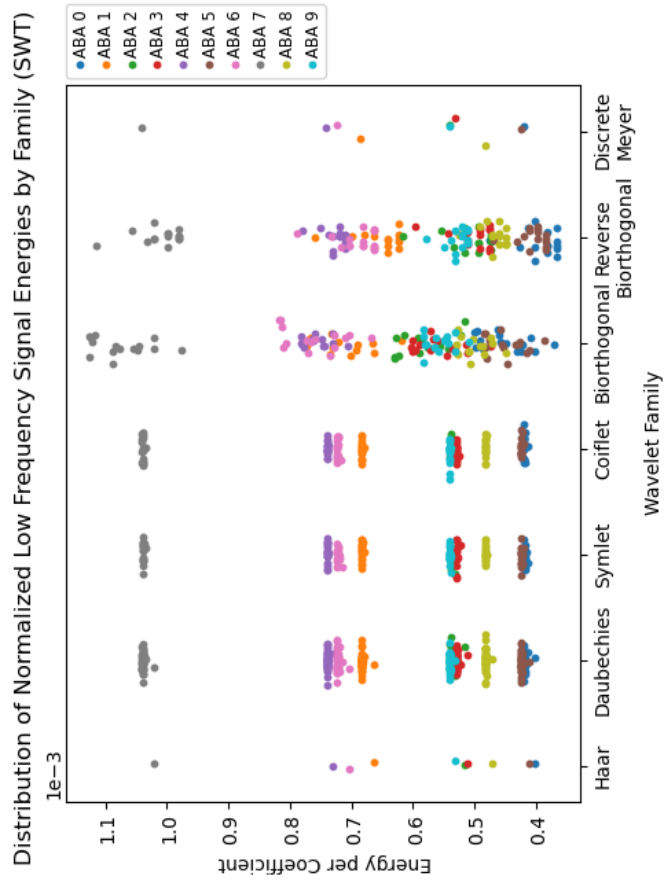
From this information, it was hypothesized that a transform using a Reverse Biorthogonal wavelet basis may be useful for precisely locating high frequency information, which can be used to locate the position of track characteristics like switches and interlocks in accelerometer signals.

There was a similar need to investigate the performance of these wavelet families on low frequency information as well, to determine what transform (if any) is best at that. The same methodology was used, however the first order approximation coefficients (cA1) were used instead of cD1, corresponding to frequency information approximately in the range of $[0, f_N/2]$.

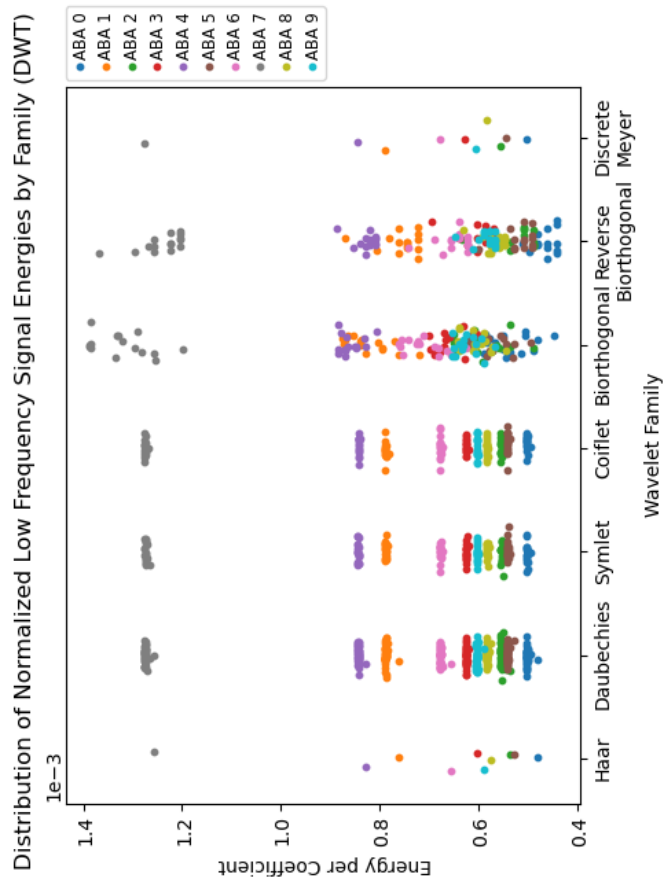
Looking at the coefficient signal energies of these in Figure 5.14, two noteworthy trends appeared. One was the terraced clusters of points at specific energies. This was most present in the Daubechies, Symlet, and Coiflet wavelets, but was similarly present in the Haar and Discrete Meyer wavelets. These clusters corresponded to the 10 ABA segments having their own baseline energy that seems to be captured irrespective of wavelet basis from these families. That is, when each datapoint was colour coded according to its associated ABA segment data as in Figure 5.15, each cluster corresponds to one segment sample. This suggested that any low frequency information is going to be captured to similar degrees if chosen by any wavelet in these families. The Biorthogonal and Reverse Biorthogonal families show similar clustering but not as pronounced as in the other families. When comparing the DWT and SWT analyses, similar levels of clustering were shown in both. This characteristic suggested that the DWT may be sufficient for capturing information appropriately, and the added redundancy/time-invariance of the SWT is unnecessary.

Based on the high frequency signal information, the Reverse Biorthogonal family was likely to showcase the most unique high frequency response. From the low frequency information, all wavelet families seemed to capture low frequency information to similar degrees. As a result, the Reverse Biorthogonal family of wavelets was selected for more in-depth investigation.

The high frequency performance of the Reverse Biorthogonal family was investigated for both the DWT and SWT. There are 15 wavelets in the Reverse Biorthogonal family, which were sorted from highest to lowest RMS energy response to observe any patterns in their response. For both the DWT and SWT, five of the six highest RMS energy responses were `rbio3.1`, `rbio3.3`, `rbio3.5`, `rbio3.7`, and `rbio3.9`. Biorthogonal and Reverse Biorthogonal families generate different wavelet and scaling functions for decomposition and synthesis from the transform coefficients. The naming convention for these families informs the number of vanishing moments in the decomposition filter first, and the synthesis filter second. For instance, the wavelet `rbio3.1` is designed so that the decomposition filter has three vanishing moments, while the analysis filter has one. This pattern suggests that something characteristic of the three vanishing moments in the decomposition filters is dictating this behaviour. These wavelet and scaling decomposition filters are shown in Figure 5.16.



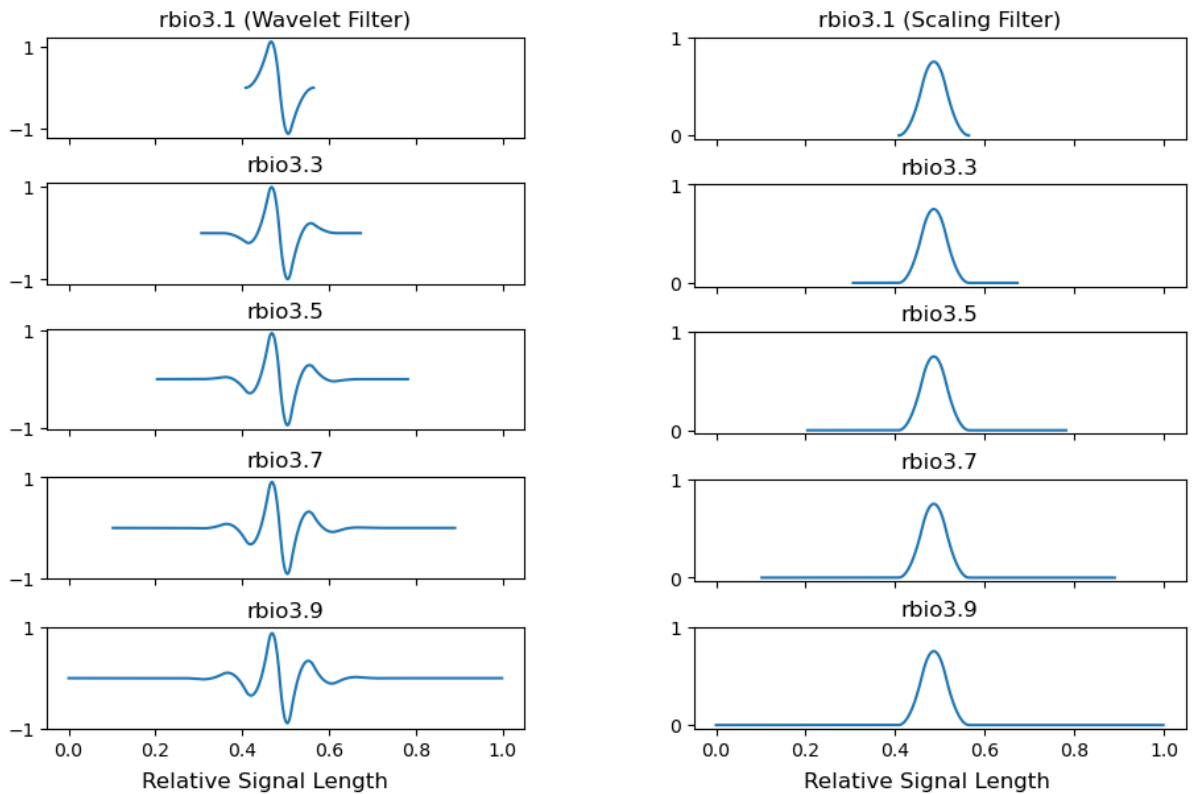
(a) Colour-coded DWT energies according to acceleration sample.



(b) Colour-coded SWT energies according to acceleration sample.

Figure 5.15: DWT and SWT energies colour-coded according to which of the 10 ABA segments it is from. Note the terracing in some families, suggesting that all wavelets in that family are equally sensitive to the cA1 coefficients.

The plots in the figures that follow are colour coded to distinguish the contents of the plots. Blue plots are either raw or reconstructed signals in time. Orange plots are coefficient magnitudes, while green plots are smoothed coefficient magnitudes for the purpose of denoising. Figures 5.17 to 5.22 can be found at the end of this section.



(a) Wavelet decomposition function associated with the rbio3.x subfamily.

(b) Scaling decomposition function associated with the rbio3.x subfamily.

Figure 5.16: Wavelet and Scaling decomposition functions of Reverse Biorthogonal wavelet family with 3 vanishing moments.

Figure 5.17 shows the high frequency information derived from the DWT using the rbio3.1 wavelet when applied to an ABA segment sample. The original signal, as well as its high frequency reconstruction (i.e. just the cD1 coefficients), are visualized in the top two plots, while the magnitude of the coefficients are displayed in the bottom. The large coefficient spikes at approximately $t = 18$ seconds, $t = 205$ seconds, and $t = 391$ seconds were of interest, and an inspection of the signals was undertaken at these points. At $t = 18$ seconds and $t = 391$ seconds, there are two main impacts that occur approximately a second apart. This is different from the signature at $t = 205$ seconds, which is a longer, more violent signal. Speculation based on the track charts suggest that the $t = 18$ and $t = 391$ second impacts are from a switch, a smaller track feature, while the $t = 205$ second impact is from passing through a more extended interlock. These signal characteristics are present in the SWT version of the transforms as well, as shown in Figure 5.18. The consistency between transforms affirmed the suspicion that the DWT was sufficient for this application due to the high sampling rate of the signal. Coefficients from a level j compress information in the neighbourhood of $\text{len}(\text{wavelet}) \times 2^j$

points. The longest wavelet being tested is the `rbio3.9` which has wavelet filter length of 20 points. For $j < 4$ and sampling frequency $f_s = 3225$ Hz, features are encoded into time windows of less than 0.1 seconds, which is sufficiently accurate for this application.

These timestamps were not sufficiently unique from the rest of the signal to confidently identify these track characteristics though. There is a very high noise floor and throughout the signal there are similar peaks that cannot confidently be connected to other track characteristics. Seeing that the expectation is that these switches and interlocks would be present throughout the frequency spectrum, there was a need to inspect `cD1` and `cA1` together. Figure 5.19 presents the raw acceleration signal, the `cD1` and `cA1` coefficients, and the product of these coefficients at each index in the signal in red. This effectively denoises the signal to obtain points where both `cD1` and `cA1` demonstrate high magnitude. This more clearly aligned with the suspected locations in the track. It is possible that this technique could be further improved by decomposing the signal into more subbands to amplify the points where the subbands agree. Again, these results can't be confirmed at this point, but this reinforced the idea that the track information is encoded in the accelerometer data and is obtainable with the wavelet transform.

Precisely locating the low frequency information requires a higher level decomposition, as it can fall in any of the subbands within `cA1`. The maximum depth of decomposition for a given signal and wavelet is defined by `pywt` as the maximum level where at least one coefficient is uncorrupted by edge effects due to signal extension. This can be calculated as:

$$\text{max_level} = \lfloor \log_2 \left(\frac{\text{len}(\text{signal})}{\text{len}(\text{wavelet}) - 1} \right) \rfloor. \quad (5.3)$$

As previously mentioned, 30-100 mm corrugation with vehicle velocities of 60-90 km/h, manifests in ABA data between 170-830 Hz. The sampling frequency of the ABAs is $f_s = 3225$ Hz, meaning the Nyquist frequency is $f_N = 1612.5$ Hz. The frequency band associated with a given decomposition level $1 \leq j \leq \text{max_level}$ can be calculated as follows:

$$\text{cD}j = \left[\frac{f_N}{2^j}, \frac{f_N}{2^{j-1}} \right] \quad (5.4)$$

$$\text{cA}j = \left[0, \frac{f_N}{2^j} \right]. \quad (5.5)$$

From this equation, nearly all corrugation-related information for the estimated parameters will be found in `cDj` for $2 \leq j \leq 4$. For this signal of length 1.3×10^6 , and the longest wavelet `rbio3.9`, the maximum decomposition level is 16, which suggests the information of interest is recoverable from the decomposition. The bandwidth of `cD2` to `cD4` is 101-806 Hz, while the bandwidth of `cA1` is 0-806 Hz. As a result, it is possible that using just `cA1` may provide adequate spectral *magnitude* information at approximately one half of the com-

putational cost. Note that these computational cost savings are in exchange for any spectral resolution within the bandwidth of $cA1$, as everything is compressed into these first level approximation coefficients. However if the spectral noise is roughly constant throughout the signal, then the higher amplitude sections in the approximation coefficients will point to sections of poor rail condition, and investigating multiple levels of detail coefficients in the future may be unnecessary for rail maintenance work.

Investigating the $cA1$ coefficients in Figure 5.19 was deemed too noisy, so these were denoised slightly with a moving average filter of length 100 samples (0.03 second windows). Zooming in on certain sections of this signal illustrates some interesting features. The smoothed coefficients demonstrate some quasi-periodic peaks of varying amplitude, specifically from roughly $t = 254$ seconds to $t = 266$ seconds in Figure 5.20a. Convolving two real, oscillating signals will produce several peaks as they pass in and out of phase with each other, similar to what is seen here, due to the lack of imaginary component for the phase to pass into. This quasi-periodicity possibly indicates that there is a periodic resonance being detected by this transform. Due to the lack of accompanying rail condition data, it cannot be determined whether it is just an artifact of the vehicle dynamics, or if it is a manifestation of the vibration experienced by the vehicle due to rail condition. A signature of similar quasi-periodicity can be found from $t = 116$ to $t = 130$ seconds shown in Figure 5.20b, but not constantly throughout the segment, suggesting that it may only exist at certain points and is not systemic throughout collection.

Splitting $cA1$ into $cD2$, $cD3$, and $cD4$ yielded new information about these quasi-periodic peaks, shown in Figure 5.21 and Figure 5.22. The quasi-periodicity previously seen in $cA1$ seemed to be most closely related to the quasi-periodicity in $cD4$. Additionally, there seems to possibly be some quasi-periodicity in $cD2$ and $cD3$. Again, these cannot be confirmed due to the lack of temporally linked rail condition data, but by inspection, there may be characteristics of the rail being observed by the ABA data.




Figure 5.17: Raw z acceleration signal, the reconstruction from only the cD1 coefficients (DWT), and the cD1 coefficient magnitudes.




Figure 5.18: Raw z acceleration signal, the reconstruction from only the cD1 coefficients (SWT), and the cD1 coefficient magnitudes.


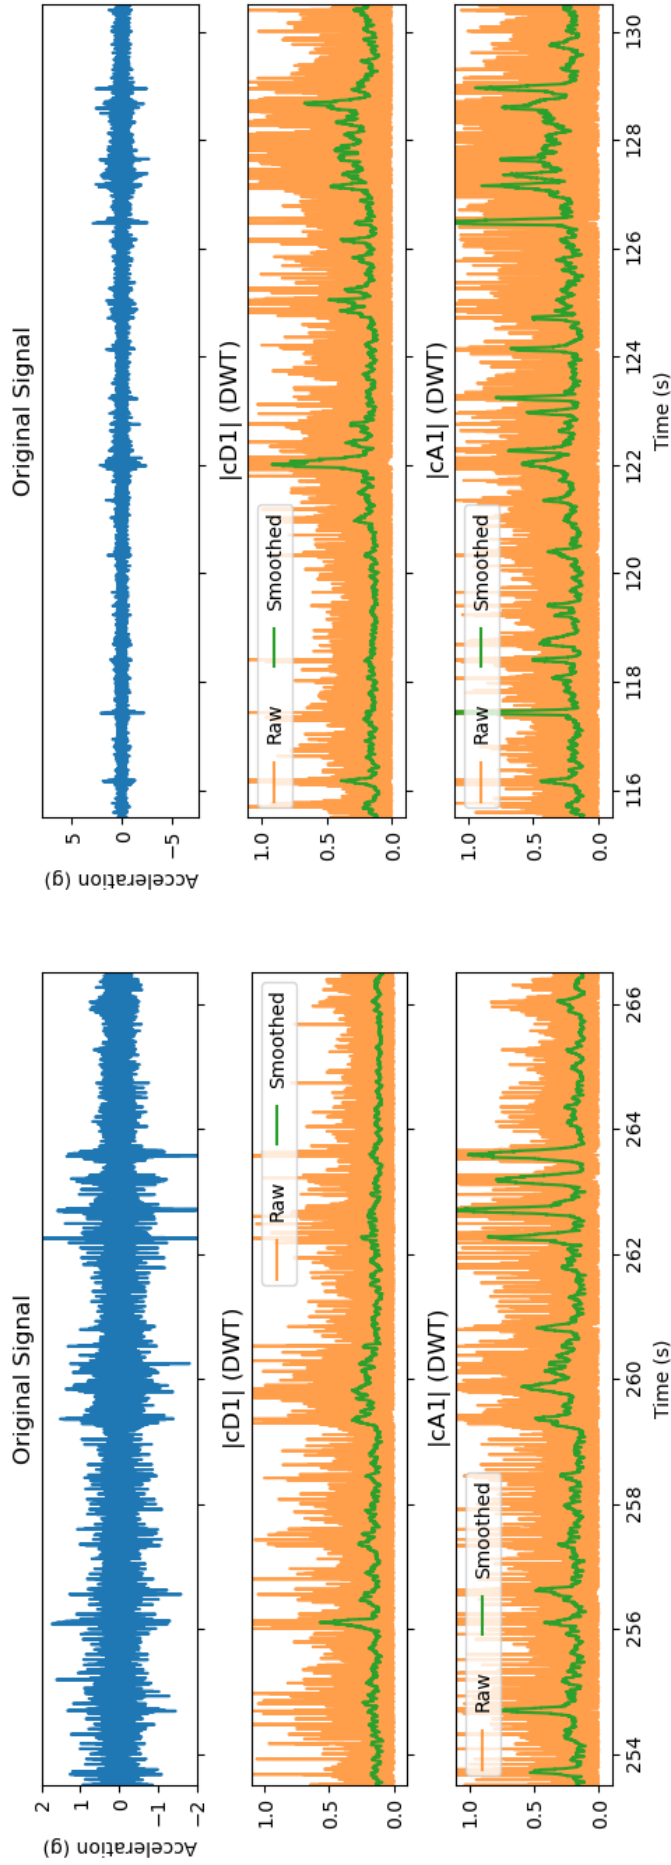


Figure 5.19: Raw z acceleration signal, the untouched as well as smoothed (100 sample/0.03 second window) $cD1$ and $cA1$ coefficients (DWT), and the index-by-index product of the $cD1$ and $cA1$ coefficient magnitudes. Peaks in this plot indicate points where both the low and high frequency information shows a high amplitude, suggesting the location of switches or interlocks.



(a) $|cA1|$ coefficient zoom from $t = 254$ seconds to $t = 266$ seconds.

(b) $|cA1|$ coefficient zoom from $t = 116$ seconds to $t = 130$ seconds.

Figure 5.20: Quasi-periodic peaks in $cA1$ coefficients.




Figure 5.21: $|cD2|$, $|cD3|$, and $|cD4|$ coefficient zoom from $t = 254$ seconds to $t = 266$ seconds.




Figure 5.22: $|cD2|$, $|cD3|$, and $|cD4|$ coefficient zoom from $t = 116$ seconds to $t = 130$ seconds.

5.3 Chapter Summary and Next Steps

ABA data is shown to be a valuable tool due to its non-invasive installation, and robust data collection capabilities. However, the data has its flaws. There is a requirement for a significant amount of post-processing to rectify the usable parts of the data collected, and even more to glean something informative from it. This work has demonstrated the usability of ABA data to monitor rail state qualitatively, but has also demonstrated how further work is required to better understand what is encoded in the data.

Fundamentally, the ABA units used in this work struggle to locate position if the vehicle is moving. Pre-processing locates the ABA when it is stopped in a station, but between stations there is no way to track the location of the device with any certainty. Augmenting the data collection to include positional data could be done in two ways, both with drawbacks. The addition of a GNSS module to the ABA would theoretically solve this problem, but GNSS is unreliable in dense urban areas due to the presence of ample electronic interference, as well as physical obstructions such as tunnels, bridges, and tall buildings which can interfere with GNSS communication. An alternate solution would be to connect the ABA to on-board sensors and gauges like the odometer and speedometer, however this makes the ABA significantly more invasive, and requires proper interfacing with every vehicle in the fleet, at every property a company may service. These issues leave dead reckoning as the only option at this stage.

Identifying important characteristics of the ABA data with the wavelet transform, a signal processing technique which provides variable resolution in time and frequency, has been shown to be qualitatively possible. This preliminary work showed that the application of the DWT is suited to this information extraction, and demonstrates sensitivity to the different kinds of artifacts present in a vibration signal. However, this work has not demonstrated empirical connection between vibration data and observable rail condition due to the lack of paired data in this study. As such, the outcomes of this stage of research have been purely qualitative. This work has shown that the variable resolution of the transform lends nicely to this problem due to the nature of information encoded in ABA signals. The ability to identify sharp impulses with short time durations helps identify spatial track features like switches, while the identification of long time rumbles due to maintenance issues like corrugation is equally desirable.

Continuation of this work is necessary to understand the underlying cause of certain ABA behaviours, as well as to improve the software responsible for locating the vehicle in space. One possible solution for improving the location certainty is the addition of other data streams like rotation and velocity data, and the unification of these data streams through a Kalman filter [52]. This is a signal processing technique meant to combine sensor data with dynamic models of what is being monitored to maximize certainty in an estimate. Additionally, there is space for machine learning with this ABA data. Convolutional Neural Networks (CNNs) could identify encoded sequences derived from the wavelet transform, and compare them with measured track condition to derive a relationship. This can help property maintainers identify the state of rail from just the vibration data,

at a specific spatial location identified from ABA data.

Chapter 6

Conclusion and Future Work

This chapter presents a summary of key findings and contributions of this thesis, and makes recommendations for future extensions to this research.

6.1 Summary of Key Findings and Contributions

This thesis presented the application of various signal processing techniques to real world rail transportation data, with the intent of supplementing present day property maintenance with methodologies and tools not widely used in industry. The ability for property maintainers to incorporate a wide array of data streams which will provide more accurate and up to date information on the condition of a rail system will lead to an increase in ride performance quality, as well as streamlining the maintenance workflow.

Corrugation and rail surface roughness are the predominant focus for this work. In terms of assessing the surface of the rail, a long spatial signal obtained from a CAT measurement is very informationally dense. RCI_x aims to condense the condition of a long segment of rail into one representative sample, with the intention of simplifying rail roughness assessment. By focusing on short pitch corrugation, the representative block RMS technique demonstrated growth as a function of wear in a repeatable way. This was extended with $RCI_x(\lambda)$ to generalize from property to property by filtering for the corrugation needs of a particular location. It was demonstrated that selecting a proper filtration band improves the performance of $RCI_x(\lambda)$ by up to 22.4%, and an ample waiting period after grinding influenced an increase of the median ΔRCI_x value by up to 18 μm . The analysis of corrugation as a spatial signal was shown to provide insight for property maintainers to better understand the state of their rail.

It is known that measuring corrugation directly is a time-consuming and intensive process, so measuring corrugation indirectly if possible is desirable. This work has demonstrated that wayside noise can be related to corrugation as a byproduct of vehicle dynamics with rough rail, suggesting that rail roughness can be obliquely observed by measuring wayside noise. LA_{eq} was shown to mimic corrugation growth, but also demonstrated

problems in less reliable systems, due to its potential sensitivity to other noise generating phenomena. This led to the definition of $Lv_{\text{eq}}(\lambda, v)$, a noise weighting meant to focus on corrugation of a particular wavelength for noise measurements where the passing vehicle was travelling at a particular velocity. It was demonstrated that this more precise wayside noise weighting scheme showed a reliable response to grinding, and successfully echoed the LA_{eq} growth pattern. This thesis showed that the relationship between LA_{eq} and RCI_x could be explained as a linear relationship with R^2 values between 0.64 and 0.97. $Lv_{\text{eq}}(\lambda, v)$ showed an improvement over LA_{eq} at a property with other underlying phenomena and demonstrated an accurate response to grinding rate of 100%, as opposed to LA_{eq} which showed a 66.6% rate. This usage of wayside noise spectra as proxy for corrugation monitoring gives property maintainers a useful tool for roughly measuring corrugation at a location with as little intervention as possible.

An extension of indirect corrugation measurement is the use of on-board vibration data to measure the effect of rail condition on vehicle dynamics. This work demonstrated that ABA data provides extensive information that needs to be clarified through signal processing techniques after measurement. Due to the lack of displacement information, the need to spatially identify constant track characteristics such as interlocks and switches is desirable. It was shown that the wavelet transform is a tool that has the potential to capture low frequency corrugation information, as well as full spectrum impulses to locate the vehicle in space. This technique showed a qualitative difference in characteristics shown throughout vibration signals, suggesting it was sensitive to various vibrations experienced by the train. Due to a lack of paired data, it is not possible to quantify these sensitivities, but the analysis of the wavelet transform coefficients showed promise. The wavelet transform as a tool provides sufficient time-frequency resolution trade-off to identify the different characteristics present in the vibration data in a way that more traditional time-frequency analyses like the spectrogram cannot due to the variable resolution of the transform.

6.2 Suggestions for Future Work

The work presented in this thesis leaves significant room for improvement and iteration. Some of the questions left unanswered due to the scope of this work may become easier to answer given the presented work.

6.2.1 Corrugation and Rail Roughness

- Repeat the RCI_x and $RCI_x(\lambda)$ work at a third property for further validation. There is desire to test in a property that has issues with both short- and long-pitch corrugation, to better understand the comparative sensitivity of the metric to different corrugation bandwidths.
- Incorporate some level of unsupervised or semi-supervised learning methodology to quantify corrugation for a given measurement. The sequential aspect of a corrugation signal suggests that something structurally similar to a recurrent neural network (RNN) or a Long Short-Term Memory (LSTM) model may be

applicable, or co-opting techniques used by large language models (LLMs) or natural language processing (NLP) models to better understand sequential characteristics in a corrugation sample.

- Incorporate computer vision analysis into this work by using the V-CAT. This would incorporate visual information with respect to profile as well as longitudinal corrugation data to better understand the overall wheel-rail contact shape.
- This work assumes that corrugation grows as a function of usage, which is shown to broadly be true. However, it assumes that all tonnage is equivalent regardless of how it is applied (i.e. ten 60-tonne vehicles impart the same wear as twenty 30-tonne vehicles). There may be value in investigating what effect the tonnage of a single load instance has on corrugation growth as a function of accumulated tonnage. It may be the case that corrugation growth is a function of both cumulative MGT, but also average vehicle load size.

6.2.2 Noise

- Incorporate wayside $Lv_{eq}(\lambda, v)$ methodology into an on-board noise measurement systems. By mounting noise measurement hardware to the outside of the vehicle, then noise data can be used to monitor corrugation by proxy throughout the system. Assuming the vibrations of the vehicle do not affect the ability for noise data to be collected reliably, this would extend the noise data collection capabilities of property maintainers.
- Investigate extending $Lv_{eq}(\lambda, v)$ methodology to other rail phenomena such as those in Table 4.3. For reference, these were rolling noise, flat spots on wheels, ground-borne rail vibrations, and flanging noise. If the $Lv_{eq}(\lambda, v)$ methodology can capture these noises from simple spectrum filtering, then this further improves the robustness of the metric to extend beyond just corrugation detection.
- Investigate the effect of ambient sampling conditions on the noise measurement. This includes, but is not limited to, weather, climate, surrounding structure density, nearness to water and forest, and nearness to other traffic. If, for instance, winter weather dampens sound through snow on the ground, then noise will be underreported by the SLM.
- This thesis presents its findings largely through comparison of deterministic scalars due, in part, to the lack of data availability. An alternative methodology when comparing corrugation with wayside noise would be to treat them as stochastic processes. This motivates an investigation of the relationship between corrugation and the distribution of wayside noise samples, as opposed to selecting a representative deterministic sample of the corrugation outputs, and the overall distribution of spectrum information from these samples can be used to summarize noise. This extends to obtaining an unbiased representation of the data. For instance, the data collected at Property A for instance, all 14-24 wayside noise samples were

collected on the same day, meaning they all would be biased according to the weather of the day (i.e. humidity, wind, snow on the ground). Instead, this can be mitigated by sampling 2-4 times per day, once per day in the week surrounding the corrugation measurement date. This is more demanding of the transit agency collecting the data, but it may prove valuable if it removes a potential source of bias in data collection.

6.2.3 Vibration

- Incorporate additional data streams such as gyroscope, GNSS, speedometer, and odometer data, and unify these data streams with Kalman filtering to improve the measurement confidence. It is believed that the data throughput can be reduced if necessary as the current sampling rate is very high. This is arguably the highest priority continuation, as the limitations of this work stem from a lack of spatial information and confidence.
- Perform a more rigorous wavelet testing regimen to identify the most appropriate wavelet for vibration data. This can include the construction of customized wavelets based on properties of the suspension as well as the IMU characteristics. This extends to the possible application of using the CWT on a minimal number of scale values to focus on specific frequencies/pseudofrequencies in the data spectrum.
- Perform paired analysis with spatially and temporally linked corrugation data or other rail surface defect data. This would inform whether the oscillating peaks present in the lower frequency wavelet transform components were due to corrugation, or simply the dynamics of the wheel-rail interface. Additionally, these peaks may not have been due to corrugation, but due to another rail surface defect like squats. Having spatial information about the density of squats throughout the system would allow the investigation of ABA sensitivity to squats and squat-type defects.
- As indicated earlier, this thesis presents its findings largely through comparison of deterministic scalars due, in part, to the lack of data availability. As with noise, an alternative methodology when comparing corrugation with vibration would be to treat them as stochastic processes. For vibration data, assuming the issues with spatially locating the data do not exist, the vibration information collected from multiple vehicles can be amalgamated to summarize the vibration response throughout a particular station-to-station segment.

Appendix A

Wavelet Transform Theory

This appendix provides a more thorough overview of the wavelet transform. The presentation herein is largely a summary of [44], specifically Chapter 4, Chapter 5, and Chapter 7. This is by no means a complete overview of the wavelet transform, but it provides enough mathematical background to understand the basic principles and characteristics of the wavelet transform.

APPENDIX A.1 CONTINUOUS DOMAIN

This section focuses on the theory underpinning the continuous wavelet transform. This section defines the constraints of a wavelet family, and explicitly introduces the definition of the continuous wavelet transform. It goes on to prove the variable resolution of the wavelet transform, as well as the low frequency approximation of a wavelet family known as the scaling function.

A.1.1 CONSTRUCTING A FAMILY OF WAVELETS

The short-time Fourier transform provides a signal representation in time and space by projecting a signal onto a family of windowed sinusoids. These sinusoids have uniform resolution in time as well as frequency. The wavelet transform is an alternative time-frequency representation of a signal, which projects the signal onto a family of wavelets which have variable resolution.

The continuous wavelet transform is a linear time-frequency transform meant to determine the presence of a family of waveforms called wavelets at various points in a signal $f(t)$'s duration. These waveforms individually may have very small time and/or frequency span, but collectively will provide an accurate representation of $f(t)$ throughout its duration.

This transform is designed to project a function in time $f(t)$ onto a family (or set) of wavelets $\{\psi_\gamma\}$. Gamma

denotes the arbitrary number of parameters for defining the family. Each wavelet in the family is a function $\psi \in \mathbf{L}^2(\mathbb{R})$, referred to as the mother wavelet, and each wavelet has zero average:

$$\int_{-\infty}^{+\infty} \psi(t) dt = 0. \quad (\text{A.1})$$

Each wavelet $\psi(t)$ is normalized such that $\|\psi\| = 1$, and is centred at $t = 0$ (i.e. has a time support across $[-t_0, t_0]$). We denote the Fourier transform of $\psi(t)$ as $\hat{\psi}(\omega)$. The wavelet admissibility condition requires that:

$$C_\psi = \int_0^{+\infty} \frac{|\hat{\psi}(\omega)|^2}{\omega} d\omega < \infty. \quad (\text{A.2})$$

Guaranteeing this integral to be finite necessitates that $\hat{\psi}(0) = 0$, and that $\hat{\psi}(\omega)$ is continuously differentiable [44, 53].

It is desirable that a wavelet have limited time and frequency support, necessitating the existence of the entire family of wavelets. As a result, a wavelet is expected to be zero or near-zero outside of a time window $-t' < 0 < t'$. Each wavelet in the family is obtained by parametrizing the mother wavelet according to a scaling parameter s and a translation parameter u :

$$\psi_{u,s}(t) = \frac{1}{\sqrt{s}} \psi\left(\frac{t-u}{s}\right), \quad u \in \mathbb{R}, \quad s \in \mathbb{R} > 0. \quad (\text{A.3})$$

So long as a mother wavelet satisfies these conditions, a valid family of wavelets can be constructed from it. Some known continuous wavelet families include the complex Morlet wavelet family, the complex Gaussian wavelet family, and the Ricker wavelet family.

A.1.2 DEFINING THE CONTINUOUS WAVELET TRANSFORM

The wavelet transform of a signal $f(t) \in \mathbf{L}^2(\mathbb{R})$ at time u and scale s is given by the following inner product:

$$Wf(u, s) = \langle f, \psi_{u,s} \rangle = \int_{-\infty}^{+\infty} f(t) \frac{1}{\sqrt{s}} \psi^*\left(\frac{t-u}{s}\right), \quad (\text{A.4})$$

where ψ^* denotes the complex conjugate of ψ .

Equation (A.4) can be rewritten as the following convolution product:

$$Wf(u, s) = \int_{-\infty}^{+\infty} f(t) \frac{1}{\sqrt{s}} \psi^*\left(\frac{t-u}{s}\right) = f * \bar{\psi}_s(u), \quad (\text{A.5})$$

where

$$\bar{\psi}_s(t) = \frac{1}{\sqrt{s}} \psi^*\left(\frac{t-u}{s}\right), \quad (\text{A.6})$$

and $f * g$ denotes the convolution between two functions f and g .

The Fourier transform of $\bar{\psi}_s(t)$ is:

$$\widehat{\bar{\psi}}_s(\omega) = \sqrt{s}\hat{\psi}^*(s\omega), \quad (\text{A.7})$$

where

$$\hat{\psi}^*(\omega) = \hat{\psi}(-\omega). \quad (\text{A.8})$$

Both Equation (A.7) and Equation (A.8) can be derived from the Fourier transform time-scaling and complex conjugate properties respectively.

The Fourier transform defines the zero frequency spectrum component of a signal $g(t)$ to be $\hat{g}(0) = \int_{-\infty}^{+\infty} g(t)dt$. Equation (A.1) enforces that this condition must be 0 for $g(t) = \psi(t)$. Because the 0 frequency component is 0 as necessitated by Equation (A.2), the transfer function $\hat{\psi}$ defines a band-pass filter. Therefore, the convolution Equation (A.4) computes the wavelet transform at all times u with dilated band-pass filters according to the scaling parameter s . Figure A.1 illustrates the effect of changing s on the Ricker wavelet in time, as well as its frequency response.

Additionally, Figure A.1 shows the change in wavelet resolution as a function of s . Both the centre frequency and bandwidth of the passband are proportional to $1/s$, while the oscillating support of the wavelet in time is proportional to s . This indicates that the resolution in time and the resolution in frequency are inversely proportional. This is not unlike the short-time Fourier transform (STFT), where there is a time-frequency resolution trade-off. In the case of the wavelet transform however, there is variable resolution for different parts of the spectrum due to the bandwidth scaling of s .

A.1.3 ANALYTIC WAVELETS, AND TIME-FREQUENCY RESOLUTION

The Ricker wavelet is an example of a real wavelet. It has no imaginary component in time, and as a result has a two-sided frequency response. Complex or analytic wavelets also exist, having one-sided frequency responses. Explicitly, a function or wavelet is said to be analytic if its Fourier transform is zero for negative frequencies. For example, a sinusoidal function $f(t) = \cos(\omega_0 t)$ has a two-sided Fourier transform $\hat{f}(\omega) = \pi[(\delta(\omega - \omega_0) + \delta(\omega + \omega_0))]$. Compared to a complex exponential with the same angular frequency $f(t) = e^{j\omega_0 t}$ that has the Fourier transform $\hat{f}(\omega) = 2\pi\delta(\omega - \omega_0)$, which is strictly zero for negative frequencies, and has twice the amplitude in the positive frequencies. Similar to the Fourier transform with complex exponentials, an analytic wavelet can separate the phase and amplitude information of a signal. Additionally, analytic wavelets help to prove the inverse time-frequency resolution [54].

Let ψ denote a wavelet that is explicitly analytic. The wavelet transform remains as presented in Equation (A.4). As ψ is centred at $t = 0$, $\psi_{u,s}$ is centered at $t = u$. It can be shown with substitution $v = \frac{t-u}{s}$ and $dv = \frac{dt}{s}$ that




Figure A.1: A sample Ricker wavelet with its frequency response. The effect of changing s can be seen in both domains. Note that as $s \rightarrow \infty$, the peak of the frequency response approaches 0. By utilizing all $0 < s$, the entire frequency spectrum $(-\infty, \infty)$ rad/sec of a signal can be represented.

$$\begin{aligned}
 \int_{-\infty}^{+\infty} (t-u)^2 |\psi_{u,s}(t)|^2 dt &= \int_{-\infty}^{+\infty} (vs)^2 |\psi(v)|^2 \left(\frac{1}{\sqrt{s}}\right)^2 dv s \\
 &= \int_{-\infty}^{+\infty} s^2 v^2 |\psi(v)|^2 dv \\
 &= \int_{-\infty}^{+\infty} s^2 t^2 |\psi(t)|^2 dt \\
 &= s^2 \sigma_t^2.
 \end{aligned} \tag{A.9}$$

Here, $\sigma_t^2 = \int_{-\infty}^{+\infty} t^2 |\psi(t)|^2 dt$ defines the energy variance in time of the wavelet. This suggests that there is a baseline energy variance in time σ_t^2 determined by the wavelet family, and this variance increases as s does.

Let $\hat{\psi}$ denote the Fourier transform of the analytic wavelet ψ . Because ψ is analytic, its Fourier transform is 0 for negative frequency. Therefore, its centre frequency η is defined as the expected value of its Fourier transform:

$$\eta = \frac{1}{2\pi} \int_0^{+\infty} \omega |\hat{\psi}(\omega)|^2 d\omega. \tag{A.10}$$

According to Fourier transform properties, the Fourier transform of $\psi_{u,s}$ dilates $\hat{\psi}$ by a factor of $\frac{1}{s}$ with an imparted phase according to the shift u . The centre frequency of $\hat{\psi}_{u,s}$ is therefore $\frac{\eta}{s}$. The frequency energy spread around $\frac{\eta}{s}$ can be calculated similarly to the time energy spread about u .

$$\frac{1}{2\pi} \int_0^{+\infty} \left(\omega - \frac{\eta}{s} \right)^2 |\hat{\psi}_{u,s}(\omega)|^2 d\omega = \frac{\sigma_\omega^2}{s^2}, \quad (\text{A.11})$$

where $\sigma_\omega^2 = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \omega^2 |\hat{\psi}(\omega)|^2 d\omega$. This indicates that the frequency spread is defined by the wavelet and decreases as s increases.

These spreads define boxes centered at $(u, \frac{\eta}{s})$ for a wavelet $\psi_{u,s}$. These boxes, referred to as Heisenberg boxes [55, 56], are of size $s\sigma_t \times \frac{\sigma_\omega}{s}$ in the time and frequency dimensions respectively. These Heisenberg boxes denote the resolution boundaries of a wavelet in both the time and frequency domains, indicative of the range of time and frequency information encoded in the transform centered at $(u, \frac{\eta}{s})$. Because the dimensions of a Heisenberg box are defined as $s\sigma_t \times \frac{\sigma_\omega}{s}$, the area of all Heisenberg boxes derived from the same mother wavelet are equal to $\sigma_t \sigma_\omega$ irrespective of s , but the resolution in time and frequency, as well as the centre frequency of the box, is dependent on s . Because of this fixed area of a box, it is impossible to narrow both the frequency and time ranges at the same time, leading to this variable resolution at different frequencies.

A.1.4 THE SCALING FUNCTION

Representing a signal with all $s > 0$ would require an infinite number of wavelets and so is unfeasible. This is solved by the use of a scaling function ϕ for low-frequency approximation of $f(t)$ at scale $s \geq s_0$. The modulus of its Fourier transform is defined as

$$|\hat{\phi}(\omega)|^2 = \int_1^{+\infty} |\hat{\psi}(s\omega)|^2 \frac{ds}{s} = \int_\omega^{+\infty} \frac{|\hat{\psi}(\xi)|^2}{\xi} d\xi. \quad (\text{A.12})$$

It can be verified that $\|\phi\| = 1$, and Equation (A.2) dictates that

$$\lim_{\omega \rightarrow 0} |\hat{\phi}(\omega)|^2 = C_\psi. \quad (\text{A.13})$$

The low frequency approximation of a function $f(t)$ using the scaling function $\phi(t)$ is calculated as:

$$Lf(u, s_0) = \langle f, \phi_{u,s_0} \rangle = \int_{-\infty}^{+\infty} f(t) \frac{1}{\sqrt{s_0}} \phi^* \left(\frac{t-u}{s_0} \right) dt = f * \bar{\phi}_{s_0}(u) \quad (\text{A.14})$$

The Ricker wavelet $\psi(t)$, its frequency response $\hat{\psi}(\omega)$, and its scaling function frequency response $\hat{\phi}(\omega)$ can be seen in Figure A.2.




Figure A.2: The Ricker wavelet $\psi(t)$, $\hat{\psi}(\omega)$, and its $\hat{\phi}(\omega)$. Note that the scaling function's time domain representation $\phi(t)$ exists but is not visualized due to having to be generated from the inverse Fourier transform of $\hat{\phi}(\omega)$. The $\phi(t)$ associated with the Ricker wavelet does not have a closed form expression to my knowledge.

APPENDIX A.2 DISCRETE DOMAIN

This section formalizes the theory underpinning the discrete wavelet transform and its time-invariant sister transform the stationary wavelet transform. Neither of these transforms are simply a discretized version of the continuous wavelet transform for a linear selection of values for s . Instead, the scales used are defined according to a cascading series of filters with appropriate downsampling to highlight the an entire frequency spectrum while maintaining variable resolution.

A.2.1 DISCRETIZING WAVELET THEORY

The resolution and centre frequency of the continuous wavelet transform is determined by the scaling parameter $s > 0$. While in the continuous domain s can take on any positive real value, discretizing the transform in a consistent fashion requires limiting the value of s . In particular, the discrete wavelet transform adopts a dyadic resolution selection, where the scaling parameter s is discretized as powers of 2, explicitly $s = 2^j$, $j = 1, 2, 3, \dots$. Therefore, for a discrete mother wavelet $\psi[n]$, the discrete wavelet scaled by 2^j and normalized appropriately is defined as:

$$\psi_j[n] = \frac{1}{\sqrt{2^j}} \psi\left[\frac{n}{2^j}\right]. \quad (\text{A.15})$$

The wavelet transform of a discretely sampled signal $f[n]$ of length N can be calculated using a circular convolution as follows:

$$Wf[n, 2^j] = \sum_{m=0}^{N-1} f[m] \psi_j^*[m-n] = f \circledast \bar{\psi}_j[n], \quad (\text{A.16})$$

where $\bar{\psi}_j[n] = \psi_j^*[-n]$.

By sampling the s parameter discretely in this way, we obtain centre frequencies associated with a level j . Let $\hat{\psi}[k]$ denote the discrete Fourier transform of $\psi[n]$. This discrete frequency domain signal denotes a bandpass filter analogous to the continuous domain. The centre frequency η for level $j + 1$ is half the centre frequency of the level j according to a discretized version of Equation (A.10) when dilating the frequency response:

$$\eta = \frac{1}{2\pi} \sum_{k=0}^{\infty} k |\hat{\psi}[k]|^2. \quad (\text{A.17})$$

A discrete scaling function $\phi[n]$ provides a low frequency representation for scales $j > J$, defined as:

$$\phi_J[n] = \frac{1}{\sqrt{2^J}} \phi\left[\frac{n}{2^J}\right]. \quad (\text{A.18})$$

Similarly, the low frequency approximation is calculated using a circular convolution:

$$Lf[n, 2^J] = \sum_{m=0}^{N-1} f[m] \phi_J^*[m-n] = f \circledast \bar{\phi}_J[n], \quad (\text{A.19})$$

where $\bar{\phi}_J[n] = \phi_J^*[-n]$.

Together, $\psi_j[n]$, $j = 1, 2, 3, \dots, J$ and $\phi_J[n]$ form the basis defining the passbands of a discretized wavelet transform up to the Nyquist frequency f_N . An example of how these passbands may manifest in a spectrum is shown in Figure A.3 for $J = 3$.




Figure A.3: Passbands for a 3 level discretized wavelet decomposition. f_N denotes the Nyquist sampling rate of the original signal.

For a signal $f[n]$ of length N , and a wavelet $\psi[n]$ of length M , this dyadic sampling methodology holds for all levels $j < J'$ such that the dilated $\psi_j[n]$ does not exceed the length of the signal $f[n]$. That is, $N \geq 2^{J'-1}M$. The decomposition can be continued with circular convolution, but if the wavelet length exceeds the signal length then the entire decomposition for that level becomes corrupted due to signal wrapping.

A.2.2 SUBBAND ENCODING

The discretized wavelet transform can be implemented using cascading filter banks in an effort to improve efficiency of the algorithm, similar to how the Fast Fourier Transform (FFT) improves on the efficiency of the Discrete Fourier Transform (DFT) [57].

$\hat{\psi}_1$ in Figure A.3 has a passband that covers the entire upper half of the frequency spectrum $[f_N/2, f_N)$. Because we know the transform covers the entirety of the spectrum, it follows that $\hat{\psi}_2$, $\hat{\psi}_3$, and $\hat{\phi}_3$ collectively cover $[0, f_N/2)$. Therefore, if we were to divide the frequency spectrum in half at $f_N/2$, we would separate $\hat{f}_h = [\hat{\psi}_1]$ from $\hat{f}_l = [\hat{\psi}_2, \hat{\psi}_3, \hat{\phi}_3]$. It follows that f_h and f_l exist as the time-domain representations of signals with frequency responses \hat{f}_h and \hat{f}_l respectively. The intention of the discrete wavelet transform is to reuse the filters that decomposed the signal into \hat{f}_l and \hat{f}_h , to decompose \hat{f}_l into its own high and low frequency information.

It is clear that \hat{f}_h can be obtained from \hat{f} with a highpass filter $h[n]$ with a lower cutoff at $f_N/2$, as the Nyquist frequency necessarily provides an upper cutoff at f_N . Similarly, \hat{f}_l could be obtained with a lowpass filter $g[n]$ with an upper cutoff at $f_N/2$. It is intuitive that in order to recover $\hat{\psi}_2$ from f_l you would need a bandpass filter cutting off at $f_N/4$ and $f_N/2$, however f_l already has the upper cutoff at $f_N/2$ when it was convolved with $g[n]$. Therefore, $\hat{\psi}_2$ could be recovered from f_l using a new highpass filter with lower cutoff at $f_N/4$.

Instead of constructing a new filter for each level we can manipulate the signal f_l to appropriately reuse the highpass filter $h[n]$. If we were to downsample f_l by a factor of 2, that is $f'_l[n] = f_l[2n]$, the Nyquist frequency of f'_l would be $f'_N = f_N/2$. This now means that $\hat{\psi}_2$ from the original spectrum still covers $[f_N/4, f_N/2)$ of \hat{f} , but it covers $[f'_N/2, f'_N)$ in the spectrum for f'_l . As a result, we can reuse the same highpass filter $h[n]$ at $f'_N/2$ to recover $\hat{\psi}_2$.

This principle drives the implementation of a discretized wavelet transform. The combination of filtering and downsampling allows the simulation of unique bandpass filters by reusing the same two lowpass and highpass filters $g[n]$ and $h[n]$. For an input signal $f[n]$, we define the outputs of these two filters to be:

$$f[n] \otimes h[n] = f_h[n] \tag{A.20}$$

$$f[n] \otimes g[n] = f_l[n]. \tag{A.21}$$

The act of downsampling these signals by a factor of two yield what are called the first level detail coefficients (cD1) and first level approximation coefficients (cA1) respectively. Explicitly,

$$cD1 = f_h[2n] \quad (A.22)$$

$$cA1 = f_l[2n]. \quad (A.23)$$

These coefficients correspond to the convolution of the original signal $f[n]$ with the lowpass and highpass filters $g[n]$ and $h[n]$, and represent signals in the temporal domain. This process can be repeated, taking convolutions of approximation coefficients cA^j with $g[n]$ and $h[n]$ at increasingly downsampled levels j . Each filter cascade increments the coefficient decomposition level, so $(cA1 \otimes h)[2n] = cD2$, and $(cA1 \otimes g)[2n] = cA2$. Each level j is downsampled by an additional factor of 2, so for a signal $f[n]$ of length N , the length of coefficients at level j is $N/2^j$. Figure A.4 shows how Figure A.3 fits into the coefficient bandwidths of this highpass/lowpass filtering scheme.

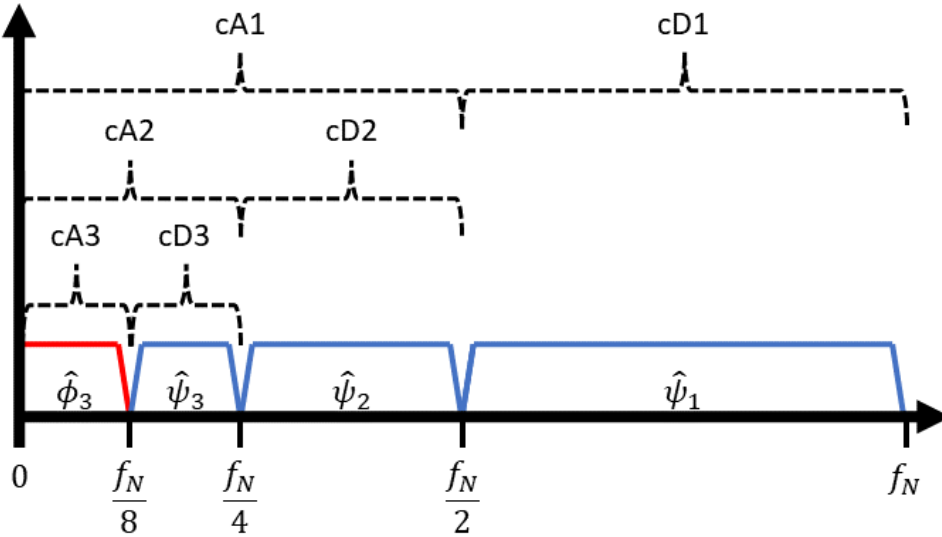


Figure A.4: Figure A.3 under the coefficient umbrellas of cA^j and cD^j bandwidths for a level j .

One crucial condition that must be satisfied is that $g[n]$ and $h[n]$ are quadrature mirror filters (QMF) of each other. Consider their Z-transforms $G(z)$, and $H(z)$. $H(z)$ is quadrature mirror to $G(z)$ if $H(z) = G(-z)$. This makes their frequency responses, $H(e^{j\omega}) = \hat{h}(\omega)$ and $G(e^{j\omega}) = \hat{g}(\omega)$ on the range $[0, \pi]$, mirrored across $\omega = \frac{\pi}{2}$. When the angular frequency is normalized such that $f_N = \pi$, their frequency responses are reflected across $f_N/2$. Additionally, they may be constrained such that:

$$|G(e^{j\omega})|^2 + |H(e^{j(\pi-\omega)})|^2 = 1, \quad (A.24)$$

meaning there is unity power gain (i.e. the summation of power of the lowpass and highpass filter is 1 at all frequencies). An example of quadrature mirror filters with unity power gain is visualized in Figure A.5.




Figure A.5: Two filters that are quadrature mirror to each other. Note these filters also satisfy the unity power gain condition.

A.2.3 DEFINING THE DISCRETE WAVELET TRANSFORM

The previous section outlines the theory for bandpass filtering a signal with cascading highpass and lowpass filters, which is the fundamental theory for the discrete wavelet transform implementation. The block diagram of this approach has already been presented in Figure 5.9 in Chapter 5.




Figure A.6: Figure 5.9 repeated for convenience. Block diagram of DWT cascading filter algorithm. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.

The circular convolution operation in the time domain between a signal $f[n]$ of length N and a high pass filter $h[n]$ of length $M < N$, can be implemented as a matrix-vector product $Ax = b$. We define A to be an $N \times N$ matrix, where each row is a circular shift of the reversed filter $h[-n]$, zero-padded to be of length N . The vector x is defined to be an $N \times 1$ column vector containing the signal $f[n]$, while b is the resulting convolution containing $f_h[n]$. For convenience, we assume $N = 2^k$ for some non-negative integer k .

$$\begin{bmatrix} h[M-1] & h[M-2] & h[M-3] & \cdots & h[0] & 0 & 0 & 0 & \cdots \\ 0 & h[M-1] & h[M-2] & \cdots & h[1] & h[0] & 0 & 0 & \cdots \\ 0 & 0 & h[M-1] & \cdots & h[2] & h[1] & h[0] & 0 & \cdots \\ & & & \vdots & & & & & \\ & & & & & & & & \end{bmatrix} \begin{bmatrix} f[0] \\ f[1] \\ \vdots \\ f[N-1] \end{bmatrix} = \begin{bmatrix} f_h[0] \\ f_h[1] \\ \vdots \\ f_h[N-1] \end{bmatrix} \quad (\text{A.25})$$

$$\text{cD1} = \begin{bmatrix} f_h[0] \\ f_h[2] \\ f_h[4] \\ \vdots \\ f_h[N-4] \\ f_h[N-2] \end{bmatrix}. \quad (\text{A.26})$$

It should be clear that Equation (A.25) contains many wasteful instances of multiplication by 0 that can be removed to improve performance. The rote matrix-vector product has complexity of $O(N \times N \times 1) = O(N^2)$. By removing the zero products and only performing the M multiplications with values $h[m]$, $0 \leq m < M$, the complexity of the convolution can be reduced to $O(MN)$. In general, $M \ll N$ so it is appropriate to bound convolution as an operation by the signal length $O(N)$ to obtain f_h . This complexity cannot be improved further, but the overall time of the operation can be halved by omitting the multiplications that would be removed by downsampling after. By not computing the odd indices of the f_h and instead shifting the filter by 2 indices at each row, we can obtain cD1 from half of a circular convolution.

This extends to the calculation of cA1, as well as cD j and cA j for all decomposition levels. When calculating the coefficients for level j , $f[n]$ is replaced with cA($j-1$).

This decomposition halves the input signal length at each level, meaning there can be no more than $\log_2(N)$ levels of the decomposition. However, the decomposition becomes invalid once the length of the filter exceeds the length of the input signal or coefficients, that is $M > \frac{N}{2^j}$ for a decomposition of level $j+1$. As such, the discrete wavelet decomposition for a signal of length N is $O(N \log \frac{N}{M})$.

A.2.4 SHIFT-INVARIANCE (THE STATIONARY WAVELET TRANSFORM)

The discrete wavelet transform has the shortcoming that it is not a shift-invariant transform. For a continuous transform T that is applied to a signal $f(t)$, $Tf(t) = g(t)$. We say T is time-invariant, if shifting $f(t)$ by a constant value t_0 will shift the output $g(t)$ by the same value without altering the signal in any way. Formally, $Tf(t-t_0) = g(t-t_0)$ for a time-invariant transform T . For the discretized case of a transform T_d , input sequence $f[n]$, and output sequence $g[n]$, T_d is referred to as shift-invariant if $T_d f[n-n_0] = g[n-n_0]$. The discrete wavelet transform does not have this property due to the downsampling step.

Time-variance of the discrete wavelet transform can be proven by an example. Downsampling removes the odd-indexed steps of the circular convolutions. As a result, circularly shifting $f[n]$ by one index in either direction means the previously removed indices are retained, while the previously retained indices are removed. One may surmise that the transform is shift-invariant for even shifts, however this does not hold. Shifting $f[n]$ by two indices shifts $cA1$ by 1 index after downsampling, propagating the previous shift-invariance problem one level down the decomposition. The logical continuation is that a signal of length $N = 2^k$, if zero-padded infinitely, is shift-invariant for shifts of size N . This is the smallest shift that will not produce an odd-numbered shift at some decomposition level in the transform.

Shift- or time-invariance is a potentially desirable property of a transform as it means the way in which a signal was captured will not affect the performance of the transform. The discrete wavelet transform of a captured audio signal will look different if the microphone started capturing 5 seconds before audio started, than if it captured 20 seconds before audio started. The stationary wavelet transform is an augmentation of the discrete wavelet transform that preserves shift-invariance of the transform at the expense of more demanding computation. The block diagram of this approach has been presented in Figure 5.12 in Chapter 5.




Figure A.7: Figure 5.12 repeated for convenience. Block diagram of SWT cascading filter algorithm. Adapted from [51], licensed under a Creative Commons Attribution 4.0 International License.

The same general cascading filter algorithm is used, whereby $h[n]$ and $g[n]$ are used to repeatedly filter each level of the decomposition. Downsampling at each level of the discrete wavelet transform halves the Nyquist frequency, subsequently halving the upper and lower cutoff frequencies of each band. However, downsampling is also the cause of the time-variance in the transform. Therefore, the downsampling step needs to be replaced by something that similarly changes the frequency response of the filter, while not changing the sampling rate or signal length.

Upsampling the filters $h[n]$ and $g[n]$ by inserting a zero between each filter value attempts to solve this for the stationary version of the transformation through what is called the *algorithme à trous*. The act of upsampling in this way doubles the length of the signal without changing the spectral content of the filter values. Now, the same spectral content is spread out over twice the period, meaning the frequency response has been halved. Essentially, this yields the numerical approximation of Equation (A.15). For a filter $h[n]$ of length M , we let $\hat{h}[n]$ denote $h[n]$ upsampled by a factor of 2, to be a filter of length $2M$. Explicitly,

$$\hat{h}[n] = \begin{bmatrix} h[0] & 0 & h[1] & 0 & \dots & h[M-1] & 0 \end{bmatrix} \quad (\text{A.27})$$

We let $\hat{h}[n]$ denote $\dot{h}[n]$ with zeros inserted in the same manner. For j dots, we insert $2^j - 1$ zeros between each index of the original filter. An example of this upsampling and the effect it has on the signal can be seen in Figure A.8.




Figure A.8: The effect of inserting zeros on the frequency response of the db10 wavelet.

Figure A.8 shows that the lower cutoff frequency has been reduced by a factor of 2 from $\frac{\pi}{2}$ to $\frac{\pi}{4}$ but the upper cutoff is not at $\frac{\pi}{2}$ as is desired for the cD2 coefficients. Recall that the convolution with $\hat{h}[n]$ is performed on the cA1 coefficients, which were obtained from convolving the signal $f[n]$ with the low-pass filter $g[n]$. These two filters combined have the desired cutoff frequencies of $\frac{\pi}{4}$ and $\frac{\pi}{2}$, as shown in Figure A.9.

The zero insertion methodology in the stationary wavelet transform replaces the frequency cutoff manipulation aspect of downsampling in the discrete wavelet transform, but it does not alone solve the issue of time-invariance. Time-invariance is achieved by performing full circular convolutions with the zero-inserted filters, as opposed to the downsampled convolution described in Equations (A.20) to (A.23). Formally,

$$f[n] \otimes h[n] = f_h[n] = \text{cD1} \quad (\text{A.28})$$

$$f[n] \otimes g[n] = f_l[n] = \text{cA1} \quad (\text{A.29})$$

in the context of the stationary wavelet transform. The next decomposition level coefficients are calculated as

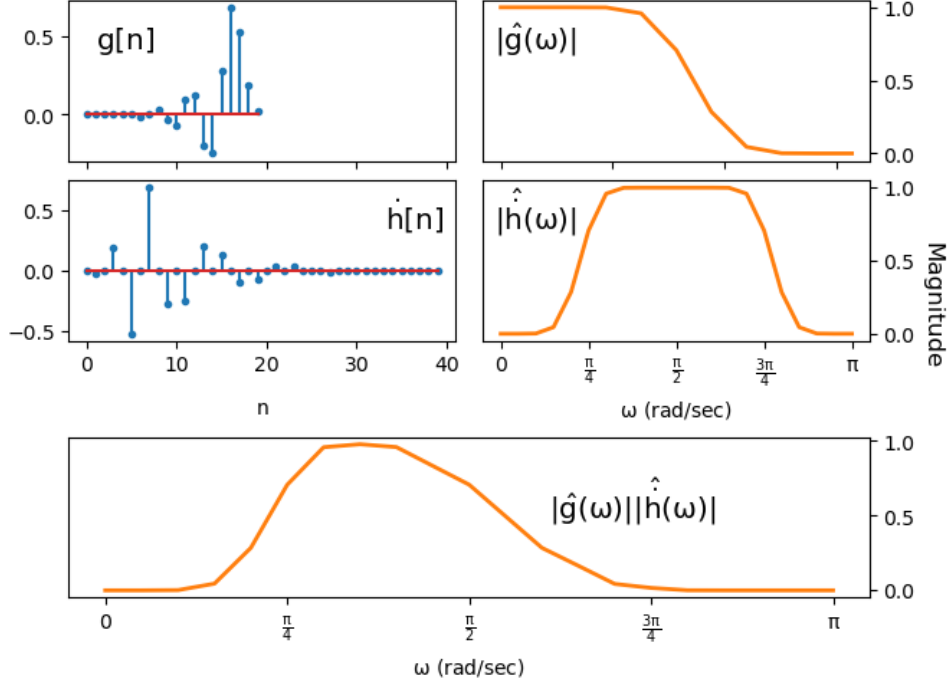


Figure A.9: Obtaining cD2 from the second level decomposition using the db10 wavelet. These coefficients are obtained by first filtering with $g[n]$ then by $\hat{h}[n]$, which has the effective frequency response shown in the bottom plot.

$$cA1 \otimes \hat{h}[n] = cD2 \quad (A.30)$$

$$cA1 \otimes \hat{g}[n] = cA2. \quad (A.31)$$

The removal of downsampling combined with the insertion of zeros into the filters yield the time-invariant stationary wavelet transform.

These improvements come at a computational cost. In the discrete wavelet transform, the convolution at each level was reduced by a factor of 2 due to downsampling, whereas in the stationary wavelet transform, these convolutions are all $O(N)$ operations for a signal $f[n]$ of length N . The maximum transformation depth is still determined by how many zero insertions can be done until the zero-inserted filter is longer than the signal length N . The maximum level J is determined as $M2^J \leq N$ meaning $J \leq \log(\frac{N}{M})/\log(2)$, where M is the length of the original filter. It follows then that the stationary wavelet transform has a time complexity of $O(N\log(\frac{N}{M}))$. While this is still the same complexity as the non-stationary version of the transform, it is slower due to the lack of decimation. In the previous version of the transform, a decomposition of a signal of length N at step j would involve a convolution with a signal of length $\frac{N}{2^j}$ as the length of the coefficient vector would be halved at each step. In the stationary version however, the coefficient length at each level is always length N . As a result, each decomposition level j for the stationary wavelet transform takes 2^j times longer than the equivalent level j in the discrete wavelet transform.

Bibliography

- [1] *Light Rail Transit – City of Hamilton*. <https://www.hamilton.ca/city-council/plans-strategies/city-projects/light-rail-transit>. Accessed: Nov 30, 2023.
- [2] *Le projet – Le tramway de Québec*. <https://tramwaydequebec.info/projet/>. Accessed: Nov 30, 2023.
- [3] *Quarterly and Annual Totals by Mode*. Tech. rep. Washington, The United States: American Public Transit Association, 2023.
- [4] S. L. Grassie and J. Kalousek. “Rail Corrugation: Characteristics, Causes and Treatments”. In: *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit* 207.1 (2009), pp. 57–68. doi: 10.1243/PIME_PROC_1993_207_227_02. eprint: https://doi.org/10.1243/PIME_PROC_1993_207_227_02. url: https://doi.org/10.1243/PIME_PROC_1993_207_227_02.
- [5] X. Du et al. “Rail Corrugation of High-Speed Railway Induced by Rail Grinding”. eng. In: *Shock and vibration* 2021 (2021), pp. 1–14. issn: 1070-9622.
- [6] E. E. Magel. *Rolling contact fatigue : a comprehensive review*. eng. Washington, DC: U.S. Department of Transportation, Federal Railroad Administration, Office of Research and Development, 2011.
- [7] R. Schwarzenberger and R. Stock. “Rail milling as the next level of rail maintenance technology in Australia”. In: *Global Railway Review* (Nov. 2019).
- [8] J. W. Palese, T. Euston, and A. M. Zarembski. “Use of Profile Indices for Quality Control of Grinding”. In: *Annual Conference and Exposition*. 2004.
- [9] J. Kalousek and E. Magel. *Achieving a balance: The ”magic” wear rate*. eng. 1997.
- [10] *Railway applications - Track - Acceptance of works - Part 3: Acceptance of reprofiling rails in track*. Standard. Brussels, Belgium: European Committee for Standardization, Jan. 2012.
- [11] A. Galaitsis and E. Bender. “Wheel/rail noise—Part V : Measurement of wheel and rail roughness”. eng. In: *Journal of sound and vibration* 46.3 (1976), pp. 437–451. issn: 0022-460X.

- [12] E. Magel and K. Oldknow. "Quality Indices for Managing Rail through Grinding". In: *Proceedings of the 11th International Conference on Contact Mechanics and Wear of Rail/Wheel Systems (CM 2018)*. Ed. by Z. Li and A. Núñez. Delft, The Netherlands, Sept. 2018, pp. 658–67.
- [13] *CAT Product Datasheet*. Available at <https://www.railmeasurement.com/wp-content/uploads/CAT-product-data-sheet.pdf>. RailMeasurement. Stadthagen, Germany, Sept. 2022.
- [14] *Acoustics – Railway applications – Measurement of noise emitted by railbound vehicles*. Standard. Geneva, Switzerland: International Organization for Standardization, Aug. 2013.
- [15] L. Cohn and W. McColl. *L10 versus Leq—A User’s Perspective* /. eng. West Conshohocken, Pa: ASTM International, 1979. ISBN: 9780803147577.
- [16] *Acoustics - Preferred reference values for acoustical and vibratory levels*. Standard. Geneva, Switzerland: International Organization for Standardization, May 2015.
- [17] *Electroacoustics - Sound level meters - Part 2: Pattern evaluation tests*. Standard. Geneva, Switzerland: International Electrotechnical Commission, Sept. 2013.
- [18] T. Sethuramalingam and A. Vimalajuliet. "Design of MEMS based capacitive accelerometer". In: *2010 International Conference on Mechanical and Electrical Technology*. 2010, pp. 565–568. DOI: 10.1109/ICMET.2010.5598424.
- [19] K. Hempelmann and K. Knothe. "An extended linear model for the prediction of short pitch corrugation". In: *Wear* 191.1-2 (Jan. 1996). ISSN: 0043-1648.
- [20] H. Xiao et al. "Initiation and development of rail corrugation based on track vibration in metro systems". In: *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit* 232.9 (2018), pp. 2228–2243. DOI: 10.1177/0954409718768956. eprint: <https://doi.org/10.1177/0954409718768956>. URL: <https://doi.org/10.1177/0954409718768956>.
- [21] *Mechanical vibration - Ground-borne noise and vibration arising from rail systems, Part 32: Measurement of dynamic properties of the ground*. Standard. Geneva, Switzerland: International Organization for Standardization, Dec. 2015.
- [22] *Mechanical vibration - Measurement of vibration generated internally in railway tunnels by the passage of trains*. Standard. Geneva, Switzerland: International Organization for Standardization, Sept. 2016.
- [23] *Railway applications - Rolling stock equipment - Shock and vibration tests*. Standard. Geneva, Switzerland: International Electrotechnical Commission, May 2010.
- [24] *Mechanical vibration and shock - Evaluation of human exposure to whole-body vibration - Part 4: Guidelines for the evaluation of the effects of vibration and rotational motion on passenger and crew comfort in fixed-guideway transport systems*. Standard. Geneva, Switzerland: International Organization for Standardization, July 2010.

- [25] J. Kawasaki and K. Youcef-Toumi. “Estimation of rail irregularities”. In: *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*. Vol. 5. 2002, 3650–3660 vol.5. DOI: 10.1109/ACC.2002.1024495.
- [26] K. Kang. “A study of the long wave track defect analysis for high speed railways”. In: *J. Korean Soc. Railway* 8.2 (2005).
- [27] J. S. Lee et al. “A Mixed Filtering Approach for Track Condition Monitoring Using Accelerometers on the Axle Box and Bogie”. In: *IEEE Transactions on Instrumentation and Measurement* 61.3 (2012), pp. 749–758. DOI: 10.1109/TIM.2011.2170377.
- [28] *Uncertainty of measurement - Part 3: Guide to the expression of uncertainty in measurement (GUM:1995)*. Standard. Geneva, Switzerland: International Organization for Standardization, Sept. 2009.
- [29] P. D. Groves. “GNSS: User Equipment Processing and Errors”. In: *Principles of GNSS, Inertial, and Multi-sensor Integrated Navigation Systems (2nd Edition)*. Artech House, 2013. Chap. 9, pp. 349–436. ISBN: 978-1-60807-005-3. URL: <https://app.knovel.com/hotlink/khtml/id:kt011M48R1/principles-gnss-inertial/gnss-user-equipment-processing>.
- [30] T. Shirai and N. Kubo. “RTK-GPS Reliability Improvement in Dense Urban Areas”. In: *Japan Society of Aeronautical Space Sciences* 60 (Mar. 2012), pp. 40–47. DOI: 10.2322/jjsass.60.40.
- [31] NovAtel. *IMU Errors and Their Effects*. Rev A. NovAtel. Feb. 2014.
- [32] P. T. Torstensson and M. Schilke. “Rail corrugation growth on small radius curves—Measurements and validation of a numerical prediction model”. In: *Wear* 303.1 (2013), pp. 381–396. ISSN: 0043-1648. DOI: <https://doi.org/10.1016/j.wear.2013.03.029>. URL: <https://www.sciencedirect.com/science/article/pii/S0043164813002184>.
- [33] P. S. Sroba and M. Roney. “Rail Grinding Best Practices”. In: *Proceedings of the AREMA 2003 Annual Conference*. Chicago, The United States, Oct. 2003.
- [34] M. Ishida, T. Moto, and M. Takikawa. “The effect of lateral creepage force on rail corrugation on low rail at sharp curves”. In: *Wear* 253.1 (2002). CM2000 S.I., pp. 172–177. ISSN: 0043-1648. DOI: [https://doi.org/10.1016/S0043-1648\(02\)00096-0](https://doi.org/10.1016/S0043-1648(02)00096-0). URL: <https://www.sciencedirect.com/science/article/pii/S0043164802000960>.
- [35] J. Paragreen. “Principles of Wheel-Rail Friction Management”. WRI. 2022. URL: <https://wri.wheel-rail-seminars.com/archives/2022/pc-papers/presentations/PC07%20Jonathan%20Paragreen%20-%20WRI%202022%20Principles%20Friction%20Management%20-%20J%20Paragreen%20V2.pdf>.

- [36] D. T. Eadie et al. "The effects of top of rail friction modifier on wear and rolling contact fatigue: Full-scale rail-wheel test rig evaluation, analysis and modelling". In: *Wear* 265.9 (2008). Contact Mechanics and Wear of Rail/Wheel Systems - CM2006, pp. 1222–1230. ISSN: 0043-1648. DOI: <https://doi.org/10.1016/j.wear.2008.02.029>. URL: <https://www.sciencedirect.com/science/article/pii/S0043164808001695>.
- [37] R. M. Pugsley. *Pugsley's dead reckoning*. eng. New York: R. M. Pugsley, 1916.
- [38] M. R. Parodkar, M. Sheta, and S. Chokkadi. "Automotive Dead Reckoning-Implementation of Fusion Algorithm on Raspberry Pi". In: *Smart Sensors Measurements and Instrumentation*. Singapore: Springer Singapore, 2021, pp. 189–209.
- [39] R. J. Babb. "Navigation of unmanned underwater vehicles for scientific surveys". In: *Symposium on Autonomous Underwater Vehicle Technology*. 1990, pp. 194–198. DOI: 10.1109/AUV.1990.110456.
- [40] D. Wyatt. *Aircraft Flight Instruments and Guidance Systems*. 1st Edition. London, England: Routledge, Aug. 2014. URL: <https://doi-org.um1.idm.oclc.org/10.4324/9781315858975>.
- [41] R. Zanetti and R. Bishop. "Precision Entry Navigation Dead-Reckoning Error Analysis: Theoretical Foundations of the Discrete-Time Case". In: Aug. 2007.
- [42] S. Bai, J. Thibault, and D. D. McLean. "Dynamic data reconciliation: Alternative to Kalman filter". In: *Journal of Process Control* 16.5 (2006), pp. 485–498. ISSN: 0959-1524. DOI: <https://doi.org/10.1016/j.jprocont.2005.08.002>. URL: <https://www.sciencedirect.com/science/article/pii/S095915240500096X>.
- [43] R. L. (L. Eubank. *A Kalman filter primer*. eng. Statistics, textbooks and monographs ; v. 186. Boca Raton, Fla: Chapman & Hall/CRC, 2006. ISBN: 0-429-11759-0.
- [44] S. G. Mallat. *A Wavelet Tour of Signal Processing*. eng. 2nd ed. San Diego: Academic Press, 1999. ISBN: 1-281-07081-5.
- [45] E. Sejdic, I. Djurovic, and L. Stankovic. "Quantitative Performance Analysis of Scalogram as Instantaneous Frequency Estimator". In: *IEEE Transactions on Signal Processing* 56.8 (2008), pp. 3837–3845. DOI: 10.1109/TSP.2008.924856.
- [46] G. R. Lee et al. "PyWavelets: A Python package for wavelet analysis". In: *Journal of Open Source Software* 4.36 (2019), p. 1237. DOI: 10.21105/joss.01237. URL: <https://doi.org/10.21105/joss.01237>.
- [47] K. Gröchenig. "The Short-Time Fourier Transform". In: *Foundations of Time-Frequency Analysis*. Boston, MA: Birkhäuser Boston, 2001, pp. 37–58. ISBN: 978-1-4612-0003-1. DOI: 10.1007/978-1-4612-0003-1_4. URL: https://doi.org/10.1007/978-1-4612-0003-1_4.
- [48] B. P. Lathi. *Linear Systems and Signals, Second Edition*. New York, The United States: Oxford University Press, 2005.

- [49] X. Zhang et al. “The attenuated Ricker wavelet basis for seismic trace decomposition and attenuation analysis”. eng. In: *GEOPHYSICAL PROSPECTING* 68.2 (2020), pp. 371–381. ISSN: 0016-8025.
- [50] M. Elgendi et al. “Detection of Heart Sounds in Children with and without Pulmonary Arterial Hypertension—Daubechies Wavelets Approach”. In: *PLOS ONE* 10.12 (Dec. 2015), pp. 1–22. DOI: 10 . 1371 / journal . pone . 0143146. URL: <https://doi.org/10.1371/journal.pone.0143146>.
- [51] L.-Y. Shao et al. “Data-Driven Distributed Optical Vibration Sensors: A Review”. In: *IEEE Sensors Journal* PP (Sept. 2019), pp. 1–1. DOI: 10 . 1109 / JSEN . 2019 . 2939486.
- [52] V. Kordic. *Kalman Filter*. Ed. by V. Kordic. IntechOpen, 2010. ISBN: 953-51-5912-7.
- [53] V. Giurgiutiu. “Chapter 14 - Signal Processing and Pattern Recognition for Structural Health Monitoring with PWAS Transducers”. In: *Structural Health Monitoring with Piezoelectric Wafer Active Sensors (Second Edition)*. Ed. by V. Giurgiutiu. Second Edition. Oxford: Academic Press, 2014, pp. 807–862. ISBN: 978-0-12-418691-0. DOI: <https://doi.org/10.1016/B978-0-12-418691-0.00014-9>. URL: <https://www.sciencedirect.com/science/article/pii/B9780124186910000149>.
- [54] J. A. Hogan and J. D. Lakey. *Time–Frequency and Time–Scale Methods Adaptive Decompositions, Uncertainty Principles, and Sampling*. eng. 1st ed. 2005. Applied and Numerical Harmonic Analysis. Boston, MA: Birkhäuser Boston, 2005. ISBN: 1-281-14809-1.
- [55] M. H. Song et al. “Spectral analysis based on wavelet transform maxima; identification of multi-order stratigraphic boundaries and cycles”. eng. In: *Mathematical geosciences* 53.5 (2021), pp. 969–997. ISSN: 1874-8961.
- [56] R. Marsalek, J. Pomenkova, and S. Kapounek. “A Wavelet-Based Approach to Filter Out Symmetric Macroeconomic Shocks”. In: *Computational Economics* 44 (Dec. 2013). DOI: 10 . 1007 / s10614 - 013 - 9403 - x.
- [57] G. Strang. “Wavelets”. In: *American Scientist* 82.3 (1994), pp. 250–255. ISSN: 00030996. URL: <http://www.jstor.org/stable/29775194> (visited on 11/26/2023).