

A Physically Informed Data-Driven Approach to Analyze Human Induced Vibration in Civil Structures

Ellis C. Kessler

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Mechanical Engineering

Pablo A. Tarazaga, Co-chair
Serkan Gugercin, Co-chair
Rodrigo Sarlo
Robert West
Andrew Kurdila

May 12, 2021
Blacksburg, Virginia

Keywords: Smart Buildings, Vibration, Wave Propagation, Gait Analysis, Data-Driven
Modelling

Copyright 2021, Ellis C. Kessler

A Physically Informed Data-Driven Approach to Analyze Human Induced Vibration in Civil Structures

Ellis C. Kessler

(ABSTRACT)

With the rise of the Internet of Things (IoT) and smart buildings, new algorithms are being developed to understand how occupants are interacting with buildings via structural vibration measurements. These vibration-based occupant inference algorithms (VBOI) have been developed to localize footsteps within a building, to classify occupants, and to monitor occupant health. This dissertation will present a three-stage journey proposing a path forward for VBOI research based on physically informed data-driven models of structural dynamical systems.

The first part of this dissertation presents a method for extracting temporal gait parameters via underfloor accelerometers. The time between an occupant's consecutive steps can be measured with only structural vibration measurements with a similar accuracy to current gait analysis tools such as force plates and in-shoe pressure sensors. The benefit of this, and other VBOI gait analysis algorithms, is in their ease of use. Gait analysis is currently limited to a clinical setting with specialized measurement systems, however VBOI gait analysis provides the ability to bring gait analysis to any building.

VBOI algorithms often make some simplifying assumptions about the dynamics of the building in which they operate. Through a calibration procedure, many VBOI algorithms can learn some system parameters. However, as demonstrated in the second part of this dissertation, some commonly made assumptions oversimplify phenomena present in civil structures such as: attenuation, reflections, and dispersion. A series of experimental and theoretical investigations show that three common assumptions made in VBOI algorithms are unable to account for at least one of these phenomena, leading to algorithms which are more accurate under certain conditions.

The final part of this dissertation introduces a physically informed data-driven modelling technique which could be used in VBOI to create a more complete model of a building. Continuous residue interpolation (CRI) takes FRF measurements at a discrete number of testing locations, and creates a predictive model with continuous spatial resolution. The fitted CRI model can be used to simulate the response at any location to an input at any other location. An example of using CRI for VBOI localization is shown.

A Physically Informed Data-Driven Approach to Analyze Human Induced Vibration in Civil Structures

Ellis C. Kessler

(GENERAL AUDIENCE ABSTRACT)

Vibration-based occupant inference (VBOI) algorithms are an emerging area of research in smart buildings instrumented with vibration sensors. These algorithms use vibration measurements of the building's structure to learn something about the occupants inside the building. For example the vibration of a floor in response to a person's footstep could be used to estimate where that person is without the need for any line-of-sight sensors like cameras or motion sensors. The storyline of this dissertation will make three stops:

- The first is the demonstration of a VBOI algorithm for monitoring occupant health.
- The second is an investigation of some assumptions commonly made while developing VBOI algorithms, seeking to shed light on when they lead to accurate results and when they should be used with caution.
- The third, and final, is the development of a data-driven modelling method which uses knowledge about how systems vibrate to build as detailed a model of the system as possible.

Current VBOI algorithms have demonstrated the ability to accurately infer a range of information about occupants through vibration measurements. This is shown with a varied literature of localization algorithms, as well as a growing number of algorithms for performing gait analysis. Gait analysis is the study of how people walk, and its correlation to their health. The vibration-based gait analysis procedure in this work demonstrates extracting distributions of temporal gait parameters, like the time between steps.

However, many current VBOI algorithms make significant simplifying assumptions about the dynamics of civil structures. Experimental and theoretical investigations of some of these assumptions show that while all assumptions are accurate in certain situations, the dynamics of civil structures are too complex to be completely captured by these simplified models.

The proposed path forward for VBOI algorithms is to employ more sophisticated data-drive modelling techniques. Data-driven models use measurements from the system to build a model of how the system would respond to new inputs. The final part of this dissertation is the development of a novel data-driven modelling technique that could be useful for VBOI. The new method, continuous residue interpolation (CRI) uses knowledge of how systems vibrate to build a model of a vibrating system, not only at the locations which were measured, but over the whole system. This allows a relatively small amount of testing to be used to create a model of the entire system, which can in turn be used for VBOI algorithms.

Acknowledgments

The work in this dissertation is supported in part by the National Science Foundation via grant no. DGE-1545362, UrbComp (Urban Computing): Data Science for Modeling, Understanding, and Advancing Urban Populations. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Contents

List of Figures	ix
List of Tables	xv
1 Introduction	1
1.1 Smart Buildings and VBOI	1
1.1.1 VBOI Algorithms	3
1.2 Background for VBOI Algorithms	4
1.2.1 Input/Output relationship based algorithms	5
1.2.2 Energy based algorithms	6
1.2.3 Waveguide based algorithms	6
1.2.4 Common assumptions in VBOI	7
1.3 Data-Driven Modelling for VBOI	8
1.3.1 VBOI Algorithm Calibration	8
1.3.2 Opportunity for structural dynamics data-driven models	9
1.4 Contributions	10
1.5 Outline of Dissertation	11
2 Gait Analysis via Underfloor Accelerometers	12
2.1 Introduction	12
2.1.1 Traditional Gait Analysis Techniques	13
2.2 Outline of Gait Analysis Process	14
2.2.1 Participants	14
2.2.2 Experiment and Data Processing	15
2.3 Gait Parameters from Structural Measurements	17
2.3.1 Effect of number of sensors	17

2.3.2	Analysis of Normal Gait	18
2.4	Comparison with traditional gait analysis tools	22
2.4.1	Participants	22
2.4.2	Equipment and Test Set-up	22
2.4.3	Data Analysis	23
2.4.4	Results and Discussion	23
2.5	Discussion	25
2.6	Conclusion	25
3	Investigation of assumptions in VBOI	27
3.1	Introduction	27
3.2	Current vibration-based occupant inference algorithms	29
3.3	Methods	30
3.3.1	Series 1: Observations of vibration in buildings	31
3.3.2	Series 2: Detailed investigations of vibration in a beam-column structure	32
3.4	Assumption 1: Steady-state standing wave behavior	34
3.4.1	Observation of standing wave quality through modal analysis	34
3.4.2	Sources of variability in FRFs and modal parameters	37
3.4.3	Effect of energy loss on FRFs	40
3.4.4	Summary of assumption	43
3.5	Assumption 2: Infinite plate with uniform attenuation	43
3.5.1	Observation of attenuation and reflections in Goodwin Hall	44
3.5.2	Wave propagation through a column	46
3.5.3	Parameters affecting mode conversion and reflections	48
3.5.4	Summary of assumption	50
3.6	Assumption 3: Neglected or simplified dispersion relations	50
3.6.1	Observations of dispersion in buildings	51
3.6.2	Quantification of dispersion relations	52
3.6.3	Parameters that affect dispersion relations	54

3.6.4	Summary of assumption	55
3.7	Conclusion	55
4	Modelling Second Order Systems	58
4.1	Introduction	58
4.1.1	Overview of Characterizing Second Order Systems	58
4.2	Theory of Second Order Systems	60
4.2.1	Partial Differential Equations	60
4.2.2	State Space Model	61
4.2.3	Transfer Function	61
4.3	Building Predictive Models	64
4.3.1	Admissible functions in the Rayleigh-Ritz Method	65
4.3.2	Choice of collocation locations	68
4.3.3	Model order of Finite Element models	70
4.3.4	Rational Approximation	71
4.4	Conclusion	72
5	A Novel Data-Driven Method to Build Continuous Predictive Models	74
5.1	Introduction	74
5.2	Comparison to Current Methods for Modelling Second Order Systems	76
5.2.1	Comparison to analytical methods	76
5.2.2	Comparison to data-driven methods	77
5.3	Continuous Residue Interpolation Algorithm	78
5.3.1	Setting up the optimization problem	78
5.3.2	Solving the optimization problem	80
5.3.3	Interpolated Points and State Space	80
5.3.4	Implementation Considerations	81
5.4	Applications	84
5.4.1	4-DOF system	84

5.4.2	Euler-Bernoulli Beam	86
5.4.3	Marimba Bar	89
5.5	CRI for VBOI	94
5.5.1	Two-floor system	94
5.5.2	CRI fit	95
5.5.3	CRI-FEEL Localization	97
5.6	Conclusion	99
6	Conclusion	101
6.1	Summary of Contributions	101
6.2	Areas of Future Work	102
6.2.1	Future of Vibration-based Gait Analysis	102
6.2.2	Future of Vibration-based Occupant Inference	103
6.2.3	Future of Continuous Residue Interpolation	103
	Bibliography	105
	Appendices	115
	Appendix A Admissible Function Details	116
	Appendix B CRI Optimization Details	121
	Appendix C CRI Code	123
C.1	CRI Example	124
C.2	CRI Function	126
C.3	CRI2SS Function	131

List of Figures

1.1	Left: A picture of Goodwin Hall. Right: A diagram of the fourth floor of Goodwin with sensor locations. The number on each sensor location refers to the number of axes measured. The single axis measurements are all in the vertical direction, and an example of how one of these sensor is shown in the bottom right.	1
1.2	Floor vibration (b) measured by two underfloor accelerometers as an occupant walks down a hallway from left to right (a). Two individual footstep responses (c,d) are shown of these events comparing a sensor near the impact and a sensor roughly 42 feet away.	2
1.3	Graphical representation of how input forces from footsteps induces a response in the building's structure.	5
2.1	(a) Diagram of all sensors along the hallway used for testing. (b) An example of an acceleration history from a single trial with all steps detected. The trial comes from the sensor marked in blue in (a).	16
2.2	An example of an acceleration measurement containing two steps, the corresponding signal-to-noise-ratio (SNR), and the detected step times based on the peaks in SNR.	16
2.3	Step intervals for one trial each of Participant 12 (left) and Participant 16 (right). (a) and (b) correspond to approach (i), (c) and (d) correspond to approach (ii), and (e) and (f) correspond to approach (iii). The distributions using each approach are shown as histograms, while normal probability distribution fits to each histogram are also shown for reference. Beside each fit are the mean, μ , and standard deviation, σ	18
2.4	Step intervals and relative step lengths for all participants. Females in the left plot (a), and males in the right plot (b). Marker size is related to the maximum energy in a step from each participant.	20
2.5	Comparison of cycle times and step signal energy for Participant 12 (a, b), and Participant 16 (c, d), respectively. Walking trials are shown in blue, brisk walking in orange, and running in yellow. Fitted normal distributions are plotted for cycle times along with the mean, μ , and standard deviation, σ . Fitted Rayleigh distributions are plotted for step energy along with the mean.	21

2.6	Examples of detecting the heel strikes of two consecutive steps with (top) Pedar, and (bottom) accelerometers. Each detected heel strike is shown with a +.	24
2.7	Time between heel strikes for three different participants using three different measurement techniques. The mean is shown as the darker bar, as well as +1 SD as the lighter bar.	24
3.1	Floor vibration (b) measured by two underfloor accelerometers as an occupant walks down a hallway from left to right (a). Two individual footstep responses (c,d) are shown of these events comparing a sensor near the impact and a sensor roughly 42 feet away.	28
3.2	(a) Structural layout of the floor in Goodwin Hall under the hallway being studied. The 17 underfloor sensor locations, 10 surface mounted sensor locations, and 3 impact locations are shown. Additionally, structural information about the building such as locations of columns, beams, and the walls of the hallway are shown. For reference, the distance between surface mounted sensors (excluding impact location 1) is 10 feet. (b) Diagram showing relationship between surface mounted sensors, underfloor sensors, and the floor.	32
3.3	Image (a) and diagram (b) of the beam-column structure. The 14 measurement locations are shown in the diagram. For each measurement location, tri-axial measurements were taken on three faces: for example at location B2 there were measurements at the middle line on the top, front, and bottom faces simultaneously.	33
3.4	Accelerance FRFs for the 17 underfloor hallway sensors (a) and coherence of selected sensors (b). Darker colors are for sensors closer to the impact location (location 1) shown in Figure 3.2	35
3.5	Estimated mode shapes for the four modes from OMA (a-d) and EMA (e-h). The phase of each measurement point is plotted on a circle around each mode shape, showing the complexity of the mode.	37
3.6	Variation of estimated natural frequencies (a) and damping ratios (b) for the four modes with average hallway RMS acceleration energy.	38
3.7	Demonstration of reciprocity between two sensors 5 feet apart (a) and 25 feet apart (b). First extracted mode shapes for the surface mounted sensors with each of the three impact locations (c). For each mode shape, a triangle shows where the impact location was for the FRFs from which the mode shape was extracted.	40

3.8	Diagram of clamped-clamped Euler-Bernoulli rectangular beam (a) with 100 elements connected to the environment at each node by a spring and damper to simulate energy loss. The following parameters were used in the beam model: Total length $L = 30$ m, width $w = 1$ m, height $h = 0.5$ m, Young's modulus $E = 25$ GPa, density $\rho = 2700$ kg/m ³ , spring stiffness $k = 3 \times 10^5$ N/m. Results for three values of damping: $c = 3 \times 10^2$ kg · s (b), $c = 3 \times 10^3$ kg · s (c), $c = 1.5 \times 10^4$ kg · s (d). For each value of damping the time response to a half-sine input, FRFs, and first extracted mode shape for selected nodes are shown. For mode shapes, a triangle denotes the corresponding input location.	42
3.9	Visualization of acceleration attenuation with distance for surface sensors, underfloor sensors, and ground floor sensors (a). Acceleration attenuation with distance for surface floor sensors only, separated into 0-30 Hz and 30-100 Hz bands (b). RMS acceleration is normalized by the RMS acceleration at 0 distance from the impact, RMS_0	45
3.10	Diagram of sensors on the fourth floor and third floor near a column (a). Spectrograms of measurements at each sensor from an impact on the fourth floor.	46
3.11	Visualization of cross section motion in the x and y directions (z motion is ignored as it was negligible). The energy and relative motion direction for each sensor on the beam and column structure are shown. Energy is calculated by integrating FRFs to remove influence of differences in hammer impacts, while relative phase is determined by cosine similarity of time series since all measurements at each location were taken during the same impact. Inset are time series measurements for location B2, to illustrate how the acceleration changes between the faces as well as schematics of what various wave modes would look like.	48
3.12	Diagram of the Abaqus model (a). Amplitude ratios for changing slenderness ratio (b), and amplitude ratios for changing stiffness ratio (c). Amplitude ratios for A_0 waves staying in the beam, both A_0 and S_0 waves in the column, as well as the total magnitude in the column are shown.	50
3.13	Diagram of sensors on the surface of floor, and under the I-beam (a). Spectrograms of measurements from the surface (b,c) and under the I-beam (d,e) for an impact from an instrumented hammer. One sensor for the surface and I-beam are close to the impact (b,d) while the other is 40-42 feet away (c,e) so the effects of dispersion can be seen.	52
3.14	Spectrograms from an impact on a ground floor, measured by sensors placed on the surface of the concrete floor 15 feet away from each other.	53

3.15	Dispersion curves measured in the beam and column structure, as well as the fourth floor of Goodwin (a). The group velocities (c_g) for the first three symmetric and antisymmetric Lamb modes of a plate with $\nu = 0.2$, normalized by the material's shear speed c_s and the plate thickness H (b).	54
4.1	A graphical comparison of selected methods for characterizing second order systems.	59
4.2	Diagram of beam with pinned boundary conditions at both ends.	62
4.3	A diagram of the Meirovitch bar system.	65
4.4	Comparison of errors for mode 1 and mode 5 of the bar system with various choices of admissible functions.	67
4.5	Condition number of the $\mathbf{M}$ matrix for various choices of admissible functions.	68
4.6	Comparison of polynomial interpolation accuracy with uniformly spaced points and Chebyshev nodes. Polynomial order 8 is used for both types of spacing, and for both functions.	69
4.7	Diagram of a three element finite element model of the bar system.	70
4.8	Comparison of the error for mode 1 and mode 5 using RRM and FE modelling at the same model orders.	71
4.9	Two vector fitting models, one with model order 2 and one with model order 10. The Meirovitch bar has 5 modes in this frequency range, thus the model order 10 fit is able to recreate the FRFs well as seen on the right.	72
5.1	Similarity between steps in CRI and three analytical methods.	77
5.2	Diagram of the 4-DOF system showing inputs $f_{1-4}(t)$ and outputs $y_{1-4}(t)$ and system properties. Inputs and outputs where there are virtual sensors for the simulated experiment are shown in blue. A schematic of the full transfer function $H_F(s) \in \mathbb{C}^{4 \times 4}$ is also shown, where the measured entries $H(s) \in \mathbb{C}^{3 \times 2}$ are shown in blue and entries which will be interpolated are shown in red.	85
5.3	Residual error after each iteration of the CRI optimization problem for the 4-DOF system in Application 1.	85
5.4	Frequency response functions of the 4-DOF system, the CRI model, and the error between the system and CRI. Both the six measured and fitted FRFs (left) and the ten interpolated FRFs (right) are shown.	86

5.5	Diagram of the clamped-clamped Euler-Bernoulli beam used in Application 2. The system properties, stiffness, E , inertia, I , density, ρ , and area, A are uniform along the beam's length.	86
5.6	Frequency response functions of the beam system, the CRI model, and the error between the system and CRI. The measured and fitted FRFs (left) and a selection of interpolated FRFs (right) are shown.	87
5.7	Error plots showing error between the system FRFs and the CRI model FRFs for all combinations of input and output locations for two polynomial orders. Lighter colors are higher error. For each order, measured input locations are shown with vertical dotted lines, while measured output locations are shown with horizontal lines. FRFs where lines intersect are the measured and fitted FRFs, while the rest of the plot represent interpolated CRI FRFs.	88
5.8	Comparison of errors for modes 1-4 for CRI, RRM, and collocation for three different polynomial orders.	89
5.9	Diagram of the experimental setup for the marimba bar test. The bar is suspended on tensioned ropes that run through the bar. The velocity of the top face is measured with a PSV-400 LDV system. Input force is provided by a PCB impact hammer.	90
5.10	The 2D basis functions used for fitting the CRI model of the marimba bar. The basis functions are products of 1D basis functions for each dimension: $\phi_{ij}(x, y) = \phi_i(x)\phi_j(y)$	90
5.11	Measured FRFs and the corresponding CRI model FRFs for one measured input. Each mode has an arrow above, corresponding to the mode groups in Figure 5.12.	91
5.12	Visualization of the 11 continuous CRI modes. The modes are separated into three groups: boundary condition modes that are dominated by rotation and translation, flexural modes which are dominated by deflection along the y -direction, and torsional modes which also have deflection along the x -direction.	92
5.13	Left: A matrix showing the error between all combinations of input and output locations. Right: The errors for all output locations (one column of the matrix on the right) mapped onto the surface of the marimba bar for input 7 (fitted) and input 11 (interpolated).	93
5.14	Comparison of velocity responses simulated with the true measured transfer function, and CRI interpolated transfer function for two cases. Top: fitted input and interpolated output. Bottom: both interpolated input and interpolated output. Each response is also shown with a zoomed view of the beginning of the response to the right.	94

5.15	Diagram of the two-floor system for which CRI based localization will be built for the second floor. The input and output locations from the simulated test are shown.	95
5.16	The 15 measured FRFs from the simulated test, and the fitted FRFs from the CRI model. Fits assuming real modes and complex modes are shown.	96
5.17	Overall error for all possible FRFs: $H(\omega) \in \mathbb{C}^{101 \times 101}$. Fits assuming real modes and complex modes are shown.	96
5.18	Comparison of CRI-FEEL and FEEL localization for a simulated impact at $x = 0$	98
5.19	CRI-FEEL localization for all possible input locations on the second floor.	99
C.1	Output figures from the example 4DOF system from the code in Section C.1.	123

List of Tables

1.1	Summary of selected VBOI localization algorithms. The number of sensors, sensor type (A=accelerometer, G=geophone), and the distance error reported are included.	3
2.1	Participant information for all 17 participants.	14
2.2	Results for each participant at each speed. Cadence is time between steps in seconds, L is the average step length in inches, and v is the average velocity of the participant in ft/s.	19
2.3	Results for right cycle time, c_r , and left cycle time, c_l for all cases. p -values are also shown from two-sample t-tests comparing the left and right distributions.	20
3.1	Frequencies (f_i), damping ratios (ζ_i), and Mode Complexity Factor (MCF) comparison between EMA and OMA. MACs shown are between the EMA and OMA estimated mode shape for each mode order.	36
3.2	Mean and standard deviation ($x \pm \sigma_x$) for the estimated natural frequencies (f_i) and damping ratios (ζ_i) of the four modes along with the correlation coefficients ($r_{E,x}$) for each estimated modal parameter with average hallway acceleration RMS (RMS). Correlations with an asterisk (r^*) are statistically significant with a value $p \leq 0.05$	38
3.3	Summary of each assumption, providing relevant algorithms, sources of error when making the assumptions, considerations which affect the applicability of the assumption, advantages to the assumption, and recommendations from the authors for each assumption in light of the investigations in this work.	57
4.1	Each of the nine sets of admissible functions, along with the computational cost to generate each function, and if the set meets the geometric boundary condition at $x = 0$, BC_0 , and/or the natural boundary condition at $x = 1$, BC_1 , and if the set qualifies as quasi-comparison functions, QCF.	67
4.2	Maximum polynomial interpolation error for multiple polynomial orders for two functions. The same functions and polynomial spacing from Figure 4.6.	70
5.1	Polynomial order, number of outputs N and number of inputs M used in the comparisons in Figure 5.7 and Figure 5.8.	89

List of Abbreviations

CRI	Continuous Residue Interpolation
FE	Finite Element
FRF	Frequency Response Function
GRF	Ground Reaction Force
IMMAT	Impedance Matched Multi-Axis Testing
IoT	Internet of Things
PDE	Partial Differential Equation
TDoA	Time Difference of Arrival
VBOI	Vibration-based occupant inference

Chapter 1

Introduction

This chapter introduces vibration-based occupant inference (VBOI) and motivates the application of physically informed data-driven models to VBOI algorithms. A summary of contributions and an outline of this dissertation are also provided.

1.1 Smart Buildings and VBOI

Society is becoming more connected with the rise of the “Internet of Things” (IoT) [77]. One branch of IoT research has been with smart buildings: buildings with integrated sensors which allow a new level of insight into how the building responds to its environment [33]. One example of a smart building is Goodwin Hall on Virginia Tech’s campus, which is the most instrumented public building in the world for vibrations. Goodwin Hall has 225




Figure 1.1: Left: A picture of Goodwin Hall. Right: A diagram of the fourth floor of Goodwin with sensor locations. The number on each sensor location refers to the number of axes measured. The single axis measurements are all in the vertical direction, and an example of how one of these sensor is shown in the bottom right.

high sensitivity accelerometers permanently mounted to the building's steel structure which are capable of measuring the building's vibratory response in real time. Figure 1.1 shows an image of Goodwin Hall along with an example of an accelerometer mounted to the underside of an I-beam on the fourth floor. Goodwin has served as a test bed for research ranging from structural health monitoring [97, 98] to VBOI algorithms [2, 9, 61, 88, 101, 117, 118].




Figure 1.2: Floor vibration (b) measured by two underfloor accelerometers as an occupant walks down a hallway from left to right (a). Two individual footstep responses (c,d) are shown of these events comparing a sensor near the impact and a sensor roughly 42 feet away.

Figure 1.1 also shows a layout of sensors on the fourth floor of Goodwin Hall. On this floor, there is a hallway that is particularly densely instrumented with sensors configured to measure the vertical acceleration of the floor. This hallway will be used in multiple studies throughout this dissertation. Figure 1.2 shows measurements from two sensors in this hallway as a single person walks down the hallway. The overall signal containing multiple steps in

Figure 1.2b shows that sensors closer to a given footstep see higher response magnitudes. The zoomed-in plots of single footstep events in Figure 1.2c-d shows the acceleration response to a single footstep from the two sensors.

1.1.1 VBOI Algorithms

Vibration-based occupant inference (VBOI) algorithms utilize structural vibration measurements like those in Figure 1.2 to infer information about the building’s occupants. VBOI algorithms have been developed for the purposes of energy conservation [79], occupant safety [59], and tracking occupant health [39, 61].

One well-researched VBOI application is that of localization [1, 2, 7, 30, 34, 71, 88, 90, 117]. These algorithms use vibration measurements from multiple sensors to estimate the location of footsteps, learning where the building’s occupants are. Vibration based localization fills a gap of indoor localization where GPS is less reliable, and may struggle to pinpoint which floor an occupant is on [108]. Other indoor localization methods such as cameras and motion sensors rely on line-of-sight, which may require many sensors to cover the entire building. Additionally, some localization methods like cameras, or device tracking with wifi present more privacy concerns than vibration based localization [19, 104]. For these reasons, localization has been one of the most researched VBOI applications. Localization algorithms can be broken up into four families as shown in Table 1.1: those that work by observing energy attenuation over distance, those that compare wave arrival times, those that observe dispersion of waves, and those that utilize predictive models. These localization algorithms have been demonstrated with different sensor configurations and in different buildings, but in general, current localization algorithms can estimate footstep locations with an accuracy between 0.15-1 meter.

Table 1.1: Summary of selected VBOI localization algorithms. The number of sensors, sensor type (A=accelerometer, G=geophone), and the distance error reported are included.

Family	Method	# Sensors	Type	Error (m)	Ref
Energy	Weighted Sensor Coordinates	6	A	0.85	[1]
	Energy Attenuation over Distance	4	A	0.25	[2]
Time	Time Difference of Arrival (TDoA)	12	A	0.70	[88, 90]
	Sign Only TDoA	9	A	0.46	[7]
	Wavelet Transformed TDoA	9	G	0.34	[71]
Dispersion	Warped Frequency Transform Correlation	15	A	0.15	[117]
Model Based	Error-Domain Model Falsification	4	A	n/a	[34]
	Force Estimation and Event Localization	4	A	n/a	[30]

Another VBOI application is classification. The classification of occupants into male and female groups has been demonstrated with high accuracy [9], as well as the identification of individuals [80]. Classification has also been applied to occupant safety, where the classification of firearm type has been demonstrated with only vibration measurements of the platform on which the shooter is standing [59]. The combination of classifiers and localization algorithms paints the picture of identifying threats within a building, being able to track the locations of suspects and building occupants, and passing this information on in real time to emergency services. Classification and localization could also be combined to develop algorithms to track building pattern usage for energy savings. It is estimated that with better occupant based energy management in buildings could lead to 10-40% savings on energy [75].

Another area of VBOI research is in health applications such as gait analysis [39, 61]. Gait analysis is the study of how people walk, or their gait, and how this can relate to their health [96]. For example, gait variability is linked to fall risk in elderly adults [49, 50]. However, gait analysis is often restricted to infrequent visits to specialists with specialized tools [23]. There is a growing body of work in VBOI algorithms for gait analysis, that could be applied to passively record gait information to track the health of occupants in a nursing home for example. One of the contributions in this dissertation is the demonstration of extraction of temporal gait parameters from structural vibration measurements [61]. This work, along with other VBOI algorithms such as estimating force magnitudes [39], point to a future possibility of moving gait analysis out of the clinics and into the real world.

1.2 Background for VBOI Algorithms

In order for VBOI algorithms to be able to infer information about occupants, there must be some information transferred from the occupants into the building. When people walk, they input a force into the floor known as the ground reaction force (GRF) [91]. During walking, the GRF often exhibits a bimodal shape with two peaks, and a maximum force around 1.25 times the person's bodyweight [106]. The GRF is dominated by low frequency content [106], and therefore the response is also dominated by low frequency waves [37]. The force input from the GRF into the floor is filtered through the floor's structural dynamics, resulting in a response as shown in Figure 1.3. There are multiple ways that VBOI algorithms view the dynamics in order to infer information about occupants from the vibration response. An overview of three common types of VBOI algorithms along with how they model occupant induced vibration in buildings is provided in this section.




Figure 1.3: Graphical representation of how input forces from footsteps induces a response in the building's structure.

1.2.1 Input/Output relationship based algorithms

One method for extracting information from a measured response is by using a model which relates inputs to outputs. The transfer function, for example, relates the Fourier transforms of the input and output:

$$X(s) = H(s)F(s) \quad (1.1)$$

where $X(s)$ and $F(s)$ are the Fourier transforms of the response and input force respectively, and $H(s)$ is the transfer function between the input location of the footfall and the output location of the sensor. Often in this dissertation, instead of the transfer function, $H(s)$, we will discuss the frequency response function (FRF), $H(\omega)$. The FRF defines the same input/output relationship, and is the same as the transfer function but only sampled on the imaginary axis where $s = i\omega$. There are examples of VBOI algorithms using FRFs from Finite Element models [34] as well as directly measuring FRFs [30]. However, accurate FE models are not always available, and algorithms built on measured FRFs are limited to the locations at which the FRFs were measured. Beyond algorithms using transfer functions, there have also been VBOI algorithms which use regression to build a relationship between inputs and outputs [39].

Classification algorithms also leverage the input/output relationship to group occupants. Depending on how a person walks, the magnitude and frequency content of their input force $F(\omega)$ will change. If there are similarities between the input forces within a group, then there will also be similarities in the measured response $X(\omega)$. This is the principle that underpins VBOI classification. However, the relationship between $F(\omega)$ and $X(\omega)$ is only consistent when $H(\omega)$ does not change. This means that classification algorithms only work in the same location that they are trained, and do not generalize well to new locations which have different transfer functions.

1.2.2 Energy based algorithms

Some other VBOI algorithms view occupant induced vibration from an energy perspective. Spikes in signal energy in a building are indications of occupant activity and can be used for coarse occupant detection [79]. Some algorithms also assume a model of energy attenuation with distance. For instance, it may be assumed that a source event with energy E_0 will decay exponentially while travelling away from the source location:

$$E_i(r) = E_0 e^{-\beta r}, \quad (1.2)$$

where E_i is the measured energy at a distance r from the source, and β is a decay coefficient that is intrinsic to the building. Localization algorithms have been built around this idea of energy attenuation, where energy observations are made from a set of sensors, and an optimization problem is solved to find the source energy and source location [2]. Another heuristic energy-based localization algorithm is built on the idea of higher measured energy signals being closer to the footstep by performing a weighted average of sensor locations to estimate the event location [1].

1.2.3 Waveguide based algorithms

A final group of methods we will discuss are those that treat occupant induced vibration as propagating waves. These algorithms focus on the travelling of waves generated by events such as footsteps, and view the dynamics as wave propagation through a waveguide. For example, comparing the difference in wave arrival times at a group of sensors is how Time-Difference of Arrival (TDoA) localization algorithms work. TDoA assumes that waves are travelling at a constant speed in the medium, and therefore sensors which detect the wave earlier are closer to the source. A group of time difference comparisons from a set of sensors allows the estimation of the source location via multilateration [88]. TDoA was developed and is generally applied to non-dispersive media such as air. However, elastic waves in solid media undergo dispersion: different frequency components of waves travel at different speeds [44]. In order to address the challenge of dispersion in TDoA, various modifications have been made to the TDoA algorithm for use in VBOI, a selection of which are mentioned in Table 1.1 [7, 71].

VBOI algorithms which explicitly consider the dispersion relation, $\mathcal{D}(\omega, k)$, of the waveguide have also been developed. Methods such as the warped frequency transform can be used to leverage the dispersion of waves to build a highly accurate localization algorithm [117], however the dispersion relation of the system must be known. Different wave modes have different dispersion relations [44]. For example, considering a floor as a thin beam and starting from the partial differential equation (PDE) for an Euler-Bernoulli beam leads to the dispersion relation:

$$\mathcal{D}(\omega, k) = -\omega^2 + \frac{EI}{\rho A} k^4 = 0, \quad (1.3)$$

which relates the wave frequency ω and the wave number k , and depends on the system properties: stiffness E , moment of inertia I , density ρ , and area A . From the dispersion relation, the phase velocity c_p can be defined:

$$c_p \equiv \frac{\omega}{k} = \pm \sqrt{\frac{EI}{\rho A}} k. \quad (1.4)$$

This phase velocity will determine how the wave morphs as it propagates. Higher frequency waves (with a higher wave number and shorter wave length) have a higher velocity, and travel faster. This can be seen in Figure 1.2c-d in the difference between the responses at sensors near and far from the same event. All waves arrive at the same time at a sensor near the footstep. However, at a sensor further away, the higher frequency waves have been travelling faster and therefore arrive before lower frequency waves. This has the effect of spreading out the wavefront, and making it hard to define an arrival time for TDoA.

1.2.4 Common assumptions in VBOI

It is difficult to account for the complex phenomena present in buildings such as reflections, attenuation, and dispersion all at once. For this reason, each of the types of VBOI algorithms mentioned in this section make some assumptions in order to build a simpler model of how occupant inputs turn into measured responses. Three common assumptions are:

Assumption 1: Steady-state standing wave behavior is induced by footsteps

Assumption 2: Floors behave like infinite plates with uniform attenuation

Assumption 3: Dispersion relations for waves are neglected or simplified

In Chapter 3 these three assumptions will be examined in-depth, however each are briefly introduced here.

Assumption 1 This assumption allows the use of linear methods such as modal superposition, and modal analysis for FRF measurement. The use of measured FRFs [30] makes the assumption that waves induced in the testing procedure propagate out to boundaries, reflect, and generate standing waves in the building. However, with high attenuation within the building and non-linear effects of joints the assumption of consistent steady-state behavior is not always observed in buildings.

Assumption 2 Algorithms that propose a model to account for attenuation over distance [2] make the opposite assumption to Assumption 1. Instead of relying on reflections to generate steady-state standing waves, reflections are ignored and only attenuation is considered. This assumption operates as if there are no boundaries, and waves are propagating through an infinite plate. However, walls, columns, and floor geometry do create reflections in buildings causing attenuation to never perfectly follow a monotonically decreasing model.

Assumption 3 Dispersion presents a significant challenge for time-based VBOI algorithms. Some algorithms ignore dispersion [90], while still achieving viable results. Other methods that deal with dispersion only consider waves in a small frequency range [71], or assume a simple dispersion model such as thin-plate theory [117]. However, anything short of accounting for the true dispersion relation of the building is missing some of the information in the response.

Each assumption makes a simplification of one or more of the main challenges in modelling complex systems like civil structures: reflections, attenuation, and dispersion. The ideal situation for VBOI is to have a model of a real world building which can account for all three of these phenomena and their interactions accurately. However, in practice this is nearly impossible.

1.3 Data-Driven Modelling for VBOI

Most current VBOI algorithms could be considered to have a data-driven model of the building system at their core. This comes in the form of a preliminary calibration stage, or an adaptive refinement of system parameters. However, due to the assumptions made in developing these algorithms, the calibration stage only builds a simplified model of the system. There is a large literature of data-driven modelling for dynamical systems which may be able to further the current state of the art in VBOI algorithms.

1.3.1 VBOI Algorithm Calibration

It is common for VBOI algorithms to need a calibration stage before the algorithm can perform its inference. For example, a TDoA algorithm needs to have a wave speed to transform from time to distance [90]. Energy methods also require the calibration of an attenuation term [2], although calibration-free methods are also possible [1]. If dispersion is to be explicitly considered, the dispersion relation must be provided to the algorithm. Dispersion requires multiple things in the calibration step: the wave mode must be chosen and system parameters must be calibrated. For example, Woolard et al. [117] chose to use thin-plate theory, but also demonstrated that the localization accuracy is dependent on

system parameter calibration as well. In this same work Woolard et al. motivated their simplification from thick-plate theory to thin-plate theory, but this choice once again was part of the calibration. In Chapter 3 we will show that the dispersion relation for Lamb waves [44] may be more appropriate.

If FRFs from a FE model are used, it is likely that the FE model will need to be calibrated to match measured modal parameters of the system. If measured FRFs are directly used for a VBOI algorithm, then that measurement takes on the calibration role itself.

Some other VBOI algorithms remove the calibration step by utilizing an adaptive methodology. Instead of providing their TDoA variant with a wave speed, Mirshekari et al. [71] estimate a new wave speed for each event. The wave speed is varied such that the variance of the multilateration is minimized. Despite the removal of a prerequisite calibration step, these adaptive methods still may make the same underlying assumptions.

1.3.2 Opportunity for structural dynamics data-driven models

While current VBOI algorithms contain calibration procedures which learn some parameters of the building, there is still an opportunity to use more sophisticated data-driven modelling to get more complete models of the building. With better models, we can make fewer assumptions, and therefore have fewer limitations in scope and generalizability of VBOI algorithms. Some current data-driven methods that are relevant will be discussed, and a new physically-informed data-driven modelling technique which will be demonstrated in Chapter 5 will be introduced.

Taking a transfer function approach to relating inputs and outputs has the potential to be the most general starting point to build a model that can account for all complexities of building systems. By measuring FRFs, all information about attenuation, reflections, and dispersion are automatically incorporated into the measured data. There are rational approximation methods which can fit a state space model to measured FRFs [11, 48, 67, 74]. This state space model could then be used in a VBOI algorithm such as a localization algorithm. However, one weakness of measured FRFs are that they may change with changes in the system, such as non-linear amplitude effects or changes in system mass or stiffness properties. To address these problems, there are also parametric rational approximation methods which could be used to build a data-driven model that can account for changes in system parameters. A parametric state space model can be represented as

$$\dot{\mathbf{x}}(t) = \mathbf{A}(p)\mathbf{x}(t) + \mathbf{B}(p)\mathbf{u}(t); \quad \mathbf{y}(t) = \mathbf{C}(p)\mathbf{x}(t), \quad (1.5)$$

where $\mathbf{y}(t) \in \mathbb{R}^N$ are the responses, $\mathbf{u}(t) \in \mathbb{R}^M$ are the inputs, $\mathbf{x}(t) \in \mathbb{R}^\rho$ are the system states, and $\mathbf{A}(p) \in \mathbb{R}^{\rho \times \rho}$, $\mathbf{B}(p) \in \mathbb{R}^{\rho \times M}$, and $\mathbf{C}(p) \in \mathbb{R}^{N \times \rho}$ are the state matrices which may depend on some parameter p [18, 63].

Another limitation of using measured FRFs for VBOI is that they are tied to the locations which were tested. Here, we could look to another current data-driven technique of mode expansion. There are various methods that use high dimensional mode shapes to expand low dimensional measured data across a full spatial domain [24, 78]. These expansion methods, however, are restricted to expanding measured responses or mode shapes. It would be more helpful to build a predictive model with continuous spatial resolution directly from measure transfer functions.

The final contribution in this dissertation is a data-driven method conceived to overcome the spatial limitation of using measured FRFs for VBOI. Continuous residue interpolation (CRI) is designed to take FRF measurements at a discrete number of testing locations, and create a predictive model with continuous spatial resolution. The fitted CRI model can be used to simulate the response at any location to an input at any other location. This method provides a data-driven model that can be used to build VBOI algorithms such as localization.

1.4 Contributions

There are three main contributions which will be presented in this document:

Contribution 1: Gait analysis via underfloor accelerometers

Contribution 2: Investigation of assumptions in VBOI

Contribution 3: Continuous Residue Interpolation

The first contribution is the demonstration of extracting gait parameters from underfloor accelerometer measurements. This along with other VBOI gait analysis works paves the way for widespread adoption of passive occupant health monitoring.

Through the experience of developing VBOI algorithms, and reviewing the VBOI literature, the common assumptions made in VBOI algorithms along with their limitations became clear. Therefore, the second contribution is a detailed investigation of how well these assumptions match a real world building, and when they break down.

Finally, the third contribution is a new physically-informed data-driven modelling technique, inspired by the desire to be able to make a model of complex physical systems that can capture the system behavior over a continuous spatial domain with fewer compromises. The story of this dissertation is brought full circle by demonstrating how this data-driven technique can be applied to a simulated building to create a VBOI localization algorithm.

1.5 Outline of Dissertation

In Chapter 2 a VBOI algorithm for extracting gait parameters is detailed (Contribution 1).

In Chapter 3 a detailed investigation of current common assumptions in VBOI is presented (Contribution 2).

Chapter 4 gives some background on current methods for creating predictive models of dynamical systems. This chapter provides background and motivation for the final contribution.

In Chapter 5 a physically-informed data-driven modelling technique is introduced and demonstrated (Contribution 3).

Finally, Chapter 6 summarizes all contributions and takes a look forward at the future of VBOI and continued role data-driven modelling in improving VBOI algorithms.

Chapter 2

Gait Analysis via Underfloor Accelerometers

This chapter investigates the viability of a novel, low-cost approach to gait analysis based on floor mounted accelerometers.

2.1 Introduction

In the proposed approach to gait analysis, a section of a corridor is instrumented with accelerometers; these sensors detect the floor vibration resulting from footsteps of a participant walking in the corridor. The study takes place in Virginia Tech’s Goodwin Hall, which is instrumented with 225 highly sensitive accelerometers. One hallway is particularly dense with accelerometers with over 15 sensors mounted to measure vertical acceleration in a 115 ft span. The high density of sensors will allow methods with many or few accelerometers to be compared using the same data. Information learned from this test-bed could then be applied to inform future smart infrastructure layouts whose aim is to analyze the gait of the building’s occupants. The work in this chapter expands the scope of human-building interaction by monitoring and keeping track of occupants’ health. The ability to reliably extract various gait features is investigated by studying the vibratory signals detected by the sensors. Signal-energy based algorithms detect the heel-strike of each step during trials. From the detected footsteps gait parameters such as the average stride length, the time between steps, and the step signal energy are calculated. Distributions for these parameters are investigated, and the variation of these parameters between the left and right foot was compared for seventeen participants in the study. The distribution of temporal parameters using the proposed approach is also compared to traditional gait analysis measurement techniques.

Previously, buildings instrumented with accelerometers were also used to detect footsteps, localize events [71, 88, 101], and classify various sources of vibrations [9, 118]. The gait analysis method presented in this chapter is distinct from those performing gait analysis with inertial sensors because it does not rely on a wearable device. Wearable devices suffer from the necessity of the user to remember to use it and make sure it is charged and operating [87]. In this chapter, temporal gait parameters are estimated from sensors which are a

permanent fixture in the building, and therefore are not reliant on the user to ensure they are functioning properly. Goodwin Hall is an instrumented test-bed that has allowed researchers to develop algorithms to monitor the condition of the building [98], to develop localization algorithms for noisy-dispersive environments [88], and to study human-structure interaction for classification [9]. This chapter takes the study of human-structure interaction a step further and seeks to use building mounted sensors to detect occupant health by translating footsteps into gait parameters.

2.1.1 Traditional Gait Analysis Techniques

Gait analysis plays a pivotal role in diagnosing and monitoring human health [51]. Traditionally, in a clinical setting, physical therapists and activity therapists observe deviations in gait parameters to assess the physical condition of patients [114]. Many studies have demonstrated the significance of human gait in monitoring human health with conditions such as diabetics [99], rehabilitation [8, 31], pathological diseases [26, 116], neurological abnormalities [25, 42], alcohol intake [57] and aging [116]. Although the significance of gait analysis in monitoring human health condition is evident from years of research, penetration of technology into gait analysis has been mostly limited to research labs [120]. Specialists such as physical therapists evaluate the quality of gait by observing a participant walking in a clinic. Such a subjective approach to evaluating gait has many limitations including keeping track of minor changes in gait over time and estimating the progress of intervention techniques [96, 119]. Recent technological advancements have made it possible to quantitatively monitor the kinematic motion of humans. Multiple three-dimensional motion capture systems accurately capture human kinematics to precisely track the marked locations over time [102]. Such measurement techniques determine the orientations and trajectories of multiple joints in the human body that provide the capability to quantitatively monitor multiple features of human gait. Since establishing such a system in a clinical setup is very expensive and time-consuming to test each participant [110]; this has given rise to the development of low-cost wearable or non-wearable technology focused towards gait analysis. Techniques involving image processing with high-speed cameras [73] and Microsoft Kinect [86] has shown considerable promise in observing some gait features in a lab setting.

Additionally, multiple studies have also shown that wearable inertial sensors have the potential to detect falls and walking patterns [15, 16, 72]. Although most of these techniques do not provide the complete kinematic information that a 3D motion capture system does, they are easier to implement outside of a laboratory setting. This allows monitoring of useful features more often than relying on time-consuming and invasive full motion capture systems.

2.2 Outline of Gait Analysis Process

This section will present a procedure for using one or multiple underfloor accelerometers to extract gait parameters from real measurements of 17 participants walking down a hallway. First, trials were run for each participant at three different speeds. The acceleration measurements were then used to detect steps, and the times of these detected steps were used to extract the gait parameters.

2.2.1 Participants

A total of 17 healthy participants aged between 18 and 27 took part in the study, 9 female and 8 male. The participants' mean height is $\mu = 67.8$ in ($\sigma = 4.3$ in) and mean weight is $\mu = 165$ lb ($\sigma = 39$ lb). The participants took part in the study after informed consent. Further information such as height, weight, leg length, dominant foot, and sex about the participant demographics are available in Table 2.1.

Table 2.1: Participant information for all 17 participants.

Participant	Sex(M/F)	Height (inch)	Weight (lbs)	Leg length (inch)
Participant 1	M	67.50	173.0	39.00
Participant 2	M	73.20	250.6	41.00
Participant 3	M	71.30	196.8	41.75
Participant 4	F	68.50	159.0	41.00
Participant 5	M	69.00	165.0	40.00
Participant 6	F	62.75	165.0	37.25
Participant 7	M	74.00	165.4	43.00
Participant 8	F	70.00	127.4	37.00
Participant 10	F	62.00	215.0	38.00
Participant 11	F	63.00	141.0	38.00
Participant 12	M	75.00	200.8	43.50
Participant 13	M	65.00	98.2	40.00
Participant 14	F	66.00	117.0	39.00
Participant 15	F	61.50	123.6	36.50
Participant 16	F	65.0	141.6	39.50
Participant 18	F	67.50	199.0	39.00
Participant 19	M	71.75	170.0	41.00

Data from participant 9 and 17 were discarded as a result of high noise in the data.

2.2.2 Experiment and Data Processing

This experiment was conducted at Goodwin Hall, a building which is instrumented with 225 high sensitivity accelerometers on the Virginia Tech campus. A layout of the location of the sensors in a fourth-floor hallway is shown in Figure 2.1 (a). In the present study, participants walked along a marked 115 ft portion of this hallway while 17 accelerometers which are permanently mounted under the floor of the hallway recorded acceleration data at 32,768 Samples/sec. For each trial, participants lined up at the starting line. A hammer strike to the floor marked the beginning of the trial, and the participant began walking with their dominant foot first. As the participant passed the finish line, another hammer strike marked the end of the trial (the magnitude of the hammer strikes are easily distinguished from any steps or noise and are used to break apart the data sets). After completing one direction, the participant lined up at the new start line and the process was repeated. A total of four trials were recorded for each participant at three self selected speeds: a walking speed, a brisk walking speed, and a running speed. Participants performed all trials in socks to remove the effect of shoe sole on the acceleration responses. Sample data from a trial is shown in Figure 2.1(b). All the trials were conducted during weekends and evenings so that the noise from ambient activity in the building was minimized.

As a participant walked towards an accelerometer, the amplitude of the floor vibrations due to heel-strikes gradually increased to a maximum and then reduced as the participant moved away. This is seen in Figure 2.1(b), and is expected as the waves propagating away from the heel-strike lose energy exponentially with distance [2]. As a result of this phenomenon, not all accelerometers may capture all heel-strikes in a 115 ft long hallway.

Previous literature has shown that seismic sensors such as accelerometers perceive heel-strikes as impulse contact, and toe-slaps as sliding (friction) contact with the floor [37, 38, 95]. As a result, these studies have seen a distinctive peak in the vibration signal during each heel-strike. There are many approaches to detect heel-strikes on such vibration signatures that are typical of a heel-strike [90]. In this work, a signal energy based custom algorithm identifies steps by searching for peaks in energy. The highest acceleration energy occurs immediately following the heel-strike. Finding where all the energy peaks are in the signal gives the approximate timing of the beginning of each step. An example of an entire acceleration signal with steps detected is included in Figure 2.1 (b).

The algorithm broke up the acceleration signal from a single sensor into windows approximately 0.025 seconds long. The root mean square (RMS) value represents signal energy as a measure acceleration in a window. The ratio of each window's RMS was taken with the window directly preceding it to find the instant when a heel-strike occurs. A heel-strike is a high energy event, which is directly preceded by very low signal energy noise. Therefore the ratio of a window's energy with the window before will be high at each heel-strike. Figure 2.2 shows an example acceleration signal containing two steps, along with the window energy ratio over time. The ratio represents the signal-to-noise ratio (SNR) of the signal, and the steps were found by searching for peaks in the SNR and correlating it to the time vector.




Figure 2.1: (a) Diagram of all sensors along the hallway used for testing. (b) An example of an acceleration history from a single trial with all steps detected. The trial comes from the sensor marked in blue in (a).

Since each participant started with their dominant foot, each heel-strike can be identified as a left or a right foot impact.




Figure 2.2: An example of an acceleration measurement containing two steps, the corresponding signal-to-noise-ratio (SNR), and the detected step times based on the peaks in SNR.

2.3 Gait Parameters from Structural Measurements

2.3.1 Effect of number of sensors

Figure 2.3 shows the distributions of step intervals (time between a step and the following step) for two different participants: Participant 12 was a tall male who was walking at a brisk pace, while Participant 16 was a short female who was traveling at a walking pace. These two cases were chosen because the steps in the trial for Participant 12 created a higher magnitude of acceleration (due to participant height, weight, and speed), while the trial for Participant 16 had less acceleration magnitude. Comparing methods for both of these cases will ensure that the step detection method works for trials with high and low signal-to-noise-ratios. For each participant, all three approaches were used to find the distribution of step intervals over the full trial data and the half trial data.

As Figure 2.3 shows, when the whole trial was used there is a larger spread (standard deviations of 0.048s, 0.039s, 0.034s for Participant 12 for each method) compared to the when the middle half of the trial was used (0.013s, 0.025s, 0.015s for the same trials). The standard deviation of the half trial is less than the full trail because the data artifacts resulting from starting or ending the trial are avoided. The variation of the step interval estimates based on half interval is higher when data from all 17 sensors is considered as it included estimates from sensor farther away with accelerometer signals containing higher noise. With data from the half trial, there is no significant difference in estimating step-interval values from the single sensor (Approach i) and the data put together by choosing the best sensor for each step (Approach iii).

Using these observations it is clear that in this application only one sensor is needed over the course of the middle half of the trial. Using only one sensor is also the simplest case, and easiest to implement. Therefore, Approach i was chosen to report all statistics on step interval or stride cycle time in this work. It should be noted that the dynamics of the floor on which the sensor is installed will have an effect on how far away from that sensor it can accurately detect a footstep. In this case, one sensor was used in the middle of a span of about 50 feet. Further modeling and testing would be needed to extrapolate these results to other surfaces.




Figure 2.3: Step intervals for one trial each of Participant 12 (left) and Participant 16 (right). (a) and (b) correspond to approach (i), (c) and (d) correspond to approach (ii), and (e) and (f) correspond to approach (iii). The distributions using each approach are shown as histograms, while normal probability distribution fits to each histogram are also shown for reference. Beside each fit are the mean, μ , and standard deviation, σ .

2.3.2 Analysis of Normal Gait

Once the steps were located, all gait parameters were estimated and are tabulated in Table 2.2. The step length and step interval parameters are shown in Figure 2.4, with female participants on the left and male participants on the right. The data from the single sensor (Approach i) was used to estimate these gait parameters from all trials. For each participant, the step interval estimates and the step length estimates for all trials were averaged at each of the three speeds. The error bars for the step interval estimates represent the standard deviation of the distribution. In Figure 2.4, the parameters for each participant at each speed is shown connected by lines. The size of each marker in Figure 2.4 represents the highest ratio of the step signal energy to noise floor for that participant at that speed.

Figure 2.5 shows the cycle time and step signal energy for the left and right legs of Participant

Table 2.2: Results for each participant at each speed. Cadence is time between steps in seconds, L is the average step length in inches, and v is the average velocity of the participant in ft/s.

Participant	Walking			Brisk Walking			Running		
	Cadence (s)	L (in)	v (ft/s)	Cadence (s)	L (in)	v (ft/s)	Cadence (s)	L (in)	v (ft/s)
1	0.51 ± 0.012	27	4.25	0.44 ± 0.017	31	5.53	0.36 ± 0.018	40	8.72
2	0.6 ± 0.007	33	4.38	0.5 ± 0.007	38	5.86	0.4 ± 0.01	54	10.36
3	0.59 ± 0.029	25	3.5	0.49 ± 0.008	35	5.62	0.35 ± 0.008	52	11.66
4	0.51 ± 0.012	26	3.98	0.47 ± 0.008	28	4.83	0.35 ± 0.005	36	8.13
5	0.49 ± 0.01	27	4.39	0.42 ± 0.006	35	6.37	0.33 ± 0.008	47	10.94
6	0.58 ± 0.023	24	3.3	0.45 ± 0.01	27	4.74	0.35 ± 0.01	35	7.78
7	0.48 ± 0.011	30	4.98	0.42 ± 0.011	34	6.68	0.36 ± 0.012	49	10.42
8	0.51 ± 0.012	30	4.62	0.44 ± 0.011	33	5.97	0.36 ± 0.012	46	10.04
10	0.54 ± 0.017	22	3.34	0.4 ± 0.008	32	6.26	0.33 ± 0.012	40	9.31
11	0.48 ± 0.013	23	3.87	0.42 ± 0.011	27	5.15	0.31 ± 0.013	32	8.21
12	0.53 ± 0.012	33	4.9	0.46 ± 0.011	38	6.48	0.37 ± 0.005	63	12.71
13	0.51 ± 0.014	26	4.09	0.46 ± 0.012	30	5.22	0.36 ± 0.012	50	10.56
14	0.52 ± 0.011	27	4.24	0.43 ± 0.011	31	5.74	0.39 ± 0.013	45	8.69
15	0.51 ± 0.012	24	3.72	0.41 ± 0.012	27	5.15	0.31 ± 0.011	32	8.23
16	0.48 ± 0.01	27	4.45	0.37 ± 0.008	29	6.28	0.32 ± 0.005	33	8.02
18	0.52 ± 0.015	30	4.46	0.48 ± 0.011	33	5.35	0.35 ± 0.012	35	7.85
19	0.55 ± 0.026	26	3.78	0.45 ± 0.017	31	5.59	0.36 ± 0.018	51	11.03

12 and 16. The cycle time and step signal energy were calculated for all speeds, and the mean estimates and the corresponding standard deviations are presented along with the fitted distributions. Figure 2.5 (a) and (c) visually compares the cycle time estimates and spread for the left and the right feet. Additionally, to quantitatively monitor the acceleration signals generated due to the force of the step impacts, Figure 2.5 (b) and (d) compares the resulting ratio of the signal energy to noise floor for the left and the right feet.

The distributions for cycle times were calculated for each foot of all participants at each speed. A two-sample t-test was carried out to compare the left and right foot at each combination. The null hypothesis was that the two distributions had equal means and variances. The results of all cycle time distributions and two-sample t-tests are shown in Table 2.3.

As seen in Figure 2.4, as each participant traveled at a faster pace, the step interval time went down and the relative step length went up as expected. This shows that to go faster, all participants took longer strides at quicker intervals which agrees with previous research on walking speed patterns in adults [45]. Another observation is that as participants traveled faster, there was more signal energy in their steps; represented as the size of each marker, where larger signifies more energy. This is also expected as more force is being imparted into the floor as the participants are running as compared to walking [76]. Although in this

Table 2.3: Results for right cycle time, c_r , and left cycle time, c_l for all cases. p -values are also shown from two-sample t-tests comparing the left and right distributions.

Participant	Walking			Brisk Walking			Running		
	c_r (s)	c_l (s)	p	c_r (s)	c_l (s)	p	c_r (s)	c_l (s)	p
1	1.02 ± 0.045	1.03 ± 0.046	0.95	0.88 ± 0.041	0.88 ± 0.038	0.92	0.72 ± 0.048	0.72 ± 0.057	0.85
2	1.19 ± 0.02	1.18 ± 0.037	0.71	1.01 ± 0.023	1.01 ± 0.023	0.89	0.8 ± 0.027	0.8 ± 0.018	0.56
3	1.15 ± 0.119	1.16 ± 0.147	0.92	0.98 ± 0.014	0.98 ± 0.019	0.96	0.68 ± 0.015	0.68 ± 0.019	0.98
4	1.03 ± 0.044	1.03 ± 0.05	0.76	0.94 ± 0.027	0.94 ± 0.023	0.87	0.69 ± 0.024	0.69 ± 0.027	0.97
5	0.99 ± 0.048	0.98 ± 0.018	0.72	0.83 ± 0.016	0.83 ± 0.017	0.79	0.65 ± 0.024	0.65 ± 0.024	0.9
6	1.14 ± 0.126	1.14 ± 0.122	0.95	0.9 ± 0.032	0.9 ± 0.037	0.9	0.69 ± 0.022	0.69 ± 0.024	0.98
7	0.96 ± 0.026	0.96 ± 0.02	0.71	0.83 ± 0.018	0.83 ± 0.02	0.94	0.72 ± 0.014	0.72 ± 0.013	0.78
8	1.03 ± 0.039	1.04 ± 0.04	0.76	0.87 ± 0.035	0.87 ± 0.03	0.92	0.71 ± 0.024	0.71 ± 0.015	0.92
10	1.07 ± 0.109	1.07 ± 0.096	0.81	0.8 ± 0.034	0.8 ± 0.032	0.99	0.67 ± 0.048	0.67 ± 0.055	0.96
11	0.96 ± 0.052	0.96 ± 0.056	0.91	0.83 ± 0.036	0.83 ± 0.034	0.97	0.62 ± 0.048	0.63 ± 0.069	0.79
12	1.06 ± 0.019	1.06 ± 0.017	0.92	0.91 ± 0.028	0.91 ± 0.016	0.69	0.73 ± 0.018	0.73 ± 0.018	0.81
13	1.06 ± 0.072	1.05 ± 0.061	0.74	0.93 ± 0.047	0.94 ± 0.046	0.72	0.72 ± 0.02	0.72 ± 0.032	0.79
14	1.03 ± 0.029	1.03 ± 0.023	0.69	0.86 ± 0.022	0.86 ± 0.021	1	0.78 ± 0.027	0.78 ± 0.017	0.84
15	1.01 ± 0.072	1.01 ± 0.086	0.74	0.83 ± 0.03	0.83 ± 0.026	0.88	0.61 ± 0.025	0.61 ± 0.022	0.8
16	0.96 ± 0.039	0.96 ± 0.042	0.9	0.74 ± 0.028	0.74 ± 0.034	0.73	0.64 ± 0.018	0.64 ± 0.017	0.92
18	1.05 ± 0.064	1.05 ± 0.076	0.79	0.96 ± 0.034	0.96 ± 0.053	0.79	0.68 ± 0.046	0.68 ± 0.045	0.98
19	1.1 ± 0.117	1.11 ± 0.119	0.82	0.89 ± 0.065	0.89 ± 0.066	0.58	0.71 ± 0.065	0.71 ± 0.057	0.8

study the extracted gait parameters were not compared to another measurement system, the agreement with accepted trends from the literature shows promise for this method.




Figure 2.4: Step intervals and relative step lengths for all participants. Females in the left plot (a), and males in the right plot (b). Marker size is related to the maximum energy in a step from each participant.

A powerful opportunity provided by this method is the ability to quantitatively compare gait parameters for each leg. Having trials containing 20 or more steps allowed comparisons between legs to be carried out quickly. In a healthy participant, it would be expected that there should be no difference in the time between consecutive right foot steps or consecutive left foot steps (cycle time) [21]. The cycle time distributions presented in Figure 2.5 for each

foot have equivalent means down to the hundredth of a second for all speeds. The standard deviations are also very similar between each foot for all cases. Furthermore, out of all 52 combinations of participant and speed, the lowest p-value from the two-sample t-tests was over 0.5. Therefore, the null hypothesis is never rejected unless a very high threshold is chosen. Since the null hypothesis was that the two distributions have the same mean and variance, this means that there is no statistically significant difference between the cycle times for the left and right foot of any participant.

Similarly, it would be expected that there is no difference in the forces imparted on the floor by each foot. No difference in the forces imparted translates to no difference in the signal energy measured in the floor acceleration response induced by each foot. As shown in Figure 2.5 there is more spread in the step signal energy distributions because it is sensitive to the distance between the step impact and the sensor as can be expected due to the structural dynamics of the floor. Even with this limitation, the mean of the distribution for step signal energy is still very similar in all cases between the two feet. The standard deviation was not considered for step energy because the distributions are so heavily skewed.




Figure 2.5: Comparison of cycle times and step signal energy for Participant 12 (a, b), and Participant 16 (c, d), respectively. Walking trials are shown in blue, brisk walking in orange, and running in yellow. Fitted normal distributions are plotted for cycle times along with the mean, μ , and standard deviation, σ . Fitted Rayleigh distributions are plotted for step energy along with the mean.

2.4 Comparison with traditional gait analysis tools

The purpose of this section is to determine if floor mounted accelerometers are capable of capturing meaningful gait data that is comparable to other state of the art measurement techniques. If accelerometers are able to collect meaningful information about an individual's gait then they can potentially be instrumented in a setting that would be less invasive to a patient's everyday life patterns. It is hypothesized that there will be no significant difference in the time measured between consecutive steps by floor mounted accelerometers and other established gait analysis measurement systems.

2.4.1 Participants

There were three participants in this study. Two participants were female, while one was male. All three were healthy college aged participants with no gait abnormalities. The weights of the participants ranged from 627 N to 680 N.

2.4.2 Equipment and Test Set-up

Three different systems were used to measure the time between consecutive footsteps for each participant: in-shoe pressure sensors, floor-mounted accelerometers, and force plates.

The in-shoe pressure sensor system used was the Pedar-X system from Noble (Saint Paul, Minnesota). The pressure sensors had been calibrated up to 6 bar. For each participant, the correct sized insole was put in each shoe, and all participants wore the same style of shoe: Nike Air Pegasus 26s. Once the pressure sensors were put in each shoe, the system was calibrated for zero pressure by the participant alternatively picking each foot off the ground. For the Pedar system, data was recorded at 100 Hz.

Floor mounted accelerometers were also used to measure time between heel strikes. Four PCB 393B04 accelerometers (PCB Piezotronics, Depew, New York) were secured to the ground with wax. These accelerometers have a sensitivity of around 1000 mV/g; each accelerometer has a calibrated sensitivity marked from the factory. The four accelerometers were placed with two on each side of the path participants walked along. The accelerometers were centered 18 inches around the middle of the path, and accelerometers on the same side of the path were 48 inches apart. The two sides were offset along the direction the participants walked by 24 inches. The four accelerometers were connected to a PCB 482C signal conditioner, and then to a National Instruments data acquisition box (National Instruments, Austin, Texas). The system was set to measure at 19,200 Hz.

Three AMTI (Watertown, MA) Gen 5 force plates which are 600 mm by 600 mm were positioned in a row in the line where participants walked. All force plates were zeroed with

no load before measurements were taken. The force plates recorded data at a sampling frequency of 1,920 Hz.

2.4.3 Data Analysis

The accelerometer data was analyzed in the same way as other data in this chapter. The analysis of the Pedar and force plate data followed the same process, but were different than the accelerometer data. The raw data for each was used to detect heel strikes, no filtering was done because the signals were clean and any filtering would obscure the heel strike. For the Pedar system, the total force of each foot was output from the software. A heel strike was defined as when the force on a foot went from zero force (during swing phase) to greater than 10 N. A force threshold is commonly applied, and in fact is built into commercial softwares such as Visual-3D. This threshold method was also used for the force plate data, however there was an additional concern for the force plate data. A time between heel strikes data point was only recorded within a trial if two consecutive force plates surpassed the threshold without overlap from an incomplete footstep on one or more force plates.

2.4.4 Results and Discussion

Figure 2.6 shows an example of detecting two consecutive heel strikes using the Pedar and accelerometer data. Once consecutive heel strikes were detected from a given measurement system, the time between them was calculated. For the three participants there was small differences between the three measurement techniques as shown in Figure 2.7. Participant 2 showed a slightly higher time between heel strikes in accelerometers, however the mean was still within the standard deviations of the other collection techniques.

For all of the participants, the means of all three methods were within the standard deviations of the others. Participant 3 showed means that were nearly identical to one another. Across all of the methods, force plates showed the greatest standard deviation, due to the fact that there were only 3 force plates in a line, so for some participants timing between heel strikes had fewer than 10 total data points. Several times as well, force plate data had to go unused, as there were instances of a footstep overlapping plates. Pedar proved to be a very reliable and consistent method for detecting heel strikes with very little fluctuation in the measured time between heel strikes, which was expected since the Pedar system has been shown to be repeatable [92]. In some instances, the standard deviation for time between heel strikes for accelerometers was less than that of Pedar while their averages were comparable. Future fall patients showed variability in stance times of 30 ms [50], which is less than the standard deviations for time between heel strikes for all measurement systems in this work. With stance time being double the time of heel strikes, it can be assumed the stance time variability would be larger than the time between heel strikes, but not orders of magnitude. This gives confidence that accelerometers would be able to detect variability in stance times




Figure 2.6: Examples of detecting the heel strikes of two consecutive steps with (top) Pedar, and (bottom) accelerometers. Each detected heel strike is shown with a +.




Figure 2.7: Time between heel strikes for three different participants using three different measurement techniques. The mean is shown as the darker bar, as well as +1 SD as the lighter bar.

given the proper implementation strategy.

Floor mounted accelerometers appear to be a suitable measurement technology to capture

meaningful gait parameters. They show little to no difference in mean time between heel strikes from other common measurement techniques and have less variability in some participants. Accelerometers also appear to have much easier implementation strategies for cost and logistical reasoning than Pedar and force plate technologies. Accelerometers may encounter problems from the environment it is monitoring due to floor composition and environmental noise, however these can be mitigated with filtering.

2.5 Discussion

The structural vibration measurement based gait analysis process presented will not be able to fully replace traditional gait analysis, nor is that the goal of this chapter. Traditional methods such as 3D motion capture are capable of measuring joint kinematics with sub-millimeter accuracy, which is unlikely to be surpassed by any other method. However, gait analysis in instrumented buildings provide the possibility of supplementing current traditional methods. Most gait analysis currently taking place requires expensive equipment in a laboratory setting (3D motion capture, in floor force plates, etc.), or the supervision of trained professionals at a physical therapy office. In the case of a person recovering from a joint replacement or injury, there is no method of continuously monitoring patients between visits to specialists. A cheap, easily installed system such as that proposed in this chapter would be capable of continuously monitoring this patient's recovery in their home. Gait parameters such as cycle time and step energy could be continuously monitored, as well as the gait signature being monitored to ensure it is moving in the right direction. Another possible use-case for this type of gait analysis would be in elderly care facilities. A system continuously monitoring the health of those living in one of these facilities could track for changes in gait patterns of individuals which could be precursors for things such as fall likelihood, or types of abnormal gait correlated with Parkinson's (hypokinetic gait). In these types of cases, an automated gait analysis system such as that proposed could provide valuable information to monitor between specialist visits or monitor for signs indicating intervention is needed.

2.6 Conclusion

This chapter has demonstrated that certain gait parameters can be extracted using measurements from floor mounted accelerometers. In the hallway which was used, a single sensor was able to accurately detect the time of step occurrences over about 50 feet centered around the sensor. After detecting the occurrence of steps, step intervals, average step length, and step energies can be estimated from the acceleration time signals. Distributions for these parameters were investigated, and the variation of these parameters between the left and right foot was compared. As expected, no statistically significant difference was found between

the cycle times of the two legs.

Although this method of gait analysis shows promise, it is not without its limitations. Currently, only the heel-strike event can be detected. The timing of the toe off event is difficult to distinguish in the floor acceleration signal, meaning that at the current time this method cannot find stance time or swing time, only the entire time between the beginning of steps. Despite these limitations, there are clear advantages. Needing only a single sensor, it can be implemented in a cost effective manner. It also does not require any special set up, and therefore can be readily implemented on any floor. It is also not constrained to a short distance: this work shows a single sensor allows trials up to 50 feet while more sensors could increase that distance. Longer trials allow data to be collected more quickly since there is less time spent resetting and changing directions. Compared to traditional methods of gait analysis, this method requires much less data analysis and processing time, and can provide gait parameters with similar accuracy.

Floor mounted accelerometers present a novel method to approach gait analysis. The cost effectiveness and ease of implementation show promise. Once set up, the system also requires little post processing of data, and can automatically generate distributions of gait parameters and comparisons between feet.

Chapter 3

Investigation of assumptions in VBOI

This chapter provides an experimental and theoretical investigation of three assumptions commonly made in current VBOI algorithms.

3.1 Introduction

With the rise of the Internet of Things (IoT) and instrumented buildings, new techniques have been proposed for inferring information about building occupants through structural vibration measurements. Such techniques will be referred to in this work as vibration-based occupant inference (VBOI) algorithms. As a person walks in a building, each footstep impact generates elastic waves, which can be observed by an array of sensors, as shown by an example in Figure 3.1. The VBOI field has made advances in processing these types of measurements to detect and localize occupants [2, 71, 90]; however, processing approaches have primarily relied on heuristic algorithms, which make significant simplifying assumptions about wave propagation in floor structures. While these approaches can be calibrated to work in specific scenarios, large-scale implementation will require an improved understanding of the complexities of wave propagation in floor systems.

There are three main physical phenomena that dominate wave propagation behavior in finite structures: attenuation, reflections, and dispersion. By studying the current literature, the authors found that a large majority of VBOI algorithms make at least one of three main simplifying assumptions about these phenomena:

Assumption 1: Steady-state standing wave behavior is induced by footsteps

Assumption 2: Floors behave like infinite plates with uniform attenuation

Assumption 3: Dispersion relations for waves are neglected or simplified

Examples of each of these assumptions will be discussed in Section 3.2, and each assumption will be investigated in-depth in Sections 3.4-3.6 respectively.

However, it can be observed that real wave propagation in standard floor systems deviates significantly from these assumptions. This is illustrated by the floor vibrations observed in




Figure 3.1: Floor vibration (b) measured by two underfloor accelerometers as an occupant walks down a hallway from left to right (a). Two individual footstep responses (c,d) are shown of these events comparing a sensor near the impact and a sensor roughly 42 feet away.

Figure 3.1. In Figure 3.1b, it can be seen that although the acceleration signal magnitude generally decays with footstep distance from the sensor, the rate of decay may not be uniform (violating assumption 2). In Figures 3.1c and 3.1d, the shape of the vibration response observed for the same footstep at two different locations is significantly altered by reflections and dispersion (violating assumption 2 and some simplifications under assumption 3). These phenomena and others (such as non-linear joints or inhomogeneous materials) can potentially introduce large variability into algorithms based on simplified wave propagation behavior, hence there is a need to understand how phenomena manifest and when assumptions may or may not be justified. This would lead to methods that improve the performance of existing algorithms, making them more widely applicable and boosting the potential impact of the

field as a whole.

This chapter first reviews the current literature in VBOI algorithms. Next, it presents the methodology for a series of experiments which provide results for three investigations. Each investigation targets one of the assumptions mentioned above and is presented as an independent section. The objective of these investigations is to survey where current assumptions are acceptable and where future developments could be made to improve the generalization of these algorithms. Each investigation section will follow a similar format. First, the assumption is tested directly through a vibration experiment in a building. Next, a follow-up experiment is conducted to quantify the underlying physical phenomenon or phenomena impacting the assumption. Thirdly, a theoretical model is presented to understand the observed behavior and determine which parameters are the most influential. Finally, the findings are summarized and the implications for VBOI algorithms are discussed.

3.2 Current vibration-based occupant inference algorithms

One of the most well-researched topics in VBOI to-date is the localization of occupants based on the propagating waves created by their footstep inputs. Localization algorithms use a spatial array of sensors to measure vibration generated by a footstep and infer the impact location. The majority of localization algorithms can be broadly categorized into three groups: model-based methods, energy-based methods, and time-of-arrival based methods.

Model-based methods using Finite Element (FE) analysis or measured transfer functions have been applied to the localization problem [28, 34, 93]. These methods operate under the assumption that steady-state behavior is present, and therefore techniques such as modal superposition, linear amplitude scaling, or steady-state testing methods can be used. However, large attenuation rates may prevent floor systems from reaching steady-state in response to footsteps. That is, the energy is dissipated before the floor can reach standing wave behavior, meaning that transient behavior dominates the response. Furthermore, slight changes to the floor system dynamics due to operational or environmental conditions may render calibrated models or measured transfer functions invalid.

Energy-based algorithms leverage the observation that the waves attenuate with distance, allowing various signal energies from an array of sensors to estimate the location of the source [1, 2]. These algorithms assume that attenuation takes some form, such as exponential decay with distance, in order to infer the most likely source location for a footstep. Therefore, they assume uniform attenuation as would be the case in an infinite plate. However, with sensors at different distances from walls, columns, and other boundaries which create reflections, it is difficult to define a time interval over which the energy should be measured to most closely reflect “infinite plate” behavior. For this reason, algorithms of this type also assume that there are no reflections or that reflections can be separated from the incoming wave.

Time Difference of Arrival (TDoA) is a commonly implemented time-of-arrival based framework for localization which compares differences in the time that a wave reaches an array of sensors to localize the source [7, 71, 90]. TDoA was originally developed for non-dispersive media, such as radio waves in air, and therefore does not have a built-in method to account for dispersion. In the presence of dispersion, different frequency components of the signal travel at different speeds and arrive at different times. Therefore, it is difficult to define a distinct arrival time. Multiple modified versions of the TDoA algorithm have been presented to overcome these challenges, such as in [7] where only comparison is made on whether the wave arrived at one sensor before or after another sensor, and in [71] where only a narrow bandwidth is considered by using the wavelet transform to isolate frequencies with a consistent propagation speed. However, by making simplifying assumptions about the dispersion relation (or ignoring it), some information is lost. For example, by considering only narrow bandwidths, all other frequency information which could help average out noise is discarded. With a greater fundamental understanding of the dispersion relation, it is possible that the full signal could be considered with no information loss. For example in [117] dispersion compensation techniques are used directly for localization, which considers all wave frequencies. However, the dispersion relation of a thin-plate is used, which may simplify the true dispersion relation.

Another research topic in VBOI is the classification of occupants. Classification of individuals was shown in [80], while in [9] occupants were classified into male and female groups. Classification algorithms use time and/or frequency domain features of the structural response (such as those in Figures 3.1c and 3.1d) to determine which class to place a new footstep event into. The guiding principle of these classification algorithms is that the input forces from footsteps are consistent within a class, and this consistency will enforce similar responses in the floor from multiple footsteps by that class. This allows a classification algorithm to group together responses from the same class. However the presence of reflections can change the response depending on the location of the footstep, and dispersion alters the response as it travels. These challenges lead to an implicit assumption that there are no reflections or no dispersion if a single trained classifier is attempted to be generalized beyond the input and output locations where it was trained.

Other works have used footstep induced vibrations to demonstrate applications such as occupancy detection for energy conservation [79, 89] or health applications such as gait analysis [39, 61] and fall detection [29]. Once again, all of these applications must work around the fundamental challenges presented by wave propagation phenomena which often includes making simplifying assumptions.

3.3 Methods

The following investigations draw upon two series of experiments: observations of wave propagation and vibration in buildings on Virginia Tech’s campus and an in-depth experiment on


a connected concrete beam-column structure. In this section these experiments are discussed to give background to the data used to support investigations of each assumption.

3.3.1 Series 1: Observations of vibration in buildings

The main setting for this series of experiments is Goodwin Hall. Goodwin Hall is a 160,000 square foot five story building on Virginia Tech's campus. Goodwin serves as a test-bed for various instrumented building research topics ranging from structural health monitoring [98] to VBOI algorithms [2, 9, 61, 90, 117]. There are 225 high sensitivity accelerometers permanently mounted to the steel structure of the building, allowing continuous monitoring of overall building vibration as well as the vibration of the floors in response to occupant activity. The upper floors of Goodwin are made up of a concrete slab poured over a corrugated steel deck which is suspended on supporting beams as shown in Figure 3.2b. One hallway on the fourth floor has 17 permanent underfloor sensors, and an additional 10 sensors were mounted to the top surface of the floor for this study. A schematic of the structure of the floor and the layout of sensors is shown in Figure 3.2a. This hallway was subjected to impacts from an instrumented hammer at three locations in order to observe how waves propagate through the floor. Throughout this study, this structure is referred to as a suspended-slab floor system.

Some of the same experiments were also conducted in a second location to investigate the assumptions on a ground floor. While the upper floor was a suspended concrete slab, the ground floor is a ground-bearing concrete slab which may cause different wave propagation behavior. One impact location was considered for the ground floor, with sensors spaced out every 3 feet from the impact location (once again in a straight line similarly to the surface mounted sensors in Goodwin). The ground floor test was carried out in a different building, however the same instrumented hammer and surface mounted accelerometers was used. The results from the upper floor of Goodwin are the main focus of the investigations, while results from this ground floor are presented occasionally to contrast how the upper floor and ground floor behave with respect to the assumptions.

An instrumented hammer was used to simulate human-induced inputs. The maximum dynamic force imparted into the ground during walking is the heel-strike which often has low frequency content (up to 75 Hz) and has a maximum force of up to 1.25 times the person's body-weight [106]. However, various conditions such as gait speed and shoe material can have a significant impact on the forces during walking [91, 115]. It has been observed that the force of the heel-strike generates structural vibration in the 0-500 Hz band [37]. In the case of the occupant in Figure 3.1, the vast majority of the energy generated by their footsteps is concentrated below 200 Hz. A soft rubber tip was used on the hammer in this study to simulate low frequency human-induced inputs, which provided excitation up to 400 Hz. Additionally, the maximum force from the hammer impacts was consistently around 10,000 N, which is about 10 times the bodyweight of a 100 kg (220 lb) person.



3.3.2 Series 2: Detailed investigations of vibration in a beam-column structure

The second series of experiments involved a separate structure. Controlled experiments were carried out on a beam and column combined structure made out of reinforced concrete. Both




Figure 3.3: Image (a) and diagram (b) of the beam-column structure. The 14 measurement locations are shown in the diagram. For each measurement location, tri-axial measurements were taken on three faces: for example at location B2 there were measurements at the middle line on the top, front, and bottom faces simultaneously.

the beam and column have a square cross section, with a width and height of 20 inches. This structure was selected because it allows access to all faces of the beam and column, is made of reinforced concrete similar to the floor in many civil structures, and affords the opportunity to examine how waves interact with a support column. The beam spans at least 24 feet on either side of a junction, where the column goes 12 feet to the ground and also continues upwards to the next floor. An image and diagram of the measurement locations on the beam are shown in Figure 3.3. Two experiments were run on this structure. In both experiments an instrumented hammer was used to impact the top of the beam at the location shown in Figure 3.3b, 6 inches from location B1. The instrumented hammer used in this study was smaller than the hammer used in the first series of tests, and had a harder tip. It input energy up to 1000 Hz, in contrast to the 400 Hz in the building experiments.

In the first experiment, a roving sensor setup was used so that tri-axial measurements of the top surface, front surface, and bottom surface could be measured simultaneously as shown in Figure 3.3b. This was designed to study how the structure was moving in order to understand the type of waves propagating, e.g. if the top and bottom surfaces moved in phase as would be expected of transverse waves. All 14 locations in Figure 3.3b were used to track the movement of the beam's cross-section and see how waves transfer from the beam into the column.

A second experiment was carried out with sensors placed to track the dispersion of waves in both the beam and the column. For this test, it was necessary to have multiple sensors measuring in the same direction with some distance between them. This allows the measurement of the time it takes waves of different frequencies to travel between the sensors. This allows the time-of-arrival method to be used to calculate wave speed. For example, to measure

the dispersion of transverse waves within the beam, spectrograms of vertical acceleration at locations B1 and B4 could be used.

3.4 Assumption 1: Steady-state standing wave behavior

When waves propagate fast and long enough to reach boundaries, reflect, and reach steady-state they are called standing waves. In floor systems, however, high attenuation may prevent steady-state response from being reached. The presence of standing waves enables some simplifications for analyzing human induced vibrations. For example, a publication (SCI-P354) from the Steel Construction Institute in the United Kingdom recommends that steady-state floor vibration magnitude should be estimated during floor design via modal superposition for low frequency floors with a fundamental frequency below 10 Hz [107]. This steady-state response is intended for continuous walking signals with multiple footsteps spaced by a low walking frequency, and SCI-P354 recommends an additional transient analysis for single footfalls for both low and high frequency floors. If floor vibration response to single footfalls is similarly explained by modal superposition, algorithms such as model-based localization can be implemented through inverse dynamic analysis. This is seen in localization algorithms which use FE models, for example [34]. However, as shown in this investigation, the steady-state assumption and corresponding simplifications are not always valid in structurally complex floors such as those in Goodwin Hall.

3.4.1 Observation of standing wave quality through modal analysis

Under the steady-state, standing wave assumption, there are three modal parameters that affect the frequency response function (FRF), which relates inputs (footstep forces) to steady-state response: natural frequencies, damping ratios, and mode shapes. For a linear time invariant (LTI) system under steady-state vibration, these modal parameters should be consistent and independent of input location. In addition, mode shapes (with a few exceptions) should also demonstrate standing wave behavior, that is, measurement points should vibrate in-phase (0°) or completely out-of-phase (180°). In order to test the assumption that standing waves are present, these modal parameters were extracted through both experimental modal analysis (EMA) and operational modal analysis (OMA) [62].

For the EMA, average FRFs were estimated from five hammer impacts at impact location 1 as discussed in Section 3.3. These FRFs are shown in Figure 3.4a where the darker colors are from locations closer to the impact. The effect of attenuation can be seen as a progressive decrease in FRF magnitude as distance from the impact increases (colors get lighter). Despite this attenuation, at frequencies lower than approximately 30 Hz, there are consistent peaks present in a majority of the FRFs, a clear characteristic of reflections and




Figure 3.4: Accelerance FRFs for the 17 underfloor hallway sensors (a) and coherence of selected sensors (b). Darker colors are for sensors closer to the impact location (location 1) shown in Figure 3.2

the presence of global mode shapes. These are highlighted by the dotted lines. At frequencies above approximately 30 Hz, some peaks are present, but fewer are consistent across FRFs. Nevertheless, attenuation is still present to a similar degree as lower frequencies.

The signal coherence, shown in Figure 3.4b, generally declines in quality further from the impact, likely due to a decrease in signal-to-noise ratio. This degradation is less pronounced at high frequencies, indicating a more consistent input to output relationship. However, it is not possible to draw conclusion about steady-state behavior from coherence alone.

For the OMA, it was important to find a time period where the hallway under study experienced as much ambient excitation as possible. During an unplanned fire alarm, everyone in the building was required to evacuate. As the occupants of the building walked towards the exits, there was a period of approximately 10 minutes where the hallway experienced continuous distributed input from walking occupants. Covariance-Driven Stochastic Subspace Identification [84] and automated modal selection based on clustering [94] have been successfully applied in Goodwin Hall [98], and were used here on the set of 17 underfloor

hallway sensors to estimate the modal parameters in this study.

Modal parameters were extracted via EMA and OMA in the range from 6-13 Hz, where global FRF peaks were observed. For EMA, the modes were extracted from the hammer-impact FRFs via a least-squares rational approximation using the vector fitting algorithm [32, 47, 48]. For OMA, the same automated process used in Goodwin Hall by Sarlo et al. [98] was used to estimate modal parameters. The parameters estimated for the same four modes from each method are compared in Table 3.1. The estimated frequencies of each mode are similar, with differences under 10% between the methods. The frequencies estimated by OMA are all lower than those of the corresponding modes from EMA. The estimated damping ratio is always higher in the OMA, although all damping estimates from both methods are between 1% and 8%.

Table 3.1: Frequencies (f_i), damping ratios (ζ_i), and Mode Complexity Factor (MCF) comparison between EMA and OMA. MACs shown are between the EMA and OMA estimated mode shape for each mode order.

Mode	EMA			OMA			
	f_i	ζ_i	MCF	f_i	ζ_i	MCF	MAC
1	6.39	0.044	3.8	5.94	0.058	3.8	0.93
2	8.88	0.013	14.8	8.11	0.030	31.8	0.53
3	10.67	0.037	47.8	10.13	0.044	29.6	0.59
4	12.64	0.045	64.2	11.88	0.076	22.4	0.44

A visual comparison of the mode shapes and the phase of each location are shown in Figure 3.5. The mode shapes follow the pattern one would expect, with the first mode showing completely in-phase motion and higher orders exhibiting an increasing number of nodal lines (lines with zero amplitude). Due to the layout of the sensors mounted in this hallway, only the part of the mode shapes which develop along the length of the hallway are seen. Since the floor appears to be acting as a single structure and vibrating together, full plate-like modes for the floor would be expected which would exhibit deflection shapes in the direction perpendicular to the hallway as well, however this part of the mode shapes is not visible. A closer look at the phases of each measurement location in Figure 3.5 shows that some mode shapes are complex (e.g., EMA modes three and four, OMA modes two, three, and four). In other words, they do not exhibit the pure in-phase or out-of-phase behavior expected from standing waves which would be observed as points limited only to opposite sides of the phase circle. However, some of the mode estimates have a spread of phases suggesting that they are complex. This is reinforced by a high Mode Complexity Factor [54], as shown in Table 3.1. Nevertheless, there is good visual agreement between the EMA and OMA estimated mode shapes. However, the modal assurance criterion [83] (MAC, Table 3.1) between the mode shape estimates shows that there are some differences between the two methods. The




Figure 3.5: Estimated mode shapes for the four modes from OMA (a-d) and EMA (e-h). The phase of each measurement point is plotted on a circle around each mode shape, showing the complexity of the mode.

MAC for the first mode is high, while the MAC values get progressively lower for higher modes.

The results show some characteristics of successful modal analysis (e.g., global FRF peaks, intuitive mode shapes). However, the discrepancies in modal parameters obtained from both EMA and OMA tests as well as the complex nature of several mode shapes indicate that clear modal behavior cannot be guaranteed for such structures. Although noise could be partially to blame, it not a sufficient explanation as the coherence for a majority of sensors was strong. The following sections will explore additional possible sources for these discrepancies.

3.4.2 Sources of variability in FRFs and modal parameters

Various operational and environmental conditions may have an effect on modal parameters and could account for some of the differences between the EMA and OMA results. One such example which will be investigated is the excitation level of the hallway. A higher vibration amplitude observed during the OMA measurements, resulting from high occupant activity, is likely to contribute to increased stick-slip friction between joints as well as floor slab and beams, thus increasing the observed damping ratio [109]. In order to see the effect of vibration amplitude on modal parameters, a period of two hours surrounding the fire-alarm evacuation was considered. The same OMA procedure was carried out on this time period with 10 minute windows and an overlap of 90%. Figures 3.6a and 3.6b shows how




Figure 3.6: Variation of estimated natural frequencies (a) and damping ratios (b) for the four modes with average hallway RMS acceleration energy.

the natural frequencies and damping ratios for the first four modes vary with the root mean square (RMS) acceleration along the hallway. This time period had a mix of normal activity before the evacuation, high activity during the evacuation, and almost no occupant activity while the building was evacuated. Figure 3.6 shows that the first and third modes were not able to be estimated during some windows, while the second and fourth were estimated for almost all windows.

Table 3.2: Mean and standard deviation ($\bar{x} \pm \sigma_x$) for the estimated natural frequencies (f_i) and damping ratios (ζ_i) of the four modes along with the correlation coefficients ($r_{E,x}$) for each estimated modal parameter with average hallway acceleration RMS (RMS). Correlations with an asterisk (r^*) are statistically significant with a value $p \leq 0.05$.


Mode	OMA		Correlation	
	$\bar{f}_i \pm \sigma_{f_i}$	$\bar{\zeta}_i \pm \sigma_{\zeta_i}$	r_{RMS,f_i}	r_{RMS,ζ_i}
1	5.88 ± 0.14	0.068 ± 0.022	0.60^*	-0.01
2	8.05 ± 0.05	0.024 ± 0.008	0.36^*	0.46^*
3	10.11 ± 0.18	0.051 ± 0.023	0.26^*	0.08
4	11.71 ± 0.12	0.045 ± 0.020	0.37^*	0.71^*

All four modes showed a range of estimated frequencies and damping ratios over the time period as seen in Table 3.2. All modes showed a positive correlation between acceleration RMS and frequency, meaning higher hallway accelerations led to higher natural frequencies. The second and fourth modes also showed a positive correlation between RMS and damping ratio, meaning that higher hallway accelerations led to higher damping observations. It is


noted that in the case of the fourth mode, there is an order of magnitude change in damping values observed (1 to 10%). Other factors such as temperature and mass loading may also affect the modal parameters of the floor. Since the EMA and OMA were performed some time apart, these factors are also likely contributing to the differences between the methods.

In addition to operational conditions, the input location may have an effect on the measured FRF if steady-state is not achieved. For the modal analysis of an LTI system reciprocity should hold, i.e., the transfer function between two locations is the same regardless of which is the input and which is the output. In addition, the mode shapes should not change with input location. The attenuation seen in the FRFs in Figure 3.4 and the complexity of the extracted mode shapes suggests that waves are not reaching steady-state at all frequencies. In order to test for reciprocity and the effects of impact location on mode shapes estimates, further impacts were carried out at locations 5 and 25 feet from the original impact location. These additional impacts at locations 2 and 3 (as shown in Figure 3.2) were also directly next to a surface mounted sensor, allowing a reciprocity comparison between two locations. Figures 3.7a and 3.7b confirm the existence of reciprocity at both 5 and 25 feet from the original impact location.


However, Figure 3.7c shows that the first mode shape, extracted from the hammer-impact FRFs, changes with the impact location. In fact, the estimated mode shape changes such that its maximum displacement is always at the same location as the input. This presents a significant problem for model updating. If measured mode shapes are used to update a FE model, the updated model would be dependent on what input location(s) are chosen for the modal analysis. This dependence on input location also may explain some of the discrepancies seen in the MAC values between EMA and OMA as well as complex mode shape behavior. The following section presents a theoretical explanation for this phenomenon.



(a) Reciprocity between two locations 5 feet apart



(b) Reciprocity between two locations 25 feet apart



(c) First mode for each location in (a-b)

Figure 3.7: Demonstration of reciprocity between two sensors 5 feet apart (a) and 25 feet apart (b). First extracted mode shapes for the surface mounted sensors with each of the three impact locations (c). For each mode shape, a triangle shows where the impact location was for the FRFs from which the mode shape was extracted.

3.4.3 Effect of energy loss on FRFs

It is hypothesized that energy loss at the floor supports is a contributing reason that clear modal estimates are hard to obtain in floor structures. In order to investigate the attenuation seen in the FRFs in Figure 3.4 and input location dependency seen in Figure 3.7c, a FE model of an Euler-Bernoulli beam was created as shown in Figure 3.8a. This beam is connected by a series of springs and viscous dampers at each of its nodes to a grounded environment. The dampers were adjusted uniformly to simulate different amounts of energy loss and observe the interaction of attenuation, reflections, and FRFs. At each damping level, the beam was subjected to an excitation from the same half-sine pulse. As shown in the responses in Figures 3.8b-3.8d, fewer reflections were observed as the damping coefficient was increased and more energy was lost to the supports. This is reflected clearly in the time duration of the response (left column). Since the energy loss caused less reflections, the FRFs of the system

(middle column) were affected. These FRFs are calculated directly from the system matrices and are therefore not impacted by noise. Nevertheless, with high levels of reflections, there were clear peaks in the FRFs corresponding to each mode, while with no reflections the behavior of the FRFs was dominated by attenuation similar to that seen in Goodwin Hall above 30 Hz in Figure 3.4.

In the right column, it is also observed that as more energy is lost, the peak of the first extracted mode shape shifts towards the impact location. For a LTI system operating at steady state it would be expected that the mode shapes are invariant to input location. However, extracting the first mode shape from the FRFs using the same process as before, the mode shape does not remain invariant. For the case with the lowest damping, the first mode shape is almost identical for each input location. However, as the damping increasing and more energy is lost in this model, the peak of the mode shape shifts towards the input location. This mirrors the effect of input location on the first mode shape in Figure 3.7c. This model shows how energy loss to the environment causes problems if traditional steady state methods such as modal analysis are used. The dependency on input location means that calibrating a model or using mode shapes to expand spatial resolution are more difficult.




Figure 3.8: Diagram of clamped-clamped Euler-Bernoulli rectangular beam (a) with 100 elements connected to the environment at each node by a spring and damper to simulate energy loss. The following parameters were used in the beam model: Total length $L = 30 \text{ m}$, width $w = 1 \text{ m}$, height $h = 0.5 \text{ m}$, Young's modulus $E = 25 \text{ GPa}$, density $\rho = 2700 \text{ kg/m}^3$, spring stiffness $k = 3 \times 10^5 \text{ N/m}$. Results for three values of damping: $c = 3 \times 10^2 \text{ kg} \cdot \text{s}$ (b), $c = 3 \times 10^3 \text{ kg} \cdot \text{s}$ (c), $c = 1.5 \times 10^4 \text{ kg} \cdot \text{s}$ (d). For each value of damping the time response to a half-sine input, FRFs, and first extracted mode shape for selected nodes are shown. For mode shapes, a triangle denotes the corresponding input location.

3.4.4 Summary of assumption

As displayed by the complexity of the extracted mode shapes, as well as the dependence of mode shape on input location, standing waves are not always generated in floor structures by the level of forces common in occupant walking. The hypothesis for this behavior, supported by a theoretical model, is significant loss of energy at the floor supports. Despite this finding, there is a cutoff frequency (for this floor around 30 Hz) below which there are more apparent reflections. At these low frequencies, FE models may still be useful, especially if they explicitly take into account attenuation over distance and energy loss. Furthermore, some important expectations of FRFs such as reciprocity are not always affected by energy loss or lack of standing wave behavior, strengthening the case for data-driven FRF or transfer function-based approaches to VBOI. Nevertheless, confounding variables such as vibration amplitudes were also shown to alter at least some modal parameters. Changing operational and environmental conditions of a building would complicate the use of both computational and data-driven models as there is not one fixed true model of the system. In order to fully account for effects caused by attenuation, input location, activity level, temperature, and many more factors would require further assumptions or an unreasonable amount of testing. When using calibrated models, some care should be taken to ensure that results are independent of excitation location, excitation amplitude and environmental conditions.

3.5 Assumption 2: Infinite plate with uniform attenuation


If steady-state behavior cannot be assumed and transient behavior dominates, it may be preferable to implement methods based on wave propagation. A common assumption made in this case is that waves in a building propagate as they would through an infinite plate. Under this assumption one may ignore the influence of boundaries, which produce effects such as wave reflection, mode conversion and localized energy dissipation. Therefore, wave attenuation can be idealized as a constant rate of decay. This is common with energy-based localization algorithms which may, for example, model energy decay as an exponential function of distance [1, 2]. It is obvious that real floor systems are far from infinite plates, with a variety of components that may constitute boundaries: walls, furniture, supporting beams and columns, and material interfaces. Thus, it is important to consider the degree to which these boundaries influence wave propagation in real systems, particularly with respect to attenuation behavior. This section begins by considering attenuation behavior as a whole before exploring wave reflection and conversion behavior at a support column. This helps to illustrate when boundaries effects may be the most significant and impact the infinite plate assumption.

3.5.1 Observation of attenuation and reflections in Goodwin Hall


In order to experimentally observe how energy changes with distance in a real floor system, five impacts were carried out at location 1 in Figure 3.2. Additionally, three impacts were carried out at one location on a ground floor of another building to compare attenuation on an upper floor and a ground floor. If the infinite plate assumption is true, there should be a uniform attenuation rate with distance. For each impact, the RMS of the acceleration signal received by an array of sensors was calculated, and normalized by the RMS of the signal from the sensor placed directly next to the impact. Figure 3.9a shows how RMS acceleration attenuates as a function of distance from the impact. For the impacts on the upper floor, the RMS from surface and underfloor accelerometers matches well at similar distances from the impact, suggesting that the concrete slab and supporting beams move together. However, the effect of boundaries is apparent as the RMS does not decay at a consistent rate with distance. For example, RMS increases moving from a distance of -11 to -14 meters from the impact source. In addition, the decay rate is not consistent in both directions. The inset plot of Figure 3.9a shows the response on the surface of the floor at two sensors which are both the same distance (1.5 meters) from the impact. Although they are the same distance from the source and would therefore have undergone equal amounts of attenuation, one shows about twice the signal RMS and a significantly different time-history acceleration response. This shows that the rate of attenuation observed at points the same distance from the source in Goodwin Hall is sensitive to the location of the sensor, likely due to different distances to boundaries causing reflections.

Unlike the upper floor, the ground floor demonstrates quite a consistent rate of decay. This is also plotted in Figure 3.9a. However, this rate is much higher and consequently the RMS decays to the noise floor at 10 meters away from the impact, while for the upper floors there was an observable signal as far as 20 meters from the impact. Considering the distributed support provided by the ground-bearing slab construction versus the local supports in the suspended slab, higher energy dissipation is the most likely explanation for the higher attenuation rate. Because waves attenuate faster and therefore do not travel as far before decaying to the noise floor, a significantly denser sensor network would be necessary on the ground floor to ensure equivalent signal to noise levels as the upper floors.

Let us now consider the upper floor again, but in a band-limited sense. Similar to the results in Section 3.4, it is possible to observe different behavior at low and high frequencies. Figure 3.9b shows the RMS attenuation of the surface-mounted sensors, but in this case, band-limited to 0-30 Hz and 30-100 Hz. At high frequencies, the attenuation rate is more uniform. This is likely due to a higher influence of reflections at low frequencies as was seen in the FRFs in Figure 3.4. From Section 3.4 and the results of uneven attenuation, it is clear that reflections arising from boundaries play an important role in wave propagation in civil structures. The following sections will demonstrate how boundaries can lead to these reflections.



(a)



(b)

Figure 3.9: Visualization of acceleration attenuation with distance for surface sensors, underfloor sensors, and ground floor sensors (a). Acceleration attenuation with distance for surface floor sensors only, separated into 0-30 Hz and 30-100 Hz bands (b). RMS acceleration is normalized by the RMS acceleration at 0 distance from the impact, RMS_0 .

3.5.2 Wave propagation through a column

Columns are one of the most common boundaries encountered in a floor system. It has been observed that occupant-induced vibrations not only transmit waves along floor slabs and floor beams, but also through columns to other floors [68]. In particular, waves traveling within a floor may undergo mode conversion where some percentage of the wave is converted into waves transmitting through a column to another floor while some percentage of the incoming wave is reflected back. The former mechanism can be seen in Figure 3.10 which shows spectrograms from vertical sensors on floor beams on the fourth floor and on the third floor of Goodwin Hall, but both attached near the same column. Figure 3.10b shows that very quickly after the impact on the fourth floor (time 0), a wave is observed by the column sensor on the fourth floor. Approximately 5-10 ms later, the wave is observed by the third floor sensor (Figure 3.10c). It can be seen that although the wave's magnitude is lower at the third floor, its time-frequency content is essentially unchanged. This indicates that some of the floor's vibration is partially "lost" through mode conversion into the building's columns and back out to different floors. The rest may be either reflected, transmitted, or dissipated. In addition, the consistency of the time-frequency behavior indicates that dispersion does not play a role in the propagation of waves vertically through these columns.




Figure 3.10: Diagram of sensors on the fourth floor and third floor near a column (a). Spectrograms of measurements at each sensor from an impact on the fourth floor.

In order to explore wave mode conversion, transmission, and reflection at the beam-column boundary, a roving sensor setup was employed on the beam-column structure which allowed

tri-axial acceleration measurements to be taken at the centerline of the beam or column on the top, front, and bottom faces simultaneously (see Figure 3.3). This allows the measurement of the motion of cross section at each location to be obtained. Figure 3.11 shows the acceleration energy and relative motion direction of each face of the cross section at each location, ignoring out-of-plane motion (z -direction). For example, the inset plots show that at location B2 all accelerations in the y -direction have a similar magnitude, and are all in-phase. At the same location, there is acceleration in the x -direction of opposite phase on the top and bottom of the beam, while there is very little horizontal (x) acceleration in the middle of the front face. This combination of in phase transverse motion and out of phase horizontal motion at each surface is typical of antisymmetric Lamb waves, especially antisymmetric wave mode (A_0) which exists at all frequencies (as opposed to the higher order, evanescent antisymmetric wave modes) [44]. For reference, schematics are provided in Figure 3.11 for what deflection shapes would be seen for wave modes. After the junction, there is still the same pattern indicative of antisymmetric Lamb waves along the beam, but with a smaller magnitude. On the other hand, in the column there appear to be antisymmetric Lamb waves and also symmetric Lamb waves. The symmetric Lamb waves are also likely the first mode, the S_0 mode, since they also exist at all frequencies [44]. The antisymmetric Lamb waves are still seen as in phase transverse motion, although this time the transverse direction is the x -direction. The symmetric Lamb waves are seen as an offset in the y -direction, which would not be seen in the antisymmetric Lamb waves. Finally, some percentage of the waves appear to be reflected as well. Destructive interference from these waves are seen in the motion at location B5 which has significantly lower acceleration than the previous locations B1-B4.




Figure 3.11: Visualization of cross section motion in the x and y directions (z motion is ignored as it was negligible). The energy and relative motion direction for each sensor on the beam and column structure are shown. Energy is calculated by integrating FRFs to remove influence of differences in hammer impacts, while relative phase is determined by cosine similarity of time series since all measurements at each location were taken during the same impact. Inset are time series measurements for location B2, to illustrate how the acceleration changes between the faces as well as schematics of what various wave modes would look like.

3.5.3 Parameters affecting mode conversion and reflections

Waves propagating which encounter an interface such as a different material or boundary can undergo a process called “mode conversion.” For example, shear waves in a semi-infinite media which encounter the interface of a free surface at an angle can undergo mode conversion which creates both shear and pressure waves [44]. In the case of the beam-column structure we are considering, the junction of the supporting column and beam presents a complex interface. Since we have seen one wave mode in the beam giving rise to two wave modes in the column, we use the term mode conversion for the generation of waves in the column. The interface in our case is more complicated than a single free surface, thus there is also the avenue for waves to continue in the beam beyond the column and also to reflect back along the beam. In cases of waves propagating into interfaces, the literature defines amplitude ratios to quantify the degree to which different wave modes are converted or transmitted at or beyond an interface. The amplitude of each resulting wave is divided by the incoming wave amplitude giving the amplitude ratio for that wave mode. A higher amplitude ratio means

that a higher percentage of the incoming wave has been transformed into that wave mode. In our case we will have three ratios which we can measure: the ratio of waves converted into A_0 waves in the column, converted into S_0 waves in the column, and transmitted beyond the column as A_0 waves in the beam. The amplitude ratio for reflections cannot be directly measured, but can be inferred.

It is also important to understand what influences the ratio of waves which are present after the interface in both the beam and column, and what ratio are reflected. To this end, a 2D Abaqus finite element model was created which matched the geometry and material of the beam and column structure, and was subjected to an impulse force input to simulate the hammer hit. Figure 3.12 shows the amplitude ratios for the outgoing A_0 waves in the beam beyond the junction, and the A_0 and S_0 waves travelling down the column as well as the total acceleration (A_0 and S_0 waves together) in the column for the experimental data and for the model. These ratios were calculated with respect to the relative beam span to thickness ratio (slenderness ratio, Figure 3.12b), as well as the relative column to beam stiffness (stiffness ratio, Figure 3.12c). As the column material was made stiffer, there were less of all wave modes present, suggesting that a stiff joint leads to more reflections and less transmission to other floors through columns. As the slenderness ratio changed, different waves exhibited different changes in amplitude ratios. With a very thin beam, more acceleration is transmitted across the joint, while with a thick beam there was more acceleration converted into S_0 waves in the column. The A_0 waves in the column were the least sensitive to changes in the beam thickness. Also shown in Figure 3.12 are the experimentally observed amplitude ratios which match well with the finite element model. The ratio of waves going beyond the column and going into the column matched very well, while there is a slight difference in the amount of A_0 and S_0 waves in the column between the experiment and model.

Buildings such as Goodwin Hall have thin floors relative to the span between columns (low slenderness ratio) and stiff steel columns (high stiffness ratio). If Goodwin were to follow the trend seen in this structure, there would be high levels of reflection, high transmission across boundaries (staying within the same floor) with a low amount of conversion into columns (going between floors). These expectations are qualitatively supported by the observations in Goodwin Hall. It was seen that there were high levels of reflections in Figure 3.9, and also that there is limited transmission between floors which did not exhibit much dispersion in Figure 3.10. The lack of dispersion supports the idea that S_0 waves (or pressure type waves) play a role in mode conversion transferring waves between floors.




Figure 3.12: Diagram of the Abaqus model (a). Amplitude ratios for changing slenderness ratio (b), and amplitude ratios for changing stiffness ratio (c). Amplitude ratios for A_0 waves staying in the beam, both A_0 and S_0 waves in the column, as well as the total magnitude in the column are shown.

3.5.4 Summary of assumption

Overall, the data shows that for some real floor systems it may be unreasonable to assume an infinite plate and hence a uniform wave attenuation rate, which would impact the estimation quality of energy-based localization methods. In some cases there may be frequency ranges where this assumption is more viable than others. In the case of Goodwin Hall, above the cutoff frequency of around 30 Hz the attenuation rate is more uniform and the effect of boundaries are less apparent. It has also been shown that the geometry and basic material properties (e.g., stiffness) of floor components will have an effect on the degree to which these boundaries produce reflection, transmission or mode conversion. Additionally, different floor types have different attenuation behaviors. For example, a ground floor with its distributed ground-bearing support show a much more even attenuation, in line with the infinite plate assumption. If an infinite plate or other idealized structure is to be assumed, it is important to understand how the building geometry impacts wave propagation, and if there are frequency ranges where the assumption is more viable.

3.6 Assumption 3: Neglected or simplified dispersion relations

Dispersion presents a significant challenge for VBOI algorithms which depend on time difference of arrival (TDoA) or other time-domain calculations. Because dispersion manifests as different propagation speeds for different frequencies in a waveform, a single arrival time

is difficult to define. In fact, the waveform morphs as it travels and it is not consistent among measurement locations, thus any type of time domain analysis is challenging. Certain methodologies make assumptions to either simplify or avoid the effects of dispersion. As an example of the former, some works have made the assumption that a floor acts as a thin-plate and therefore has a relatively simple dispersion relation [117]. As an example of the latter, some works have considered only a narrow frequency band to effectively isolate a single wave propagation velocity and simplify the time of arrival calculation [81].

Because various types of simplifications are possible, this assumption should be considered in a broader sense than those in the previous sections. In some cases, it may be similar to Assumption 2 in assuming uniform floor properties. However, it is considered separately since its motivating phenomenon is different. When wave propagation speed is of primary concern, as is the case with TDoA localization, the presence of reflections and attenuation is not as important as dispersion. Therefore, this assumption is focused only on simplifications of the dispersion relation which can be made separately from the infinite plate assumption in the previous section.

3.6.1 Observations of dispersion in buildings

In order to observe dispersion in buildings, the responses measured by sensors at different distances from a hammer impact were investigated. Figure 3.13 shows spectrograms of the acceleration measured at four locations on the fourth floor of Goodwin Hall. Two locations were on the surface of the floor slab (b and c) and two were on I-beam supports approximately under each of the surface locations (d and e, respectively). The surface and support sensors were approximately the same horizontal distance apart (40 ft and 42 ft, respectively) and this difference in position (2 ft) is not expected to influence the effects of dispersion significantly. The impact occurred very close to sensors b and d. Comparing the spectrograms between surface and support sensors at the same location (b with d and c with e) shows that the acceleration magnitude is similar on the surface and on the support beams, which agrees with previous results that the concrete and support beams move together. Comparing the measurements close and far from the impact (b with c and d with e) clearly shows dispersive effects over the travelled distance, where the low frequency content arrives later than high frequency content. There also appears to be a levelling off of the propagation speed at high frequencies, whereby the time of arrival for frequencies above 200 Hz largely does not change. This same dispersion relation, with slow wavespeed at low frequencies increasing to a maximum speed at high frequencies, is exhibited by A_0 Lamb waves [44].

The type of waves which propagate on a ground-bearing slab floor could be different than those on an suspended slab floor like those shown in Figure 3.13, and therefore could experience dispersion differently. Therefore, a test was also conducted with sensors on the surface of a ground-bearing slab floor in another building. Figure 3.14 shows that the waves propagating in this ground floor show high attenuation over a distance of 15 feet. As seen




Figure 3.13: Diagram of sensors on the surface of floor, and under the I-beam (a). Spectrograms of measurements from the surface (b,c) and under the I-beam (d,e) for an impact from an instrumented hammer. One sensor for the surface and I-beam are close to the impact (b,d) while the other is 40-42 feet away (c,e) so the effects of dispersion can be seen.

previously in Assumption 2, the attenuation on a ground-bearing floor is higher than on a suspended-slab floor. This is especially pronounced at low frequencies characteristic of human induced vibration, with very high attenuation seen under 100 Hz on the ground floor. Nevertheless, it appears that dispersion occurs on both floor types, although the low frequency waves which were very sensitive to dispersion in the suspended slab floor attenuate rapidly on the ground-bearing slab. In the frequency range of footstep induced vibrations (0-100 Hz) dispersion appears to dominate in suspended slabs while attenuation dominates in ground-bearing slabs.

3.6.2 Quantification of dispersion relations

The deformation of the beam-column structure in Figure 3.11 shows that the wave propagation appears to be dominated by Lamb waves. Antisymmetric waves appear to be present in the beam, while both antisymmetric and symmetric waves appear to be present in the column. This section will seek to quantify the dispersion curves of these waves and compare them to the expected dispersion curves for Lamb waves. In order to observe dispersion, multiple sensors with some distance between them must observe the response to the same input. For example, to measure the dispersion curve of the A_0 waves in the beam, sensors measuring in the y -direction at the middle of the beam's front face at location B1 and B4 measured the response to the same impact. Spectrograms were taken of these results, sim-




Figure 3.14: Spectrograms from an impact on a ground floor, measured by sensors placed on the surface of the concrete floor 15 feet away from each other.

ilar to those in Figure 3.13, and the time of arrival of each frequency was found. The time of arrival was defined as the earliest time that the magnitude at a certain frequency made it within a tolerance (for example -6 dB) of the maximum magnitude for that frequency. The group velocity was then calculated using the differences in arrival times and the known distance between the sensors [117]. A similar process was used to measure the dispersion curves for both the A_0 and S_0 waves in the column. On the column, sensors were placed at the middle of the front face at locations C1 and C25 in both the x -direction and y -direction. In the middle of this face, all motion in the y -direction should be from the S_0 waves, while all motion in the x -direction should be from the A_0 waves. Therefore, the dispersion curves for each wave mode in the column can be calculated as well.

Figure 3.15a shows the group velocity of A_0 waves in the beam, both A_0 and S_0 in the column, and A_0 waves within the fourth floor of Goodwin Hall. In all cases, estimating group velocity at low frequencies was difficult; possibly due to the long wavelength compared to the distance between sensors. The group velocity of the symmetric Lamb waves in the column was also difficult to estimate because of how quickly the wave was propagating, and the relatively short distance over which it was possible to measure.

Despite the noise in the experimental measurements, the form of these dispersion curves matches well with the theory for A_0 and S_0 waves. This can be observed by comparing them to the analytical dispersion curves for A_0 and S_0 waves in a plate shown in Figure 3.15b. Because the dispersion curve is sensitive to plate thickness H and shear speed c_s , the wave frequency was normalized by these parameters as is noted in the x -axis. Since these parameters vary among our test case, we restrict ourselves to comparing only the shapes of these curves. In this manner, we can observe that the A_0 waves in the beam, column, and in Goodwin Hall have a low wavespeed at low frequency, which plateaus at high frequencies. This pattern is expected from A_0 waves as shown in Figure 3.15b for a theoretical plate. Furthermore, the S_0 waves in the column start at a high wavespeed at low frequencies and decrease with increasing frequency. This also matches with the analytical model for S_0 waves, although this measurement had significant noise.

Figure 3.15b also shows the dispersion curves for some possible dispersion simplifications: transverse waves following thin-plate theory, and pressure waves. The thin-plate theory simplification for transverse waves shows a similar trend to A_0 waves at very low frequencies, but has large discrepancies at higher frequencies. On the other hand, the constant group velocity of pressure waves appears to be a good approximation of the dispersion curve for S_0 waves at very low frequencies. However, as we have seen, S_0 waves are not a common mode of vibration induced on a horizontal slab by vertical impacts (i.e. footsteps). These modes only tend to occur in columns after mode conversion. Another possible simplification would be to only consider A_0 waves above a certain frequency, which could be chosen as the frequency where the group velocity plateaus. However, the frequencies of the plateaus in Figure 3.15a show that this simplification may only be viable at a frequency too high for use in occupant inference due to the low frequency nature of waves induced during walking.




Figure 3.15: Dispersion curves measured in the beam and column structure, as well as the fourth floor of Goodwin (a). The group velocities (c_g) for the first three symmetric and antisymmetric Lamb modes of a plate with $\nu = 0.2$, normalized by the material's shear speed c_s and the plate thickness H (b).

3.6.3 Parameters that affect dispersion relations

Figure 3.15b shows the analytical dispersion curves for the A_0 and S_0 vibration modes of an infinite plate which were calculated following Viktorov [113]. While the dispersion curves calculated in Figure 3.15a follow expected trends, only the very low frequency portion of the analytical curves in Figure 3.15b is seen. The general trend is always the same for the same wave mode, but the shear wave speed (c_s) of the material, the Poisson's ratio of the material (ν), and the thickness of the plate (H) will influence the final mapping between frequency

and group velocity. The effect of increasing the shear wave speed can be seen, as the A_0 waves in the column reach higher velocity than those in the beam. This is partly because the shear speed increases with compressive stress [52] developed in the column, leading to the waves plateauing at a higher velocity. If the beam and column are made of the same material and have the same cross-section, it would be expected for their dispersion curves to be the same. The beam and column do have the same cross-section, however it is possible there is more reinforcement in the concrete of the column, which combined with the effect of the compressive stress on the shear wave speed, give the beam and column different dispersion curves. Similarly, changes in the Poisson's ratio or thickness of the plate would lead to changes in the dispersion curve.

It is also important to note that the curves plotted in Figure 3.15b are for an isotropic plates of uniform thickness. In a real building, there may be interfaces where material properties or geometry (such as slab thickness) change, and the dispersion curves would be expected to change accordingly.

3.6.4 Summary of assumption

Dispersion is a complex phenomenon which cannot always be easily simplified. Simplifications such as thin-plate theory may be viable at very low frequencies, while the assumption of little dispersion effects may be viable at higher frequencies. There are many factors which affect the specific dispersion relation, and will therefore alter the viability of certain simplifications on a case-by-case situation. The wave mode determines what the overall shape of that curve is, while material properties and geometry parameters affect the final dispersion curve. Additionally, considering floors which have non-uniform material properties or changing geometry would further complicate the observed dispersion effects on wave propagation. In this section we have demonstrated a simple experimental method for obtaining approximate dispersion characteristics. It is recommended that such an analysis be performed in order to determine what level of dispersion simplification is necessary for implementing effective VBOI in each specific structure.

3.7 Conclusion

This chapter provides a unique categorization system for VBOI algorithms, particularly those focused on footstep localization, organized by the principal simplifying assumptions made about floor dynamic behavior. So far, such assumptions have been applied in an ad hoc manner, without significant considerations for the type of floor structure or the physical mechanisms driving wave propagation. Therefore, the goal of this study was to provide an experimental and theoretical basis for assessing the limitations of these key assumptions, in order to develop algorithms which generalize to a wide range of structures.

Three principal assumptions were identified: (1) Steady-state behavior, (2) infinite plate behavior, and (3) neglected/simplified dispersion behavior. This categorization enabled a targeted investigation of each of these assumptions through a series of experiments on typical floor systems and a concrete beam-column structure. Based on the experimental observations, theoretical models were presented to explain the primary phenomena impacting the validity of each assumption. These models can serve as a framework for justifying the use of certain algorithms over others depending on the application. Table 3.3 provides a summary of the results from each investigation, including the main sources of error for each assumption along with key considerations and recommendations for future works making these or similar assumptions about vibrations in civil structures. This table is meant to be a quick reference of the main conclusions, while the previous sections for each assumption have more detailed summaries of each investigation.

It is important to note that the goal of this chapter was not to quantify the impact of the assumptions on VBOI algorithm performance, this will be left for future work. It is the authors' belief that a clear fundamental understanding of wave propagation in floors is required before pursuing algorithmic approaches. In addition, it is noted that the set of floors considered in this study (a suspended-slab upper floor and a ground-supported slab), while good test structures for the phenomena discussed, only represent a small subset of possible floor systems. Future work should expand to a wide variety of floor systems in order to better understand the generalizability of the concepts discussed here. This work provides a set of tests and validations that can be applied to a variety of floor systems and help quantify the effect of assumptions on various algorithms. This should ultimately allow the field of vibration-based occupant inference to improve performance and expand its possible range of applications.

Table 3.3: Summary of each assumption, providing relevant algorithms, sources of error when making the assumptions, considerations which affect the applicability of the assumption, advantages to the assumption, and recommendations from the authors for each assumption in light of the investigations in this work.

	Assumption 1: Steady-state behavior	Assumption 2: Infinite plate behavior	Assumption 3: Simplified dispersion behavior
Relevant Algorithms:			
Localization	Model-based [30, 34, 93]	Energy-based [2]	Time-based [7, 71, 90]
Other	Floor design [107]	Classification [9, 80] Energy-based gait analysis [39, 61]	Dispersion compensation [117] Temporal gait analysis [61]
Sources of Error	High attenuation Changing modal parameters	Reflections	Complex dispersion
Considerations	Amount of attenuation Excitation conditions (ex. RMS accel) Environmental conditions (ex. temperature)	Floor materials Floor thickness Size, density and geometry of boundaries (walls, supports, joints, etc.)	Wave mode Floor materials Floor thickness
Suitable Cases	Low frequencies	High frequencies Ground-bearing floors	Lamb wave theory
Recommendations	Simple techniques such as modal superposition are not capable of representing all floor systems. FE models or data-driven FRF methods may be useful provided they take into account attenuation and energy loss. Modal testing can help find frequency ranges where these techniques are most useful.	Reflections arising from boundaries such as walls and columns cannot always be ignored. Observing vibration decay from hammer impacts can show frequency ranges where there are fewer reflections, making this assumption more viable.	Lamb wave theory has been shown to model dispersion of transverse waves in floors well (especially the A_0 wave mode). Otherwise, spectrograms from hammer impacts may be used to justify simplifications over a chosen frequency range.

Chapter 4

Modelling Second Order Systems

This chapter reviews current methods for modelling second order systems, setting the scene and giving context for the next chapter, which will introduce a novel data-driven modelling method for creating predictive models of second order systems.

4.1 Introduction

Vibrating systems are governed by second order differential equations. Mechanical vibrations are of interest in diverse fields from the aerospace industry [27] to civil structures [79, 98]. Predictive models which have a fine spatial resolution or are continuous in space are often used to simulate responses to various inputs. For example, in the Impedance Matched Multi-Axis Testing (IMMAT) methodology, predictive simulations with a range of combination of input locations are used to optimize environmental testing procedures [27]. Environmental testing is commonly used to ensure that structures can handle hostile environments such as aerodynamic loading of a missile in flight [27]. In VBOI, predictive models have been applied to the problem of localization [30, 34]. Using simulations [34] or solving an inverse problem using FRFs [30] allows an estimation of the most likely place for the input. The localization accuracy is directly tied to the spatial resolution of the FRFs or model used, so having a continuous model is desired.

4.1.1 Overview of Characterizing Second Order Systems

There are many methods for characterizing second order systems. In general, the more information that is known or measured about the system, the better the system is understood. Figure 4.1 shows this principle graphically. With full knowledge of the system and access to the partial differential equation (PDE), analytical methods are capable of creating continuous predictive models of the system with high accuracy [53, 112], as shown in the top right quadrant of Figure 4.1.

However there is not always full knowledge of all geometry, material properties, and boundary conditions to fully define a system. By measuring the output or input and outputs of a system modal analysis can be used to estimate eigenvalues and eigenvectors of a system [4], the same system properties that define the system built by analytical methods. Models




Figure 4.1: A graphical comparison of selected methods for characterizing second order systems.

can be fit to the data measured by modal analysis to build data-driven predictive models via rational approximation algorithms [11, 48, 67]. The limitation of data-driven methods, however, is that the knowledge gained about the system is limited to the locations where data is measured. There are some ways to measure full field outputs, such as digital image correlation, or to expand from discrete measurements to full field outputs like dynamic vibration expansion [25, 78]. However, currently, there is no method to take discrete system measurements and generate a continuous predictive model of the system.

The next chapter will introduce a novel method, continuous residue interpolation (CRI), which is designed to do just this. By measuring FRFs between inputs and outputs of a system at a discrete set of locations, an optimization process is used to fit a continuous predictive model of the entire system. Figure 4.1 shows how CRI compares to other current methods, filling a gap of creating a continuous predictive model from discretely measured input and output data.

This chapter is dedicated to reviewing theory of second order systems and some methods for building models of second order systems. Section 4.2 is dedicated to exploring concepts relevant to continuous predictive models: partial differential equations, state space models, and transfer functions. Section 4.3 will explore some aspects of current modelling methods which are relevant to the novel method CRI which will be fully explored in Chapter 5.

4.2 Theory of Second Order Systems

This section provides a brief overview of the theory necessary for modelling second order systems. The focus is on PDEs, state space models, and especially the idea of transfer functions and FRFs.

4.2.1 Partial Differential Equations

Vibrating systems can be described with second order partial differential equations (PDEs). For example, the PDE of an Euler-Bernoulli beam without damping is:

$$m(x)\frac{\partial^2 w(x,t)}{\partial t^2} + \frac{\partial^2}{\partial x^2} \left(E(x)I(x)\frac{\partial^2 w(x,t)}{\partial x^2} \right) = F(x,t), \quad (4.1)$$

where $m(x)$ is the linear density, $E(x)$ is the material stiffness, $I(x)$ is the moment of inertia, $w(x,t)$ is the out-of-plane displacement of the beam, and $F(x,t)$ is the force input to the beam. Second order PDEs can be discretized into mass, stiffness, and damping matrices:

$$\mathbf{M}\ddot{w}(t) + \mathbf{D}\dot{w}(t) + \mathbf{K}w(t) = F(t). \quad (4.2)$$

The mass matrix, $\mathbf{M}$, and stiffness matrix, $\mathbf{K}$, can be built via spectral methods such as the Rayleigh Ritz method (RRM), or via the finite element (FE) method. However, it is often difficult to model the damping matrix, $\mathbf{D}$, from first principles. One common approach for including damping in the system is to apply proportional damping:

$$\mathbf{D} = \alpha\mathbf{M} + \beta\mathbf{K}, \quad (4.3)$$

where the damping matrix is proportional to the mass and stiffness matrices [55]. Another approach is to define a damping ratio, ζ_i , for each of the decoupled equations of motion:

$$\ddot{r}_i(t) + 2\zeta_i\omega_i\dot{r}_i(t) + \omega_i^2 r_i(t) = \tilde{f}_i(t). \quad (4.4)$$

Here, the matrix equations of motion from Equation 4.2 are decoupled into one equation of motion for each of the modal coordinates $r_i(t)$, where ω_i is the angular natural frequency of the mode, and $\tilde{f}_i(t)$ is the modal force [55].

4.2.2 State Space Model

It is often helpful to turn the second order equations of motion into ordinary first order differential equations to make simulations easier. From Equation 4.2, a state space model can be defined

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t); \quad \mathbf{y}(t) = \mathbf{C}\mathbf{x}(t), \quad (4.5)$$

where $\mathbf{x}(t) \in \mathbb{R}^\rho$ are the states of the system, $\mathbf{A} \in \mathbb{R}^{\rho \times \rho}$, $\mathbf{B} \in \mathbb{R}^{\rho \times M}$, $\mathbf{C} \in \mathbb{R}^{N \times \rho}$, $\mathbf{u}(t) \in \mathbb{R}^M$ is the input (forcing) term, and $\mathbf{y}(t) \in \mathbb{R}^N$ are the measured outputs of the system. The state space matrices can be defined as:

$$\mathbf{x} = \begin{bmatrix} w \\ \dot{w} \end{bmatrix}, \quad \mathbf{A} = \begin{bmatrix} 0 & \mathbf{I} \\ -\mathbf{M}^{-1}\mathbf{K} & -\mathbf{M}^{-1}\mathbf{D} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 0 \\ \mathbf{M}^{-1} \end{bmatrix}, \quad \mathbf{C} = [\mathbf{I} \quad 0]. \quad (4.6)$$

A diagonalized state space model is equivalent to the decoupled equations of motion in Equation 4.4. The eigenvalue decomposition of $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ can be used to define a new state space model

$$\dot{\mathbf{x}}(t) = \mathbf{\Lambda}\mathbf{x}(t) + \tilde{\mathbf{B}}\mathbf{u}(t); \quad \mathbf{y}(t) = \tilde{\mathbf{C}}\mathbf{x}(t), \quad (4.7)$$

where $\mathbf{\Lambda}$ is diagonal, $\tilde{\mathbf{C}} = \mathbf{C}\mathbf{V}$ has columns $\tilde{\mathbf{c}}_j$ and $\tilde{\mathbf{B}} = \mathbf{V}^{-1}\mathbf{B}$ has rows $\tilde{\mathbf{b}}_j^\top$. The diagonal entries of $\mathbf{\Lambda}$ are the poles of the system λ_i , and each mode has a complex-conjugate pair of poles $\lambda_j, \bar{\lambda}_j = -\omega_j \left(\zeta_j \pm i\sqrt{1 - \zeta_j^2} \right)$.

4.2.3 Transfer Function

While state space models are useful for time-domain simulations, transfer functions can be used equivalently for frequency domain simulations. The transfer function between one input location and output location, $h(s) \in \mathbb{C}$, relates Fourier transform of the measured output $y(t)$, $y(s) \in \mathbb{C}$, and the Fourier transform of the input $u(t)$, $u(s) \in \mathbb{C}$: $h(s) = y(s)/u(s)$. Transfer functions are also particularly useful due to their property of being state invariant. There can be different state space realizations for the same system (as seen with Equation 4.5 and Equation 4.7), however these different system representations will still have the same transfer function. Transfer functions can also be measured experimentally, called frequency response functions (FRFs) when measured along the imaginary axis at $s = i\omega$. The field of modal analysis is built upon using measured FRFs to update system models or to extract modal parameters.

The transfer function can be defined directly from PDEs, or from the state space representation of the system. First, we will consider the PDE for a uniform beam without damping:

$$m \frac{\partial^2 w(x_{\text{out}}, t)}{\partial t^2} + EI \frac{\partial^2}{\partial x^2} \left(\frac{\partial^2 w(x_{\text{out}}, t)}{\partial x^2} \right) = F(x_{\text{in}}, t), \quad (4.8)$$

where the linear density m , stiffness E , and inertia I are constant along the beams length. The beam is defined along $x \in \mathbb{R}$ on $[0, L]$, and the input $F(x_{\text{in}}, t)$ and output $w(x_{\text{out}}, t)$ are measured at one point each. If both ends of the beam are pinned, the boundary conditions are:

$$w(0, t) = w(L, t) = w''(0, t) = w''(L, t) = 0. \quad (4.9)$$




Figure 4.2: Diagram of beam with pinned boundary conditions at both ends.

Figure 4.2 shows a diagram of the pin-pin beam. For this beam, Fahy et al.¹ in their book include a derivation which shows that there is a closed form solution for the scalar FRF $h(\omega) \in \mathbb{C}$ between input $F(x_{\text{in}}, t)$ and output $w(x_{\text{out}}, t)$ with both the input and the output at the beam's midpoint, $x_{\text{in}} = x_{\text{out}} = L/2$:

$$h(\omega) = \frac{m}{4EI k_b^3} (\tan(k_b L/2) + \tanh(k_b L/2)), \quad (4.10)$$

where $k_b = (\omega^2 m / EI)^{1/4}$ [40]. While there are closed form solutions for a good number of special cases of FRFs [14], it is more common to consider a mode summation expansion [69] of the FRF. For the same beam, the FRF between any two points on the bar can be written as

$$h(\omega) = \sum_{j=1}^{\infty} \frac{u_j(x_{\text{in}}) u_j(x_{\text{out}})}{(i\omega - \lambda_j)(i\omega - \bar{\lambda}_j)}, \quad (4.11)$$

¹In Chapter 2.3.2 of [40], Fahy et al. derive the FRF for the beam's velocity (the mobility), which has been transformed to the FRF for the beam's displacement (the receptance) here for consistency.

where the eigenvalues of the PDE, $\lambda_j \in \mathbb{C}$ are the poles of the system, and the eigenfunctions of the system $u_j(x) \in \mathbb{R}$ are commonly referred to as the mode shapes. The mode shapes in this example of a beam without damping are real valued, and in fact any second order system without damping or with proportional damping will have real modes and can use this same FRF formulation [55]. A multi-input multi-output (MIMO) system with M inputs and N outputs will have a matrix valued FRF $H(\omega) \in \mathbb{C}^{N \times M}$:

$$H(\omega) = \sum_{j=1}^{\infty} \frac{\mathbf{u}_{j,\text{out}} \mathbf{u}_{j,\text{in}}^{\top}}{(i\omega - \lambda_j)(i\omega - \bar{\lambda}_j)}, \quad (4.12)$$

where $\mathbf{u}_{j,\text{out}} \in \mathbb{R}^N$ and $\mathbf{u}_{j,\text{in}} \in \mathbb{R}^M$ are vectors of the mode shape $u_k(x)$ sampled at the output locations and input locations respectively. However, if the system does not have proportional damping, the FRF must instead be written as:

$$H(\omega) = \sum_{j=1}^{\infty} \left(\frac{\mathbf{w}_{j,\text{out}} \mathbf{w}_{j,\text{in}}^{\top}}{i\omega - \lambda_j} + \frac{\bar{\mathbf{w}}_{j,\text{out}} \bar{\mathbf{w}}_{j,\text{in}}^{\top}}{i\omega - \bar{\lambda}_j} \right), \quad (4.13)$$

where the modal vectors $\mathbf{w}_{j,\text{out}} \in \mathbb{C}^N$ and $\mathbf{w}_{j,\text{in}} \in \mathbb{C}^M$ are now complex [55], coming from a complex valued function $w_j(x) \in \mathbb{C}$ sampled at the output locations and input locations respectively. These FRF formulations all demonstrate an important property of many second order systems, reciprocity. Reciprocity means that the FRF between two points is the same no matter which is the input location and which is the output location: $h_{12}(\omega) = h_{21}(\omega) = y_1(\omega)/u_2(\omega) = y_2(\omega)/u_1(\omega)$ [56]. This property is usually observed in second order systems, although some cases such as gyroscopic or circulatory forces can create a system which does not exhibit reciprocity [100, 105]. We will focus on systems which do display reciprocity.

The transfer function can also be defined from the state space representation. The matrix transfer function $H(s) \in \mathbb{C}^{N \times M}$ defined from the state space matrices in Equation 4.5 is

$$H(s) = \mathbf{C}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}. \quad (4.14)$$

The transfer function can also be built from the diagonalized state space model in Equation 4.7. In this case, the inverse of $(s\mathbf{I} - \Lambda)$ is diagonal, allowing the transfer function to be written in a pole-residue form:

$$H(s) = \sum_{j=1}^{\rho} \frac{\tilde{c}_j \tilde{b}_j^{\top}}{s - \lambda_j}, \quad (4.15)$$

where the poles λ_j are in the denominator, and there is a matrix residue $R_j = \tilde{c}_j \tilde{b}_j^{\top}$ in the numerator. Since we are interested in second order systems, the poles of the system come

in complex conjugate pairs, provided there are no real poles (an assumption that will be made in this chapter and the next). The pole-residue form can be rewritten considering the complex conjugate nature of the poles as:

$$H(s) = \sum_{j=1}^{\rho/2} \left(\frac{R_j}{s - \lambda_j} + \frac{\bar{R}_j}{s - \bar{\lambda}_j} \right). \quad (4.16)$$

This transfer function is equivalent to the second order FRF in Equation 4.13. The split complex conjugate poles can also be brought together for the equation:

$$H(s) = \sum_{j=1}^{\rho/2} \frac{\tilde{Q}_j s + \tilde{R}_j}{(s - \lambda_j)(s - \bar{\lambda}_j)}, \quad (4.17)$$

where $\tilde{Q}_j, \tilde{R}_j \in \mathbb{R}^{N \times M}$ are real matrices. The similarity between this transfer function and the FRF in Equation 4.12 become clear if $\tilde{Q}_j = 0$. This is the case when the mode shapes are real, and the residue matrix $\tilde{R}_j = \mathbf{u}_{j,\text{out}} \mathbf{u}_{j,\text{in}}^\top$, while $\tilde{Q}_j = 0$. Equation 4.12 will provide the starting point for the CRI method which is developed in detail in the next chapter.

4.3 Building Predictive Models

Many methods build predictive models by discretizing the PDEs of the system into matrices. Some methods build mass and stiffness matrices, such as the Rayleigh Ritz method (RRM) and the finite element (FE) method. This is the same process of discretizing PDEs as mentioned going from Equation 4.1 to Equation 4.2. With the discretized $\mathbf{M}$ and $\mathbf{K}$ matrices, an eigenvalue problem such as

$$(\mathbf{K} - \omega^2 \mathbf{M})\mathbf{v} = 0 \quad (4.18)$$

can be solved where ω are the natural frequencies and $\mathbf{v}$ are the eigenvectors which define the mode shapes.

In this section, some methods will be applied to a non-uniform bar with a spring at the free end as used by Meirovitch [70], and shown in Figure 4.3.

The system properties of the bar are:




Figure 4.3: A diagram of the Meirovitch bar system.

$$\begin{aligned}
 EA(x) &= \frac{6EA}{5} \left(1 - \frac{1}{2}x^2\right), \\
 m(x) &= \frac{6m}{5} \left(1 - \frac{1}{2}x^2\right), \\
 k &= EA(L),
 \end{aligned} \tag{4.19}$$

and the spring constant is chosen such that the boundary conditions are:

$$\begin{aligned}
 w(0) &= 0, \\
 w'(L) + w(L) &= 0.
 \end{aligned} \tag{4.20}$$

Three analytical methods and one data-driven method will be applied to the bar. One aspect for each method will be the focus of the discussion, as these aspects will be points of comparison for the next chapter's method CRI:

Rayleigh-Ritz Method: Choosing admissible functions for discretizing the PDE.

Collocation: Choosing points to minimize polynomial interpolation error.

Finite Element Method: High model order required for FE models.

Rational Approximation: Least squares fit to measured FRFs.

4.3.1 Admissible functions in the Rayleigh-Ritz Method

The Rayleigh-Ritz method (RRM) is widely used to predict natural frequencies and natural modes of dynamical systems [20, 53, 58]. The RRM discretizes the PDE by converting it into a weak form using some trial functions—often referred to as admissible functions in the RRM literature. The PDE is integrated over the admissible, giving a discrete mass and discrete

stiffness matrix. The elements of $\mathbf{M} \in \mathbb{R}^{\rho/2 \times \rho/2}$ and $\mathbf{K} \in \mathbb{R}^{\rho/2 \times \rho/2}$ for the bar example are built using the integrals

$$\begin{aligned}\mathbf{M}_{j,n} &= \int_0^L m(x) \phi_j(x) \phi_n(x) dx \\ \mathbf{K}_{j,n} &= \int_0^L EA(x) \frac{d\phi_j(x)}{dx} \frac{d\phi_n(x)}{dx} dx + k\phi_j(L)\phi_n(L)\end{aligned}\tag{4.21}$$

where $\phi_j(x)$ for $j = 1, \dots, \rho/2$ are the admissible functions [70]. The eigenvalue problem in Equation 4.18 is then solved to give the estimated natural frequencies, ω_j , and the eigenvector $\mathbf{v}_j \in \mathbb{R}^N$ which defines the approximated mode shape. For mode shape $u_j(x)$, the approximated mode shape $\hat{u}_j(x)$ is

$$\hat{u}_j(x) = \sum_{n=1}^{\rho/2} v_{j,n} \phi_n(x).\tag{4.22}$$

There has been much research put into the choice of admissible functions for RRM. The only requirements for a set of functions to be admissible is that they must be complete in energy and match the geometric boundary conditions [53]. This second requirement may be circumvented with penalty methods, however that is not the focus in this section, and we will focus on admissible functions meeting both requirements. Many families of admissible functions have been proposed, often with the goal of accelerating convergence [3, 13, 53, 58, 70]. Another consideration when choosing admissible functions is conditioning of the eigenvalue problem which can lead to numerical rounding errors.

Nine sets of admissible functions will be applied to the Meirovitch bar. The convergence as well as conditioning of the eigenvalue problem will be compared for each. The nine sets are summarized in Table 4.1, and more detailed explanations and equations are provided in Appendix A.

Each set of admissible functions were used to discretize the Meirovitch bar with increasing numbers of admissible functions. The error of modes 1 and 5 are compared in Figure 4.4. The true mode shape was chosen to be the mode shape from using 30 BC-Cheby admissible functions (modes 1-5 had all converged with this many admissible functions). The error for the mode is defined as

$$e_{j,i} = \|\mathbf{u}_{j,i} - \mathbf{u}_{j,30}\|_2\tag{4.23}$$

where $e_{j,i}$ is the error of the j th mode with i admissible functions, $\mathbf{u}_{j,i}$ is the normalized mode shape with i admissible functions, and $\mathbf{u}_{j,30}$ is the normalized converged mode shape both sampled at 200 linearly spaced points from 0 to 1.

Table 4.1: Each of the nine sets of admissible functions, along with the computational cost to generate each function, and if the set meets the geometric boundary condition at $x = 0$, BC_0 , and/or the natural boundary condition at $x = 1$, BC_1 , and if the set qualifies as quasi-comparison functions, QCF.

Admissible Function Set		cost	BC_0	BC_1	QCF
Comparison Functions	(CF)	non-linear solve	✓	✓	-
Admissible Functions	(AF)	-	✓	×	×
Quasi-Comparison Functions	(QCF)	-	✓	×	✓
Trig+Poly QCF	(QCF-Poly)	-	✓	×	✓
Monomial	(Monomial)	-	✓	×	✓
Chebyshev with Forced BC	(Forced Cheby)	-	✓	×	✓
Jacobi	(Jacobi)	-	✓	×	✓
Boundary Characteristic Orth. Poly.	(BCOP)	Gram-Schmidt	✓	✓	-
Boundary Condition Cheby. Poly.	(BC-Cheby)	linear solve	✓	✓	-




Figure 4.4: Comparison of errors for mode 1 and mode 5 of the bar system with various choices of admissible functions.

As seen in Figure 4.4, there is a wide range in the convergence rate of the admissible functions. The two slowest converging were the trigonometric admissible functions, and the boundary condition orthogonal polynomials. The transcendental comparison functions and trigonometric admissible functions augmented with the monomial x converged faster, but were not the fastest. The boundary condition Chebyshev polynomials and the quasi-comparison functions were the fastest converging choices for both modes. In the higher mode, the boundary condition Chebyshev functions converged slower than the quasi-comparison functions, however the BC-Cheby functions did not slow the rate of convergence like all other types other than quasi-comparison functions.

Convergence speed is not the only consideration when choosing admissible functions for RRM. Some choices for admissible functions lead to an ill-conditioning of the eigenvalue

problem and large numerical errors. Figure 4.5 shows the condition number of the matrix $\mathbf{M}$. The eigenvalues v from the generalized eigenvalue problem in Equation 4.18 which is used are used to define the mode shape as shown in Equation 4.22 are orthogonal with respect to the mass matrix: $\mathbf{v}_i^\top \mathbf{M} \mathbf{v}_j = \delta_{ij}$. Therefore, if the matrix $\mathbf{M}$ is poorly conditioned, it can lead to numerical problems while solving the eigenvalue problem for the discretized PDE. As shown in Figure 4.5, one of the fastest converging choices, the QCF, also has a condition number orders of magnitude higher than the others. This shows the same pattern as monomials, fast convergence at the cost of a high condition number which may lead to numerical errors.




Figure 4.5: Condition number of the $\mathbf{M}$ matrix for various choices of admissible functions.

The combination of convergence speeds and condition numbers suggest that orthogonal polynomials are a good choice for admissible functions in the RRM. However, the difference between BCOP and BC-Cheby suggests that the choice of orthogonal polynomials plays a role. The polynomial order of BCOP is higher than necessary due to the generating function—as discussed in Appendix A—while each polynomial order for BC-Cheby is as low as possible.

4.3.2 Choice of collocation locations

The collocation method uses polynomial interpolation to turn a discrete eigenvector from an eigenvalue problem into a continuous mode shape [112]. The operator of the PDE is discretized at a choice of collocation locations, and resulting eigenvalue problem

$$(\mathbf{I} - \omega^2 \mathbf{L}) \mathbf{u} = 0, \quad (4.24)$$

is solved where ω is the natural frequency, $\mathbf{u}$ is the mode shape at the collocation locations, and $\mathbf{L}$ is the discretized operator of the PDE. Because of the reliance on polynomial inter-

polation, the accuracy of the method depends on the choice of points where the operator is discretized.

Polynomial interpolation is subject to Runge's phenomenon: interpolation accuracy with uniformly spaced points gets worse with higher polynomial order for some functions. However, if the interpolation points are grouped more densely near the boundaries, this phenomenon can be avoided [112]. Figure 4.6 shows polynomial interpolation for the fifth mode of Meirovitch's bar, as well as the Runge function

$$f(x) = \frac{1}{1 + 16x^2}, \quad (4.25)$$

for both uniformly spaced support points and Chebyshev nodes. Chebyshev nodes are an example of a spacing which groups points near the boundaries, and is known to give better results for polynomial interpolation [12], and is therefore sometimes used to choose collocation locations [112].




Figure 4.6: Comparison of polynomial interpolation accuracy with uniformly spaced points and Chebyshev nodes. Polynomial order 8 is used for both types of spacing, and for both functions.

Using the same functions from Figure 4.6 at multiple polynomial orders, Table 4.2 shows the utility of choosing Chebyshev nodes for polynomial interpolation over uniformly spaced points. The Runge function demonstrates Runge's phenomenon, where the maximum error grows with polynomial order for the uniformly spaced points. Interpolation with the Chebyshev nodes, however, are able to fit the Runge function better with increasing polynomial order. The fifth mode of the Meirovitch bar does not show increasing error with uniformly spaced points, but the interpolation has a lower error at all orders with the Chebyshev nodes. This suggests that in physical systems, although Runge's phenomenon of increasing error may not be encountered, the interpolation is still aided by choosing Chebyshev nodes.

Table 4.2: Maximum polynomial interpolation error for multiple polynomial orders for two functions. The same functions and polynomial spacing from Figure 4.6.

Poly Order	$u_5(x)$		$f(x)$	
	Uniform	Cheby	Uniform	Cheby
8	2.41e-1	1.30e-1	0.761	0.164
10	5.88e-2	1.81e-2	1.238	0.096
12	9.80e-3	1.70e-3	2.029	0.055
14	1.10e-3	1.00e-5	3.244	0.033
16	1.12e-4	4.87e-6	5.630	0.021

4.3.3 Model order of Finite Element models

Analytical methods like the RRM and collocation are capable of very high accuracy with low model orders [53, 112]. However, these methods are often restricted to relatively simple systems like beams and bars with well defined boundary conditions. When attempting to model a system with multiple connected members or complex boundary conditions the application of these methods becomes more difficult.

The finite element (FE) method, on the other hand, is a more general method; with FE modelling, it is easy to combine multiple components, and to apply a wide range of boundary conditions. This often comes at the cost of model order. In order to compare FE and RRM, an FE model of the Meirovitch bar is created with an increasing number of elements. For each additional element, the model order rises by two. This shows the fundamental difference between FE and the other methods discussed: for the FE method, the model order is tied to spatial resolution, whereas for other analytical methods it is not. Figure 4.7 shows a diagram of the FE model of the beam broken into three elements. Each element is considered to be a uniform bar, with the physical properties of the true bar at the element's midpoint.




Figure 4.7: Diagram of a three element finite element model of the bar system.

Figure 4.8 shows a comparison of the convergence for $u_1(x)$ and $u_5(x)$ with the RRM using two sets of admissible functions, and with the FE method. This example demonstrates the relatively slow convergence of FE models, especially in comparison to the RRM with well chosen admissible functions.




Figure 4.8: Comparison of the error for mode 1 and mode 5 using RRM and FE modelling at the same model orders.

4.3.4 Rational Approximation

An influential data-driven method for building predictive models of dynamical systems is rational approximation. Rational approximation assumes that the transfer function (or FRF) can be expressed as a rational function:

$$\tilde{H}(s) = \frac{n(s)}{d(s)} = \frac{\sum_{i=1}^k \frac{\beta_i}{s - \sigma_i}}{\sum_{i=1}^k \frac{\alpha_i}{s - \sigma_i}}, \quad (4.26)$$

where $n(s)$ is the numerator function and $d(s)$ is the denominator function. A popular tool in rational approximation is to use the Barycentric representation of rational polynomials as shown in Equation 4.26, where the unknowns are $\{\beta_j\}$, $\{\alpha_j\}$, and $\{\sigma_j\}$ [48, 74]. The goal of these algorithms is to find the closest matching rational approximation of the system's transfer function. Some rely on least-squares fit and/or interpolation to measured transfer function data [11, 48, 67, 74], while some other algorithms require the ability to evaluate the transfer function [46]. We will focus on those that include a least-squares fit to measured FRFs. While some of these rational approximation methods were initially formulated for scalar functions, extensions have been made to matrix valued functions as seen in multi-input multi-output (MIMO) systems [36, 43].

One particularly relevant rational approximation method is called vector fitting [32, 47, 48]. Vector fitting leverages an iterative scheme using the Barycentric representation in Equation 4.26 to find the poles of the system, and then in a separate step finds the residues using the pole-residue formulation from Equation 4.15. This two step process highlights the similarities between rational approximation and modal identification. As Allemang et. al show in reviewing many common modal identification methods [4], these also commonly take

a two stage approach: first identifying poles and modal participation factors, then extracting modal vectors. There are many other rational approximation methods. Some of these focus on interpolation of transfer function data [67], while others combine interpolation and least-squares [74]. Figure 4.9 shows how a vector fitting model can be fit using a simulated experiment on the Meirovitch bar [70]. Five output locations were measured at $x = j/5L$, with $j = 1, \dots, 5$, with the input location at the free end. The FRFs were measured at 1000 uniformly spaced frequencies over the frequency range. Rational approximation is capable of creating an accurate model even at low model orders, however, any resulting predictive model such as the fit in Figure 4.9 is restricted to the tested input and output locations.




Figure 4.9: Two vector fitting models, one with model order 2 and one with model order 10. The Meirovitch bar has 5 modes in this frequency range, thus the model order 10 fit is able to recreate the FRFs well as seen on the right.

4.4 Conclusion

This chapter has reviewed the theory of second order systems and examined aspects from four methods for building predictive models. When CRI is introduced in the next Chapter, the concepts of the state space model and especially transfer functions will be important in building the data-driven model. The focus of the discussion from each of the modelling methods will also provide useful context and comparison for CRI in the next chapter.

In the Rayleigh-Ritz method, a set of admissible functions are used to discretize the PDE. The choice of these functions has an effect on how quickly the discretized model converges to the true system, and if poorly chosen can lead to an ill-conditioned problem which leads to numerical errors. Orthogonal sets of polynomials are a good choice for quick convergence, while staying well conditioned. Additionally, if more information about boundary conditions can be incorporated, similar accuracy can be achieved with fewer functions. As will be seen in Chapter 5, CRI will also use basis functions to approximate the modes of a system, and orthogonal polynomials will be chosen for the same reasons that they work well in the RRM.

The choice of locations in collocation methods similarly has an effect on model accuracy. Since collocation relies on polynomial interpolation, large errors can arise from rapid oscillations near the boundaries if the interpolation points are not well chosen. Chebyshev nodes provide a good choice of interpolation points for avoiding this problem, known as Runge's phenomenon. Since CRI is a data-driven method built on measured data, there is a similar choice when choosing where to measure inputs and outputs on the system. Therefore the lessons learned from collocation will also be useful.

It was shown also that in the finite element method, spatial resolution is tied to model order. This leads to slow convergence of the model and low accuracy compared to methods where model order is separated from spatial resolution. The purpose of CRI is to mimic the path taken by analytical methods like the RRM from a data-driven angle. The basis functions used by CRI, like in the RRM, lead to a spatial resolution not dependent on model order, allowing an accurate model to be created with a low model order.

Finally, the idea of creating data-driven predictive models from measured FRF data was discussed through the idea of rational approximation. While rational approximation and other data-driven methods exist and have been applied to second order systems, the restriction of spatial resolution to only measured locations limits their usefulness in some cases. The next chapter will introduce a data-driven technique designed specifically to take measured FRFs from second order systems, and create a predictive model with continuous spatial resolution.

Chapter 5

A Novel Data-Driven Method to Build Continuous Predictive Models

This chapter introduces a novel data-driven method to build predictive models of dynamical systems with continuous spatial resolution. The proposed method, continuous residue interpolation (CRI), is demonstrated on two simulated systems and one experimentally measured system, before the application of CRI to VBOI localization is demonstrated.

5.1 Introduction

Simulations of dynamical systems are useful in VBOI [30, 34], as well as environmental testing [27]. It is sometimes necessary to have a model of a system that can be used for predictive simulations at many input and output locations. For example, the Impedance Matched Multi-Axis Testing (IMMAT) methodology uses a model to optimize environmental testing procedures by considering many combinations of input locations to provide optimal excitation to match environmental conditions [27]. Models with continuous (or fine) spatial resolution may be built by discretizing continuous partial differential equations (PDEs) with spectral methods, finite elements, or other numerical methods. For a linear time invariant (LTI) system, the resulting models can be written as a system of ordinary differential equations:

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t); \quad \mathbf{y}(t) = \mathbf{C}\mathbf{x}(t), \quad (5.1)$$

where $\mathbf{x}(t) \in \mathbb{R}^\rho$ are the states of the system; $\mathbf{A} \in \mathbb{R}^{\rho \times \rho}$, $\mathbf{B} \in \mathbb{R}^{\rho \times M}$, $\mathbf{C} \in \mathbb{R}^{N \times \rho}$ are constant; $\mathbf{u}(t) \in \mathbb{R}^M$ is the input (forcing) term; and $\mathbf{y}(t) \in \mathbb{R}^N$ are the measured outputs of the system. This is the state space representation for the model with ρ states, M inputs, and N outputs.

However, high quality state space models are not always available. For example, a system may have intricate boundary conditions that are difficult to model, or there may be materials which have unknown properties. If an already built system is of interest, it is unlikely that a high fidelity finite element model will be available, and the exact geometry needed to create such a model may not be available either. In the absence of an accurate state space model

built from first principles, it is possible to build a data-driven model from transfer function evaluations [11, 48, 74]. The transfer function $H(s)$,

$$H(s) = \mathbf{C}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B} \quad (5.2)$$

of the dynamical system in Equation 5.1 contains the information from the state space model, and therefore captures the dynamics of the system. However, data-driven models are limited to the locations where the transfer function was measured, thus creating a model with fine spatial resolution would require an infeasible amount of testing.

The goal of this Chapter is to present a methodology that creates an approximated predictive model of a physical dynamical system with continuous spatial resolution for both inputs and outputs using only discrete testing data. Continuous spatial resolution will be achieved by assuming that modes can be approximated by a linear combination of smooth continuous basis functions. In a second order system these modes are the linear modes of a system, and are related to the residue in the pole-residue formulation:

$$H(s) = \sum_{i=1}^{\rho} \frac{R_i}{s - \lambda_i}. \quad (5.3)$$

The poles of the system, λ_i , do not change with spatial location of inputs and outputs in a LTI second order system. Only the residues, $R_i \in \mathbb{C}^{N \times M}$, are affected by the spatial location through the modes. Since the residues are separated into a partial fraction expansion, each of the modes can be fitted separately by observing the residue of the measured transfer functions. The basis functions will be used to expand the residues, and a quadratic optimization problem will be set up to minimize the transfer function approximation error at the measured input/output location pairs. After fitting the model, the basis functions can be sampled at any combination of input and output locations to create the interpolated model.

The method presented in this chapter, continuous residue interpolation (CRI) will be proposed in order to take sparse spatial transfer function measurements and create a predictive model with continuous resolution for both inputs and outputs. Section 5.2 will explore some connections between CRI and other current methods for modelling second order systems. Then, the CRI algorithm will be detailed in Section 5.3. Finally in Section 5.4, a group of applications of CRI will be presented in order to show how it is applied and to demonstrate the viability of the method.

5.2 Comparison to Current Methods for Modelling Second Order Systems

Before going more into the specifics of the proposed method, this section will explore connections to other second order system modelling methods. The transfer function formulation which will be used in CRI and its connection to state space models was shown in Chapter 4. In this section, comparisons of the proposed method with current analytical and current data-driven methods will be made.

5.2.1 Comparison to analytical methods

The goal of CRI is to provide a single data-driven method that mirrors the path taken by spectral methods to build a predictive model of a second order system. In particular, two spectral methods which bear similarity to CRI are the Rayleigh-Ritz method (RRM) and collocation. In RRM, a set of admissible functions are used to build mass and stiffness matrices from the PDE [53]. A generalized eigenvalue problem is then solved from those mass and stiffness matrices to find the stationary points which approximate the natural frequencies and modes. In collocation methods, a number of collocation points are chosen over which the PDE operator is discretized, and once again an eigenvalue problem is solved [112]. There are clear parallels between CRI and these methods. The choice of a set of continuous functions in RRM is similar to the basis functions used to approximate the modes in CRI. Also, the choice of a discrete set of spatial locations in collocation is similar to the choice of input and output locations to measure the transfer function which is in turn used to fit the CRI model. However, instead of being able to directly solve an eigenvalue problem related to the PDE, CRI uses the transfer function measurements to approximate the same poles and modes as those eigenvalue problems. Figure 5.1 shows the similarities between the procedures for CRI, RRM, and collocation.

Another similarity between CRI and spectral methods is that the model order is not tied to the spatial resolution like it is with finite element (FE) methods. While FE methods provide a more general tool than spectral methods, making it easier to incorporate complex geometries and boundary conditions, it is well known that spectral methods can provide the same accuracy as FE model with a much lower model order [70]. CRI takes this a step further; the model order for CRI is fixed by the number of modes in the measured frequency range, and the accuracy of the model can be increased within that frequency range by testing more locations and using more basis functions without increasing the model order. The sizes of models created with CRI, RRM, collocation, and the FE method are also compared in Figure 5.1. In RRM the model order is twice the number of admissible functions used, and in collocation it is twice the number of collocation points. These are generally higher than the model order for CRI which is only twice the number of modes in the frequency range, although all are much smaller than the model order of a finite element model. In general,

	Data Driven	Analytical		
	CRI	RRM	Collocation	FEM
Begin with	Physical System	PDE, BC	PDE, BC	Geometry, BC
	↓	↓	↓	↓
Choose	Basis Functions Testing Locations	Admissible Functions	Collocation Points	Element Type Mesh
	↓	↓	↓	↓
Interrogate System	Test System $H(s) \in \mathbb{C}^{N \times M}$	Integrate PDE $M, K \in \mathbb{R}^{N_k \times N_k}$	Discretize Operator $L \in \mathbb{R}^{N_k \times N_k}$	Build $M, K \in \mathbb{R}^{N_{\text{DOF}} \times N_{\text{DOF}}}$ from Elements
	↓	↓	↓	↓
Eigenvalue Problem	LS fit to $H(s)$	$(K - \omega_{\text{RR}}^2 M)v_{\text{RR}} = 0$	$(\mathbb{I} - \omega_{\text{Co}}^2 L)u_{\text{Co}} = 0$	$(K - \omega_{\text{FEM}}^2 M)u_{\text{FEM}} = 0$
	↓	↓	↓	↓
Continuous Mode Shapes	$U_{\text{CRI}} = \Phi V$	$U_{\text{RR}} = \Phi V_{\text{RR}}$	Polynomial Interpolate U_{Co}	Shape Function Interpolate U_{FEM}
	↓	↓	↓	↓
Model Order	$2N_m$	$2N_k$	$2N_k$	$2N_{\text{DOF}}$

Figure 5.1: Similarity between steps in CRI and three analytical methods.

$$N_m < N_k \ll N_{\text{NDOF}}.$$

Analytical and numerical methods require not only knowledge of the PDE, but also complete knowledge of the geometry of the system, boundary conditions, and all material properties. Although they are capable of building very accurate models when provided all of this information, in some cases knowledge is lacking in at least one of these areas [41, 82]. This limitation has led to the growing role of data-driven methods in characterizing and modelling dynamical systems.

5.2.2 Comparison to data-driven methods

Data-driven methods have been around for a long time, and have gained increased traction in recent years. Connections to data-driven methods which operate on measured transfer function data or mode shapes are particularly relevant to the discussion of CRI.

Classical modal identification algorithms are data-driven system-identification methods. These algorithms generally operate on frequency domain or time domain data, and those that use frequency domain data (transfer function measurements) usually solve two consecutive least squares problems to identify poles and modal participation factors, then the modal vectors [4]. As shown in Section 5.2, the residue depends on both the input and output location. By solving for participation factors and modal vectors separately, often modal identification algorithms ignore the input locations in the modal vector. Instead, CRI will seek to fit the entire residue with a quadratic optimization problem. In this way, the input location information is kept and the focus is put on fitting the whole residue which is more directly observed in the transfer function as shown in Equation 4.16.

Another related family of data-driven methods are dynamic vibration expansion (DVE)

methods. These are used to expand discretely measured vibration responses to full-field responses. Some of the most used methods in this area require finite element mode shapes [6, 78], however it is also possible to do non-model based expansion [24]. Furthermore, there is evidence that DVE methods are still viable without a correlated finite element model which relaxes the reliance on previous models [66]. However, to the authors knowledge at this time, all current vibration expansion methods operate directly on responses or extracted mode shapes. The proposed CRI method operates on a similar principle to these methods in expanding mode shapes. However, performing DVE on responses does not help with predictive simulations, and performing DVE on mode shapes requires the mode shapes to be estimated; relying on a separate modal identification algorithm.

There is also a large literature into rational approximation of dynamical systems in the field of mathematics and model order reduction. Many of these algorithms rely on the interpolation and/or a least-squares fit to measured transfer function data [11, 48, 67, 74]. However, any resulting predictive model is restricted to the tested input and output locations. These tend to be quite similar to modal identification algorithms, although they can be applied to a wider range of problems such as model order reduction [5], and parametric systems [18, 63]. The goal of CRI is in a way to parameterize a rational approximation problem, with the parameter being space. However the needs of CRI are different from parametric rational approximation methods in literature. Parametric rational approximation algorithms generally focus on parameterizing $\mathbf{A}(p)$, whereas in order to expand spatial resolution it would be more important to parameterize $\mathbf{B}(p)$ and $\mathbf{C}(p)$ [10, 18, 63].

5.3 Continuous Residue Interpolation Algorithm

In this section, the CRI algorithm will be introduced. We will start by setting up a quadratic optimization problem, and then discussing how it can be optimized. Finally, some implementation considerations will be discussed.

5.3.1 Setting up the optimization problem

The core idea of the CRI algorithm is that the continuous mode shape functions, $u_i(\mathbf{r}) \in \mathbb{R}$, can be approximated as a linear combination of basis functions:

$$\tilde{u}_i(\mathbf{r}) = \sum_{j=1}^d v_{i,j} \phi_j(\mathbf{r}) \quad (5.4)$$

where $v_{i,j}$ are the coefficients corresponding to the basis functions $\phi_j(\mathbf{r})$, $j = 1, \dots, d$ for the approximated mode $\tilde{u}_i(\mathbf{r})$, and $\mathbf{r} \in \mathbb{R}^n$ is a general spatial vector for a problem with n spatial

dimensions. The approximation to the spatially sampled modal vectors can be written as the product of a matrix and vector:

$$\tilde{\mathbf{u}}_{i,\text{out}} = \Phi_{\text{out}} \mathbf{v}_i \quad \text{and} \quad \tilde{\mathbf{u}}_{i,\text{in}} = \Phi_{\text{in}} \mathbf{v}_i, \quad (5.5)$$

where the coefficient vector $\mathbf{v}_i = [v_{i,1}, \dots, v_{i,d}]^\top$ is the same for the input and output modal vectors, and the interpolation matrices are

$$\Phi_{\text{out}} = \begin{bmatrix} \phi_1(\mathbf{r}_{o_1}) & \cdots & \phi_d(\mathbf{r}_{o_1}) \\ \phi_1(\mathbf{r}_{o_2}) & \cdots & \phi_d(\mathbf{r}_{o_2}) \\ \vdots & \ddots & \vdots \\ \phi_1(\mathbf{r}_{o_N}) & \cdots & \phi_d(\mathbf{r}_{o_N}) \end{bmatrix} \quad \text{and} \quad \Phi_{\text{in}} = \begin{bmatrix} \phi_1(\mathbf{r}_{i_1}) & \cdots & \phi_d(\mathbf{r}_{i_1}) \\ \phi_1(\mathbf{r}_{i_2}) & \cdots & \phi_d(\mathbf{r}_{i_2}) \\ \vdots & \ddots & \vdots \\ \phi_1(\mathbf{r}_{i_M}) & \cdots & \phi_d(\mathbf{r}_{i_M}) \end{bmatrix}. \quad (5.6)$$

The output interpolation matrix $\Phi_{\text{out}} \in \mathbb{R}^{N \times d}$ consists of the basis functions sampled at the output locations $\{\mathbf{r}_{o_1}, \dots, \mathbf{r}_{o_N}\}$ and the input interpolation matrix $\Phi_{\text{in}} \in \mathbb{R}^{M \times d}$ consists of the basis functions sampled at the input locations $\{\mathbf{r}_{i_1}, \dots, \mathbf{r}_{i_M}\}$.

Plugging in the approximated mode vectors into Equation 4.12, the transfer function approximation is:

$$\tilde{H}(s) = \sum_{i=1}^{N_m} \frac{\Phi_{\text{out}} \mathbf{v}_i \mathbf{v}_i^\top \Phi_{\text{in}}^\top}{(s - \lambda_i)(s - \bar{\lambda}_i)}. \quad (5.7)$$

The goal will be to minimize the error between the measured transfer function $H(s)$ and the approximation $\tilde{H}(s)$. The optimization will search for optimal $\mathbf{v}_i$, while the poles $\{\lambda_i\}$ are assumed to be known. However, $H(s)$ and $\tilde{H}(s)$ are both matrices, and we want to be able to consider more than one frequency at a time, making setting up the optimization problem more difficult. Therefore, first we will vectorize the transfer function approximation at each frequency:

$$\tilde{F}_k = \sum_{i=1}^{N_m} (\Phi_{\text{in}} \otimes S_{i,k})(\mathbf{v}_i \otimes \mathbf{v}_i), \quad (5.8)$$

where $\tilde{F}_k$ is an approximation of the vectorized measured transfer function at the frequency s_k , $F_k = \text{vec}(H(s_k))$, and $S_{i,k} = \Phi_{\text{out}} / [(s_k - \lambda_i)(s_k - \bar{\lambda}_i)]$ is a matrix which resembles a Cauchy matrix that contains the output interpolation matrix and frequency information. The vectorization identities $\text{vec}(AXB) = B^\top \otimes A \text{vec}(X)$ and $\text{vec}(\mathbf{x}\mathbf{x}^\top) = \mathbf{x} \otimes \mathbf{x}$ are leveraged for simplification. This vectorization also separates the constant parts of the problem (the interpolation matrices and denominator) and the part to be optimized (the coefficient vectors). This will make finding the Jacobian for optimization easier.

The final optimization problem, which is quadratic in the vectors $\mathbf{v}_i$, is then

$$\min_{\{\mathbf{v}_1, \dots, \mathbf{v}_{N_m}\}} \left\| F - \tilde{F}(\mathbf{v}_1, \dots, \mathbf{v}_{N_m}) \right\|_2, \quad (5.9)$$

where F_k , $\tilde{F}_k$ and $S_{i,k}$ are stacked for all frequencies into F , $\tilde{F}$, and S_i , and the other parts of the sum are left at their original size. We end up with a final sizes of $\tilde{F}$, $F \in \mathbb{C}^{NMN_s}$ and $S_i \in \mathbb{C}^{NMN_s \times d}$ where N_s is the number of frequency samples.

5.3.2 Solving the optimization problem

The optimization problem in Equation 5.9 is what we want to solve to find the mode coefficient vectors. It is quadratic, so it cannot be solved in a single step. There are a range of iterative methods that could be employed [60]. Before choosing any optimization method, an objective function for the error must be defined:

$$J = \left(F - \sum_{i=1}^{N_m} (\Phi_{\text{in}} \otimes S_i)(\mathbf{v}_i \otimes \mathbf{v}_i) \right)^* \left(F - \sum_{j=1}^{N_m} (\Phi_{\text{in}} \otimes S_j)(\mathbf{v}_j \otimes \mathbf{v}_j) \right), \quad (5.10)$$

where $(*)$ represents the complex conjugate. All of the mode coefficient vectors can be combined into a single vector $\mathbf{x} = [\mathbf{v}_1^\top \dots \mathbf{v}_{N_m}^\top]^\top$, and the objective function can be minimized using iterative methods. For example, straight forward Newton iterations could be used:

$$\mathbf{x}_{(k+1)} = \mathbf{x}_{(k)} - \frac{1}{\alpha} \text{Re} \left[\left(\frac{\partial^2 J}{\partial \mathbf{x} \partial \mathbf{x}^\top} \right)^{-1} \frac{\partial J}{\partial \mathbf{x}} \right], \quad (5.11)$$

where the Jacobian and second derivative are used to help the iterations converge faster, $\alpha \geq 1$ is a damping term that can be used to control the step size, and the vector $\mathbf{x}$ is kept real since the modes are assumed to be real. For the derivation of the derivatives, see Appendix B. Other iterative optimization schemes could also be used, which may or may not also incorporate the same derivative information. Appendix C includes MATLAB code for setting up this optimization problem and finding $\{\mathbf{v}_j\}$, as well as example code for applying CRI to a measured system.

5.3.3 Interpolated Points and State Space

After the optimization problem has converged, a model can be built with any input and output locations. The poles and coefficient vectors remain constant, so to interpolate to

new input and output locations, new interpolation matrices $\hat{\Phi}_{\text{out}}$ and $\hat{\Phi}_{\text{in}}$ simply need to be defined:

$$\hat{\Phi}_{\text{out}} = \begin{bmatrix} \phi_1(\mathbf{r}_{o_1}) & \cdots & \phi_d(\mathbf{r}_{o_1}) \\ \phi_1(\mathbf{r}_{o_2}) & \cdots & \phi_d(\mathbf{r}_{o_2}) \\ \vdots & \ddots & \vdots \\ \phi_1(\mathbf{r}_{o_{N_2}}) & \cdots & \phi_d(\mathbf{r}_{o_{N_2}}) \end{bmatrix} \quad \text{and} \quad \hat{\Phi}_{\text{in}} = \begin{bmatrix} \phi_1(\mathbf{r}_{i_1}) & \cdots & \phi_d(\mathbf{r}_{i_1}) \\ \phi_1(\mathbf{r}_{i_2}) & \cdots & \phi_d(\mathbf{r}_{i_2}) \\ \vdots & \ddots & \vdots \\ \phi_1(\mathbf{r}_{i_{M_2}}) & \cdots & \phi_d(\mathbf{r}_{i_{M_2}}) \end{bmatrix}, \quad (5.12)$$

where N_2 and M_2 are the new number of output and input locations and the same d basis functions are used. The interpolated transfer function $\hat{H}(s) \in \mathbb{C}^{N_2 \times M_2}$ is:

$$\hat{H}(s) = \sum_{j=1}^{N_m} \frac{\hat{\Phi}_{\text{out}} \mathbf{v}_j \mathbf{v}_j^\top \hat{\Phi}_{\text{in}}^\top}{(s - \lambda_j)(s - \bar{\lambda}_j)}. \quad (5.13)$$

With the new interpolation matrices, simulations can be run directly with the interpolated transfer function in Equation 5.13 using fast Fourier transforms (FFTs) of inputs and inverse FFTs of outputs. A state space model can also be created:

$$\dot{\mathbf{x}}(t) = \hat{\mathbf{A}}\mathbf{x}(t) + \hat{\mathbf{B}}\mathbf{u}(t); \quad \mathbf{y}(t) = \hat{\mathbf{C}}\mathbf{x}(t), \quad (5.14)$$

where the state space matrices are defined as: $\hat{\mathbf{A}} = \text{diag}(\lambda_1, \bar{\lambda}_1, \dots, \lambda_{N_m}, \bar{\lambda}_{N_m})$, $\hat{\mathbf{B}} = V^\top \hat{\Phi}_{\text{in}}^\top$, and $\hat{\mathbf{C}} = -\frac{i}{2} \hat{\Phi}_{\text{out}} V \text{Im}(\hat{\mathbf{A}})^{-1}$ where $V = [\mathbf{v}_1, \mathbf{v}_1, \dots, \mathbf{v}_{N_m}, \mathbf{v}_{N_m}]$. These relationships come from relating $\tilde{R}_i$ and R_i in Equation 4.16.

5.3.4 Implementation Considerations

Before exploring some applications of CRI, we will look at some of the choices that go into implementing using CRI.

Pole estimation In order for the CRI algorithm to run, it must be provided with the poles of each mode to be fitted. If the poles are unknown, they must be estimated. Modal identification and rational approximation are well researched [4, 43]; providing many choices for pole estimation. For example, least squares algorithms could be used such as Polymax [85] or vector fitting [48]. Straight forward heuristic methods could also be used, such as peak picking for natural frequencies with the 3dB method or circle fitting methods for damping ratios. The poles can be defined by a natural frequency, ω_j , and a damping ratio, ζ_j : $\lambda_j, \bar{\lambda}_j = -\omega_j(\zeta_j \pm i\sqrt{1 - \zeta_j^2})$. Pole estimates could be further refined within CRI by being included in the optimization problem:

$$\min_{\{\mathbf{v}_1, \dots, \mathbf{v}_{N_m}, \lambda_1, \dots, \lambda_{N_m}\}} \left\| F - \tilde{F}(\mathbf{v}_1, \dots, \mathbf{v}_{N_m}, \lambda_1, \dots, \lambda_{N_m}) \right\|_2. \quad (5.15)$$

This would require extra terms to be fitted and changes to the Jacobian and updating procedure, but would limit the reliance on the accuracy of external algorithms.

Choosing output and input locations Choosing input and output locations for testing will have real world considerations. For instance, sensors or inputs can only be placed on the exterior surface of a hollow body. Inputs from a hammer or shaker also cannot obstruct a laser measurement system and therefore cannot get too close to measurement locations. Nevertheless, there are still some guiding principles that may be helpful for choosing locations. Since CRI is based on interpolation, grouping output measurement locations near to the edges of boundary conditions is helpful. For example, Chebyshev nodes decrease polynomial interpolation error [12] and are often chosen when using collocation [112].

Input locations, on the other hand, would be best chosen so that all modes are provided excitation. Some works have shown the utility of the discrete empirical interpolation method (DEIM [22] or its variant Q-DEIM [35]) in sensor placement [65]. Applying the Q-DEIM algorithm to the basis functions provides a quantitative method which should ensure that energy is put into any mode that is similar to the basis functions.

The recommendations for output and input locations assume there are more output locations than input locations: $N > M$. However, if $M > N$ then the advice could be flipped: input locations chosen as Chebyshev nodes and output locations chosen with DEIM.

Number of basis functions There is a limited amount of testing that can be done to provide transfer function measurements to the CRI algorithm. With M measured input locations and N measured output locations, the residues have $N + M - 1$ degrees of freedom. This is because of how the rank 1 residue is formed:

$$\tilde{R}_i = \mathbf{u}_{i,\text{out}} \mathbf{u}_{i,\text{in}}^\top = \begin{bmatrix} u_{o1} u_{i1} & \cdots & u_{o1} u_{iM} \\ u_{o2} u_{i1} & \cdots & u_{o2} u_{iM} \\ \vdots & \ddots & \vdots \\ u_{oN} u_{i1} & \cdots & u_{oN} u_{iM} \end{bmatrix}, \quad (5.16)$$

where $\tilde{R}_i \in \mathbb{R}^{N \times M}$, $\mathbf{u}_{i,\text{out}} \in \mathbb{R}^N$ has entries $\{u_{o1}, \dots, u_{oN}\}$, and $\mathbf{u}_{i,\text{in}} \in \mathbb{R}^M$ has entries $\{u_{i1}, \dots, u_{iM}\}$. The ambiguity of scaling between $\mathbf{u}_{i,\text{out}}$ and $\mathbf{u}_{i,\text{in}}$ leads to the loss of a degree of freedom, which traditionally may require the measurement of a driving point FRF [103]. The basis function expansion removes the need for a driving point FRF by sampling the values for the input and output mode vectors from the same smooth basis functions. However, having more basis functions than degrees of freedom may lead to issues in the

optimization, so it is desired that the basis functions converge quickly to the modes since adding more basis functions requires more testing.

Choosing basis functions The similarities between admissible functions in the RRM and the basis functions used in CRI were noted in Section 5.2. Although they play a similar role for approximating the modes of the system, some requirements for admissible functions in the RRM do not exist for CRI. In the RRM the admissible functions must match the geometric boundary conditions and be complete [53, 111]. However, in CRI there is no need for the basis functions to match the boundary conditions, as these can be learned by the optimization problem as long as the basis functions are not more restricted than the mode shapes. It is only necessary for the basis functions to span as much of the modes to be fitted as possible to get a good approximation. This is clear from the definition of the approximated mode:

$$\tilde{u}_i(\mathbf{r}) = \sum_{j=1}^d v_{i,j} \phi_j(\mathbf{r}) \quad (5.17)$$

where the approximation error only comes from the part of the mode shape not in the set of d basis functions.

There has been a lot of research put into the effect of admissible functions on convergence in the RRM. It is known that polynomials converge quickly [53], can be well scaled if sets of orthogonal polynomials are chosen [20], and that the accuracy is tied only to polynomial order [17]. Therefore orthogonal polynomials are a natural choice due to their rapid convergence and conditioning properties. Further, some boundary conditions can be imposed on the polynomials [3, 58], thus increasing the polynomial order without increasing the number of basis functions.

Complex Modes The CRI algorithm as described in this section has assumed that the mode shapes of the system are real. In reality, in systems without specific damping types, the mode shapes are complex valued. The complex modal vectors can be expanded using interpolation matrices Φ_{out} and Φ_{in} in the same manner they were when assuming real modes. This results in the approximated transfer function:

$$\hat{H}(s) = \sum_{j=1}^{N_m} \left(\frac{\Phi_{\text{out}} \mathbf{v}_j \mathbf{v}_j^{\top} \Phi_{\text{in}}^{\top}}{s - \lambda_j} + \frac{\Phi_{\text{out}} \bar{\mathbf{v}}_j \bar{\mathbf{v}}_j^{\top} \Phi_{\text{in}}^{\top}}{s - \bar{\lambda}_j} \right), \quad (5.18)$$

where the mode coefficient vectors $\mathbf{v}_j$ are now complex. From here, the optimization problem can be set up in a similar way. The difference is that each mode will have two residues (which are conjugates of one another), and need two different $S_{j,k}$ matrices:

$$\tilde{F}_k = \sum_{j=1}^{N_m} \left((\Phi_{\text{in}} \otimes S_{j,k}^{(\mathbf{v})})(\mathbf{v}_j \otimes \mathbf{v}_j) + (\Phi_{\text{in}} \otimes S_{j,k}^{(\bar{\mathbf{v}})})(\bar{\mathbf{v}}_j \otimes \bar{\mathbf{v}}_j) \right), \quad (5.19)$$

where $S_{j,k}^{(\mathbf{v})} = \Phi_{\text{out}}/(s_k - \lambda_j)$ and $S_{j,k}^{(\bar{\mathbf{v}})} = \Phi_{\text{out}}/(s_k - \bar{\lambda}_j)$. This modification doubles the memory space used by the $S_{j,k}$ matrices, which is already the most memory intensive part of the algorithm. Therefore, if real modes are expected, defining $S_{j,k}$ as presented in Equation 5.8 saves time and memory.

Updating a CRI model If, after fitting a CRI model, it is decided that the tested spatial resolution does not provide sufficient accuracy, DEIM and its variants may again provide the answer. By using DEIM to select the next most important locations given the already measured input and output locations, the accuracy of the model can be increased with a minimum number of added locations.

If the spatial resolution is not the problem, it is also possible that new basis functions should be chosen or more basis functions could be added. For example, if incorrect boundary conditions were applied to the initial choice of basis functions, the accuracy may be greatly helped by using less restricted basis functions and rerunning the CRI algorithm.

5.4 Applications

In this section, three applications of CRI will be demonstrated with increasing levels of complexity. These applications range from a discrete simulated system to a real-world tested system with complex boundary conditions.

5.4.1 4-DOF system

The first application is a 4 degree-of-freedom (DOF) system. A simulated experiment is carried out where outputs are measured at masses 2-4 and inputs at masses 1-2. A diagram of the system as well as a representation of the transfer function matrix and its entries is shown in Figure 5.2.

A CRI model is fitted to the measured transfer function, $H(s) \in \mathbb{C}^{3 \times 2}$, over a frequency range which contains all of the system's poles. The poles are provided to the algorithm, along with the identity matrix, I_4 , as basis functions which can recreate all of the discrete modes. This choice of basis functions gives: $\Phi_{\text{out}} = [e_2 \ e_3 \ e_4]$ and $\Phi_{\text{in}} = [e_1 \ e_2]$ where e_j is a unit vector with only a nonzero entry in the j th row. The error at each iteration of the optimization is shown in Figure 5.3. After 7 iterations, the approximation to the measured transfer functions has converged, and as seen in Figure 5.4 they are matched to machine




Figure 5.2: Diagram of the 4-DOF system showing inputs $f_{1-4}(t)$ and outputs $y_{1-4}(t)$ and system properties. Inputs and outputs where there are virtual sensors for the simulated experiment are shown in blue. A schematic of the full transfer function $H_F(s) \in \mathbb{C}^{4 \times 4}$ is also shown, where the measured entries $H(s) \in \mathbb{C}^{3 \times 2}$ are shown in blue and entries which will be interpolated are shown in red.

precision. Figure 5.4 also shows that the interpolated transfer functions (the 10 transfer functions not originally measured) are also matched with the same accuracy. Although some of these transfer functions could be inferred via reciprocity, four of them would not have originally been available in a traditional rational approximation of the measured data. This shows the capability of the CRI method to expand beyond measured points.




Figure 5.3: Residual error after each iteration of the CRI optimization problem for the 4-DOF system in Application 1.




Figure 5.4: Frequency response functions of the 4-DOF system, the CRI model, and the error between the system and CRI. Both the six measured and fitted FRFs (left) and the ten interpolated FRFs (right) are shown.

5.4.2 Euler-Bernoulli Beam

The previous application gave a simulated example with a discrete amount of information. In this example, a CRI model will be fit to a model of a clamped-clamped Euler-Bernoulli with the PDE from Equation 4.1. A diagram of the beam is shown in Figure 5.5. This will demonstrate the ability of CRI to perform spatial transfer function interpolation, as well as highlight some limitations.




Figure 5.5: Diagram of the clamped-clamped Euler-Bernoulli beam used in Application 2. The system properties, stiffness, E , inertia, I , density, ρ , and area, A are uniform along the beam's length.

In this simulated example, the poles are once again known, and a frequency range which contains the first 5 natural frequencies is chosen. These 5 modes will be approximated with Chebyshev polynomials as basis functions with multiple polynomial orders to show the effect of the number of basis functions on the CRI model. Figure 5.6 shows errors for the measured FRFs and some interpolated FRFs with polynomial order 7. It can be seen that, especially for the interpolated FRFs, there is higher error at higher frequencies. This is a combination




Figure 5.6: Frequency response functions of the beam system, the CRI model, and the error between the system and CRI. The measured and fitted FRFs (left) and a selection of interpolated FRFs (right) are shown.

of the fact that higher order polynomials are needed for the higher frequency modes being fitted, and also error coming from higher modes which are not fitted.

In order to see a more full picture of the CRI model's accuracy, Figure 5.7 shows the error for all combinations of input and output locations. The error for each pair of input and output locations is defined as:

$$\|h_{i,j}(\omega) - \hat{h}_{i,j}(\omega)\|_2, \quad (5.20)$$

where $h_{i,j}(\omega) \in \mathbb{C}^{N_s}$ is a vector of the true FRF and $\hat{h}_{i,j}(\omega) \in \mathbb{C}^{N_s}$ is a vector of the interpolated FRF between input j and output i measured at linearly spaced frequencies over the entire frequency range shown in Figure 5.6. 200 linearly spaced spatial locations along the beam were considered as both input and output locations. In the first error plot it can be seen that with a low polynomial order, the FRFs are matched very well at the measured locations (where the vertical lines for the two input locations and the horizontal lines for the four output locations meet), and there is higher error away from measurement locations. However, with a higher polynomial order, there is almost no correlation with the measured locations. Instead, the error shows a pattern with six peaks in each spatial direction, corresponding to the lowest order mode not being fitted, the sixth mode, which has six peaks. In order to increase the polynomial order, the number of input and output locations was increased to keep the optimization problem stable.

As seen in Figure 5.7, the error for interpolation in CRI goes down with increasing polynomial order. Figure 5.8 compares error in mode shape approximation for CRI and two analytical methods, RRM and collocation. The error is defined as




Figure 5.7: Error plots showing error between the system FRFs and the CRI model FRFs for all combinations of input and output locations for two polynomial orders. Lighter colors are higher error. For each order, measured input locations are shown with vertical dotted lines, while measured output locations are shown with horizontal lines. FRFs where lines intersect are the measured and fitted FRFs, while the rest of the plot represent interpolated CRI FRFs.

$$\|\mathbf{u}_j - \tilde{\mathbf{u}}_j\|_2 \quad (5.21)$$

where $\mathbf{u}_j \in \mathbb{R}^{200}$ is the modal vector from the RRM with 30 admissible functions (at which time modes 1-4 had converged), and $\tilde{\mathbf{u}}_j \in \mathbb{R}^{200}$ is the modal vector approximation from each method. For each of these, the same Chebyshev polynomial order are used for a fair comparison with CRI. It can be seen that for low polynomial orders (with well chosen input and output locations), CRI is capable of fitting modes with the same amount of error as common analytical methods. At high polynomial order, there is a point when CRI can no longer fit the first or second modes much better. This is due to the fact that the higher order modes that are not being fitted still have some influence at lower frequencies, and add noise to the optimization. Table 5.1 shows the number of output locations, N , and number of input locations, M , used for the CRI fit at each polynomial order.

The first application demonstrated how the optimization problem at the heart of CRI can accurately learn mode shapes from measured data, and expand to combinations of input and output locations that were not measured. This second application went further, showing how CRI can perform spatial interpolation of continuous modes of a beam system. The accuracy of FRF interpolation in CRI depends on the number of measurement locations and the number of basis functions. With enough measurement locations, CRI can demonstrate equal accuracy in approximating mode shapes with analytical methods up to a point. Due to the limited frequency range of the measurements, influence from modes outside of the measured frequency range put a limit on the possible accuracy of a CRI model.




Figure 5.8: Comparison of errors for modes 1-4 for CRI, RRM, and collocation for three different polynomial orders.

Table 5.1: Polynomial order, number of outputs N and number of inputs M used in the comparisons in Figure 5.7 and Figure 5.8.

Polynomial Order	N	M
7	6	2
9	8	3
11	10	4

5.4.3 Marimba Bar

For the final application, FRFs will be measured experimentally from a marimba bar. A marimba bar was chosen so that CRI can be demonstrated over 2 spatial dimensions on a system with unknown material properties and complex boundary conditions. The rosewood material of the bar is not isotropic, and the tensioned ropes that run through the bar make applying appropriate boundary conditions difficult. It is important to note that the tensioned ropes are not meant to simulate free boundary conditions, but instead represent the realistic boundary conditions of how marimba bars are mounted on marimbas via tensioned ropes of the same size. Therefore the marimba bar provides a good example of some characteristics in real world systems where CRI would be helpful. The velocity of 30 locations on the surface of the bar was measured with a laser doppler vibrometer (LDV) system, while input forces were measured at 11 input locations with a small modal hammer. Five FRFs were averaged at each pair of input and output locations. A diagram of the experimental setup is given in Figure 5.9.

The spatial domain for the CRI model will be the top surface of the marimba bar. This will require 2D basis functions for the mode approximation. A total of 18 basis functions will




Figure 5.9: Diagram of the experimental setup for the marimba bar test. The bar is suspended on tensioned ropes that run through the bar. The velocity of the top face is measured with a PSV-400 LDV system. Input force is provided by a PCB impact hammer.

be used, which are all products of 1D Chebyshev polynomials. Along the short direction, a polynomial order of 2 is used so that flexural as well as torsional shapes can be fitted. Along the long direction, a polynomial order of 9 is used so that there are enough polynomials to approximate the flexural mode shapes. The basis functions are shown in Figure 5.10.




Figure 5.10: The 2D basis functions used for fitting the CRI model of the marimba bar. The basis functions are products of 1D basis functions for each dimension: $\phi_{ij}(x, y) = \phi_i(x)\phi_j(y)$.

In order to fit the CRI model, the poles must be provided. Unlike the first two applications, in this case the poles are not known. Therefore it is necessary to take the measured transfer functions to estimate the poles. In this work, this was done via rational approximation with vector fitting [32, 47, 48], but any number of methods to estimate the natural frequencies and damping ratios could be used. After the pole estimation, the CRI model can be fitted. 27 of the output locations and 9 of the input locations were used to fit the model, which reserves 3 output locations and 2 input locations for testing the accuracy of the interpolation with measured data.




Figure 5.11: Measured FRFs and the corresponding CRI model FRFs for one measured input. Each mode has an arrow above, corresponding to the mode groups in Figure 5.12.

Figure 5.11 shows the measured data and fit for one of the training input locations. In the measured frequency range, there are 11 natural frequencies. Each of the fitted CRI continuous mode shapes are shown in Figure 5.12. The modes are separated into three main categories. Boundary-driven modes that consist of mainly translation and rotation of the bar as a rigid body. Flexural modes with constant deflection along the short axis of the bar, and torsional modes with deflection along both the short axis and long axis of the bar. Interestingly, there are boundary-driven modes at a higher frequency than the first flexural mode. This demonstrates the complex boundary conditions enforced by the tensioned ropes which is automatically learned by the CRI fitting procedure.

Figure 5.13 shows a spatial visualization of the error for FRFs for two input locations, as well as a summary of all fitted and interpolated errors. The figure on the top right shows the error between input 7 (the same input from Figure 5.11) and the fitted and interpolated output locations. It can be seen that there is some spatial pattern to the error, with higher error across the short direction of the bar near where the ropes support the bar. It is possible that the boundary conditions cause extra dynamics in this location that the basis functions were




Figure 5.12: Visualization of the 11 continuous CRI modes. The modes are separated into three groups: boundary condition modes that are dominated by rotation and translation, flexural modes which are dominated by deflection along the y -direction, and torsional modes which also have deflection along the x -direction.

not able to capture. The error for the interpolated input, input 11, is slightly higher than the fitted input, input 7. However, the error summary on the left shows that the interpolated error is on the same order as the fitted error for all inputs and outputs. The bottom three rows are interpolated output locations, while the final two columns are interpolated input locations. Each point in these final rows and columns represent a FRF that would not be available with a classical rational approximation of the measured data. The CRI model is able to interpolate the FRFs at these locations while keeping a similar level of accuracy to the fitted data, and can create FRFs for any pair of input and output locations on the top surface of the bar.




Figure 5.13: Left: A matrix showing the error between all combinations of input and output locations. Right: The errors for all output locations (one column of the matrix on the right) mapped onto the surface of the marimba bar for input 7 (fitted) and input 11 (interpolated).

The goal of CRI is to create a continuous model, which can then be used to run predictive simulations. Therefore, simply looking at the FRF errors is not the end of the line for a true application of CRI. In the case of a marimba bar, it may be desirable to predict the sound produced by an impact at different locations on the bar in order to optimize the striking location for some sound quality criteria. Figure 5.14 shows estimated and true velocity responses for two pairs of impact locations. The estimated response is the velocity predicted by a simulation with the CRI model, while the true response is the response to the same input simulated using the measured transfer function. The top response is simulated with an input location which was fitted, but an output location which was not. The bottom response has both an interpolated input and interpolated output. In both cases, the estimated response is very similar to the true response. This shows that the CRI model is capable of creating a useful predictive model of a real-world system.




Figure 5.14: Comparison of velocity responses simulated with the true measured transfer function, and CRI interpolated transfer function for two cases. Top: fitted input and interpolated output. Bottom: both interpolated input and interpolated output. Each response is also shown with a zoomed view of the beginning of the response to the right.

5.5 CRI for VBOI

The applications in the previous section showed how CRI can be used to build an interpolated model of a system. This section will present a simulated example of how a CRI model can be applied to a vibration-based occupant inference algorithm. As discussed in Chapter 3, there are often significant simplifying assumptions made in many VBOI algorithms. CRI may provide an avenue to build sophisticated models of building systems of interest in VBOI without making assumptions about attenuation, dispersion, or reflections. In fact, as long as we can measure the FRFs, no further assumptions about the system are required for fitting a CRI model. The interpolated model can then be used to build a VBOI algorithm with a spatial resolution that is finer than the originally measured FRFs. In this section, we will build a CRI model of a simulated floor system, and use that model for localization.

5.5.1 Two-floor system

A simulated experiment on a simplified building model will be used to demonstrate the usefulness of CRI for VBOI algorithms. Three important characteristics of measured FRFs from Goodwin Hall in Chapter 3 will be emulated by the simplified model:

1. Modally dense at low frequencies
2. Energy loss through damping and transfer between floors
3. Complex mode shapes

In order to recreate these features, two coupled beams were designed into a simulated two floor building as shown in Figure 5.15. The system consists of two 100 element finite element beams with Euler-Bernoulli elements. The ground floor is more massive, and is on an elastic foundation of springs and dampers at all nodes. There are “columns” consisting of springs and dampers between the ground floor and the second floor. The coupling of these two beams gives rise to modally dense FRFs, matching the first characteristic of Goodwin FRFs. The elastic foundation and column dampers simulate the energy loss seen in Goodwin, as well as energy transfer between floors. Finally, the damping in the two-floor system was chosen specifically to be non-proportional: making the modes of the system complex, like those seen in Goodwin Hall.




Figure 5.15: Diagram of the two-floor system for which CRI based localization will be built for the second floor. The input and output locations from the simulated test are shown.

5.5.2 CRI fit

Since the two-floor system was designed with non-proportional damping, the mode shapes are complex. In Section 5.3.4 the extension of CRI to complex modes was discussed, and will be demonstrated here. A CRI fit assuming real modes was fit to the two-floor system using the FRFs, $H(\omega) \in \mathbb{C}^{5 \times 3}$, measured at the locations shown in Figure 5.15. A second CRI model was also created where optimization problem was modified to allow complex modes. The poles for all 10 modes with natural frequencies under 75 Hz were included in the CRI fit, and 6 Chebyshev polynomials were chosen as basis functions.




Figure 5.16: The 15 measured FRFs from the simulated test, and the fitted FRFs from the CRI model. Fits assuming real modes and complex modes are shown.

Figure 5.16 shows the fitted FRFs from both the real and complex fit. As expected for a system with complex modes, the CRI fit considering complex modes fits the measured data from the simulated experiment significantly better. The error from the FRFs between every combination of input and output locations on the second floor also showed that the complex fit was better over the whole beam, as seen in Figure 5.17. The pattern in the error in Figure 5.17 comes from the first mode outside of the measured frequency range (the 11th mode). This is also indicated in Figure 5.16—especially in the complex fit FRFs—where the error relative to the FRF magnitudes increases significantly at the very high end of the measured frequency range as the influence of the next mode increases.




Figure 5.17: Overall error for all possible FRFs: $H(\omega) \in \mathbb{C}^{101 \times 101}$. Fits assuming real modes and complex modes are shown.

5.5.3 CRI-FEEL Localization

Using the CRI interpolated FRFs, a VBOI localization algorithm will be built following the Force Estimation and Event Localization (FEEL) algorithm [30]. The original FEEL algorithm was demonstrated with measured FRFs between 4 input locations and 4 output locations on one floor of a building. While there was a high accuracy in localization using the measured FRFs, the spatial resolution was limited because only the measured input locations could be considered as candidate locations. However, by using interpolated FRFs from a CRI augmented model, the spatial resolution can be expanded to consider any input location. The combination of CRI and FEEL into CRI-FEEL localization will be demonstrated on the two-floor system.

To implement CRI-FEEL, first the true system FRFs $h_{i,j}(\omega)$ between input location j and output location i are used to simulate the response, $y_i(\omega)$, at each of the 5 output locations to an input force $f(t)$:

$$y_i(\omega) = \text{FFT}(f(t)) h_{i,j}(\omega). \quad (5.22)$$

The input force used was an impulse-like Gaussian pulse with a maximum force of 1000 N. The next step in the FEEL localization algorithm is to estimate the force from each of the measured responses at a number of candidate input locations. Here, the interpolated FRFs $\hat{h}_{i,j}(\omega)$ are used to estimate the force, $\hat{f}_{i,j}(t)$:

$$\hat{f}_{i,j}(t) = \text{IFFT} \left(\frac{y_i(\omega)}{\hat{h}_{i,j}(\omega)} \right). \quad (5.23)$$

If FRFs from the true location are used, there should be agreement between the estimated force from each measured locations. The maximum estimated force from each of the five measured responses are compared to see which candidate input location has the lowest variance in the estimated force. The maximum force, as defined by Davis et al. [30] is $\hat{f}_{i,j}^{\max} = \max(\hat{f}_{i,j}(t)) - \min(\hat{f}_{i,j}(t))$, which accounts for a bias in the estimated force. The coefficient of variation (CV) will be compared for each candidate input location:

$$\text{CV}_j = \frac{\sigma \left(\begin{matrix} \hat{f}_{1,j}^{\max} & \hat{f}_{2,j}^{\max} & \dots & \hat{f}_{5,j}^{\max} \end{matrix} \right)}{\mu \left(\begin{matrix} \hat{f}_{1,j}^{\max} & \hat{f}_{2,j}^{\max} & \dots & \hat{f}_{5,j}^{\max} \end{matrix} \right)}, \quad (5.24)$$

where $\sigma(\cdot)$ is the standard deviation and $\mu(\cdot)$ is the mean of the max force estimates. The focus of this section is not to investigate the accuracy of this method itself, but to highlight how CRI can be used to expand the spatial resolution of a current VBOI algorithm. In this same manner, CRI could be applied to other methods to augment locations which

lack measurements. Figure 5.18 shows an example of the CV for all 101 candidate input locations on the second floor, where the true input was at $x = 0$. The CV at the true input location is the lowest, and the estimated force is very close to the true force of 1000 N. Figure 5.18 also highlights the benefit of using the interpolated FRFs: all possible input locations can be considered, whereas without CRI only the three measured input locations could be considered.




Figure 5.18: Comparison of CRI-FEEL and FEEL localization for a simulated impact at $x = 0$.

Figure 5.19 shows the process of using CRI-FEEL localization for all possible input locations on the second floor. The true input location always has a low CV, although for some points a candidate location near the true input location has an even lower CV. While the localization is not perfectly accurate, all but one case is less than four nodes from the true input location: an accuracy within 4% of the length of the beam. Figure 5.19 also shows that the maximum force estimate from the estimated input location is generally close to the true maximum force of 1000 N.




Figure 5.19: CRI-FEEL localization for all possible input locations on the second floor.

5.6 Conclusion

We have presented a method called CRI for creating a continuous predictive model from a discrete amount of measured data. It works by setting up an optimization problem to minimize the error between a transfer function approximation and measured transfer function data. After this fitting of the model, the basis functions used to approximate the modes can be resampled at any location to give interpolated transfer functions. This method can have some limitations based on the chosen locations, basis functions, and frequency range. In particular there is a limit to the accuracy due to higher order modes not in the measured range. Nevertheless, through a series of applications it has been shown that CRI can be successfully applied to real world systems. The marimba example shows that CRI can learn

material properties and boundary conditions.

A simulated example of building a CRI model to develop a VBOI algorithm was also presented. A simulated experiment measuring 5 outputs and 3 inputs on the second floor of a simplified building model was used to fit a CRI model. The interpolated FRFs were then used in combination with the FEEL localization method [30] to create a VBOI localization algorithm with high accuracy and fine spatial resolution.

Future work for CRI could continue to look into modifications such as considering complex mode shapes, more sophisticated optimization schemes (because local minima may present more difficulties with more complex systems), and more rigorous experimental design rules for choosing locations and basis functions.

Chapter 6

Conclusion

We will conclude by summarizing the three contributions in this dissertation, and discuss future work for each.

6.1 Summary of Contributions

In Chapter 2 a technique for extracting temporal gait parameters via underfloor, structurally mounted accelerometers was demonstrated. It was shown that even with a sparse sensor infrastructure, accurate temporal parameters can be measured. Parameter distributions from the structural measurements match well with traditional gait analysis techniques of floor-mounted force plates and in-shoe pressure sensors.

In Chapter 3 three common assumptions made in VBOI algorithms were investigated. The dynamics in civil structures are complex, with the interaction of attenuation, reflections, and dispersion. Each assumption simplifies at least one of these phenomena, and while each assumption has some situations where they are viable, they each have some major drawbacks. Through a combination of experimental and theoretical investigations of each assumption, some recommendations were made for when each assumption is most accurate, and what to look out for if future researchers are thinking of using such assumptions in developing new VBOI algorithms.

In Chapter 5 a novel data-driven method for creating a predictive model of a dynamical system was presented called continuous residue interpolation (CRI). CRI works by choosing a set of basis functions to expand the residues of the FRF, and fitting a model to the data with a quadratic optimization problem. Once the model is fitted, the basis functions can be resampled to get an interpolated model between any input and output location. The CRI method was demonstrated on some simulated systems as well as experimental data of a marimba bar, demonstrating the ability to generate useful FRFs at locations not originally measured. Finally, a CRI model was also used to develop an accurate VBOI localization algorithm on a simulated system.

6.2 Areas of Future Work

The path taken by this dissertation in moving from developing a specific VBOI algorithm, to testing assumptions commonly made throughout the VBOI literature, to developing a data-driven modelling technique has left room for future work at each step. However, each step has provided results pointing to promising directions for the future work of the previous contribution.

6.2.1 Future of Vibration-based Gait Analysis

The contribution in Chapter 2 demonstrated gait analysis of healthy individuals and was restricted to temporal gait parameters. This leaves two clear paths forward to expand vibration-based gait analysis:

Extension beyond temporal parameters There has been research into extracting force based gait parameters [30, 39], which combined with temporal parameters would provide a more complete gait analysis system. In particular, the Force Estimation and Event Localization (FEEL) method presented by Davis et al. provides a good framework for extracting both GRF measurements as well as footstep location data [30]. A system capable of measuring temporal parameters, force parameters, and spatial parameters would provide some of the capabilities currently available in gait analysis only in controlled environments with expensive equipment such as 3D motion capture systems and in-floor force plates. However, FEEL requires measured FRFs between an input and output location to get accurate force and location estimates. In order to be able to use FEEL across an entire floor of a building would require an extremely large amount of testing. An interesting future area of research would be using CRI to generate a full model of the floor with much less testing, yet allowing FEEL to operate with a model of the entire floor without spatial resolution limitations.

Gait analysis of individuals with abnormal gait Another important area of future work is validating the ability of vibration-based gait analysis to measure indicators of abnormal gait. In a scenario where such gait analysis is used passively to monitor changes in gait, for example in a nursing home, the ability to detect gait asymmetries or other indicators of abnormal gait are important. In Chapter 2 it was shown that there is not a significant difference in temporal parameters for healthy individuals between their left and right legs, which is expected for normal gait. Similar symmetry checks can also run with force parameters [39]. If such a gait analysis system is to ever be deployed, it is important that there is robust validation of the system with the populations whose gait would be monitored for abnormalities.

6.2.2 Future of Vibration-based Occupant Inference

To date, the field of VBOI has been dominated by solitary works where an algorithm is proposed to infer some information like localization, classification, or gait analysis, and the new algorithm is validated on a data set (which is often measured just for that study) within one article. The human activity benchmark [64] challenges this trend, because it invites multiple algorithms to compete on the same data. However, at least in the opinion of the author, the VBOI field is still too focused on algorithm development with little fundamental research or consideration for the complex structural dynamics of the civil structures where VBOI algorithms operate.

It is the context of this state of the field that led to the contributions in Chapter 3. Hopefully future VBOI research can build on a better fundamental understanding of the dynamics of civil structures, and fewer assumptions can be made, allowing algorithms to be more accurate and more generalizable to other structures. Data-driven methods may be particularly useful as they can learn more about the system than simple calibration procedures. As mentioned, CRI may be combined with model-based localization methods like FEEL to create a localization algorithm with a relatively small amount of testing. Hammer impacts could also be used to learn the dispersion curve within a floor, and that information could be used in an improved TDoA localization algorithm. Dispersion based localization algorithms like in [117] could also improve localization accuracy. A final idea for future VBOI work comes from the fact that each assumption from Chapter 3 is accurate some of the time, and sometimes the assumptions are at odds. For instance, Assumption 1 is more accurate at low frequencies while Assumption 2 is more accurate at high frequencies. Creating an ensemble model of multiple simplified models which make assumptions may be another way to create more robust VBOI algorithms.

6.2.3 Future of Continuous Residue Interpolation

Chapter 5 introduced CRI, and in that Chapter some implementation considerations were discussed. One area for future work on CRI is in investigating best practices for implementing CRI:

Basis Functions From the investigation in Chapter 4 on the RRM, and through preliminary experiments with CRI, orthogonal polynomial sets such as Chebyshev polynomials are a natural choice for basis functions in CRI. Polynomials converge quickly to the mode shape, and choosing orthogonal polynomials gives good conditioning to the interpolation matrices, making the optimization easier. Future work should look at the utility of applying boundary conditions to the basis functions, for which there are options as shown in Chapter 4. Another possibility worth exploring is the use of the modes of a FE model as basis functions instead of continuous functions, similarly to the SEREP method [78].

Input Locations Chebyshev nodes are an obvious choice for the output locations due to the connection with polynomial interpolation [12], however there is not such an obvious connection for choosing input locations to minimize CRI error. DEIM was used in Chapter 5 and performed well. However, other methods for choosing input locations like observability methods could also be tested. Finally, while driving point FRFs are not required for CRI, specific tests could be run to see if they help create more accurate models.

Frequency Resolution The connection between frequency resolution and CRI model accuracy should also be investigated. Some methods like impact testing or chirp excitation require a uniform frequency spacing, but methods like sine testing could be used to target specific frequencies—like concentrating measurements near natural frequencies— and decrease the size of data required to fit a good model.

Model Convergence Finally, there is the question of how to know if a CRI model has converged, or if more testing or basis functions are needed. Here it is possible that cross validation methods could be used, where some data is held back as testing data similarly to the marimba bar test at the end of Chapter 5

Along with the previously mentioned considerations that are worth exploring, there are also extensions of the CRI algorithm that are areas for future work:

Complex Modes While briefly discussed in Chapter 5, the extension of the optimization problem to consider systems with complex modes—like civil structures exhibit as shown in Chapter 3—should be further researched.

Parametric models Extending CRI to be able to create parametric models is also an area for future work. Chapter 3 also showed that civil structures are sensitive to some parameters, like damping ratios that depend on excitation level.

Finally, future work for CRI will be using interpolated models in real world scenarios. For example building a CRI model of the floor in Goodwin Hall to create a localization algorithm, or using CRI and IMMAT to optimize an environmental testing procedure.

Bibliography

- [1] Sa'ed Alajlouni and Pablo Tarazaga. A new fast and calibration-free method for footstep impact localization in an instrumented floor. *Journal of Vibration and Control*, February 2019.
- [2] Sa'ed Alajlouni, Mohammad Albakri, and Pablo Tarazaga. Impact localization in dispersive waveguides based on energy-attenuation of waves with the traveled distance. *Mechanical Systems and Signal Processing*, 105:361–376, May 2018.
- [3] Berkan Alanbay and Rakesh K. Kapania. On the Use of Classical Jacobi Orthogonal Polynomials in the Ritz Method. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Kissimmee, Florida, January 2018. American Institute of Aeronautics and Astronautics.
- [4] R.J. Allemang and D.L. Brown. A Unified Matrix Polynomial Approach to Modal Identification. *Journal of Sound and Vibration*, 211(3):301–322, April 1998.
- [5] A. C. Antoulas, D. C. Sorensen, and S. Gugercin. A survey of model reduction methods for large-scale systems. *Contemporary Mathematics*, 280:193–219, 2001.
- [6] Peter Avitabile and Pawan Pingle. Prediction of Full Field Dynamic Strain from Limited Sets of Measured Data. *Shock and Vibration*, 19(5):765–785, 2012.
- [7] R. Bahroun, O. Michel, F. Frassati, M. Carmona, and J. L. Lacoume. New algorithm for footstep localization using seismic sensors in an indoor environment. *Journal of Sound and Vibration*, 333(3):1046–1066, February 2014.
- [8] Richard Baker. Gait analysis methods in rehabilitation. *Journal of neuroengineering and rehabilitation*, 3(1):4, 2006.
- [9] D. Bales, P. A. Tarazaga, M. Kasarda, D. Batra, A. G. Woolard, J. D. Poston, and V. V. N. S. Malladi. Gender Classification of Walkers via Underfloor Accelerometer Measurements. *IEEE Internet of Things Journal*, 3(6):1259–1266, December 2016.
- [10] Peter Benner, Serkan Gugercin, and Karen Willcox. A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems. *SIAM Review*, 57: 483–531, June 2015.
- [11] Mario Berljafa and Stefan Güttel. The RKFIT Algorithm for Nonlinear Rational Approximation. *SIAM Journal on Scientific Computing*, 39(5):A2049–A2071, January 2017.

- [12] Jean-Paul Berrut and Lloyd N. Trefethen. Barycentric Lagrange Interpolation. *SIAM Review*, 46(3):501–517, January 2004.
- [13] Rama B Bhat. Vibration of beams using novel boundary characteristic orthogonal polynomials satisfying all boundary conditions. *Advances in Mechanical Engineering*, 7(4):1687814015578355, April 2015.
- [14] Richard Evelyn Donohue Bishop, R. E. D. Bishop, and D. C. Johnson. *The Mechanics of Vibration*. Cambridge University Press, June 1979.
- [15] M. Bocian, J. M. W. Brownjohn, V. Racic, D. Hester, A. Quattrone, and R. Monnickendam. A framework for experimental determination of localised vertical pedestrian forces on full-scale structures using wireless attitude and heading reference systems. *Journal of Sound and Vibration*, 376:217–243, August 2016.
- [16] M. Bocian, J. M. W. Brownjohn, V. Racic, D. Hester, A. Quattrone, L. Gilbert, and R. Beasley. Time-dependent spectral analysis of interactions within groups of walking pedestrians and vertical structural motion using wavelets. *Mechanical Systems and Signal Processing*, 105:502 – 523, 2018.
- [17] R. E. Brown and M. A. Stone. On the use of polynomial series with the Rayleigh-Ritz method. *Composite Structures*, 39(3):191–196, November 1997.
- [18] Andrea Carracedo Rodriguez and Serkan Gugercin. The p-AAA algorithm for data driven modeling of parametric dynamical systems. *arXiv e-prints*, 2003: arXiv:2003.06536, March 2020.
- [19] Ylenia Cascone, Maria Ferrara, Luigi Giovannini, and Gianluca Serale. Ethical issues of monitoring sensor networks for energy efficiency in smart buildings: a case study. *Energy Procedia*, 134:337–345, October 2017.
- [20] S. Chakraverty. *Vibration of Plates*. CRC Press, December 2008.
- [21] Charalambos P Charalambous. Walking patterns of normal men. In *Classic Papers in Orthopaedics*, pages 393–395. Springer, 2014.
- [22] Saifon Chaturantabut and Danny C. Sorensen. Nonlinear Model Reduction via Discrete Empirical Interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, January 2010.
- [23] Shanshan Chen, John Lach, Benny Lo, and Guang-Zhong Yang. Toward pervasive gait analysis with wearable sensors: A systematic review. *IEEE journal of biomedical and health informatics*, 20(6):1521–1537, 2016.
- [24] Y. Chen, P. Logan, P. Avitabile, and J. Dodson. Non-Model Based Expansion from Limited Points to an Augmented Set of Points Using Chebyshev Polynomials. *Experimental Techniques*, 43(5):521–543, October 2019.

- [25] Ying-Ju Chen, Fu-Chou Cheng, Meei-Ling Sheu, Hong-Lin Su, Chun-Jung Chen, Jason Sheehan, and Hung-Chuan Pan. Detection of subtle neurological alterations by the Catwalk XT gait analysis system. *Journal of neuroengineering and rehabilitation*, 11(1):62, 2014.
- [26] Veronica Cimolin and Manuela Galli. Summary measures for clinical gait analysis: a literature review. *Gait & posture*, 39(4):1005–1010, 2014.
- [27] P. M. Daborn, C. Roberts, D. J. Ewins, and P. R. Ind. Next-Generation Random Vibration Tests. In Randall Allemang, editor, *Topics in Modal Analysis II, Volume 8*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 397–410, 2014.
- [28] Benjamin Davis and Juan M. Caicedo. Impact Force Estimation And Event Localization, October 2017.
- [29] Benjamin T. Davis, Juan M. Caicedo, Scott Langevin, and Victor Hirth. Use of Wireless Smart Sensors for Detecting Human Falls through Structural Vibrations. In Tom Proulx, editor, *Civil Engineering Topics, Volume 4*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 383–389. Springer, 2011.
- [30] Benjamin T. Davis, Juan M. Caicedo, and Victor A. Hirth. Force Estimation and Event Localization (FEEL) of Impacts Using Structural Vibrations. *Journal of Engineering Mechanics*, 147(3):04020154, March 2021.
- [31] Joel A DeLisa. *Gait analysis in the science of rehabilitation*. Diane Publishing, 1998.
- [32] Dirk Deschrijver, Michal Mrozowski, Tom Dhaene, Senior Member, and Daniel De Zutter. Macromodeling of Multiport Systems Using a Fast Implementation of the Vector Fitting Method. *IEEE Microwave and Wireless Components Letters*, pages 383–385, 2008.
- [33] Bing Dong, Vishnu Prakash, Fan Feng, and Zheng O’Neill. A review of smart building sensing system for better indoor environment control. *Energy and Buildings*, 199: 29–46, September 2019.
- [34] Slah Drira, Yves Reuland, Sai G. S. Pai, Hae Young Noh, and Ian F. C. Smith. Model-Based Occupant Tracking Using Slab-Vibration Measurements. *Frontiers in Built Environment*, 5, 2019.
- [35] Zlatko Drmac and Serkan Gugercin. A New Selection Operator for the Discrete Empirical Interpolation Method—Improved A Priori Error Bound and Extensions. *SIAM Journal on Scientific Computing*, 38, May 2015.
- [36] Z. Drmač, S. Gugercin, and C. Beattie. Vector Fitting for Matrix-valued Rational Approximation. *SIAM Journal on Scientific Computing*, 37(5):A2346–A2379, January 2015.

- [37] Alexander Ekimov and James M. Sabatier. Vibration and sound signatures of human footsteps in buildings. *The Journal of the Acoustical Society of America*, 120(2):762–768, August 2006.
- [38] Alexander Ekimov and James M Sabatier. Human motion characterization. Technical report, MISSISSIPPI UNIV UNIVERSITY NATIONAL CENTER FOR PHYSICAL ACOUSTICS, 2009.
- [39] Jonathon Fagert, Mostafa Mirshekari, Shijia Pan, Pei Zhang, and Hae Young Noh. Characterizing Structural Changes to Estimate Walking Gait Balance. In Shamim Pakzad, editor, *Dynamics of Civil Structures, Volume 2*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 333–335, 2019.
- [40] Frank Fahy and Paolo Gardonio. *Sound and Structural Vibration*. Academic Press, second edition edition, December 2006.
- [41] Wang Fengquan and Chen Shiyu. A Method to Determine the Boundary Conditions of the Finite Element Model of a Slender Beam Using Measured Modal Parameters. *Journal of Vibration and Acoustics*, 118(3):474–478, July 1996.
- [42] James R Gage. Gait analysis. An essential tool in the treatment of cerebral palsy. *Clinical orthopaedics and related research*, pages 126–134, March 1993.
- [43] Ion Victor Gosea and Stefan Güttel. Algorithms for the rational approximation of matrix-valued functions. *arXiv:2003.06410 [cs, math]*, March 2020.
- [44] Karl F. Graff. *Wave motion in elastic solids*. Ohio State University Press, June 1975.
- [45] D. W. Grieve and Ruth J. Gear. The Relationships Between Length of Stride, Step Frequency, Time of Swing and Speed of Walking for Children and Adults. *Ergonomics*, 9(5):379–399, September 1966.
- [46] S. Gugercin, A. C. Antoulas, and C. Beattie. H2 Model Reduction for Large-Scale Linear Dynamical Systems. *SIAM Journal on Matrix Analysis and Applications*, 30(2):609–638, January 2008.
- [47] B. Gustavsen. Improving the pole relocating properties of vector fitting. *IEEE Transactions on Power Delivery*, 21(3):1587–1592, July 2006.
- [48] B. Gustavsen and A. Semlyen. Rational approximation of frequency domain responses by vector fitting. *IEEE Transactions on Power Delivery*, 14(3):1052–1061, July 1999.
- [49] J. M. Hausdorff, H. K. Edelberg, S. L. Mitchell, A. L. Goldberger, and J. Y. Wei. Increased gait unsteadiness in community-dwelling elderly fallers. *Archives of Physical Medicine and Rehabilitation*, 78(3):278–283, March 1997.

- [50] J. M. Hausdorff, D. A. Rios, and H. K. Edelberg. Gait variability and fall risk in community-living older adults: a 1-year prospective study. *Archives of Physical Medicine and Rehabilitation*, 82(8):1050–1056, August 2001.
- [51] Diana Hodgins. The importance of measuring human gait. *Medical Device Technology*, 19(5):42–44, 2008.
- [52] Nelson N. Hsu. Acoustical birefringence and the use of ultrasonic waves for experimental stress analysis. *Experimental Mechanics*, 14(5):169–176, May 1974.
- [53] Sinniah Ilanko, Luis E. Monterrubio, and Yusuke Mochida. *The Rayleigh-Ritz Method for Structural Analysis*. ISTE Ltd and John Wiley & Sons, Inc., November 2014.
- [54] M. Imregun and D. J. Ewins. Complex Modes - Origins and Limits. In *Proceedings of the 13th International Modal Analysis Conference*, page 496, 1995.
- [55] D. J. Inman. Theoretical Models for Modal Analysis. In Júlio M. M. Silva and Nuno M. M. Maia, editors, *Modal Analysis and Testing*, NATO Science Series, pages 227–240. Springer Netherlands, Dordrecht, 1999.
- [56] Daniel J. Inman. *Vibration with Control*. John Wiley & Sons, November 2006.
- [57] EC Jansen, HH Thyssen, and J Brynskov. Gait analysis after intake of increasing amounts of alcohol. *Zeitschrift für Rechtsmedizin*, 94(2):103–107, 1985.
- [58] Subrat Kumar Jena, Snehashish Chakraverty, and Francesco Tornabene. Buckling Behavior of Nanobeams Placed in Electromagnetic Field Using Shifted Chebyshev Polynomials-Based Rayleigh-Ritz Method. *Nanomaterials*, 9(9):1326, September 2019.
- [59] M. Kasarda, Pablo A. Tarazaga, M. Embree, S. Gugercin, A. G. Wooldard, B. A. Joyce, and J. M. Hamilton. Detection and Identification of Firearms Upon Discharge Using Floor-Based Accelerometers. 2016.
- [60] C. T. Kelley. *Iterative Methods for Optimization*. Society for Industrial and Applied Mathematics, January 1999.
- [61] Ellis Kessler, Vijaya VN Sriram Malladi, and Pablo A. Tarazaga. Vibration-based gait analysis via instrumented buildings. *International Journal of Distributed Sensor Networks*, 15(10), 2019.
- [62] Ellis Kessler, Vijaya Malladi, Rodrigo Sarlo, Luke Martin, and Pablo Tarazaga. Comparison of Modal Parameters of a Concrete Slab Floor from EMA and OMA. In *Dynamics of Civil Structures, Volume 2*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 237–240, January 2021.

- [63] Sanda Lefteriu, Athanasios C. Antoulas, and Antonio C. Ionita. Parametric Model Reduction in the Loewner framework. *IFAC Proceedings Volumes*, 44(1):12751–12756, January 2011.
- [64] Ramin Madarshahian, Juan M. Caicedo, and Diego Arocha Zambrana. Benchmark problem for human activity identification using floor vibrations. *Expert Systems with Applications*, 62:263–272, November 2016.
- [65] Krithika Manohar, Bingni W. Brunton, J. Nathan Kutz, and Steven L. Brunton. Data-Driven Sparse Sensor Placement for Reconstruction. *IEEE Control Systems*, 38(3):63–86, June 2018.
- [66] R. Mayes, L. Ankers, and P. Daborn. Predicting System Response at Unmeasured Locations. *Experimental Techniques*, 44(4):457–474, August 2020.
- [67] A. J. Mayo and A. C. Antoulas. A framework for the solution of the generalized realization problem. *Linear Algebra and its Applications*, 425(2):634–662, September 2007.
- [68] Mauro S. Maza, Mohammad I. Albakri, V. V. N. S. Malladi, and Pablo A. Tarazaga. On Wave Propagation in Smart Buildings. In Shamim Pakzad, editor, *Dynamics of Civil Structures, Volume 2*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 237–242, 2020.
- [69] Leonard Meirovitch. *Fundamentals of Vibrations*. McGraw-Hill, 2001.
- [70] Leonard Meirovitch and Moon K. Kwak. Convergence of the classical Rayleigh-Ritz method and the finite element method. *AIAA Journal*, 28(8):1509–1516, August 1990.
- [71] Mostafa Mirshekari, Shijia Pan, Jonathon Fagert, Eve M. Schooler, Pei Zhang, and Hae Young Noh. Occupant localization using footstep-induced structural vibration. *Mechanical Systems and Signal Processing*, 112:77–97, November 2018.
- [72] Alvaro Muro-De-La-Herran, Begonya Garcia-Zapirain, and Amaia Mendez-Zorrilla. Gait analysis methods: An overview of wearable and non-wearable systems, highlighting clinical applications. *Sensors*, 14(2):3362–3394, 2014.
- [73] Akinori Nagano, Masahiro Fujimoto, Shoma Kudo, and Ryosuke Akaguma. An image-processing based technique to obtain instantaneous horizontal walking and running speed. *Gait & posture*, 51:7–9, 2017.
- [74] Yuji Nakatsukasa, Olivier Sète, and Lloyd N. Trefethen. The AAA Algorithm for Rational Approximation. *SIAM Journal on Scientific Computing*, 40(3):A1494–A1522, January 2018.
- [75] Tuan Anh Nguyen and Marco Aiello. Energy intelligent buildings based on user activity: A survey. *Energy and Buildings*, 56:244–257, January 2013.

- [76] J. Nilsson and A. Thorstensson. Ground reaction forces at different speeds of human walking and running. *Acta Physiologica Scandinavica*, 136(2):217–227, June 1989.
- [77] Jeretta Horn Nord, Alex Koohang, and Joanna Paliszkiewicz. The Internet of Things: Review and theoretical framework. *Expert Systems with Applications*, 133:97–108, November 2019.
- [78] John O’Callahan, Peter Avitable, and Robert Reimer. System Equivalent Reduction Expansion Process (SEREP). *Proceedings of the 7th International Modal Analysis Conference*, pages 29–37, January 1989.
- [79] Shijia Pan, Amelie Bonde, Jie Jing, Lin Zhang, Pei Zhang, and Hae Young Noh. BOES : Building Occupancy Estimation System Using Sparse Ambient Vibration Monitoring. In *SPIE Smart Structures and Materials + Nondestructive Evaluation and Health Monitoring*, 2014.
- [80] Shijia Pan, Tong Yu, Mostafa Mirshekari, Jonathon Fagert, Amelie Bonde, Ole J. Mengshoel, Hae Young Noh, and Pei Zhang. FootprintID: Indoor Pedestrian Identification Through Ambient Structural Vibration Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 1(3):89:1–89:31, September 2017.
- [81] Shijia Pan, Mostafa Mirshekari, Jonathon Fagert, Ceferino Gabriel Ramirez, Albert Jin Chung, Chih Chi Hu, John Paul Shen, Pei Zhang, and Hae Young Noh. Characterizing human activity induced impulse and slip-pulse excitations through structural vibration. *Journal of Sound and Vibration*, 414:61–80, February 2018.
- [82] Varun Panwar, Prashant Gupta, Ashok K Bagha, and Nitin Chauhan. A Review on studies of Finite Element Model Updating and Updating of Composite Materials. *Materials Today: Proceedings*, 5(14, Part 2):27912–27918, January 2018.
- [83] Miroslav Pastor, Michal Binda, and Tomáš Harčarik. Modal Assurance Criterion. *Procedia Engineering*, 48:543–548, January 2012.
- [84] Bart Peeters and Guido De Roeck. Reference-based stochastic subspace identification for output-only modal analysis. *Mechanical Systems and Signal Processing*, 13(6): 855–878, November 1999.
- [85] Bart Peeters, Herman Van der Auweraer, Patrick Guillaume, and Jan Leuridan. The PolyMAX Frequency-Domain Method: A New Standard for Modal Parameter Estimation? *Shock and Vibration*, 11:395–409, January 2004.
- [86] Alexandra Pfister, Alexandre M West, Shaw Bronner, and Jack Adam Noah. Comparative abilities of Microsoft Kinect and Vicon 3D motion capture for gait analysis. *Journal of medical engineering & technology*, 38(5):274–280, 2014.

- [87] Lukasz Piwek, David A. Ellis, Sally Andrews, and Adam Joinson. The Rise of Consumer Health Wearables: Promises and Barriers. *PLOS Medicine*, 13(2):e1001953, February 2016.
- [88] J. D. Poston, J. Schloemann, R. M. Buehrer, V. V. N. S. Malladi, A. G. Woolard, and P. A. Tarazaga. Towards indoor localization of pedestrians via smart building vibration sensing. In *2015 International Conference on Location and GNSS (ICL-GNSS)*, pages 1–6, June 2015.
- [89] Jeffrey D. Poston, R. Michael Buehrer, and Pablo A. Tarazaga. A Framework for Occupancy Tracking in a Building via Structural Dynamics Sensing of Footstep Vibrations. *Frontiers in Built Environment*, 3, 2017.
- [90] Jeffrey D. Poston, R. Michael Buehrer, and Pablo A. Tarazaga. Indoor footstep localization from structural dynamics instrumentation. *Mechanical Systems and Signal Processing*, 88:224–239, May 2017.
- [91] V. Racic, A. Pavic, and J. M. W. Brownjohn. Experimental identification and analytical modelling of human walking forces: Literature review. *Journal of Sound and Vibration*, 326(1):1–49, September 2009.
- [92] A. K. Ramanathan, P. Kiran, G. P. Arnold, W. Wang, and R. J. Abboud. Repeatability of the Pedar-X® in-shoe pressure measuring system. *Foot and Ankle Surgery*, 16(2):70–73, June 2010.
- [93] Yves Reuland, Sai G. S. Pai, Slah Drira, and Ian F. C. Smith. Vibration-Based Occupant Detection Using a Multiple-Model Approach. In Juan Caicedo and Shamim Pakzad, editors, *Dynamics of Civil Structures, Volume 2*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 49–56, 2017.
- [94] Edwin Reynders, Jeroen Houbrechts, and Guido De Roeck. Fully automated (operational) modal analysis. *Mechanical Systems and Signal Processing*, 29:228–250, May 2012.
- [95] James M Sabatier and Alexander E Ekimov. Ultrasonic methods for human motion detection. Technical report, MISSISSIPPI UNIV UNIVERSITY NATIONAL CENTER FOR PHYSICAL ACOUSTICS, 2006.
- [96] Michael Saleh and George Murdoch. In defence of gait analysis. Observation and measurement in gait assessment. *Bone & Joint Journal*, 67(2):237–241, 1985.
- [97] Rodrigo Sarlo and Pablo Tarazaga. Modal Parameter Uncertainty Estimates as a Tool for Automated Operational Modal Analysis: Applications to a Smart Building. pages 177–182. January 2019.

- [98] Rodrigo Sarlo, Pablo A. Tarazaga, and Mary E. Kasarda. High resolution operational modal analysis on a five-story smart building under wind and human induced excitation. *Engineering Structures*, 176:279 – 292, 2018.
- [99] Zimi Sawacha, Guarneri Gabriella, Giuseppe Cristoferi, Annamaria Guiotto, Angelo Avogaro, and Claudio Cobelli. Diabetic gait and posture abnormalities: a biomechanical investigation through three dimensional gait analysis. *Clinical biomechanics*, 24(9):722–728, 2009.
- [100] J.T. Sawicki and G. Genta. MODAL UNCOUPLING OF DAMPED GYROSCOPIC SYSTEMS. *Journal of Sound and Vibration*, 244(3):431–451, July 2001.
- [101] Javier Schloemann, VVN Sriram Malladi, Americo G Woolard, Joseph M Hamilton, R Michael Buehrer, and Pablo A Tarazaga. Vibration event localization in an instrumented building. In *Experimental Techniques, Rotating Machinery, and Acoustics, Volume 8*, pages 265–271. Springer, 2015.
- [102] Daniel Schmitt, Alexander Vap, and Robin M Queen. Effect of end-stage hip, knee, and ankle osteoarthritis on walking mechanics. *Gait & posture*, 42(3):373–379, 2015.
- [103] Brian Schwarz and Mark Richardson. Scaling Mode Shapes Obtained from Operating Data. *SOUND AND VIBRATION*, page 5, 2003.
- [104] Tim Sharpe. Ethical issues in domestic building performance evaluation studies. *Building Research & Information*, 47(3):318–329, April 2019.
- [105] Rong C. Shieh. Effects of Damping and Circulatory Forces on Dynamic Instability of Gyroscopic Conservative Continuous Systems. *Journal of Structural Mechanics*, 11(2): 197–213, January 1983.
- [106] Sheldon R. Simon, Igor L. Paul, Joseph Mansour, Michael Munro, Peter J. Abernethy, and Eric L. Radin. Peak dynamic force in human gait. *Journal of Biomechanics*, 14(12):817–822, January 1981.
- [107] Andy Smith, S. J. Hicks, and P. J. Devine. Design of Floors for Vibration: A New Approach. *Steel Construction Institute*, February 2009.
- [108] Andrey Soloviev and Jeff Dickman. Extending GPS carrier phase availability indoors with a deeply integrated receiver architecture. *IEEE Wireless Communications*, 18(2): 36–44, 2011.
- [109] Spence Seymour M. J. and Kareem Ahsan. Tall Buildings and Damping: A Concept-Based Data-Driven Model. *Journal of Structural Engineering*, 140(5), May 2014.
- [110] Erik Stone, Marjorie Skubic, Marilyn Rantz, Carmen Abbott, and Steve Miller. Average in-home gait speed: Investigation of a new metric for mobility and fall risk assessment of elders. *Gait & posture*, 41(1):57–62, 2015.

- [111] Joel Storch and Gilbert Strang. Paradox lost: Natural boundary conditions in the Ritz–Galerkin method. *International Journal for Numerical Methods in Engineering*, 26(10):2255–2266, 1988.
- [112] Lloyd N. Trefethen. *Spectral Methods in MATLAB*. Society for Industrial and Applied Mathematics, 2000.
- [113] Igor A Viktorov. *Rayleigh and Lamb Waves: Physical Theory and Applications*. Transl. from Russian. Plenum press, 1967.
- [114] Michael W Whittle. Clinical gait analysis: A review. *Human Movement Science*, 15(3):369–387, 1996.
- [115] Michael W Whittle. Generation and attenuation of transient impulsive forces beneath the foot: a review. *Gait & Posture*, 10(3):264–275, December 1999.
- [116] David A Winter. *Biomechanics and motor control of human gait: normal, elderly and pathological*. University of Waterloo Press, 1991.
- [117] Americo G Woolard and Pablo A Tarazaga. Applications of dispersion compensation for indoor vibration event localization. *Journal of Vibration and Control*, 24(21):5108–5117, 2018.
- [118] Americo G Woolard, Sa’ed Alajlouni, VVN Sriram Malladi, and Pablo A Tarazaga. Classification of event location using matched filters via on-floor accelerometers. In *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2017*, volume 10168. International Society for Optics and Photonics, 2017.
- [119] Tishya AL Wren, Susan A Rethlefsen, Bitte S Healy, K Patrick Do, Sandra W Dennis, and Robert M Kay. Reliability and validity of visual assessments of gait using a modified physician rating scale for crouch and foot contact. *Journal of Pediatric Orthopaedics*, 25(5):646–650, 2005.
- [120] Tishya AL Wren, George E Gorton, Sylvia Ounpuu, and Carole A Tucker. Efficacy of clinical gait analysis: A systematic review. *Gait & posture*, 34(2):149–153, 2011.

Appendices

Appendix A

Admissible Function Details

This appendix gives detailed explanations for all nine sets of admissible functions used in Chapter 4.

Transcendental Comparison Functions (CF) match all geometric and natural boundary conditions, and are trigonometric functions possibly combined with hyperbolic functions (for example in a beam, hyperbolic functions are sometimes needed) [70]. An example of comparison functions for the bar example would be sinusoidal functions:

$$\phi_j(x) = \sin(\beta_j x) \quad \text{for } j = 1, 2, \dots, \quad (\text{A.1})$$

where the β_j term must be solved such that the natural boundary condition at the free end is met. However, in some cases solving a transcendental system to ensure all boundary conditions are met may become time consuming or infeasible.

Trigonometric Admissible Functions (AF) seek to overcome the necessity to solve a transcendental equation to generate comparison functions by only matching geometric boundary conditions [70]. This allows a simpler set of admissible functions using only trigonometric functions (provided they are still complete in energy [111]). An example of trigonometric admissible functions are sinusoids of increasing frequency, without regard for the natural boundary condition of the bar:

$$\phi_j(x) = \sin\left(\left(j\pi - \frac{\pi}{2}\right)x\right) \quad \text{for } j = 1, 2, \dots \quad (\text{A.2})$$

Meirovitch observed that for the bar, this set of admissible functions has poor convergence, possibly due to the inability to match the natural boundary condition at the free end with a finite number of functions [70]. This is because all of the functions have a non-zero value at the end, but a zero derivative, making satisfying the equation $w'(L) + w(L) = 0$ impossible with a finite number of these admissible functions.

Quasi-Comparison Functions (QCF) were introduced by Meirovitch to solve the problem of needing an infinite number of admissible functions to match natural boundary con-

ditions [70]. In order to be able to satisfy the natural boundary condition, a second set of sinusoids are introduced, giving the set of functions:

$$\phi_j(x) = \sin\left(\frac{j\pi}{2}x\right) \quad \text{for } j = 1, 2, \dots \quad (\text{A.3})$$

Now, there are functions with a non-zero derivative at the free end, allowing a finite number of functions to match the natural boundary condition. The term quasi-comparison function was introduced to represent this bridge between admissible functions and comparison functions [70]. While introducing this second set of functions was shown to increase convergence speed, we will show that it also introduced a conditioning problem which was not present with the well-scaled admissible function set.

Trig+Poly QCF (QCF-poly) augment a trigonometric set of admissible functions with low order polynomials. It has been originally stated that the aim is to combine fast convergence of polynomials with the well conditioned property of trig functions [53]. Although they have not been explicitly called quasi-comparison functions, they are capable of matching all boundary conditions with a finite number of functions, and therefore meet the same goal as Meirovitch's quasi-comparison functions. For example, in our bar system, we may augment the trigonometric admissible function set with the monomial x :

$$\begin{aligned} \phi_1(x) &= x \\ \phi_{j+1}(x) &= \sin\left(\left(j\pi - \frac{\pi}{2}\right)x\right) \quad \text{for } j = 1, 2, \dots \end{aligned} \quad (\text{A.4})$$

As will be shown, the convergence of these functions is greatly improved over the trigonometric admissible functions, although not as much as the Meirovitch QCFs. However, these QCF-poly functions do not introduce the same level of conditioning problems that the QCFs do in the bar example.

The remaining sets of functions are various sets of polynomials which demonstrate the fast convergence of polynomials as well as methods for making them more well conditioned.

Monomials are included as a reference to compare with other polynomial sets. For the bar example, monomial admissible functions can be made from all monomials of order 1 and above:

$$\phi_j(x) = x^j \quad \text{for } j = 1, 2, \dots \quad (\text{A.5})$$

As has been well documented, monomials demonstrate fast convergence and poor conditioning [53].

Chebyshev Polynomials with Forced Boundary Conditions (Forced Cheby) is one set of polynomials applied to the RRM to solve the conditioning problem of monomials [20, 58]. This set starts with Chebyshev polynomials T_j (or shifted Chebyshev polynomials T_j^* defined on the interval from 0 to 1). These polynomials are well scaled on their own, and can be weighted to meet the geometric boundary conditions and become usable admissible functions for the RRM. The constants p and q define the weighting function in the admissible functions:

$$\phi_j(x) = x^p(1-x)^q T_{j-1}^*(x) \quad \text{for } j = 1, 2, \dots \quad (\text{A.6})$$

For the beam example, the constants $p = 1$ and $q = 0$ force the shifted Chebyshev polynomials to meet the geometric boundary condition of $w(0) = 0$.

Jacobi Polynomials (Jacobi) have also been applied to the RRM [3]. Jacobi polynomials are defined by the orthogonality relationship:

$$\int_{-1}^1 w(x) P_n^{(\alpha, \beta)}(x) P_m^{(\alpha, \beta)}(x) dx = \frac{2^{\alpha+\beta+1}}{2n + \alpha + \beta + 1} \frac{\Gamma(n + \alpha + 1) \Gamma(n + \beta + 1)}{\Gamma(n + 1) \Gamma(n + \alpha + \beta + 1)} \delta_{mn}, \quad (\text{A.7})$$

with the weight function $w(x) = (1-x)^\alpha(1+x)^\beta$. Alanbay et al. present a method to create a set of polynomials which are orthogonal and meet the requirements to be admissible functions for the RRM [3]. The functions are defined as

$$\phi_j(x) = (1-x)^{\alpha/2}(1+x)^{\beta/2} P_{j-1}^{(\alpha, \beta)}(x) \quad \text{for } j = 1, 2, \dots, \quad (\text{A.8})$$

such that the admissible functions themselves are orthogonal. For the Meirovitch bar, the values of $\alpha = 2$ and $\beta = 0$ can be used to impose a pinned-free boundary condition on the polynomials. These admissible functions look very similar to the forced Chebyshev, and can be shifted to be defined from 0 to 1 in the same way. The difference is that these polynomials are orthogonal over the inner product used to create the mass matrix, where Chebyshev polynomials are orthogonal over a different inner product. This means that with these Jacobi admissible functions, the mass matrix for a uniform bar would be diagonal.

Boundary Characteristic Orthogonal Polynomials (BCOP) provide another way to generate orthogonal polynomials, which is capable of incorporating more general boundary conditions and weighting functions [13]. These functions, as presented by Bhat, use a modified Gram-Schmidt procedure to generate orthogonal polynomials which match all boundary conditions, and will create a diagonal mass matrix [20]. For the bar example, the BCOP procedure can create the admissible functions

$$\begin{aligned}
\phi_1(x) &= 2x^2 - 3x \\
\phi_2(x) &= (g(x) - B_j) \phi_1(x) \\
\phi_{j+1}(x) &= (g(x) - B_j) \phi_j(x) - C_j \phi_{j-1}(x) \quad \text{for } j = 1, 2, \dots,
\end{aligned} \tag{A.9}$$

with generating function $g(x) = x^2 - 2x$ and constants $B_j = \langle g\phi_j, \phi_j \rangle / \langle \phi_j, \phi_j \rangle$ and $C_j = \langle g\phi_j, \phi_{j-1} \rangle / \langle \phi_{j-1}, \phi_{j-1} \rangle$ to ensure the set remains orthogonal under the inner product $\langle u, v \rangle = \int_0^L m(x)u(x)v(x) dx$ with $m(x)$ being a weight function to ensure mass orthogonality.

The results from the preceding sets of polynomials show that the polynomial order is all that matters for the convergence speed, and the difference between the first three sets of polynomials is only in the conditioning of the problem. This result is expected, as has been explained in the literature [17]. However, the BCOP are able to match all boundary conditions, yet show poor convergence compared to other polynomials. The generating function used in the modified Gram-Schmidt process leads to a higher polynomial order higher than the other polynomial sets, therefore being unable to create an arbitrary polynomial of the highest order like the other polynomial sets, and slowing convergence.

The final set of polynomial functions is a novel method for creating a set of polynomials which match all boundary conditions, but have the lowest possible polynomial order at each step to keep fast convergence. By enforcing one more boundary condition in the bar example than the previous Chebyshev or Jacobi polynomials, these have a polynomial order one higher, and therefore achieve a similar accuracy with one less function than these polynomial sets.

Boundary Condition Chebyshev Polynomials (BC-Cheb) are a linear combination of Chebyshev polynomials which meet all boundary conditions of the system. This can be written as,

$$\phi_j(x) = \alpha_O T_O^*(x) + \sum_{j=0}^{O-1} \alpha_j T_j^*(x) \tag{A.10}$$

where O is the polynomial order, α_j are the coefficients for each of the shifted Chebyshev polynomials $T_j^*(x)$, and $|\alpha_O| > 0$. The coefficients come from solving a linear system of equations

$$\mathbf{P}\alpha = \mathbf{b}, \tag{A.11}$$

which ensures all boundary conditions are met. In this equation, $\mathbf{A}$ is a matrix of Chebyshev functions sampled at locations in the boundary conditions and $\mathbf{P} \in \mathbb{R}^{N_{BC} \times O}$ if all

boundary conditions are zero or $\mathbf{P} \in \mathbb{R}^{N_{BC} \times O+1}$ if there are non-zero boundary conditions. The vector $\mathbf{b} \in \mathbb{R}^{N_{BC}}$ is the values of boundary conditions if there are non-zero boundary conditions, or they are the negative of $T_O^*(x_{BC})$ if all boundary conditions are zero to avoid a trivial solution. In the bar example, $\mathbf{b} = [0 \quad 1 \quad -1]^\top$ and the columns of $\mathbf{P}$, $p_j = [T_{j-1}^*(0) \quad T_{j-1}^*(1) \quad T_{j-1}'(1)]^\top$ while the maximum order for each admissible function is $O = 2, 3, \dots$. If all boundary conditions were zero, then $\mathbf{b} = -p_O$ and that column is excluded from $\mathbf{P}$.

Appendix B

CRI Optimization Details

This appendix provides derivations for the derivatives used in the iterative optimization for CRI.

For the Newton iterations we need $\frac{\partial J}{\partial \mathbf{z}}$ and $\frac{\partial^2 J}{\partial \mathbf{z} \partial \mathbf{z}^\top}$. These derivatives will be made up of different parts corresponding to each mode:

$$\frac{\partial J}{\partial \mathbf{z}} = \begin{bmatrix} \frac{\partial J}{\partial \mathbf{v}_1} \\ \vdots \\ \frac{\partial J}{\partial \mathbf{v}_{N_m}} \end{bmatrix}, \quad (\text{B.1})$$

and

$$\frac{\partial^2 J}{\partial \mathbf{z} \partial \mathbf{z}^\top} = \begin{bmatrix} \frac{\partial^2 J}{\partial \mathbf{v}_1 \partial \mathbf{v}_1^\top} & \cdots & \frac{\partial^2 J}{\partial \mathbf{v}_1 \partial \mathbf{v}_{N_m}^\top} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 J}{\partial \mathbf{v}_{N_m} \partial \mathbf{v}_1^\top} & \cdots & \frac{\partial^2 J}{\partial \mathbf{v}_{N_m} \partial \mathbf{v}_{N_m}^\top} \end{bmatrix}. \quad (\text{B.2})$$

First we will expand J :

$$\begin{aligned} J = F^* F - \sum_{i=1}^{N_m} (\mathbf{v}_i^\top \otimes \mathbf{v}_i^\top) (\Phi_{\text{in}}^\top \otimes S_i^*) F - F^* \sum_{j=1}^{N_m} (\Phi_{\text{in}} \otimes S_j) (\mathbf{v}_j \otimes \mathbf{v}_j) \\ + \sum_{i=1}^{N_m} \sum_{j=1}^{N_m} (\mathbf{v}_i^\top \otimes \mathbf{v}_i^\top) (\Phi_{\text{in}}^\top \otimes S_i^*) (\Phi_{\text{in}} \otimes S_j) (\mathbf{v}_j \otimes \mathbf{v}_j) \end{aligned} \quad (\text{B.3})$$

Next we will take the derivative of J with respect to one of the mode coefficient vectors $\mathbf{v}_i$. This will require the property of a Kronecker product: $\frac{\partial}{\partial \mathbf{a}} \mathbf{a} \otimes \mathbf{a} = \mathbf{a} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{a}$. The derivative is then:

$$\frac{\partial J}{\partial \mathbf{v}_i^\top} = (\mathbf{v}_i^\top \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{v}_i^\top) \left[(\Phi_{\text{in}}^\top \otimes S_i^*) F + \sum_{j=1}^{N_m} (\Phi_{\text{in}}^\top \otimes S_i^*) (\Phi_{\text{in}} \otimes S_j) (\mathbf{v}_j^\top \otimes \mathbf{v}_j^\top) \right], \quad (\text{B.4})$$

and the second derivative is:

$$\frac{\partial^2 J}{\partial \mathbf{v}_i^\top \partial \mathbf{v}_j} = (\mathbf{v}_i^\top \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{v}_i^\top)(\Phi_{\text{in}}^\top \otimes S_i^*)(\Phi_{\text{in}} \otimes S_j)(\mathbf{v}_j \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{v}_j). \quad (\text{B.5})$$

Computational Efficiency Considerations

Some things can be precomputed to save time during iterations. Firstly, the term $(\Phi_{\text{in}}^\top \otimes S_j^*)F$ can be computed for each mode at the beginning and stored to be accessed during each iteration. Secondly, the term $(\Phi_{\text{in}}^\top \otimes S_i^*)(\Phi_{\text{in}} \otimes S_j)$ requires two large Kronecker products which add computational time. The property $(A \otimes B)(C \otimes D) = (AC \otimes BD)$ can be used to rewrite this term as $(\Phi_{\text{in}}^\top \Phi_{\text{in}} \otimes S_i^* S_j)$, where once again a large computation savings can be made by precomputing at least the $S_i^* S_j$ term.

With these modifications in mind we can rewrite our derivatives, where quantities that are precomputed are shown in [blue](#):

$$\frac{\partial J}{\partial \mathbf{v}_i^\top} = (\mathbf{v}_i^\top \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{v}_i^\top) \left[(\Phi_{\text{in}}^\top \otimes S_j^*)F + \sum_{j=1}^{N_m} (\Phi_{\text{in}}^\top \Phi_{\text{in}} \otimes S_i^* S_j)(\mathbf{v}_j^\top \otimes \mathbf{v}_j^\top) \right], \quad (\text{B.6})$$

$$\frac{\partial^2 J}{\partial \mathbf{v}_i^\top \partial \mathbf{v}_j} = (\mathbf{v}_i^\top \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{v}_i^\top)(\Phi_{\text{in}}^\top \Phi_{\text{in}} \otimes S_i^* S_j)(\mathbf{v}_j \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{v}_j). \quad (\text{B.7})$$

For the same reasons, an increase in efficiency can be found in the calculation of the residual error which may be used to check for convergence. The original equation, following from the model, would be:

$$Residual = F - \sum_{i=1}^{N_m} (\Phi_{\text{in}} \otimes S_i)(\mathbf{v}_i \otimes \mathbf{v}_i), \quad (\text{B.8})$$

which can be rewritten as:

$$Residual = F - \sum_{i=1}^{N_m} (\Phi_{\text{in}} \mathbf{v}_i \otimes S_i \mathbf{v}_i), \quad (\text{B.9})$$

in order to reduce the number of operations.

Appendix C

CRI Code

This appendix includes MATLAB code to run the CRI algorithm, as well as an example of how CRI can be applied to a 4DOF system like that shown in Chapter 5. The function in Section C.2 directly applies the CRI algorithm detailed in Chapter 5. This code is designed to be readable, and to accurately apply the algorithm. However it is missing some of the safeguards that would accompany a finalized version of code. Despite this, when properly applied this code can be used to fit a CRI model to measured FRFs from any system. Also included in Section C.3 is a function to take a CRI fit and create a state space model. Figure C.1 shows the output of the example code in Section C.1, applying CRI to a 4DOF system.




Figure C.1: Output figures from the example 4DOF system from the code in Section C.1.

C.1 CRI Example

```

close all;

%% Set up 4DOF system
N = 4;
M = eye(N);
K = 1e3*(2*eye(N) - diag(ones(1,N-1),1) - diag(ones(1,N-1),-1));
K(N,N) = K(N,N)/2;
D = M + 1e-4*K;
% D(1) = D(1)*3; % Uncomment for complex modes

%% Create State Space
A = [zeros(N) inv(M); -(K), -inv(M)*D];
B = [zeros(N); inv(M)];
C = [eye(N) zeros(N)];

[V,L] = eig(A);
poles = diag(L);
[1,inds] = sort(abs(poles));
poles = poles(inds(1:2:end)); % One for each conjugate pair

Ns = 1000;
w = linspace(0,ceil(1.5*max(abs(poles))),Ns);
H = freqresp(ss(A,B,C,zeros(N)),w);

%% Fit CRI model
opts.complex = 0; % Set to 1 for complex modes
opts.tol = 1e-15;
opts.alpha = 1;
opts.random = 0;

P = eye(4);
[v,rmse] = CRI(H(2:4,1:2,:),w,poles,P(2:4,:),P(1:2,:),ones(size(w)),opts);
[At,Bt,Ct,Dt] = CRI2SS(poles,v,P,P,opts);
Ht = freqresp(ss(At,Bt,Ct,Dt),w);

%% Check accuracy
H = reshape(H,N*N,Ns); Ht = reshape(Ht,N*N,Ns);

figure;
for k = [2:4 6:8]
    ll=semilog(w/2/pi,abs(H(k,:)), "Color",.8*[1 1 1],"LineWidth",1.25);
    hold on;

```

```

    l2=semilogy(w/2/pi,abs(Ht(k,:)),"-.k","LineWidth",1.25);
    l3=semilogy(w/2/pi,abs(H(k,:)-Ht(k,:)),"-.b");
end
legend([l1 l2 l3],{"True FRFs","Fitted FRFs","Error"},"location","east");
xlim([0 12]); ylim([1e-20 1e0]); grid on;
xlabel("Frequency (Hz)"); ylabel("FRF Magnitude (m/N)");
title("Fitted FRFs");

figure;
for k = [1 5 9:16]
    l1=semilogy(w/2/pi,abs(H(k,:)),"Color",.8*[1 1 1],"LineWidth",1.25);
    hold on;
    l2=semilogy(w/2/pi,abs(Ht(k,:)),"-.k","LineWidth",1.25);
    l3=semilogy(w/2/pi,abs(H(k,:)-Ht(k,:)),"-.r");
end
legend([l1 l2 l3],{"True FRFs","Interp. FRFs","Error"},"location","east");
xlim([0 12]); ylim([1e-20 1e0]); grid on;
xlabel("Frequency (Hz)"); ylabel("FRF Magnitude (m/N)");
title("Interpolated FRFs");

```

C.2 CRI Function

```

function [v,rmse] = CRI(H,w,poles,Po,Pi,weight,opts,v0)
% function [v,rmse] = CRI(H,w,poles,Po,Pi,weight)
% function [v,rmse] = CRI(H,w,poles,Po,Pi,weight,opts)
% function [v,rmse] = CRI(H,w,poles,Po,Pi,weight,opts,v0)
%
% Inputs:
%   H: Frequency response matrix [p x q x Ns]
%   w: Angular frequency (rad/s) [1 x Ns]
%   poles: Poles to be used in fit [d x 1]
%           One of each complex-conjugate pole
%   Po: Output fitting matrix [p x r]
%   Pi: Input fitting matrix [q x r]
%   weight: Weight for each frequency [1 x Ns] or [1 x 1]
%           If scalar, weight is 1 for all frequencies
%   opts: Struct with options
%   v0: Initial coefficient matrix [r x d]
%
% Options:
%   iter_max: Maximum number of iterations
%   tol: Tolerance for slope at convergence
%   complex:
%       0 -> Fit real modes
%       1 -> Fit complex modes
%   alpha: Damping term to prevent oscillations in iterations
%   random:
%       0 -> Deterministic optimization starting point
%       1 -> Random optimization starting point
%
% Output:
%   v: Matrix of coefficients for each mode [r x d]
%   rmse: Vector of relative error at each step [iter+1 x 1]

%% Set up all options
def.iter_max = 100; % Default maximum number of iterations
def.tol = 1e-7; % Default tolerance for slope at convergence
def.complex = 0; % Default is to fit with real coefficients
def.alpha=1; % Default alpha
def.random=0; % Default to deterministic starting point

if nargin<7 % If no opts are given, use defaults
    opts=def;

```

```

else % Otherwise merge default values into opts
    fn=fieldnames(def);
    for opt = transpose(fn(~ismember(fn,fieldnames(opts))))
        opts.(opt{1}) = def.(opt{1});
    end
end

iter_max = opts.iter_max; tol = opts.tol; alpha = opts.alpha; % Apply opts
if opts.complex < 1, complex = 0; else, complex = 1; end



---


%% Display starting message
disp("Setting up the optimization...");



---


%% Sizes
p = size(Po,1); % Number of outputs
q = size(Pi,1); % Number of inputs
r = size(Po,2); % Number of basis functions
d = length(poles); % Number of modes
Ns = length(w); % Number of samples



---


%% Initialize mode coefficient matrix
if nargin < 8
    v = eyeish(r,d); % Initialize v
    if opts.random, v = orth(rand(r,r))*v; end % Randomize starting point
else
    v = v0; % Use provided starting point
end

if complex
    poles = [poles; conj(poles)]; d = d*2; v = [v conj(v)];
    if nargin < 8
        v = v*(3-3i)/mean(abs(poles)); % Adjust scaling of v
    end
end



---


%% Create the vectorized F and pole matrix S
F = zeros(p*q*Ns,1); % Vectorized measured transfer function
for k = 1:q
    F_col = weight.*squeeze(H(:,k,:));
    F((k-1)*p*Ns+1:k*p*Ns) = F_col(:); % Vectorize the FRF
end

S = zeros(p*Ns,d*q); % Pole matrix

```

```

for k = 1:d
    if ~complex
        s = weight./((1i*w-poles(k)).*(1i*w-conj(poles(k)))); % Real S
    else
        s = weight./(1i*w-poles(k)); % Complex S
    end
    s = repmat(s,p,1); s = s(:); % Vectorize the denominator of S
    sn = repmat(Po,Ns,1); % rep Po for numerator of S
    S(:, (k-1)*r+1:k*r) = sn.*s;
end

%% Precompute some large products that collapse dimensions *Important*
inds = reshape(1:r*d,r,d); % Initialize indices for separating modes
V = zeros(1,r*r*d); vinds = reshape(1:r*r*d,r*r,d); % Precompute for dJ
for k = 1:d
    V((k-1)*r*r+1:k*r*r) = ctranspose(F)*kron(Pi,S(:,inds(:,k)));
end
V = ctranspose(V); S2 = ctranspose(S)*S;

%% Newton iterations
rmse = []; rmse = addRMSE(rmse,F,S,v,Pi,inds); % Initialize rmse
for iter = 1:iter_max
    disp(strcat(" Iteration ",num2str(iter))); % Display progress
    dJ = diffJ(S2,V,v,Pi,inds,vinds); % Take the derivative of J
    ddJ = ddiffJ(S2,v,Pi,inds); % Take the second derivative of J
    dv = reshape(pinv(ddJ)*dJ,r,d); % Find the direction for v to move

    if ~complex
        v = v - real(dv)/alpha; % Update real v
    else
        v = v - dv/alpha; % Update complex v
        v(:, (d/2+1):end) = conj(v(:,1:d/2)); % Clamp the coefficients
    end

    rmse = addRMSE(rmse,F,S,v,Pi,inds); % Update rmse
    if size(rmse,2) > 2
        if abs(mean(diff(rmse(end-2:end)))) < tol % Check slope
            if complex
                v = v(:,1:d/2); % Only return one vector per mode
            end
            disp("Iterations complete"); break; % Finish loop
        end
    end
end
end

```

```

%% Loop functions
% Update the rmse of the fit
function rmse = addRMSE(rmse,F,S,v,Pi,inds)
    R = F;
    for m = 1:size(inds,2) % Subtract each mode from R to get residual
        R = R - kron(Pi*v(:,m),S(:,inds(:,m))*v(:,m));
    end
    rmse(1,size(rmse,2)+1) = norm(R)/norm(F); % Track relative error
end

% Take the first derivative of J
function dJ = diffJ(SS,V,v,Pi,inds,vinds)
    dJ = zeros(numel(inds),1);
    for m = 1:size(inds,2)
        cm = v(:,m);
        vm = -V(vinds(:,m),:);
        for n = 1:size(inds,2) % Add contributions to vm from each mode
            cn = v(:,n);
            PiPi = transpose(Pi)*Pi;
            vm = vm + kron(PiPi,SS(inds(:,m),inds(:,n)))*kron(cn,cn);
        end
        dJm = k_sum(transpose(cm))*vm;
        dJ(inds(:,m)) = dJm;
    end
end

% Take the second derivative of J
function ddJ = ddiffJ(SS,v,Pi,inds)
    ddJ = zeros(numel(inds),numel(inds));
    for m = 1:size(inds,2)
        cmt = transpose(v(:,m));
        for n = 1:size(inds,2)
            cn = v(:,n);
            PiPi = transpose(Pi)*Pi;
            ddJmn = k_sum(cmt)*kron(PiPi,SS(inds(:,m),inds(:,n)))*k_sum(cn);
            ddJ(inds(:,m),inds(:,n)) = ddJmn;
        end
    end
end

%% Helper Functions
% Kronecker sum of vector x
function k_s = k_sum(x)

```

```

    l = length(x);
    k_s = kron(eye(l),x) + kron(x,eye(l));
end

% Eyeish: Deterministic low condition number matrix
function A = eyeish(n,m)
    d1 = max(n,m); d2 = min(n,m);
    if d2 == 1, A = ones(n,m); return; end
    A = zeros(d2,d1); e = eye(d2);
    for j = 1:d2
        col = interp1(1:d2,e(j,:),linspace(1,d2,d1));
        A(j,:) = col;
    end
    if n > m, A = transpose(A); end; A = A./vecnorm(A);
end
end

```

C.3 CRI2SS Function

```

function [A,B,C,D] = CRI2SS(poles,v,Po,Pi,opts)
% function [A,B,C,D] = CRI2SS(poles,v,Po,Pi)
% function [A,B,C,D] = CRI2SS(poles,v,Po,Pi,opts)
%
% Inputs:
%     poles:  System poles
%     v:      Coefficient vectors
%     Po:     Output interpolation matrix
%     Pi:     Input interpolation matrix
%     opts:   Options structure
%
% Options:
%     complex:
%         0 -> Interpolate real modes
%         1 -> Interpolate complex modes

if nargin < 5, opts.complex=0; end
A = [transpose(poles); ctranspose(poles)]; A = diag(A(:));

V = zeros(size(v,1),size(A,1));
for k = 1:size(V,2)
    V(:,k) = v(:,ceil(k/2));
    if opts.complex && mod(k,2) == 0
        V(:,k) = conj(V(:,k));
    end
end

if opts.complex
    C = Po*V;
else
    C = -1i*Po*V*inv(imag(A))/2;
end
B = transpose(V)*transpose(Pi);

D = zeros(size(C,1),size(B,2));
end

```