

Carnegie Mellon University

CARNEGIE INSTITUTE OF TECHNOLOGY

THESIS

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF Doctor of Philosophy

TITLE

Data-Driven, Sparsity-Based Matched Field

Processing for Structural Health Monitoring

PRESENTED BY

Joel Harley

ACCEPTED BY THE DEPARTMENT OF

Electrical and Computer Engineering

José Moura

ADVISOR, MAJOR PROFESSOR

5/1/14

DATE

Jelena Kovacevic

DEPARTMENT HEAD

5/1/14

DATE

APPROVED BY THE COLLEGE COUNCIL

Vijayakumar Bhagavatula

DEAN

5/1/14

DATE

Data-Driven, Sparsity-Based Matched Field Processing for Structural Health Monitoring

Submitted in partial fulfillment for the requirements for
the degree of
Doctor of Philosophy
in
Electrical and Computer Engineering

Joel B. Harley

M.S., Electrical and Computer Engineering, Carnegie Mellon University
B.S., Electrical Engineering, Tufts University

Carnegie Mellon University
Pittsburgh, PA

May 2014

Keywords: guided waves, Lamb waves, structural health monitoring, localization, acoustic emission, acousto-ultrasonics, compressed sensing, matched field processing, basis pursuit, scale transform, temperature compensation

*dedicated to my family,
for teaching me to be independent and reach for the stars,
and to my wife Jen,
for supporting me from afar for nearly six years*

Abstract

This dissertation develops a robust, data-driven localization methodology based on the integration of matched field processing with compressed sensing ℓ_1 recovery techniques and scale transform signal processing. The localization methodology is applied to an ultrasonic guided wave structural health monitoring system for detecting, locating, and imaging damage in civil infrastructures. In these systems, the channels are characterized by complex, multi-modal, and frequency dispersive wave propagation, which severely distort propagating signals. Acquiring the characteristics of these propagation mediums from data represents a difficult inverse problem for which, in general, no readily available solution exists. In this dissertation, we build data-driven models of these complex mediums by integrating experimental guided wave measurements with theoretical wave propagation models and ℓ_1 sparse recovery methods from compressed sensing. The data-driven models are combined with matched field processing, a localization framework extensively studied for underwater acoustics, to localize targets in complex, guided wave environments. The data-driven matched field processing methodology is then refined, through the use of the scale transform, to achieve robustness to environmental variations that distort guided waves. Data-driven matched field processing is experimentally applied to an ultrasound structural health monitoring system to detect and locate damage in aluminum plate structures.

Acknowledgements

I would like to acknowledge each member of the my committee for their support, insights, and guidance: Professors José M. F Moura, Jacobo Bielak, Rohit Negi, Soumya Kar, and William A. Kuperman. For Professor Negi, thank you teaching me the groundwork in estimation and optimization used in this dissertation. For Professor Bielak, thank you for teaching me the physics of wave propagation and discussing applications of my ideas in seismology and geoengineering. For Professor Kar, thank you for your career advice and discussions on sensor network aspects of my work. For Professor Kuperman, thank you for sharing your vast knowledge on underwater acoustics and discussing applications for my ideas in that domain.

Above all, I would like to acknowledge my graduate advisor, Professor Moura. Thank you for teaching me to write technical papers and present to technical audience. Thank you for pushing me to work hard, meet deadlines, and write well. Thank you for teaching me about academia and collaborating with other researchers. Finally, thank you for demanding excellence in everything. This dissertation would not have been possible without your guidance and support.

I would like to thank my colleagues, collaborators, and mentors throughout my PhD career. I would like to thank Professors Irving J. Oppenheim, David W. Greve, Mario Bergés, James H. Garrett, Lucio Soibelman, Mark Patton, Marton H. Altschul, and Jelena Kováček for teaching me structural health monitoring, waves, devices, machine learning, and the countless other skills that I have learned as a graduate student. I particularly want to acknowledge Professor Oppenheim and Professor Greve for being steadfast mentors on research and teaching. To my fellow graduate student collaborators, Nicholas O'Donoghue, Yujie Ying, Chang Liu, Peng Gong, Jun Shi, Aurora Schmidt, Jia Guo, Qing Shi, and Siheng Chen: it has been an honor to work with such intelligent and motivated individuals.

I would further like to acknowledge my collaborators at the Westinghouse Electric Company, including John P. Lareau, Mike Burke, Tom Nenno, and John Duryea.

I would particularly like to thank Warren Junker from the Westinghouse Electric Company for imparting his help, advice, and practical knowledge from working in the nondestructive testing industry for many years.

I am also grateful for the amazing support from other staff and faculty at Carnegie Mellon University, including Samantha Goldstein, Elaine Lawrence, Dave Fortna, Tara Moe, and Claire Bauerle. I would particularly like to recognize Carol Patterson, who helped get anything and everything done, and Susan Farrington, an invaluable advisor to the ECE Graduate Student Organization. Furthermore, I would like to thank Hilary Schuldt for helping me prepare for a career as an educator.

I would like to acknowledge my officemates and colleagues with whom I have had the pleasure to work alongside for many years: Nicholas O'Donoghue, Basil AsSadhan, Bernardo Pires, Augusto Santos, Evgeny Toropov, Jonathan Mei, and Yuan Chen. I also would like to thank all of the other past and present members of José Moura's group, including Soumya Kar, Divyanshu Vats, Marek Telgarsky, Usman Khan, Kyle Anderson, Aurora Schmidt, June Zhang, Joya Deri, Sérgio Pequeto, Dusan Jakovetic, Dragana Bajovic, Subhro Das, Stephen Kruzick, Liangyan Gui, Matthew Baron, with whom I have been able to talk with, learn from, and make every day a little bit more interesting.

Finally, I would like to acknowledge the other graduate students I have worked alongside at Carnegie Mellon University. In particular, I would like to thank the ECE Graduate Student Organization officers: Matthew Beckler, Shah Amini, Johnathan Becker, Jon Smereka, Milda Ziztye, Stephen Siena, Tianyu Gu, and Christian Koehler. I also thank all the students in Porter Hall for their support and advice, including Akshay Rajhans, Rohan Chabukswar, Nikos Arechiga, Kyri Baker, Jhi-Young Joo, Andrew Hsu, Nipun Popli, Jon Donadee, Javad Mohammadi, and Anit Sahu.

This work was also supported by a National Defense Science and Engineering Graduate Fellowship, sponsored by the Department of Defense and Office of Naval Research, and by a National Science Foundation Graduate Research Fellowship under Grant No. 0946825.

Contents

List of Figures	13
List of Tables	18
1 Introduction	19
1.1 Motivation	19
1.2 Challenges	20
1.2.1 Challenge 1: Uncertainty	21
1.2.2 Challenge 2: Complexity	23
1.2.3 Challenge 3: Variability	25
1.3 Model Framework	26
1.3.1 Lamb wave model	26
1.3.2 Sensor model	28
1.4 Dissertation Outline	29
2 Learning Data-Driven Models: Sparse Wavenumber Processing	31
2.1 Motivation	31
2.2 The frequency-wavenumber space	33
2.3 Sparse Recovery by Basis Pursuit	35
2.3.1 Sparse Wavenumber Analysis	36
2.3.2 Sparse Wavenumber Synthesis	38
2.3.3 Debiasing Results	39
2.3.4 Comparison with least squares	39
2.4 Simulation and experimental methodologies	41
2.4.1 Data collection	42
2.4.2 Simulation models	42
2.5 Simulation results and discussion	43
2.5.1 Scenario 1: Unbounded plate	46
2.5.2 Scenario 2: Multipath	47
2.5.3 Scenario 3: Multipath with window	48

2.5.4	Scenario 4: Multipath and position noise with window	49
2.5.5	Comparison with least squares	51
2.6	Experimental results and discussion	52
2.6.1	Sparse wavenumber analysis	53
2.6.2	Sparse wavenumber denoising	54
2.7	Conclusion	54
3	Leveraging Guided Wave Complexity: Data-Driven Matched Field Processing	57
3.1	Motivation	57
3.2	Data collection	59
3.2.1	Calibration data	59
3.2.2	Test data	60
3.2.3	Model data	60
3.3	Data-driven matched field processing	61
3.3.1	Building the data-driven model	62
3.3.2	Coherent data-driven matched field processor	63
3.3.3	Incoherent data-driven matched field processing	65
3.4	Asymptotic behavior	66
3.4.1	Single path scenario	68
3.4.2	Multipath scenario	69
3.5	Comparison with clustered sensor topologies	70
3.6	Conclusions	72
4	Coherent Data-Driven Matched Field Proc.: Acousto-Ultrasonic Localization	75
4.1	Motivation	75
4.2	Data-driven calibration	76
4.3	Experimental methodology	78
4.3.1	Method parameters	79
4.3.2	Comparison with delay-and-sum	80
4.4	Experimental results and discussion	82
4.4.1	Figures of merit	82
4.4.2	One scatterer	83
4.4.3	Two scatterers	86
4.5	Conclusions	88
5	Incoherent Data-Driven Matched Field Proc.: Acoustic Emission Localization	91
5.1	Motivation	91
5.2	Data-driven calibration	93
5.3	Experimental methodology	94
5.3.1	Method parameters	95
5.3.2	Comparison with multilateration	97
5.4	Experimental results and discussion	99
5.4.1	Multipath study	99
5.4.2	Noise study	102

5.5	Conclusions	104
6	Refining for Environmental Variations: The Scale Transform	105
6.1	Motivation	105
6.2	Scale transform signal processing	107
6.2.1	Problem Formulation	107
6.2.2	The Scale Transform	108
6.2.3	The Fast Mellin Transform	110
6.2.4	Maximization of the Scale Cross-Correlation Function	111
6.3	Performance of Scale Transform Methods	114
6.3.1	Scale-invariant correlation (SIC) method	114
6.3.2	Iterative scale transform (IST) method	115
6.3.3	SIC/IST combination	116
6.3.4	Comparison with the optimal signal stretch (OSS) method	117
6.4	Results: Single sensor compensation	119
6.4.1	Experimental methodology	119
6.4.2	Estimation accuracy	120
6.4.3	Computational cost	121
6.5	Results: Data-driven matched field integration	126
6.5.1	Scale transform temperature compensation	126
6.5.2	Experimental methodology	127
6.5.3	Results and discussion	128
6.6	Conclusions	130
7	Conclusions and Future Work	131
7.1	Extension to new applications	133
7.2	Integration with data science	133
7.3	Application to imaging modalities	134
	Appendix A The restricted nullity property	137
	Appendix B Proof of (3.23): single path scenario (coherent processor)	139
B.1	Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$	139
B.2	Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$	139
B.3	Target-to-artifact ratio	140
	Appendix C Proof of (3.23): single path scenario (incoherent processor)	141
C.1	Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$	141
C.2	Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$	141
C.3	Target-to-artifact ratio	142
	Appendix D Proof of (3.25): multipath scenario (coherent processor)	143
D.1	Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$	143
D.2	Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$	144
D.3	Target-to-artifact ratio	145

Appendix E Proof of (3.25): multipath scenario (incoherent processor)	147
E.1 Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$	147
E.2 Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$	148
E.3 Target-to-artifact ratio	149
Bibliography	151

List of Figures

1.1	Dispersion curves for the waves in an aluminum plate with a Poisson ratio of 0.35 and density of 2700 kg/m^3 . Solid lines represent asymmetric wave modes (denoted by A0, A1, and A2) and dotted lines represent antisymmetric wave modes (denoted by S0, S1, and S2).	22
1.2	An example time-domain measurement from a Lamb wave ultrasonic source with different frequency bandwidths and different preprocessing. (a) Wideband, 500 kHz bandwidth measurement. (b) Narrowband 50 kHz bandwidth measurement. (c) Narrowband 50 kHz bandwidth measurement after envelop extraction.	24
1.3	Excitations signals corresponding to measurements in Figure 1.2. (a) Wideband, 500 kHz bandwidth chirp excitation. (b) Narrowband 50 kHz bandwidth Gaussian excitation.	24
1.4	Coordinate axes (r, θ, z) for the plate geometry considered in this dissertation. . . .	27
2.1	Dispersion curves from Figure 1.1 uniformly discretized across 100 wavenumber values with amplitudes generated from the Lamb wave's theoretical displacement amplitudes.	34
2.2	Sensor configuration for simulations and experiments. Each diamond represents the location of a sensor on a plate.	41
2.3	Wideband and narrowband time responses between sensors 14 and 6, labeled in Figure 2.2. Each narrowband response is a Gaussian filtered response with center frequency of 180 kHz and 120 kHz bandwidth. (a) Wideband simulation of a plate with windowed, multipath effects. (b) Wideband experimental measurement with window. (c) Narrowband simulation of a plate with windowed, multipath effects. (e) Narrowband experimental measurement with window.	44
2.4	The sparse wavenumber analysis recovered frequency-wavenumber magnitudes $ \mathbf{v}(\omega) $ for four simulations: (a) unbounded plate, (b) plate with multipath effects, (c) plate with a windowed response and multipath effects, (d) plate with a windowed response, multipath effects, and sensor location uncertainty.	45

2.5	Narrowband direct path responses synthesized by sparse wavenumber denoising (thin, dark line) and corresponding narrowband measured responses (broad, light line) between sensor 14 and sensor 6 for four simulations: (a) unbounded plate, (b) plate with multipath effects, (c) plate with a windowed response and multipath effects, (d) plate with a windowed response, multipath effects, and sensor location uncertainty.	45
2.6	The correlation coefficient between 1000 randomly chosen responses generated from sparse wavenumber synthesis and the corresponding true responses for 3 different simulation scenarios: (a) unbounded plate, (b) plate with multipath effects, (c) plate with windowed, multipath effects.	47
2.7	The correlation coefficient between 1000 randomly chosen responses generated from sparse wavenumber synthesis and the corresponding true responses for a simulation of a plate with windowed, multipath effects and sensor position noise. . . .	50
2.8	The magnitude of the least-squares frequency-wavenumber representation of simulation data of an (a) unbounded plate and (b) plate with with a windowed response and multipath effects.	51
2.9	The average frequency magnitude responses for a simulation with windowed multipath effects and experimental data.	52
2.10	(a) The magnitude of the recovered frequency-wavenumber representation $ \mathbf{v}(\omega) $ for the data with light gray lines illustrating the theoretical dispersion curves for the experiment. (b) The synthesized narrowband direct path responses generated from sparse wavenumber denoising (thin, dark line) and corresponding narrowband measured responses (broad, light line) between sensor 14 and sensor 6 for our experiment.	53
3.1	The frequency-wavenumber representation (dispersion curves) of guided wave data from an aluminum plate, recovered by sparse wavenumber analysis.	61
3.2	The average ambiguity ratio of the localization results versus signal-to-interference ratio η for $M = 10$ (dark lines) and 25 (light lines) sensors. Solid lines show results for a random topology while dotted lines illustrate results for a clustered topology.	72
3.3	The average ambiguity ratio of the localization results versus the number of sensors in the system M for signal-to-interference ratios of $\eta = 5$ dB (dark lines) and -9 dB (light lines). Solid lines show results for a random topology while dotted lines illustrate results for a clustered topology.	72
4.1	A diagram of the aluminum plate used for our experimental setup. Each square indicates a sensor used to transmit and receive signals and the filled circles indicate the locations of the two holes, drilled through the plate, that are used in the experiments.	78
4.2	Example broadband and narrowband signals from the calibration and test data. The test data represents residue signals after baseline subtraction.	79
4.3	An example plot of the normalized ambiguity function from the data-driven model as function of distance. The dotted line represents the Gaussian curve fit to the data.	83

4.4	Ambiguity functions resulting from applying the data-driven, delay, and envelope-delay models to localize a single 0.5 cm hole in an aluminium plate. The figures show the entire 1.22 m by 1.22 m plate. The squares and circles denote the locations of the sensors and the hole, respectively.	85
4.5	Ambiguity functions resulting from applying the data-driven, delay, and envelope-delay models to localize a single 0.5 cm hole in an aluminium plate. The figures show a 4 cm by 4 cm region around the scatterer. The squares and circles denote the locations of the sensors and the hole center, respectively. The cross denotes the estimated location of the hole.	85
4.6	Ambiguity functions resulting from applying the data-driven, delay, and envelope-delay models to localize two 0.75 cm holes in an aluminium plate. The figures show the entire 1.22 m by 1.22 m plate. The squares and circles denote the locations of the sensors and the holes, respectively.	87
4.7	Ambiguity functions resulting from applying the data-driven model to localize two 0.75 cm holes in an aluminium plate. The figures show the two 4 cm by 4 cm regions around each scatterer. The squares and circles denote the locations of the sensors and the hole centers, respectively. The crosses denote the estimated locations of each hole.	88
5.1	Positions of each sensor location (open squares) and acoustic emission(AE) test location (closed circles) on the experimental plate setup. In the experiments, each acoustic emission location is tested separately.	95
5.2	A recovered 3 kHz to 19 kHz frequency-wavenumber, or dispersion curve, representation from measured Lamb wave data.	96
5.3	An example measured Lamb wave signal (a) without additional processing, (b) with narrowband filtering around 10 kHz, (c) with additional white, Gaussian noise, (d) with additional white Gaussian noise and narrowband filtering.	97
5.4	The acoustic emission locations estimated by incoherent data-driven matched field processing (DDMFP, crosses) and multilateration (MLat, plus signs). The true origin of each event is shown as an open circle.	99
5.5	The average localization error, in terms of distance from true location, across the 10 acoustic emission event locations. The mean (dotted lines) and median (solid lines) errors are shown for the incoherent data-driven matched field processor (DDMFP) and multilateration (MLat).	101
5.6	The incoherent data-driven ambiguity functions for a specific acoustic emission location (marked by an open circle) with (a) 4, (b) 7, and (c) 10 sensors. The sensors are marked as white, open squares.	101
5.7	The sample probability of detection versus signal-to-noise ratio for the incoherent data-driven matched field processor (DDMFP) and multilateration (MLat) with 5 sensors (solid lines) or 8 sensors (dotted lines).	103
5.8	The incoherent data-driven ambiguity function measured with 8 sensors and a signal-to-noise ratio of 3 dB. The plus signs (+) show 100 different estimated acoustic emission locations from multilateration.	103

6.1	(a) An experimentally obtained guided wave signal measured at a temperature of 20.5°C. (b) An experimentally obtained guided wave signal measured at a temperature of 7.8°C (c) The scale transform magnitude of the signal in Figure 6.1a. (d) The scale transform magnitude of the signal in Figure 6.1b. (e) The scale cross-correlation between of the signals in Figure 6.1a and Figure 6.1b. For convenience, the horizontal axis is zoomed into the region of interested.	109
6.2	A demonstration of exponential sampling compared with uniform sampling. (a) A continuous sine wave. (b) A uniformly sampled sine wave. (c) An exponentially sampled sine wave.	112
6.3	Block diagram for computing the scale cross-correlation function between $x(t)$ and $s(t)$	113
6.4	The experimental setup shows the aluminum plate and the two PZT (lead zirconate titanate) transducers used to monitor the plate.	119
6.5	Estimation error ($\hat{\alpha} - \alpha$) of SIC, SIC/IST, search OSS, and search/iterative OSS under simulated ideal conditions. In the simulation, experimental measurements were numerically stretched by a factor of 1.001278 and compared with the original signal to estimate the stretch factor.	121
6.6	Comparison between the stretch factor estimate of experimental data, computed using SIC/IST, with ambient temperature. Results show a close correlation between values.	122
6.7	The average, normalized computation time used to compute the search OSS and SIC stretch factor estimates. The search OSS methodology is illustrated for multiple stretch factor ranges $0.008 \geq R \geq 0.024$. Data is fit with curves to illustrate trends.	123
6.8	The average, normalized computation time used to compute the search/iterative OSS and SIC/IST stretch factor estimates. The search/iterative OSS methodology is illustrated for multiple stretch factor ranges $0.008 \geq R \geq 0.024$. Data is fit with curves to illustrate trends.	123
6.9	The average, normalized computation time used to compute the SIC/IST stretch factor estimate after truncating the scale transform domain to various lengths. The value ρ represents the percentage of scale domain retained. Data is fit with curves to illustrate trends.	124
6.10	The average absolute change in the stretch factor estimate of SIC/IST after truncating the scale transform domain to various lengths. The value ρ represents the percentage of scale domain utilized.	125
6.11	The number of iterations used when computing the SIC/IST and search/iterative OSS methods.	125
6.12	The incoherent data-driven ambiguity functions of a 0.5 m by 0.5 m region to localize a 5 cm cylindrical mass on an aluminum plate for three temperature scenarios: (a) no temperature change, (b) temperature change without compensation, and (c) temperature change with compensation. The squares and circles denote the locations of the sensors (only 2 out of 16 sensors fall in the region shown) and the mass, respectively.	129

6.13 The incoherent data-driven ambiguity functions of a 0.06 m by 0.06 m region to localize a 5 cm cylindrical mass on an aluminum plate for three temperature scenarios: (a) no temperature change, (b) temperature change without compensation, and (c) temperature change with compensation. The squares and circles denote the locations of the sensors (only 2 out of 16 sensors fall in the region shown) and the mass, respectively. The crosses denote the estimated locations of the mass. . . . 130

List of Tables

4.1 Figures of merit used to evaluate the localization performance of each model. The figures of merit include the localization accuracy (Acc.), half-amplitude width resolution (Res.), and average peak-to-artifact ratio (PAR). 84

6.1 The computational complexity of each temperature compensation discussed. 117

Introduction

1.1 Motivation

Matched field processing is a widely studied methodology in underwater acoustics for localizing targets in complex environments. Match field processing [1, 2] utilizes a wavefield propagation of the medium, which represents the Green's function between any two arbitrary points in the environment, to find the most likely location of an acoustic source from experimental measurements. Traditionally, these wavefield propagation models are obtained analytically, often by numerical integration of the wave equation, or Helmholtz equation, for environment of interest. Although these approaches work well in theory, they often face significant of complexity, uncertainty, and variability in the environment. In many scenarios, accurate numerical models are not available, or we do not have enough accurate knowledge about the environment's geometry, its physical parameters, and its boundary conditions. This often occurs due to varying environmental and operational conditions [3–5].

This dissertation is concerned with overcoming these challenges of uncertainty, complexity, and variability associated with implementing matched field processing. We accomplish this by developing a new methodology that we refer to as data-driven model for matched field processing. Data-driven matched field processing estimates from data an accurate, theoretical multi-modal and dispersive propagation model of the actual environment that is then used by matched field processing to detect and localize acoustic sources and structural damage. The data-driven model is estimated from data by combining fundamental knowledge about the propagation en-

environment with compressed sensing and ℓ_1 sparse signal recovery. We integrate this data-driven model with the coherent and incoherent matched field processors and demonstrate, both analytically and experimentally, the performance characteristics of the resulting data-driven matched field processors.

We evaluate our methods with ultrasonic guided waves in structural health monitoring applications. Structural health monitoring systems are designed to track large, physical structures for damage and degradation. For example, these systems are of interest to observe the structural integrity of transportation, power, and resource distribution network infrastructures. They can be employed to monitor bridges, pipelines, airplanes, and many other structures. They allow operators to reduce the costs associated with testing and maintaining large infrastructures and help to prevent catastrophic failures that have costly consequences. In the following section, we discuss the significant challenges, which we address in the dissertation, for designing and utilizing guided wave structural health monitoring systems.

1.2 Challenges

Ultrasonic guided waves are an attractive tool for monitoring large, physical structures because they possess low attenuation due to being guided by the boundaries of the structure, and they are highly sensitive to small variations in the structure, such as damage. However, there are significant challenges for analyzing guided waves. The geometry of the structure makes them complex signals that are difficult to interpret without very accurate knowledge of the medium. The waves are also highly sensitive to outside environmental effects, such as temperature. We discuss these challenges in the following subsections.

Guided waves are high-frequency waves (where the wavelength is much smaller than the physical dimensions of the environment under test) that are "guided" by the geometry of a structure, or waveguide. For example, a long and wide plate with finite thickness is a waveguide for acoustic waves. Acoustic waves reflect off of the top and bottom boundaries of the plate, which guide the waves in the direction of the plate's length and width. Guided waves in a plate are commonly known as Lamb waves [6]. Other well known guided waves include Rayleigh and Love waves that are commonly studied in seismology [7], shallow water Perekis guided

waves that are studied in underwater acoustics [8], and electromagnetic waves in transmission lines that are used in electronics and power applications [9].

We address the challenges of uncertainty, complexity, and variability by developing an efficient method that constructs environmental models using physical principles, measured data, and signal processing methods. We illustrate our technique with wave waves. We combine fundamental physical principles that allow us to understand, interpret, and leverage the complex characteristics of guided waves, with measured data, from which we can learn specific unknown parameters of the propagation environment. We learn these propagation characteristics with compressed sensing and integrate them with matched field processing and the scale transform to design and develop an effective ultrasound structural health monitoring system. The next section addresses the challenges considered in this dissertation.

1.2.1 Challenge 1: Uncertainty

Due to the geometry of physical structures, ultrasonic guided waves generally exhibit complex behavior. We consider a large plate and, for simplicity, assume it has infinite length, infinite width, and finite thickness. For this plate, there is well-developed theory that describes the behavior of acoustic waves as they travel through the structure [6]. As waves travel through the plate, the waves interact with the top and bottom surfaces and form interference patterns across the plate's thickness. This constructive and destructive interference forms a number of discrete wave modes. Each wave mode propagates with a unique velocity that varies as a function of frequency, a phenomenon known as dispersion. This multimodal and dispersive behavior significantly distorts the waves as they travel through the waveguide.

The multimodal and dispersive properties of a guided wave are often characterized by their dispersion curves. The dispersion curves of a medium illustrate how the wavenumber, or sometimes phase velocity, of the guided waves changes as a function of frequency [10]. Figure 1.1 illustrates the dispersion curves for Lamb waves in an aluminum plate of finite thickness and infinite length and width. In these plots, each line represents the wavenumber of a propagating Lamb mode as a function of frequency. At the lowest frequencies, two most always propagate in the plate: the zeroth antisymmetric mode (A0) and the zeroth symmetric mode (S0). As

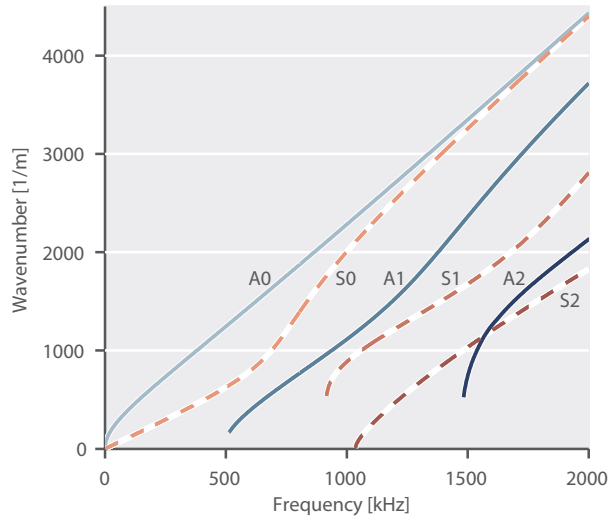


Figure 1.1: Dispersion curves for the waves in an aluminum plate with a Poisson ratio of 0.35 and density of 2700 kg/m^3 . Solid lines represent asymmetric wave modes (denoted by A0, A1, and A2) and dotted lines represent antisymmetric wave modes (denoted by S0, S1, and S2).

frequency increases, more modes are introduced at particular “cut-off” frequencies. The total number of modes is finite, but at higher frequency, the number of grows very large.

Lamb wave modes are usually grouped into two families based on their resonant properties: symmetric modes and antisymmetric modes [11], as illustrated in Figure 1.1. As suggested by the names, symmetric modes are characterized by evenly symmetric wave displacements through the plate’s thickness, while antisymmetric modes are characterized by oddly symmetric wave displacements through the plate’s thickness. We note, however, in this dissertation that we do need nor will distinguish between mode families for purposes other than discussion.

Many techniques have been proposed to recover the multimodal and dispersive velocities from observed data. Among these methods are the two dimensional discrete Fourier transform (2D-DFT) [10, 12], various time-frequency analysis methods [13–15], time-domain matching pursuit approaches [16, 17], and other model-based strategies [18, 19]. While these methods are often effective to analyze guided waves, they are generally not accurate enough to be used to build an effective model of the physics [20]. Furthermore, many of these approaches fail when the plate no longer has an infinite length and infinite width and data is corrupted by unmodeled

multipath interference, which distorts measurements. We refer to the multipath interference as “unmodeled” since it is not accounted for in the simple waveguide model.

In chapter 2 of this dissertation, we develop and demonstrate a method known as sparse wavenumber analysis that utilizes optimization techniques from compressed sensing [21–23] to accurately, efficiently, and robustly recover the dispersion curves of experimental guided wave data, even after corrupted by unmodeled multipath interference. We demonstrate that, from the recovered dispersion curves, we can synthesize guided wave measurements based on simulated and experimental data. The synthesized data is our data-driven model. With simulated data, we show that the synthesized data can match true measurements with accuracies greater than 97.5 %. With experimental data, we achieve similar results.

1.2.2 Challenge 2: Complexity

Due to the distorting effects of multimodal propagation, dispersion, and unmodeled multipath interference, guided wave time traces are very complex. Figure 1.2a illustrates an example of guided wave measurements after traveling from a single source through a 1.2 m by 1.2 m by 0.284 cm aluminum plate. The wideband, 500 kHz bandwidth chirp excitation is shown in Figure 1.3a. In a simple medium, without multiple modes or dispersion, we would only observe delayed replicas of the excitation signals at different points in the structure. From those replicas, we could readily determine the wave’s time-of-flight information. However, we observe in Figure 1.2a that the measured response does not resemble the excitation, quite to the contrary. Therefore, traditional analysis methods that rely on the similarity between the received and transmitted signals are not sufficient for interpreting guided wave data.

To preserve the shape of the excitation with guided waves, researchers often analyze only a narrow band of possible frequencies. Figure 1.2b shows a time trace response from a narrowband excitation for the same sensor setup used in 1.2a. The narrowband, 50 kHz bandwidth Gaussian excitation is shown in Figure 1.3b. While the narrowband measurements are visually simpler, they still only moderately resemble the excitation due to dispersion and unmodeled multipath interference. Also, due to dispersion, the phase information is unreliable because different frequencies travel at different speeds. Therefore, researchers also often remove the phase information by only analyzing the envelope of the narrowband response, which is shown




Figure 1.2: An example time-domain measurement from a Lamb wave ultrasonic source with different frequency bandwidths and different preprocessing. (a) Wideband, 500 kHz bandwidth measurement. (b) Narrowband 50 kHz bandwidth measurement. (c) Narrowband 50 kHz bandwidth measurement after envelop extraction.




Figure 1.3: Excitations signals corresponding to measurements in Figure 1.2. (a) Wideband, 500 kHz bandwidth chirp excitation. (b) Narrowband 50 kHz bandwidth Gaussian excitation.

in Figure 1.2c. Using the signal envelope and an estimate of the guided waves group velocity around the narrow band of frequencies, researchers have employed several methods, such as delay-and-sum localization [24, 25], time reversal [26–28], multilateration [29–31], and probabilistic methods [32, 33] to detect and locate damage in physical structures.

These approaches remove a large amount of valuable information by restricting the data to a narrow band of frequencies and removing phase information. The analysis of guided waves

could be made much more precise and effective by utilizing this knowledge. Using the dispersion curves or specific knowledge about how waves propagates, we can more effectively use the phase information. For example, [34, 35] have developed methods to remove multimodal and dispersive effects from measured data and to perform high resolution localization with time reversal methods [26–28]. These methodologies generally assume *a priori*, accurate dispersion curve knowledge, which is generally not true.

In chapters 3, 4, and 5 in this dissertation, we demonstrate a sparsity-based, data-driven matched field processing methodology that combines sparse wavenumber analysis (discussed challenge 1) with a matched field processing framework [1, 2]. Matched field processing is a model-based localization methodology that has been extensively studied and applied in underwater acoustics. Data-driven matched field processing allows us to leverage the complexity of guided wave signals to perform accurate localization by incorporating data-derived multimodal and dispersive properties within the model. To detect damage, we apply data-driven matched field processing to the difference between a prior baseline signal and a current measurement. We compare our data-driven matched field processing results with traditional structural health monitoring localization methods and demonstrate scenarios in which we can achieve more than a 48 times improvement in localization resolution and more than a 21 times improvement in localization accuracy.

1.2.3 Challenge 3: Variability

For our second challenge, we integrate our data-driven model with matched field processing to locate damage that occurs between two points in time. When the environment changes significantly between those points in time, the baseline subtraction step in matched field processing will not yield reliable data to use. This is because many environmental parameters, such temperature [36] or applied stress [37, 38], have a direct effect on the material properties, and therefore the dispersion curves, of the structure under test.

Temperature is one of the most prominent of these effects [39]. Variations in temperature are known to change the Young’s modulus of a material, therefore affecting the velocity of the guided waves. Due to the complex, multi-modal, dispersive, and the multi-path behavior of guided waves, these variations in velocity can be approximated as a uniform time-scaling

or time-stretching effect on the received signals [40, 41]. While different modes often do not vary at the same rates [3, 42], the aggregate affect on all the multipath components resembles a stretching effect. This model has been experimentally demonstrated to be accurate for sufficiently small changes in velocity [3, 40–43].

To improve the detection and localization of damage in structures, the literature has proposed several methods, such as local peak coherence [41] and optimal signal stretch (OSS) [40], to compensate for temperature by estimating the stretch factor and properly adjusting the measured signal. Local peak coherence [41, 42, 44] is a stretch factor estimation technique based on approximating a stretching operation as a series of time-dependent delay operations. Local peak coherence can be computed very quickly, but is sensitive to other effects that do not uniformly stretch the signal over time. These changes could come from the formation of damage or other environmental and operational effects [4]. In contrast, OSS [40] correlates each observed signal with a library of stretched replicas of a single baseline. This approach is accurate and robust to damage and other variations but computationally inefficient.

In chapter 6 of this dissertation, we present a robust and computationally efficient method for temperature compensation based on scale transform signal processing methods. We compare these results with OSS and demonstrate up to a 6.9 times improvement in computational speed. We then demonstrate how we can combine our previous localization methods with our scale transform temperature compensation method to achieve accurate localization of damage with robustness to environmental changes.

1.3 Model Framework

In this section, we provide a brief derivation of the Lamb wave model used throughout this dissertation. We always assume that our sensors are surface mounted transducers that cannot distinguish between displacements in multiple directions, and instead observe some linear combination of each displacement component.

1.3.1 Lamb wave model

We consider a plate of finite depth, infinite length and width, and traction-free surfaces. We assume cylindrical coordinates such that the z -axis is aligned with the height of the plate. The




Figure 1.4: Coordinate axes (r, θ, z) for the plate geometry considered in this dissertation.

r -axis and the θ -axis represent the distance and angle from the origin and form a plane parallel to the plate's surface. Figure 1.4 illustrates this coordinate system.

Under these conditions, the general phasor solution for displacements originating from $r = 0$ and oscillating in the r -direction U_r and z -direction U_z , can be expressed in the form [6]

$$\begin{aligned} U_r(r, z, \omega) &= \sum_m B_m(z, \omega) H_1^{(1)}(k_m(\omega)r) \\ U_z(r, z, \omega) &= \sum_m C_m(z, \omega) H_0^{(1)}(k_m(\omega)r), \end{aligned} \quad (1.1)$$

where $H_\nu^{(1)}(\cdot)$ is the ν -th order Hankel function of the first kind [45] (also known as a Bessel function of the third kind) and $k_m(\omega)$ is a frequency dependent, i.e., dispersive, wavenumber for wave mode m .

To simplify the expressions in (1.1), we assume our observations occur in the “far-field”. For sufficiently large values of $k_m(\omega)r$, the Hankel function can be approximated as [45]

$$H_\nu^{(1)}(k_m(\omega)r) \approx \sqrt{\frac{2}{\pi k_m(\omega)r}} e^{j(k_m(\omega)r - (1/2)\nu\pi - (1/4)\pi)} \quad (1.2)$$

This far-field approximation is valid for roughly $r > 2\lambda_m(\omega)$ [6], where $\lambda_m(\omega) = (2\pi)/k_m(\omega)$ is the wavelength of mode m . Therefore, the Lamb wave's displacement can be approximately represented as

$$\begin{aligned} U_r(r, z, \omega) &\approx \sum_m \sqrt{\frac{1}{k_m(\omega)r}} \tilde{B}_m(z, \omega) e^{jk_m(\omega)r} \\ U_z(r, z, \omega) &\approx \sum_m \sqrt{\frac{1}{k_m(\omega)r}} \tilde{C}_m(z, \omega) e^{jk_m(\omega)r} \end{aligned} \quad (1.3)$$

where

$$\begin{aligned}\tilde{B}_m(z, \omega) &= \sqrt{\frac{2}{\pi}} B_m(z, \omega) e^{-j(3/4)\pi} \\ \tilde{C}_m(z, \omega) &= \sqrt{\frac{2}{\pi}} C_m(z, \omega) e^{-j(1/4)\pi} .\end{aligned}\tag{1.4}$$

In (1.1) and (1.3), we represent the wave propagation in the r - and z -directions as a summation of guided wave modes. For both displacement functions, each mode m is characterized by a frequency dependent wavenumber $k_m(\omega)$. Figure 1.1 illustrates each theoretical $k_m(\omega)$ function, which can be computed numerically [11], for a Lamb wave propagating in a 0.284 cm thick aluminum plate with a Poisson ratio of 0.35 and density of 2700 kg/m³. These curves are known as the dispersion curves of the medium.

Lamb waves will also usually wave evanescent wave modes such that the wavenumber $k_m(\omega)$ for each mode m is imaginary and attenuate quickly with $k_m(\omega)r$. However, we assume the effects from these evanescent wave modes to be small for sufficiently large $k_m(\omega)r$ values. In our model, we also assume that the propagating waves experience no structural damping due to internal friction.

1.3.2 Sensor model

In our experiments, we use surface mounted PZT (lead zirconate titanate) transducers for transmitting and measuring guided Lamb waves. We choose our sensor locations so that each sensor lies on the same surface $z = z_0$ of the plate. We also assume that our sensors, which convert displacements on the plate surface into an electrical voltage, cannot uniquely distinguish displacements in the r or z directions. Instead the measured voltage signal $X(r, \omega)$ can be represented as a linear combination of each displacement

$$X(r, \omega) = T(\omega)R(\omega)(\alpha U_r(r, z_0, \omega) + \beta U_z(r, z_0, \omega)) ,\tag{1.5}$$

where α and β are arbitrary constants and $T(\omega)$ and $R(\omega)$ represent the transfer functions of the transmitting and receiving transducers respectively [19]. Note that we assume that each transducer can be characterized by the same transfer functions $R(x)$ and $T(x)$. This has been shown to be approximately true for comparable sensors [19].

By plugging the approximate displacement functions in (1.3) into (1.5), we can represent the signal between any two transducers as

$$X(r, \omega) = \sum_m \sqrt{\frac{1}{k_m(\omega)r}} G_m(\omega) e^{jk_m(\omega)r}, \quad (1.6)$$

where r is the distance between a given pair of transducers. In (1.6), the coefficient

$$G_m(\omega) = T(\omega)R(\omega) \left(\alpha \tilde{B}_m(z_0, \omega) + \beta \tilde{C}_m(z_0, \omega) \right) \quad (1.7)$$

represents the frequency dependent complex amplitude for mode m in the measured response.

While (1.6) represents a continuous model of the plate across space and frequency, measured data is sampled across a finite subset of sensors or points in space. Therefore, we represent a collection of M measurements by an $M \times 1$ vector

$$\mathbf{x}(\omega) \approx \begin{bmatrix} X(r_1, \omega) & X(r_2, \omega) & \dots & X(r_M, \omega) \end{bmatrix}^T, \quad (1.8)$$

where r_1, \dots, r_M are known distances between points in the environment. We highlight that this relationship is approximate because most measurements are usually corrupted by some form of error, such as random measurement noise, coherent multipath interference from boundaries, or imperfect modeling of the sensors or system. For most data acquisition and processing systems, we would also use a discrete, finite collection of time samples or frequencies, but for this dissertation, we consider a continuous time and frequency domain without loss of generality.

1.4 Dissertation Outline

In the following chapters, we use the Lamb wave model and sensor model described in the previous sections to recover the dispersion curves and build a propagation model of Lamb waves in a particular aluminum pipe, leverage that model to localize damage or acoustic emissions, and then refine the framework to account for temperature variations in the environment. In Chapter 2, we derive sparse wavenumber analysis and sparse wavenumber synthesis to recover the dispersion curves and generate the propagation model. In Chapter 3, we integrate sparse wavenumber analysis with matched field processing to derive data-driven matched field processing, and we analyze the asymptotic characteristics of data-driven matched field processing. In Chapter 4, we apply coherent data-driven matched field processing to localize holes in

the aluminum plate. In Chapter 5, we apply incoherent data-driven matched field processing to locate spontaneous acoustic emission on an aluminum plate. In Chapter 6, we demonstrate and integrate a method based on the scale transform to compensate for variations in temperature in an the aluminum plate. Together, the content of these chapters present a high resolution, robust structural health monitoring framework. In Chapter 7, we demonstrate how this framework can be extended to new structures, environments, and applications.

Learning Data-Driven Models: Sparse Wavenumber Processing

2.1 Motivation

In guided wave structural health monitoring, researchers are concerned with detecting and locating damage in complex propagation media. These media are often characterized by multimodal and frequency dispersive behavior [46]. This implies that the received measurements can be expressed as a sum of wave modes that travel at different frequency dependent velocities. Often, each velocity may also vary as a function of environmental and operational effects, such as temperature [3, 5]. Furthermore, most guided wave systems in physical infrastructures exhibit a significant degree of multipath effects from physical boundaries. As a result of all of these effects, accurately characterizing and analyzing guided waves is very challenging.

To successfully detect or locate damage, many methods rely on incoherent signal processing techniques, such as envelope extraction [24], to detect and locate areas with damage. Incoherent signal processing techniques can be used to compensate for the distorting effects of dispersion, but these techniques remove potentially useful information, particular phase information, from the measured data. In contrast, accurate estimates of the multimodal and frequency dispersive behavior of a medium could be used to greatly improve current methods [34].

In this chapter, we propose a new methodology that we refer to as sparse wavenumber analysis, which is based on compressed sensing [21–23], to accurately recover the multimodal and

dispersive properties of guided waves in a plate, also known as Lamb waves. Compressed sensing concerned with accurately and efficiently recovering signals and systems that admit sparse signals representations, which contain mostly zeroes [21, 47, 48]. We show that Lamb waves have a sparse representation in the frequency-wavenumber domain and, through this representation, we can accurately recover the parameters of Lamb waves with ℓ_1 optimization techniques. We use these recovered parameters to “denoise” the measurements, removing random noise and multipath signal interference, and generate data-driven models that can predict the measured response between any other two points on the plate. We refer to this process as sparse wavenumber synthesis.

In the following sections, we outline and discuss our methodology for using sparse representations to recover the multimodal and frequency dispersive behavior of Lamb waves. This work was originally presented in references [20, 49, 50]. In Section 2.2, we formulate a general model for Lamb wave behavior. We then discretize this model across the wavenumber domain and pose the problem of estimating the multimodal and frequency dispersive behavior of Lamb waves as a discrete inverse problem that compressed sensing can accurately solve. In Section 2.3, we formulate our sparse wavenumber analysis and sparse wavenumber sparse synthesis techniques, and we briefly review compressed sensing and the basis pursuit algorithm that we use to solve the Lamb wave inverse problem.

In Section 2.4, we discuss our experimental and simulation methodology, which considers 17 ultrasonic transducers randomly arranged across a finite plate. In Section 2.5 and Section 2.6, we show that sparse wavenumber analysis can accurately recover the frequency-wavenumber representation of Lamb waves from simulation and experimental data that are both corrupted by multipath interference. We then show that sparse wavenumber synthesis can be used to successfully remove the multipath interference. Our simulations show that we achieve a correlation coefficient of 0.99 between the true direct path response and the denoised signal. We also apply sparse wavenumber synthesis to the simulated data to predict the response between two arbitrary points in the plate. We show that these predicted responses and the true direct path responses can achieve correlation coefficients greater than 0.97.

2.2 The frequency-wavenumber space

Given the model $X(r, \omega)$ in (1.6), we can fully characterize Lamb waves with knowledge of $k_m(\omega)$ and $G_m(\omega)$ for each mode m . While $k_m(\omega)$ and $G_m(\omega)$ may be derived theoretically, the resulting values are usually not precise due to uncertainty in the properties of the sensors and medium. The complex amplitude $G_m(\omega)$, for example, is significantly affected by the transducers [51] and their bond with the medium [52]. The wavenumber $k_m(\omega)$ is also affected by properties of the medium, which can change with environmental conditions such as temperature [3, 5].

In this section, we present an approach to estimate $k_m(\omega)$ and $G_m(\omega)$ from our data $\mathbf{x}(\omega)$ and the continuous propagation model $X(r, \omega)$. To accomplish this, we discretize $X(r, \omega)$ across the frequency-wavenumber (ω, κ) space. In the continuous frequency-wavenumber space, we can rewrite the expression for Lamb waves in (1.6) as

$$V(\kappa, \omega) = \begin{cases} G_m(\omega) & \text{if } \kappa = k_m(\omega) \text{ for any } m \\ 0 & \text{otherwise} \end{cases}. \quad (2.1)$$

Since the true wavenumber values $k_m(\omega)$ are unknown to us for analysis, we choose to discretize the domain across N possible wavenumber values $\kappa_1, \kappa_2, \dots, \kappa_N$, specifying a sampling interval to achieve some desired resolution. By discretizing the wavenumber domain, we can express (2.1) as a $N \times 1$ wavenumber vector

$$\mathbf{v}(\omega) = \begin{bmatrix} V(\kappa_1, \omega) & V(\kappa_2, \omega) & \dots & V(\kappa_N, \omega) \end{bmatrix}^T \quad (2.2)$$

that varies with frequency.

Figure 2.1 illustrates an example of $\mathbf{v}(\omega)$ coarsely sampled across the wavenumber domain. Each column of Figure 2.1 represents $\mathbf{v}(\omega_i)$, at some given frequency ω_i , discretized over 100 wavenumbers. The shading represents the magnitude of $V(\omega, \kappa)$ at each frequency-wavenumber pair, generated from the Lamb wave's theoretical displacement amplitudes [6]. Therefore, the location of each non-zero component in Figure 2.1 is an estimate of $k_m(\omega)$ and the corresponding value at each location is an estimate of the associated $G_m(\omega)$. Note that the amplitudes shown in Figure 2.1 will likely not correspond to experimental results due to the many factors influencing each mode's amplitude, such as the transducers and their bond with the medium.




Figure 2.1: Dispersion curves from Figure 1.1 uniformly discretized across 100 wavenumber values with amplitudes generated from the Lamb wave's theoretical displacement amplitudes.

The two vectors $\mathbf{x}(\omega)$ and $\mathbf{v}(\omega)$ now illustrate two representations for our model and data: the frequency-distance representation, which we directly measure, and the frequency-wavenumber representation, which specifies $k_m(\omega)$ and $G_m(\omega)$. From (1.6), the two representations are linearly related by

$$\mathbf{x}(\omega) = \mathbf{D}_r \mathbf{A} \mathbf{D}_\kappa \mathbf{v}(\omega), \quad (2.3)$$

where \mathbf{D}_r and \mathbf{D}_κ are $M \times M$ and $N \times N$ diagonal matrices, respectively, defined by

$$\mathbf{D}_r = \text{diag} \left[r_1^{-1/2}, \dots, r_M^{-1/2} \right] \quad (2.4)$$

$$\mathbf{D}_\kappa = \text{diag} \left[\kappa_1^{-1/2}, \dots, \kappa_N^{-1/2} \right] \quad (2.5)$$

and \mathbf{A} is a $M \times N$ generalized Vandermonde matrix defined by

$$\mathbf{A} = \left[e^{j\kappa_j r_i} \right]_{ij}. \quad (2.6)$$

In (2.6), r_1, r_2, \dots, r_M represent the M known distances associated with each measurement and $\kappa_1, \kappa_2, \dots, \kappa_N$ specify the N chosen wavenumbers to represent our data across the frequency-wavenumber domain.

Therefore, (2.3) represents a discrete, linear inverse problem for which our goal is to identify the vector $\mathbf{v}(\omega)$ at each frequency that uniquely represents the multimodal and dispersive properties of the entire medium. In many situations, the number of sampled wavenumbers N will be larger than the number of measurements M . In these circumstances, \mathbf{A} represents an underdetermined system, and therefore, there exists many vectors $\mathbf{v}(\omega)$ that satisfy (2.3). However, as illustrated in (2.1) and Figure 2.1, most elements in $\mathbf{v}(\omega)$ are zero, i.e., the vector is sparse. Therefore, we use sparse recovery methods to solve this inverse problem.

Compressed sensing and sparse recovery techniques are concerned with recovering, from a small number of measurements, signals or systems that admit sparse representations [23]. In the following section, we briefly discuss compressed sensing and basis pursuit optimization, which we will use to recover $\mathbf{v}(\omega)$. In Section 2.5 and Section 2.6, we apply basis pursuit to simulated and experimental data to illustrate that we can recover $\mathbf{v}(\omega)$ both accurately and robustly.

2.3 Sparse Recovery by Basis Pursuit

In this section, we discuss how sparse recovery techniques from compressed sensing, and specifically basis pursuit, is used to recover $\mathbf{v}(\omega)$. By using sparsity, we show that we can accurately recover the medium's frequency-wavenumber representation even when the system is underdetermined [53].

If our system is noise-free and $\mathbf{v}(\omega)$ is sparse, the unique solution to (2.3) can be described as the maximally sparse vector, the vector with the fewest non-zero elements, that satisfies $\mathbf{x}(\omega) = \mathbf{D}_r \mathbf{A} \mathbf{D}_\kappa \mathbf{v}(\omega)$. This solution can be represented as the solution to the optimization problem [23]

$$\begin{aligned} \mathbf{v}(\omega) &= \arg \min_{\mathbf{v}(\omega)} \|\mathbf{v}(\omega)\|_0 \\ \text{s.t. } \mathbf{x}(\omega) &= \mathbf{D}_r \mathbf{A} \mathbf{D}_\kappa \mathbf{v}(\omega) , \end{aligned} \quad (2.7)$$

where the ℓ_0 -pseudo-norm $\|\cdot\|_0$ is defined as the number of non-zero elements in a vector. This optimization is solved by a combinatorial search [23], which quickly becomes intractable for large N . However, a convex relaxation of this optimization problem, known as basis pursuit [22, 54], can often be used in place of (2.7).

2.3.1 Sparse Wavenumber Analysis

The basis pursuit algorithm relaxes (2.7) by replacing the ℓ_0 -norm with an ℓ_1 -norm,

$$\|\mathbf{v}\|_1 = \sum_{n=1}^N |v_n| . \quad (2.8)$$

The use of an ℓ_1 -norm reformulates the intractable optimization in (2.7) as a convex problem [54,55]. This implies that the sparse solution $\mathbf{v}(\omega)$ can be recovered computationally fast using a convex program.

Most implementations of basis pursuit use a normalized sensing matrix Φ such that each column has a unit ℓ_2 -norm [23]. By applying this normalization to our sensing matrix $\mathbf{D}_r \mathbf{A} \mathbf{D}_\kappa$, we get

$$\Phi = \frac{1}{\sqrt{M} \|\mathbf{D}_r\|_F} \mathbf{D}_r \mathbf{A} , \quad (2.9)$$

where $\|\cdot\|_F$ is the Frobenius norm, such that

$$\|\mathbf{D}_r\|_F = \sqrt{\sum_m^M r_m^{-1}} . \quad (2.10)$$

To reduce variations in our solution across each frequency, we also normalize the energy of our measured signal

$$\mathbf{x}_n(\omega) = \frac{\mathbf{x}(\omega)}{\|\mathbf{x}(\omega)\|_2} \quad (2.11)$$

so that $\mathbf{x}_n(\omega)$ has a unit ℓ_2 -norm at each frequency. The normalized frequency-wavenumber representation of our data $\mathbf{v}_n(\omega)$ can then be expressed through basis pursuit as [22, 54, 56]

$$\begin{aligned} \mathbf{v}_n(\omega) &= \arg \min_{\mathbf{v}_n(\omega)} \|\mathbf{v}_n(\omega)\|_1 \\ \text{s.t. } \mathbf{x}_n(\omega) &= \Phi \mathbf{v}_n(\omega) . \end{aligned} \quad (2.12)$$

The unnormalized frequency-wavenumber representation can then be retrieved from

$$\mathbf{v}(\omega) = \|\mathbf{x}(\omega)\|_2 \|\mathbf{D}_r\|_F \mathbf{D}_\kappa \mathbf{v}_n(\omega) . \quad (2.13)$$

However, $\mathbf{x}(\omega)$ is usually corrupted by errors or noise. Therefore, we elect to use a robust implementation of basis pursuit represented as the unconstrained optimization problem [53, 54]

$$\mathbf{v}_n(\omega) = \arg \min_{\mathbf{v}_n(\omega)} \|\Phi \mathbf{v}_n(\omega) - \mathbf{x}_n(\omega)\|_2^2 + \tau \|\mathbf{v}_n(\omega)\|_1 \quad (2.14)$$

for a particular $\tau > 0$. As with the standard basis pursuit solution, we can compute the unnormalized frequency-wavenumber representation $\mathbf{v}(\omega)$ using (2.13). The optimization in (2.14) is often known as basis pursuit denoising [53, 54] or the ℓ_1 penalized least-squares method [57]. In (2.14), a large τ value will cause the optimization to place greater emphasis on making $\mathbf{v}(\omega)$ sparse than fitting the linear model $\Phi \mathbf{v}_n(\omega) = \mathbf{x}_n(\omega)$. In contrast, a small τ will place greater emphasis on fitting the linear model.

Note that if τ is too small, the solution converges to a least-squares solution without sparsity, and if τ is too large, the ℓ_1 penalty dominates the cost function and the zero vector becomes the optimal solution. Therefore, it is necessary to choose an appropriate value of τ . For this dissertation, we use simulation data to investigate the performance and robustness of our approach with τ in Section 2.5. Based on these results, we choose a reasonable value of τ for our experimental data. Alternatively, there exist a number of algorithms designed to automatically obtain an optimal τ value given additional information, such as the number of non-zero values in the final solution. These techniques include the ‘‘In-Crowd’’ algorithm [58], homotopy methods [59], and the spectral projected gradient for ℓ_1 minimization approach [60].

The general effectiveness of basis pursuit for a particular sensing matrix is also often evaluated by satisfaction of the restricted isometry property (RIP) with a sufficiently small constant δ_r [48]

$$(1 - \delta_r) \|\mathbf{v}\|_2^2 \leq \|\Phi \mathbf{v}\|^2 \leq (1 + \delta_r) \|\mathbf{v}\|_2^2, \quad (2.15)$$

where the solution \mathbf{v} has r non-zero values. Testing if a given Φ satisfies the RIP condition is, in general, a combinatorial and intractable problem [61]. However, it has been shown both theoretically and numerically [62] that matrices of complex exponentials, such as our choice of Φ , with a collection of random exponents r_1, r_2, \dots, r_M , will satisfy the restricted isometry condition with high probability when M is sufficiently large [62]. These results imply that our setup is well suited for sparse recovery techniques as long as we use a sufficient number of

measurements. In Section 2.5, we show that our performance increases monotonically with the number of measurements M .

We refer to the process of obtaining this frequency-wavenumber representation $\mathbf{v}(\omega)$ for our application as sparse wavenumber analysis. Intuitively, we are transforming a frequency-distance signal into a sparse frequency-wavenumber signal. The optimization in (2.14) represents the analysis step of the transformation. In the following subsection, we discuss the synthesis step of the transformation.

2.3.2 Sparse Wavenumber Synthesis

Given a frequency-wavenumber representation $\mathbf{v}(\omega)$, we can now generate a data-driven model of Lamb waves in our plate. That is, we can use the data-derived dispersion curves in $\mathbf{v}(\omega)$ to synthesize a Lamb wave response $X(r, \omega)$ between any two points in the medium. This is accomplished by solving the forward problem in (2.3). To do this, we consider a new collection of \widehat{M} distances $\widehat{r}_1, \dots, \widehat{r}_{\widehat{M}}$ for which we want to estimate or predict $X(\widehat{r}_1, \omega), \dots, X(\widehat{r}_{\widehat{M}}, \omega)$.

We construct an $\widehat{M} \times N$ $\widehat{\mathbf{A}}$ matrix and an $\widehat{M} \times \widehat{M}$ $\widehat{\mathbf{D}}_r$ matrix based on our model from (1.6)

$$\widehat{\mathbf{A}} = \left[\sqrt{\frac{1}{\kappa_j \widehat{r}_i}} e^{j\kappa_j \widehat{r}_i} \right]_{ij} \quad (2.16)$$

$$\widehat{\mathbf{D}}_r = \text{diag} \left[\widehat{r}_1^{-1/2}, \dots, \widehat{r}_{\widehat{M}}^{-1/2} \right], \quad (2.17)$$

and synthesize the response at each frequency by a matrix-vector multiplication

$$\widehat{\mathbf{x}}(\omega) = \widehat{\mathbf{D}}_r \widehat{\mathbf{A}} \mathbf{D}_\kappa \mathbf{v}(\omega) \quad (2.18)$$

$$= \begin{bmatrix} \widehat{X}(r_1, \omega) & \widehat{X}(r_2, \omega) & \dots & \widehat{X}(r_{\widehat{M}}, \omega) \end{bmatrix}^T. \quad (2.19)$$

We refer to this process as sparse wavenumber synthesis. By coupling sparse wavenumber analysis with sparse wavenumber synthesis, we are able to actively predict wave responses between any points in the medium with only a small subset of measurements. In Section 2.5, we use sparse wavenumber synthesis to accurately predict signals in a simulated environment.

We refer to the special case $\widehat{\mathbf{D}}_r \widehat{\mathbf{A}} \cong \mathbf{D}_r \mathbf{A}$, for which the synthesized signals correspond to the same distances as the measured signals used in sparse wavenumber analysis, as *sparse wavenumber denoising*. Intuitively, the “noise” in the each measurement, whether from random

measurement noise or multipath reflections, is not sparse in the frequency-wavenumber domain. As a result, the recovered frequency-wavenumber representation is robust to these errors and any signal synthesized from this representation shows a significant reduction in this “noise.” In Section 2.5 and Section 2.6, we show that sparse wavenumber denoising effectively removes multipath interference in both simulated and experimental data.

2.3.3 Debiasing Results

It is known that for sufficiently large values of τ , the basis pursuit denoising solution $\mathbf{v}(\omega)$ has a multiplicative bias [63, 64]. This is a result of the $\tau\|\mathbf{v}_n(\omega)\|_1$ term’s dependence on the scalar amplitude of $\mathbf{v}_n(\omega)$ as well as its sparsity. To unbiased our results, we estimate the multiplicative term at each frequency $\mu(\omega)$ by minimizing the squared error between the measured signal $\mathbf{x}(\omega)$ and its sparse wavenumber denoised version

$$\begin{aligned}\hat{\mu}(\omega) &= \arg \min_{\mu(\omega)} \|\mu(\omega)\mathbf{D}_r\mathbf{A}\mathbf{D}_\kappa\mathbf{v}(\omega) - \mathbf{x}(\omega)\| \\ &= \frac{(\mathbf{v}(\omega))^H\mathbf{D}_\kappa\mathbf{A}^H\mathbf{D}_r\mathbf{x}(\omega)}{\|\mathbf{D}_r\mathbf{A}\mathbf{D}_\kappa\mathbf{v}(\omega)\|_2^2},\end{aligned}\tag{2.20}$$

where $(\cdot)^H$ denotes the Hermitian transpose of a vector or matrix. The multiplicative bias estimate $\hat{\mu}(\omega)$ is then multiplied to the frequency-wavenumber representation $\mathbf{v}(\omega)$ and any synthesized signals $\hat{\mathbf{x}}(\omega)$.

2.3.4 Comparison with least squares

In Section 2.5, we compare sparse wavenumber analysis to a similar and computationally fast least-squares approach. In the least-squares formulation, we solve the optimization

$$\mathbf{v}_{\ell_2}(\omega) = \arg \min_{\mathbf{v}(\omega)} \|\mathbf{v}(\omega)\|_2\tag{2.21}$$

$$\text{s.t. } \mathbf{x} = \mathbf{A}\mathbf{v}(\omega).\tag{2.22}$$

As a closed form solution, $\mathbf{v}_{\ell_2}(\omega)$ is represented as

$$\mathbf{v}_{\ell_2}(\omega) = (\mathbf{A}^H\mathbf{A})^{-1}\mathbf{A}^H\mathbf{x}(\omega) = \mathbf{A}^\dagger\mathbf{x}(\omega).\tag{2.23}$$

Note that this approach is similar to sparse wavenumber analysis with $\tau = 0$ since neither promote sparsity. For the special case in which the discretized wavenumbers $\kappa_1, \kappa_2, \dots, \kappa_N$ are

uniformly sampled with spacing $\Delta\kappa$, the sensors are aligned in a uniform, linear array with spacing

$$\Delta d = \frac{2\pi}{N\Delta\kappa} , \quad (2.24)$$

and the excitation source is positioned on either end of the linear array, we can represent our sensing matrix as

$$\mathbf{A} = \mathbf{F}^H \mathbf{D}_F^H \quad (2.25)$$

$$\mathbf{A}^\dagger = \frac{1}{N} \mathbf{D}_F \mathbf{F} \quad (2.26)$$

$$\mathbf{D}_F = \text{diag} \left(e^{-j\kappa_1 \Delta d}, \dots, e^{-j\kappa_N \Delta d} \right) . \quad (2.27)$$

In (2.25), the matrix \mathbf{F}^H is an unnormalized $M \times N$ Fourier frame, representing the M -point inverse discrete Fourier transform, and \mathbf{F} is an $N \times M$ Fourier frame, representing the N -point discrete Fourier transform.

In nondestructive testing, this special case is commonly used to measure dispersion curves and is known as the two-dimensional discrete Fourier transform (2D-DFT) technique [10, 12]. However, given a small number of sensors, the 2D-DFT will generally perform poorly at recovering the frequency-wavenumber representation [20]. This is because when P sensors are arranged as a uniform, linear array, the number of unique distances that can be measured is $M = P - 1$. In contrast, an arbitrary sensor topology can measure $M = P(P - 1)/2$ unique distances. Furthermore, in guided wave structural health monitoring, we often desire the sensors to be sparsely separated in space to interrogate areas that are as large as possible. For these reasons, we choose to compare our approach to the more general least squares solution in (2.23). In Section 2.5, we show that, by promoting sparsity, we can achieve much more reliable results through our proposed sparse wavenumber analysis.

Time-frequency approaches [13–15] and related time domain matching pursuit methods [16, 17] are also often used to estimate dispersion curves. However, these techniques are not inherently array-based, and as a result, are usually sensitive to noise and multipath interference. Furthermore, these approaches compute, as a function of frequency, the group velocity rather than wavenumber or phase velocity. Since the group velocity is defined as the derivative of

frequency with respect to wavenumber $v_g = d\omega/dk$, the wavenumber or phase velocity representations may not be recoverable from the group velocity. For these reasons, we do not consider these techniques to be appropriate comparisons to our approach.

Note that since sparsity is not promoted by either the least-squares and time-frequency approaches, additional algorithms are often necessary to extract each dispersion curve from the respective representations [65]. We illustrate this in Section 2.5. This introduces an additional layer of complexity with an additional possibility of error.

2.4 Simulation and experimental methodologies

In this section, we discuss the setup of our experiment and simulations, which are designed similarly. We consider a 0.284 cm thick, 1.22 m long, and 1.22 m wide aluminum plate. On the surface of the plate, we consider a collection of 17 randomly placed sensors. Sensors are placed randomly so that no particular bias affects the results and because, as discussed in Section 2.2, random sensor placement assures good properties for sparse recovery methods [62]. Figure 2.2 illustrates the location of each sensor. In our experiment, each sensor is a 0.7 cm by 0.8 cm by 0.2 mm PZT (lead zirconate titanate) transducer permanently bonded to the plate's surface. In

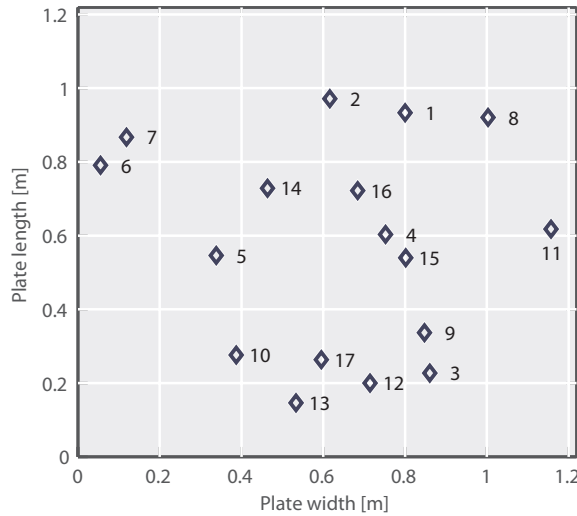


Figure 2.2: Sensor configuration for simulations and experiments. Each diamond represents the location of a sensor on a plate.

our simulations, we model each sensor as an ideal point source or point receiver.

To recover $\mathbf{v}(\omega)$ with basis pursuit denoising in (2.14), we use CVX, a MATLAB package specifically designed for solving convex programs [66, 67].

2.4.1 Data collection

To collect data, we iteratively use each sensor as the transmitter and receive signals from each of the other 16 transducers. This provides a total of 272 different measurements for 136 unique distances. For the scenarios considered in this chapter, we found that the performance of our methods to be approximately equal whether analyzing all of the measurements or only a subset with unique distances. However, if the data is corrupted by significant random measurement noise, the additional redundant information can improve recovery performance.

During each iteration, the transmitter excites a $10\ \mu\text{s}$ linear frequency modulated chirp with a 3 dB bandwidth between 0 Hz and 2 MHz. This excitation choice allows us to collect data across a large band of frequencies. Through filtering we can extract narrowband signals, which are more commonly used in structural health monitoring, from each chirp response [68]. At the receivers, which are synchronized with the transmitter, waves are measured with a 10 MHz sampling rate.

2.4.2 Simulation models

To simulate the wave response in the plate, we use the “far-field” model in (1.6). The dispersion curves shown in Figure 1.1 are used to determine $k_m(\omega)$ for each mode m . To aid the presentation of our results in Section 2.5, we choose the complex amplitudes to be $G_m(\omega) = 1$ for each m . This allows us to more readily compare the amplitudes in frequency-wavenumber plots.

We simulate multipath signals by computing the distance of the direct path as well as each additional path that travels from the transmitter to the receiver after reflecting from any boundary once. Each boundary is considered to be ideally rigid such that all incident wave energy is reflected but the amplitude is negated. Although this model only approximates the first few multipath signals in the experiment, it is sufficient to illustrate the effects of multipath on our methodology.

We also simulate errors arising from modeling physical sensors, which extend in space, as point sources and point receivers. We simulate this error by introducing uncertainty in each sensor’s position. We accomplish this by perturbing the expected distances r_1, \dots, r_M for each measurement with “position noise” so that we have inaccurate knowledge about the each sensor’s location. While this does not perfectly model the effect of each sensor’s geometry on the measured responses, it illustrates the approximate effect of such error on our methodology.

2.5 Simulation results and discussion

Through simulations, we evaluate the effectiveness of our methodology for four different scenarios. 1) We consider only the direct path signal, which simulates a plate of unbounded length and width. 2) We consider both the direct and multipath arrivals as described in Section 2.4. 3) We consider the direct and multipath arrivals but window the signals to remove late arrivals. 4) We again consider the windowed, multipath responses but with additional sensor position noise.

In the third and fourth scenarios, we apply a rectangular window with an exponential taper. The exponential taper begins after the arrival of a hypothetical signal traveling at a group velocity of 2000 m/s and reduces the signal amplitude by 3 dB after approximately 28 μ s. Note that this window removes multipath as well as some direct arrivals. In the fourth scenario, we perturb each expected distance by a uniform random variable with a horizontal range from -0.3 cm to 0.3 cm and a vertical range from -0.35 cm to 0.35 cm, which is 0.1 cm smaller in each dimension than each 0.7 cm by 0.8 cm transducer.

Figure 2.3 depicts a single time response between sensor 14 and sensor 6, as denoted in Figure 2.2, for the third simulation scenario and our experiment. The top traces represent the simulated or measured wideband signals and the bottom traces represent the same signals filtered by a narrowband Gaussian filter with a center frequency of 180 kHz and a 120 kHz bandwidth. As verification of our simulations, we note that simulated narrowband signals in Figure 2.3(c) and experimental narrowband signals in Figure 2.3(d) have a close resemblance.

We evaluate our four simulation scenarios with three different metrics: 1) the visual accuracy of the recovered dispersion curves in the frequency-wavenumber domain, 2) the accuracy of




Figure 2.3: Wideband and narrowband time responses between sensors 14 and 6, labeled in Figure 2.2. Each narrowband response is a Gaussian filtered response with center frequency of 180 kHz and 120 kHz bandwidth. (a) Wideband simulation of a plate with windowed, multipath effects. (b) Wideband experimental measurement with window. (c) Narrowband simulation of a plate with windowed, multipath effects. (e) Narrowband experimental measurement with window.

synthesizing a noise-free, direct path response with sparse wavenumber denoising, and 3) the accuracy of predicting the direct path responses between random points on the plate with sparse wavenumber synthesis.

For the first metric, we recover the frequency-wavenumber representation $\mathbf{v}(\omega)$ of our data by applying sparse wavenumber analysis. We uniformly discretize the wavenumber space across 3000 samples between 0 m^{-1} and 4500 m^{-1} . We also consider 2000 frequencies uniformly spaced between 0 Hz and 2 MHz. In Figure 2.4, we show the resulting frequency-wavenumber representations for each scenario. Note that we apply a Gaussian blur and unsharp mask filter [69] to improve visibility for the reader. We also show a magnified segment of the A0 mode centered at 180 kHz to illustrate that each curve is approximately one or two pixels wide. For convenience, the figures are normalized so that the maximum value is associated with 0 dB.

For the second metric, we use $\mathbf{v}(\omega)$ with sparse wavenumber denoising to synthesize the direct path response between sensor 14 and sensor 6, as denoted in Figure 2.2, and filter the results around a center frequency of 180 kHz and a 120 kHz wide bandwidth. In Figure 2.5,

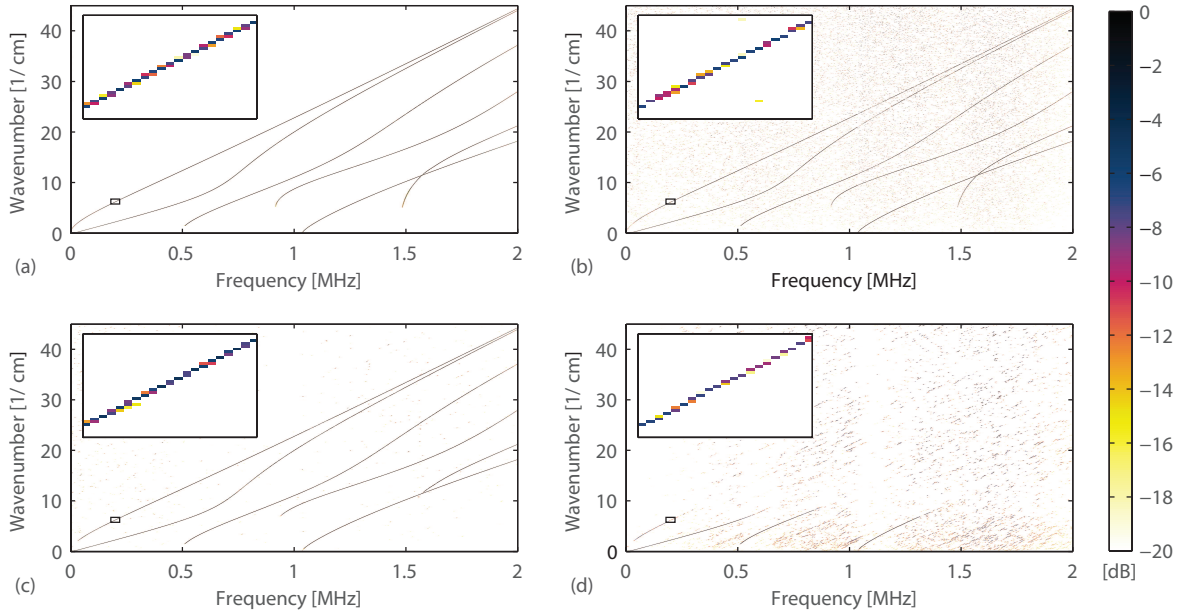


Figure 2.4: The sparse wavenumber analysis recovered frequency-wavenumber magnitudes $|\mathbf{v}(\omega)|$ for four simulations: (a) unbounded plate, (b) plate with multipath effects, (c) plate with a windowed response and multipath effects, (d) plate with a windowed response, multipath effects, and sensor location uncertainty.

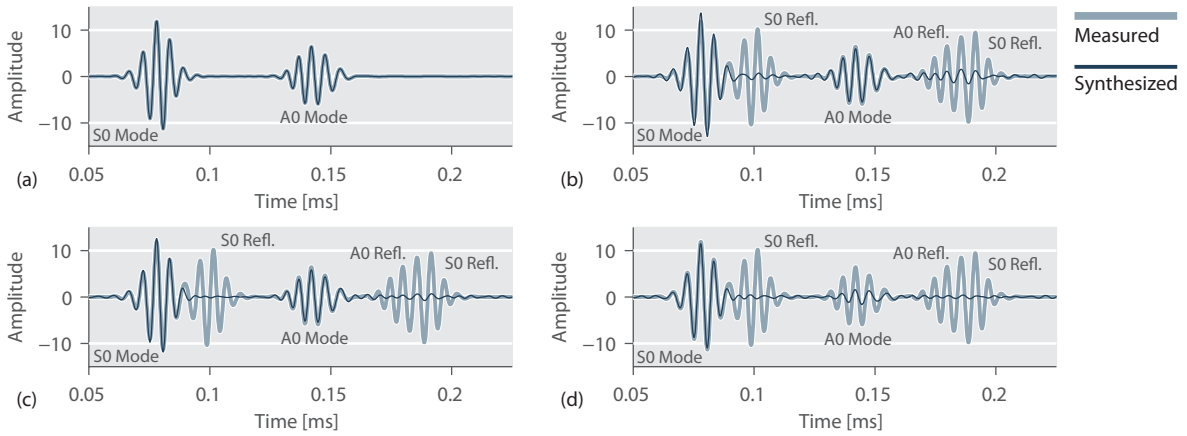


Figure 2.5: Narrowband direct path responses synthesized by sparse wavenumber denoising (thin, dark line) and corresponding narrowband measured responses (broad, light line) between sensor 14 and sensor 6 for four simulations: (a) unbounded plate, (b) plate with multipath effects, (c) plate with a windowed response and multipath effects, (d) plate with a windowed response, multipath effects, and sensor location uncertainty.

we compare these synthesized responses (shown as a thin, dark curve) with the corresponding measured responses (shown as a broad, light curve).

For the third metric, we use $\mathbf{v}(\omega)$ with sparse wavenumber synthesis to predict the direct path responses corresponding to 1000 randomly chosen pairs of points on the plate. We compare the synthesized signals $\hat{\mathbf{x}}(\omega)$ with the true signals $\mathbf{x}(\omega)$ by computing the correlation coefficient σ across every frequency and measurement

$$\sigma = \frac{\sum_{q=1}^Q \mathbf{x}(\omega_q)^H \hat{\mathbf{x}}(\omega_q)}{\sqrt{\sum_{q=1}^Q \|\mathbf{x}(\omega_q)\|_2^2} \sqrt{\sum_{q=1}^Q \|\hat{\mathbf{x}}(\omega_q)\|_2^2}}. \quad (2.28)$$

In (2.28), Q is the number of discrete frequencies being analyzed. When processing a windowed data set, we also window the true signal in the same manner. We compute the correlation coefficient as a function of the number of sensors, uniformly spanning from 5 to 17, and a function of the basis pursuit denoising parameter τ , logarithmically spanning from $\tau = 0.004$ to $\tau = 10$. To ease the computational cost, we recover $\mathbf{v}(\omega)$ for 200 different frequencies, uniformly spanning from 0 Hz to 2 MHz, instead of 2000 frequencies used for the other metrics.

2.5.1 Scenario 1: Unbounded plate

Sparse wavenumber analysis

Figure 2.4(a) illustrates the magnitude of the frequency-wavenumber representation $|\mathbf{v}(\omega)|$ for the simulated unbounded plate with $\tau = 0.01$. In the figure, the non-zero values in the frequency-wavenumber representation overlap very well with the true dispersion curves.

Sparse wavenumber denoising

Figure 2.5(a) compares the synthesized and measured narrowband response between sensor 14 and sensor 6 on the unbounded plate. In the plot, the estimated and measured responses are found to be nearly identical. The two curves have a correlation coefficient of approximately 1.000.

Sparse wavenumber synthesis

Figure 2.6(a) illustrates our capability to predict the wideband response of random locations in the simulated, unbounded plate. For a large range of τ values, up to approximately $\tau =$

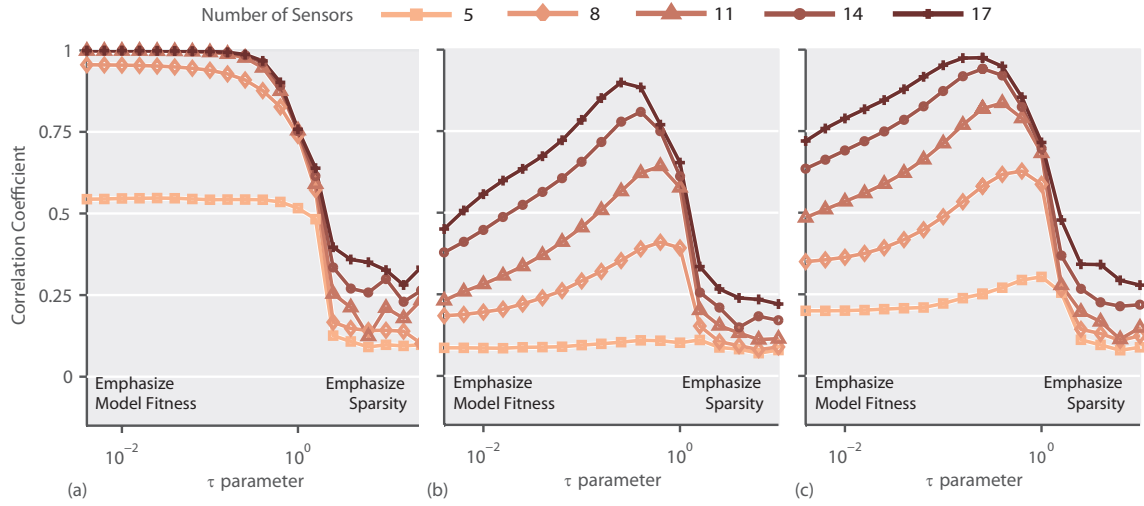


Figure 2.6: The correlation coefficient between 1000 randomly chosen responses generated from sparse wavenumber synthesis and the corresponding true responses for 3 different simulation scenarios: (a) unbounded plate, (b) plate with multipath effects, (c) plate with windowed, multipath effects.

0.2512, we achieve correlation coefficients greater than 0.99 for $M \geq 11$. Note that as we approach $\tau = 0$, we approach a least squares estimate, which does not emphasize any sparsity. So for some sufficiently small τ , $\mathbf{v}(\omega)$ will eventually lose its sparse structure. We illustrate a similar least squares solution at the end of this section.

Overall, the first simulation scenario shows our methodology to be very accurate and robust in noise-free conditions.

2.5.2 Scenario 2: Multipath

Sparse wavenumber analysis

Figure 2.4(b) depicts the frequency-wavenumber representation magnitude $|\mathbf{v}(\omega)|$ for the simulated plate with multipath and a basis pursuit parameter of $\tau = 0.25$. The sparse curves in Figure 2.4(b) again overlap well the true dispersion curves. However, the additional multipath generates an abundant amount of speckle-like noise in the figure.

Sparse wavenumber denoising

Figure 2.5(b) shows the measured and synthesized narrowband responses between sensor 14 and sensor 6 in a simulated plate with multipath interference. In the figure, our denoised signal very accurately reconstructs the direct path modes at approximately $80 \mu\text{s}$ and $140 \mu\text{s}$ while only weakly reconstructing the multipath responses at $100 \mu\text{s}$, $180 \mu\text{s}$, and $190 \mu\text{s}$. The synthesized response in Figure 2.5(b) has correlation coefficient of 0.966 with the true direct path response.

Sparse wavenumber synthesis

Figure 2.6(b) shows the correlation coefficient between the wideband responses predicted by sparse wavenumber synthesis and true responses of 1000 random locations in the simulated plate with multipath. For $M = 17$, we achieve correlation coefficients greater than 0.75 for approximately $0.1 < \tau < 0.631$. For $\tau = 0.2512$, we achieve a correlation coefficient of 0.899.

Therefore, the second simulation scenario shows that our methodology remains accurate in the presence of additive multipath interference.

2.5.3 Scenario 3: Multipath with window

Sparse wavenumber analysis

Figure 2.4(c) shows the magnitude of the frequency-wavenumber representation $|\mathbf{v}(\omega)|$, with $\tau = 0.25$, for the simulated plate with multipath and windowed responses. In general, the recovered curves in $\mathbf{v}(\omega)$ overlap with the true dispersion curves very well. Due to the window, a small number of frequency-wavenumber pairs with large slopes (small group velocities) are no longer visible in our estimate. However, for most frequency-wavenumber pairs, we observe much less noise relative to Figure 2.4(b).

Sparse wavenumber denoising

Figure 2.5(c) compares the measured and synthesized narrowband and windowed responses with multipath interference between sensor 14 and sensor 6. Compared with results in Figure 2.5(b), the synthesized response is noticeably less influenced by the multipath responses at $100 \mu\text{s}$, $180 \mu\text{s}$, and $190 \mu\text{s}$. The correlation coefficient between the denoised response in

Figure 2.5(c) and the true direct path was found to be 0.992, a significant improvement over the 0.966 correlation coefficient achieved without the window.

Sparse wavenumber synthesis

Figure 2.6(c) shows the accuracy of sparse wavenumber synthesis for predicting the wide-band, windowed responses corresponding to 1000 random locations in the simulated plate with multipath. Compared with the results in Figure 2.6(b), we observe a significant improvement in performance and robustness. For $M = 17$, we observe correlation coefficients greater than 0.75 for approximately $0.00631 < \tau < 0.5$. For $\tau = 0.2512$, we achieve a correlation coefficient of 0.975, a significant improvement over 0.899 achieved without the window.

Therefore, the third simulation scenario shows that utilizing a window significantly improves our technique's accuracy and robustness in the presence of additive multipath interference.

2.5.4 Scenario 4: Multipath and position noise with window

Sparse wavenumber analysis

Figure 2.4(d) shows the magnitude of the frequency-wavenumber representation $|\mathbf{v}(\omega)|$, with $\tau = 0.25$, of the simulated plate with multipath, windowed responses with sensor position noise. In the figure, our frequency-wavenumber curves follow the true dispersion curves closely up until wavenumbers between 700 m^{-1} and 1200 m^{-1} . Between these values, the amplitude of each mode declines. The wavelength range corresponding to these wavenumbers is 0.5 cm to 0.9 cm. This corresponds roughly to the size of the magnitude of the sensor position noise, which perturbs each expected sensor location by at most ± 0.461 cm from their original true location.

For this test, we use a basis pursuit denoising parameter of $\tau = 0.25$ to compare with simulation scenarios 2 and 3. We note, however, that due to sensor position noise, the data incurs larger errors at higher frequencies, which have a greater number higher wavenumber components. Therefore, the optimal choice of τ gradually increases with frequency from $\tau = 0.25$ to $\tau = 0.6$. However, our tests indicate that the overall performance does not vary significantly within this range of τ . Visually, increasing τ will help to reduce the noise in Figure 2.4.

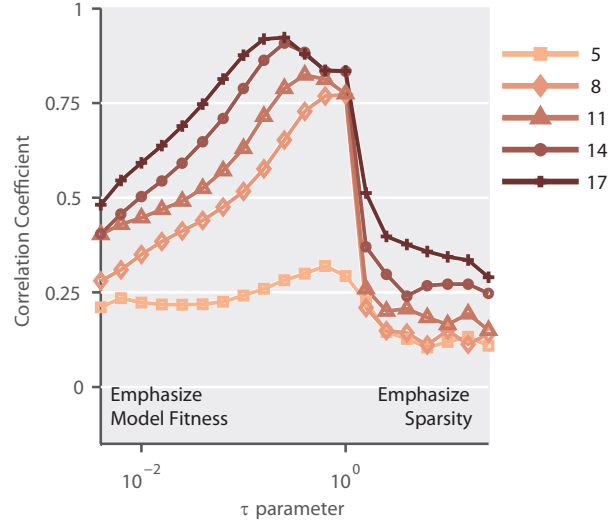


Figure 2.7: The correlation coefficient between 1000 randomly chosen responses generated from sparse wavenumber synthesis and the corresponding true responses for a simulation of a plate with windowed, multipath effects and sensor position noise.

Sparse wavenumber denoising

Figure 2.5(d) shows the measured and synthesized narrowband responses between sensor 14 and sensor 6 in the presence of multipath interference and position noise. We observe a strong correlation between the measured and synthesized responses. However, we also observe a significant reduction on the A0 mode's amplitude. This can be attributed to the A0 mode's higher wavenumber at 180 kHz. The denoised response has a correlation coefficient of 0.9144 with the true direct path, windowed response.

Sparse wavenumber synthesis

Figure 2.7 shows our accuracy of predicting the windowed responses of 1000 random locations from multipath data with sensor position noise. However, unlike Figure 2.6(a), Figure 2.6(b), and Figure 2.6(c), in which we considered the full wideband signal, we now consider only the response in the narrow band of frequencies between 120 kHz and 240 kHz. We choose these frequencies since they are not significantly affected by the thresholding effect across wavenumber, shown in Figure 2.4(d). For $M = 17$, we observe correlation coefficients greater than 0.75 for $0.03981 < \tau < 1.0$. For $\tau = 0.2512$, we achieve a correlation coefficient

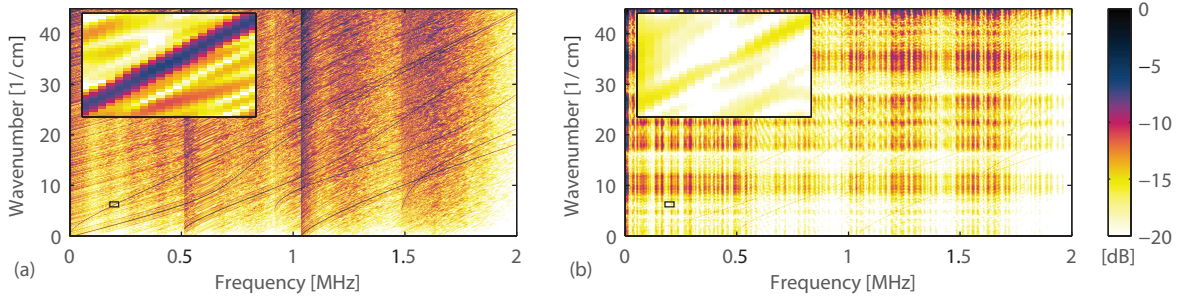


Figure 2.8: The magnitude of the least-squares frequency-wavenumber representation of simulation data of an (a) unbounded plate and (b) plate with with a windowed response and multipath effects.

of 0.926.

Our final simulation scenario suggests that while we may be restricted by the geometry of our sensors, we can still accurately recover the frequency-wavenumber representation of our data, remove multipath interference, and predict Lamb wave responses across certain wavenumber ranges. Across all of the test scenarios, we observe reasonable results for values τ in the neighborhood of $0.1 < \tau < 0.6$. We find this to be a reasonable range for our experiment as well.

2.5.5 Comparison with least squares

Figure 2.8(a) and Figure 2.8(b) illustrates the least-squares, or generalized 2D-DFT, frequency-wavenumber representation of an unbounded plate simulation (scenario 1) and a simulation with a windowed multipath response (scenario 3), respectively. As the figures show, the least-squares representations do not represent well the true dispersion curves of the medium. In Figure 2.8(a), the dispersion curves are visible but with a significant amount of additional “noise”, which physically does not represent any property of the medium. After introducing multipath reflections in Figure 2.8(b), the dispersion curves become significant overpowered by this noisy information.

Overall, since the least-squares solution does not promote sparsity in the model, it does not recover the medium’s frequency-wavenumber representation with high precision. As a result, our signal synthesis and denoising approaches perform poorly. This is supported by Figure 2.6 since the least squares solution is similar to sparse wavenumber analysis for $\tau = 0$.

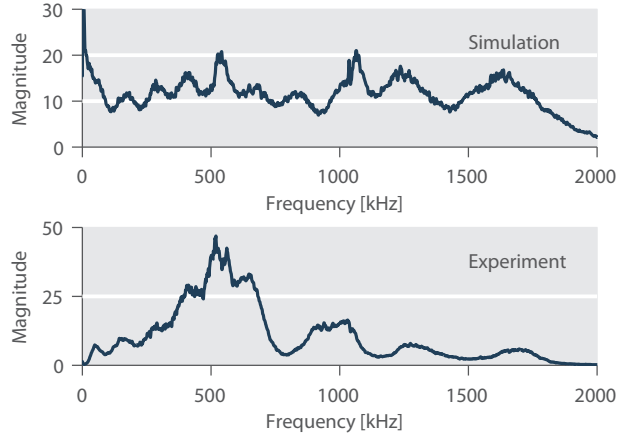


Figure 2.9: The average frequency magnitude responses for a simulation with windowed multipath effects and experimental data.

2.6 Experimental results and discussion

In this section, we apply sparse wavenumber analysis and sparse wavenumber denoising to experimental data. For this data, we apply the same window as in the third simulation scenario in Section 2.5. We also apply a window to remove the first $10 \mu\text{s}$ of each signal to eliminate any electrical cross-talk. We compute the basis pursuit denoising solution in (2.14) using 3000 uniformly spaced samples in the wavenumber domain across between 0 m^{-1} and 4500 m^{-1} . We sample across 2000 frequencies from 0 Hz to 2 MHz.

In the simulation results, we achieved excellent performance with a basis pursuit denoising parameter τ around the neighborhood of 0.25. However, we anticipate a greater degree of noise in the experimental data, so we choose to increase τ to 0.50. Testing has shown both $\tau = 0.25$ and $\tau = 0.50$ achieve good results, but $\tau = 0.5$ achieves greater robustness to multipath interference.

In contrast with the simulations, the experimental data is also affected by each sensor's frequency response, which have natural bandpass characteristics. Figure 2.9 illustrates the average magnitude frequency response for the windowed experimental data and windowed simulation data. In the figure, the experimental results are most prominent around 500 kHz whereas the simulation data is relatively flat across the entire range of interest.

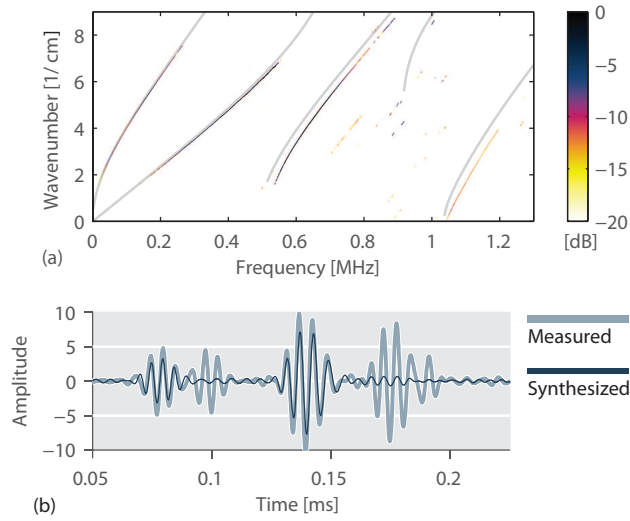


Figure 2.10: (a) The magnitude of the recovered frequency-wavenumber representation $|\mathbf{v}(\omega)|$ for the data with light gray lines illustrating the theoretical dispersion curves for the experiment. (b) The synthesized narrowband direct path responses generated from sparse wavenumber de-noising (thin, dark line) and corresponding narrowband measured responses (broad, light line) between sensor 14 and sensor 6 for our experiment.

2.6.1 Sparse wavenumber analysis

Figure 2.10(a) shows the magnitude of the frequency-wavenumber representation $|\mathbf{v}(\omega)|$ for our experimental data. As with the simulations, we use a Gaussian blur and unsharp mask filter [69] to broaden lines and ease visibility for the reader. Each estimated curve is generally only one or two pixels wide. We also crop our results to focus on the frequency-wavenumber region of most interest. Frequencies greater than 1300 kHz and wavenumbers greater than 900 m^{-1} are only characterized by weak noise.

In Figure 2.10(a), we overlay the theoretical dispersion curves used in our simulations in light gray. The results show that the frequency-wavenumber values follow similar trends as their corresponding dispersive curves, although the curves do not exactly overlap. However, this is expected since the theoretical dispersion curves are derived from approximate values for the material properties of the aluminum plate.

Similar to the simulations with sensor position error in Figure 2.10(d), the experimental data's frequency-wavenumber representation shows a rapid decrease in magnitude for wavenum-

bers between 700 m^{-1} and 900 m^{-1} . This range corresponds to wavelengths in the range of 0.7 cm to 0.9 cm. As with the simulations, these wavelength are on the order of the transducer's dimensions, which are 0.7 cm by 0.8 cm.

We also observe in Figure 2.10(a) that the S1 mode is never visible and the A2 mode vanishes for frequencies larger than 1200 kHz. Decreasing the lower bound on the magnitudes in Figure 2.10(a) does not reveal any additional structure. We note, however, that the S1 mode has a small group velocity around 900 kHz that is also not visible in our simulations in Figure 2.4(d). Additionally, the measured frequency spectrum in Figure 2.9 shows weak magnitudes for frequencies greater than 1100 kHz.

If we apply the least-squares solution in (2.23) to the experimental data, we observe a result similar to the simulation in Figure 2.8(b). Since the least-squares result does not promote sparsity, its frequency-wavenumber domain is dominated by a significant amount of “noise” and the dispersion curves cannot be clearly identified visible.

2.6.2 Sparse wavenumber denoising

Figure 2.10(b) compares the measured narrowband, experimental response between sensor 14 and sensor 6, as denoted in Figure 2.2, with its denoised response. The responses are shown for a band of frequencies from 120 kHz to 240 kHz. The results show that our method successfully removes the multipath interference from the experimental data. Although we observe a small reduction in the amplitude of the S0 and A0 modes at $80 \mu\text{s}$ and $140 \mu\text{s}$, respectively, we observe a significant reduction in the multipath responses at $100 \mu\text{s}$, $180 \mu\text{s}$, and $190 \mu\text{s}$.

Overall, our experimental results verify our simulations results. We computed the frequency-wavenumber representation for experimental Lamb wave data and demonstrated that our experiments can be reasonable well modeled by with additive multipath interference and sensor position noise. We also used the frequency-wavenumber representation to successfully remove multipath components from the data.

2.7 Conclusion

In this chapter, we presented a compressed sensing based methodology, which we call sparse wavenumber analysis, to accurately recover the frequency-wavenumber representation of Lamb

waves. Using simulation and experimental data, we showed that we could use this representation to accurately represent Lamb wave dispersion curves. We then utilized this frequency-wavenumber representation to remove multipath components from simulated and experimental data, through a method we refer to as sparse wavenumber denoising, and to predict Lamb wave responses in a simulated plate, through a method we term sparse wavenumber synthesis.

We tested our methodology with 4 simulations of increasing complexity. We showed that, even with data corrupted by multipath interference and modeling error, we can very accurately recover the frequency-wavenumber representation of the medium. We quantified this accuracy by comparing 1000 synthesized signals corresponding to 1000 randomly chosen pathes, generated using sparse wavenumber synthesis, with their true direct path responses. We showed that, with only multipath interference, we could achieve a correlation coefficient of 0.975 between the true and estimated responses across a wide range of frequencies from 0 MHz to 2 MHz. After incorporating sensor position noise, which simulates errors from each sensor's extended geometry, we could still achieve a correlation coefficient of 0.926 between the true and estimated responses across a narrow band of frequencies from 120 kHz to 240 kHz.

We applied sparse wavenumber analysis to accurately recover the frequency-wavenumber representation of experimental data. We showed the recovered frequency-wavenumber representation closely followed theoretical dispersion curves. We then used this representation to accurately remove the multipath components from the experimental time traces. Furthermore, we visually demonstrated a good correspondence between our experimental and simulation results.

Leveraging Guided Wave Complexity: Data-Driven Matched Field Processing

3.1 Motivation

Matched field processing is a general framework that utilizes models of complex propagation environments to estimate the location of targets with high resolution and accuracy [1, 2]. In underwater acoustics, matched field processing has been studied extensively to improve the performance of systems operating in acoustically complex underwater environments [2, 70–74]. It has also been applied to problems in seismology [75], radar and electromagnetic propagation [76, 77], and nondestructive evaluation [78, 79]. Matched field processing is also closely related to time reversal processing [26, 27], which utilizes direct measurements, rather than a model, to learn the environment and improve localization performance. In guided wave structural health monitoring, where guided waves are used to detect and locate damage in large structures, matched field processing is an attractive tool due to the complex multimodal and dispersive characteristics of these propagation environments.

To implement matched field processing, we require a model of the responses that we expect to measure from a given target location. Historically, guided wave structural health monitoring localization methods have used a single mode model with constant and equal group and phase velocities [24, 25]. As a result, significant preprocessing is often necessary to simplify the measured data to approximately match the model. This is accomplished by filtering measurements

around a narrow band of frequencies to isolate a single mode and then extracting the signal's envelope to remove phase distortions due to differences in the wave's group and phase velocities. Due to these simplifications, these approaches are prone to poor resolutions and large artifacts.

In underwater acoustics, a very different approach is taken. The model is instead carefully constructed by formulating and solving, numerically, the wave equation for the geometry and characteristics of the medium [1]. This approach can provide very accurate and highly resolved localization when the model is correct, but very poor localization performance when the model is incorrect [80]. Therefore, this approach requires significant knowledge of the material properties and geometry of the medium, which are often not precisely known. Although researchers have developed algorithms [81,82] to estimate these properties from calibration data, estimating these parameters generally requires solving nonlinear and nonconvex optimization problems, such as simulated annealing [83] and Monte Carlo sampling [81, 83], across many unknown parameters. For these reasons, the localization process is often computationally expensive and unreliable.

This chapter presents a new alternative approach. We assume that the wave equation for the environment of interest is solved by a modal solution. We construct the model directly from our calibration data through a method known as sparse wavenumber analysis [49]. We refer to the model as a sparsity-based, data-driven model because it is constructed from the sparse properties of the measured data, which are recovered by using tools from compressed sensing [53]. We then integrate this data-driven model with matched field processing to derive a sparsity-based, data-driven matched field processing methodology.

Unlike the signal simplification strategies that assume the waves to have a single velocity, our approach allows us to use all of the modes and dispersion in the measurements to improve localization accuracy and resolution. Unlike model-based parameter estimation strategies, sparse wavenumber analysis is performed using basis pursuit optimization [53, 54], a convex process. Therefore, the resulting model is globally optimal and can be computed quickly and reliably. This process can be applied to many problems with modal solutions, such as multilayered plates [84], pipe waves [6], Rayleigh waves [7], and Pekeris shallow water channels [8]. In the following two chapters, we specifically consider Lamb wave data from an aluminum plate.

In this chapter, we develop and analyze data-driven matched field processing by integrating sparse wavenumber analysis with matched field processing. This work was originally presented in references [50, 85, 86]. We integrate sparse wavenumber analysis with both the coherent matched field processor [74, 87] and the incoherent matched field processing [1], both of which are commonly applied in the literature. We also analytically deriving the asymptotic behavior of data-driven matched field processing and show that it provides accurate localization even in the presence of unmodeled multipath interference. In Chapter 4 and Chapter 5, we respectively apply the new coherent data-driven matched field processor and the incoherent data-driven matched field processor to localize a passive scatterer and an acoustic emission event on an aluminum plate.

3.2 Data collection

We consider three collections of data: calibration data, test data, and model data for our data-driven matched field processing methodology. Each collection of data is represented at Q discrete frequencies, $\omega_1, \omega_2, \dots, \omega_Q$. The calibration data represents physical measurements of waves that have travelled known distances. The test data represents physical measurements, potentially after removing baseline information, that have travelled unknown distances and potentially interacted with a scatterer in the medium. The model data represents synthesized signals, generated from the calibration data, that predict what the test data will look like given a scatterer at a chosen location.

3.2.1 Calibration data

We represent the calibration data at frequency ω_q as a $D \times 1$ vector

$$\begin{aligned} \mathbf{y}_q &= [X(d_1, \omega_q) \cdots X(d_D, \omega_q)]^T + \mathbf{n}_q^{(y)} \\ &= \mathbf{\Phi}(\mathbf{d})\mathbf{v}_q + \mathbf{n}_q^{(y)}, \end{aligned} \quad (3.1)$$

where D is the number of signals measured. The calibration data \mathbf{y}_q represents the signals transmitted and received between pairs of sensors with known distances $\mathbf{d} = [d_1, \dots, d_D]^T$ between them. The vector $\mathbf{n}_q^{(y)}$ represents errors in our measurements due to random noise

or unmodeled multipath interference from boundaries that is not included in the ideal plate framework in (1.6).

The $D \times N$ matrix $\Phi(\mathbf{d})$ represents a linear mapping between the measured frequency-distance signal \mathbf{y}_q and a discretized frequency-wavenumber, or dispersion curve, representation \mathbf{v}_q . The matrix $\Phi(\mathbf{d})$ can be derived from (1.6) as a matrix of complex exponentials [49]

$$\begin{aligned}\Phi(\mathbf{d}) &= \rho(\mathbf{d}) \left[d_i^{-1/2} e^{-j\kappa_j d_i} \right]_{ij} \\ \rho(\mathbf{d}) &= \left(\sum_{m=1}^D |d_m|^{-1} \right)^{-1/2}.\end{aligned}\tag{3.2}$$

The normalization constant $\rho(\mathbf{d})$ is included so that each column of $\Phi(\mathbf{d})$ has a unit ℓ_2 norm. Each element of \mathbf{v}_q is a weight assigned to each column, or basis vector, in $\Phi(\mathbf{d})$. Each column in $\Phi(\mathbf{d})$ then corresponds to one of N possible wavenumbers $\kappa_1 \dots \kappa_N$. The minimum number of wavenumbers N required to fully represent \mathbf{y}_q is equal to the number of wave modes in the data, i.e., $\kappa_1 = k_1(\omega_q), \dots, \kappa_N = k_N(\omega_q)$. If we choose N to be much larger than the number of wave modes in the signal, all of the additional elements in \mathbf{v}_q that do not correspond to a mode will have values of zero.

3.2.2 Test data

We represent the test data at frequency ω_q as an $M \times 1$ measurement vector

$$\begin{aligned}\mathbf{x}_q &= [X(r_1^*, \omega_q) \dots X(r_M^*, \omega_q)]^T + \mathbf{n}_q^{(x)} \\ &= \Phi(\mathbf{r}^*) \mathbf{v}_q + \mathbf{n}_q^{(x)}.\end{aligned}\tag{3.3}$$

In (3.3), each value in the vector $\mathbf{r}^* = [r_1^*, \dots, r_M^*]^T$ represents the true, unknown distance travelled from a transmitter to a scatterer and then to a receiving sensor. The matrix $\Phi(\mathbf{r}^*)$ is defined the same as in (3.2) and the vector $\mathbf{n}_q^{(x)}$ represents error in the data due to random noise or unmodeled multipath interference.

3.2.3 Model data

We represent the model data at frequency ω_q and for chosen distances $\mathbf{r} = [r_1, \dots, r_M]^T$ as an $M \times 1$ vector function

$$\hat{\mathbf{x}}_q(\mathbf{r}) = \left[\hat{X}(r_1, \omega_q) \dots \hat{X}(r_M, \omega_q) \right]^T, \tag{3.4}$$

where $\hat{X}(r_i, \omega_q)$ represents an estimate of the true response $X(r_i, \omega_q)$ from (1.6) with $1 \leq i \leq M$. The model data represents a collection of expected responses from a scatterer with no random noise or unmodeled multipath interference. Each element of \mathbf{r} corresponds to the total travel distance from a transmitting sensor to the damage scatterer and then to a receiving sensor. In Section 3.3, we discuss how to obtain $\hat{\mathbf{x}}_q(\mathbf{r})$ using sparse wavenumber analysis and sparse wavenumber synthesis.

3.3 Data-driven matched field processing

In this section, we present data-driven matched field processing by integrating matched field processing with a sparsity-based, data-driven model generated from sparse wavenumber analysis and sparse wavenumber synthesis. We assume our sensors are randomly distributed across the plate's surface. In Section 3.4, we analytically demonstrate the asymptotic benefits of using data-driven matched field processing with random sensor placement.

To generate a data-driven model, we apply two processes: sparse wavenumber analysis and sparse wavenumber synthesis. Sparse wavenumber analysis uses compressed sensing algorithms [53] to recover a sparse frequency-wavenumber representation $\hat{\mathbf{v}}_q$ from the calibration data \mathbf{y}_q . This process leverages the knowledge that, when the chosen number of wavenumbers N in (3.1) is very large, the dispersion curves \mathbf{v}_q are sparse, or contain mostly zeros. Alternative sparsity-based approaches for dispersion estimation [88, 89] have been discussed in other




Figure 3.1: The frequency-wavenumber representation (dispersion curves) of guided wave data from an aluminum plate, recovered by sparse wavenumber analysis.

applications but have not been extended for use in broadband localization. Sparse wavenumber synthesis then uses the recovered dispersion curves $\hat{\mathbf{v}}_q$ to predict how waves propagate between any two points in the medium.

3.3.1 Building the data-driven model

As discussed in Chapter 2, sparse wavenumber analysis recovers the frequency-wavenumber, or dispersion curve, representation $\hat{\mathbf{v}}_q$ of the medium by applying a basis pursuit denoising, or lasso, optimization [53, 54] from compressed sensing to the calibration data \mathbf{y}_q such that

$$\hat{\mathbf{v}}_q = \arg \min_{\mathbf{v}_q} \|\Phi(\mathbf{d})\mathbf{v}_q - \mathbf{y}_q\|_2^2 + \tau \|\mathbf{v}_q\|_1, \quad (3.5)$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ are defined as the ℓ_1 norm, the absolute sum of all elements, and the ℓ_2 norm, the Euclidean distance, of a vector, respectively. By utilizing the knowledge that the true \mathbf{v}_q vector is sparse, basis pursuit denoising allows us to compute an accurate estimate $\hat{\mathbf{v}}_q$ with high resolution even when the matrix $\Phi(\mathbf{d})$ is underdetermined, i.e., $N > M$. Figure 3.1 illustrates an example frequency-wavenumber representation recovered from a collection of experimental plate data. The regularization parameter τ is used to tune the method's robustness to error and is chosen to be 0.5 based on results from prior work [49, 85].

From the recovered frequency-wavenumber representation $\hat{\mathbf{v}}_q$, we can then synthesize model data corresponding to our chosen distances \mathbf{r} through sparse wavenumber synthesis by solving the forward problem

$$\hat{\mathbf{x}}_q(\mathbf{r}) = \mu_q \Phi(\mathbf{r}) \hat{\mathbf{v}}_q, \quad (3.6)$$

where the scalar μ_q is a factor

$$\mu_q = \frac{\mathbf{y}^H \Phi(\mathbf{d}) \hat{\mathbf{v}}_q}{\|\Phi(\mathbf{d}) \hat{\mathbf{v}}_q\|_2^2} \quad (3.7)$$

that adjusts for a multiplicative bias in $\hat{\mathbf{v}}_q$ introduced by basis pursuit denoising. The function $\hat{\mathbf{x}}_q(\mathbf{r})$ represents the sparsity-based, data-driven model of the medium since it represents our predicted response between any two points, as defined by \mathbf{r} , and is based on our sparse, data-derived dispersion curves $\hat{\mathbf{v}}_q$.

3.3.2 Coherent data-driven matched field processor

Matched field processing localizes a target by comparing test data \mathbf{x}_q with model data $\hat{\mathbf{x}}_q(\mathbf{r})$. This comparison is implemented by a matched field processor [1] and the output of the processor is an ambiguity function $b(\mathbf{r})$. If the value of the ambiguity function, at a given \mathbf{r} , is large, the model and measurements are closely “matched.” For a single target, the target’s estimated location $\hat{\mathbf{r}}$ is defined by

$$\hat{\mathbf{r}} = \arg \max_{\mathbf{r}} b(\mathbf{r}) . \quad (3.8)$$

When measured across a large number of locations in a grid, the ambiguity function forms an “image” of the medium.

In this section, we integrate our data-driven model in (3.6) with the coherent matched field processor [73, 74], one of the most widely studied and applied processors throughout the literature. In this next section, we integrate our data-driven model in (3.6) with the also widely implemented incoherent matched field processor [1]. While we focus on these processors, the data-driven model can be applied to any other matched field processor [71, 72, 74] and can be applied to many other localization algorithms [90, 91].

The ambiguity function of the coherent matched field processor can be derived from the solution to the minimization problem [73]

$$\hat{\mathbf{r}} = \arg \min_{\mathbf{r}, \beta} \sum_{m=1}^Q \|\mathbf{x}_q - \beta \hat{\mathbf{x}}_q(\mathbf{r})\|_2^2 , \quad (3.9)$$

where \mathbf{x}_q is the test data in (3.3) and $\hat{\mathbf{x}}_q(\mathbf{r})$ is the model data in (3.4). The complex-valued coefficient β represents the unknown amplitude of the measured signal. The coherent processor is best applied when the measured signal’s frequency response is well known up to a single multiplicative factor.

By solving for and substituting the closed-form solution of β into (3.9) and expanding the results, the coherent processor can be expressed as a maximization problem, in the same form

as (3.16), with an ambiguity function defined by

$$\begin{aligned}\hat{\mathbf{r}} &= \arg \max_{\mathbf{r}} b(\mathbf{r}) \\ b(\mathbf{r}) &= \frac{\left| \sum_{q=1}^Q \mathbf{x}_q^H \hat{\mathbf{x}}_q(\mathbf{r}) \right|^2}{\sum_{q=1}^Q \|\hat{\mathbf{x}}_q(\mathbf{r})\|_2^2}.\end{aligned}\quad (3.10)$$

In this form, the coherent processor is represented by an inner product between the test data \mathbf{x}_q and the model data $\hat{\mathbf{x}}_q(\mathbf{r})$. Intuitively, the coherent processor propagates a time-reversed replica of each received signal (or the complex conjugate of \mathbf{x}_q in the frequency domain) backward into the medium modeled by $\hat{\mathbf{x}}_q(\mathbf{r})$. For this reason, the coherent processor may also be referred to as a time reversal processor [26, 27].

If we substitute \mathbf{x}_q in (3.10) with its matrix representation in (3.3) and $\hat{\mathbf{x}}_q(\mathbf{r})$ in (3.10) with the data-driven model as expressed in (3.6), we derive the coherent data-driven matched field ambiguity function. For a noise-free ($\mathbf{n}_q^{(x)} = \mathbf{0}$) system, this ambiguity function is expressed as

$$b(\mathbf{r}) = \frac{\left| \sum_{q=1}^Q \mathbf{v}_q^H \Phi^H(\mathbf{r}^*) \Phi(\mathbf{r}) \hat{\mathbf{v}}_q \right|^2}{\sum_{q=1}^Q \|\Phi(\mathbf{r}) \hat{\mathbf{v}}_q\|_2^2}, \quad (3.11)$$

where \mathbf{v}_q represents the *true* frequency-wavenumber representation at frequency ω_q and $\hat{\mathbf{v}}_q$ is the recovered dispersion curve representation from sparse wavenumber analysis. We rewrite (3.11) in a compact matrix/vector. We stack each of the wavenumbers vectors

$$\mathbf{v} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_Q] \quad (3.12)$$

$$\hat{\mathbf{v}} = [\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \dots, \hat{\mathbf{v}}_Q], \quad (3.13)$$

and form a block diagonal matrix with the propagation matrices

$$\overline{\Phi}(\mathbf{r}) = \begin{bmatrix} \Phi(\mathbf{r}) & \mathbf{0} & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \Phi(\mathbf{r}) \end{bmatrix}. \quad (3.14)$$

The ambiguity function is then written as

$$b(\mathbf{r}) = \frac{\left| \mathbf{v}^H \overline{\Phi}^H(\mathbf{r}^*) \overline{\Phi}(\mathbf{r}) \hat{\mathbf{v}} \right|^2}{\|\overline{\Phi}(\mathbf{r}) \hat{\mathbf{v}}\|_2^2}. \quad (3.15)$$

When $\mathbf{r} = \mathbf{r}^*$, the value of the ambiguity function is dependent on how well the frequency-wavenumber representations \mathbf{v} and $\hat{\mathbf{v}}$ match. When $\mathbf{r} \neq \mathbf{r}^*$, the value of the ambiguity function becomes dependent on the coherence between the columns of $\Phi(\mathbf{r}^*)$ and $\Phi(\mathbf{r})$. We utilize this form of the coherent processor for asymptotic analysis in Section 3.4.

3.3.3 Incoherent data-driven matched field processing

We now integrate the dispersion curves from sparse wavenumber analysis $\hat{\mathbf{v}}_q$ with the incoherent matched field processor to localize an acoustic source. As with the coherent matched field processor, the target's estimated location $\hat{\mathbf{r}}$ for a single target is defined by

$$\hat{\mathbf{r}} = \arg \max_{\mathbf{r}} b(\mathbf{r}) . \quad (3.16)$$

Similarly again, when measured across a large number of locations in a grid, the ambiguity function for the incoherent matched field processor forms an “image” of the medium.

Given the measured test data \mathbf{x}_q and the model data $\hat{\mathbf{x}}_q(\mathbf{r})$, the incoherent matched field processor can be defined by a least-squares optimization such that the optimal location estimate $\hat{\mathbf{r}}$ is [73]

$$\hat{\mathbf{r}} = \arg \min_{\mathbf{r}, \beta_1, \dots, \beta_Q} \sum_{q=1}^Q \|\mathbf{x}_q - \beta_q \hat{\mathbf{x}}_q(\mathbf{r})\|^2 . \quad (3.17)$$

Unlike the coherent processor, where there is a single unknown constant β , the incoherent processor minimizes the error between the data \mathbf{x}_q and the model $\hat{\mathbf{x}}_q(\mathbf{r})$ with an unknown factor complex-valued β_q at each frequency. By optimizing each β_q value, the minimization chooses an optimal frequency-domain representation for the data. Intuitively, this means that the process optimally shifts the model in time (by changing frequency domain phase characteristics) and optimally alters the frequency domain amplitude to match the measured signal. For this reason, we do not need prior information about the shape or timing of the acoustic source to localize it.

We simplify the optimization problem in (3.17) by finding the closed-form solution for each β_q . We then plug these results into (3.17) to represent the original optimization as a maximiza-

tion problem [73, 74]

$$\begin{aligned}\hat{\mathbf{r}} &= \arg \max_{\mathbf{r}} b(\mathbf{r}) \\ b(\mathbf{r}) &= \sum_{q=1}^Q \frac{|\mathbf{x}_q^H \hat{\mathbf{x}}_q(\mathbf{r})|^2}{\|\hat{\mathbf{x}}_q(\mathbf{r})\|^2},\end{aligned}\tag{3.18}$$

where $b(\mathbf{r})$ is known as the ambiguity function of the incoherent matched field processor. Computationally, this representation is much easier to solve than (3.17) since we reduced the number of unknown variables. In the new form (3.18), we are maximizing the inner product between the data \mathbf{x}_q and the model $\hat{\mathbf{x}}_q(\mathbf{r})$ at each frequency. The magnitude of the inner product is squared and normalized such that $b(\mathbf{r}) = 1$ when the data \mathbf{x}_q perfectly matches the model $\hat{\mathbf{x}}_q(\mathbf{r})$. By squaring the magnitude of the inner product at each frequency, the incoherent processor in (3.18) is invariant to phase shifts or time delays in the data.

We derive the incoherent data-driven matched field processor by substituting the expressions for our data-derived model $\hat{\mathbf{x}}_q(\mathbf{r})$ in (3.6) and the guided wave data \mathbf{x}_q in (3.3) into the ambiguity function in (3.18) such that

$$b(\mathbf{r}) = \sum_{q=1}^Q \frac{|\mathbf{v}_q^H \Phi^H(\mathbf{r}^*) \Phi(\mathbf{r}) \hat{\mathbf{v}}_q|^2}{\|\Phi(\mathbf{r}) \hat{\mathbf{v}}_q\|^2},\tag{3.19}$$

where \mathbf{v}_q represents the *true* dispersion curves of the medium and $\hat{\mathbf{v}}_q$ is the recovered dispersion curves from sparse wavenumber analysis. Notice that this result is similar to (3.15), except we sum over each frequency independently. Therefore, our analysis of both the coherent data-driven matched field processor and the incoherent data-driven matched field processor will be similar. In the following section, we derive the asymptotic behavior of these localization methods.

3.4 Asymptotic behavior

In this section, we analyze the asymptotic behavior, as a function of the number of measurements M , of the coherent data-driven matched field ambiguity function. We analyze this behavior when the measured data \mathbf{x}_q is ideal, i.e. $\mathbf{n}_q^{(x)} = \mathbf{0}$, and when \mathbf{x}_q is corrupted by unmodeled multipath interference. We do not consider the scenario where \mathbf{x}_q is corrupted by

random noise since matched field processing is already well known to be robust to Gaussian measurement noise [1].

Our analysis shows that, even under arbitrary multipath interference, we can theoretically guarantee accurate localization in the medium with a sufficient number of measurements. To show this, we utilize results from compressed sensing and the study of random matrices. The core assumption used by these results is that the sensors are distributed randomly across the medium and therefore $\Phi(\mathbf{r})$ has a random structure.

Our analysis focuses around two properties of random matrices: the restricted isometry property (RIP) and what we introduce and refer to as the restricted nullity property (RNP). If the sensors are randomly distributed across the medium, then the matrix $\Phi(\mathbf{r})$ represents a random matrix of complex exponentials with weighted rows. From compressed sensing, it has been well established that random matrices, such as $\Phi(\mathbf{r})$, satisfy RIP [62, 92, 93]

$$(1 - \delta_s)\|\mathbf{v}_q\|_2^2 \leq \|\Phi(\mathbf{r})\mathbf{v}_q\|_2^2 \leq (1 + \delta_s)\|\mathbf{v}_q\|_2^2, \quad (3.20)$$

with a small $\delta_s \geq 0$ and s non-zero components in \mathbf{v}_q . Intuitively, a matrix that satisfies RIP with a small δ_s is “nearly unitary,” i.e., $\Phi^H(\mathbf{r})\Phi(\mathbf{r}) \approx \mathbf{I}$. Although computing the restricted isometry constant δ_s for a particular matrix is usually an intractable problem [94], the RIP constant δ_s for random matrices is known to decrease in probability as the number of measurements M increases [61, 95, 96]. Furthermore, it has been shown that, for matrices of bounded random vectors, such as $\Phi(\mathbf{r})$, the constant δ_s converges to zero with rate of $1/\sqrt{M}$ or faster [96]. Therefore this convergence occurs relatively quick.

We now introduce RNP, which is an extension of RIP. It states that if $\Phi(\mathbf{r}^*)$, $\Phi(\mathbf{r})$, and $(1/\sqrt{2})(\Phi(\mathbf{r}^*) + \Phi(\mathbf{r}))$ each satisfy RIP with small constants, then the matrices also satisfy the RNP inequality

$$-2\delta'_s\|\mathbf{v}_q\|_2^2 \leq \mathbf{v}_q^H \Phi^H(\mathbf{r}^*)\Phi(\mathbf{r})\mathbf{v}_q \leq 2\delta'_s\|\mathbf{v}_q\|_2^2 \quad (3.21)$$

with a small constant $\delta'_s \geq 0$ and s non-zero components in \mathbf{v}_q . The proof of this property can be found in Appendix A. Intuitively, two matrices that satisfy RNP with a small δ'_s have “nearly uncorrelated” columns, i.e. $\mathbf{v}_q^H \Phi^H(\mathbf{r}^*)\Phi(\mathbf{r})\mathbf{v}_q \approx 0$. For the simplicity of our analysis,

we assume the RIP and RNP constants are equal. This occurs when the RIP constants of $\Phi(\mathbf{r}^*)$, $\Phi(\mathbf{r})$, and $(1/\sqrt{2})(\Phi(\mathbf{r}^*) + \Phi(\mathbf{r}))$ are all equal.

To satisfy RNP, the distance vectors \mathbf{r}^* and \mathbf{r} should be randomly distributed and sufficiently different. Under these conditions, the phase components from each element of $\Phi(\mathbf{r}^*)$, $\Phi(\mathbf{r})$, and $\Phi^H(\mathbf{r}^*)\Phi(\mathbf{r})$ will have a high variance and can be treated as a circularly uniform random variable. As a result, the columns of $\Phi(\mathbf{r}^*)$ and $\Phi(\mathbf{r})$ will be highly uncorrelated.

However, if we finely sample the localization grid, this will not be true for all \mathbf{r} vectors corresponding to points on the grid. Due to the uncertainty principle [97], we know that the width, or resolution, of the peak lobe in the ambiguity function is inversely proportional to the signal's width in the wavenumber domain. In this region of space near the target, where \mathbf{r}^* and \mathbf{r} are similar, our matrices will not satisfy RNP. As a rule of thumb, if the distances in \mathbf{r}^* and \mathbf{r} are more than two times the maximum wavelength apart and random in distribution, then the phase components of each element in $\Phi(\mathbf{r}^*)$ and $\Phi(\mathbf{r})$ will be effectively random and the matrices will satisfy RNP. In Chapter 4, we show that this ambiguous region is approximately 0.68 cm to 0.9 cm in diameter for the coherent matched field processor applied to our experiments.

3.4.1 Single path scenario

In this subsection, we derive the asymptotic behavior of the coherent data-driven ambiguity function for ideal measurements where $\mathbf{n}_q^{(x)} = \mathbf{0}$. In previous work [49], it has been shown that, when $\Phi(\mathbf{d})$ satisfies RIP, sparse wavenumber analysis can accurately recover the frequency-wavenumber representation with a sufficient number of measurements. Therefore, we assume for our analysis that the frequency-wavenumber representation is correctly recovered up to a constant scalar factor

$$\widehat{\mathbf{v}}_q \cong \mathbf{v}_q . \quad (3.22)$$

We can ignore the unknown scalar factor since it does not change the ambiguity function when plugged into (3.15).

By using the assumption in (3.22) and applying the RIP and RNP inequalities in (3.20) and (3.21), respectively, to the either coherent or incoherent data-driven matched field ambiguity function in (3.15) or (3.19), respectively, we derive a lower bound for the ambiguity function's

target-to-artifact amplitude ratio

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{(1 - \delta_s)^2}{4\delta_s^2}. \quad (3.23)$$

Note that this result is slightly stronger than the results originally presented in [50]. This inequality is true for both the coherent and incoherent processors, although the value δ_s is different for the two processors. This relationship is proven in detail in Appendix B and Appendix C for the coherent and incoherent processors, respectively. The result in (3.23) illustrates the behavior of artifacts in the ambiguity function. It shows that, as the RIP constant $\delta_s < 1$ decreases, the ambiguity function's target-to-artifact amplitude in (3.23) increases monotonically, improving localization performance. Therefore, as we increase the number of measurements M , the target-to-artifact amplitude in (3.23) is guaranteed to decrease in probability.

3.4.2 Multipath scenario

We now consider a scenario in which the measured data is corrupted by unmodeled multipath interference from the boundaries of the medium. Under this scenario, the test data is expressed as

$$\mathbf{x}_q = \eta_0 \Phi(\mathbf{r}^*) \mathbf{v}_q + \sum_{\ell=1}^L \eta_\ell \Phi(\mathbf{r}_\ell) \mathbf{v}_q, \quad (3.24)$$

where \mathbf{r}^* represents the location of the target and $\mathbf{r}_1, \dots, \mathbf{r}_L$ represent the distance vectors associated with the L additional paths. The term η_0 represents the strength of the direct path signal and η_ℓ for $1 \leq \ell \leq L$ represents the strength of the additional signal corresponding to path ℓ . We assume the distance vector associated with each path \mathbf{r}_ℓ is sufficiently different from \mathbf{r}^* such that $\Phi(\mathbf{r}_\ell)$ and $\Phi(\mathbf{r}^*)$ satisfy RNP in (3.21) for all ℓ . This is a reasonable assumption since the reflected paths will generally travel much further distances than the direct path.

As with the single path scenario, we assume $\hat{\mathbf{v}}_q \cong \mathbf{v}_q$ and then apply both the RIP and RNP inequalities in (3.20) and (3.21) to the coherent and incoherent data-driven matched field ambiguity functions to derive a lower bound for the target-to-artifact amplitude ratio

$$b(\mathbf{r}^*) \geq \frac{\eta^2}{1 + \eta^2} \left(\frac{(1 - \delta_s)^2}{4\delta_s^2} - \frac{1}{\eta^2} \right), \quad (3.25)$$

where η is defined as the signal-to-interference ratio

$$\eta = \frac{\eta_0}{\sum_{\ell=1}^L \eta_\ell} . \quad (3.26)$$

Note that this result is slightly stronger than the results originally presented in [50]. As in the single path scenario, this inequality is true for both the coherent and incoherent processors, although with different δ_s values. The proof for this is shown in Appendix D and Appendix E for the coherent and incoherent processors, respectively. The result shows that as the signal-to-interference ratio η approaches infinity, the worst-case target-to-artifact amplitude ratio (3.25) converges to the single path bound in (3.23). The result also illustrates that, although unmodeled multipath interference negatively affects the target-to-artifact amplitude ratio, the bound still increases monotonically as $\delta_s < 1$ decreases. Therefore, as in the ideal signal scenario, the target-to-artifact amplitude ratio will increase in probability as the number of measurements M increases.

In terms of localization, this implies that, for any arbitrary multipath interference and any single target location in the medium, there exists a sufficient number of measurements M such that the region around the target contains the largest value in the ambiguity function with a high probability. Said in another way, the data-driven matched field processor is asymptotically unbiased to within a small ambiguous region around the target, where RNP does not apply. Note that targets near the boundaries of a medium will have a significantly weaker signal-to-interference ratio η than targets located in the center of a medium due to the stronger unmodeled multipath. Therefore, the number of measurements M necessary for data-driven matched field processing to locate a target may vary as a function of the target's location.

3.5 Comparison with clustered sensor topologies

One key assumption used to analyze the asymptotic behavior of the data-driven matched field processors is that the sensors are sufficiently distant, or sparsely spaced, from each other and randomly placed in order to satisfy RIP and RNP. To illustrate the necessity for this assumption, we simulate localization scenarios using both a sparsely spaced and clustered topology. We consider a 2 m by 2 m two-dimensional region with source located in the area of

interest. We assume the source transmits a continuous, single-frequency waveform with a single wavenumber of $k(\omega) = 500$ cycles/m (a wavelength of approximately 1.25 cm). To perform matched field processing, we chose a grid spacing of 2 cm in the horizontal and vertical directions.

The receiving sensor locations were chosen randomly from independent uniform distributions in the horizontal and vertical directions. We generated interference by assuming that the boundaries of the 2 m by 2 m region acted as perfect reflectors and by using ray-tracing procedures to determine the distance traveled by each path from the source to each sensor. We simulated every path that interacts with a boundary up to five times between the source and each sensor. We distributed energy equally among each of these interference signals.

We performed a Monte Carlo simulation with 50 different random receiving sensor permutations for a varying number of sensors and varying levels of signal-to-interference ratio. The ambiguity ratio $b(\mathbf{r}_0)/b(\mathbf{r})$ is computed by finding the average ratio between the value at the source and the maximum value within the the remaining grid points.

Figure 3.2 illustrates the target-to-artifact ratio, which we also refer to as the ambiguity ratio, as a function of the signal-to-interference ratio. Figure 3.3 shows the ambiguity ratio as a function of the number of sensors. The solid lines illustrate a random sensor topology that satisfies RIP and RNP with small constants. The dotted lines represent a cluster sensor topology, in which a cluster of sensors are randomly placed in a small space and do not satisfy RIP or RNP with small constants. The sensors are clustered to be positioned within 2 wavelengths of each other. These figures demonstrate a clear improvement in localization performance and much larger ambiguity ratio $b(\mathbf{r}^*)/b(\mathbf{r})$ when RIP and RNP are satisfied.

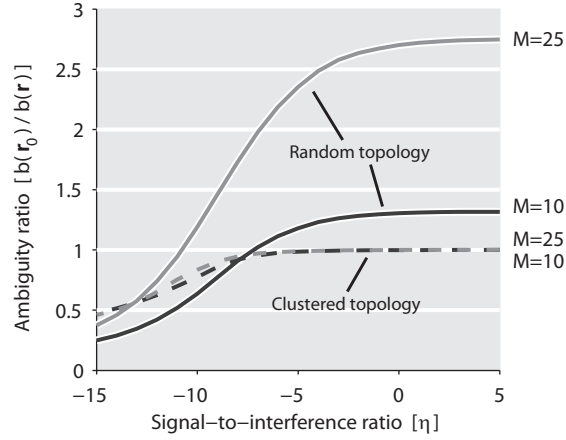


Figure 3.2: The average ambiguity ratio of the localization results versus signal-to-interference ratio η for $M = 10$ (dark lines) and 25 (light lines) sensors. Solid lines show results for a random topology while dotted lines illustrate results for a clustered topology.




Figure 3.3: The average ambiguity ratio of the localization results versus the number of sensors in the system M for signal-to-interference ratios of $\eta = 5$ dB (dark lines) and -9 dB (light lines). Solid lines show results for a random topology while dotted lines illustrate results for a clustered topology.

3.6 Conclusions

In this chapter, we developed data-driven matched field processing, which combines matched field processing with sparse wavenumber analysis to locate damage for structural health monitoring applications. Analytically, we demonstrated data-driven matched field processing to have good asymptotic localization properties. Given a sufficient number of measurements, the data-

driven approach guarantees accurate localization to within a small region. This was shown to be true even in the presence of arbitrary unmodeled multipath interference. Our methodology also offers an alternative to traditional matched field processing techniques that often use complex and rigid models of the propagation environment. Data-driven matched field processing uses a relatively simple multimodal framework to generate a model directly from measured data through fast, convex optimization.

Coherent Data-Driven Matched Field Processing: Acousto-Ultrasonic Localization

4.1 Motivation

Due to their sensitivity to damage and capability of quickly interrogating large structures, guided waves have been an attractive tool for structural health monitoring. However, most guided waves are multimodal and dispersive in nature. This implies that a propagating wave pulse can be represented as a superposition of multiple wave modes that each travel at several different frequency dependent velocities [46]. These properties reduce the effectiveness of most conventional localization techniques. Since many localization algorithms rely on accurate timing information, phase distortions caused by a difference in a wave's group velocity and phase velocity, can be particularly problematic.

To mitigate the effects of a guided wave's complexity on localization techniques, many structural health monitoring methods preprocess data. These preprocessing steps often filter data and then extract an envelope of each measured signal to localize the damage [24, 25, 33]. The filtering usually extracts a narrow band of frequencies with a single, dominant wave mode and relatively little dispersion. The envelope extraction is then performed to reduce the additional phase distortion associated with dispersion. Overall, these preprocessing steps improve localization accuracy by utilize specific features of signals, but at the cost of significantly reducing resolution.

Other localization methods, matched field processing for example, utilize the complex characteristics of the medium to improve localization performance [1, 2, 70–74]. However, these approaches require a good model of the propagation environment and are well known to be sensitive to model errors and environmental variations [2]. In structural health monitoring, known theoretical and numerical models for many guided waves mediums are also available. However, these models rely heavily on properties of the medium that may not be precisely known and may vary over time with changes in environmental and operational conditions [3, 5, 52]. As a result, techniques are necessary to build appropriate and accurate models.

In this chapter, we use sparse wavenumber analysis [49], a method for recovering the frequency-wavenumber representation of data (presented in Chapter 2), and coherent data-driven matched field processing, a localization method that combines sparse wavenumber analysis with matched field processing (presented in Chapter 3), to localize a passive scatterer that reflects ultrasonic guided waves. We demonstrate the methodology’s performance experimentally by localizing two holes in an aluminum plate with data collected in a laboratory. We use the coherent processor, which assumes we can accurately predict the frequency characteristics of the measured signals, because our sensors are synchronized and we model each hole as an ideal point scatterer. Due to these conditions, the frequency characteristics of the reflections from each hole should match those recovered by sparse wavenumber analysis.

In our laboratory experiments, we consider a monitoring problem in which sensors actively search for scatterers in the medium [24]. We compare coherent data-driven matched field processing with a delay-and-sum based approach, which is commonly used in structural health monitoring research [24, 25, 33]. In our results, we achieve a 5 times smaller localization error and a 49 times finer resolution than the delay-and-sum method. We also demonstrate that data-driven matched field processing can localize multiple nearby scatters, which is generally not possible with the delay-and-sum method due to its poor resolution.

4.2 Data-driven calibration

Data-driven matched field processing consists of calibration and matched field localization. We use calibration data from the structure to accurately recover the dispersion curves of the

medium with sparse wavenumber analysis [49]. We then integrate the dispersion curve knowledge with matched field processing to locate a target in the medium. For the *in situ* system considered in this chapter, we do not use a reference specimen to collect the calibration data. The calibration data is collected from the same sensors that are used to collect the test data. This is because sparse wavenumber analysis is robust to noise and to unmodeled multipath interference (e.g., from damage, boundaries, or other inhomogeneities in the structure), as previously shown in chapter 2. Therefore, a reference specimen is not necessary.

The calibration data \mathbf{y}_q , originally discussed and derived in Section 3.2.1, is collected by transmitting and receiving an ultrasonic pulse between all possible pairs of sensors in the system. Therefore, if there are P sensors on the structure, we can collect $P(P-1)/2$ unique calibration measurements with which to build a model. For the localization results discussed in Section 4.4, the calibration data and test data are collected with the same sensors and in the same manner.

The test data \mathbf{x}_q , originally discussed and derived in Section 3.2.2, in this chapter is represented by the difference between a collection of baseline measurements, prior to the introduction of damage, and the current set of measurements, which may contain damage. This baseline subtraction step is often done to remove signals that are not related to the growth of damage and is commonly applied to guided wave structural health monitoring [25, 91]. If the medium has not been changed due to environmental or operation variations, the result of background subtraction will contain the Lamb wave response from damage in the aluminum plate corresponding to travel total distances of $\mathbf{r}^* = [r_1^* \dots r_M^*]$.

The model data $\hat{\mathbf{x}}_q(\mathbf{r})$, originally discussed and derived in Section 3.2.3, represents the signals predicted by sparse wavenumber synthesis for the distances contained in the vector \mathbf{r} . To localize damage in a structure, we can generate model data for many \mathbf{r} vectors, each corresponding to a possible scatterer at different locations on a grid, and then apply coherent matched field processing to compare the data-driven model $\hat{\mathbf{x}}_q(\mathbf{r})$ with the test data \mathbf{x}_q . The coherent data-driven matched field processor is discussed in depth in Chapter 3. In the following sections, we define our specific choice of \mathbf{r} for locating a damage in a large structure, and we demonstrate the performance of data-driven matched field processing.

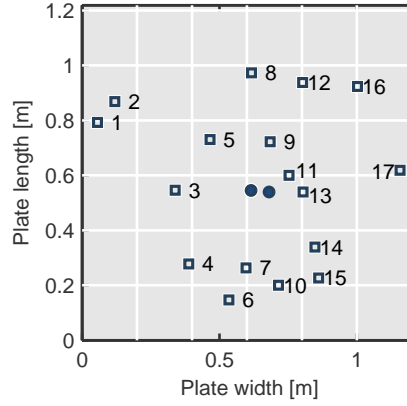


Figure 4.1: A diagram of the aluminum plate used for our experimental setup. Each square indicates a sensor used to transmit and receive signals and the filled circles indicate the locations of the two holes, drilled through the plate, that are used in the experiments.

4.3 Experimental methodology

We test data-driven matched field processing on a 1.22 m by 1.22 m by 0.28 cm aluminum plate. To transmit and receive Lamb wave signals, we bonded 17 PZT (lead zirconate titanate) transducers with dimensions of 0.7 cm by 0.8 cm by 0.2 mm across the top surface of the plate. The positions of each sensor were chosen at random to satisfy RIP, as discussed in Section IV. Figure 4.1 illustrates the locations of each sensor. Note that several sensors are located near the edges of the plate, so we measure significant unmodeled multipath interference from reflections.

The experiment was performed in two steps. First, we collected the baseline, calibration data \mathbf{y}_q on the undamaged plate by transmitting and measuring a 10 μs linear chirp from 0 Hz to 2 MHz between each pair of transducers. This resulted in 272 measurements with 136 unique distances. This data was used to compute the frequency-wavenumber representation of the medium by sparse wavenumber analysis. We then drilled a 0.5 cm hole near the center of the plate and collected an additional 272 measurements in the same manner. Additional measurements were taken after the hole was expanded to 0.75 cm and after a second 0.75 cm hole was drilled 6.5 cm away from the first. The hole locations are shown in Figure 4.1. For each experiment, the test data \mathbf{x}_q is represented by the difference between the data measured with each hole present and a baseline data set with no hole present. While we also use the baseline signals for our calibration data \mathbf{y}_q , we can alternatively use any of the other measured signals due to the

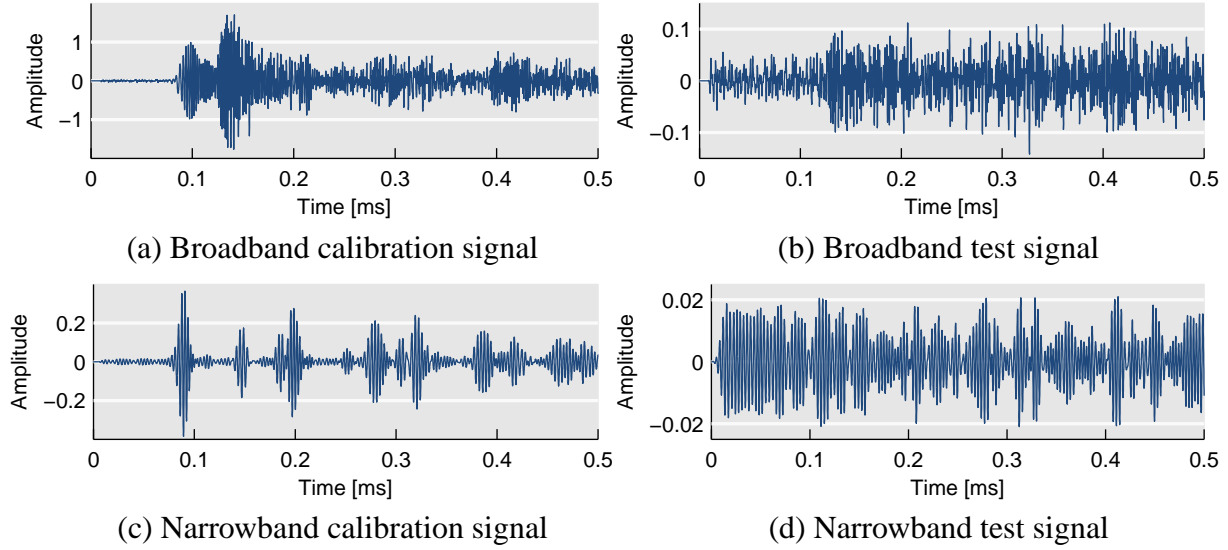


Figure 4.2: Example broadband and narrowband signals from the calibration and test data. The test data represents residue signals after baseline subtraction.

robustness of sparse wavenumber analysis.

Figures 4.2a and 4.2b illustrate signals in the calibration data and baseline subtract test data with the single 0.5 cm hole, respectively. Before plotting, the signals are pulse compressed [98] by correlating them with the transmitted chirp. Figures 4.2c and 4.2d show the same signals but filtered around a frequency of 300 kHz with a bandwidth of 120 kHz. These figures illustrate the complexity of a broadband Lamb wave signal and show that the test data in Figures 4.2b and 4.2d is heavily buried in noise and multipath interference.

4.3.1 Method parameters

To generate the model data $\hat{\mathbf{x}}(\mathbf{r})$, we define \mathbf{r} with respect to a grid of points. We denote (x_r, y_r) to be the coordinates of a specific point on the grid and $(\tilde{x}_i, \tilde{y}_i)$ to be the coordinates of each sensor $1 \leq i \leq P$. To localize a scatterer, we approximate the damage as an ideal point reflector. Therefore, the distance traveled from one sensor to a point on the grid and then to a second sensor is defined by

$$\mathbf{r} = \left\{ \sqrt{(\tilde{x}_i - x_r)^2 + (\tilde{y}_i - y_r)^2} + \sqrt{(\tilde{x}_j - x_r)^2 + (\tilde{y}_j - y_r)^2} \mid i \neq j \text{ and } i, j \in [1, P] \right\}. \quad (4.1)$$

Note that we assume the transmit and receiver sensors are different.

We apply sparse wavenumber analysis to the calibration data with 1000 wavenumbers, uniformly spanning from $\kappa = 0 \text{ m}^{-1}$ to $\kappa = 1000 \text{ m}^{-1}$. Figure 3.1 shows an example frequency-wavenumber representation for the plate. Before computing the frequency-wavenumber representation, as described in (3.5), we apply a window to the data to remove signals that have apparent group velocities, relative to the distance between sensors, of less than 2000 m/s. This has been demonstrated to improve performance in previous work [49].

Like many structural health monitoring localization methods, we also apply a window to the test data \mathbf{x}_q to reduce unmodeled multipath interference appearing late in the signal's time domain [25]. Note that since the target's location is unknown, these windows may remove reflected signals originating from the scatterer. For our experiments, we window the test data with the same 2000 m/s velocity window used by the calibration data. Compared with a fixed time window, this approach adjusts window duration according to the positions of each sensor.

When performing the localization, we utilize 60 frequencies uniformly sampled between 60 kHz and 780 kHz. Note that one strength of data-driven matched field processing is its capability to use wideband signals with many modes and dispersive effects, such as shown in Figure 4.2b. This is in contrast with many traditional localization methods, such as delay-and-sum techniques, that must use a narrow band of frequencies to isolate a signal with single mode and a single group velocity, such as shown in Figure 4.2d.

4.3.2 Comparison with delay-and-sum

We compare our data-driven matched field processing framework with a delay-and-sum localization technique, which is commonly applied to structural health monitoring problems [24,25]. This methodology uses a single mode model with constant and equal phase and group velocities. For Lamb waves, this modeling assumption is generally applicable only around a narrow band of frequencies. Therefore, before applying delay-and-sum localization, we filter the test data around 300 kHz with a bandwidth of 120 kHz, as illustrated in Figures 4.2c and 4.2d.

Specifically, the delay model associated with delay-and-sum localization can be represented

by an $M \times 1$ vector

$$\hat{\mathbf{x}}_q^{\text{dly}}(\mathbf{r}) = S_f(\omega_q) \begin{bmatrix} e^{-j\omega_q r_1/v_g} & \dots & e^{-j\omega_q r_M/v_g} \end{bmatrix}^T. \quad (4.2)$$

where $S_f(\omega)$ is the narrowband-filtered, transmitted signal and the parameter v_g represents the narrowband group velocity of the wave. We let $v_g = 5110.83$ m/s, which is computed from the slope of the zeroth symmetric mode of the frequency-wavenumber representation, the second curve shown in Figure 3.1, at 300 kHz. In the time domain, each complex exponential term in (4.2) represents a delay by r_i/v_g , the expected travel time of a wave traveling with a group velocity v_g .

Due to phase errors resulting from dispersion, many Lamb wave delay-and-sum localization methods do not compare the raw model data and test data, but instead compare the envelopes of each data set [24]. This envelope-delay model is defined by the $M \times 1$ vector

$$\hat{\mathbf{x}}_q^{\text{env}}(\mathbf{r}) = \mathcal{F} \left\{ \left| \mathcal{F}^{-1} \left\{ \hat{\mathbf{x}}_q^{\text{dly}}(\mathbf{r}) \right\} + j\mathcal{H} \left\{ \mathcal{F}^{-1} \left\{ \hat{\mathbf{x}}_q^{\text{dly}}(\mathbf{r}) \right\} \right\} \right| \right\}, \quad (4.3)$$

where $\mathcal{F}\{\cdot\}$, $\mathcal{F}^{-1}\{\cdot\}$, and $\mathcal{H}\{\cdot\}$ represent the discrete Fourier transform, the inverse discrete Fourier transform, and the Hilbert transform, respectively, that operate over the time or frequency ω_q domain associated with each vector.

We consider both the standard delay model $\hat{\mathbf{x}}_q^{\text{dly}}$ and the envelope-delay $\hat{\mathbf{x}}_q^{\text{env}}(r)$. To perform localization, we integrate these models with the coherent matched field processor in (3.10), choosing $\hat{\mathbf{x}}_q = \hat{\mathbf{x}}_q^{\text{dly}}$ or $\hat{\mathbf{x}}_q = \hat{\mathbf{x}}_q^{\text{env}}$. The pairing of the delay-envelope model and the coherent matched field processor is very similar to the delay-and-sum localization methods commonly used in other guided wave structural health monitoring work [24, 25].

As with our data-driven approach, the delay and envelope-delay models are implemented over 60 frequencies. However, since the delay model requires a narrow band of frequencies, it uses 60 frequencies uniformly spanning from 180 kHz to 420 kHz. Also, since the envelope-delay model removes most high frequency components in the envelope extraction process, it uses 60 frequencies uniformly spanning from 0 kHz to 120 kHz. We use a wide band of frequencies for the data-driven model to emphasize that it is not restricted to narrowband signals like other localization methods. As discussed in Section 4.4, we observe little difference in our experimental results when the data-driven model uses the same frequencies as the delay model.

4.4 Experimental results and discussion

In this section, we apply coherent data-driven matched field processing in (3.15) to our experimental data. We then compare our results with the delay model and delay-envelope model in (4.2) and (4.3), respectively. To analyze the relative strength of each method, we consider a normalized ambiguity function

$$b_n(\mathbf{r}) = \frac{b(\mathbf{r}) - \min_{\mathbf{r}} b(\mathbf{r})}{\max_{\mathbf{r}} (b(\mathbf{r}) - \min_{\mathbf{r}} b(\mathbf{r}))} . \quad (4.4)$$

This normalization scales all values of the ambiguity function $b(\mathbf{r})$ to be between 0 and 1.

4.4.1 Figures of merit

We evaluate the performance of each model through three metrics: 1) the localization accuracy, 2) the localization resolution, and 3) the average peak-to-artifact ratio. The localization accuracy tests each method's capability of correctly estimate the target's position. If we assume the true coordinates of the target are (x_r^*, y_r^*) and the estimated coordinates of the target are (\hat{x}_r, \hat{y}_r) , then the localization error r_{err} is defined by

$$r_{err} = \sqrt{(x_r^* - \hat{x}_r)^2 + (y_r^* - \hat{y}_r)^2} . \quad (4.5)$$

We define the localization resolution as the half-amplitude width of the main lobe around the maximum value (\hat{x}_r, \hat{y}_r) in the normalized ambiguity function $b_n(\mathbf{r})$. We determine the half-amplitude width of the main lobe by plotting the values of $b_n(\mathbf{r})$ as a function of distance from the maximum value (\hat{x}_r, \hat{y}_r) . We then perform a least-square fit of a Gaussian curve with maximum value of 1, an unknown minimum value of $1 - c$, and an unknown decay coefficient a ,

$$f(r) = ce^{-ar^2} + (1 - c) , \quad (4.6)$$

where r is the distance from (\hat{x}_r, \hat{y}_r) . The Gaussian fit is chosen because it generally provides a good match to the data. The half-amplitude width σ is then directly computed from the resulting fitting coefficients as

$$\sigma = 2\sqrt{a \log\left(\frac{c - 1/2}{c}\right)} . \quad (4.7)$$




Figure 4.3: An example plot of the normalized ambiguity function from the data-driven model as function of distance. The dotted line represents the Gaussian curve fit to the data.

In general, a small half-amplitude width is desirable because it indicates less ambiguity concerning the location of the target. The half-amplitude also represents a measure of the region for which RNP does not apply, which is discussed in Section 3.4. Therefore, a small half-amplitude width implies that multiple targets may be localized near each another.

Since the maximum of the normalized ambiguity function is 1, we define the average peak-to-artifact ratio as c from (4.6), one minus the minimum value of the Gaussian fit. A large c value denotes greater confidence in the localization due to fewer artifacts. Figure 4.3 illustrates the ambiguity function, for the 0.5 cm hole, plotted as a function of distance from the maximum value and its Gaussian curve fit. While other researchers have similarly used an exponential least-squares fit [90] to evaluate localization performance, we found the Gaussian curve to better characterize our results.

Table 1 shows the resulting figures of merit from our experiments. Resolutions are not provided in Table 1 if localization accuracy is poor, more than 4 hole diameters away from the scatterer. The following subsections discuss these results in greater detail.

4.4.2 One scatterer

First we consider the problem of localizing a single 0.5 cm hole in an aluminum plate. Figures 4.4a–4.4c illustrate the ambiguity functions for the data-driven model, the standard delay model, and the envelope-delay model, respectively. These figures show the ambiguity function over the entire plate. For each of the plots, the light colors represent large values of

Model/Test	Acc. [cm]	Res. [cm]	PAR [dB]
<i>Pitch-catch: One 0.5 cm hole</i>			
Data-driven	0.36	0.68	19.46
Delay	10.87	–	14.93
Envelope-delay	1.72	33.10	5.68
<i>Pitch-catch: One 0.75 cm hole</i>			
Data-driven	0.14	0.78	16.42
Delay	63.76	–	13.95
Envelope-delay	1.72	34.15	5.60
<i>Pitch-catch: Two 0.75 cm holes</i>			
Data-driven (hole 1)	0.13	0.88	17.60
Data-driven (hole 2)	0.41	0.90	19.93

Table 4.1: Figures of merit used to evaluate the localization performance of each model. The figures of merit include the localization accuracy (Acc.), half-amplitude width resolution (Res.), and average peak-to-artifact ratio (PAR).

the ambiguity function and dark colors represent small values. The color code is chosen to be readable in both grayscale and color and is used in all of the following figures. Figures 4.5a–4.5c show a magnified 4 cm by 4 cm region of the ambiguity function around the scatterer. In each plot, the PZT sensors on the plate are denoted by squares and the hole is represented by a circle. If located within the bounds of each magnified image, the crosses represent the maximum value of each respective ambiguity function.

The data-driven method illustrates a clear singular peak with a maximum value located 0.36 cm away from the hole’s measured center, a value 28.0% smaller than the hole’s diameter. The main lobe’s half-amplitude width measures 0.69 cm, approximately the same size as the smallest dimension of each sensor. The data-driven method also achieves an average peak-to-artifact ratio of 19.46 dB and the artifacts gradually weaken as we move further from the target. Note that if we use the same narrowband data as used by the delay model, we observe a relatively insignificant change in performance. The localization error and half-amplitude width improves slightly to 0.24 cm and 0.56 cm, respectively, and the peak-to-artifact ratio weakens to 18.48 dB.

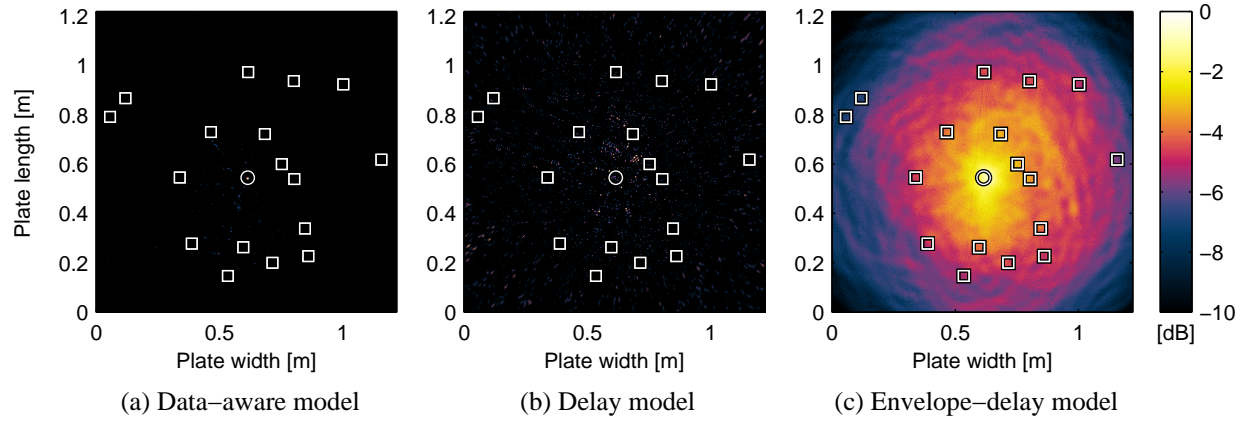


Figure 4.4: Ambiguity functions resulting from applying the data-driven, delay, and envelope-delay models to localize a single 0.5 cm hole in an aluminium plate. The figures show the entire 1.22 m by 1.22 m plate. The squares and circles denote the locations of the sensors and the hole, respectively.

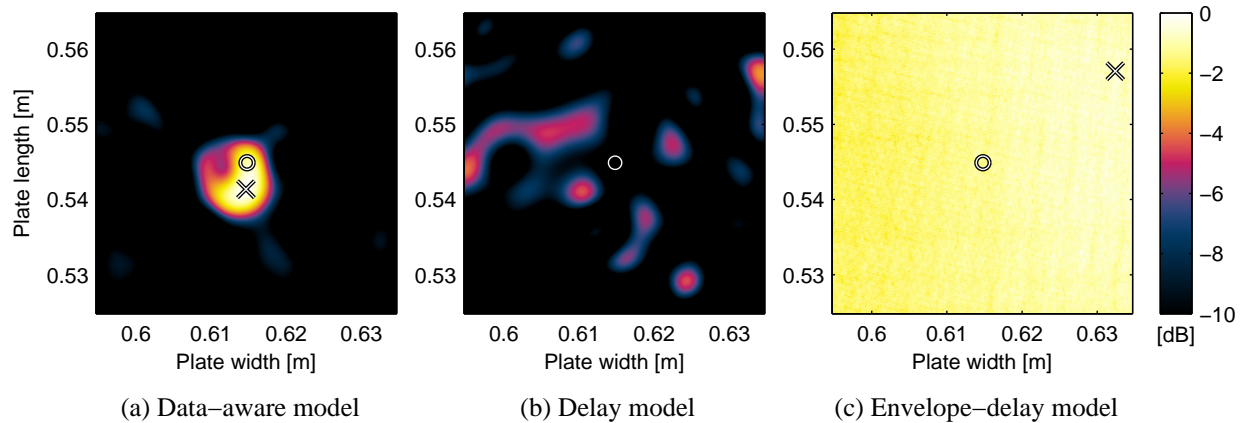


Figure 4.5: Ambiguity functions resulting from applying the data-driven, delay, and envelope-delay models to localize a single 0.5 cm hole in an aluminium plate. The figures show a 4 cm by 4 cm region around the scatterer. The squares and circles denote the locations of the sensors and the hole center, respectively. The cross denotes the estimated location of the hole.

Compared with our wideband data-driven approach, the delay and envelope-delay models achieve poor performance. Due to phase distortions, the delay model fails to localize the scatterer. The delay model's maximum value is located 10.87 cm away from the measured hole center, resulting in a localization error 30.2 times greater than the wideband data-driven method.

The ambiguity function also has many large artifacts throughout the ambiguity function, resulting in an average peak-to-artifact ratio of 14.93 dB, which is 2.84 times smaller than the coherent data-driven matched field processor.

For the envelope-delay model, the maximum value of the ambiguity function is approximately 1.72 cm away from the measured hole center. This result is superior to the standard delay model since it compensates for phase distortions, but it remains worse than the data-driven approach, which has a 4.8 times smaller error. Due to the envelope extraction, the envelope-delay model has a poor half-amplitude width. The main lobe covers almost the entire plate, and, as a result, the wideband data-driven method has a 48.7 times smaller resolution. The large main lobe also results in a poor average peak-to-artifact ratio of 5.68 dB, which is 23.9 times worse than the coherent data-driven matched field processor.

When we increase the hole size from 0.5 cm to 0.7 cm, we observe that the data-driven method's half-amplitude width increases and its peak-to-artifact ratio decreases slightly. However, as shown in Table 4.1, the data-driven results do not change significantly. In contrast, the accuracy of the delay model changes substantially after enlarging the hole. This shows that coherent data-driven matched field processor is relatively stable to small, structural changes in the environment.

4.4.3 Two scatterers

As shown in the previous results, the coherent data-driven matched field processor achieves significantly better accuracy and resolution than the delay and envelope-delay approaches. These results suggest that the data-driven method should be capable of localizing multiple scatterers near each other. In this subsection, we test this hypothesis by using the data-driven approach to localize two 0.75 cm holes separated by a distance of 6.5 cm, approximately 8.7 hole diameters.

We localize a second scatterer by subtracting the Gaussian fit $f(r)$ in (4.6) from the ambiguity function and re-normalizing the result using (4.4) to obtain a new ambiguity function. If we defined $b_n(r)$ as the normalized ambiguity as function of distance from the initial maximum value (\hat{x}_r, \hat{y}_r) and $f(r)$ is the Gaussian fit, then the new normalized ambiguity function $b'_n(r)$ is




Figure 4.6: Ambiguity functions resulting from applying the data-driven, delay, and envelope-delay models to localize two 0.75 cm holes in an aluminium plate. The figures show the entire 1.22 m by 1.22 m plate. The squares and circles denote the locations of the sensors and the holes, respectively.

defined by

$$\begin{aligned}
 b'(r) &= b_n(r) - f(r) \\
 b'_n(r) &= \frac{b'(r) - \min_r b'(r)}{\max_r (b'(r) - \min_r b'(r))} ,
 \end{aligned} \tag{4.8}$$

The new maximum value of $b'_n(r)$ and its associated location on the grid is then defined as the location of the second hole. Alternative multi-target localization approaches, such as the CLEAN algorithm [99], have been proposed to detect multiple targets with matched field processing. In this chapter, we use our iterative approach to quantify the resolution and obtain the figures of merit for the two largest regions in the ambiguity function.

Figures 4.6a–4.6c illustrate the matched field ambiguity function for the the data-driven model, delay model, and the envelope-delay model, respectively, across the entire plate for the two scatterer scenario. Figures 4.7a and 4.7b show the magnified 4 cm by 4 cm regions around each of the two scatterer locations for the data-driven model only. The two hole locations are denoted by circles, the transducers are shown as squares, and the estimated scatterer locations are illustrated as crosses.

Compared with the data-driven model's single scatterer ambiguity function in Figure 4.4a, we visually observe slightly larger artifacts around each hole but good localization of both scat-




Figure 4.7: Ambiguity functions resulting from applying the data-driven model to localize two 0.75 cm holes in an aluminium plate. The figures show the two 4 cm by 4 cm regions around each scatterer. The squares and circles denote the locations of the sensors and the hole centers, respectively. The crosses denote the estimated locations of each hole.

terers. The maximum value of the ambiguity function corresponds to the new hole and achieves a localization error of 0.41 cm, approximately 45.3% smaller than the hole diameter, and a half-amplitude width of 0.88 cm, about 20.0% larger than the hole diameter. After removing the first peak, the second peak corresponds to the location of the original hole from the previous scenario and has a localization error of 0.13 cm, about 82.7% smaller than the hole diameter, and a half-amplitude width of 0.90 cm.

In Figures 4.6b and 4.6c, we see that the delay and envelope-delay models fail to accurately localize both scatterers. As in the single scatterer scenario, the delay method localizes neither scatterer accurately. For the envelope-delay model, a single main lobe covers both holes, causing the result to appear as a single scatterer.

4.5 Conclusions

This chapter implemented the coherent data-driven matched field processing localization methodology, initially discussed in Chapter 3, to localize two holes in an aluminum plate structure. To accurately accomplish this, data-driven matched field processing leverages a data-driven model, which is generated by sparse wavenumber analysis. Sparse wavenumber analysis and sparse wavenumber synthesis are discussed in depth in Chapter 2.

Using experimental guided wave data from the aluminum plate, we compared the performance of data-driven matched field processing with a delay-and-sum localization approach, which is commonly applied in the literature. Compared with this approach, we achieved a 4.8 times improvement in localization accuracy and a 48.7 times improvement in resolution for a single scatterer. We then demonstrated that data-driven matched field processing successfully localizes two different scatterers located 6.5 cm, or 8.7 hole diameters, apart from each other while the delay-and-sum methods could not.

Overall, the data-driven matched field processing framework achieves an improved localization performance over the other methods because it incorporates and obtains information about the medium directly from data. The delay and envelope-delay methods use simple models to represent the complex multimodal and dispersive characteristics of the medium and use filtering and preprocessing to make the data fit the model. In contrast, data-driven matched field processing utilizes the data-derived wideband, multimodal, and dispersive characteristics to localize the scatterers.

Incoherent Data-Driven Matched Field Processing: Acoustic Emission Localization

5.1 Motivation

Acoustic emissions analysis plays a large role in nondestructive evaluation and structural health monitoring systems. An acoustic emission is commonly defined as a spontaneous, transient wave that is generated in a localized region of space. An acoustic emission event may be the result of crack formation or growth from stresses and strains in the material [100, 101] or may be generated from strong, external, and potentially damaging impact forces incident on a structure [28, 29]. For these reasons, detecting and locating acoustic emission sources is a widely applied tool for detecting and locating damage in structures, including concrete structures [102], steel bridge girders [103], diesel engines [104], laminates and composites [28, 105, 106], and various plate-like media [30, 105, 107–109].

In plate-like structures, acoustic emission events generate guided waves that radiate outward from the damage location. We measure these events using piezoelectric sensors distributed across the plate's surface and process the data for information about the acoustic emission's origin. Due to the structure's geometry, plate waves, also known as Lamb waves, propagate with multimodal and dispersive characteristics that distort the waves as they travel through the medium. Mathematically, we can represent a sensor's measured voltage signal $X(r, \omega)$ in the frequency (ω) and (r) domain by the plate framework expressed in (1.6).

In (1.6) r is the distance traveled by the wave resulting from the acoustic emission. In this representation, $S(\omega)$ is the signal emitted by the source and $G_m(\omega)$ and $k_m(\omega)$ are the frequency-dependent amplitude and wavenumber functions of each mode m at angular frequency ω . As illustrated in (1.6), an acoustic event generates multiple wave modes that each propagate with a different frequency-dependent wavenumber, or phase velocity $v_p(\omega) = \omega/k_m(\omega)$. We refer to the collection of wavenumber functions $k_m(\omega)$ for all m as the dispersion curves of the medium [10]. An example set of theoretical dispersion curves for Lamb waves is shown in Figure 1.1.

The multimodal and dispersive wave behavior represents a significant challenge for guided wave analysis because the multiple modes and dispersion continuously alter both the envelope and phase of the guided waves as they propagate through the medium. We generally do not know *a priori* the true dispersion curves for a structure because they vary with material and environmental properties, such as density and temperature. Therefore, we do not precisely know how the guided waves evolve as they travel throughout the medium, and we need to either learn the dispersion characteristics or compensate for them when performing data analysis.

Traditional acoustic emission localization methods determine the origin of the measured signal through triangulation [32, 107, 110] or multilateration [29–31] techniques. These methods are performed in two steps. First, the wave’s time-of-arrival or time-difference-of-arrival, relative to some anchor sensor, is determined for each sensor. This is usually computed by peak detection by finding the maximum value of each measured signal. However, to compensate for dispersion’s distorting effects, the data is usually first simplified by narrowband filtering [24, 111], sometimes through application of the continuous wavelet transform [29, 105], and extracting each signal’s envelope [31, 107]. The time-of-arrival values are then input into a least-squares optimization to estimate the acoustic emission origin.

Although triangulation and multilateration work well in ideal conditions, they can be unreliable when the data is corrupted by significant multipath interference and noise [112]. These effects commonly cause time-of-arrival estimation errors since the signal’s maximum envelope value can be significantly affected by small signal variations, leading to large errors in the localization result. Due to this sensitivity, adding more sensors into a system, which may introduce

new errors, is not always beneficial and can degrade localization performance.

In this chapter, we apply data-driven matched field processing [50] to overcome the challenges associated with triangulation and multilateration methods. Data-driven matched field processing uses data-derived dispersion curves of the medium, extracted from a relatively small set of calibration data [49], to optimally locate an acoustic emission source. Unlike triangulation and multilateration, data-driven matched field processing utilizes all of the data to perform localization and requires no filtering or envelope extraction to simplify the signal, resulting in strong robustness to multipath errors and noise.

In Chapter 5, data-driven matched field processing was used to locate a hole in an aluminum plate using active guided wave interrogation methods, where the initial excitation is well known. In this chapter, we extend the framework to acoustic emission localization, in which the excitation's time trace is unknown and the time of excitation is unknown. Due to those conditions, we use the incoherent data-driven matched field processor and apply it to experimental measurements of acoustic emission events in a laboratory. We evaluate and compare the performance of the incoherent data-driven matched field processor with a conventional technique, known as multilateration, for different levels of multipath error and additive noise. With 8 sensors, we demonstrate a 13 dB, or a factor of 20 times, improvement in mean accuracy and an 11 dB, or a factor of 12.5 times, improvement in robustness to noise when compared with multilateration.

5.2 Data-driven calibration

Data-driven matched field processing consists of two steps, data-driven calibration followed by matched field localization. In the first step, we acquire calibration data from the structure and use it to accurately recover the dispersion curves of the medium. This is accomplished through sparse wavenumber analysis [49], as discussed in Chapter 2. We then integrate the dispersion curve knowledge with matched field processing. In this section, we briefly discuss the data-driven calibration procedure.

The calibration data \mathbf{y}_q , originally discussed and derived in Section 3.2.1, is assumed to be collected prior to the acoustic emission localization. We can collect the calibration data in several ways. For an *in situ* system with transmission capabilities, the calibration data can

be gathered by transmitting and receiving guided waves between each of the sensors on the structure. If the system contains P sensors, this approach will generate $P(P - 1)/2$ unique measurements, allowing us to obtain a significant amount of data with relatively few sensors. This is the approach used in this chapter. Alternatively, we can manually excite and receive guided wave measurements from throughout the structure with a combination of manually positioned ultrasonic probes and/or *in situ* sensors. This approach can be advantageous when the *in situ* sensors cannot transmit information or when we want to collect more calibration data than available with the former approach.

The measured test data \mathbf{x}_q , originally discussed and derived in Section 3.2.2, from each acoustic emission is represented as a collection of guided Lamb wave responses that travel $\mathbf{r}^* = [r_1^* \dots r_M^*]$ distances from the acoustic emission origin and each sensor in the system. Unlike the damage localization system in Chapter 4, we do not need to perform baseline subtraction to remove prior known information. Instead, we process the raw data.

The model data $\hat{\mathbf{x}}_q(\mathbf{r})$, originally discussed and derived in Section 3.2.1, represents signals generated from our data-driven model for distances \mathbf{r} . Similar to the damage localization scenario in Chapter 4, we generate model data for many \mathbf{r} vectors, each corresponding to different source locations on a grid, and apply incoherent matched field processing to compare each data-driven model vector $\hat{\mathbf{x}}_q(\mathbf{r})$ with the test data \mathbf{x}_q . In the following sections, we demonstrate the performance of incoherent data-driven matched field processing.

5.3 Experimental methodology

We test and evaluate incoherent data-driven matched field processing on a 1.22 m by 1.22 m by 0.28 cm aluminum plate. For generating calibration data and measuring acoustic events, we attached 15 *in situ* 0.7 cm by 0.8 cm by 0.09 mm PZT (lead zirconate titanate) transducers on the surface, near the perimeter of the plate. The specific locations of each transducer are shown in Figure 5.1. This choice of sensor locations allows the system to effectively monitor most of the plate's surface for acoustic emissions. However, reflections from the nearby plate boundaries complicate the data analysis.

To generate an acoustic emission, we drop a 9.5 mm diameter ball bearing from a fixed




Figure 5.1: Positions of each sensor location (open squares) and acoustic emission(AE) test location (closed circles) on the experimental plate setup. In the experiments, each acoustic emission location is tested separately.

height of 6.7 cm above the plate. We repeat this procedure at 10 different locations across the plate, shown in Figure 5.1. During the experiments, we observed no ball bearing bounces. We use a ball bearing drop rather than the more traditional pencil lead break for two reasons. First, the ball bearing drop is acoustically stronger. This allows us to perform a noise study, by artificially adding noise, with little additional noise from the experiment.

Second, the ball bearing drop has better repeatability. We test incoherent data-driven matched field processing with 15 sensors but only 8 input channels. For each position in Figure 5.1, we repeat the same acoustic emission experiment 5 times for two different sets of 8 transducers, keeping one transducer common between both sets. We then combine the data sets that have most similar measurements from the common transducer. For the chosen data sets, the correlation coefficient between the common sensor measurements is always greater than 0.99, indicating good repeatability.

5.3.1 Method parameters

In all of our experiments, the calibration data is acquired by having each of the 15 transducers transmit a 0 kHz to 30 kHz linear chirp pulse of 1 ms duration and by measuring the response at each of the other transducers. The entire system transmits and receives data with a 10 MHz sampling rate. Before applying sparse wavenumber analysis to generate the disper-




Figure 5.2: A recovered 3 kHz to 19 kHz frequency-wavenumber, or dispersion curve, representation from measured Lamb wave data.

sion curves shown in Figure 5.2, the calibration data is windowed with an exponential taper to remove any signals traveling with apparent group velocities less than 700 m/s. Figure 5.2 illustrates the resulting recovered dispersion curves. While the results in Figure 5.2 do not display perfectly smooth dispersion curves as found in Chapter 4, possibly due to additional multipath error, we demonstrate that Figure 5.2 is still a good representation of the medium through our experiments.

For incoherent data-driven matched field processing, we assume an acoustic emission originates from a pair of coordinates (x_e, y_e) and our P sensors are positioned at coordinates (x_i, y_i) for $i = 1, 2, \dots, P$, the distances between the acoustic emission source and each sensor are contained in the vector

$$\mathbf{r}_e = \left\{ \sqrt{(x_e - x_i)^2 + (y_e - y_i)^2} \mid i \neq j \text{ and } i \in [1, P] \right\}. \quad (5.1)$$

The vector $\hat{\mathbf{x}}_q(\mathbf{r} = \mathbf{r}_e)$ will then represent a model of the measurements that we would observe from an acoustic emission originating at (x_e, y_e) . By changing the coordinates of (x_e, y_e) , we can generate a new model $\hat{\mathbf{x}}_q(\mathbf{r} = \mathbf{r}_e)$ of an acoustic emission at any another point in the medium. By combining multiple models across a set grid on the structure, we create a discrete wave propagation model $\hat{\mathbf{x}}_q(\mathbf{r})$, which varies as a function of grid location, or \mathbf{r} , and can be integrated with matched field processing.

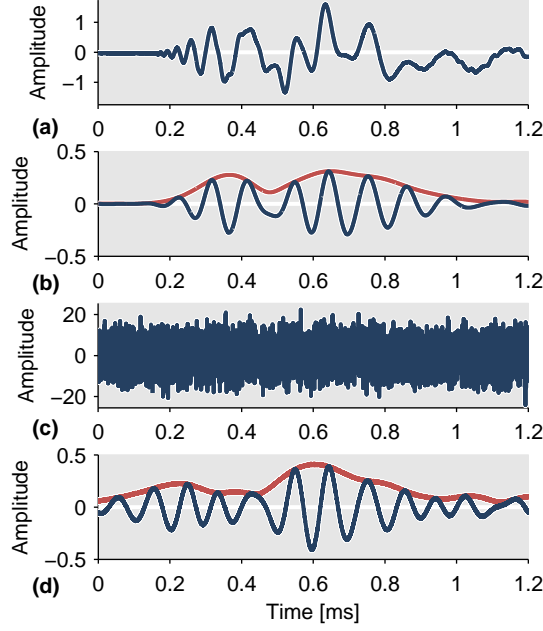


Figure 5.3: An example measured Lamb wave signal (a) without additional processing, (b) with narrowband filtering around 10 kHz, (c) with additional white, Gaussian noise, (d) with additional white Gaussian noise and narrowband filtering.

To localize an acoustic emission with incoherent matched field processing, we use $Q = 21$ frequencies uniformly spaced between 4 kHz and 24 kHz and apply no additional processing. For multilateration, we apply a narrowband Gaussian filter to the data with a 3 dB bandwidth between 5 kHz and 15 kHz. This is done to isolate a single group velocity. We choose the 5 kHz to 15 kHz frequency band because the strongest frequencies are found in this range. The group velocity used by multilateration is computed from the slope of the dispersion curve recovered by sparse wavenumber analysis in Figure 5.2. We use a group velocity of 1142.4 m/s. Examples of the typical signals used for localization are found in Figure 5.3.

5.3.2 Comparison with multilateration

The multilateration localization method uses time-of-arrival estimates from each measurement to find the acoustic emission origin. Conventionally, these methods apply a narrowband filter [29, 111] to reduced distortion from dispersion followed by envelope extraction [31, 107] to simplify the measured data. If $x_f(t)$ represents a measured and narrowband filtered signal,

the envelope can be extracted using the expression

$$\tilde{x}_f(t) = |x_f(t) + j\mathcal{H}\{x_f(t)\}|, \quad (5.2)$$

where j is the imaginary number and $\mathcal{H}\{\cdot\}$ is the Hilbert transform. Once the envelope has been extracted, the maximum value of the data envelope is used to measure the signal's time-of-arrival.

The challenge with this approach is that multipath signals and noise can easily result in an incorrect time-of-arrival estimate. Figure 5.3 shows examples from experimental data where this occurs. Figure 5.3a shows a measured acoustic emission signal before filtering. Figure 5.3(b) shows the filtered signal and demonstrates that the maximum value, at 0.64 ms, is not associated with the direct arrival time, at 0.36 ms. Figure 5.3(c) shows the measured data from Figure 5.3a after being corrupted by wideband white, Gaussian noise. Figure 5.3(d) shows the corrupted signal after filtering. We can see that the noise again changes the estimate arrival time to be 0.59 ms.

Multilateration uses each time-of-arrival estimate to locate the acoustic emission source through a least-square optimization. If we assume t_1, t_2, \dots, t_{P-1} represent the time-of-arrival estimates from $P - 1$ sensors and t_a represents the time-of-arrival for a single anchor sensor, required if the excitation time is not known *a priori*, then multilateration estimates the coordinate source location (\hat{x}, \hat{y}) as the solution to [30]

$$\begin{aligned} (\hat{x}, \hat{y}) &= \arg \min_{x,y} \|u(x, y)\|_2^2 \\ \text{s.t. } u(x, y) &= \sum_{p=1}^P ((x_a - x_p)^2 + (y_a - y_p)^2 - (x - x_p)^2 + (y - y_p)^2) - (t_a v_g - t_p v_g), \end{aligned} \quad (5.3)$$

where v_g is the assumed known group velocity of the waves around the filter frequency.

Compared with data-driven matched field processing, multilateration is more sensitive to multipath and noise. Instead of potentially utilizing the entire measured signal, multilateration optimization only uses P time-of-arrival estimates to locate the acoustic emission source. If one of those estimates is incorrect, it will have a large effect on the result. Therefore, multilateration does not significantly improve with more sensors. We demonstrate this sensitivity experimentally in Section 5.

When we solve multilateration optimization, we use an interior point algorithm and select the center of the plate as its initial condition. We also constrain the optimization so that the result must be located within the plate's boundaries.

5.4 Experimental results and discussion

With the experimentally measured data, we conduct two studies. In the first study, we investigate the performance of data-driven matched field processing when locating acoustic emissions at different points on the plate that experience different levels of multipath error. This study demonstrates data-driven matched field processing's robustness to multipath signals. In the second study, we corrupt the data with various level of white, Gaussian noise to demonstrate data-driven matched field processing's robustness to random noise. We also investigate how the data-driven ambiguity function $b(\mathbf{r})$ varies as a function of the number of sensors used and the signal-to-noise level.

5.4.1 Multipath study

We evaluate data-driven matched field processing at each of the ten acoustic emission locations shown in Figure 5.1. At each location, we estimate the optimal location with sets of 3 to 15 different sensors. The numbers in Figure 5.1 indicate which sensors are in each set such that

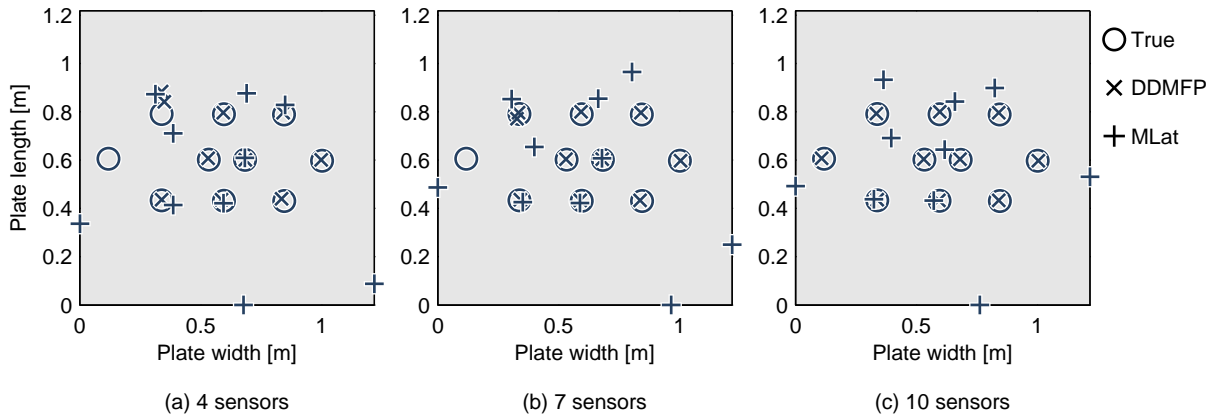


Figure 5.4: The acoustic emission locations estimated by incoherent data-driven matched field processing (DDMFP, crosses) and multilateration (MLat, plus signs). The true origin of each event is shown as an open circle.

a sensor is a member of a set if its number is less than or equal to the total number of sensors in that set.

Figure 5.4 summarizes and compares the localization performance of the incoherent data-driven matched field processor and multilateration for sets of 4, 7, and 10 sensors. We observe that data-driven matched field processing performs much better than multilateration. If we define an accurate localization to be any estimate less than 0.8 cm, the length of one transducer, away from the true location, then the data-driven matched field processor successfully locates 8, 9, and 10 of the acoustic sources with 4, 7, and 10 sensors, respectively. In contrast, multilateration never accurately locates more than 3 acoustic emissions and only accurately estimates 2 acoustic emission locations with 10 sensors. Multilateration degrades with more sensors because the additional sensors introduce more incorrect time-of-arrival estimates due to the multipath.

Figure 5.5 illustrates the mean and median localization error, in terms of distance from the true locations, as a function of the number of sensors. For multilateration, the localization error improves slightly as the number of sensors increases but is significantly biased due to multipath error. In contrast, the mean and median error for the incoherent data-driven matched field processor is always below the multilateration error and remains consistently below 0.8 cm for 8 or more sensors.

Figure 5.6 demonstrates how the incoherent data-driven ambiguity function, for a single acoustic emission location, evolves as we increase the number of sensors. For all three plots, the maximum value correctly estimates the acoustic emission's origin. As we increase the number of sensors, we see the localization uncertainty significantly decreases. In Figure 5.6(c), where most of the image is relatively dark, we can confidently state that the acoustic emission is only limited to a single point. This demonstrates, in contrast with the multilateration results from Figure 5.4, that the incoherent data-driven matched field processor benefits greatly from additional sensor information.

Overall, these results have demonstrated that the incoherent data-driven matched field processor can withstand a significant amount of multipath interference in the data and becomes more robust to multipath as we introduce more sensors into the medium. In contrast, multilater-

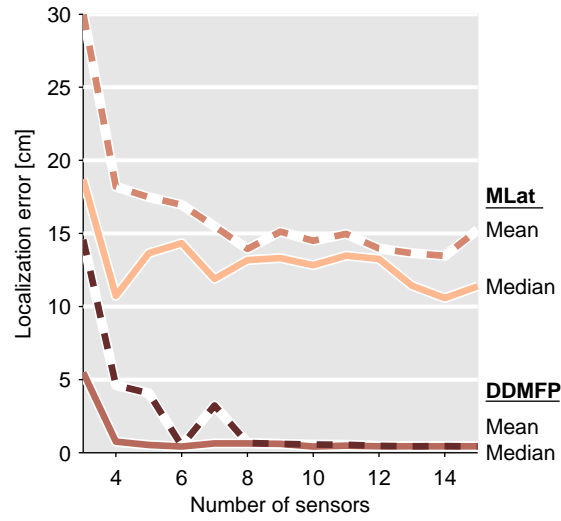


Figure 5.5: The average localization error, in terms of distance from true location, across the 10 acoustic emission event locations. The mean (dotted lines) and median (solid lines) errors are shown for the incoherent data-driven matched field processor (DDMFP) and multilateration (MLat).

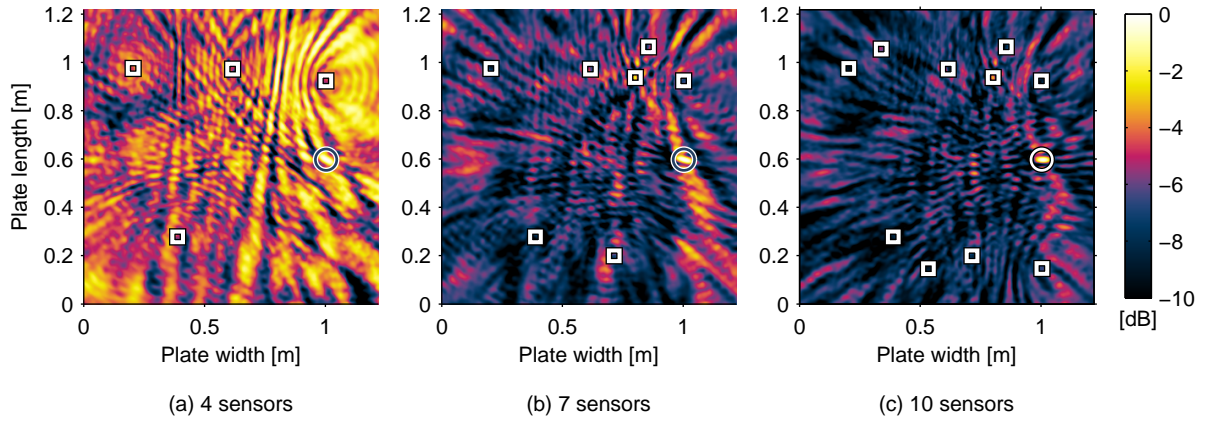


Figure 5.6: The incoherent data-driven ambiguity functions for a specific acoustic emission location (marked by an open circle) with (a) 4, (b) 7, and (c) 10 sensors. The sensors are marked as white, open squares.

ation is sensitive to multipath interference and may suffer from additional sensors because they may introduce incorrect time-of-arrival estimates and significantly alter the final localization result.

5.4.2 Noise study

This study evaluates the performance of data-driven matched field processing as we increase the amount of noise in the data. We corrupt the data with additive white, Gaussian noise at a variety of different average signal-to-noise ratios. If the measured signals, before filtering, at P sensors are represented in the time domain as $x_1(t), \dots, x_P(t)$, then we define the average signal-to-noise ratio (SNR) as

$$SNR = \frac{1}{PQ\sigma^2} \sum_{p=1}^P \int_0^\infty |x_p(t)|^2 dt, \quad (5.4)$$

where Q is the number of frequencies used by the incoherent data-driven matched field processor and σ^2 is the noise variance. We corrupt the data with 100 different instantiations of noise for each signal-to-noise ratio considered. As in the previous study, we claim that the localization is accurate when the estimate is less than 0.8 cm, the length of a PZT sensor, away from the true location. Based on this metric, we find the sample probability of detection by computing the ratio of the number of accurate localizations versus the total number of trials.

We test the incoherent data-driven matched field processor's robustness to noise by estimating the location of the sixth (from left-to-right then top-to-bottom in Figure 5.1) acoustic emission with both a set of 5 and 8 sensors. In a noise-free scenario, both data-driven matched field processing and multilateration accurately locate the acoustic emission.

Figure 5.7 illustrates the resulting sample probability of detection as a function of signal-to-noise ratio. The solid and dotted lines represent the 5 and 8 sensor results, respectively. Curves further to the left on the plot indicate better robustness to noise. Compared with multilateration, incoherent data-driven matched field processing achieves a 7 dB, or 5 times, and a 11 dB, or 12.5 times, improvement in signal-to-noise robustness with 5 and 8 sensors, respectively. Furthermore, we observe that multilateration's robustness worsens when the number of sensors is increased. As similarly demonstrated by the multipath study, this effect occurs because additional sensors introduce additional opportunity for time-of-arrival estimation errors.

Figure 5.8 shows the incoherent data-driven matched field processing ambiguity function for a 3 dB signal-to-noise ratio with 8 sensors. Figure 5.3(c) and 5.3(d) illustrate examples of measured data with this level of additive noise. At 3 dB, data-driven matched field processing

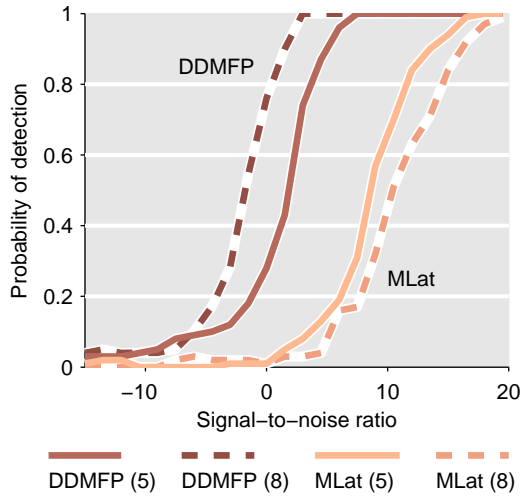


Figure 5.7: The sample probability of detection versus signal-to-noise ratio for the incoherent data-driven matched field processor (DDMFP) and multilateration (MLat) with 5 sensors (solid lines) or 8 sensors (dotted lines).




Figure 5.8: The incoherent data-driven ambiguity function measured with 8 sensors and a signal-to-noise ratio of 3 dB. The plus signs (+) show 100 different estimated acoustic emission locations from multilateration.

has a 100 percent localization accuracy while multilateration has a 3 percent accuracy. The plus signs (+) in Figure 5.8 show the 100 different location estimates by multilateration. We can see that multilateration estimates are close in some instances but the noise is almost always strong enough to cause significant error.

Overall, this study demonstrates that the incoherent data-driven matched field processor is, compared with multilateration, strongly robust to multipath interference and noise. Since the incoherent data-driven matched field processing utilizes the dispersive properties of the medium to locate the source, it does not have to rely on the maximum amplitude values like multilateration. For these reasons, the incoherent data-driven matched field processor uses more information than multilateration and its performance can improve as we introduce additional sensors that provide even more information.

5.5 Conclusions

This chapter implemented the incoherent data-driven matched field processor, which can localize an acoustic source without knowledge of the acoustic signal's timing or shape characteristics, for guided wave acoustic emission localization. Unlike conventional acoustic emission localization methods, the incoherent data-driven matched field processor performs localization as one step and utilizes, through sparse wavenumber analysis, all of the dispersive behavior of the guided waves over a wide range of frequencies to localize an acoustic emission event. As a result, the incoherent data-driven matched field processor is robust to multipath interference and random noise that is commonly encountered in acoustic emission measurements.

We conducted two studies with 10 sets of experimental acoustic emission data to demonstrate the effects of multipath interference and random noise on the incoherent data-driven matched field processor and a conventional acoustic emission localization technique known as multilateration. The multipath study demonstrated that, unlike multilateration, the incoherent data-driven matched field processor's accuracy improves consistently as more sensors are added to the system. With 8 or more sensors, the incoherent data-driven matched field processor successfully localized all 10 acoustic emission events. In contrast, multilateration only successfully located 2 of the 10 events with 10 sensors.

In the noise study, we demonstrated that the incoherent data-driven matched field processor's robustness to noise also improves as additional sensors are added to the system. In contrast, multilateration degrades with additional sensors due to its sensitivity to error. With 5 and 8 sensors, the incoherent data-driven matched field processor achieved 5 times and 12.5 times, respectively, better noise performance than multilateration.

Overall, the studies show that the data-driven matched field processing is a robust tool for localizing acoustic sources in complex media. In future work, we plan to further integrate sparse wavenumber analysis and its data-driven model with more complex structures and other localization processors. We also plan to explore the theoretical performance bounds of these data-driven methodologies.

Refining for Environmental Variations: The Scale Transform

6.1 Motivation

In laboratory settings, guided wave ultrasonic inspection methods are being extensively explored for testing large structures. Guided ultrasonic waves are popular because they propagate through the thickness of the structure and can travel over long distances with little attenuation [46]. This allows sensors to interrogate large areas all at once. However, guided waves are inherently multi-modal and dispersive in their propagation [113]. In addition, structures' boundaries generate reflections and exchange energy between wave modes [114]. These effects make the interpretation of measured data difficult and necessitate the use of baseline measurements.

Under ideal conditions, damage can be detected by performing baseline subtraction or time-domain correlation between a measured signal and a baseline measurement [42]. Baseline subtraction and correlation techniques attempt to remove effects from static sources, such as reflecting boundaries, but are impractical under realistic conditions. When environmental or operational conditions change, the propagating medium and guided wave behavior also change. Therefore, simple baseline comparison methods are unable to distinguish damage from benign environmental and operational effects.

One of the most ubiquitous environmental properties to affect signals is temperature [39]. Variations in temperature alter the velocity of the guided waves [38]. This effect on velocity

can be attributed to temperature's influence on thermal expansion and the Young's modulus of the material [36]. Furthermore, small changes in temperature have been shown to cause large changes to the residual signal observed after baseline subtraction [42]. For these reasons, it is important that structural health monitoring systems adjust for the effects of temperature.

Several techniques have already been developed to compensate for temperature in guided waves. These compensation methods can be divided into two categories: data-driven methods and model-driven methods. In this chapter, we focus on model-driven techniques based on an approximate model of the effects of temperature on ultrasonic waves. This is in contrast to data-driven approaches, where temperature compensation is achieved by comparing new observations with collections of previously obtained data. Although we focus on model-driven approaches, these techniques can usually be combined with data-driven methods to improve the overall compensation performance [115].

The methods explored in this chapter assume a stretch-based model for temperature. A change in temperature $T_\alpha\{\cdot\}$ on an ultrasonic signal $x(t)$ is approximated by a time-stretch

$$T_\alpha\{x(t)\} \approx x(\alpha t) \quad (6.1)$$

for small stretch factors α . This time-stretch model is attributed to the temperature's effect on the wave velocity. A change in the wave velocity causes later wave arrivals to have proportionally larger time shifts, creating a stretch-like effect. Although the time-stretch model is not generally true, it is approximately true for diffuse waves, which are highly sensitive to temperature thermal fluctuations [3, 41, 116]. Temperature changes also distort the diffuse wave fields due to the different sensitivities of shear and longitudinal waves, but we neglect these effects [38, 117].

Methods that use a time-stretch model include the local peak coherence [42, 44] and optimal signal stretch (OSS) [39, 40, 115] techniques. Given small variations in temperature and ideal conditions, these approaches work well. However, previous work has shown local peak coherence to be sensitive to errors associated with the approximate stretch-based model as well as other effects [118]. Although OSS tends to be more robust to model error, it suffers from a high computational complexity.

In this chapter, we present a new methodology for stretch-based, model-driven temperature compensation algorithms based on the scale transform. This work was originally presented in references [5, 118–120]. In the scale transform domain, we can directly manipulate the stretch factor of signals and compute quantities invariant to changes in that stretch factor. We discuss three algorithms for temperature compensation based on these scale domain tools: the scale-invariant correlation (SIC) method, the iterative scale transform (IST) method, and the combined SIC/IST method. We demonstrate these algorithms, through analysis and experimental tests, to be more computationally efficient than currently available techniques.

6.2 Scale transform signal processing

In this section, we express the solution to the temperature compensation problem as an optimization problem. We then demonstrate how this optimization problem may be solved with the scale transform.

6.2.1 Problem Formulation

According to the model in (6.1), we can compensate for temperature by stretching the distorted signal by a factor of $1/\alpha$. Since α is usually unknown, it must be estimated. We define the optimal stretch factor estimate $\hat{\alpha}$ between two signals $x(t)$ and $s(t)$ as that which minimizes their normalized squared error

$$\hat{\alpha} = \arg \min_{\alpha} \int_0^{\infty} \left| \frac{x(t)}{\sigma_x} - \frac{s(\alpha t)}{\sigma_s/\sqrt{\alpha}} \right|^2 dt. \quad (6.2)$$

The normalization factors in (6.2) are defined as

$$\begin{aligned} \sigma_x^2 &= \int_0^{\infty} |x(t)|^2 dt \\ \sigma_s^2/\alpha &= \int_0^{\infty} |s(\alpha t)|^2 dt \end{aligned} \quad (6.3)$$

so that the energy of each term, $x(t)/\sigma_x$ and $s(\alpha t)/(\sigma_s/\sqrt{\alpha})$, in (6.2) is equal to 1. This makes the processing robust when we lack knowledge of the energy of the recorded signal $x(t)$.

Expanding the square in (6.2) yields

$$\begin{aligned}
\hat{\alpha} &= \arg \min_{\alpha} \int_0^{\infty} \frac{|x(t)|^2}{\sigma_x^2} + \frac{|s(\alpha t)|^2}{\sigma_s^2/\alpha} - \frac{x(t)s(\alpha t)}{\sigma_x(\sigma_s/\sqrt{\alpha})} dt \\
&= \arg \min_{\alpha} 2 - \int_0^{\infty} \frac{x(t)s(\alpha t)}{\sigma_x(\sigma_s/\sqrt{\alpha})} dt \\
&= \arg \max_{\alpha} \frac{\sqrt{\alpha}}{\sigma_x \sigma_s} \Phi_{xs}(\alpha) .
\end{aligned} \tag{6.4}$$

In this formulation, $\hat{\alpha}$ is expressed as the stretch factor that maximizes the inner product between the energy normalized signals $x(t)/\sigma_x$ and $s(\alpha t)/(\sigma_s/\sqrt{\alpha})$. In (6.4),

$$\Phi_{xs}(\alpha) = \int_0^{\infty} x(t)s(\alpha t) dt \tag{6.5}$$

is the scale cross-correlation function between $x(t)$ and $s(t)$ [121], which, we note, is not a function of time, but a function of the stretch factor α .

6.2.2 The Scale Transform

In this subsection, we briefly review the scale transform and discuss a couple of properties relevant to solving (6.4). The scale transform is defined as [121]

$$\mathcal{S} \{x(t); c\} = X(c) = \int_0^{\infty} x(t) t^{-jc-1/2} dt , \tag{6.6}$$

and the inverse scale transform is

$$\mathcal{S}^{-1} \{X(c); t\} = x(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} X(c) t^{jc-1/2} dc . \tag{6.7}$$

The scale transform is a special variant of the Mellin transform, which has several stretch-invariant properties and has been used in applications such as the classifying of ships from radar signals [122] and the interpreting of speech waveforms [123]. The Fourier-Mellin transform, which is derived from the Mellin transform of the magnitude of the Fourier transform, is also popular in several engineering fields [124–133] due to its invariant properties to both shifting and stretching in time.

To serve as an example, Figure 6.1a and 6.1b illustrate two experimentally measured signals, truncated to 2 ms, at different temperatures. Figure 6.1c and 6.1d show the magnitude of

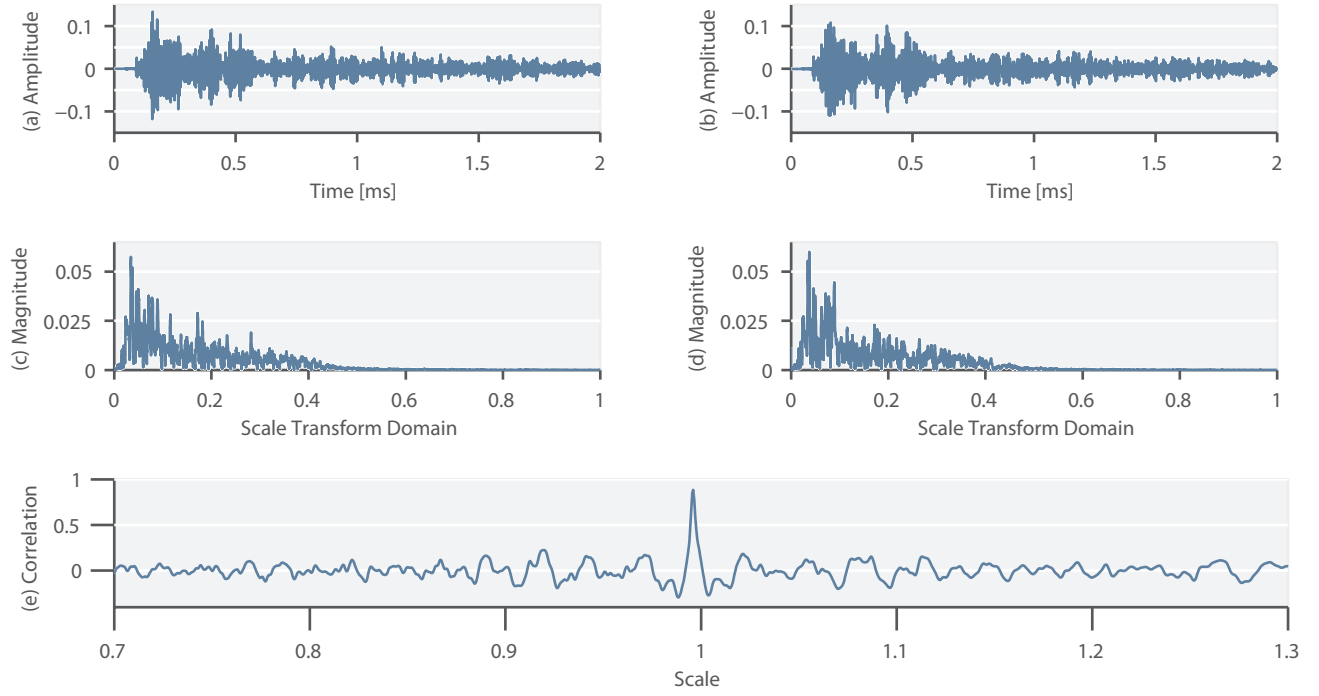


Figure 6.1: (a) An experimentally obtained guided wave signal measured at a temperature of 20.5°C. (b) An experimentally obtained guided wave signal measured at a temperature of 7.8°C. (c) The scale transform magnitude of the signal in Figure 6.1a. (d) The scale transform magnitude of the signal in Figure 6.1b. (e) The scale cross-correlation between of the signals in Figure 6.1a and Figure 6.1b. For convenience, the horizontal axis is zoomed into the region of interested.

the scale transforms of each measured signal. Figure 6.1e shows the normalized scale cross-correlation function, from (6.5), of the two guided wave signals. Note that, for many physical signals, the majority of the energy is often concentrated early in the scale transform domain. This is analogous to a frequency “low pass” signal, but in the scale transform domain [122, 133].

Intuitively, the Mellin and scale transforms are analogous to the Laplace and Fourier transforms but with delay or time-shifting operations replaced by stretching or time-scaling operations. Like the Fourier transform, the scale transform satisfies Parseval’s theorem,

$$\int_0^\infty x^*(t)s(t) dt = \frac{1}{2\pi} \int_{-\infty}^\infty X^*(c)S(c) dc, \quad (6.8)$$

where $(\cdot)^*$ denotes complex conjugate. This implies that signal energy is conserved between

the time and scale transform domains,

$$\int_0^\infty |x(t)|^2 dt = \frac{1}{2\pi} \int_{-\infty}^\infty |X(c)|^2 dc. \quad (6.9)$$

When applying the Fourier transform, a time delay operation translates to a phase shift in the frequency domain. Analogously, an energy-preserving stretch operation translates to a phase shift in the scale transform domain [121]

$$\mathcal{S} \{ \sqrt{\alpha} x(\alpha t); c \} = X(c) e^{j c \ln(\alpha)}. \quad (6.10)$$

We refer to this stretching as energy-preserving since

$$\int_0^\infty |\sqrt{\alpha} x(\alpha t)|^2 dt = \int_0^\infty |x(t)|^2 dt. \quad (6.11)$$

6.2.3 The Fast Mellin Transform

In this subsection, we briefly discuss the fast Mellin transform, a computationally efficient algorithm used to compute the Mellin or scale transform of a signal. Through algebraic manipulation and a change of variables such that $t = e^\tau$ and $dt = e^\tau d\tau$, the scale transform in (6.6) may be expressed as

$$\begin{aligned} \mathcal{S} \{ x(t); c \} &= \int_{-\infty}^\infty x(e^\tau) e^{(1/2)\tau} e^{-j c \tau} d\tau \\ &= \mathcal{F} \{ e^{(1/2)\tau} x(e^\tau); c \}. \end{aligned} \quad (6.12)$$

The expression in (6.12) shows that the scale transform is the Fourier transform of the exponentially time-skewed signal multiplied by an exponential envelope. With access to the fast Fourier transform, (6.12) provides a very simple and practical implementation for the scale transform. This implementation is known as the fast Mellin transform [122, 129, 134, 135].

From this expression, we may also represent the inverse scale transform as

$$\mathcal{S}^{-1} \{ X(c); t \} = e^{-(1/2) \ln(t)} \mathcal{F}^{-1} \{ X(c); \ln(t) \}. \quad (6.13)$$

In practice, we do not always compute the inverse scale transform. Performing the $t = \ln(\tau)$ substitution has several numerical complications. Instead, we can often extract the same information using the inverse Fourier transform. We demonstrate this in the following subsection.

Since most discrete signals are uniformly sampled in time, computation of $x(e^\tau)$ requires that we resample, or interpolate, the signal onto an exponential axis. This is often accomplished using cubic spline interpolation since it is computationally fast and provides a close approximation to sinc interpolation [136, 137]. For this application, however, we found linear interpolation to also provide good results. While linear interpolation is faster, both methods can be computed in linear time.

To compute the exponential axis, we assume the first sample of each data record does *not* correspond to time $t = 0 = e^{-\infty}$. This allows $x(e^\tau)$ to be finite in length. We also assume the sampling period to be 1 so that the first sample of each data record corresponds to time $t = 1 = e^0$. This allows $x(e^\tau)$ to be causal. Since uniformly sampling a signal $x(nT)$ with period T is equivalent to stretching a signal by a factor T , this assumption can be corrected for in the scale transform domain by applying a phase shift.

To properly resample the signal, we also need to ensure that no information is lost. This condition requires that the minimum sampling rate satisfies the Nyquist sampling criteria [138]. To ensure that the Nyquist sampling criteria is satisfied, the length of the resampled signal must be greater than or equal to $N \ln(N)$, where N is the length of the original truncated signal [134, 139, 140]. Figure 6.2 demonstrates the exponential resampling of a sine wave.

6.2.4 Maximization of the Scale Cross-Correlation Function

In this subsection, we discuss two strategies for maximizing the scale cross-correlation function. In Section 6.3, we will discuss the numerical benefits of these approaches.

Maximization in the Stretch Factor Domain α

The scale cross-correlation function $\Phi_{xs}(\alpha)$ of two signals, $x(t)$ and $s(t)$, can be computed from the scale transform as [121, 141]

$$\begin{aligned}\Phi_{xs}(\alpha) &= \int_0^\infty x^*(t) s(\alpha t) dt \\ &= \mathcal{S}^{-1} \{ X^*(c) S(c); \alpha \} .\end{aligned}\tag{6.14}$$




Figure 6.2: A demonstration of exponential sampling compared with uniform sampling. (a) A continuous sine wave. (b) A uniformly sampled sine wave. (c) An exponentially sampled sine wave.

To prove (6.14), we first substitute (6.13) into (6.14), apply the Fourier transform's convolution theorem, and finally set $\tau = \ln(t)$ to show

$$\begin{aligned}
 \mathcal{S}^{-1} \{X^*(c)S(c); \alpha\} &= e^{-(1/2)\ln(\alpha)} \mathcal{F}^{-1} \{X^*(c)S(c); \ln(t)\} \\
 &= \int_{-\infty}^{\infty} e^{\tau} x^*(e^{\tau}) s(e^{\tau+\ln(\alpha)}) d\tau \\
 &= \int_0^{\infty} x^*(t) s(\alpha t) dt \\
 &= \Phi_{xs}(\alpha) .
 \end{aligned} \tag{6.15}$$

The result in (6.14) computes $\Phi_{xs}(\alpha)$, a function of α , directly from the scale transform representations, $X(c)$ and $S(c)$, without requiring us to stretch the signal $s(t)$. The solution to the optimization problem in (6.4) can therefore be expressed as the normalized maximum of (6.14), i.e.

$$\hat{\alpha} = \arg \max_{\alpha} \frac{\sqrt{\alpha}}{\sigma_x \sigma_s} \mathcal{S}^{-1} \{X^*(c)S(c); \alpha\} \tag{6.16}$$

We also note that the scale-invariant correlation coefficient between $x(t)$ and $s(t)$ is defined by

$$\phi_{xs} = \max_{\alpha} \frac{\sqrt{\alpha}}{\sigma_x \sigma_s} \mathcal{S}^{-1} \{X^*(c)S(c); \alpha\} , \tag{6.17}$$




Figure 6.3: Block diagram for computing the scale cross-correlation function between $x(t)$ and $s(t)$.

where ϕ_{xs} is normalized such that $\phi_{xs} = 1$ when $x(t)$ is a stretched replica of $s(t)$.

As previously mentioned, we normally do not compute the inverse scale transform due to numerical complications. Therefore, to avoid this complication when computing $\Phi_{xs}(\alpha)$, we further manipulate (6.16) by substituting the inverse scale transform relationship in (6.13) and expressing the maximization as

$$\begin{aligned}
 \hat{\alpha} &= \arg \max_{\alpha} \frac{\sqrt{\alpha}}{\sigma_x \sigma_s} e^{-(1/2) \ln(\alpha)} \mathcal{F}^{-1} \{X^*(c)S(c); \ln(\alpha)\} \\
 &= \arg \max_{\alpha} \frac{1}{\sigma_x \sigma_s} \mathcal{F}^{-1} \{X^*(c)S(c); \ln(\alpha)\} \\
 &= \arg \max_{\alpha} \frac{1}{\sigma_x \sigma_s} \Psi_{xs}(\ln(\alpha)) .
 \end{aligned} \tag{6.18}$$

In (6.18), $\Psi_{xs}(\ln(\alpha))$ is now a function of the log-stretch factor. Since the natural logarithm is monotonic, the maximum with respect to α is equivalent to the maximum with respect to $\ln(\alpha)$. So, to avoid the inverse scale transform in (6.16), we instead compute the optimal stretch factor and scale-invariant correlation coefficient as

$$\hat{\alpha} = \exp \left(\arg \max_{\ln(\alpha)} \Psi_{xs}(\ln(\alpha)) \right) \tag{6.19}$$

$$\phi_{xs} = \max_{\ln(\alpha)} \frac{1}{\sigma_x \sigma_s} \Psi_{xs}(\ln(\alpha)) . \tag{6.20}$$

Figure 6.3 provides a block diagram illustrating each step for computing $\Psi_{xs}(\ln(\alpha))$.

Maximization in the Scale Transform Domain c

By applying Parseval's theorem in (6.8) and the time-stretching property in (6.10) to the scale-cross correlation function $\Phi_{xs}(\alpha)$, we can express the optimal stretch estimate in the scale

transform domain as

$$\begin{aligned}\hat{\alpha} &= \arg \max_{\alpha} \frac{\sqrt{\alpha}}{\sigma_x \sigma_s} \Phi_{xs}(\alpha) \\ &= \arg \max_{\alpha} \frac{1}{2\pi \sigma_x \sigma_s} \int_{-\infty}^{\infty} X^*(c) S(c) e^{jc \ln(\alpha)} dc\end{aligned}\quad (6.21)$$

In this formulation, we are able to directly change the stretch factor α by altering the phase of the scale transform. This is accomplished by multiplying either $X(c)$ or $S(c)$ by a complex exponential.

6.3 Performance of Scale Transform Methods

In this section, we present three algorithms for maximizing the scale cross-correlation function based on the formulations discussed in Section 6.2. We discuss each algorithm's implementation, resolution, and computational complexity. We then compare these algorithms with current methods and show that the scale transform based techniques are less computationally complex for a given resolution or signal length.

6.3.1 Scale-invariant correlation (SIC) method

The scale-invariant correlation (SIC) method maximizes $\Psi_{xs}(\ln(\alpha))$ directly in the log-stretch factor domain $\ln(\alpha)$ as expressed in (6.19). Note that, in practice, by sampling $x(t)$ and $s(t)$ in the time domain and truncating the signal to a length of N samples, the scale transform representations, $X(c)$ and $S(c)$, are represented only by a finite number of values. Since (6.19) then computes $\Psi_{xs}(\ln(\alpha))$ as the inverse Fourier transform of $X^*(c)S(c)$, we know that we can also only evaluate $\Psi_{xs}(\ln(\alpha))$ for a finite, discrete set of stretch factors α .

The resolution of the set of stretch factors is defined by the sampling interval of $\Psi_{xs}(\ln(\alpha))$. Assuming a unitary sampling period, $x(t)$ is defined over the range $1 \leq t \leq N$ and $x(e^\tau)$ is defined for $0 \leq \tau \leq \ln(N)$. Then, since $x(e^\tau)$ is of length $N \ln(N)$, the interval between each sample must be $1/N$. Therefore, since $\Psi_{xs}(\ln(\alpha))$ is related to $x(e^\tau)$ and $s(e^\tau)$ by a Fourier transform followed by an inverse Fourier transform, it must also have a sampling interval of $1/N$.

This implies that the smallest measurable deviation from $\alpha = 1$ is

$$\Delta\alpha = \exp\left(\pm\frac{1}{N}\right). \quad (6.22)$$

For sufficiently large values of N , we can approximate $\Delta\alpha$ by a first order Taylor series approximation to get

$$\Delta\alpha \approx 1 \pm \frac{1}{N}. \quad (6.23)$$

So the resolution of SIC is approximately $1/N$. Therefore, SIC is limited in resolution. However, we will show that we can improve this resolution by combining SIC with an iterative optimization approach.

We now calculate the computational complexity of the SIC method. To compute SIC, we need to exponentially resample $x(t)$ and $s(t)$, compute their Fourier transforms, compute an inverse Fourier transform, and then find the maximum of the result. Maximizing $\Psi_{xs}(\ln(\alpha))$ and exponentially resampling $x(t)$ can both be computed in linear time. If we assume $x(e^\tau)$ to be of length $N \ln(N)$, then the computational complexity of computing its Fourier transform, using the fast Fourier transform algorithm, is $\mathcal{O}(N \ln(N) \log(N \ln(N)))$, or $\mathcal{O}(N \ln(N) \log(N))$ after simplifying. Since this is the most computationally expensive operation in SIC, the computational complexity of SIC is also $\mathcal{O}(N \ln(N) \log(N))$.

6.3.2 Iterative scale transform (IST) method

The iterative scale transform (IST) method maximizes the scale cross-correlation function $\Phi_{xs}(\alpha)$ by phase shifting $X^*(c)$ or $S(c)$ in the scale transform domain c as shown in (6.21). Solving this optimization problem iteratively in the scale transform domain c allows IST to have a very high precision. However, as shown in Figure 6.1e, the scale cross-correlation is not (globally) convex, but is locally convex around multiple maxima. In the next subsection, we address this issue further.

To compute the stretch factor estimate $\hat{\alpha}$ using IST, we first compute the scale transforms $X^*(c)$ and $S(c)$. As with SIC, the complexity of these operations is $\mathcal{O}(N \ln(N) \log(N))$. We then choose an initial guess for α , multiply $S(c)$ (or $X^*(c)$) by $e^{jc \ln(\alpha)}$, and then compute the inner product between $X^*(c)$ and $S(c)e^{jc \ln(\alpha)}$. Each of these operations has a linear complexity.

This process of choosing an α , applying a phase shift, and computing an inner product is then repeated for different values of α by a convex optimization algorithm until the inner product converges to a maximum value. The complexity of most convex optimization algorithms, neglecting special cases, is $\mathcal{O}(M^2)$ where M is the number of parameters to optimize across [55]. For this application, we only optimize across one variable α , so $M = 1$ and the complexity is constant. Therefore, the complexity of the optimization procedure is $\mathcal{O}(N \ln(N))$ for each iteration, where $N \ln(N)$ is the number of samples in the scale transform domain.

We can also improve the computational speed of IST by taking advantage of the structure of the scale transform. The majority of the energy in a signal is often located early in the scale transform domain. Therefore, we can truncate a large portion of the domain with little loss of information. As a result, the cost of the iterative algorithm becomes $\mathcal{O}(\rho N \ln(N))$, where ρ represents the percentage of the scale transform domain retained after truncation. In Section 6.4, we demonstrate that we can reduce ρ to 0.25 with only a small change to the scale estimate.

6.3.3 SIC/IST combination

As previously discussed, IST is a very precise estimation strategy but only if the result converges to the global maximum. In contrast, SIC requires no assumption of convexity but has a finite resolution. By combining these two methods, we can have highly precise estimates and guarantee convergence to the global maximum. This is done by using SIC to generate the initial stretch factor estimate $\hat{\alpha}$ for IST. In general, the SIC estimate will lie within the locally convexity region around the global maximum of the scale cross-correlation function $\Phi_{xs}(\alpha)$.

Note however that if N is small enough such that SIC cannot adequately resolve the main lobe of the scale cross-correlation function, then the SIC estimate may not be accurate and IST may not be guaranteed to converge to the globally optimal result. However, for sufficiently large values of N , this is not an issue. In our experimental results in Section 6.4, where $400 \leq N \leq 10000$, the problem never arose.

Since IST already computes the scale transform representations, $X(c)$ and $S(c)$, the only additional step required when combined with SIC is the computation of the inverse Fourier transform in (6.18) and maximization over $\Psi_{xs}(\alpha)$ in (6.19). The computational complexity of these operations is $\mathcal{O}(N \ln(N) \log(N))$, the same as initially computing $X(c)$ and $S(c)$. Therefore,

Methods	Resolution	Computational Complexity	Iterative Complexity
<i>Finite Resolution Methods</i>			
Search OSS	$1/N$	$\mathcal{O}(RN^2)$	–
SIC	$1/N$	$\mathcal{O}(N \ln(N) \log(N))$	–
<i>Fine Resolution Methods</i>			
Search/Iterative OSS	–	$\mathcal{O}(RN^2)$	$\mathcal{O}(N)$
SIC/IST	–	$\mathcal{O}(N \ln(N) \log(N))$	$\mathcal{O}(\rho N \ln(N))$

Table 6.1: The computational complexity of each temperature compensation discussed.

these operations do not change the overall computational complexity of IST and the complexity of SIC/IST is equivalent to the computational complexity of IST.

6.3.4 Comparison with the optimal signal stretch (OSS) method

OSS [39,40,115] is another optimization strategy for estimating the optimal stretch factor $\hat{\alpha}$ between two signals. As with our scale transform based methodology, OSS also defines $\hat{\alpha}$ as the stretch factor that minimizes the squared error between the received signal $x(t)$ and a stretched baseline signal $s(\alpha t)$. We differentiate SIC and IST with OSS by their different approaches to implementation.

OSS directly computes $s(\alpha t)$ and its associated squared error for several values of α . This time-stretching operation is computed by interpolation, which can be done in several ways. One common method is to apply truncation and zero-padding operations in the time and frequency domains [115]. This is a relatively efficient method for computing $s(\alpha t)$, but limits the resolution of α . For a fixed value of N , the smallest computable deviation from a unitary stretch of $\alpha = 1$ is approximately $1/N$.

The time-stretching operation can also be accomplished by interpolation directly in the time domain. This approach is not limited in resolution [142]. Sinc interpolation would be theoretically ideal but is computationally slow. Therefore, cubic spline and linear interpolation, which are both $\mathcal{O}(N)$ fast, may be used instead to approximate sinc interpolation.

There are then two common strategies for solving the OSS optimization problem. In this chapter, we will refer to these strategies as search OSS and iterative OSS.

Search OSS

Search OSS is a simple “brute force” optimization method that performs an optimization over a finite set of stretch values. In this procedure, a baseline library $s(\alpha t)$ is computed for different values of α in an interval of length R . The optimal estimate $\hat{\alpha}$ is then the stretch value of the library signal that minimizes the squared error with $x(t)$. This approach is similar to some data-driven approaches, such as optimal baseline selection [115]. However, in this situation the library is generated from a single baseline signal rather than a collection of measured data.

We compute the squared error between $x(t)$ and each library signal as a matrix-vector multiplication. To parallel SIC, we choose the resolution of search OSS to be $1/N$. Given that resolution, we would need a library of RN baselines to uniformly cover a range R . Therefore, the computational complexity of searching the library is $\mathcal{O}(RN^2)$ and the computational complexity of generating every baseline in the library, with linear or spline based time domain interpolation, is also $\mathcal{O}(RN^2)$.

Iterative OSS

Like IST, iterative OSS uses local convexity to compute $\hat{\alpha}$ without a baseline library. The optimal value is found by an iterative algorithm. As with IST, iterative OSS may be solved using a variety of convex optimization algorithms. Since the convex optimization has a constant complexity, the overall computational complexity of iterative OSS is given by the $\mathcal{O}(N)$ computation of $s(\alpha t)$ by linear or spline based interpolation at each iteration of the algorithm.

Search/Iterative OSS

As with SIC and IST, we can combine the search and iterative strategies to improve overall performance. Search OSS performs a coarse search and iterative OSS performs a fine search with the initial condition taken from search OSS. This combined method possesses a fine resolution and has a computational complexity of $\mathcal{O}(RN^2)$ from search OSS plus $\mathcal{O}(N)$ from iterative OSS.

Table 6.1 provides a concise summary of each method implemented in the following sections. We choose not to focus on the iterative methods alone since they will always perform poorly for sufficiently large variations in temperature. The table shows separately the computational complexity of the one-time computations and the complexity of the each iteration of the



Figure 6.4: The experimental setup shows the aluminum plate and the two PZT (lead zirconate titanate) transducers used to monitor the plate.

optimization algorithms. The results show that SIC and SIC/IST to have smaller overall computational complexities than search OSS and search/iterative OSS. In Section 6.4, we confirm this experimentally.

6.4 Results: Single sensor compensation

In this section, we discuss the experiments used to test the scale transform techniques. We discuss our physical experimental setup, signal preprocessing steps, and choice of signal excitation. We also briefly discuss the iterative convex optimization algorithms used by iterative OSS and IST.

6.4.1 Experimental methodology

In our experiment, we excite and measure guided ultrasonic waves on a thin plate under variable temperature conditions. To generate the guided waves, we used a pair of synchronized lead zirconate titanate (PZT) piezoelectric transducers permanently bonded to the surface of a 9.8 cm wide by 30.5 cm long by 0.1 cm thick aluminum plate. The aluminum plate is shown in Figure 6.4. Guided wave signals are recorded for 10 ms at a sampling rate of 1 MHz using National Instruments PXI data acquisition equipment.

For 36.3 hours between 11:30 AM and 11:50 PM of the following day, the aluminum plate was cooled and warmed by adjusting its ambient temperature. During this time, guided waves

were synchronously generated and measured every two minutes. Using a thermocouple, the ambient temperature was also measured every one minute.

At 6:18 PM on the first day, a cylindrical, steel, grease-coupled mass with a diameter of 3.8 cm and height of 4.5 cm, was placed on top of the aluminum plate to scatter waves and simulate damage. Although the mass may not perfectly simulate damage, it changes the propagation environment in ways unlike temperature. At 4:04 PM on the second day, the mass was then removed from the plate.

During data collection, the acquisition equipment applied a low-pass filter with a cutoff frequency of 500 kHz to each analog signal. After measuring each signal, a high-pass filter with a 3 kHz cutoff frequency was also applied to remove systematic low frequency noise in the system. To eliminate any phase effects introduced by the excitation signal, each measured signal was correlated with the excitation waveform.

We chose to transmit a wideband, impulsive sinc excitation with a center frequency of 250 kHz and flat spectral bandwidth of 400 kHz. The wideband excitation helps to satisfy the diffuse field limit conditions [3, 41]. Under these conditions, the approximate time-stretch temperature model is more reliable.

We implement iterative OSS and IST with three different iterative algorithms: a quasi-Newton line search method [137], an active-set method [143], and an interior-point method [55]. Our results illustrate the mean results from the three algorithms.

The first algorithm is implemented using MATLAB's *fminunc* function. The other two algorithms are implemented using MATLAB's *fmincon* function. We also constrain the active-set and interior-point methods to solutions based on the resolution of search OSS and SIC. Therefore we constrain the solution space to $-1/(2N) < \alpha - \hat{\alpha} < 1/(2N)$, where $\hat{\alpha}$ is the estimate from search OSS or SIC.

6.4.2 Estimation accuracy

We first carry out a simulation study where we stretch the experimental data by a factor of $\alpha = 1.001278$, which is comparable to stretch factors observed in our experiment. For each measurement taken, we use the unaltered signal as the single baseline for estimating α . Figure 6.5 shows the results of estimating this stretch factor with four methods: search OSS, iterative




Figure 6.5: Estimation error ($\hat{\alpha} - \alpha$) of SIC, SIC/IST, search OSS, and search/iterative OSS under simulated ideal conditions. In the simulation, experimental measurements were numerically stretched by a factor of 1.001278 and compared with the original signal to estimate the stretch factor.

OSS, SIC, and SIC/IST. The results confirm that the estimation error for search OSS and SIC stay within their resolution bounds ($\pm 1/(2N)$) while iterative OSS and SIC/IST approaches a small error (on the order of 10^{-6}) very quickly. This study shows that all four methods can successfully estimate the 1.001278 stretch factor.

We then apply our SIC/IST methodology to the experimental data. We use the first measurement, taken at 20.5°C, as a single baseline for estimating the stretch factor between itself and the other measurements. Figure 6.6 shows the the reciprocal of the stretch factor estimate $1/\hat{\alpha}$ plotted with the ambient temperature measured during the experiment. The results show a strong linear correlation between the stretch factor estimate and temperature. Applying SIC, search OSS, or search/iterative OSS yields similar results. This study verifies the reliability of the time-stretch model used by these methods.

6.4.3 Computational cost

We evaluate the computational cost of each algorithm as the average computation time required to process each measurement. We evaluate each method for 41 different record lengths




Figure 6.6: Comparison between the stretch factor estimate of experimental data, computed using SIC/IST, with ambient temperature. Results show a close correlation between values.

N , uniformly spanning from $N = 400$ to $N = 10000$. For search OSS, we consider a search space of $(1 - R/2) \geq \alpha \geq (1 + R/2)$ for a range $0.008 \geq R \geq 0.024$. Note that, as shown in Figure 6.6, the true stretch factor in our experiment varies between approximately 1 and 1.005. For search/iterative OSS and SIC/IST, we also evaluate the computational effort in terms of the average number of iterations required for convergence. We define a single iteration as a single call to the function that stretches the baseline and compares the two signals.

Since the iterative algorithm convergence rate may vary widely with N , we normalize the average computation time to be

$$\bar{\tau}_n = \frac{\tau_n}{c_n} \frac{1}{K} \sum_k c_k, \quad (6.24)$$

where τ_n and c_n represent the average computation time and average number of iterations used for records of some length specified by n . This helps to reduce anomalies from record lengths with poor convergence rates.

We first compare the computational effort required to compute search OSS and SIC. Figure 6.7 shows the average computation time for these techniques as a function of the record length N (the reciprocal of resolution). Figure 6.7 clearly shows that the computational effort of search/iterative OSS grows quadratically with the number of samples N while SIC grows




Figure 6.7: The average, normalized computation time used to compute the search OSS and SIC stretch factor estimates. The search OSS methodology is illustrated for multiple stretch factor ranges $0.008 \geq R \geq 0.024$. Data is fit with curves to illustrate trends.




Figure 6.8: The average, normalized computation time used to compute the search/iterative OSS and SIC/IST stretch factor estimates. The search/iterative OSS methodology is illustrated for multiple stretch factor ranges $0.008 \geq R \geq 0.024$. Data is fit with curves to illustrate trends.

linearly. At $N = 10,000$, SIC shows a 2.3 (for $R = 0.008$) to 6.9 (for $R = 0.024$) times improvement in computational speed over search OSS. Note that SIC's variations in computational effort largely due to MATLAB's fast Fourier transform implementation.

We now compare the computational effort required to compute search/iterative OSS and SIC/IST. Figure 6.8 shows the average computation time for these methods as a function of the record length. As with Figure 6.7, Figure 6.8 shows that the computational effort of search/iterative OSS increases quadratically with the number of samples N while SIC/IST grows linearly. At $N = 10,000$, SIC/IST shows a 1.5 (for $R = 0.008$) to 3.6 (for $R = 0.024$) times improvement in computational speed over search/iterative OSS.

We also investigate the effect of reducing the number of samples used by SIC/IST to estimate α in the scale transform domain. We truncate the scale transform domain and retain a “low pass” approximation of the domain with $\rho N \ln(N)$ samples. Here ρ represents the percentage of samples kept by SIC/IST. In Figure 6.9, we show that we can speed up SIC/IST by up to 1.3




Figure 6.9: The average, normalized computation time used to compute the SIC/IST stretch factor estimate after truncating the scale transform domain to various lengths. The value ρ represents the percentage of scale domain retained. Data is fit with curves to illustrate trends.

times by removing up to 75% of the scale transform domain information. Figure 6.8 shows that the mean absolute change in the stretch factor estimate resulting from these truncations is small. For $\rho = 0.250$, the error is 12 (at $N = 400$) to 100 (at $N = 10,000$) times below the resolution of SIC ($1/N$). For $N > 5000$, the error is comparable to or smaller than the resolutions of search/iterative OSS and SIC/IST, as measured from Figure 6.10.

Figure 6.11 shows the number of iterations required for SIC/IST and search/iterative OSS to converge, roughly 25.5 iterations for search/iterative OSS and 17.5 iterations for SIC/IST. On average, OSS requires approximately 1.45 times more iterations to converge. For clarity, we only show two curves as the results vary little for different values of R and ρ .

We note that SIC/IST also requires much less storage compared with OSS for large values of N or R . Search/iterative OSS requires a total of $\mathcal{O}(RN^2)$ doubles while SIC/IST only requires $\mathcal{O}(N \ln(N))$ doubles.




Figure 6.10: The average absolute change in the stretch factor estimate of SIC/IST after truncating the scale transform domain to various lengths. The value ρ represents the percentage of scale domain utilized.




Figure 6.11: The number of iterations used when computing the SIC/IST and search/iterative OSS methods.

6.5 Results: Data-driven matched field integration

In this section, we combine our scale transform temperature compensation method with data-driven matched field processing, originally discussed in Chapter 3 and implemented in Chapter 4 and Chapter 5. We demonstrate that integration of the scale transform temperature compensation with data-driven matched field processing can be used to improve localization results.

Prior work has shown how stretch-based temperature compensation methods can improve the performance of delay-and-sum localization methods, which were discussed in Chapter 4. Unlike delay-and-sum localization, data-driven matched field processing utilizes all the phase information of the signal, and therefore is affected by temperature in a different manner. This section integrates scale transform temperature compensation to incoherent data-matched field processing and demonstrates that the scale transform can significantly improve our localization performance in variable environmental conditions. In our demonstration, we use the incoherent processor because we assume that the frequency characteristics of signals reflecting from our scatterer, a large mass, are unknown.

6.5.1 Scale transform temperature compensation

As described in Chapter 3, data-driven matched field processing generates a data-driven model $\hat{\mathbf{x}}(\mathbf{r})$, described in Section 3.2.3, from a collection of calibration data \mathbf{y} , described in Section 3.2.1. The matched field processor then compares the model $\hat{\mathbf{x}}(\mathbf{r})$ with a collection of test data \mathbf{x} , described in Section 3.2.2. For the purpose of discussion, let $x(r_1^*, t), \dots, x(r_M^*, t)$ be the M time domain test measurements associated with true distances r_1^*, \dots, r_M^* between pairs of sensors and the scatterer, such as damage in the structure, such that

$$\mathbf{x} = [\mathcal{F}\{x(r_1^*, t)\} \cdots \mathcal{F}\{x(r_M^*, t)\}]^T. \quad (6.25)$$

For simplicity of discussion, we omit noise from the expressions in this section.

Now let $y_c(r_1^*, t), \dots, y_c(r_M^*, t)$ and $y_b(r_1^*, t), \dots, y_b(r_M^*, t)$ represent the collection of current measurements and a collection of baseline measurements taken from our sensor system. These measurements include both a signal originating from the scatterer as well as the direct

signal originating from the transmitting sensor. In traditional systems, we define the test data $x(r_m^*, t)$ for each measurement m as the difference between our current measurements $y_c(r_m^*, t)$ and baseline measurements $y_b(r_m^*, t)$ such that

$$x(r_m^*, t) = y_c(r_m^*, t) - y_b(r_m^*, t) . \quad (6.26)$$

In ideal conditions, the result of the baseline subtraction yields only the scattered signal from the damage. However, if the current measurements are taken at a temperature sufficiently different from the temperature at which the baseline data is acquired, standard baseline subtraction will yield a different result.

We use the scale transform, specifically SIC/IST as defined in Section 6.3.3, to estimate the scale that best relates our current measurements to the baseline signal. We then stretch the current measurements to best resemble the baseline data. We use the baseline data $y_b(r_m^*, t)$ as our calibration data

$$\mathbf{y} = [\mathcal{F}\{y(r_1^*, t)\} \cdots \mathcal{F}\{y(r_M^*, t)\}]^T , \quad (6.27)$$

so the stretching process will adjust the current measurements $y_c(r_m^*, t)$ to have similar velocity characteristics as the baseline data $y_b(r_m^*, t)$, the test data $x(r_m^*, t)$, and the data-driven model $\mathbf{x}(\hat{\mathbf{r}})$. The test data $x(r_m^*, t)$ is then defined by

$$x(r_m^*, t) = y_c(r_1^*, \hat{\alpha}t) - y_b(r_1^*, t) , \quad (6.28)$$

where $\hat{\alpha}$ is the optimal stretch factor between $y_b(r_m^*, t)$ and $y_c(r_m^*, t)$. We use this adjusted test data to localize the scatterer.

6.5.2 Experimental methodology

We consider a 1.22 m by 1.22 m by 0.2844 cm aluminum plate with sixteen 0.7 cm by 0.7 cm PZT (lead zirconate titanate) transducers randomly distributed on its upper surface. We use the same sensor configuration as shown in Figure 5.1. As in Chapter 4, we collect the calibration data to compute the frequency-wavenumber representation of the plate by transmitting and measuring signals between each of the transducers. This results in a total of 120 unique measurements. We use frequency wavenumber synthesis to recover the frequency-wavenumber

representation, or dispersion curves. From the frequency-wavenumber representation, we then generate the data-driven model $\mathbf{x}(\hat{\mathbf{r}})$ of the medium using the sparse wavenumber synthesis, as discussed in Chapter 2.

We use the data-driven model to then localize a metal, cylindrical mass placed on the top of the aluminum plate. The mass has a diameter of approximately 5 cm and is used to act as a weak wave scatterer as a substitute for irreversible damage. Due to the relatively large size of the mass compared to the 0.75 cm diameter holes from Chapter 4, we utilize lower frequencies (longer wavelengths) to localize it. We implement incoherent data-driven matched field processing, derived in Chapter 3, with 30 frequencies uniformly spanning from 0 kHz to 72.5 kHz.

6.5.3 Results and discussion

Each of the plots in Figure 6.12 and Figure 6.13 illustrate the results from applying the incoherent data-driven matched field processing in three separate scenarios: (1) no temperature change, (2) temperature change without compensation, and (3) temperature change with compensation. In each plot in Figure 6.12, we illustrate a 0.5 m by 0.5 m region of the plate centered around the mass location. The plots in Figure 6.13 show the magnified 6 cm by 6 cm region around the mass location. The largest value in each figure indicates the estimated location of the mass. The circles in each figure indicate the known center location of the mass, and the squares indicate sensor locations. Note that not all sensor locations are shown since they are outside of the 0.5 m by 0.5 m region. In Figure 6.13, the maximum value in each plot is marked by a cross.

Figure 6.12a and Figure 6.13a show the results when there is approximately no temperature change. In this figure, the incoherent data-driven matched field processor accurately localizes the mass center. The maximum value is located 0.21 cm away from the known mass center location, a value significantly smaller than the 5 cm diameter of the mass.

Figure 6.12b and Figure 6.13b show the results when there is temperature change (a mean scale factor $\hat{\alpha} = 1.0018$ across all sensors), but scale transform temperature compensation is not applied. In these figures, there are more artifacts throughout the image and there is no clear maximum value near the mass's center. The largest value in the region 0.5 m by 0.5 m region is located 14.93 cm away from the known mass center location, a value approximately three




Figure 6.12: The incoherent data-driven ambiguity functions of a 0.5 m by 0.5 m region to localize a 5 cm cylindrical mass on an aluminum plate for three temperature scenarios: (a) no temperature change, (b) temperature change without compensation, and (c) temperature change with compensation. The squares and circles denote the locations of the sensors (only 2 out of 16 sensors fall in the region shown) and the mass, respectively.

times larger than the 5 cm mass diameter and falls outside the region shown in Figure 6.13b. Therefore, when there is only a small temperature change, we can no longer localize the mass.

The signals used in Figure 6.12c and Figure 6.13c incorporate the same temperature change (a mean scale factor $\hat{\alpha} = 1.0018$ across all sensors) as in Figure 6.12b and Figure 6.13b. However, we now use scale transform temperature compensation to better match the current measurements with the baseline measurements. In Figure 6.12c and Figure 6.13, the maximum value is located 0.30 cm away from the mass center location, a value significantly smaller than the 5 cm diameter of the mass. Scale transform temperature compensation successfully accounted for variations due to temperature, resulting in improved localization performance.




Figure 6.13: The incoherent data-driven ambiguity functions of a 0.06 m by 0.06 m region to localize a 5 cm cylindrical mass on an aluminum plate for three temperature scenarios: (a) no temperature change, (b) temperature change without compensation, and (c) temperature change with compensation. The squares and circles denote the locations of the sensors (only 2 out of 16 sensors fall in the region shown) and the mass, respectively. The crosses denote the estimated locations of the mass.

6.6 Conclusions

In this chapter, we discussed the scale transform and its role in developing three model-based optimal temperature compensation methods: the scale-invariant correlation method, the iterative scale transform method, and a combination of the two. These methods were compared with the optimal signal stretch optimization technique, which we demonstrated to be limited in computation speed due to its need to directly stretch signals. The scale transform based techniques circumvent these limitations by computing the results in the stretch factor and scale transform domains.

We showed that the scale transform methods have an approximately linear computational complexity while the optimal signal stretch methods have quadratic complexity. We also demonstrated the scale transform methods, for signals of length 400 to 10000, to be up to 6.9 times faster than other optimal methods with equal resolutions over reasonable search spaces. Furthermore, we integrated the scale transform temperature compensation method with the incoherent data-driven matched field processor and demonstrated that the scale transform can successfully compensate for small variations in temperature and improve localization performance.

Conclusions and Future Work

In this dissertation, we presented a method for recovering the frequency-wavenumber representation, or dispersion curves, of a collection of guided wave structural health monitoring measurements through the use of compressed sensing and ℓ_1 sparse recovery methods. We refer to this process as *sparse wavenumber analysis*. We then demonstrated, through a process we refer to as *sparse wavenumber synthesis*, the capability to use the frequency-wavenumber representation to generate a data-driven model of the guided wave environment. We integrated this data-driven model with a localization framework known as matched field processing, where the data-driven model replaced computationally expensive, often unreliable, numerical models of the environment. The new data-driven matched field processing methodology was tested for localizing holes and acoustic emission events in an aluminum plate. The results were shown to be more accurate and better resolved than conventional structural health monitoring approaches. We then integrated the data-driven matched field processor with scale transform temperature compensation to achieve improved robustness to environmental and operational effects, such as temperature.

Through this work, we have made several novel contributions to the current literature.

- **SPARSE WAVENUMBER ANALYSIS** In Chapter 2, We accurately recovered dispersion curves from guided wave data that has been corrupted by multipath interference. [20, 49]
- **SPARSE WAVENUMBER SYNTHESIS** In Chapter 2, we also synthesized and predicted guided wave signals from sparse wavenumber analysis results, generating a data-driven

model of an environment. Through simulations, we predicted measurements with accuracies greater than 97%. From experimental data, we demonstrated similar results.

- **DATA-DRIVEN MATCHED FIELD PROCESSING** In Chapter 3, we integrated data-driven models from sparse wavenumber synthesis with matched field processing, yielding data-driven matched field processing, and analyzed the new methodology's asymptotic localization performance. We found data-driven matched field processing to be theoretically accurate to within a small region for a sufficient number of sensors. [50, 85, 86]
- **COHERENT DATA-DRIVEN MATCHED FIELD PROCESSING** In Chapter 4, we localized two holes in an aluminum plate using coherent data-driven matched field processing. We demonstrated a 5 times improvement in accuracy and a 49 times improvement in resolution over conventional localization methods. [50]
- **INCOHERENT DATA-DRIVEN MATCHED FIELD PROCESSING** In Chapter 5, we localized acoustic emissions on an aluminum plate using incoherent data-driven matched field processing. We demonstrated a 2 times improvement in accuracy and a 12.5 times improvement robustness to noise.
- **SCALE TRANSFORM TEMPERATURE COMPENSATION** In Chapter 6, we reduced the distorting effects of temperature on ultrasonic signals with the computational efficient scale transform. The scale transform was shown to be up to 6.9 times faster than other approaches with near identical performance. [5, 119, 120]
- **SCALE TRANSFORM INTEGRATED DATA-DRIVEN MATCHED FIELD PROCESSING** In Chapter 6, we also integrated the scale transform with data-driven matched field processing to achieve localization that is robust to environmental variations.

The combination of these signal processing methods presents a strong framework for structural health monitoring that can detect and locate damage in large, physical infrastructures. Yet, there are many ways to improve these methods and expand on them. In the following sections, we discuss several possible directions for the future of this work.

7.1 Extension to new applications

For our future work, we plan to extend the structural health monitoring framework and signal processing methods developed to generate data-driven models to more complex structures and materials that are found in modern bridges, buildings, pipelines, rail lines, airplanes, and wind turbines. Many of these structures exhibit inhomogeneous or anisotropic properties as well as unique geometric features that further complicate analysis and processing. Pipes, for example, have a periodic boundary condition that creates significant multipath behavior. In our future work, we plan to develop methods that will incorporate this additional information to produce accurate and reliable data-driven models for the structure.

Complex guided wave propagation environments are also found in many other disciplines, including seismology [7], underwater acoustics [8], power systems [9], medical ultrasound [144]. Due to relatively similar geometries and wave characteristics encountered in these fields and structural health monitoring current applications, the existing work can be directly applied to a variety of problems in these disciplines. For particularly complex environments, we can construct data-driven models from rigorous analysis and understanding of the complex wave propagation or from integrating numerical modelling techniques, such as finite element analysis. In seismology and geotechnical engineering, these methods could help improve site characterization algorithms, which are used to determine the properties and integrity of soil prior to laying the foundations for new construction. In power systems applications, these techniques can be used to help locate faults in transmission lines or wiring. In underwater acoustics, the methods can be used to better track marine vehicles and aquatic animals. In medical and biological acoustics, data-driven models could improve our estimates of material and elastic properties of biological structures, which are common indicators for diagnosing tumors, lesions, and other ailments.

7.2 Integration with data science

In chapter 6, we refined data-driven matched field processing with the scale transform to compensate for variations in environmental temperature. In controlled, laboratory experiments these methods have shown to be very effective with relatively little data. However, for uncon-

trolled, real-world environments, more data and modelling is necessary to distinguish damage from benign environmental and operation changes. Furthermore, structural damage does not form immediately. Instead, it evolves slowly over many years of wear. Fortunately, large-scale structural health monitoring systems can collect vast amounts of data for processing and analysis over these periods of time. For these practical scenarios, we need to develop computational efficient methods to model and process the large reservoirs of data.

We have explored some effective approaches for processing large sets of data through machine learning [145–148]. We plan to further utilize methods from data science and machine learning to best store, manage, process, and data-mine our large data sets. The large volumes of data will likely be measured from large-scale sensor networks. Many questions in implementing these networks will need to be addressed, including power management, communication protocols, distributed data processing, and data storage. Many researchers have begun to propose solutions to these problems through the development of new technology for sensor networks, communication systems, energy harvesting, and distributed computing, but there is still significant work necessary to improve and integrate these methods. We plan to further investigate areas through collaboration with experts in each field.

7.3 Application to imaging modalities

Data-driven models are built by populating a general, theoretical framework of a physical system with information extracted from experimental data. This approach allows us to develop a holistic understanding about a medium by utilizing general knowledge about an environment with specific knowledge from experiments. We applied data-driven models to problems in acoustics and ultrasound in order to recover velocity information from guided waves. However, these concepts have strong applicability in other imaging modalities. In medical imaging, for example, incomplete knowledge about material and geometric properties is common and creates significant challenges for image-generation [149]. Multi-modal medical imaging approaches, which utilize data from more than one modality, are being explored to address these challenges by using a variety of information obtained from different modalities. Yet in these approaches, it is still necessary to fuse this information with some underlying knowledge about

the physical system.

In these scenarios, data-driven models may help to reduce uncertainties in each imaging modality and help fuse data between multiple modalities. We intend to further investigate these potential applications for data-driven models for medical imaging, including computed tomography (CT), magnetic resonance imaging (MRI), Electroencephalography (EEG), and Electrocardiography (ECG). Overall, there are many opportunities for improving and extending the sparsity-based, data-driven matched field processing framework and data-driven modelling concepts from this dissertation to new domains and new applications.

The restricted nullity property

Our proofs utilize two properties of random matrices: the restricted isometry property (RIP), defined in (3.20), and what we refer to as the restricted nullity property (RNP), which we derive here. A matrix \mathbf{A} is said to be “nearly unitary” if it satisfies RIP with a small restricted isometry constant δ_s [62,92,93]. We use RIP to derive RNP, which considers two matrices whose columns are “nearly uncorrelated.” In this Appendix, we prove RNP for a general pair of matrices \mathbf{A} and \mathbf{B} . For the proofs in Appendix A and Appendix B, we use a problem-specific form of RNP, shown in (3.21).

We consider two arbitrary matrices \mathbf{A} and \mathbf{B} that both satisfy RIP with small constants δ_s^A and δ_s^B . We claim the columns of \mathbf{A} and \mathbf{B} are “nearly uncorrelated” if

$$\frac{1}{\sqrt{2}} (\mathbf{A} + \mathbf{B}) \tag{A.1}$$

also satisfies RIP with a small constant δ_s^{A+B} .

We can show this by plugging (A.1) into the RIP inequality in (3.20) to get

$$\begin{aligned} \frac{1}{2} \|\mathbf{A}\mathbf{v}\|_2^2 + \mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} + \frac{1}{2} \|\mathbf{B}\mathbf{v}\|_2^2 &\geq (1 - \delta_s^{A+B}) \|\mathbf{v}\|_2^2 \\ \frac{1}{2} \|\mathbf{A}\mathbf{v}\|_2^2 + \mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} + \frac{1}{2} \|\mathbf{B}\mathbf{v}\|_2^2 &\leq (1 + \delta_s^{A+B}) \|\mathbf{v}\|_2^2 . \end{aligned} \tag{A.2}$$

Subtracting each term in the inequality by $\frac{1}{2}(\|\mathbf{A}\mathbf{v}\|_2^2 + \|\mathbf{B}\mathbf{v}\|_2^2)$ provides

$$\begin{aligned} \mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} &\geq (1 - \delta_s^{A+B}) \|\mathbf{v}\|_2^2 - \frac{1}{2}(\|\mathbf{A}\mathbf{v}\|_2^2 + \|\mathbf{B}\mathbf{v}\|_2^2) \\ \mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} &\leq (1 + \delta_s^{A+B}) \|\mathbf{v}\|_2^2 - \frac{1}{2}(\|\mathbf{A}\mathbf{v}\|_2^2 + \|\mathbf{B}\mathbf{v}\|_2^2) . \end{aligned} \tag{A.3}$$

Applying the RIP inequality in (3.20) to $\|\mathbf{A}\mathbf{v}\|_2^2$ and $\|\mathbf{B}\mathbf{v}\|_2^2$ in (A.3) changes the bounds to

$$\begin{aligned} \mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} &\geq \left(1 - \delta_s^{A+B} - \frac{1}{2} [(1 + \delta_s^A) + (1 + \delta_s^B)] \right) \|\mathbf{v}\|_2^2 \\ \mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} &\leq \left(1 + \delta_s^{A+B} - \frac{1}{2} [(1 - \delta_s^A) + (1 - \delta_s^B)] \right) \|\mathbf{v}\|_2^2 . \end{aligned} \tag{A.4}$$

Finally, simplifying the expression in (A.4) provides the result

$$-2\delta'_s \|\mathbf{v}\|_2^2 \leq \mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} \leq 2\delta'_s \|\mathbf{v}\|_2^2 , \tag{A.5}$$

where $2\delta'_s = \delta_s^{A+B} + \delta_s^A/2 + \delta_s^B/2$. This expression tells us that the columns of \mathbf{A} and \mathbf{B} are “nearly uncorrelated” or $\mathbf{v}^H \mathbf{A}^H \mathbf{B} \mathbf{v} \approx 0$ when δ'_s is small, and \mathbf{v} contains only s non-zero values. Said in another way, all vectors \mathbf{v} with sparsity s are nearly in the null space of $\mathbf{A}^H \mathbf{B}$ when \mathbf{A} and \mathbf{B} satisfy this property with small δ'_s . We refer to this as the restricted nullity property (RNP) for a pair a matrices.

Proof of (3.23): single path scenario (coherent processor)

We now derive the target-to-artifact ratio of the coherent data-driven matched field processor when there is no unmodeled multipath interference. For our analysis, we assume the recovered frequency-wavenumber representation is correct up to a scaling factor, and therefore satisfies $\hat{\mathbf{v}} \cong \mathbf{v}$.

B.1 Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$

The coherent data-driven matched field processor in (3.15) at the target location $\mathbf{r} = \mathbf{r}^*$ simplifies to

$$\begin{aligned} b(\mathbf{r} = \mathbf{r}^*) &= \frac{\left| \mathbf{v}^H \overline{\Phi}^H(\mathbf{r}^*) \overline{\Phi}(\mathbf{r}) \mathbf{v} \right|^2}{\left\| \overline{\Phi}(\mathbf{r}) \mathbf{v} \right\|_2^2} \\ &= \left\| \overline{\Phi}(\mathbf{r}^*) \mathbf{v} \right\|_2^2. \end{aligned} \tag{B.1}$$

B.2 Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$

We can derive an upper bound by assuming \mathbf{r} to be sufficiently far from the target such that $\Phi^H(\mathbf{r})$ and $\Phi(\mathbf{r}^*)$ satisfy RNP. Under this condition, we can apply the RNP inequality in (A.5)

to the numerator to derive the upper bound

$$\begin{aligned} b(\mathbf{r}) &= \frac{\left| \mathbf{v}^H \overline{\Phi}^H(\mathbf{r}^*) \overline{\Phi}(\mathbf{r}) \mathbf{v} \right|^2}{\|\overline{\Phi}(\mathbf{r}) \mathbf{v}\|_2^2} \\ &\leq \frac{\|\mathbf{v}\|_2^4}{\|\overline{\Phi}(\mathbf{r}) \mathbf{v}\|_2^2} 4(\delta'_s)^2. \end{aligned} \quad (\text{B.2})$$

We now utilize part of the the RIP inequality, which states

$$\frac{\|\overline{\Phi}(\mathbf{r}) \mathbf{v}\|_2^2}{\|\mathbf{v}\|_2^2} \geq (1 - \delta_s), \quad (\text{B.3})$$

and we apply it to the denominator of (B.2). This simplifies the bound to

$$b(\mathbf{r}) \leq \frac{4(\delta'_s)^2}{1 - \delta_s} \|\mathbf{v}\|_2^2. \quad (\text{B.4})$$

B.3 Target-to-artifact ratio

Taking the ratios of (B.1) to (B.4) obtains the lower bound for the single path

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{\|\overline{\Phi}(\mathbf{r}^*) \mathbf{v}\|_2^2 (1 - \delta_s)}{\|\mathbf{v}\|_2^2 4(\delta'_s)^2}. \quad (\text{B.5})$$

By applying the RIP inequality in (B.3) once more to the numerator of this expression, the bound becomes

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{\|\overline{\Phi}(\mathbf{r}^*) \mathbf{v}\|_2^2 (1 - \delta_s)}{\|\mathbf{v}\|_2^2 4(\delta'_s)^2} \geq \frac{(1 - \delta_s)^2}{4(\delta'_s)^2}. \quad (\text{B.6})$$

Setting $\delta'_s = \delta_s$ then yields the expression in (3.23),

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{(1 - \delta_s)^2}{4\delta_s^2}. \quad (\text{B.7})$$

Proof of (3.23): single path scenario (incoherent processor)

In this appendix, we derive the target-to-artifact ratio of the incoherent data-driven matched field processor when there is no unmodeled multipath interference. As for the coherent processor, we assume the recovered frequency-wavenumber representation is correct up to a scaling factor, and therefore satisfies $\hat{\mathbf{v}} \cong \mathbf{v}$.

C.1 Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$

The coherent data-driven matched field processor in (3.15) at the target location $\mathbf{r} = \mathbf{r}^*$ simplifies to

$$\begin{aligned} b(\mathbf{r} = \mathbf{r}^*) &= \sum_{q=1}^Q \frac{|\mathbf{v}_q^H \Phi^H(\mathbf{r}^*) \Phi(\mathbf{r}) \mathbf{v}_q|^2}{\|\Phi(\mathbf{r}) \mathbf{v}_q\|_2^2} \\ &= \sum_{q=1}^Q \|\Phi(\mathbf{r}^*) \mathbf{v}_q\|_2^2. \end{aligned} \tag{C.1}$$

C.2 Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$

We can derive an upper bound by assuming \mathbf{r} to be sufficiently far from the target such that $\Phi^H(\mathbf{r})$ and $\Phi(\mathbf{r}^*)$ satisfy RNP. Under this condition, we can apply the RNP inequality in (A.5)

to the numerator to derive the upper bound

$$\begin{aligned}
b(\mathbf{r}) &= \sum_{q=1}^Q \frac{|\mathbf{v}_q^H \Phi^H(\mathbf{r}^*) \Phi(\mathbf{r}) \mathbf{v}_q|^2}{\|\Phi(\mathbf{r}) \mathbf{v}_q\|_2^2} \\
&\leq \sum_{q=1}^Q \frac{\|\mathbf{v}_q\|_2^4}{\|\Phi(\mathbf{r}) \mathbf{v}_q\|_2^2} 4(\delta'_s)^2.
\end{aligned} \tag{C.2}$$

We now apply part of the the RIP inequality, shown in (B.3), to the denominator of (C.2). This simplifies the bound to

$$b(\mathbf{r}) \leq \frac{4(\delta'_s)^2}{1 - \delta_s} \sum_{q=1}^Q \|\mathbf{v}_q\|_2^2. \tag{C.3}$$

C.3 Target-to-artifact ratio

Taking the ratios of (C.1) to (C.3) obtains the lower bound for the single path

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{\sum_{q=1}^Q \|\Phi(\mathbf{r}^*) \mathbf{v}_q\|_2^2 (1 - \delta_s)}{\sum_{q=1}^Q \|\mathbf{v}_q\|_2^2 4(\delta'_s)^2} = \frac{\|\overline{\Phi}(\mathbf{r}^*) \mathbf{v}\|_2^2 (1 - \delta_s)}{\|\mathbf{v}\|_2^2 4(\delta'_s)^2}. \tag{C.4}$$

By applying the RIP inequality in (B.3) once more to the numerator of this expression, the bound becomes

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{\|\overline{\Phi}(\mathbf{r}^*) \mathbf{v}\|_2^2 (1 - \delta_s)}{\|\mathbf{v}\|_2^2 4(\delta'_s)^2} \geq \frac{(1 - \delta_s)^2}{4(\delta'_s)^2}. \tag{C.5}$$

Note that for simplicity of analysis of the asymptotic behavior, we assume that the RIP constant δ_s for Φ is equivalent to the RIP constant for $\overline{\Phi}$. For a more comprehensive analysis, we would want to analyze each RIP constant as distinct values. To further simplify the analysis, we also set $\delta'_s = \delta_s$ and achieve the expression in (3.23),

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{(1 - \delta_s)^2}{4\delta_s^2}. \tag{C.6}$$

Proof of (3.25): multipath scenario (coherent processor)

We derive the target-to-artifact ratio of the coherent data-driven matched field processor when there is unmodeled multipath interference that is corrupting the measured data. As with the single path scenarios, we assume $\hat{\mathbf{v}} \cong \mathbf{v}$. To analyze the ambiguity function under multipath corruption, we substitute expression for the test data \mathbf{x}_q with unmodeled multipath interference in (3.24) into the coherent data-driven ambiguity function in (3.15). This results in an ambiguity function expressed by

$$b(\mathbf{r}) = \frac{\left| \mathbf{v}^H \left(\eta_0 \overline{\Phi}(\mathbf{r}^*) + \sum_{\ell=1}^L \eta_\ell \overline{\Phi}(\mathbf{r}_\ell) \right)^H \overline{\Phi}(\mathbf{r}) \mathbf{v} \right|^2}{\eta_0^2 \|\overline{\Phi}(\mathbf{r}) \mathbf{v}\|_2^2}. \quad (\text{D.1})$$

D.1 Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$

In the numerator, we apply the RNP inequality in (A.5) to each of the mismatched matrix pairs, i.e., $\mathbf{v}^H \overline{\Phi}(\mathbf{r}_\ell) \overline{\Phi}(\mathbf{r}^*) \mathbf{v}$ for $1 \leq \ell \leq L$. For simplicity, we assume each pair of matrices to have the same RNP constant. Applying this inequality results in the ambiguity function

$$b(\mathbf{r} = \mathbf{r}^*) \geq \frac{\left| \|\overline{\Phi}(\mathbf{r}^*) \mathbf{v}\|_2^2 - 2\delta'_s \eta^{-1} \|\mathbf{v}\|_2^2 \right|^2}{\|\overline{\Phi}(\mathbf{r}^*) \mathbf{v}\|_2^2}, \quad (\text{D.2})$$

where η is the signal-to-interference ratio as defined by

$$\eta = \frac{\eta_0}{\sum_{\ell=1}^L \eta_\ell} . \quad (\text{D.3})$$

We then apply a form of the triangle inequality that states that any scalar values a and b must satisfy $|a - b|^2 \geq ||a|^2 - |b|^2|$, where $a = \|\overline{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^2$ and $b = 2\delta'_s\eta^{-1}\|\mathbf{v}\|_2^2$. We assume $\delta'_s\eta^{-1}$ is sufficiently small such that $||a|^2 - |b|^2| = |a|^2 - |b|^2$. Applying the triangle inequality changes the lower bound to

$$\begin{aligned} b(\mathbf{r} = \mathbf{r}^*) &\geq \frac{\|\overline{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^4 - 4(\delta'_s)^2\eta^{-2}\|\mathbf{v}\|_2^4}{\|\overline{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^2} \\ &= \|\overline{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^2 - \frac{\|\mathbf{v}\|_2^4}{\|\overline{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^2} \left(\frac{4(\delta'_s)^2}{\eta^2} \right) . \end{aligned} \quad (\text{D.4})$$

By then applying the RIP inequality in (B.3) to the denominator, the bound changes to

$$b(\mathbf{r} = \mathbf{r}^*) \geq \|\overline{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^2 - \left(\frac{4(\delta'_s)^2}{\eta^2(1 - \delta)} \right) \|\mathbf{v}\|_2^2 . \quad (\text{D.5})$$

This expression shows us that the lower bound is equal to the maximum possible value from the single path scenario, as shown in (B.1), minus a small quantity. As δ'_s approaches zero or η approaches infinity, the bound converges to the single path condition.

D.2 Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$

When \mathbf{r} and \mathbf{r}^* are sufficiently different to satisfy RNP, the numerator of the ambiguity function in (D.1) consists of only mismatched matrix pairs. We apply the RNP inequality from (A.5) to each term in the numerator to get

$$b(\mathbf{r}) \leq \frac{\|\mathbf{v}\|_2^4}{\|\overline{\Phi}(\mathbf{r})\mathbf{v}\|_2^2} \left(4(\delta'_s)^2(1 + \eta^{-1})^2 \right) , \quad (\text{D.6})$$

where η is the signal-to-interference ratio as defined in (D.3). By then applying RIP in (B.3) to the denominator, the bound then changes to

$$b(\mathbf{r}) \leq (1 + \eta)^2 \left(\frac{4(\delta'_s)^2}{\eta^2(1 - \delta_s)} \right) \|\mathbf{v}\|_2^2 , \quad (\text{D.7})$$

This expression is equal to the single path scenario in (B.4) multiplied by $(1 + \eta)^2/\eta^2$. As η approaches infinity, (D.7) approaches the single path result in (B.4).

D.3 Target-to-artifact ratio

When we take the ratio of (D.5) to (D.7), we derive the lower bound for the peak-to-artifact ratio

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{\|\overline{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^2}{\|\mathbf{v}\|_2^2} \left(\frac{\eta^2}{1+\eta^2} \right) \left(\frac{1-\delta_s}{4(\delta'_s)^2} \right) - \frac{1}{1+\eta^2}, \quad (\text{D.8})$$

If we then apply RIP in (B.3) to this expression, the bound changes to

$$\begin{aligned} \frac{b(\mathbf{r}^*)}{b(\mathbf{r})} &\geq \left(\frac{\eta^2}{1+\eta^2} \right) \left(\frac{(1-\delta_s)^2}{4(\delta'_s)^2} \right) - \frac{1}{1+\eta^2} \\ &= \frac{\eta^2}{1+\eta^2} \left(\frac{(1-\delta_s)^2}{4(\delta'_s)^2} - \frac{1}{\eta^2} \right), \end{aligned} \quad (\text{D.9})$$

Then by setting $\delta'_s = \delta_s$, we derive the expression found in (3.25),

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{\eta^2}{1+\eta^2} \left(\frac{(1-\delta_s)^2}{4\delta_s^2} - \frac{1}{\eta^2} \right). \quad (\text{D.10})$$

Proof of (3.25): multipath scenario (incoherent processor)

In this appendix, we derive the target-to-artifact ratio of the incoherent data-driven matched field processor with unmodeled multipath interference in measured data. Again, we assume $\hat{\mathbf{v}}_q \cong \mathbf{v}_q$. We substitute expression for the test data \mathbf{x}_q for unmodeled multipath interference in (3.24) into the incoherent data-driven ambiguity function in (3.19). This results in an incoherent ambiguity function expressed by

$$b(\mathbf{r}) = \sum_{q=1}^Q \frac{\left| \mathbf{v}_q^H \left(\eta_0 \Phi(\mathbf{r}^*) + \sum_{\ell=1}^L \eta_\ell \Phi(\mathbf{r}_\ell) \right)^H \Phi(\mathbf{r}) \mathbf{v}_q \right|^2}{\eta_0^2 \|\Phi(\mathbf{r}) \mathbf{v}_q\|_2^2}. \quad (\text{E.1})$$

E.1 Ambiguity function lower bound when $\mathbf{r} = \mathbf{r}^*$

In the numerator, we apply the RNP in (A.5) to the mismatched matrix pairs, i.e., $\mathbf{v}_q^H \Phi(\mathbf{r}_\ell) \Phi(\mathbf{r}^*) \mathbf{v}_q$ for $1 \leq \ell \leq L$ and assume each pair of matrices to have the same RNP constant. After applying RNP to the incoherent ambiguity function, we get

$$b(\mathbf{r} = \mathbf{r}^*) \geq \sum_{q=1}^Q \frac{\left| \|\Phi(\mathbf{r}^*) \mathbf{v}_q\|_2^2 - 2\delta'_s \eta^{-1} \|\mathbf{v}_q\|_2^2 \right|^2}{\|\Phi(\mathbf{r}^*) \mathbf{v}_q\|_2^2}, \quad (\text{E.2})$$

where η is the signal-to-interference ratio as defined by (D.3).

As in the coherent matched field processing scenario, we now apply a form of the triangle inequality such that scalar values a and b must satisfy $|a - b|^2 \geq ||a|^2 - |b|^2|$, where $a = \|\Phi(\mathbf{r}^*)\mathbf{v}\|_2^2$ and $b = 2\delta'_s\eta^{-1}\|\mathbf{v}_q\|_2^2$. We also assume $\delta'_s\eta^{-1}$ is sufficiently small such that $||a|^2 - |b|^2| = |a|^2 - |b|^2$. Applying this inequality changes the lower bound to

$$\begin{aligned} b(\mathbf{r} = \mathbf{r}^*) &\geq \sum_{q=1}^Q \frac{\|\Phi(\mathbf{r}^*)\mathbf{v}_q\|_2^4 - 4(\delta'_s)^2\eta^{-2}\|\mathbf{v}_q\|_2^4}{\|\Phi(\mathbf{r}^*)\mathbf{v}_q\|_2^2} \\ &= \sum_{q=1}^Q \|\Phi(\mathbf{r}^*)\mathbf{v}_q\|_2^2 - \frac{\|\mathbf{v}_q\|_2^4}{\|\Phi(\mathbf{r}^*)\mathbf{v}_q\|_2^2} \left(\frac{4(\delta'_s)^2}{\eta^2} \right). \end{aligned} \quad (\text{E.3})$$

By applying the RIP inequality in (B.3) to the denominator, the ambiguity function's bound becomes

$$b(\mathbf{r} = \mathbf{r}^*) \geq \sum_{q=1}^Q \|\Phi(\mathbf{r}^*)\mathbf{v}_q\|_2^2 - \left(\frac{4(\delta'_s)^2}{\eta^2(1-\delta)} \right) \|\mathbf{v}_q\|_2^2. \quad (\text{E.4})$$

As with the coherent scenario, this expression shows us that the lower bound is equal to the maximum possible value from the single path scenario, as shown in (C.1), minus a small quantity that represents the multipath error. As δ'_s approaches zero or η approaches infinity, the bound converges to the single path condition.

E.2 Ambiguity function upper bound when $\mathbf{r} \neq \mathbf{r}^*$

When \mathbf{r} and \mathbf{r}^* are sufficiently different to satisfy RNP, the numerator of the ambiguity function in (E.1) consists of only mismatched matrix pairs. We apply the RNP inequality from (A.5) to each term in the numerator to get

$$b(\mathbf{r}) \leq \sum_{q=1}^Q \frac{\|\mathbf{v}_q\|_2^4}{\|\Phi(\mathbf{r})\mathbf{v}_q\|_2^2} \left(4(\delta'_s)^2(1+\eta^{-1})^2 \right), \quad (\text{E.5})$$

where η is the signal-to-interference ratio as defined in (D.3). By then applying RIP in (B.3) to the denominator, the bound then changes to

$$b(\mathbf{r}) \leq (1+\eta)^2 \left(\frac{4(\delta'_s)^2}{\eta^2(1-\delta_s)} \right) \sum_{q=1}^Q \|\mathbf{v}_q\|_2^2, \quad (\text{E.6})$$

This expression is equal to the single path scenario shown in (C.3) multiplied by $(1+\eta)^2/\eta^2$. As η approaches infinity, (E.6) approaches the single path result in (C.3).

E.3 Target-to-artifact ratio

When we take the ratio of (E.4) to (E.6), we derive the lower bound for the peak-to-artifact ratio

$$\begin{aligned} \frac{b(\mathbf{r}^*)}{b(\mathbf{r})} &\geq \frac{\sum_{q=1}^Q \|\Phi(\mathbf{r}^*)\mathbf{v}_q\|_2^2}{\sum_{q=1}^Q \|\mathbf{v}_q\|_2^2} \left(\frac{\eta^2}{1+\eta^2} \right) \left(\frac{1-\delta_s}{4(\delta'_s)^2} \right) - \frac{1}{1+\eta^2} \\ &= \frac{\|\bar{\Phi}(\mathbf{r}^*)\mathbf{v}\|_2^2}{\|\mathbf{v}\|_2^2} \left(\frac{\eta^2}{1+\eta^2} \right) \left(\frac{1-\delta_s}{4(\delta'_s)^2} \right) - \frac{1}{1+\eta^2}, \end{aligned} \quad (\text{E.7})$$

If we then apply RIP in (B.3) to this expression, the bound changes to

$$\begin{aligned} \frac{b(\mathbf{r}^*)}{b(\mathbf{r})} &\geq \left(\frac{\eta^2}{1+\eta^2} \right) \left(\frac{(1-\delta_s)^2}{4(\delta'_s)^2} \right) - \frac{1}{1+\eta^2} \\ &= \frac{\eta^2}{1+\eta^2} \left(\frac{(1-\delta_s)^2}{4(\delta'_s)^2} - \frac{1}{\eta^2} \right), \end{aligned} \quad (\text{E.8})$$

As with the single path, incoherent scenario in Appendix C, we assume that the RIP constant δ_s for Φ is equivalent to the RIP constant for $\bar{\Phi}$ to simplify our asymptotic analysis. By then setting $\delta'_s = \delta_s$, we derive the expression found in (3.25),

$$\frac{b(\mathbf{r}^*)}{b(\mathbf{r})} \geq \frac{\eta^2}{1+\eta^2} \left(\frac{(1-\delta_s)^2}{4\delta_s^2} - \frac{1}{\eta^2} \right). \quad (\text{E.9})$$

Bibliography

- [1] A. B. Baggeroer, “Matched field processing: Source localization in correlated noise as an optimum parameter estimation problem,” *J. Acoust. Soc. Am.*, vol. 83, no. 2, p. 571, 1988.
- [2] A. Baggeroer, W. Kuperman, and P. Mikhalevsky, “An overview of matched field methods in ocean acoustics,” *IEEE J. Ocean. Eng.*, vol. 18, no. 4, pp. 401–424, Oct. 1993.
- [3] R. L. Weaver and O. I. Lobkis, “Temperature dependence of diffuse field phase,” *Ultrasonics*, vol. 38, no. 1-8, pp. 491–4, Mar. 2000.
- [4] H. Sohn, “Effects of environmental and operational variability on structural health monitoring,” *Phil. Trans. R. Soc. A*, vol. 365, no. 1851, pp. 539–60, Feb. 2007.
- [5] J. B. Harley and J. M. F. Moura, “Scale transform signal processing for optimal ultrasonic temperature compensation,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 59, no. 10, pp. 2226 – 2236, Oct. 2012.
- [6] K. F. Graff, *Wave motion in elastic solids*, 1st ed. New York: Dover Publications, 1991.
- [7] J. D. Achenbach, *Wave propagation in elastic solids*. Amsterdam: Elsevier Science Publishers B.V., 1975.
- [8] F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt, *Computational ocean acoustics*. New York, NY: Springer New York, 2011.
- [9] P. C. Magnusson, A. Weisshaar, V. K. Tripathi, and G. C. Alexander, *Transmission Lines and Wave Propagation, Fourth Edition*. Boca Raton: CRC Press, 2000.
- [10] D. Alleyne, “A two-dimensional Fourier transform method for the measurement of propagating multimode signals,” *J. Acoust. Soc. Am.*, vol. 89, no. 3, pp. 1159–1168, Sep. 1991.
- [11] J. L. Rose, *Ultrasonic Waves in Solid Media*, 1st ed. Cambridge: Cambridge University Press, 2004.

- [12] W. Gao, C. Glorieux, and J. Thoen, "Laser ultrasonic study of Lamb waves: determination of the thickness and velocities of a thin plate," *Int. J. Eng. Sci.*, vol. 41, no. 2, pp. 219–228, Jan. 2003.
- [13] W. H. Prosser, M. D. Seale, and B. T. Smith, "Time-frequency analysis of the dispersion of Lamb modes," *J. Acoust. Soc. Am.*, vol. 105, no. 5, pp. 2669–2676, May 1999.
- [14] M. Niethammer, L. Jacobs, J. Qu, and J. Jarzynski, "Time-frequency representation of Lamb waves using the reassigned spectrogram," *J. Acoust. Soc. Am.*, vol. 107, no. 5 Pt 1, pp. L19–24, May 2000.
- [15] F. Li, G. Meng, L. Ye, Y. Lu, and K. Kageyama, "Dispersion analysis of Lamb waves and damage detection for aluminum structures using ridge in the time-scale domain," *Meas. Sci. Technol.*, vol. 20, no. 9, p. 095704, Sep. 2009.
- [16] A. Raghavan and C. E. S. Cesnik, "Guided-wave signal processing using chirplet matching pursuits and mode correlation for structural health monitoring," *Smart Mater. Struct.*, vol. 16, no. 2, pp. 355–366, Apr. 2007.
- [17] W. Yuemin, "Guided waves modes identification in pipes detection by application of the matching pursuit method," in *IEEE International Conference on Electronic Measurement & Instruments*, vol. 4, Chengdu, Aug. 2011, pp. 50–53.
- [18] J. S. Hall and J. E. Michaels, "A model-based approach to dispersion and parameter estimation for ultrasonic guided waves," *J. Acoust. Soc. Am.*, vol. 127, no. 2, pp. 920–930, Feb. 2010.
- [19] ———, "Model-based parameter estimation for characterizing wave propagation in a homogeneous medium," *Inverse Probl.*, vol. 27, no. 3, p. 035002, Mar. 2011.
- [20] J. B. Harley, A. C. Schmidt, and J. M. F. Moura, "Accurate sparse recovery of guided wave characteristics for structural health monitoring," in *Proc. of the IEEE International Ultrasonics Symposium*. Dresden: IEEE, Oct. 2012, pp. 158–161.
- [21] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [22] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, 2008.
- [23] M. A. Davenport, M. F. Duarte, Y. C. Eldar, and G. Kutyniok, "Introduction to compressed sensing," in *Compressed Sensing: Theory and Applications*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge: Cambridge University Press, 2012, ch. 1, pp. 1–68.
- [24] J. E. Michaels, "Detection, localization and characterization of damage in plates with an in situ array of spatially distributed ultrasonic sensors," *Smart Mater. Struct.*, vol. 17, no. 3, p. 035035, Jun. 2008.

- [25] T. Clarke and P. Cawley, "Enhancing the defect localization capability of a guided wave SHM system applied to a complex structure," *Struct. Health Monit.*, vol. 10, no. 3, pp. 247–259, Jun. 2010.
- [26] W. A. Kuperman, W. S. Hodgkiss, H. C. Song, T. Akal, C. Ferla, and D. R. Jackson, "Phase conjugation in the ocean: Experimental demonstration of an acoustic time-reversal mirror," *J. Acoust. Soc. Am.*, vol. 103, no. 1, pp. 25–40, Jan. 1998.
- [27] J. M. F. Moura and Y. Jin, "Time reversal imaging by adaptive interference canceling," *IEEE Trans. Signal Process.*, vol. 56, no. 1, pp. 233–247, Jan. 2008.
- [28] F. Ciampa and M. Meo, "Impact detection in anisotropic materials using a time reversal approach," *Struct. Health Monit.*, vol. 11, no. 1, pp. 43–49, Jan. 2011.
- [29] —, "Acoustic emission source localization and velocity determination of the fundamental mode A0 using wavelet analysis and a Newton-based optimization technique," *Smart Mater. Struct.*, vol. 19, no. 4, p. 045027, Apr. 2010.
- [30] L. De Marchi, A. Marzani, N. Speciale, and E. Viola, "A passive monitoring technique based on dispersion compensation to locate impacts in plate-like structures," *Smart Mater. Struct.*, vol. 20, no. 3, p. 035021, Mar. 2011.
- [31] A. Perelli, L. De Marchi, A. Marzani, and N. Speciale, "Acoustic emission localization in plates with dispersion and reverberations using sparse PZT sensors in passive mode," *Smart Mater. Struct.*, vol. 21, no. 2, p. 025010, Feb. 2012.
- [32] E. Dehghan Niri and S. Salamone, "A probabilistic framework for acoustic emission source localization in plate-like structures," *Smart Mater. Struct.*, vol. 21, no. 3, p. 035009, Mar. 2012.
- [33] E. B. Flynn, M. D. Todd, P. D. Wilcox, B. W. Drinkwater, and a. J. Croxford, "Maximum-likelihood estimation of damage location in guided-wave structural health monitoring," *P. Roy. Soc. A*, vol. 467, no. 2133, pp. 2575–2596, Apr. 2011.
- [34] P. D. Wilcox, "A rapid signal processing technique to remove the effect of dispersion from guided wave signals," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 50, no. 4, pp. 419–27, Apr. 2003.
- [35] K. Xu, D. Ta, P. Moilanen, and W. Wang, "Mode separation of Lamb waves based on dispersion compensation method," *J. Acoust. Soc. Am.*, vol. 131, no. 4, pp. 2714–22, Apr. 2012.
- [36] A. Raghavan and C. E. Cesnik, "Effects of elevated temperature on guided-wave structural health monitoring," *J. Intel. Mat. Syst. Str.*, vol. 19, no. 12, pp. 1383–1398, May 2008.
- [37] A. D. Degtyar, "Wave propagation in stressed composites," *J. Acoust. Soc. Am.*, vol. 104, no. 4, p. 2192, Oct. 1998.

- [38] K. Salama and C. K. Ling, "The effect of stress on the temperature dependence of ultrasonic velocity," *J. Appl. Phys.*, vol. 51, no. 3, pp. 1505–1509, Mar. 1980.
- [39] G. Konstantinidis, P. D. Wilcox, and B. W. Drinkwater, "An investigation into the temperature stability of a guided wave structural health monitoring system using permanently attached sensors," *IEEE Sensors J.*, vol. 7, no. 5, pp. 905–912, May 2007.
- [40] T. Clarke, F. Simonetti, and P. Cawley, "Guided wave health monitoring of complex structures by sparse array systems: Influence of temperature changes on performance," *J. Sound Vib.*, vol. 329, no. 12, pp. 2306–2322, Jun. 2010.
- [41] Y. Lu and J. E. Michaels, "A methodology for structural health monitoring with diffuse ultrasonic waves in the presence of temperature variations," *Ultrasonics*, vol. 43, no. 9, pp. 717–31, Oct. 2005.
- [42] J. E. Michaels and T. E. Michaels, "Detection of structural damage from the local temporal coherence of diffuse ultrasonic signals," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 52, no. 10, pp. 1769–1782, Oct. 2005.
- [43] A. J. Croxford, P. D. Wilcox, Y. Lu, J. Michaels, and B. W. Drinkwater, "Quantification of environmental compensation strategies for guided wave structural health monitoring," in *Proc. SPIE*, 2008, pp. 69 350H.1–69 350H.11.
- [44] Y. Lu and J. E. Michaels, "Feature extraction and sensor fusion for ultrasonic structural health monitoring under changing environmental conditions," *IEEE Sensors J.*, vol. 9, no. 11, pp. 1462–1471, Sep. 2009.
- [45] F. W. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, *NIST Handbook of Mathematical Functions*. Cambridge: Cambridge University Press, 2010.
- [46] P. Cawley, "Practical long range guided wave inspection – managing complexity," *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 22, no. 657, pp. 22–40, 2003.
- [47] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, Aug. 2006.
- [48] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *C. R. Math*, vol. 346, no. 9-10, pp. 589–592, May 2008.
- [49] J. B. Harley and J. M. F. Moura, "Sparse recovery of the multimodal and dispersive characteristics of Lamb waves," *J. Acoust. Soc. Am.*, vol. 133, no. 5, pp. 2732–2745, May 2013.
- [50] ———, "Data-driven matched field processing for Lamb wave structural health monitoring," *J. Acoust. Soc. Am.*, vol. 135, no. 3, p. 1231, Mar. 2014.
- [51] A. Raghavan, "Guided-wave structural health monitoring," Ph.D. dissertation, The University of Michigan, 2007.

- [52] X. P. Qing, H.-L. Chan, S. J. Beard, T. K. Ooi, and S. A. Marotta, "Effect of adhesive on the performance of piezoelectric elements used to monitor structural health," *Int. J. Adhes. Adhes.*, vol. 26, no. 8, pp. 622–628, Dec. 2006.
- [53] D. Donoho, M. Elad, and V. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 6–18, Jan. 2006.
- [54] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 43, no. 1, pp. 129–159, 1998.
- [55] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge: Cambridge University Press, Jun. 2009.
- [56] D. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, Nov. 2001.
- [57] J. F. Claerbout, "Robust modeling with erratic data," *Geophysics*, vol. 38, no. 5, pp. 826–844, Oct. 1973.
- [58] P. R. Gill, A. Wang, and A. Molnar, "The in-crowd algorithm for fast basis pursuit denoising," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 4595–4605, Oct. 2011.
- [59] D. Malioutov, M. Cetin, and A. Willsky, "Homotopy continuation for sparse signal representation," in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 5, no. 1. Philadelphia, PA: IEEE, Mar. 2005, pp. 733–736.
- [60] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Sci. Comput.*, vol. 31, no. 2, pp. 890–912, Jan. 2009.
- [61] J. Blanchard, C. Cartis, and J. Tanner, "Decay properties of restricted isometry constants," *IEEE Signal Process. Lett.*, vol. 16, no. 7, pp. 572–575, Jul. 2009.
- [62] S. Kunis and H. Rauhut, "Random sampling of sparse trigonometric polynomials, II. orthogonal matching pursuit versus basis pursuit," *Found. Comput. Math.*, vol. 8, no. 6, pp. 737–763, Aug. 2008.
- [63] H. Zou, "The adaptive lasso and its oracle properties," *J. Amer. Statist. Assoc.*, vol. 101, no. 476, pp. 1418–1429, Dec. 2006.
- [64] N. Meinshausen and P. Bühlmann, "High-dimensional graphs and variable selection with the lasso," *Ann. Stat.*, vol. 34, no. 3, pp. 1436–1462, Jun. 2006.
- [65] I. Zorych and Z.-H. Michalopoulou, "Particle filtering for dispersion curve tracking in ocean acoustics," *J. Acoust. Soc. Am.*, vol. 124, no. 2, pp. EL45–50, Aug. 2008.
- [66] M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110.

- [67] ———, “CVX: Matlab software for disciplined convex programming, version 1.21,” Apr. 2011.
- [68] J. E. Michaels, S. J. Lee, J. S. Hall, and T. E. Michaels, “Multi-mode and multi-frequency guided wave imaging via chirp excitations,” in *Proc. of SPIE Conference on Health Monitoring of Structural and Biological Systems*, vol. 7984, San Diego, CA, Mar. 2011, pp. 79 840I–79 840I–11.
- [69] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*. Upper Saddle River, NJ: Prentice Hall, 2001.
- [70] J. L. Krolik, “Matched-field minimum variance beamforming in a random ocean channel,” *J. Acoust. Soc. Am.*, vol. 92, no. 3, pp. 1408–1419, 1992.
- [71] Z.-H. Michalopoulou, “Robust multi-tonal matched-field inversion: A coherent approach,” *J. Acoust. Soc. Am.*, vol. 104, no. 1, pp. 163–170, Jul. 1998.
- [72] C. Debever and W. A. Kuperman, “Robust matched-field processing using a coherent broadband white noise constraint processor,” *J. Acoust. Soc. Am.*, vol. 122, no. 4, pp. 1979–1986, Oct. 2007.
- [73] W. Mantzel, J. Romberg, and K. Sabra, “Compressive matched-field processing,” *J. Acoust. Soc. Am.*, vol. 132, no. 1, pp. 90–102, Jul. 2012.
- [74] S. E. Dosso and M. J. Wilmut, “Maximum-likelihood and other processors for incoherent and coherent matched-field localization,” *J. Acoust. Soc. Am.*, vol. 132, no. 4, pp. 2273–2285, Oct. 2012.
- [75] D. B. Harris and T. Kvaerna, “Superresolution with seismic arrays using empirical matched field processing,” *Geophys. J. Int.*, vol. 182, no. 3, pp. 1455–1477, Sep. 2010.
- [76] M. Papazoglou and J. Krolik, “Matched-field estimation of aircraft altitude from multiple over-the-horizon radar revisits,” *IEEE Trans. Signal Process.*, vol. 47, no. 4, pp. 966–976, Apr. 1999.
- [77] P. Gerstoft, D. Gingras, L. Rogers, and W. Hodgkiss, “Estimation of radio refractivity structure using matched-field array processing,” *IEEE Trans. Antennas Propag.*, vol. 48, no. 3, pp. 345–356, Mar. 2000.
- [78] G. Turek and W. A. Kuperman, “Applications of matched-field processing to structural vibration problems,” *J. Acoust. Soc. Am.*, vol. 101, no. 3, p. 1430, Mar. 1997.
- [79] R. K. Ing and M. Fink, “Ultrasonic imaging using spatio-temporal matched field (STMF) processing—applications to liquid and solid waveguides,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 48, no. 2, pp. 374–386, Mar. 2001.
- [80] A. Tolstoy, “Applications of matched-field processing to inverse problems in underwater acoustics,” *Inverse Problems*, vol. 16, no. 6, pp. 1655–1666, Dec. 2000.

- [81] P. Gerstoft, "Inversion of seismoacoustic data using genetic algorithms and a posteriori probability distributions," *J. Acoust. Soc. Am.*, vol. 95, no. 2, pp. 770–782, 1994.
- [82] S. E. Dosso, P. L. Nielsen, and M. J. Wilmut, "Data error covariance in matched-field geoacoustic inversion," *J. Acoust. Soc. Am.*, vol. 119, no. 1, p. 208, 2006.
- [83] N. M. Shapiro and M. H. Ritzwoller, "Monte-Carlo inversion for a global shear-velocity model of the crust and upper mantle," *Geophys. J. Int.*, vol. 151, no. 1, pp. 88–105, Oct. 2002.
- [84] M. Lowe, "Matrix techniques for modeling ultrasonic waves in multilayered media," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 42, no. 4, pp. 525–542, Jul. 1995.
- [85] J. B. Harley and J. M. F. Moura, "Broadband localization in a dispersive medium through sparse wavenumber analysis," in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Vancouver, BC, May 2013, pp. 4071–4075.
- [86] J. B. Harley, C. Liu, I. J. Oppenheim, and J. M. Moura, "High Resolution Localization with Lamb Wave Sparse Wavenumber Analysis," in *Proc. of the International Workshop on Structural Health Monitoring*, F.-K. Chang, Ed., no. 732, Stanford, CA, Sep. 2013.
- [87] Z.-H. Michalopoulou and M. Porter, "Matched-field processing for broad-band source localization," *IEEE J. Ocean. Eng.*, vol. 21, no. 4, pp. 384–392, 1996.
- [88] S. Aeron, S. Bose, H.-P. Valero, and V. Saligrama, "Broadband dispersion extraction using simultaneous sparse penalization," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 4821–4837, Oct. 2011.
- [89] G. Chardon, A. Leblanc, and L. Daudet, "Plate impulse response spatial interpolation with sub-Nyquist sampling," *J. Sound Vib.*, vol. 330, no. 23, pp. 5678–5689, Nov. 2011.
- [90] J. Hall and J. E. Michaels, "Minimum variance ultrasonic imaging applied to an in situ sparse guided wave array," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 57, no. 10, pp. 2311–2323, Oct. 2010.
- [91] R. M. Levine and J. E. Michaels, "Model-based imaging of damage with Lamb waves via sparse reconstruction," *J. Acoust. Soc. Am.*, vol. 133, no. 3, pp. 1525–34, Mar. 2013.
- [92] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A Simple Proof of the Restricted Isometry Property for Random Matrices," *Constr. Approx.*, vol. 28, no. 3, pp. 253–263, Jan. 2008.
- [93] H. Rauhut, "Stability results for random sampling of sparse trigonometric polynomials," *IEEE Trans. Inf. Theory*, vol. 54, no. 12, pp. 5661–5670, Dec. 2008.
- [94] A. M. Tillmann and M. E. Pfetsch, "The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing," *IEEE Trans. Inf. Theory*, pp. 1248–1259, Feb. 2013.

- [95] J. D. Blanchard, C. Cartis, and J. Tanner, “Compressed sensing: How sharp is the restricted isometry property?” *SIAM Review*, vol. 53, no. 1, pp. 105–125, Jan. 2011.
- [96] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices,” in *Compressed Sensing, Theory and Applications*, Y. Eldar and G. Kutyniok, Eds. Cambridge: Cambridge University Press, 2012, ch. 5, pp. 210–268.
- [97] D. Gabor, “Theory of communication. Part 1: The analysis of information,” *Journal of the IEE*, vol. 93, no. 26, pp. 429 – 441, 1946.
- [98] H. L. Van Trees, *Radar-Sonar Signal Processing and Gaussian Signals in Noise*, ser. Detection, Estimation, and Modulation Theory. New York: John Wiley and Sons, 2001.
- [99] H. C. Song, J. de Rosny, and W. a. Kuperman, “Improvement in matched field processing using the CLEAN algorithm,” *J. Acoust. Soc. Am.*, vol. 113, no. 3, pp. 1379–1386, Mar. 2003.
- [100] D. Eitzen and H. Wadley, “Acoustic Emission: Establishing the Fundamentals,” *J. Res. Nat. Bur. Stand.*, vol. 89, no. 1, pp. 75–100, Jan. 1984.
- [101] C. B. Scruby, “An introduction to acoustic emission,” *J. Phys. E*, vol. 20, no. 8, pp. 946–953, Aug. 1987.
- [102] G. C. McLaskey, S. D. Glaser, and C. U. Grosse, “Beamforming array techniques for acoustic emission monitoring of large concrete structures,” *J. Sound Vib.*, vol. 329, no. 12, pp. 2384–2394, Jun. 2010.
- [103] K. Holford, A. Davies, R. Pullin, and D. Carter, “Damage Location in Steel Bridges by Acoustic Emission,” *J. Intel. Mat. Syst. Str.*, vol. 12, no. 8, pp. 567–576, Aug. 2001.
- [104] P. Nivesrangan, J. Steel, and R. Reuben, “Source location of acoustic emission in diesel engines,” *Mech. Syst. Signal. Pr.*, vol. 21, no. 2, pp. 1103–1114, Feb. 2007.
- [105] N. Toyama, T. Okabe, and N. Takeda, “Lamb wave evaluation and localization of transverse cracks in cross-ply laminates,” *J. Mater. Sci.*, vol. 38, no. 8, pp. 1765–1771, Apr. 2003.
- [106] T. Kundu, S. Das, and K. V. Jata, “Detection of the point of impact on a stiffened plate by the acoustic emission technique,” *Smart Mater. Struct.*, vol. 18, no. 3, p. 035006, Mar. 2009.
- [107] S. M. Ziola, “Source Location in Thin Plates Using Crosscorrelation,” Ph.D. dissertation, Naval Postgraduate School, 1991.
- [108] J.-H. Park and Y.-H. Kim, “Impact source localization on an elastic plate in a noisy environment,” *Meas. Sci. Technol.*, vol. 17, no. 10, pp. 2757–2766, Oct. 2006.

- [109] C. Chen and F.-G. Yuan, "Impact source identification in finite isotropic plates using a time-reversal method: theoretical study," *Smart Mater. Struct.*, vol. 19, no. 10, p. 105028, Oct. 2010.
- [110] A. Tobias, "Acoustic-emission source location in two dimensions by an array of three sensors," *Non-Destructive Testing*, vol. 9, no. 1, pp. 9–12, Feb. 1976.
- [111] T. Kosel, I. Grabec, and F. Kosel, "Intelligent location of simultaneously active acoustic emission sources: Part I," *Aircr. Eng. Aerosp. Tec.*, vol. 75, no. 1, pp. 11–17, 2003.
- [112] T. Kundu, S. Das, and K. V. Jata, "Point of impact prediction in isotropic and anisotropic plates from the acoustic emission data," *J. Acoust. Soc. Am.*, vol. 122, no. 4, pp. 2057–66, Oct. 2007.
- [113] P. Wilcox, M. Lowe, and P. Cawley, "Mode and transducer selection for long range Lamb wave inspection," *J. Intel. Mat. Syst. Str.*, vol. 12, no. 8, pp. 553–565, Aug. 2001.
- [114] D. N. Alleyne and P. Cawley, "The interaction of Lamb waves with defects," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 39, no. 3, pp. 381–97, Jan. 1992.
- [115] A. J. Croxford, J. Moll, P. D. Wilcox, and J. E. Michaels, "Efficient temperature compensation strategies for guided wave structural health monitoring," *Ultrasonics*, vol. 50, no. 4-5, pp. 517–528, Apr. 2010.
- [116] R. L. Weaver, "On diffuse waves in solid media," *J. Acoust. Soc. Am.*, vol. 71, no. 6, pp. 1608–1609, 1982.
- [117] R. L. Weaver and O. I. Lobkis, "Temperature dependence of ultrasonic velocity using diffuse fields; implications for measurement of stress," in *AIP Conference Proceedings*, vol. 557, Ames, Iowa, May 2001, pp. 1480–1486.
- [118] J. B. Harley and J. M. F. Moura, "Guided wave temperature compensation with the scale-invariant correlation coefficient," in *Proc. of the IEEE International Ultrasonics Symposium*, Orlando, FL, Oct. 2011, pp. 1068 – 1071.
- [119] J. B. Harley, Y. Ying, J. M. Moura, I. J. Oppenheim, L. Sobelman, and J. H. Garrett, "Application of Mellin transform features for robust ultrasonic guided wave structural health monitoring," in *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 31, Burlington, VT, Jul. 2011, pp. 1551–1558.
- [120] J. B. Harley and J. M. F. Moura, "An efficient temperature compensation technique for guided wave ultrasonic inspection," in *Proc. of the International Workshop on Structural Health Monitoring*, Stanford, CA, Sep. 2011.
- [121] L. Cohen, "The scale representation," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3275–3292, Dec. 1993.

- [122] P. E. Zwicke and I. Kiss, "A new implementation of the Mellin transform and its application to radar classification of ships," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-5, no. 2, pp. 191–199, Mar. 1983.
- [123] T. Irino, "Segregating information about the size and shape of the vocal tract using a time-domain auditory model: The stabilised wavelet-Mellin transform," *Speech Commun.*, vol. 36, no. 3-4, pp. 181–203, Mar. 2002.
- [124] D. Casasent and D. Psaltis, "Scale invariant optical correlation using Mellin transforms," *Opt. Commun.*, vol. 17, no. 1, pp. 59–63, Apr. 1976.
- [125] R. A. Altes and L. Jolla, "The Fourier-Mellin transform and mammalian hearing," *J. Acoust. Soc. Am.*, vol. 63, no. 1, pp. 174–183, Jan. 1978.
- [126] Y. Sheng and H. H. Arsenault, "Experiments on pattern recognition using invariant Fourier-Mellin descriptors," *J. Opt. Soc. Am.*, vol. 3, no. 6, pp. 771–6, Jun. 1986.
- [127] Q. Chen, M. Defrise, and F. Deconinck, "Symmetric phase-only matched filtering of Fourier-Mellin transforms for image registration and recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 12, pp. 1156–1168, Dec. 1994.
- [128] B. S. Reddy and B. N. Chatterji, "An FFT-based technique for translation, rotation, and scale-invariant image registration," *IEEE Trans. Image Process.*, vol. 5, no. 8, pp. 1266–71, Jan. 1996.
- [129] S. Derrode and F. Ghorbel, "Robust and efficient Fourier-Mellin transform approximations for gray-level image reconstruction and complete invariant description," *Comput. Vis. Image Und.*, vol. 83, no. 1, pp. 57–78, Jul. 2001.
- [130] C. Y. Lin, M. Wu, J. A. Bloom, I. J. Cox, M. L. Miller, and Y. M. Lui, "Rotation, scale, and translation resilient watermarking for images," *IEEE Trans. Image Process.*, vol. 10, no. 5, pp. 767–82, Jan. 2001.
- [131] R. Cassinis, "Unsupervised matching of visual landmarks for robotic homing using Fourier-Mellin transform," *Robot. Auton. Syst.*, vol. 40, no. 2-3, pp. 131–138, Aug. 2002.
- [132] J. Zhang, Z. Ou, and H. Wei, "Fingerprint matching using phase-only correlation and Fourier-Mellin transforms," in *Sixth International Conference on Intelligent Systems Design and Applications*, Jinan, Oct. 2006, pp. 379–383.
- [133] J. Yang, T. Sarkar, and P. Antonik, "Applying the Fourier-modified Mellin transform (FMMT) to Doppler-distorted waveforms," *Digit. Signal Process.*, vol. 17, no. 6, pp. 1030–1039, Nov. 2007.
- [134] A. De Sena and D. Rocchesso, "A Fast Mellin and Scale Transform," *EURASIP J. Adv. Sig. Pr.*, vol. 2007, no. 1, pp. 1–10, Jan. 2007.

- [135] J. Bertrand, P. Bertrand, and J. P. Ovarlez, “The Mellin Transform,” in *Transforms and Applications Handbook*, 3rd ed., L. D. Poularikas, Ed. Boca Raton: CRC Press, 2010, ch. 12, pp. 12–1 – 12–37.
- [136] M. Unser, “Splines: a perfect fit for signal and image processing,” *IEEE Signal Process. Mag.*, vol. 16, no. 6, pp. 22–38, Nov. 1999.
- [137] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes*, 3rd ed. New York: Cambridge University Press, 2007.
- [138] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck, *Discrete-time Signal Processing*, 2nd ed. Upper Saddle River: Prentice Hall, 1999.
- [139] G. Robbins and T. Huang, “Inverse filtering for linear shift-variant imaging systems,” *Proc. IEEE*, vol. 60, no. 7, pp. 862–872, Jul. 1972.
- [140] H. Sundaram, S. Joshi, and R. Bhatt, “Scale periodicity and its sampling theorem,” *IEEE Trans. Signal Process.*, vol. 45, no. 7, pp. 1862–1865, Jul. 1997.
- [141] A. De Sena, “A computational framework for sound analysis with the Mellin and scale transform,” Ph.D. dissertation, Università di Verona, 2008.
- [142] Z. Lu, S. J. Lee, J. E. Michaels, T. E. Michaels, D. O. Thompson, and D. E. Chimenti, “On the optimization of temperature compensation for guided wave structural health monitoring,” in *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 1211, no. May 2012, Kingston, 2010, pp. 1860–1867.
- [143] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization*. New York: Academic Press Inc., 1981.
- [144] P. H. F. Nicholson, P. Moilanen, T. Kärkkäinen, J. Timonen, and S. Cheng, “Guided ultrasonic waves in long bones: modelling, experiment and in vivo application,” *Physiological Measurement*, vol. 23, no. 4, pp. 755–768, Nov. 2002.
- [145] Y. Ying, J. H. Garrett, J. Harley, I. J. Oppenheim, J. Shi, and L. Soibelman, “Damage Detection in Pipes under Changing Environmental Conditions Using Embedded Piezoelectric Transducers and Pattern Recognition Techniques,” *J. Pipeline Syst. Eng. Pract.*, vol. 4, no. 1, pp. 17–23, Feb. 2013.
- [146] Y. Ying, J. H. Garrett, I. J. Oppenheim, L. Soibelman, J. B. Harley, J. Shi, and Y. Jin, “Toward Data-Driven Structural Health Monitoring: Application of Machine Learning and Signal Processing to Damage Detection,” *J. Comput. Civil Eng.*, vol. 27, no. 6, pp. 667–680, Nov. 2013.
- [147] C. Liu, J. B. Harley, N. O’Donoghue, Y. Ying, M. Berges, M. H. Altschul, J. H. Garrett, Jr, D. Greve, J. M. F. Moura, I. J. Oppenheim, and L. Soibelman, “Ultrasonic scatterer detection in a pipe under operating conditions using singular value decomposition,” in *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 1454, Denver, CO, Jul. 2013, pp. 1454–1461.

- [148] C. Liu, J. B. Harley, M. Bergés, D. W. Greve, W. R. Junker, and I. J. Oppenheim, “A robust baseline removal method for guided wave damage localization,” in *Proc. of SPIE Conference on Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, J. P. Lynch, K.-W. Wang, and H. Sohn, Eds., San Diego, CA, Apr. 2014, p. 90611K.
- [149] C. Stubbs, M. Brenner, A. Despain, R. Henderson, D. Long, W. Press, J. Tonry, and P. Weinberger, “The Computational Challenges of Medical Imaging,” Mitre Corp. and McLean VA JASON Program Office, Tech. Rep., Apr. 2004.