2022-04-14

Physics-Guided Machine Learning in Ocean Acoustics Using Fisher Information

Michael Craig Mortenson
Brigham Young University

Physics-Guided Machine Learning in Ocean Acoustics Using Fisher Information

Michael Craig Mortenson

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Tracianne B. Neilsen, Chair
Mark K. Transtrum
Jared P. Whitehead
David P. Knobles

Department of Physics and Astronomy

Brigham Young University

ABSTRACT

Physics-Guided Machine Learning in Ocean Acoustics Using Fisher Information

Michael Craig Mortenson
Department of Physics and Astronomy, BYU
Master of Science

Waterborne acoustic signals carry information about the ocean environment. Ocean geoacoustic inversion is the task of estimating environmental parameters from received acoustic signals by matching the measured sound with the predictions of a physics-based model. A lower bound on the uncertainty associated with environmental parameter estimates, the Cramér-Rao bound, can be calculated from the Fisher information, which is dependent on derivatives of a physics-based model. Physics-based preconditioners circumvent the need for variable step sizes when computing numerical derivatives. This work explores the feasibility of using a neural network to perform geoacoustic inversion for environmental parameters and their associated uncertainties from ship noise spectrogram data. To train neural networks, a synthetic dataset is generated and tested for generalizability against 31 measurements taken during the SBCEX2017 study of the New England Mud Patch.

Keywords: underwater acoustics, geoacoustic inversion, Fisher information, Cramér-Rao bound, machine learning, deep learning, uncertainty analysis

ACKNOWLEDGMENTS

Many people deserve thanks in relation to this thesis, but a few deserve special mention. First, I would like to thank Dr. Neilsen. Thank you for accepting me on as a graduate student, and for encouraging and mentoring me throughout two years during which the world and our lives were turned upside-down by a world-wide pandemic. Next, I would like to thank Dr. Transtrum. Thank you for your patience and clarity in breaking down complex mathematical and modeling concepts, and also for your contagious (is that word acceptable on the heels of a pandemic?) excitement for computational physics. I'm grateful to the members of my committee: Dr. Whitehead, for his willingness to jump in and lend his machine learning and mathematical modeling knowledge to a new area of application, and Dr. Knobles, for many Zoom calls to discuss progress and ideas for the project. Finally, I'm grateful to my wife, Courtney, for her love and support throughout my graduate studies, and to Lincoln, who was born during the course of this work.

This research was supported by the Office of Naval Research contract number N00014-19-C-2001 and by the BYU Department of Physics and Astronomy.

Contents

Table of Contents **iv**

List of Figures **v**

List of Tables **ix**

1 Introduction **1**
1.1 Background 1

2 Fisher Information Analysis of Ocean Acoustic Models **6**
2.1 Ocean Acoustic Models and Fisher Information 6
2.2 Fisher information 9
2.3 Computing Accurate Numerical Gradients 15
2.4 Case Study 21
2.5 Conclusion 33

3 Synthetic Dataset Development **35**
3.1 Uncertainty in Parameter Estimates 36
3.2 Real-World Dataset 37
3.3 Synthetic Dataset 38
3.4 Uncertainty Labels 44

4 Deep Learning **48**
4.1 Background 48
4.2 Architecture and Training 50
4.3 Validation 51
4.4 Generalization 55

5 Conclusion **62**

Bibliography **66**

List of Figures

1.1 Workflow Overview. The lefthand side of the vertical dotted line depicts the process of generating ship noise spectrograms from the forward model ORCA. Chapter 2 shows how parameter uncertainties can be obtained from ORCA. The righthand side of the vertical dotted line depicts the process of training a neural network, ResNet18, on a dataset of synthetic spectrograms. The generation and labeling of the dataset with parameter values and uncertainties is discussed in Chapter 3. The training ResNet18 to perform regression for parameter values and uncertainties is discussed in Chapter 4. 4

2.1 Two-dimensional cost surface cross-section in the θ_1-θ_2 parameter plane showing the neighborhood about a local minimum. Notice the elliptical contours of equal cost about the minimum at the best-fit parameters $\hat{\boldsymbol{\theta}}$, marked by a white star. The semi-minor and semi-major axes of the contours correspond to the FIM eigenvector directions $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$, respectively. The inverse of the square root of corresponding eigenvalues of the FIM, λ_1 and λ_2, are proportional to the semi-minor and semi-major widths of the cost ellipses. The width of the ellipses projected onto the θ_1- and θ_2-axes are proportional to $\sqrt{I_{ii}^{-1}}$, the square root of the corresponding diagonal elements of the inverse FIM (covariance matrix). I^{-1} is calculated when finding the CRB. Finally, the widths of the ellipses through the best fit point and parallel to the parameter axes are proportional to $1/\sqrt{I_{ii}}$ the inverse of the square root of the corresponding elements of the FIM, also known as the conditional variance. 10

2.2 Variation in magnitude of Green's function over a unit step in the preconditioned parameter for two different preconditioners of $cp1_o$. The top row shows the best preconditioner $\frac{fr}{cp1_o}$ and the bottom row shows a poor preconditioner $\frac{cp1_o}{fh_{o,\mathrm{ref}}}$. The curves represent the modulus of the Green's function |g| for various source-receiver pairs. The legend gives the depth of the receiver followed by the depth of the source beneath the surface of the ocean in meters. Notice how the best preconditioner captures a smoothly varying region of local behavior with a unit step in the preconditioned parameter for all frequencies considered. The poor preconditioner (bottom row) features oscillations on a scale that would make finite-differencing difficult with a unit step size. 19

2.3 Fisher information eigenvalue magnitudes at frequencies 50–400 Hz for sand sediment environment. Notice the log-linear eigenvalue spacing characteristic of sloppy models. 24

2.4 Absolute value of entries of eigenvectors of the FIM with respect to log parameters at 100 Hz for a sand sediment. The horizontal axis shows the eigenvalue number in decreasing order that corresponds to each eigenvector of the FIM. Parameter names are shown on the vertical axis. Each vertical column represents a unit-length eigenvector in parameter space. Matrix elements with values closer to 1 and indicate eigenvectors that point strongly in along the associated axis in parameter space. . . 25

2.5 Absolute value of Pearson correlation matrix for a sand sediment at 100 Hz. The correlation matrix shows parameter relationships. In many cases, strong correlations occur between parameters from the same layer or parameters that mark a change at a layer interface. 26

2.6 Relative uncertainty, σ_{CRB}, associated with the CRB and the relative uncertainty associated with the conditional variance, σ_{cond}, for each log parameter for a sand sediment environment at 100 Hz. A solid black line has been placed at a value of 1 indicating the point at which the uncertainty in the parameter is equal to the magnitude of the parameter's reference value. In general, the deeper the parameter is beneath the surface of the ocean, the larger the relative uncertainty. Also, relative uncertainty tends to increase with frequency. All relative uncertainty values are based on a transmission loss threshold of 1 dB re 1 μPa in Eq. 2.9. 27

2.7 Relative uncertainty values associated with the Cramér-Rao bound (σ_{CRB}) and the conditional variance (σ_{cond}) for each of the model parameters at 50–400 Hz for a mud seabed. Relative uncertainty values based on an assumed resolution of 1 dB re 20 μPa for transmission loss. Bars for σ_{CRB} shown in blue with circles and σ_{cond} shown in orange with diagonal cross-hatching. 31

2.8 Relative uncertainty values associated with the Cramér-Rao bound (σ_{CRB}) and the conditional variance (σ_{cond}) for each of the model parameters at 50–400 Hz for a sand seabed. Relative uncertainty values based on an assumed resolution of 1 dB re 20 μPa for transmission loss. Bars for σ_{CRB} shown in blue with circles and σ_{cond} shown in orange with diagonal cross-hatching. 32

3.1 SBCEX2017 site in the New England Mud Patch. The deployment locations of VLA1 and VLA2 are marked for reference as well as the shipping lanes. The real-world SOO spectrogram data used for generalization was obtained using the setup. 39

3.2 SOO Spectrogram from the Viking Bravery recorded on VLA 1 at approximately 33 m above the seafloor. The "bathtub" striation pattern typical of SOO spectrograms is due to the shifting of interference patterns as the ship approaches and then leaves CPA.[1] 40

4.1 Validation results for the simultaneous regression of parameter values and parameter uncertainties. Subplots are broken out by parameter and predicted uncertainty bin. Points are colored according to labeled uncertainty bin. The axes of each plot are parameter prediction on the vertical axis and true value on the horizontal axis. For an exact regression, the points should align along the diagonal of each subplot. . . 54

4.2 Generalization results for neural networks trained to estimate single parameter values. Each of the 31 SOO spectrograms from SBCEX2017 are represented. The VLA from which the measurements were obtained are listed beside the ship names. Reference values from literature are indicated with dotted lines. For each parameter, the trained ResNet-18 networks tend to overestimate compared to published values. The network does not appear to prefer certain ship spectrograms the simultaneous regression of parameter values and parameter uncertainties. 58

4.3 Generalization results for neural networks trained to simultaneously estimate parameter values and parameter uncertainty labels. Each of the 31 SOO spectrograms from SBCEX2017 are represented. The VLA from which the measurements were obtained are listed beside the ship names. Reference values from literature are indicated with dotted lines. The predicted uncertainty bin is indicated by the color of each point. For $cp1$ and $ap1$ the network tends to overestimate parameter values compared with published estimates. The thickness of the top sediment layer h predicts a wide spread of values centered on the published value of 9.2 m. The best performing parameter, however, is $\rho 1$, for which over half of the network's predicted parameter values fall within the range of values published in literature. . 60

List of Tables

2.1 Preconditioners $\phi(\theta)$ for physical parameters θ used to obtain accurate derivatives, constructed according to physics-based preconditioner design principles. Parameters are defined as follows: h, layer thickness, cp, compressional sound speed, ρ, density, and αp compressional attenuation coefficient. The numerals 1 and 2 denote the top and bottom of a layer, respectively. The subscript abbreviations o, sed, bas, and ref stand for ocean, sediment, basalt basement, and reference value, respectively. . . . 17

2.2 Reference parameter values for case study ocean model. The simplified ocean models consist of 3 layers: 75-m water column over a 35-m sediment layer of either mud, clay, silt, sand, or gravel over 100-m basement layer of basalt. Compressional sound speeds cp are assumed to be constant in the water column and basement. In the sediment layer, the compressional sound speed is assumed to increase linearly at 1 m/s per meter moving deeper into the sediment. All parameters except h are defined in the model by the value at the top and bottom of the layer. Values of cp, ρ, and αp from *Computational Ocean Acoustics* Table 1.3 unless otherwise noted.[2] 22

3.1 Bounds on geoacoustic parameters used to generate environments for the synthetic dataset. Sampled values were constrained to remain constant or increase with depth. 43

3.2 Relative uncertainty bins for the labeling of SOO spectrograms. For improved weight convergence in the neural network and for ease in interpretation, each bin is associated with a unit domain. 46

4.1 Validation error analysis results from single and dual regression for four geoacoustic parameters: layer thickness h, sound speed at the top of the sediment $cp1$, density at the top of the layer $\rho 1$, and attenuation at the top of the layer $\alpha 1$. The dual regression includes the results for simultaneous regression for the uncertainty labels, σ with the associated parameter subscripted. 53

4.2 Generalization results performing regression with ResNet-18 on the 31 measured SOO spectrograms from SBCEX2017. The mean and standard deviation are presented for both the single parameter regression case and the dual regression parameter and parameter uncertainty case. Representative nominal parameter values from the SBCEX2017 literature are included for context.[3,4] 56

Chapter 1

Introduction

1.1 Background

As sound travels through and interacts with an environment, it becomes encoded with information about that environment. This environmental influence is the reason the human ear can differentiate the sound of someone singing in the shower and the sound of someone singing in Carnegie Hall. Via networks of neurons in the brain, the acoustic signal is parsed and analyzed to determine that, yes, the acoustic signal captured by the ear sounds like it bounced off a small space walled with hard, porcelain tiles. This process, which the human brain carries out with seemingly little effort, turns out to be difficult to replicate in practice with computational tools. The challenge of developing such computational tools is an area of active research.

The Office of Naval Research, whose funding supports this work, is specifically interested in determining details about the shallow ocean environment from waterborne sound. Here, the term shallow ocean refers to depths of less than 200 m, typically found along continental shelves. There are many methods of learning about the shallow ocean seabed environment including coring experiments and sonar. Sonar comes in two varieties: active sonar, in which a ping signal is

generated and information about the environment is recovered from the received echo, and passive sonar, in which ambient noise from the wind, waves, ships, biological activity, and geologic activity is used to learn information about the environment. This work is primarily concerned with passive sonar using ship noise as the source.

The general approach for estimating characteristics of an ocean environment from sound begins with a physics-based model. The goal is for the model to capture enough of the relevant physics of underwater sound propagation that when the proper environmental parameters are fed into the model, the model accurately predicts the same acoustic field as if a direct measurement had been taken. Formulated mathematically, the process of learning about an environment from signals is called *inversion* or *inference* and the problem of estimating parameters is called the *inverse* problem.[5]

Some research on inverse problems focuses on the physical model, called the *forward model*, to understand its strengths, weaknesses, structure, and behavior with respect to a given inference task. The forward model utilized in this work is ORCA, a range-independent normal mode model that predicts transmission loss, TL.[6,7]

ORCA solves the wave equation in cylindrical coordinates to obtain TL for source-receiver pairs in an ocean environment. The ocean environment is assumed to be horizontally stratified and azimuthally symmetric. Thus, ORCA is a forward model that uses frequency, geometric parameters describing source-receiver arrangements, and geoacoustic parameters describing the water column and seabed. From these parameters, ORCA calculates the Green's function over the ocean waveguide, using Eq. 23 in Westwood *et al.*[6] Conversion of the Green's function to decibels yields transmission loss, TL. The relationship between TL, the sound source level SL, and the received level RL, at a given frequency f is given by the sonar equation:

$$\mathrm{TL} = \mathrm{SL} - \mathrm{RL}. \tag{1.1}$$

For a moving source, the distance changes with time and repeated calculations of TL across frequencies can be stacked together to form a spectrogram image. Such spectrograms encode information about the simulated ocean environment, just as the spectrograms of measured acoustic signals encode information about the physical ocean environment.

The information contained in transmission loss about different input parameters can be quantified by studying ORCA as a forward model. With an understanding of this information content, it becomes possible to predict the limitations of using ORCA as a forward model for geoacoustic parameter inversions. Chapter 2 deals primarily with how numerical methods can be used to find the information content of a forward model's output with respect to input parameters. As a case study, the information content of TL obtained from ORCA with respect to geoacoustic parameters is used to estimate parameter uncertainties.

Besides studying the forward model, inversion research also focuses on the tools and procedures used to perform inference. While many tools are available to perform inversion, this work focuses on artificial neural networks, a tool that takes its inspiration from the brain. Artificial neural networks, hereafter referred to simply as neural networks, consist of interconnected computational neurons, which are parameterized nonlinear functions. A neural network can be "trained" by feeding in labeled data and adjusting network parameters (weights) so the network predictions match expected labels. When the possible label values are a discrete set, the task is called classification, and when the label values are continuous the task is called regression. Ideally, in either classification or regression, the network learns a generalizable function for consistently mapping inputs to labels such that the network is capable of predicting labels on inputs it has never seen before.[8]

In prior work by our group, both regression and classification with ship noise spectrograms have been performed. Because of the lack of labeled real-world measurements, the neural networks have been trained on synthetic spectrograms. In Van Komen *et al.*,[9] a five-layer CNN was trained to learn via regression the ship speed and distance and a value representing the type of seabed. In Escobar *et*

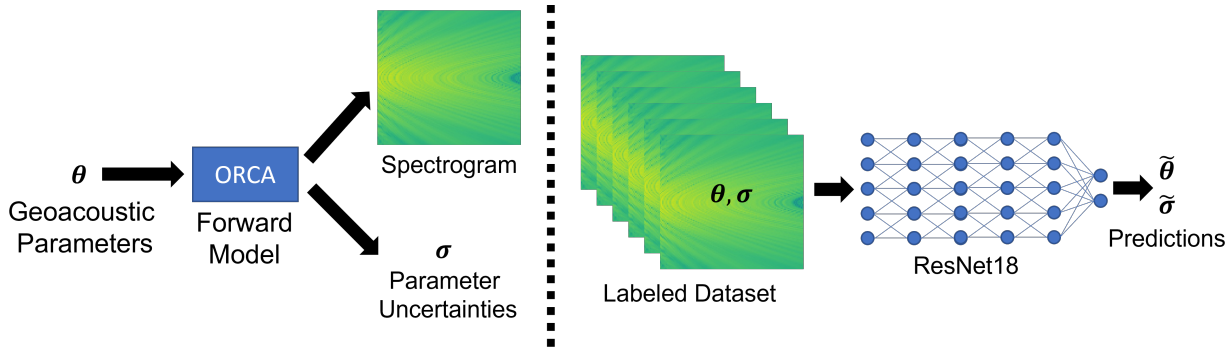

Figure 1.1 Workflow Overview. The lefthand side of the vertical dotted line depicts the process of generating ship noise spectrograms from the forward model ORCA. Chapter 2 shows how parameter uncertainties can be obtained from ORCA. The righthand side of the vertical dotted line depicts the process of training a neural network, ResNet18, on a dataset of synthetic spectrograms. The generation and labeling of the dataset with parameter values and uncertainties is discussed in Chapter 3. The training ResNet18 to perform regression for parameter values and uncertainties is discussed in Chapter 4.

al.,[10] six different network architectures were trained for classification among 34 discrete seabed types from a large number of spectrograms that have a wide variety of source-receiver geometries.

The work in this thesis explores using ship noise spectrograms in a neural network to perform regression for both geoacoustic parameters and estimated parameter uncertainties. The goal is to train a neural network that is robust enough to generalize to measured, real-world spectrograms and obtain estimates of the geoacoustic parameters and an estimate of their uncertainties. The details of measured ship spectrograms and the generating and labeling a training dataset of synthetic spectrograms is discussed in Chapter 3, along with a description of how the parameter uncertainties are found using the techniques in Chapter 2. These steps are illustrated in the left side of the workflow overview in Fig. 1.1. The right side of Fig. 1.1 depicts how the labeled data is used to train a neural network, ResNet18, which is described in Chapter 4. An exploration of the abilities and limitations of the trained network in generalizing to real-world data is also provided in Chapter 4.

In summary, this thesis seeks to expand the body of knowledge about ocean acoustic inversion and physics-guided machine learning in the following ways:

1. Develop a method for performing uncertainty analysis of parameter estimates. Uncertainty analysis is made possible through the computation of accurate gradients of the forward model with respect to individual parameters.

2. Explore how uncertainty analysis might be used to develop and label synthetic datasets for training a neural network to perform environmental parameter inversion. Uncertainty analysis is also considered as a way of evaluating potential targets for regression with a neural network.

3. Present an initial case study of training a neural network on synthetic data, both to perform regression for individual environmental parameters and to provide an estimate of the forward model uncertainty associated with parameter estimates.

The following chapters are organized as follows with relevant literature reviewed in each chapter. Chapter 2 describes how to obtain accurate gradients of a forward model with respect to parameters to enable uncertainty analysis. This chapter is an article that was submitted to the Journal of the Acoustical Society of America. Chapter 3 details the development of a synthetic dataset of spectrograms with parameter and parameter uncertainty labels. Chapter 4 covers the training of a neural network to perform regression for parameter values and parameter uncertainties. Finally, Chapter 5 concludes with a review of lessons learned and future work.

Chapter 2

Fisher Information Analysis of Ocean Acoustic Models

2.1 Ocean Acoustic Models and Fisher Information

Physics-based ocean acoustics models play a central role in inferring environmental parameters from underwater acoustic signals. Characterizing the inherent uncertainty in inferring environmental parameters is a difficult task, further complicated by the difficulty of obtaining accurate numerical derivatives of ocean acoustic models with respect to model parameters. In this work, we present a methodology for obtaining accurate numerical derivatives of the predictions of ocean acoustics models using physics-based parameter preconditioning and Richardson extrapolation, which renders possible the study of model sensitivity and parameter uncertainty via Fisher information analysis.

The task of environmental parameter inference belongs to a larger family of inverse problems.[5] From an information-theoretic perspective, an inverse problem encodes information from some signal by fitting parameters of a model.[11] In ocean acoustics inverse problems, the signal is acoustical in nature (e.g., waveforms, spectrograms, transmission loss curves, etc.) and the information to be

learned are details of the ocean environment, such as water depth, sediment density, sound speed profile, etc. Often, ocean acoustic inversion literature focuses on finding best-fit parameters without assessing model sensitivity or parameter uncertainties, which provide needed context about the trustworthiness of best-fit parameters and facilitate comparison between the resulting models and between inversion schemes.

Two powerful tools for characterizing the local behavior of a physical model in the context of inverse problems are the Fisher information matrix (FIM) and the matrix inequality known as the Cramér-Rao bound (CRB). The FIM quantifies how well the model resolves parameter estimates from environmental information in the data. The CRB places a lower bound on the variance of an unbiased estimator for the parameters of a model. If the eigenvalues of the FIM vary across many orders of magnitude with approximately log-linear spacing, the model is said to be *sloppy*, a characteristic common to many models across various fields of scientific inquiry including systems biology,[12–14] solid state physics,[11,15] and biokinematics.[16,17] Sloppy models are so-named because the predictive behavior of the model is sensitive to combinations of a few *stiff* parameters and largely insensitive to other *sloppy* parameters. As such, the FIM quantifies the local sensitivity of model predictions to changes in model parameters. Sloppy model behavior has been observed for ocean acoustic models in estimates of global sensitivity obtained using a Markov Chain Monte-Carlo approach.[18,19] One of the major contributions of the current work is a local sloppy model analysis using the FIM applied to an ocean sound propagation model.

In the literature, the context of either the experimental setup or the inversion approach typically determines how the FIM and CRB are obtained. Where circumstances permit, analytical expressions for the FIM and CRB are employed. Analytical expressions have been used for source localization and velocity estimates with a stationary array,[20] optimal sensor spacing to estimate speed and attenuation parameters from isotropic ambient noise,[21] and normal mode model Green's functions

for matched field tomography and ocean acoustic tomography.[22] See Baggeroer *et al.* for a thorough review of relevant literature from the underwater acoustics and signal processing communities.[23]

In Bayesian and other sampling-based inversion strategies applied to underwater acoustics, calculations of the FIM and CRB rely on the evaluation of (often multi-dimensional) integrals of probability distributions. These integrals are evaluated with methods ranging from naïve exhaustive grid searches[24,25] to sophisticated Markov Chain Monte-Carlo methods.[4,26–28] See the review by Gerstoft *et al.* for further details.[29]

A third approach for obtaining the FIM and CRB relies upon gradients of the physical model taken with respect to parameters. However, gradients can be difficult to obtain. Some success has been found with analytical gradients[30,31] and implicit adjoint methods[32] although both require a problem-specific gradient derivation. Finite-differencing has typically been avoided because of the difficulty in determining appropriate step sizes; too large or too small of step sizes makes the method inaccurate.[29] In this work, we present a strategy including parameter preconditioning and Richardson extrapolation that renders feasible a general finite-difference methodology for finding numerical gradients. That methodology is subsequently applied in a case study using the FIM and CRB to analyze a normal-mode propagation model for transmission loss in shallow ocean environments.

The paper is organized as follows. In Section 2.2, the FIM and CRB are derived in terms of gradients of the model, and the utility of the FIM and CRB in sensitivity and uncertainty analysis is reviewed. Section 2.3 describes the methodology for accurate gradient computation including preconditioning and Richardson extrapolation. Section 2.4 contains the case study validation of the Fisher information analysis for a variety of seabed sediment types and frequencies, and conclusions are presented in Section 2.5.

2.2 Fisher information

Solutions to the inverse problem make use of a physical model and its inputs. For example, a physical model of transmission loss, x, includes inputs for source depths $\boldsymbol{s}$, receiver depths $\boldsymbol{z}$, source-receiver range r, and model parameters $\boldsymbol{\theta}$, such that $x(\boldsymbol{z},\boldsymbol{s},r,\boldsymbol{\theta})$ can predict the transmission loss at each source-receiver pair. A forward model, $m(\boldsymbol{\theta})$, is then formed by concatenating the predictions of $x(\boldsymbol{z},\boldsymbol{s},r,\boldsymbol{\theta})$ for the different values of $\boldsymbol{s}$ and $\boldsymbol{z}$. The forward model can thus be expressed as a mapping from an N-dimensional parameter space into an M-dimensional data space, where M is the total number of data points to be fit:

$$m(\boldsymbol{\theta}) : R^N \rightarrow R^M. \tag{2.1}$$

Take, for example, a forward model for predicting transmission loss at M hydrophone locations. The predictions of the forward model for all possible parameter values generates a set of accessible model predictions. To qualify those predictions, a cost function, sometimes called an error or loss function, is defined that measures how well model prediction values match observed data. One such cost function is the least-squares cost:

$$C_\theta = \frac{1}{2}\sum_i^M (m_i(\boldsymbol{\theta}) - d_i)^2, \tag{2.2}$$

where m_i is ith prediction of the forward model corresponding to data point d_i. The input parameterization that minimizes the cost is the best fit parameterization, $\hat{\boldsymbol{\theta}}$. The graph of the cost defines an N-dimensional surface, called a cost surface, whose peaks correspond to parameter values that poorly fit the data and whose valleys correspond to parameter values that fit the data well. N-dimensional surfaces are unwieldy to visualize, so two-dimensional slices of cost surfaces are often used to gain a "topographical" intuition of the peaks and valleys of the cost surface. One such two-dimensional slice in the neighborhood of a best fit valley is shown in Fig. 2.1 for a toy model, which is discussed throughout the present section.

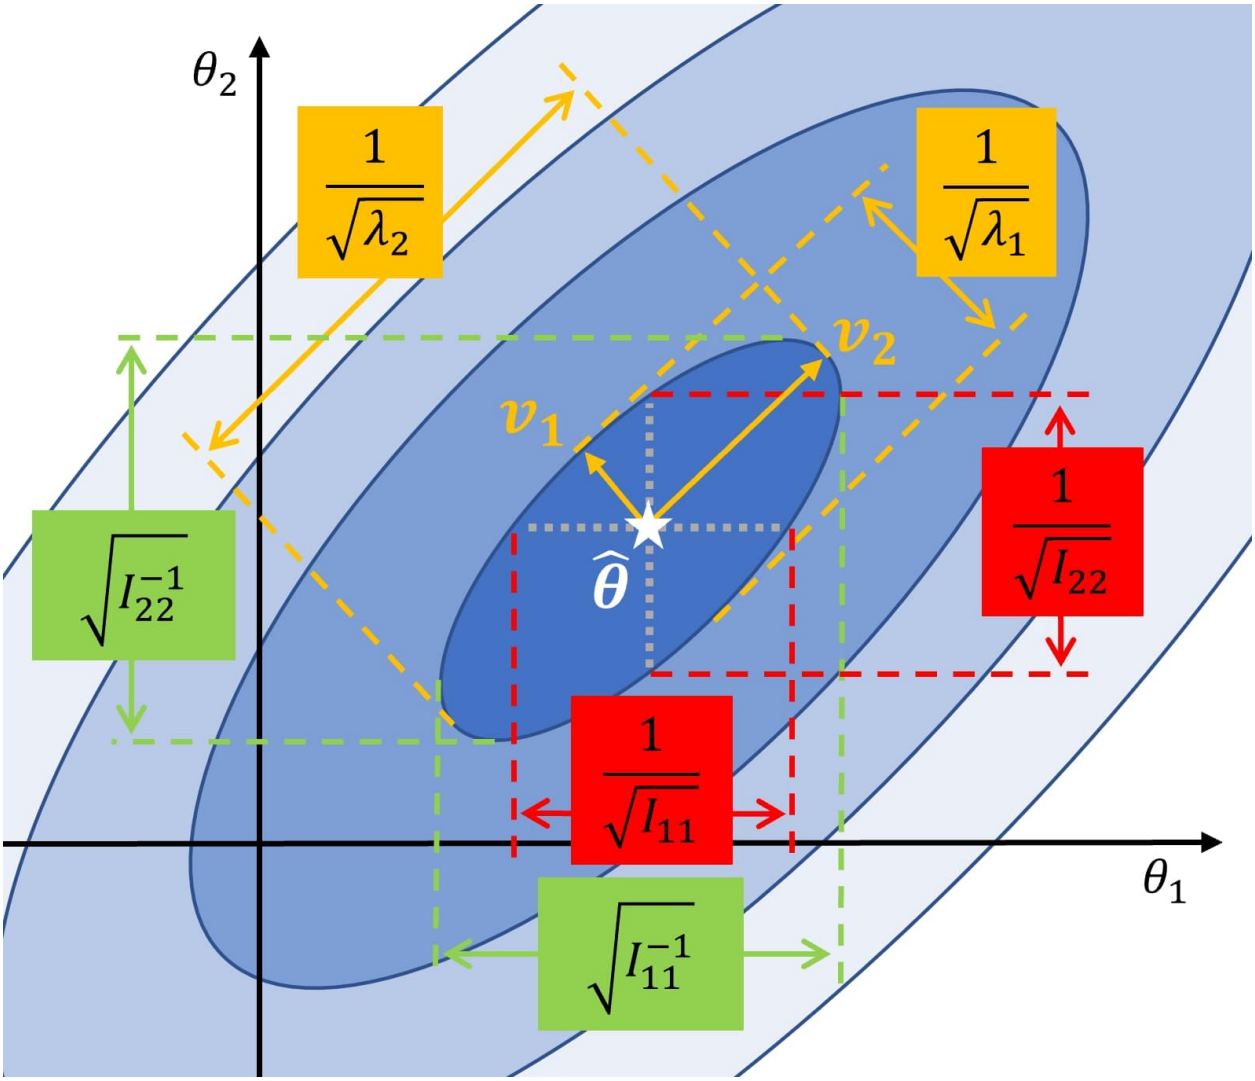

Figure 2.1 Two-dimensional cost surface cross-section in the θ_1-θ_2 parameter plane showing the neighborhood about a local minimum. Notice the elliptical contours of equal cost about the minimum at the best-fit parameters $\hat{\boldsymbol{\theta}}$, marked by a white star. The semi-minor and semi-major axes of the contours correspond to the FIM eigenvector directions $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$, respectively. The inverse of the square root of corresponding eigenvalues of the FIM, λ_1 and λ_2, are proportional to the semi-minor and semi-major widths of the cost ellipses. The width of the ellipses projected onto the θ_1- and θ_2-axes are proportional to $\sqrt{I_{ii}^{-1}}$, the square root of the corresponding diagonal elements of the inverse FIM (covariance matrix). I^{-1} is calculated when finding the CRB. Finally, the widths of the ellipses through the best fit point and parallel to the parameter axes are proportional to $1/\sqrt{I_{ii}}$ the inverse of the square root of the corresponding elements of the FIM, also known as the conditional variance.

Cost surfaces for nonlinear models, such as those found in modeling ocean acoustics, often include narrow valleys, flat plateaus, and multiple local minima. However, by zooming in on a sufficiently small region around a local minimum, the cost surface is approximately quadratic with lines of constant cost forming ellipse-shaped contours about the minimum. The axes of such ellipses provide insight into the sensitivity of model parameters.

The local geometry about a minimum is described by a Taylor series of the forward model expanded about the best fit:

$$m(\hat{\boldsymbol{\theta}}) = m(\hat{\boldsymbol{\theta}}) + J(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) + \mathscr{O}((\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^2), \tag{2.3}$$

where $J = \partial m / \partial \theta$ is the M by N Jacobian matrix of the forward model. In this linear approximation, the cost function becomes

$$C_\theta(\hat{\boldsymbol{\theta}}) = C(\hat{\boldsymbol{\theta}}) + \frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T J^T J(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})) \tag{2.4}$$

using the fact that the gradient is zero at the best fit. The Fisher information matrix, $I = J^T J$ describes the curvature of the cost surface at the best fit and quantifies local statistical properties of the inverse problem.[33]

It is instructive to consider how the singular value decomposition of J relates to the eigenvalue decomposition of I. The singular value decomposition of J is given by

$$J = U \Sigma V^T, \tag{2.5}$$

where U and V are unitary matrices and Σ is a diagonal matrix of singular values. The FIM is given by

$$I = V \Sigma^2 V^T. \tag{2.6}$$

The expression in Eq. 2.6 is an eigenvalue decomposition of I. Thus, the eigenvalues of the FIM are the squared singular values of the Jacobian.

The eigenvectors and eigenvalues of the FIM quantify the forward model's sensitivity to coordinated variation in parameter values. The eigenvectors corresponding to the largest eigenvalues of I are stiff parameter combinations, i.e., parameter combinations to which the model is most sensitive. Conversely, stiff parameter combinations are those model parameters that encode the most information about the data. In Fig. 2.1, the eigenvector corresponding to the stiffest direction $\boldsymbol{v}_1$ lies along the semi-minor axis of the ellipse and the eigenvector corresponding to the second stiffest direction $\boldsymbol{v}_2$ lies along the semi-major axis of the ellipse about the best fit.

In many models, the FIM eigenvalues exhibit an approximate log-linear spacing across many orders of magnitude, a phenomenon known as sloppiness.[13,33] Parameters most represented in eigenvectors with small eigenvalues are considered sloppy parameters. Models are insensitive to large variations in sloppy parameters, and the presence of sloppy parameters indicates the data contain insufficient information to resolve such parameters individually. The implication is that sloppy parameters are irrelevant to the model's predictions and can be removed by parameter reduction methods such as the Manifold Boundary Approximation Method.[34]

The FIM also provides insight into the statistical properties of inferred parameters. The inverse of the FIM, I^{-1}, is a covariance matrix that when placed in the context of the Cramér-Rao inequality forms a lower bound on the uncertainty of parameter estimations. In terms of the singular value decomposition of J, the Cramér-Rao bound inequality is

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) \geq \mathrm{I}^{-1} = \mathrm{V}\Sigma^{-2}\mathrm{V}^{\mathrm{T}}. \tag{2.7}$$

The lower bound on the uncertainty for estimates of individual parameters from the model is defined by the diagonal elements of I^{-1}.

Large-valued entries of I^{-1} indicate high parameter uncertainty. Often the diagonal elements I_{ii}^{-1}, are of interest, i.e., the variances in the estimate of the ith parameters, which are proportional to the length of the ellipse about the best fit projected along the θ_i direction, as shown in green for the

toy model in Fig. 2.1. The CRB inequality (Eq. 2.7) effectively states that the parameter variances found along the diagonal of I^{-1} are the smallest variances possible given a model and data.

For sloppy models, the presence of small eigenvalues of the FIM translates into large uncertainties for inferred parameters. This can result from either of two causes. First, a model might be insensitive to changes in an individual parameter, such as is the case for parameters deep beneath the seafloor in a transmission loss model at sufficiently high frequencies. Second small eigenvalues can result from strong correlations among parameters. In physical systems, correlation is typically due to physical mechanisms that manifest similar effects in model predictions. In such cases the effect of one model parameter may be cancelled by tuning another. The model is often sensitive to many of the parameters individually but is insensitive to coordinated changes in combinations of parameters. In fact, if all but one of the model parameters were fixed to their true values, the sensitivity of the model to changes in the free parameter could be ascertained. Likewise, the free parameter's uncertainty could be found if correlations were eliminated from the model.

This theoretical uncorrelated parameter uncertainty, given by $1/\sqrt{I_{ii}}$, is called the conditional variance. The conditional variance for a parameter is always less than or equal to the variance obtained from CRB analysis. Disparities between the two variances are due to correlations among parameters, and the two variances become equal for a model with uncorrelated parameters. A comparison between the two variances effectively shows the potential for improving inversion estimates if correlation effects were eliminated for the physical model via a parameter reduction technique.

In terms of the cost surface picture of the inverse problem, the conditional variance corresponds to the width of the constant-cost ellipse measured through the best fit and parallel to a parameter axis, denoted by red in Fig. 2.1. The conditional variance is the lowest theoretical bound on the variance for a parameter if all the other parameters of the model were known. This occurs in models where the physical parameters are uncoupled. Thus, in the case of a model with uncorrelated

parameters, the variance and the conditional variance of individual parameters are equal, and the axes of the ellipses of constant cost align with the parameter axes.

The correlation between parameters can also be quantified with the Pearson correlation matrix $\mathrm{Cor}(\hat{\boldsymbol{\theta}})$. The Pearson correlation matrix is calculated by normalizing the covariance matrix

$$\mathrm{Cor}(\hat{\boldsymbol{\theta}})_{\mathrm{ij}} = \frac{\mathrm{Cov}(\hat{\boldsymbol{\theta}})_{\mathrm{ij}}}{\sqrt{(\mathrm{Cov}(\hat{\boldsymbol{\theta}})_{\mathrm{ii}}\mathrm{Cov}(\hat{\boldsymbol{\theta}})_{\mathrm{jj}}}}. \tag{2.8}$$

Entries of the correlation matrix vary between -1 and 1, where a value of 1 indicates a completely positive correlation, 0 indicates no correlation, and -1 indicates a completely negative correlation between parameters. Taking the absolute value of matrix entries allows a clear picture of the magnitude of correlation in the model. An example of a Pearson correlation matrix with the absolute value applied to the entries is shown in Fig. 2.5 and discussed as part of the case study in Section 2.4.

Understanding parameter correlation can guide the improvement of optimization routines used in acoustic inversions. In optimization procedures where parameters are optimized individually, parameter correlation poses a problem by limiting the search space. Indeed, parameter correlation implies that a more natural parameterization of the model exists which uncouples parameters.[35] In this natural parameterization, the parameters are independent so that the Cramér-Rao bound and conditional variances converge.

Because the FIM is independent of the observed data d_i and depends only on derivatives of the forward model, the lower bounds on the informativity of potential experimental data can be established by analyzing the model alone. Through the CRB, the FIM establishes the best-case variance for each model parameter, determining whether it is feasible to infer individual parameters to the desired resolution from given data. Additionally, analysis tools derived from the FIM including its eigenvector decomposition, the conditional variance, and the Pearson correlation matrix, yield insight into interactions among parameters. Armed with this information, modelers can make decisions about model choice, complexity, and experimental design for intended applications.

2.3 Computing Accurate Numerical Gradients

The nonlinear nature of ocean acoustic models often makes it difficult to obtain accurate gradients from finite-differences. Finite difference estimates balance a tradeoff between round-off error when step sizes are too small and truncation error when step sizes are too large.[36] The limited accuracy of finite-difference estimates is worsened for multi-parameter functions, such as in our forward model, and when the natural step size for individual parameters may be significantly different. Furthermore, gradient-based optimization strategies get stuck on plateau regions of a cost surface where the true gradient is near zero for sloppy models,[37] a challenge compounded by inaccurate finite-difference estimates.

To mitigate these problems, we apply two techniques to improve the numerical accuracy of estimating derivatives of the forward model. First, we precondition the parameters of the forward model, so a uniform step size can be used for finite-differencing. A function, called a preconditioner, transforms the model parameters $\boldsymbol{\theta}$ into preconditioned parameters $\boldsymbol{\phi}$. Second, we apply Richardson extrapolation to the preconditioned model to obtain accurate finite difference estimates of the gradient. With gradient estimates in hand, the chain rule is applied to transform the preconditioned gradient estimates back into the initial forward model parameterization, as demonstrated explicitly hereafter.

Preconditioners can be designed many ways, but the methodology we employ relies upon three physics-based principles: 1.) Applying a preconditioner to a model parameter should result in a unitless preconditioned parameter. 2.) Quantities included in a preconditioner, either for scaling or unit cancellation, should be native to the forward model. 3.) A unit step in the preconditioned parameters should probe a region of "local behavior" in the forward model. Local behavior is used here to describe domains over which the forward model has slowly varying oscillations. We confirm that we are probing local behavior by testing that our method for calculating numerical derivatives is robust to variations in the finite difference step size. In the case of Richardson extrapolation, if

the points selected by decreasing the step size exhibit clear limiting behavior, the local behavior has been captured. Figure 2.2, which is discussed throughout this section, shows acceptable local behavior in the top row of plots and undesirable local behavior in the bottom row.

As an example of applying the physics-based preconditioner design principles, consider a simplified ocean environment with 75 m of water and a single sediment layer over a basalt basement layer. Let compressional sound speed vary linearly within the sediment layer but hold density and the compressional attenuation coefficient constant within each layer. The forward model solves for the modulus of the complex Green's function from which the transmission loss at each receiver location can be obtained. This model is used in the case study in Section 2.4, and reference values for the environmental parameters for five sediment types are shown in Table 2.1.

Consider constructing a preconditioner $\phi(cp1_o)$ for the compressional sound speed parameter at the top of the water column, $cp1_o$. Compressional sound speed has units of length per time, so, to adhere to the first preconditioner design principle, this parameter must be placed in a ratio involving quantities with units of length and time to cancel out units. What quantities should be chosen to cancel units? The second preconditioner design principle answers this question, indicating that the quantities should be native to the physics of the forward model. Since the parameter $cp1_o$ belongs to the water layer, other water layer quantities should be considered before considering sediment or basement quantities. Additionally, to ensure the preconditioner generalizes across frequencies for a unit step size, it is important that the preconditioner be a function of frequency. Taking these guidelines into account, physics-native candidates with units of length include the depth of the ocean, $h_{o,\text{ref}}$, sediment thickness, h_{sed}, the source-receiver range, r, and the characteristic wavelength of sound propagating through the ocean at a given frequency, λ_o. The only candidate quantity with units of (inverse) time is frequency, f.

With this set of candidate quantities, unitless preconditioners can begin to be constructed and tested against the third preconditioner design principle. As various unitless preconditioner

Table 2.1 Preconditioners $\phi(\theta)$ for physical parameters θ used to obtain accurate derivatives, constructed according to physics-based preconditioner design principles. Parameters are defined as follows: h, layer thickness, cp, compressional sound speed, ρ, density, and αp compressional attenuation coefficient. The numerals 1 and 2 denote the top and bottom of a layer, respectively. The subscript abbreviations o, sed, bas, and ref stand for ocean, sediment, basalt basement, and reference value, respectively.

θ	$\phi(\theta)$
h_o	$\frac{h_o r}{\lambda_o h_{o,\text{ref}}}$
$cp1_o$	$\frac{fr}{cp1_o}$
$cp2_o$	$\frac{fr}{cp2_o}$
h_{sed}	$\frac{h_{\text{sed}} r}{\lambda_{\text{sed}} h_{\text{sed,ref}}}$
$cp1_{\text{sed}}$	$\frac{fr}{cp1_{\text{sed}}}$
$cp2_{\text{sed}}$	$\frac{fr}{cp2_{\text{sed}}}$
$\rho 1_{\text{sed}}$	$\frac{\rho 1_{\text{sed}} \lambda_o}{\rho 2_{o,\text{ref}} \lambda_{\alpha p1}}$
$\rho 2_{\text{sed}}$	$\frac{\rho 2_{\text{sed}} h_{o,\text{ref}}}{\rho 1_{\text{sed,ref}} h_{\text{sed,ref}}}$
$\alpha p1_{\text{sed}}$	$\frac{\alpha p1_{\text{sed}} f \lambda_{cp1,\text{sed}}}{1000}$
$\alpha p2_{\text{sed}}$	$\frac{\alpha p2_{\text{sed}} f \lambda_{\text{cp2,sed}}}{1000}$
cp_{bas}	$\frac{cp_{\text{bas}}}{f(h_{o,\text{ref}}+h_{\text{sed,ref}})}$
ρ_{bas}	$\frac{\rho_{\text{bas}}}{\rho 2_{\text{sed}}}$
αp_{bas}	$\frac{\alpha p_{\text{bas}} f \lambda_{\alpha p,\text{bas}} (h_o+h_{\text{sed}})}{1000r}$

permutations are explored, the construction that best captures the local behavior of the model within a unit step is selected as the best preconditioner. Figure 2.2 shows the local behavior of two possible candidate preconditioners for $cp1_o$ at various frequencies. The top row of plots in Fig. 2.2 shows a preconditioner that captures local behavior and mitigates frequency dependence, while the bottom row of plots shows a preconditioner that does not achieve either effect. For $cp1_o$ the best unitless preconditioner for our model is $\frac{fr}{cp1_o}$. This choice of preconditioner leads to oscillations in the model output that vary over similar scales regardless of frequency and sets the domain such that it captures local behavior. The procedure applied in this example can be extended to the other parameters in the forward model.

One advantage of forming preconditioners according to the three physics-based principles is the potential for insights into the physics of the problem. For example, the preconditioner for $\rho 1_{\text{sed}}$, the density at the top of the seafloor sediment layer, is $\frac{\rho 1_{\text{sed}} \lambda_o}{\rho 2_{o,\text{ref}} \lambda_{\alpha p1}}$, where $\rho 2_{o,\text{ref}}$ is the density of the ocean used as a reference, $\lambda_o = cp_o/f$ is the characteristic wavelength of sound propagating as a compressional wave through the ocean, and $\lambda_{\alpha p1}$ is the characteristic length of decay at the top of the sediment layer. This last quantity is calculated as $\lambda_{\alpha p1} = \frac{1}{|\tilde{k}|}$, where $\tilde{k}$ is the complex wave number found in Eq. 23 of Westwood *et al.*[6]

The fact that this particular preconditioner for $\rho 1_{sed}$ best satisfies the physics-based preconditioner design principles implies that the ratio of densities at the interface between the ocean and the sediment is an important feature of the physics of the problem. Indeed, the ratio of densities at the interface controls acoustic impedance, the physical mechanism that determines the coupling of sound from one medium to another. Additionally, the presence of the unitless length scaling term implies that the dominant phenomena of propagation in the water and decay in the sediment play important roles in the physics of the acoustic field. Similar conclusions can be drawn by examining the other preconditioners in light of the physics. The formulation of physics-based preconditioners is helpful for gaining an intuition about the interactions of physical parameters within the model.

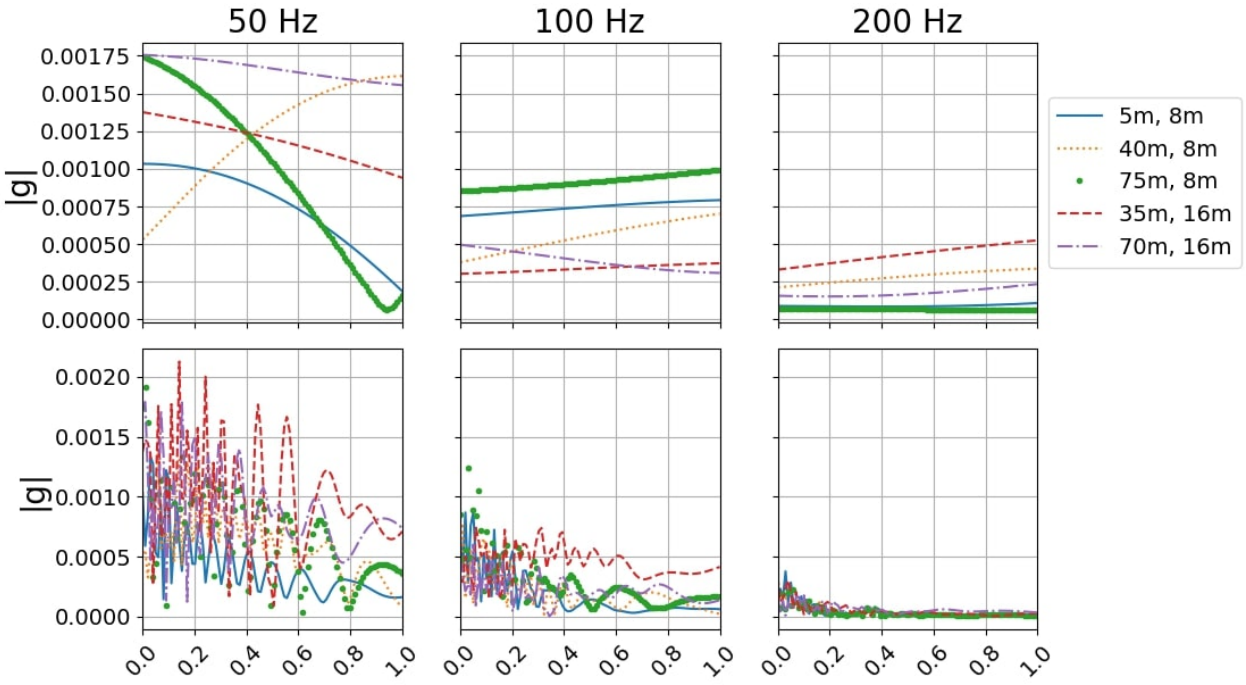

Figure 2.2 Variation in magnitude of Green's function over a unit step in the preconditioned parameter for two different preconditioners of $cp1_o$. The top row shows the best preconditioner $\frac{fr}{cp1_o}$ and the bottom row shows a poor preconditioner $\frac{cp1_o}{fh_{o,\text{ref}}}$. The curves represent the modulus of the Green's function $|\text{g}|$ for various source-receiver pairs. The legend gives the depth of the receiver followed by the depth of the source beneath the surface of the ocean in meters. Notice how the best preconditioner captures a smoothly varying region of local behavior with a unit step in the preconditioned parameter for all frequencies considered. The poor preconditioner (bottom row) features oscillations on a scale that would make finite-differencing difficult with a unit step size.

Once physics-based preconditioning improves the tractability of finite-differencing, Richardson extrapolation is applied to obtain accurate gradients. The Richardson extrapolation applies Richardson's deferred approach to the limit[38] to the estimation of numerical derivatives by using the first-order finite difference estimate at decreasing step sizes to extrapolate to the true derivative value. While the method does require many model evaluations, it effectively implements a higher order finite-difference formula that gives more accurate results. For an introduction to the technique of Richardson extrapolation see *Numerical Recipes* by Press *et al.*[36]

For this work the open-source Python package `numdifftools` is used for Richardson extrapolation to compute gradients.[39] The package `numdifftools` returns both the estimated gradient and an estimation of the error in the gradient. In this work, forward-step Richardson extrapolation used with preconditioning led to an improvement of six orders of magnitude in error estimates over using forward-step Richardson extrapolation alone. Additionally, since each preconditioner has an analytic expression, the chain rule can be applied directly to numerical derivative estimations. The existence of an analytic expression means that chain rule transformations are accurate to machine precision and are not a primary source of error in derivative calculation.

It is instructive to note that finite-differencing works best on smooth and continuous functions. Therefore, transmission loss, which features discontinuities of negative infinity at pressure nulls where the argument of the log approaches zero, is not a suitable candidate for finite-differencing. Alternatively, the modulus of the Green's function, which is used to calculate the transmission loss, is a smooth function without discontinuities, and therefore our choice for finite-differencing in the case study in Section 2.4. The chain rule can then be applied to obtain an expression for the Jacobian

$$J = \frac{\partial m}{\partial \theta} = \frac{\partial m}{\partial g}\frac{\partial g}{\partial \phi}\frac{\partial \phi}{\partial \theta}, \tag{2.9}$$

where m is transmission loss and g is the Green's function obtained from the model. The matrices $\frac{\partial m}{\partial g}$ and $\frac{\partial \phi}{\partial \theta}$ are sometimes called preconditioners because they are analogous to left and right

preconditioner matrices in linear algebra. The left preconditioner is chosen to set the threshold for error estimates, e.g., transmission loss of 1 dB re 1 μPa, and the right preconditioner is designed from physics-based principles to facilitate finite differencing.

It may be useful to compute J with respect to another parameterization. For example, taking the log of parameters removes the effect of differences in parameter scaling due to units in calculating J. Alternate parameterizations of J are achieved via the chain rule by right multiplying the expression in Eq. 2.9 by an additional right preconditioner $\frac{\partial \theta}{\partial l}$, where l is is the new parameterization. For the case study in the following section, J is calculated in this way with respect to log parameters.

2.4 Case Study

A case study of a simplified ocean acoustic model of transmission loss is presented to demonstrate the utility of Fisher information analysis. A laterally uniform shallow ocean is used with 75 m of seawater over a two-layer seabed. The seabed is composed of a 35 m sediment layer of either gravel, sand, silt, clay, or mud over a 100 m basalt basement layer. The model is parameterized by environmental parameters in each layer: h layer thickness (m), cp compressional sound speed (m/s), ρ density (g/cm^3), and αp compressional attenuation coefficient (dB/m-kHz). Values for the full 13 parameters in this case study are displayed in Table 2. For clarity, the digits 1 and 2 have been attached to parameters to indicate parameters corresponding to the top and bottom of a layer, respectively. Layer is indicated by subscript: "o" for the ocean water layer, "sed" for the sediment layer, and "bas" for basalt basement layer.

ORCA, a range-independent normal-mode model for acousto-elastic sound propagation, is used as the forward model.[6] ORCA is configured to simulate a vertical array of 15 receivers spaced equally throughout the water column and 2 acoustic sources at depths of 8 m and 16 m below the surface of the ocean at a range of 3 km from the array. For each source-receiver pair, ORCA

Table 2.2 Reference parameter values for case study ocean model. The simplified ocean models consist of 3 layers: 75-m water column over a 35-m sediment layer of either mud, clay, silt, sand, or gravel over 100-m basement layer of basalt. Compressional sound speeds cp are assumed to be constant in the water column and basement. In the sediment layer, the compressional sound speed is assumed to increase linearly at 1 m/s per meter moving deeper into the sediment. All parameters except h are defined in the model by the value at the top and bottom of the layer. Values of cp, ρ, and αp from *Computational Ocean Acoustics* Table 1.3 unless otherwise noted.[2]

Medium	h (m)	cp (m/s)	ρ (g/cm^3)	αp (dB/m-kHz)
Seawater	75	1500	1.0	-
Mud	35	1485–1520	1.6[1]	0.04[2]
Clay	35	1500-1535	1.4	0.13
Silt	35	1575-1610	1.7	0.63
Sand	35	1650-1685	1.9	0.48
Gravel	35	1800-1835	2.0	0.33
Basalt	100	5250	5.0	0.02

[1]Potty *et al.*[40] p. 11, last 2 sentences

[2]Bowles *et al.*[41] Table I., K column

computes the complex Green's function values at each receiver location for each source (see Eq. 16 in Westwood *et al.*[6]). From the modulus of the Green's function the transmission loss is calculated for each source-receiver pair, making a total of 15 x 2 forward model output values. In order to perform Fisher information analysis, the number of model predictions must be greater than or equal to the number of model parameters, otherwise the problem becomes under-specified. To encourage the computational convergence of ORCA's solvers, a non-physical, acoustically reflective mathematical halfspace layer is included beneath the basalt basement layer.

Leveraging physics-based preconditioners and Richardson extrapolation as presented in Sec. 2.3, numerical derivatives are obtained to construct a Jacobian for environments with each of the five sediment types at 50, 100, 200, and 400 Hz. A transmission loss threshold of 1 dB re. 1 μPa is selected for the left preconditioner in the Jacobian calculation (see Eq. 2.9), allowing for sensitivity analysis and parameter uncertainty analysis corresponding to a 1 dB uncertainty in transmission loss.

Using the Jacobian, the FIM and its eigenvector decomposition are calculated, as well as the covariance matrix I^{-1} and the Pearson correlation matrix. The variance and conditional variance for each parameter are also found. All results are presented in terms of log parameters to remove the effect of unit differences. Together these mathematical objects form the basis for the Fisher information analysis, the results of which are presented for a case study hereafter. For brevity and clarity, representative results are presented alongside commentary extending observations to all five seabed environments.

The eigenvalues of the FIM, shown in Fig. 2.3, span 12 to 17 orders of magnitude for a sand sediment. The log-linear spacing of the eigenvalues is characteristic of sloppy models. This sloppy model behavior is consistent across the frequencies and sediments studied. The largest eigenvalues correspond to the stiffest directions in parameter space, i.e., the parameter combinations to which

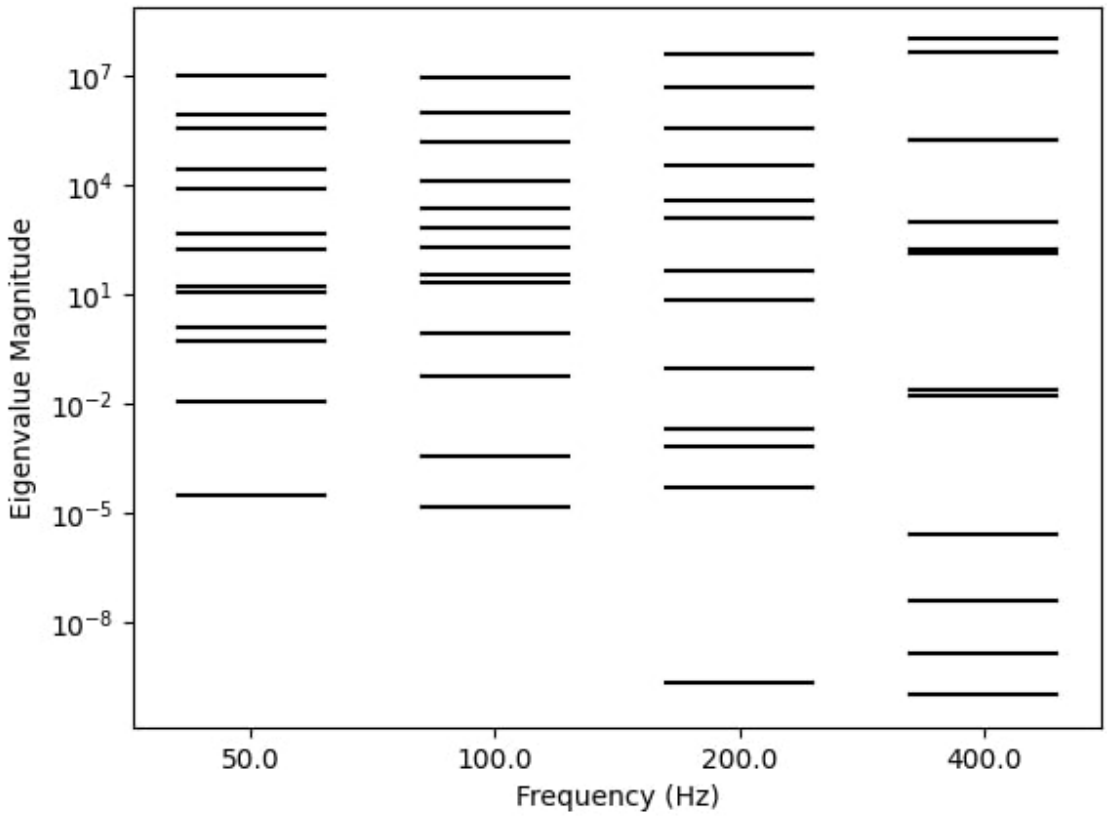

Figure 2.3 Fisher information eigenvalue magnitudes at frequencies 50–400 Hz for sand sediment environment. Notice the log-linear eigenvalue spacing characteristic of sloppy models.

the model is most sensitive. A visualization of the magnitude of the entries of the eigenvectors of the FIM in log parameters is shown in Fig. 2.4 for a sand sediment at 100 Hz.

Evidence of correlated relationships between parameters is clearly shown in the eigenvector columns of Fig. 2.4. If model parameters were uncorrelated, each eigenvector would point along a single parameter axis, and there would not be multiple elements in a column with dark colors. While some understanding of correlated parameter interaction is discernible in the FIM eigenvector visualization, the Pearson correlation matrix illustrates parameter correlations more clearly.

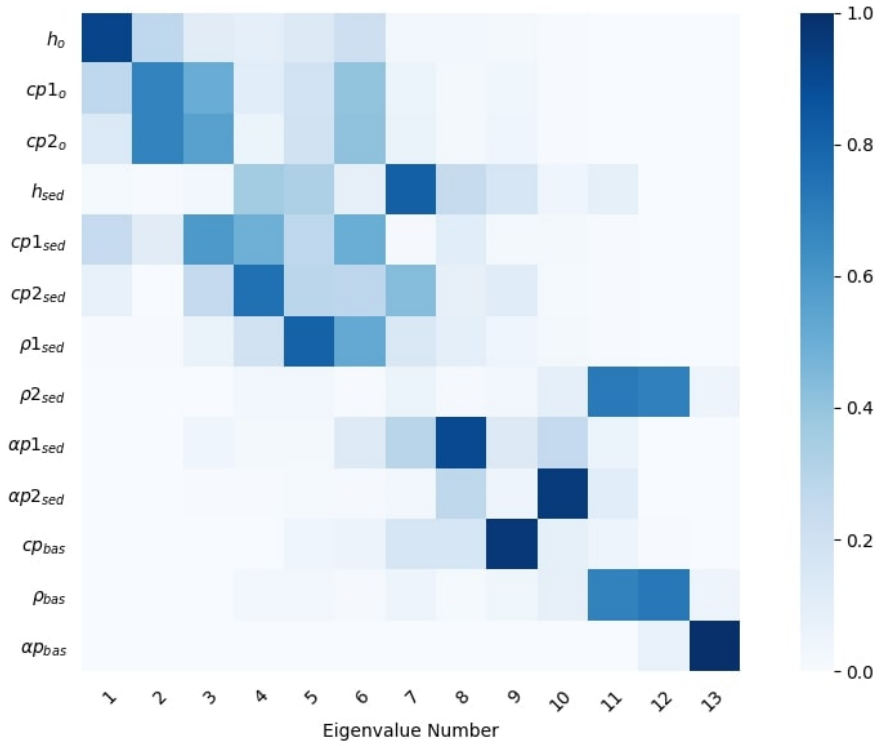

Figure 2.4 Absolute value of entries of eigenvectors of the FIM with respect to log parameters at 100 Hz for a sand sediment. The horizontal axis shows the eigenvalue number in decreasing order that corresponds to each eigenvector of the FIM. Parameter names are shown on the vertical axis. Each vertical column represents a unit-length eigenvector in parameter space. Matrix elements with values closer to 1 and indicate eigenvectors that point strongly in along the associated axis in parameter space.

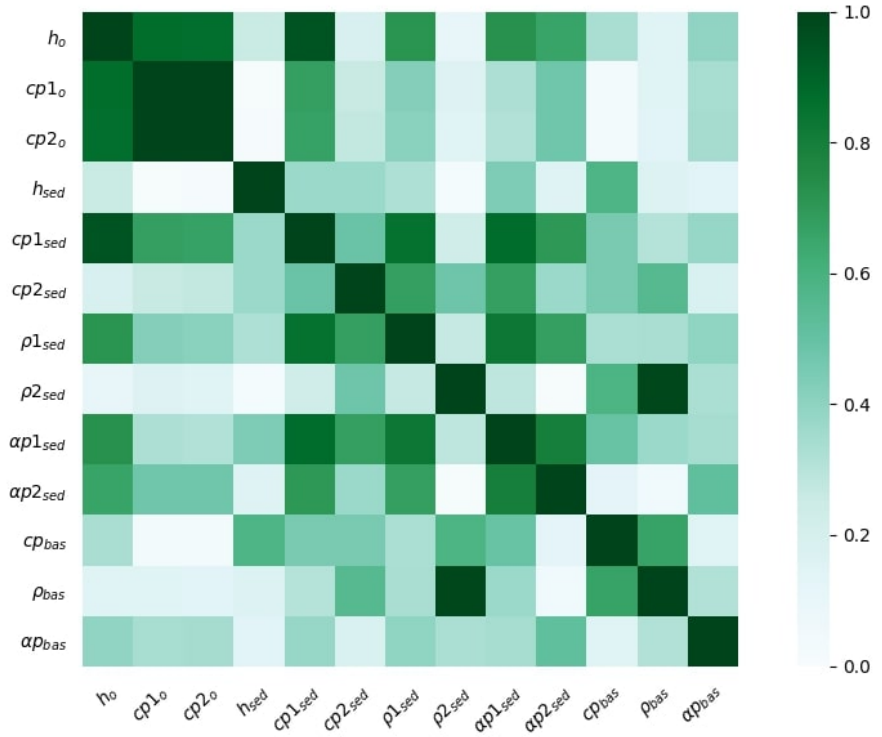

Figure 2.5 Absolute value of Pearson correlation matrix for a sand sediment at 100 Hz. The correlation matrix shows parameter relationships. In many cases, strong correlations occur between parameters from the same layer or parameters that mark a change at a layer interface.

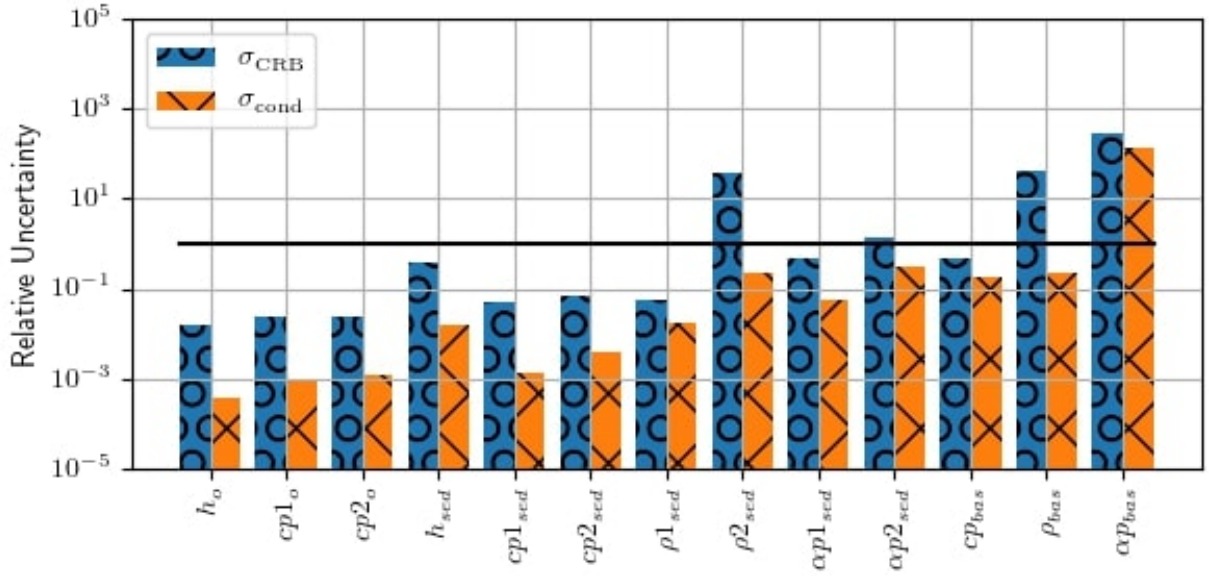

Figure 2.6 Relative uncertainty, σ_{CRB}, associated with the CRB and the relative uncertainty associated with the conditional variance, σ_{cond}, for each log parameter for a sand sediment environment at 100 Hz. A solid black line has been placed at a value of 1 indicating the point at which the uncertainty in the parameter is equal to the magnitude of the parameter's reference value. In general, the deeper the parameter is beneath the surface of the ocean, the larger the relative uncertainty. Also, relative uncertainty tends to increase with frequency. All relative uncertainty values are based on a transmission loss threshold of 1 dB re 1 μPa in Eq. 2.9.

The Pearson correlation matrix, calculated according to Eq. 2.8, uncovers parameter relationships and interactions. The absolute value of the Pearson correlation matrix for a sand sediment environment at 100 Hz is displayed in Fig. 2.5. A block of strongly correlated parameters (h_o, $cp1_o$, and $cp2_o$) is evident in the upper left corner. Parameters associated with this block are all ocean parameters, so it is reasonable that they are strongly correlated with each other but not with deeper seabed parameters, especially those describing the basalt basement layer. Similarly, $\rho 2_{\text{sed}}$ and $\rho 1_{bas}$ are strongly correlated, which matches with the physical understanding that the density ratio across the boundary between media plays a key role in acoustic propagation.

Notice the high correlation between many of the parameters, a feature of many sloppy models. This feature is present across all cases considered in this study, consistent with the results of the FIM eigenvectors. Another general trend is that as frequency increases from 50 Hz to 400 Hz, parameters deep beneath the seafloor become less correlated with other parameters. This is because higher frequencies do not penetrate as deeply into the seabed as low frequencies of the same energy. This trend is more pronounced in hard sediments (gravel, sand, silt) than in soft sediments (mud, clay) where sound propagates farther.

Working in log parameters has the benefit of removing the influence of physical units. Additionally, uncertainty in log parameters corresponds to a relative uncertainty in physical parameters, since a small variation in log parameters $\delta log(\theta)$ corresponds to $\delta\theta/\theta$. The relative uncertainty associated with the CRB, $\sigma_{\text{CRB}} = \sqrt{I_{ii}^{-1}}$, can be compared against the relative uncertainty associated with the conditional variance, $\sigma_{\text{cond}} = 1/\sqrt{I_{ii}}$, for the ith parameter. This comparison allows for better understanding parameter uncertainties, correlation, and model behavior. By way of illustration, σ_{CRB} and σ_{cond} for a sand sediment environment at 100 Hz are shown in Fig. 2.6.

As expected, Fig. 2.6 shows that the σ_{cond} are less than σ_{CRB} values. This difference is the result of parameter correlation in the model and is not surprising given the sloppy character of the model. Notice also the general trend in Fig. 2.6 that deeper parameters in the sediment have larger

relative variances. This quantitative assessment matches the physical intuition that as sound travels farther and passes through multiple media it loses energy. The trend of increasing parameter depth and uncertainty is present across the sediments and frequencies studied.

At 100 Hz, most of the parameters in the ocean layer and the sand layer have σ_{CRB} values well within 15% of the magnitude of the reference value. However, $\rho 2_{\mathrm{sed}}$ and $\alpha p 2_{\mathrm{sed}}$, both of which describe physical conditions at the bottom of the sediment layer, feature large σ_{CRB} values. The relative uncertainty associated with the conditional variance of these two parameters is still within 15% of the reference values providing further confirmation of the parameter relationships seen in the correlation matrix of Fig 2.5. This is evidence that there may exist a parameterization of the model with decreased correlation effects. In the basalt basement layer, only cp_{bas} has a relative uncertainty less than the magnitude of the reference value. The other two basement parameters have large relative uncertainties, with αp_{bas} having σ_{CRB} values on the order of 10^2.

These values can guide the design of efficient geoacoustic inversion: invert for those parameters with small σ_{CRB} values. These are the parameters for which the signal carries the most information and that will have the best resolution in an inversion. Inversions for parameters with high σ_{CRB} values yield results with high uncertainty.

By examining σ_{CRB} values across a range of frequencies and sediments, more trends emerge. For example, Fig. 2.7 shows σ_{CRB} and σ_{cond} values for a 35 m mud sediment layer at 50–400 Hz, and Fig. 2.8 shows σ_{CRB} and σ_{cond} values for a 35 m sand sediment layer over the same frequencies. The results shown in these two figures are representative of soft and hard seabed sediments, respectively.

As has been noted previously, the uncertainty increases with depth for both hard and soft seabeds. The soft seabed environment generally has lower σ_{CRB} values for deeper parameters better than the hard seabed, and that frequency impacts σ_{CRB} as well. These trends validate expected physical behaviors of underwater sound propagation.

For example, the less dense, soft sediments allow acoustic energy to propagate more deeply into the sediment than the denser hard sediments. However, the characteristic wavelength of sound at a given frequency also plays a role. For example, the σ_{CRB} value for h_{sed} changes with frequency in the mud environment. At 50 Hz in the mud environment the characteristic wavelength is on the order of the sediment layer thickness. This similarity in scales renders the model insensitive to the sediment layer thickness parameter. However, increasing the frequency to 200 Hz decreases the uncertainty with a drop of three orders of magnitude in σ_{CRB} values. This effect can be seen in several of the parameters. The identification of such "sweet spots" can improve inversions by understanding which frequencies or frequency bands are best at resolving various environmental features and in demarcating differences between environments.

The relative uncertainty plots (Fig. 2.7 and Fig. 2.8 show the impact of decay as sound energy is dispersed and absorbed in the ocean. A trend for h_{sed} in sand differs from mud, because the hardness of the sediment means that the dominant physical phenomenon in the sand layer is decay rather than propagation. Although the characteristic wavelength of sound at 200 Hz and 400 Hz is better for resolving the depth of the layer, the acoustic signal decays too much by the time it reaches the bottom of the sediment layer to relay information about the layer's depth back to the receivers. Thus, an analysis of these relative uncertainty plots reveals a trade-off between resolution reliant on characteristic wavelengths and resolution reliant on signal strength. While only two of five seabeds included in this study are shown, the results are representative of the range of outcomes found in the Fisher information analysis. Such results indicate that the model captures relevant physical behaviors of the situation and yields insight into parameter uncertainties and model sensitivities.

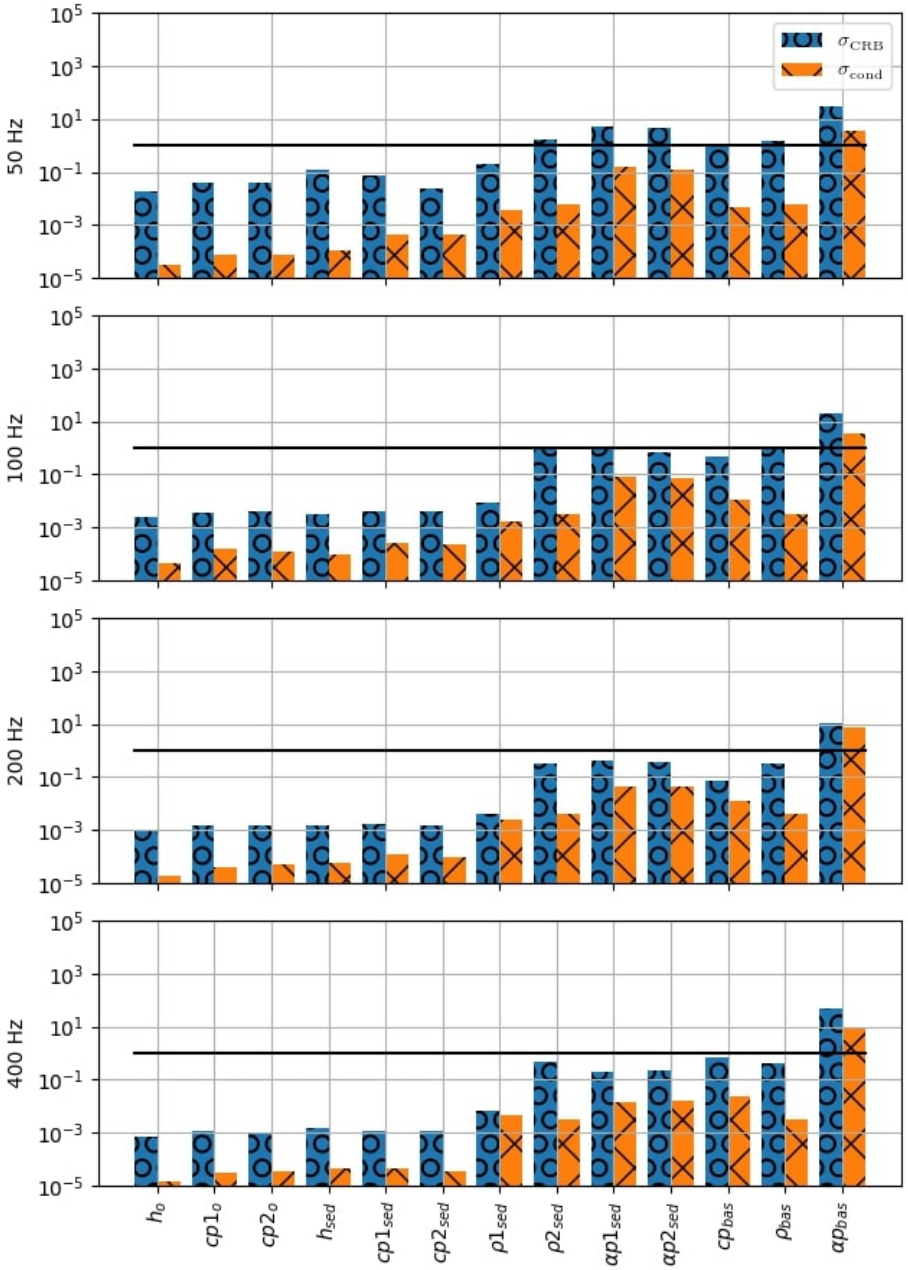

Figure 2.7 Relative uncertainty values associated with the Cramér-Rao bound (σ_{CRB}) and the conditional variance (σ_{cond}) for each of the model parameters at 50–400 Hz for a mud seabed. Relative uncertainty values based on an assumed resolution of 1 dB re 20 μPa for transmission loss. Bars for σ_{CRB} shown in blue with circles and σ_{cond} shown in orange with diagonal cross-hatching.

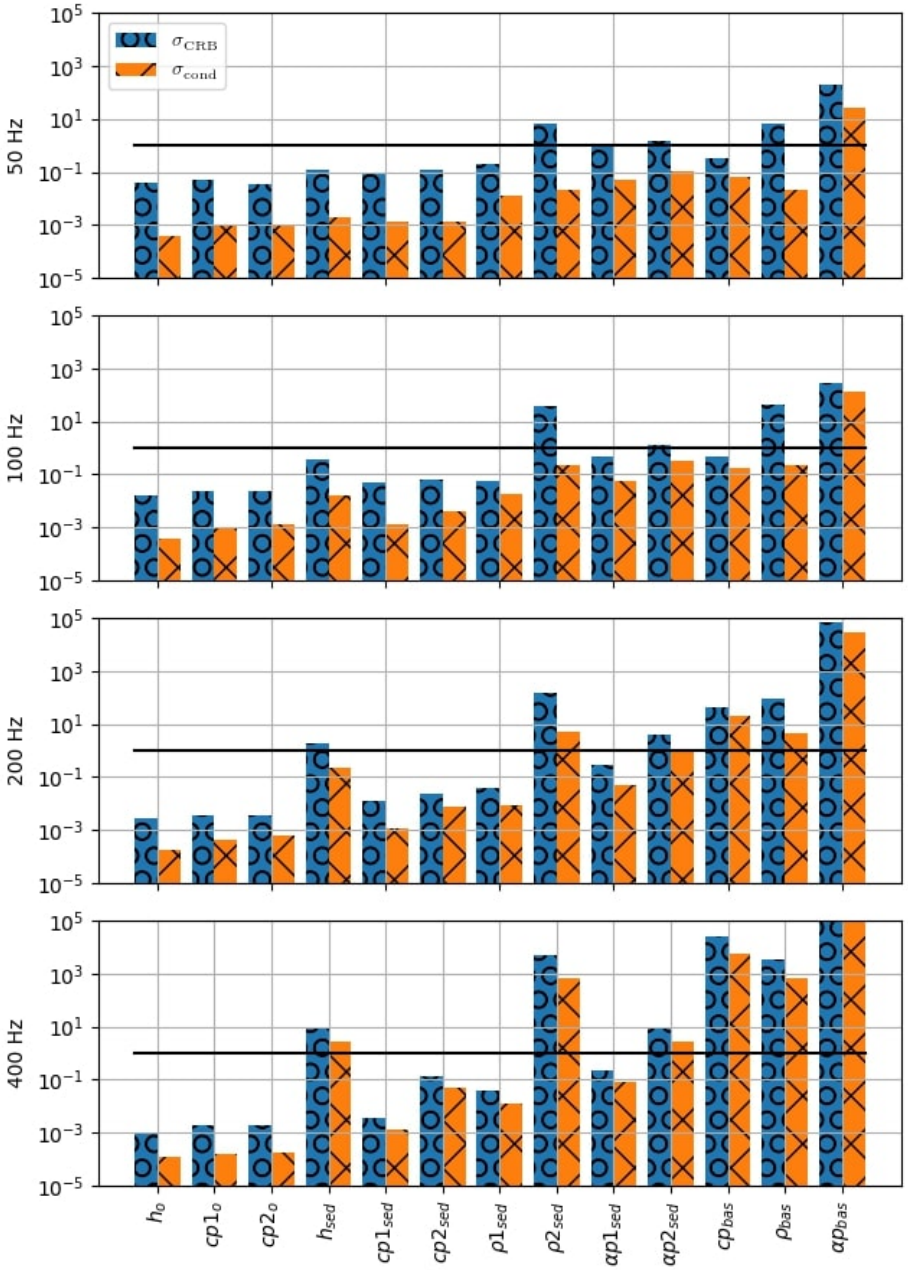

Figure 2.8 Relative uncertainty values associated with the Cramér-Rao bound (σ_{CRB}) and the conditional variance (σ_{cond}) for each of the model parameters at 50–400 Hz for a sand seabed. Relative uncertainty values based on an assumed resolution of 1 dB re 20 μPa for transmission loss. Bars for σ_{CRB} shown in blue with circles and σ_{cond} shown in orange with diagonal cross-hatching.

2.5 Conclusion

This work presents a framework for Fisher information analysis made possible by accurate numerical derivatives. The Fisher information matrix quantifies how much information a signal contains about the parameters in a model and can be used to quantify parameter uncertainty and model sensitivity. The FIM also provides insight into the sloppiness of a model and the correlation of its parameters. Understanding of parameter relationships and sloppy and stiff parameter combinations lays the groundwork for model reduction work in ocean acoustics. Additionally, for those interested in the application of machine learning tools such as neural networks and Bayesian maximum entropy to acoustic inversions, Fisher information analysis can help identify fundamental limits to parameter inference and can guide problem formulation.

While accurate numerical derivatives of ocean acoustic models have proven a difficulty in the past, we have shown that the application of physics-based preconditioners combined with Richardson extrapolation leads to accurate numerical derivative estimates. The introduction of preconditioning techniques and the use of Richardson extrapolation, while not exclusive to underwater acoustics, places powerful tools in the hands of the ocean acoustics community to move optimization and optimal experimental design forward.

As validation of the numerical derivative methodology and as demonstration of FIM analysis, a case study is presented for a transmission loss model of a simplified ocean environment with five different seabeds in the frequency range from 50–400 Hz. Results show which parameters are best resolved by the model with a given environment and frequency. Such results are useful for experimental design. Prior to performing an experiment, simulations can be run to identify the limitations of that experiment's design, i.e., understanding inherent uncertainties in parameter estimations using the CRB. Likewise, the analysis can compare the performance of various models and aid in the process of model selection. For example, the Fisher information could be used to

explore differences in the inversion capabilities of the acousto-elastic model for underwater sound propagation and the viscous grain shearing model for sound propagation.

This work is a step forward in connecting the theoretical world of mathematical models with a physics-based understanding of the ocean acoustics world. By using the Fisher information analysis, it is possible to allow physical data to influence the creation and refining of mathematical models of real-world physics that are informed by the physics, rather than by human assumption and intuition.

Chapter 3

Synthetic Dataset Development

The popular adage *garbage in, garbage out* as used in the context of computer science and machine learning refers to the idea that using poor input data decreases the odds of a useful output from a program or model. Therefore, those who design and curate datasets strive to ensure quality. Especially when training a machine learning model to perform regression, a dataset should span the space of possible inputs, thus improving the model's ability to interpolate and return reliable predictions. Machine learning models are notoriously untrustworthy at extrapolation.[42]

Physics experiments are often costly and time-intensive, so data scarcity is a common problem. This problem is only compounded by the demand from the machine learning algorithm's side for copious quantities of training data. Focusing on the dataset side of this problem, one potential panacea is to augment a small dataset by generating plausible new samples, perhaps with a Generative Adversarial Network.[43] Another option, which is the option taken in this work, is to generate completely synthetic data with a physics-based model. If the synthetic data captures enough of the underlying physics, a machine learning model trained on the synthetic dataset hopefully performs well on real-world data. However, discrepancies, such as whether the simulated data accounts for noise, whether the physics-based model is detailed enough, or whether the simulation spans the

space of real-world variability, can lead trained machine learning models to generalize poorly to real-world data.

To frame the development of a synthetic dataset, the real-world dataset from the 2017 Seabed Characterization Experiment (SBCEX2017) consisting of SOO spectrograms labeled with parameter uncertainties is introduced. The generation of synthetic SOO spectrograms with the forward model ORCA is then discussed, with details on dataset labeling, including the labeling of geoacoustic parameter uncertainties.

3.1 Uncertainty in Parameter Estimates

Uncertainty is a term that has many definitions and sources. In machine learning, uncertainty is discussed when analyzing the trustworthiness of predictions. Machine learning models, including neural networks, map an input dataset distribution to an output prediction distribution. In the case of regression, the input and output distributions are continuous. Uncertainty in predictions arises from two sources: the input data and the machine learning model.

Irreducible errors due to randomness in the input data is called *aleatoric uncertainty* from the Latin *alea*, "a game with dice".[44] Aleatory uncertainty is an inherent property of the input data distribution and is typically addressed by increasing the quantity and quality of training data. Measures of aleatory uncertainty can be directly calculated by statistical analyses of datasets.

The second source of uncertainty is the model. As a model trains its weights are adjusted to assimilate information learned from the training data. A trained model contains the "knowledge" or "experience" learned from the dataset. When a model is presented with an out-of-distribution input, it does not necessarily have the knowledge or experience to accurately produce an output prediction. This type of uncertainty is called *epistemic uncertainty* from the Greek *episteme*, "knowledge".[44] As with aleatory uncertainty, epistemic uncertainty can be improved by the addition of training

data of a higher quantity and quality for a given machine learning model. However, unlike aleatory uncertainty, epistemic uncertainty is more difficult to estimate, a difficulty which has given rise to multiple uncertainty quantification approaches such as evidential learning, Bayesian neural networks, Monte Carlo dropout variational inference, Gaussian processes, ensemble techniques, and many others. See the literature review by Abdar *et al.*[45] for more on uncertainty quantification in deep learning up through 2020.

In Chapter 2, Cramér-Rao bound, the lower bound on the uncertainty in forward model parameters, was computed. This uncertainty is neither the aleatoric uncertainty in the data nor the epistemic uncertainty of the machine learning model. Rather, the CRB should be considered a property of the *a priori* parameter distribution needed to generate data with a given aleatoric uncertainty. Chapter 2 explained how the CRB for ORCA's geoacoustic parameters is related to information content in transmission loss output. Therefore, a spectrogram of transmission loss contains information both about the parameters that led to its generation and their CRB uncertainties. These CRB parameter uncertainties are employed as an additional label for the synthetic spectrograms, as described in this chapter, in hopes that similar parameter uncertainties might be estimated for real-world data.

3.2 Real-World Dataset

In March and April of 2017, a multidisciplinary project called SBCEX2017 (Seabed Characterization Experiment 2017) studied a region of the continental shelf in the Atlantic Ocean called the New England Mud Patch. While the experiment site, which is roughly 95 km south of Martha's Vineyard, Massachusetts, is referred to as a 'Mud Patch', the seabed there is composed of a variety of sediments including mud, clay, silt, and sand. The bathymetry of the site is relatively flat and is known to have "an extensive region with comparatively thick and comparatively uniform upper

layer of fine-grained sediment".[46] SBCEX2017 occurred in March and April to take advantage of a stable, nearly isothermal, isovelocity sound speed profile in the water column.[46]

Two, sixteen-channel vertical line arrays (VLAs) were deployed by the Hodgkiss group at Scripps-Marine Physical Laboratories to record merchant vessels traveling in nearby shipping lanes.[10] A schematic of the VLAs and the shipping lanes is shown in Fig. 3.1. Passing merchant ships were identified using vessel traffic data collected by the U.S. Coast Guard through the Automatic Identification System and retrieved from the Marine Cadastre.[46,47] Each ship was labeled by its speed, measured in knots, and its CPA (closest point of approach distance in km) to the VLA. Because each passing ship represents an opportunity to study the seabed, the ships are dubbed Ships of Opportunity (SOO).[48]

The real-world SBCEX2017 dataset consists of 31 SOO spectrograms. The SOO spectrogram images were generated from signals received on a hydrophone stationed 33 m above the seafloor, roughly in the center of VLA1 and VLA2. The spectrograms depict 15-minute intervals, centered on the time at CPA, as ships approached and then sailed away from CPA. The frequency range was chosen to be 360–1500 Hz to avoid corruption from the 350 Hz tones produced by research vessels during the experiment. Of the 31 spectrograms, 13 were measured at VLA 1 and 18 were measured at VLA 2. One SOO spectrogram image of the Viking Bravery ship measured on VLA1 is shown in Fig. 3.2.

3.3 Synthetic Dataset

Training and validating a neural network on a dataset of 31 samples is infeasible, so it is necessary to generate a synthetic dataset that resembles the real-world dataset. Therefore, a synthetic dataset was developed, composed of 22,680 SOO spectrograms generated with ORCA, a range-independent normal mode model with azimuthal symmetry. ORCA calculates the transmission loss between

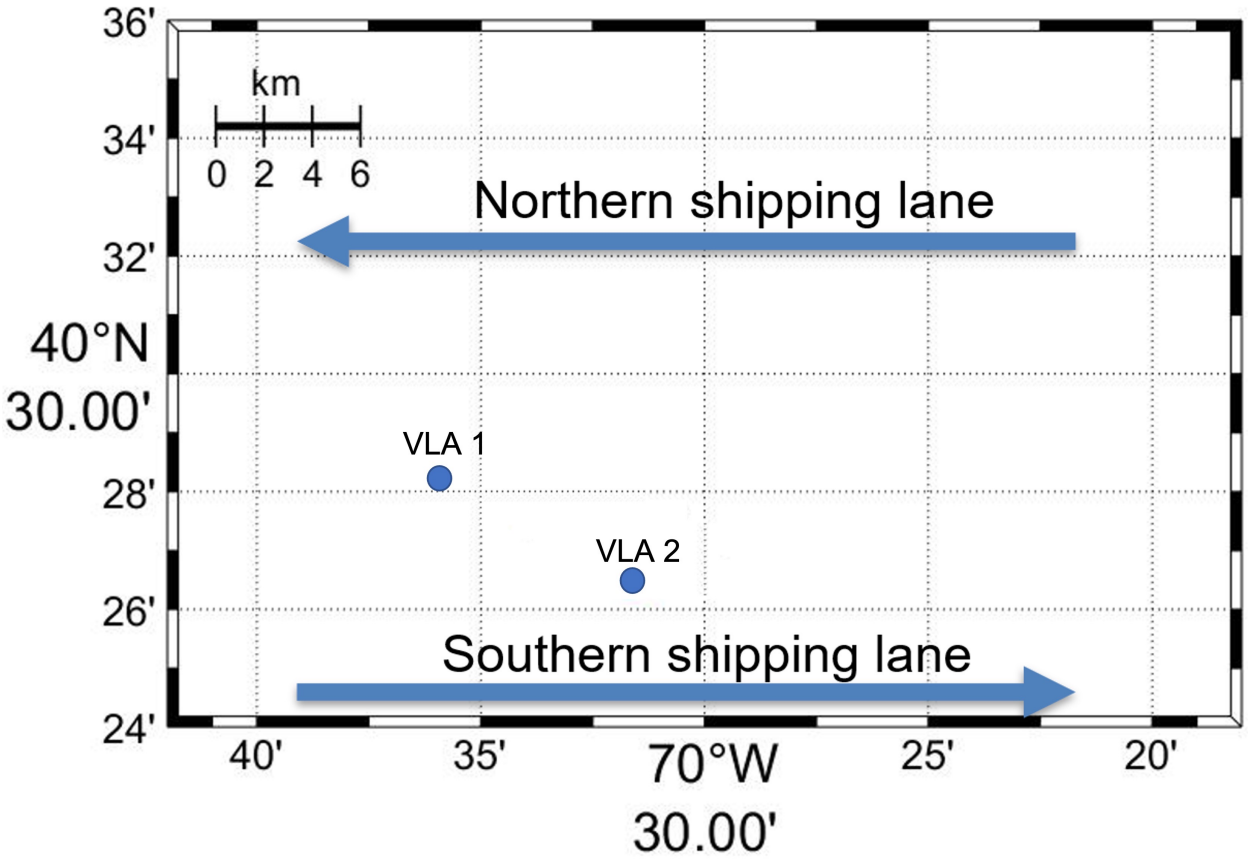

Figure 3.1 SBCEX2017 site in the New England Mud Patch. The deployment locations of VLA1 and VLA2 are marked for reference as well as the shipping lanes. The real-world SOO spectrogram data used for generalization was obtained using the setup.

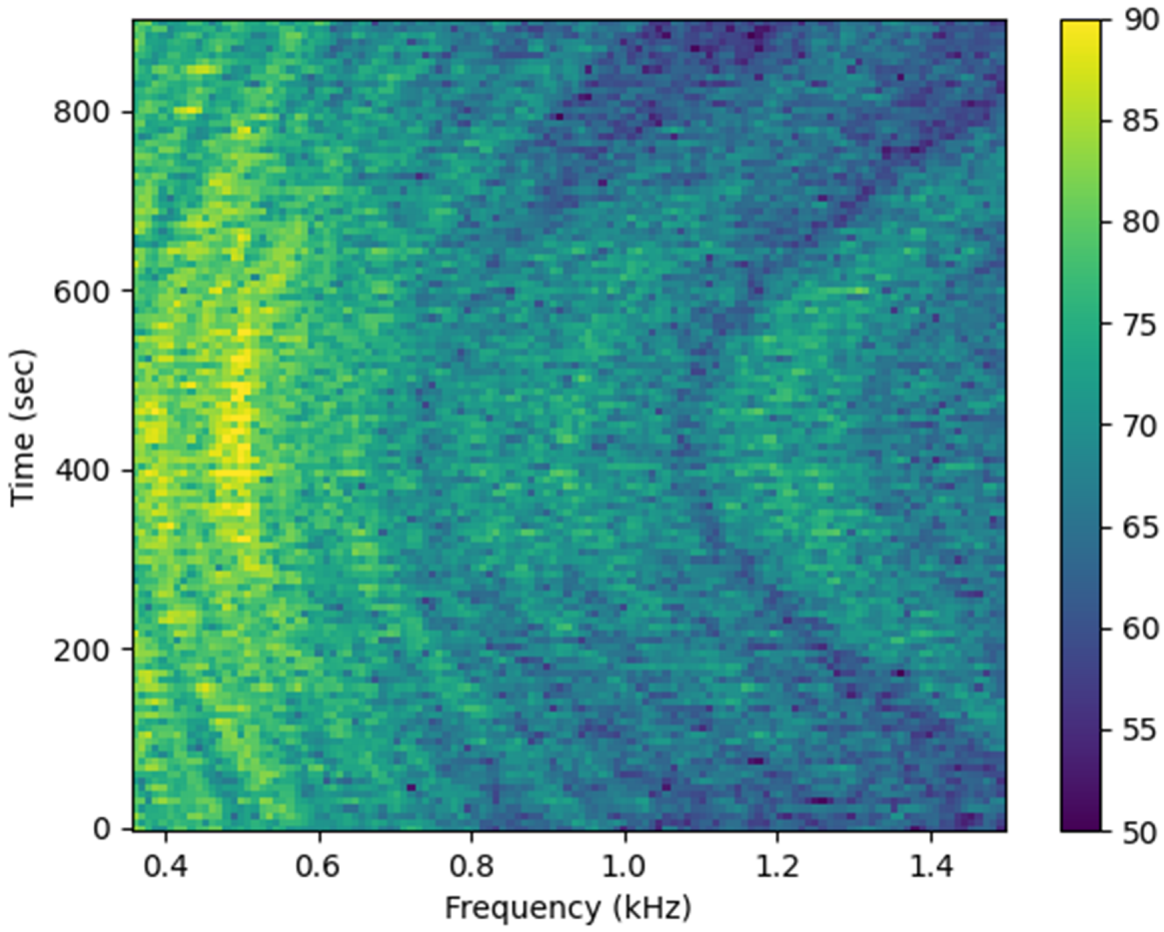

Figure 3.2 SOO Spectrogram from the Viking Bravery recorded on VLA 1 at approximately 33 m above the seafloor. The "bathtub" striation pattern typical of SOO spectrograms is due to the shifting of interference patterns as the ship approaches and then leaves CPA.[1]

each source-receiver pair. Transmission loss, TL, is the drop in sound pressure level between the source level, SL, and the received level, RL, as defined by the SONAR equation,

$$\mathrm{TL}(\mathrm{f}, \boldsymbol{\theta}, \mathrm{z_s}, \mathrm{z_r}, \boldsymbol{r}) = \mathrm{SL}(\mathrm{f}) - \mathrm{RL}(\mathrm{f}, \boldsymbol{\theta}, \mathrm{z_s}, \mathrm{z_r}, \boldsymbol{r}), \tag{3.1}$$

where f is frequency, $\boldsymbol{\theta}$ is a vector of environmental parameters, z_s is the source depth, z_r is the receiver depth, and $\boldsymbol{r}$ is a vector of source-receiver ranges. Since the source level is not a function of the environment, the uncertainty in the received levels with respect to environmental parameters is only dependent on the transmission loss. Refer to Chapter 2 for more information about ORCA. As with the SBCEX2017 data, the synthetic spectrograms cover 123 frequencies in the 360–1500 Hz band at 121 time steps for the 15 minutes around CPA (7.5 minutes on either side of CPA sampled at 7.5 second increments.)

The ocean is modeled as 75 meters of water with a surface water sound speed of 1469 m/s that increases linearly to a bottom water sound speed of 1470 m/s. Water density is held constant at 1 $\mathrm{g/cm^3}$. These ocean parameter values are consistent with those obtained from the SBCEX2017 survey and estimates of the sound speed during the experiment.[46]

The experimental area has been studied extensively. Results of a full-waveform chirp sonar survey (Fig. 3 of Wilson *et al.*[46]) shows a layer of mud on top of sand. The two-way acoustic travel time map (Fig. 2 of Wilson *et al.*[46]) indicates that the thickness of the mud varies around the measurement area up to a thickness of 12 m. For the sediment sound speed, each of the 19 published estimates, included in Wilson *et al.* Fig. 7, found different values ranging from speeds typical of mud to those typical of silt or sand, evidence of complexity in the exact layering of the sediment.[46]

In accordance with published results, the seabed was modeled with a single layer of sediment over a thicker sand or gravel layer. A fluid model is used for the top layer because shear waves are often negligible in water-saturated sediments. With this model, 140 seabed environments were generated. For each environment, the top layer was randomly generated by varying the parameters

of one of five prototypical seabed sediments ranging from fine-grained to coarse-grained: mud, clay, silt, sand, and gravel. Of the 140 generated, 40 apiece came from mud and clay and 20 apiece came from silt, sand, and gravel.

The generation process maintained ratios in the compressional sound speed and density across the water-sediment interface. The density and compressional attenuation were assigned constant values within layers. The compressional sound speed was constrained to vary linearly from top to bottom with a non-negative gradient 0–1 m/s for each meter of depth. The thickness of the sediment layer was assigned a value between 5 and 35 m.

All second layer parameters were drawn from the coarsest two sediments, sand and gravel. To improve the convergence of ORCA's solver, parameters in the second layer are constrained to values greater than or equal to the corresponding top layer sediment parameter values. Additionally, for the purposes of solver convergence, a non-physical half-space layer is included beneath 100 m of the basement layer with a compressional sound speed of 6000 m/s, a density of 5 g/cm^3, and an attenuation of 0.83 dB / m-kHz.

The bounds for the parameters and parameter ratios are drawn randomly from a uniform distribution. Bound values are drawn from Table 1.3 in *Computational Ocean Acoustics*.[2] Possible values the parameters can take on are summarized in Table 3.1. The geoacoustic parameter values used to generate each environment were saved as labels for later use in training and validation.

Table 3.1 Bounds on geoacoustic parameters used to generate environments for the synthetic dataset. Sampled values were constrained to remain constant or increase with depth.

Parameter	Mud	Clay	Silt	Sand	Gravel	
$cp_w/cp1$	[0.97, 1.0)	[1.0, 1.05)	[1.05, 1.1)	[1.1, 1.2)	[1.2, 1.22)	-
$\rho_w/\rho 1$	[1.5, 1.6)	[1.5, 1.7)	[1.4, 1.9)	[1.9,2.0)	[2.0, 2.1)	-
$cp1/cp2$	[0, 1)	[0, 1)	[0, 1)	[0, 1)	[0, 1)	-
h (m)	[5,35)	[5, 35)	[5, 35)	[5, 35)	[5,35)	75
$cp1$ (m/s)	[1425.9, 1470)	[1470, 1543.5)	[1543.5, 1617)	[1617, 1764)	[1764, 1793.4)	1469
$cp2$ (m/s)	[1425.9, 1505)	[1470, 1578.5)	[1543.5, 1652)	[1617, 1799)	[1764, 1828.4)	1470
$\rho 1 \mathrm{g/cm}^3$	[1.5, 1.6)	[1.5, 1.7)	[1.4, 1.9)	[1.9,2.0)	[2.0, 2.1)	1.0
$\rho 2 \mathrm{g/cm}^3$	[1.5, 1.6)	[1.5, 1.7)	[1.4, 1.9)	[1.9,2.0)	[2.0, 2.1)	-
$\alpha p1$ (dB/m-kHz)	[0.04, 0.13)	[0.13, 0.20)	[0.48, 0.63)	[0.33, 0.48)	[0.33, 0.63)	-
$\alpha p2$ (dB/m-kHz)	[0.04, 0.13)	[0.13, 0.20)	[0.48, 0.63)	[0.33, 0.48)	[0.33, 0.63)	-
cp_{bas} (m/s)	[1650, 1800)	[1650, 1800)	[1650, 1800)	[1650, 1800)	[1650, 1800)	-
ρ_{bas} g/cm^3	[1.9, 2.0)	[1.9, 2.0)	[1.9, 2.0)	[1.9, 2.0)	[1.9, 2.0)	-
αp_{bas} (dB/m-kHz)	[0.33, 0.63)	[0.33, 0.63)	[0.33, 0.63)	[0.33, 0.63)	[0.33, 0.63)	-

To ensure sufficient variability in the data, ship speeds, CPAs, source depths, and geoacoustic parameters were varied as well. For each of 140 seabed environments, data samples were generated that combine different source parameters: three ship source depths, nine ship speeds, and six CPAs. The three ship source depths were randomly drawn from a uniform distribution 6–12 meters below the surface of the water. A source depth of 6-12 meters models noise from the surface ship's engine and the resulting water turbulence. For each source depth, nine ship speeds are randomly sampled, three from each range of values: 8–12 kn, 12–16 kn, and 16–20 kn. Finally, six CPAs are randomly sampled, three from each range of values: 0.5–8 km and 8–15 km. This makes for a total of 140 environments x 9 source depths x 3 ship speeds x 6 CPAs = 22,680 SOO spectrograms. Source levels were randomly assigned by sampling a uniform distribution over the range of 220–240 dB referencing 1 μPa.

3.4 Uncertainty Labels

One key consideration in designing labels for datasets is how the labels will be interpreted in various contexts. For instance, from the point of view of an ocean acoustic inversion, the exact value of a parameter's uncertainty is less important than whether the uncertainty is qualitatively acceptable. For example, it does not matter whether a parameter estimate has 75% or 100% uncertainty when only uncertainties less than, say, 10% are acceptable. Uncertainty values either indicate untrustworthiness or reliability. Therefore, interpretation of uncertainties in parameter estimates has a qualitative aspect.

If the purpose of uncertainty labeling is to get a handle on the qualitative reliability of parameter estimates, it is important to note that not all uncertainty values are of equal interest nor are they treated equally in interpretation. An estimate with an associated uncertainty of 50% may be just as untrustworthy as an estimate with an associated uncertainty of 500% for a given application. Thus,

it is helpful to define ranges of interest for uncertainty values. Uncertainty value ranges of interest for this work are defined in Table 3.2.

Chapter 2 established that, at a given frequency, the information content concerning a geoacoustic parameter in a transmission loss signal can be used to estimate the CRB, a lower bound on the uncertainty in parameter estimates. For a SOO spectrogram containing 123 frequencies, 123 CRB values can be calculated for each parameter (following the same formulation set out in Chapter 2, but switching out the many receiver depths for many source-receiver ranges). However, performing regression for 123 local uncertainty labels is impractical for at least two reasons. First, the computational expensive to compute the CRB for a dataset of thousands of SOO spectrograms at 123 frequencies (across all the ranges) is prohibitively high. Secondly, and more importantly, sampling a spectrogram at every point violates the underlying independence assumption required for the calculation of traditional statistical moments like mean and standard deviation. To satisfy independence and reduce computation time, sparse sampling of CRB values is the logical choice. Thus, the following heuristic is adopted for labeling the parameter uncertainties for each SOO spectrogram with a single value.

The spectrogram is sparsely sampled at $k = 10$ frequency slices equally spaced between the spectrogram's minimum frequency, 360 Hz, and maximum frequency, 1500 Hz. For each of the selected frequencies, the CRB is calculated for each parameter of interest. On occasions where ORCA stalls at one of the ten evenly spaced frequencies, a new frequency is randomly drawn from a uniform distribution across the 360-1500 Hz band and the calculation is repeated.

Once the uncertainties are calculated at ten frequencies, they are binned. The range of possible uncertainties is divided into seven bins of interest. The bins are not divided equally, following the observation above that different uncertainty values are not interpreted in the same way. Rather, they are approximately logarithmically spaced. The bounds for each bin are given in Table 3.2.

Table 3.2 Relative uncertainty bins for the labeling of SOO spectrograms. For improved weight convergence in the neural network and for ease in interpretation, each bin is associated with a unit domain.

Percent Relative Uncertainty	Label Mapping Domain
0-0.1%	[0,1]
0.1-1%	[1,2]
1-5%	[2,3]
5-10%	[3,4]
10-25%	[4,5]
25-50%	[5,6]
50-100+%	[6,7]

A single uncertainty label assigned to each spectrogram is found as follows. The bin with the greatest number of uncertainties is determined. This bin is referred to as the "mode bin", borrowing the term "mode" from statistics where it refers to the sample value that occurs most often. It is assumed that if CRB values for a spectrogram most frequently land in a specific uncertainty bin, then a neural network asked to estimate a CRB value would also likely most often predict values that belong in the same bin. Taking the arithmetic mean of the uncertainties in the mode bin yields a representative uncertainty label for the spectrogram. If multiple bins contain the mode number of samples, the lowest uncertainty bin is used to calculate the uncertainty label.

For potential gains in training speed, the uncertainty labels are mapped onto the range from 0 to 7. Each uncertainty bin is allotted a domain between two consecutive integers onto which the uncertainty can be mapped via linear interpolation. The mapped domains for each uncertainty bin are shown in Table 3.2. The entire dataset of 22,680 synthetic spectrograms took around 4 days to generate with this labeling scheme. The largest bottleneck was the uncertainty labeling, with the

uncertainty calculation for a single SOO spectrogram taking 16-30 seconds, varying with ORCA hangups.

The decision to sample each spectrogram at $k = 10$ points is based on an exploration of various values of k to determine the least k-value that gave similar results to sampling at k=74 (every 15 Hz). It was found that k=10 matched the mode uncertainty bin found with k=74 with 93% accuracy. This sampling strategy was deemed sufficient for the purposes of qualitative uncertainty interpretation.

This uncertainty labeling process yields a single value with qualitative meaning. It is less important that the neural network exactly estimate the uncertainty label than that it returns a value within the proper uncertainty bin. Note also that the CRB labels here are a lower limit on parameter estimate uncertainties. Consistent with the case study in Chapter 2, the CRB values used in the labeling process correspond to the least possible uncertainty in a particular parameter estimate given an uncertainty in the transmission loss data of +/- 1 dB referencing 1 μPa. The uncertainty in transmission loss data is a function of the measurement apparatus and signal processing, and a dataset representing the corresponding CRB uncertainty for a different experimental setup could be easily created using the method given here.

Chapter 4

Deep Learning

4.1 Background

Neural networks are popular in tackling inverse problems across disciplines. This work tends to fall into two categories: classification and regression. Images are a popular data type for neural networks, and, where spectral data is concerned, spectrogram images are used in many disciplines. Basic feed-forward neural networks are being used to try to diagnose Alzheimer's disease from speech spectrograms[49] and classify honeybee hive state from spectrograms of recorded hive sound.[50] Convolutional neural networks, or CNNs, are being used to classify seismic events,[51] estimate the amount of grain in silos from RADAR spectrograms,[52] classify back-bead formation in gas metal arc welding from welding current spectrograms,[53] estimating synthesizer settings from spectrograms of music,[54] and estimating the joint angles of a prosthetic robotic arm from electromyograph spectrograms.[55] A recurrent neural network, RNN, was used to categorize emotion from speech spectrograms,[56] and residual neural networks, ResNets, were used to classify and estimate source parameters of gravitational waves from black hole mergers from spectrograms of electromagnetic signals.[57]

As with other domains, neural networks have gained popularity in ocean acoustic inversion applications. In ocean acoustics, inversion applications often manifest as seabed classification, source-receiver geometric parameter regression, and geoacoustic parameter regression. Seabed classification with neural networks has been performed with CNNs, for side-scan sonar data,[58] simulated time-series pressure data,[9,59,60] simulated vertical line array pressure data,[61] and ship of opportunity (SOO) spectrograms.[9,10]

Source-receiver geometric parameter regression includes estimating parameters like range, closest point of approach CPA of a moving source, and source depth. Neural networks have been used to estimate source range,[60,62–64] ship speed,[9] source depth,[65] and direction of arrival.[66]

Geoacoustic parameter regression with neural networks, which is the primary focus of this work, goes back to the '90s. Early work focused on geoacoustic parameter regression with feed-forward neural networks with synthetic data from a normal mode model[67,68] and radial basis function neural networks (RBFNN) with simulated horizontal line array data.[69] More recent work includes regression for sound speed and attenuation from long-range time-series data with general additive models,[70] regression for a single sediment layer's thickness, top and bottom sound speed, sediment density, and basement sound speed with a RBFNN,[71] and regression for layer thickness, sound speed, density, and attenuation for a single sediment layer with a large feed-forward network.[72] In all these studies, the seabed is modeled as horizontally uniform single sediment layer and, in the cases of Mao *et al.*[72] and Shen *et al.*,[71] a single basement layer.

While previous work has focused on using neural networks to estimate sediment parameters, this work proposes to simultaneously estimate parameters and the uncertainty associated with parameter estimates. With the goal of estimating parameters and uncertainties in mind, the neural network architecture and training strategies for the current work are presented. A neural network is trained and validated on the synthetic SOO spectrogram dataset covered in Chapter 3, then evaluated for generalizability with a small sample of measured data from SBCEX2017.

4.2 Architecture and Training

One of the primary challenges in training CNNs and FNNs is the vanishing gradient problem, a problem which worsens as more layers are added to a network.[8] Residual neural networks, ResNets, were first proposed by He *et al.*[73] in "Deep Residual Learning for Image Recognition" as way to combat vanishing gradients and make deeper networks more feasible through the introduction of an identity mapping skip connection between layers. In the seabed classification work by Escobar *et al.*,[10] of the size networks tested on SOO spectrogram data, the ResNet-18 architecture had the highest consistent performance in cross validation testing. Therefore, ResNet-18, so-called for its 18 layers, is used in this work.[73] A final linear layer is added with an output size equal to the number of values requested in the regression. In total the network has about 11.2 million learnable parameters.

The ResNet-18 is trained on the synthetic dataset discussed in Chapter 3. The dataset, which consists of 22,680 spectrogram images, is randomly divided with an 80-20% split into a test set (18,144 spectrograms) and a validation set (4,536 spectrograms). A mean squared error (squared L2 norm) loss criterion is imposed and the network is optimized with the AdamW optimizer, an optimizer based on Adam but with weight decay corrections.[74] An early stopping mechanism is included that stops training when the validation loss does not improve for a patience length of four epochs. Improvement in the validation loss is counted when the validation loss decreases by at least 0.001. Preliminary testing of various batch sizes indicated that a larger batch size improved the training time with ResNet-18, so a batch size of 512 is used.

Normalization is used for both labels and input images. Data labels are normalized to the domain [0,1] to encourage the convergence of network weights and decrease training time. Spectrogram images are also normalized individually to the domain of [0,1] as follows. The spectrogram, whose pixel values represent sound levels in decibels referencing 1 μPa , is centered about zero and normalized to have a standard deviation of one by subtracting the mean and dividing by standard

deviation of the spectrogram's pixels. Any pixel values outside the range of three standard deviations on either side of the mean are then reassigned a random value drawn from a uniform distribution between plus and minus three standard deviations. The result is that outlier values are filtered out. The pixel values are then linearly mapped to the domain of [0,1].

Models are trained to perform either single regression for individual parameters or simultaneous dual regression for a single parameter and its uncertainty labels (see Chapter 3 for details on the creation of uncertainty labels) for each of four geoacoustic parameters: layer thickness h, sound speed at the top of the sediment $cp1$, density at the top of the sediment $\rho 1$, and attenuation at the top of the sediment $\alpha p1$. These parameters are selected for two reasons. First, in the case study in Chapter 2, all these parameter have CRB of 50% or less for most frequencies considered. Second, these parameters have published estimates against which they can be judged for the SBCEX2017 dataset.[46] The neural network architecture and training algorithms are implemented in PyTorch version 1.5.1 and run in Python 3.6.9 on an NVIDIA Tesla T4 GPU. Under these conditions, the entire process from loading data to finishing training takes less than twenty minutes for all models.

4.3 Validation

The root mean squared error, RMSE, and mean absolute percentage error, MAPE, for prediction on the validation set are reported for both the single parameter regression and the dual parameter regression in Table 4.1. In both the single and dual regressions $cp1$ and $\rho 1$ have the lowest error measures in their parameter estimates, while $\alpha p1$ and h have large associated errors. These results are consistent with the analysis in Chapter 2 where the CRB followed the same trend. Although it remains unclear whether ResNet-18 can obtain the least error (i.e., reaching the Cramér-Rao lower bound), the qualitative trend corroborates the findings in Chapter 2.

In the dual regression case, unlike the errors in the parameter estimates, the errors in the uncertainties do not appear to follow a pattern. Most explanations of the source of errors come back to the *ad hoc* nature of the uncertainty labeling heuristic described in section 3.4. The sampling of the spectrogram points used to generate the label was not always consistent due to the need to randomly sample when ORCA failed to converge. If the regularity of the sampling is the cause of this error, trying small perturbations of the inputs rather than random sampling might be enough to fix the issue, assuming smooth variation between the original point and the perturbation. If the evenly spaced regularity of the sampling is the issue, resorting to a random sampling technique to generate the uncertainty labels may work better. If this is the case, a new study of the number of sample points needed to consistently label the uncertainty would need to be performed.

Another possible cause of the network's failure to estimate the uncertainty labels well is the binning strategy. Since the bins are not equal, it is difficult to compare the uncertainty estimates of samples in different bins. The RMSE values would suggest that $cp1$ and h are consistently a bin or more off of their labels and that $\alpha p1$ is within a bin of the true value much of the time. This is deceptive if $\alpha p1$ typically has uncertainties in the 50–100% uncertainty bin and $cp1$ is trying to correctly estimate whether the error belongs in the 0.1–1% bin or the 1–5% bin. If unequal binning is the cause of the errors, then it is unrealistic to directly compare errors between uncertainty labels in different bins.

For reference, the validation results for the dual regression plotted in Fig. 4.1. For each parameter, the predicted values are plotted against labeled parameter values. The labeled uncertainty bin is indicated by color and the predicted uncertainty bins are broken out by subplot. The neural network generally groups labeled uncertainties together, though they are often mixed across neighboring bins. The upward diagonal trend across the plots shows that the network is learning to perform regression, albeit with varying degrees of success.

Table 4.1 Validation error analysis results from single and dual regression for four geoacoustic parameters: layer thickness h, sound speed at the top of the sediment $cp1$, density at the top of the layer $\rho 1$, and attenuation at the top of the layer $\alpha 1$. The dual regression includes the results for simultaneous regression for the uncertainty labels, σ with the associated parameter subscripted.

Single Parameter Regression				
	h (m)	$cp1(m/s)$	$\rho 1$ (kg/m^3)	$\alpha p1$ (dB/m-kHz)
RMSE	8.26	27.24	0.13	0.07
MAPE	46.0%	1.3%	5.9%	29.8%
Simultaneous Dual Regression				
	h (m)	$cp1(m/s)$	$\rho 1$ (kg/m^3)	$\alpha p1$ (dB/m-kHz)
RMSE	7.34	17.56	0.16	0.07
MAPE	38.3%	0.9%	7.5%	29.0%
	σ_h	σ_{cp1}	$\sigma_{\rho 1}$	$\sigma_{\alpha p1}$
RMSE	1.19	1.28	0.32	0.62
MAPE	33.9%	44.4%	10.0%	11.2%

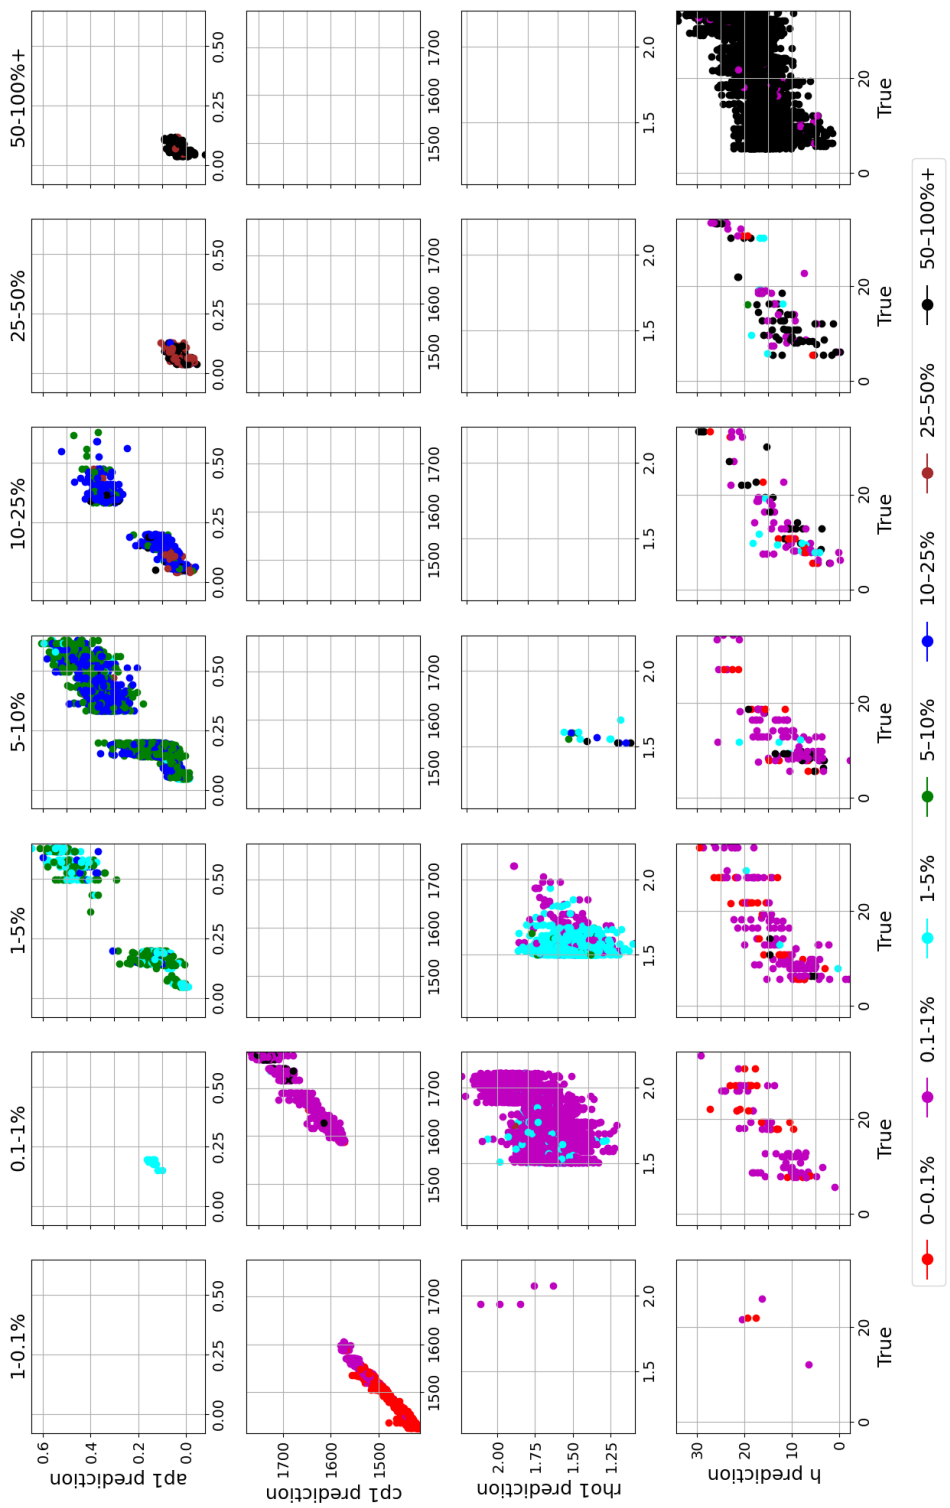

Figure 4.1 Validation results for the simultaneous regression of parameter values and parameter uncertainties. Subplots are broken out by parameter and predicted uncertainty bin. Points are colored according to labeled uncertainty bin. The axes of each plot are parameter prediction on the vertical axis and true value on the horizontal axis. For an exact regression, the points should align along the diagonal of each subplot.

4.4 Generalization

To test the generalizability of the models, measured data from the SBCEX2017 experiment is used. The data come from 31 hydrophone recordings of passing surface ships taken at one of two vertical line array locations. Information about the 31 ships is published in Table II of Escobar *et al.*[10] This work uses entries with CPA values detected on VLA1 and VLA2. To work with the trained ResNet-18 regression models, the data are extracted to have the same size as the synthetic training data: 123 frequencies between 360 and 1500 Hz by 121 time steps over 15 minutes.

Since the measured data were annotated by hand and labeled using the Marine Cadastre, it is possible that there is a small amount of human error in the labels. For example, the spectrograms might not be centered exactly on CPA. However, the translational invariance of convolutional layers should combat problems introduced by a shifted CPA. The true values for the geoacoustic parameters of the New England Mud Patch are not known precisely, so the results must be compared against the results of other inversions or surveys. Please refer to the comprehensive overview of work published on SBCEX2017 by Wilson *et al.*[46] For simplicity, two of the studies referenced in the overview are used for a comparison. Ballard *et al.*[3] measured the compressional sound speed at the top of the sediment $cp1$ at about 1480 +/- 15 m/s with an acoustic coring system. The same study measured the mean bulk density $\rho 1$ to be about 1.5 +/- 0.2 g/cm^3 . Belcourt *et al.*[4] used a trans-dimensional Bayesian approach to invert for the attenuation of the sediment in the mud patch $\alpha p1$ and came away with values ranging from 0.01 to 0.09 dB/m-kHz and a top layer thickness h of 9.2 m. These values are included in the analysis of the generalization results for comparison against the results from the regression models, shown in Table 4.2.

The ResNet-18 regression models have mixed results when fed real-world data. For the single parameter regression models, the predictions were significantly greater than the published values except for $\alpha p1$. The same trend occurs for the dual regression case, but the values are closer to the reported SBCEX2017 values. Sound speed, $cp1$, is the greatest offender of overestimating. Of the

Table 4.2 Generalization results performing regression with ResNet-18 on the 31 measured SOO spectrograms from SBCEX2017. The mean and standard deviation are presented for both the single parameter regression case and the dual regression parameter and parameter uncertainty case. Representative nominal parameter values from the SBCEX2017 literature are included for context.[3,4]

	h (m)	$cp1$ (m/s)	$\rho 1$ (kg/m^3)	$\alpha p1$ (dB/m-kHz)
Single Mean	14.07	1648.54	1.78	0.27
Single Standard Deviation	2.74	81.48	0.13	0.11
Dual Mean	9.2	1587.61	1.64	0.18
Dual Standard Deviation	4.79	79.87	0.18	0.12
SBCEX2017 Mean	9.2	1480.0	1.5	0.5
SBCEX2017 +/- Error	-	15.0	0.2	0.4

31 measured ships, only four $cp1$ values were predicted within the range of values reported in the literature. The overestimating behavior could be the result of the frequency range chosen, since, for muddy sediment, sound at the low end of the 360-1500 Hz band is known to penetrate the seafloor and is more likely to carry information about deeper layers than high frequencies. If ResNet-18 is struggling to pick up on the correct frequency band to estimate the sound speed at the top of the sediment, the network may be compromising to account for the deeper, harder sediments than were used in creating the training data.

The thickness of the top layer is estimated with the dual regression model to have a mean value equal to the reported value. It is difficult to judge the model's performance, however, because the uncertainty in the reported value is not known. The thickness of the top layer of sediment in the seabed is not uniform across the entire propagation path of the sound from ship to receiver as our model assumes. Depending on how much the reality of the New England Mud Patch differs

from a truly uniform thickness, the model may be failing to capture enough of the physics to more accurately train a model to infer parameter values.

At first glance, attenuation has the best estimates with ResNet-18, but the large range of uncertainty in the reported values, means that this could be happenstance. The Cramér-Rao bound study in Chapter 2 showed that the attenuation can vary wildly in its minimum uncertainty, which makes it a difficult candidate for inversion.

On average the models perform best when predicting $\rho 1$, estimating densities within the range of published values for approximately one in three ship spectrograms. This performance may be the result of the model capturing enough of the physics of how the sediment density at the boundary between water and sediment affects the propagation of sound. Future work might focus on whether the density ratio at the water-sediment interface or a density gradient in a sediment layer would be good candidates for regression with a trained network.

Instead of considering the overall performance represented in Table 4.2, the results can also be evaluated for each measured data sample to look for trends. Generalization results for the single regression model are shown in Fig. 4.2. The single regression models consistently overestimate the parameter values, compared with published estimates. The model does not appear to prefer certain ships over others, an indication that CPA and ship speed do not seem to influence network performance.

Generalization results for the dual regression models are show in Fig. 4.3. The results of the simultaneous regression for parameter values tend to be closer to the published parameter values than the results of the single regression models. For $cp1$ and $ap1$ the network tends to overestimate parameter values compared with published estimates. The thickness of the top sediment layer h predicts a wide spread of values centered on the published value of 9.2 m. The network appears to have learned the features of sediments with a thickness around 9 meters, but not well enough to avoid estimating, oddly enough, a negative sediment thickness for the BBC Tennessee SOO

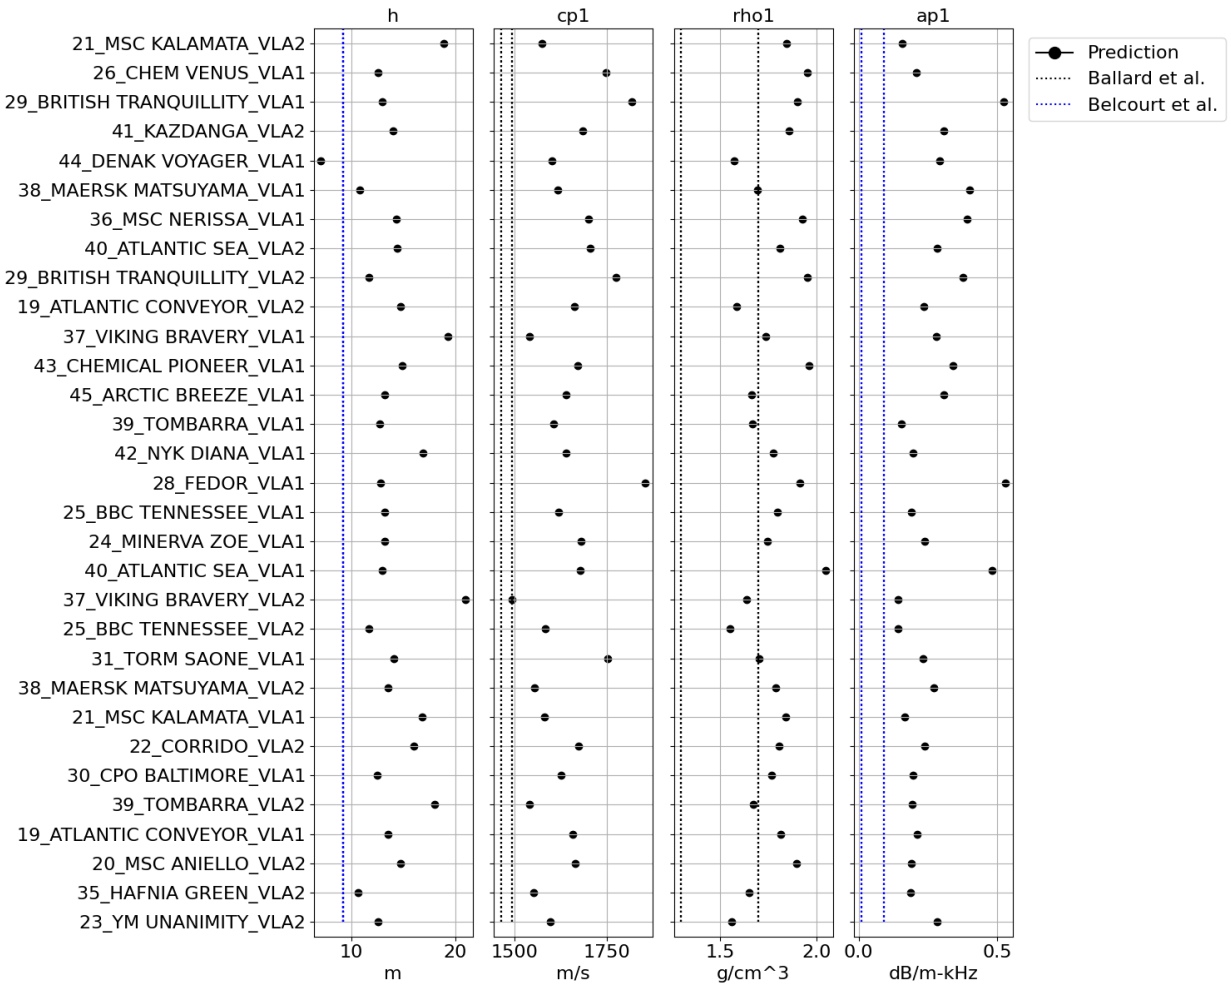

Figure 4.2 Generalization results for neural networks trained to estimate single parameter values. Each of the 31 SOO spectrograms from SBCEX2017 are represented. The VLA from which the measurements were obtained are listed beside the ship names. Reference values from literature are indicated with dotted lines. For each parameter, the trained ResNet-18 networks tend to overestimate compared to published values. The network does not appear to prefer certain ship spectrograms the simultaneous regression of parameter values and parameter uncertainties.

spectrogram from VLA2. The best performing parameter with the dual regression models, however, is $\rho 1$, for which over half of the network's predicted parameter values fall within the range of values published in the literature.

Evaluation of these generalization results point to some future research questions about the training dataset, network architecture, and training strategy. Focusing on the dataset, the best-case option would be to train a model on measured data. Since such an option is expensive and likely untenable, training with different quantities of data might prove fruitful. The quality of synthetic dataset might also be improved by using data augmentation techniques, either on measured or synthetic data.

A subtle difference between the synthetic SOO spectrograms generated by ORCA and the spectrograms extracted from the SBCEX2017 data exists. The SBCEX2017 spectrograms are generated via Fourier transform of a time-series signal, and thus represent a repackaging of a continuous source of information. ORCA-generated SOO spectrograms are created by stacking together outputs at discrete frequencies. Thus, in the SBCEX2017 data each frequency contains information from nearby frequencies while the ORCA data does not. If the spectrogram is too sparsely sampled, a network trained on synthetic data might fail to learn details about how the Fourier transform packs information from a band of frequencies into a lower-dimensional representation. A way to prevent this problem would be to use a higher resolution spectrogram. Although producing higher resolution spectrograms would slow down training and dataset development, it could lead to richer information content in each sample and a stronger trained network that is better at generalization.

On the network architecture side, future work might explore deeper ResNet architectures. In their study inferring synthesizer settings from spectrograms of music, Barkan *et al.*[54] reported that performance improved as they added layers to their CNN architecture. Of course, adding more

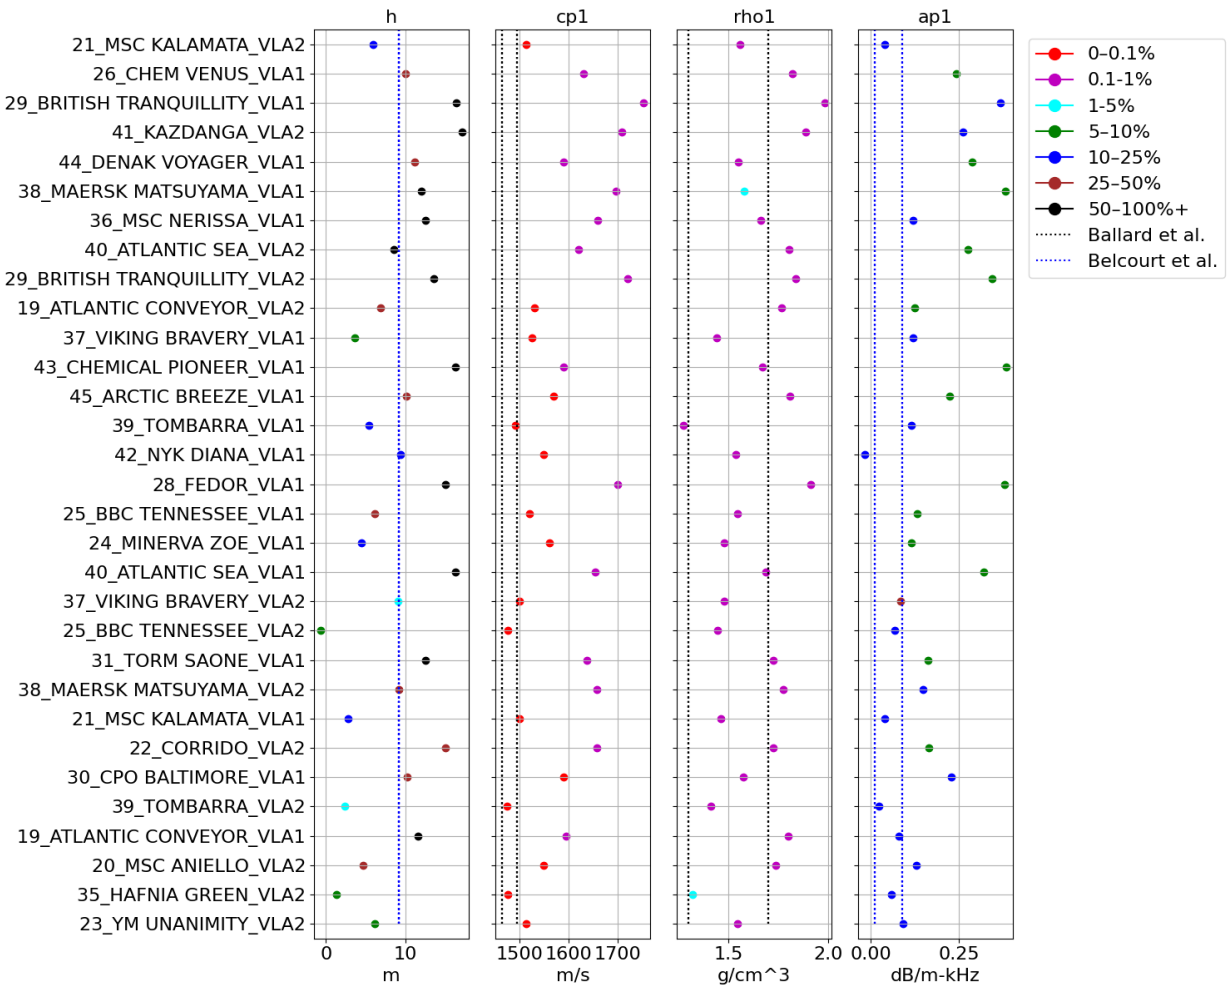

Figure 4.3 Generalization results for neural networks trained to simultaneously estimate parameter values and parameter uncertainty labels. Each of the 31 SOO spectrograms from SBCEX2017 are represented. The VLA from which the measurements were obtained are listed beside the ship names. Reference values from literature are indicated with dotted lines. The predicted uncertainty bin is indicated by the color of each point. For $cp1$ and $ap1$ the network tends to overestimate parameter values compared with published estimates. The thickness of the top sediment layer h predicts a wide spread of values centered on the published value of 9.2 m. The best performing parameter, however, is $\rho 1$, for which over half of the network's predicted parameter values fall within the range of values published in literature.

layers increases the size of the dataset needed to train the weights of a model, but there exist deeper, more powerful ResNets for this purpose.

Finally, from a training perspective, future work might consider the impact of adjusting hyperparameters. Batch size was chosen to be large based on experimentation for how long the training algorithm took to train the weights of a network, but the quality of the training was not assessed. A smaller or larger batch size might improve the resulting network. Perhaps decreasing or increasing the patience of the early stopping metric would produce a network with greater capacity for generalization. Also, other data normalization techniques exist that might improve training. If the filtering out of pixel values in the tails of the distribution causes a loss of valuable information, altering the normalization to accommodate keeping those pixels would be a potential source of performance gains.

Chapter 5

Conclusion

This thesis covers my research into some topics of interest to the ocean acoustic inversion and machine learning communities, with special interest paid to parameter uncertainty. In Chapter 2, the Cramér-Rao bound (CRB), a lower bound on the relative uncertainty of parameter estimates was calculated from the Fisher information. Accurate numerical derivatives, which rely on physics-based preconditioners and Richardson extrapolation, make it possible to calculate the Fisher information. The Fisher information also allows for forward model-based analysis of parameter correlations and model sensitivity, which can be used to design datasets and inform machine learning. Chapter 3 explored how one might use the Cramér-Rao bound to label a synthetic dataset. A key idea in this chapter was that uncertainties have utility both quantitatively and qualitatively. A general heuristic for labeling synthetic SOO spectrograms with a general uncertainty value was presented. Chapter 4 applied the synthetic dataset to training a ResNet-18 deep learning model for the purpose of generalization to 31 measured samples from SBCEX2017. The results of the generalization were mixed, but show promise given more time and work.

While this work focuses on local uncertainty, the results can be evaluated as an exploration of how local uncertainty might relate to global properties. Work in the field of sloppy models shows that the Fisher information reveals a model's hierarchy of sensitivity to parameters locally. Through

the Cramér-Rao bound, the sensitivities can be considered a hierarchy of uncertainties. These uncertainties were used to consider which parameters might make good candidates for inversion using a neural network.

Attempts at regression for the chosen candidate parameters indicate that, although the Cramér-Rao bounds predict a certain hierarchy, the results are not constrained to that hierarchy. Rather, the Cramér-Rao bound is what its name implies—a lower bound on the local uncertainty. Although $\rho 1$ for the sediment layer did not have the smallest Cramér-Rao bound, the networks trained to predict $\rho 1$ performed best in generalization to measured data. This does not mean that the Cramér-Rao bound is a poor guide for selecting candidate parameters for regression, only that it provides a limit on the potential performance. A neural network might be expected to perform regression more accurately for parameters with low CRB values than for parameters with high CRB values, but that assumes the network has been able to learn the approximate inverse function to go from data space to parameter space. Further exploration is needed to understand why, for example, $cp1$ of the first layer had low CRB values but was not predicted accurately with a trained ResNet-18. Likewise, more exploration is needed to understand why dual regression for the parameter and uncertainty label led to parameter predictions closer to the desired range of outcomes than the single regression parameter estimates.

The initial regression study in Chapter 4 showed mixed results for predicting parameter uncertainty labels with a neural network in a dual regression setting. Mixed results do not mean that uncertainty labels should be excluded from future work. Since Cramér-Rao bounds are related to the information content, performing regression for parameter uncertainties should be theoretically possible. The local nature of Cramér-Rao bound uncertainty means that a more finely sampled representation of the data in time might be helpful for a neural network to learn to predict uncertainties. Indeed, when the Cramér-Rao bound is calculated using finite differencing with Richardson extrapolation, transmission loss values of several points as the parameter value varies in the vicinity

of the point of interest are taken into account. Given that true ocean seabeds have variation in parameters, increasing the time sampling of the spectrogram could give a neural network the boost it needs to more accurately predict both parameter values and parameter uncertainties. A neural network may need access to parameter information contained in the points around the point of interest for a neural network to accurately perform regression for a local uncertainty value.

Following along with the theme of the relationship between local and global properties is the nature of data labels. Many data labels are often treated as global descriptors, as with the geoacoustic parameter values—a single value that has a relationship to the entire data sample. Other data labels are local. An image that contains a cat might be labeled as a cat despite the label only applying directly to a portion of the image. This realization prompts two potential paths for labeling spectrograms with uncertainty data. First, a global uncertainty value might be applied to an entire spectrogram. If there were a way to boil down all the local uncertainties associated with a spectrogram to a single, global uncertainty, the global uncertainty would likely make a good label for a neural network to learn. Alternatively, the second potential path would be to use a local uncertainty label that applies only to a specific region of a spectrogram—the Cramér-Rao bound at 500 Hz, for example. The model could then learn to extract that specific, local piece of information from the spectrogram image rather than having to worry about extracting the local uncertainties from all parts of the spectrogram to take a shot at estimating the uncertainty, which was likely the primary reason the network in this work struggled to estimate uncertainty labels.

If anything, this work should highlight the importance of taking the time to understand and explore the information content characteristics of both simulated and measured data with respect to a physical model. Uncertainty measures should be carefully considered. The great lesson of the uncertainty labeling and regression in this study is the corollary to “Garbage in, garbage out”, namely, “*Ad hoc* in, *ad hoc* out”. A more principled, systematic development of uncertainty labels

could lead to improved performance in dual parameter and parameter uncertainty regression with neural networks.

Bibliography

[1] K. D. Heaney, "Rapid geoacoustic characterization using a surface ship of opportunity," IEEE Journal of Oceanic Engineering **29**(1), 88–99 (2004).

[2] F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt, *Computational Ocean Acoustics* (Springer New York, New York, NY, 2011).

[3] M. S. Ballard, K. M. Lee, A. R. McNeese, P. S. Wilson, J. D. Chaytor, J. A. Goff, and A. H. Reed, "In situ measurements of compressional wave speed during gravity coring operations in the new england mud patch," IEEE Journal of Oceanic Engineering **45**, 26–38 (2020) doi: 10.1109/JOE.2019.2924560.

[4] J. Belcourt, C. W. Holland, S. E. Dosso, J. Dettmer, and J. A. Goff, "Depth-dependent geoacoustic inferences with dispersion at the new England mud patch via reflection coefficient inversion," IEEE Journal of Oceanic Engineering **45**(1), 69–91 (2020) doi: 10.1109/JOE.2019.2900115.

[5] A. Tarantola, *Inverse problem theory and methods for model parameter estimation* (Society for Industrial and Applied Mathematics, Philadelphia, 2005).

[6] E. K. Westwood, C. T. Tindle, and N. R. Chapman, "A normal mode model for acousto-elastic ocean environments," The Journal of the Acoustical Society of America **100**(6), 3631–3645 (1996) doi: 10.1121/1.417226.

[7] E. K. Westwood and R. A. Koch, "Elimination of branch cuts from the normal-mode solution using gradient half spaces," The Journal of the Acoustical Society of America **106**(5), 2513–2523 (1999) doi: 10.1121/1.428083.

[8] M. A. Nielsen, *Neural Networks and Deep Learning* (Determination Press, 2015).

[9] D. F. Van Komen, T. B. Neilsen, D. B. Mortenson, M. C. Acree, D. P. Knobles, M. Badiey, and W. S. Hodgkiss, "Seabed type and source parameters predictions using ship spectrograms in convolutional neural networks," The Journal of the Acoustical Society of America **149**(2), 1198–1210 (2021) doi: 10.1121/10.0003502.

[10] C. D. Escobar-amado, T. B. Neilsen, J. A. Castro-correa, and D. F. V. Komen, "ensemble of deep learning algorithms seabed classification from merchant ship-radiated noise using a physics-based ensemble of deep learning algorithms," **1434** (2021) doi: 10.1121/10.0005936.

[11] B. B. Machta, R. Chachra, M. K. Transtrum, and J. P. Sethna, "Parameter Space Compression Underlies Emergent Theories and Predictive Models," Science **342**(November), 604–607 (2013).

[12] K. S. Brown, C. C. Hill, G. A. Calero, C. R. Myers, K. H. Lee, J. P. Sethna, and R. A. Cerione, "The statistical mechanics of complex signaling networks : nerve growth factor signaling," Physical Biology **1**, 184–195 (2004) doi: 10.1088/1478-3967/1/3/006.

[13] K. S. Brown and J. P. Sethna, "Statistical mechanical approaches to models with many poorly known parameters," Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics **68**(2), 9 (2003) doi: 10.1103/PhysRevE.68.021904.

[14] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna, "Universally Sloppy Parameter Sensitivities in Systems Biology Models," PLoS Computational Biology **3**(10), 1871–1878 (2007) doi: 10.1371/journal.pcbi.0030189.

[15] J. J. Waterfall, F. P. Casey, R. N. Gutenkunst, K. S. Brown, C. R. Myers, P. W. Brouwer, V. Elser, and J. P. Sethna, "Sloppy-Model Universality Class and the Vandermonde Matrix," Physical Review Letters **150601**, 1–4 (2006) doi: 10.1103/PhysRevLett.97.150601.

[16] G. J. Berman and Z. J. Wang, "Energy-minimizing kinematics in hovering insect flight," Journal of Fluid Mechanics **582**, 153–168 (2007) doi: 10.1017/S0022112007006209.

[17] T. D. Sanger, "Human arm movements described by a low-dimensional superposition of principal components," Journal of Neuroscience **20**(3), 1066–1072 (2000) doi: 10.1523/jneurosci.20-03-01066.2000.

[18] M. D. Collins and W. A. Kuperman, "Inverse problems in ocean acoustics," Inverse Problems **10**(5), 1023–1040 (1994) doi: 10.1088/0266-5611/10/5/003.

[19] M. D. Collins and L. Fishman, "Efficient navigation of parameter landscapes," Journal of the Acoustical Society of America **98**(3), 1637–1644 (1995) doi: 10.1121/1.413430.

[20] H. C. Song, "Performance Bounds for Passively Locating a Moving Source of a Known Frequency in Oceanic Waveguide Using a Vertical Array," IEEE Journal of Oceanic Engineering **18**(3), 189–198 (1993) doi: 10.1109/JOE.1993.236357.

[21] S. C. Walker, C. Yardim, A. Thode, and E. Arias-Castro, "Using Fisher information to quantify uncertainty in environmental parameters estimated from correlated ambient noise," The Journal of the Acoustical Society of America **133**(4), EL228–EL234 (2013) doi: 10.1121/1.4792836.

[22] A. Baggeroer and H. Schmidt, "Cramer-Rao bounds for matched field tomography and ocean acoustic tomography," in *1995 International Conference on Acoustics, Speech, and Signal Processing*, IEEE (1995), Vol. 5, pp. 2763–2766, doi: 10.1109/ICASSP.1995.479417.

[23] A. B. Baggeroer, "The stochastic Cramér-Rao bound for source localization and medium tomography using vector sensors," **5**(141) (2017) doi: 10.1121/1.4981398.

[24] L. Wan, M. Badiey, D. P. Knobles, P. S. Wilson, and J. A. Goff, "Estimates of low-frequency sound speed and attenuation in a surface mud layer using low-order modes," IEEE Journal of Oceanic Engineering **45**(1), 201–211 (2020) doi: 10.1109/JOE.2019.2923861.

[25] E. M. Brown, Y. T. Lin, J. D. Chaytor, and W. L. Siegmann, "Geoacoustic inversion for a new England mud patch sediment using the silt-suspension theory of marine mud," IEEE Journal of Oceanic Engineering **45**(1), 144–160 (2020) doi: 10.1109/JOE.2019.2934604.

[26] J. Bonnel, S. E. Dosso, D. Eleftherakis, and N. R. Chapman, "Trans-dimensional inversion of modal dispersion data on the new England mud patch," IEEE Journal of Oceanic Engineering **45**(1), 116–130 (2020) doi: 10.1109/JOE.2019.2896389.

[27] D. Tollefsen and S. E. Dosso, "Ship source level estimation and uncertainty quantification in shallow water via Bayesian marginalization," The Journal of the Acoustical Society of America **147**(4), EL339–EL344 (2020) doi: 10.1121/10.0001096.

[28] J. Dettmer, S. E. Dosso, and C. W. Holland, "Trans-dimensional geoacoustic inversion," The Journal of the Acoustical Society of America **128**(6), 3393–3405 (2010) doi: 10.1121/1.3500674.

[29] P. Gerstoft and C. F. Mecklenbräuker, " Ocean acoustic inversion with estimation of a posteriori probability distributions ," The Journal of the Acoustical Society of America **104**(2), 808–819 (1998) doi: 10.1121/1.423355.

[30] P. Gerstoft, "Inversion of acoustic data using a combination of genetic algorithms and the Gauss-Newton approach," Journal of the Acoustical Society of America **97**(4), 2181–2190 (1995) doi: 10.1121/1.411943.

[31] R. T. Kessel, "The variation of modal wave numbers with geoacoustic parameters in layered media," The Journal of the Acoustical Society of America **102**(5), 2690–2696 (1997) doi: 10.1121/1.420322.

[32] J.-P. Hermand, M. Meyer, M. Asch, and M. Berrada, "Adjoint-based acoustic inversion for the physical characterization of a shallow water environment," The Journal of the Acoustical Society of America **119**(6), 3860–3871 (2006) doi: 10.1121/1.2197790.

[33] M. K. Transtrum, B. B. Machta, and J. P. Sethna, "Geometry of nonlinear least squares with applications to sloppy models and optimization," Physical Review E **036701**(March 2011), 1–35 (2011) doi: 10.1103/PhysRevE.83.036701.

[34] M. K. Transtrum and P. Qiu, "Model Reduction by Manifold Boundaries," Physical Review Letters **113**(9), 098701 (2014) doi: 10.1103/PhysRevLett.113.098701.

[35] T. B. Neilsen, "An iterative implementation of rotated coordinates for inverse problems," The Journal of the Acoustical Society of America **113**(5), 2574–2586 (2003) doi: 10.1121/1.1562912.

[36] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3 ed. (Cambridge University Press, 2007).

[37] M. K. Transtrum, B. B. MacHta, and J. P. Sethna, "Why are nonlinear fits to data so challenging?," Physical Review Letters **104**(6), 2–5 (2010) doi: 10.1103/PhysRevLett.104.060201.

[38] L. F. Richardson and J. A. Gaunt, "VIII. The deferred approach to the limit," Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character **226**(636-646), 299–361 (1927) doi: 10.1098/rsta.1927.0008.

[39] J. Perktold and A. Verheyleweghen, "numdifftools," (2021).

[40] G. R. Potty and J. H. Miller, "Effect of shear on modal arrival times," IEEE Journal of Oceanic Engineering **45**(1), 103–115 (2020) doi: 10.1109/JOE.2019.2925920.

[41] F. A. Bowles, "Observations on attenuation and shear-wave velocity in fine-grained, marine sediments," The Journal of the Acoustical Society of America **101**(6), 3385–3397 (1997) doi: 10.1121/1.419374.

[42] R. Yousefzadeh and X. Cao, "To what extent should we trust ai models when they extrapolate?," **1**, 1–17 (2022).

[43] A. Madhu and S. Kumaraswamy, "Data augmentation using generative adversarial network for environmental sound classification," European Signal Processing Conference, EUSIPCO (2019), Vol. 2019-September, doi: 10.23919/EUSIPCO.2019.8902819.

[44] D. Harper, "Online etymology dictionary," https://www.etymonline.com/.

[45] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, and S. Nahavandi, "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," (2020) http://arxiv.org/abs/2011.06225http://dx.doi.org/10.1016/j.inffus.2021.05.008 doi: 10.1016/j.inffus.2021.05.008.

[46] P. S. Wilson, D. P. Knobles, and T. B. Neilsen, "Guest editorial an overview of the seabed characterization experiment," IEEE Journal of Oceanic Engineering **45**(1), 1–13 (2020).

[47] B. of Ocean Energy Management (BOEM), N. Oceanic, and A. A. (NOAA), "Marine cadastre (ais 2017 03 zone19 and ais 2017 04 zone19)," https://marinecadastre.gov/ais/.

[48] C. Gervaise, B. G. Kinda, J. Bonnel, Y. Stéphan, and S. Vallez, "Passive geoacoustic inversion with a single hydrophone using broadband ship noise," J. Acoust. Soc. Am. **131**(3), 1999–2010 (2012).

[49] L. Liu, S. Zhao, H. Chen, and A. Wang, "A new machine learning method for identifying alzheimer's disease," Simulation Modelling Practice and Theory **99** (2020) doi: 10.1016/j.simpat.2019.102023.

[50] P. Mekha, N. Teeyasuksaet, T. Sompowloy, and K. Osathanunkul, "Honey bee sound classification using spectrogram image features," 205–209 (2022).

[51] R. Tibi, L. Linville, C. Young, and R. Brogan, "Classification of local seismic events in the utah region: A comparison of amplitude ratio methods with a spectrogram-based machine learning approach," Bulletin of the Seismological Society of America **109**, 2532–2544 (2019) doi: 10.1785/0120190150.

[52] H. Duysak, U. Ozkaya, and E. Yigit, "Determination of the amount of grain in silos with deep learning methods based on radar spectrogram data," IEEE Transactions on Instrumentation and Measurement **70** (2021) doi: 10.1109/TIM.2021.3085939.

[53] C. Jin, S. Park, and S. Rhee, "Spectrogram based detection algorithm for back-bead in gas metal arc welding process using convolution neural network," Journal of Welding and Joining **39**, 198–205 (2021) doi: 10.5781/jwj.2021.39.2.9.

[54] O. Barkan, D. Tsiris, O. Katz, and N. Koenigstein, "Inversynth: Deep estimation of synthesizer parameter configurations from audio signals," (2018) http://arxiv.org/abs/1812.06349.

[55] H. Qin, "Machine learning and serving of discrete field theories," Scientific Reports **10**(1), 1–15 (2020) doi: 10.1038/s41598-020-76301-0.

[56] X. Ma, Z. Wu, J. Jia, M. Xu, H. Meng, and L. Cai, "Emotion recognition from variable-length speech segments using deep learning on spectrograms," Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH **2018-September**, 3683–3687 (2018) doi: 10.21437/Interspeech.2018-2228.

[57] J. D. Alvares, J. A. Font, F. F. Freitas, O. G. Freitas, A. P. Morais, S. Nunes, A. Onofre, and A. Torres-Forne, "Gravitational-wave parameter inference using deep learning," Proceedings -

International Workshop on Content-Based Multimedia Indexing **2021-June** (2021) doi: 10.1109/CBMI50038.2021.9461893.

[58] T. Berthold, A. Leichter, B. Rosenhahn, V. Berkhahn, and J. Valerius, "Seabed sediment classification of side-scan sonar data using convolutional neural networks," 2017 IEEE Symposium Series on Computational Intelligence, SSCI 2017 - Proceedings **2018-Janua**, 1–8 (2018) doi: 10.1109/SSCI.2017.8285220.

[59] D. F. Van Komen, T. B. Neilsen, D. P. Knobles, and M. Badiey, "A convolutional neural network for source range and ocean seabed classification using pressure time-series," Proc. Meet. Acoust. **36**(1), 070004 (2019).

[60] D. F. Van Komen, T. B. Neilsen, K. Howarth, D. P. Knobles, and P. H. Dahl, "Seabed and range estimation of impulsive time series using a convolutional neural network," J. Acoust. Soc. Am. **147**(5), EL403–EL408 (2020).

[61] C. Frederick, S. Villar, and Z. H. Michalopoulou, "Seabed Classification using Physics-based Modeling and Machine Learning," arXiv **859** (2020) doi: 10.1121/10.0001728.

[62] H. Niu, E. Ozanich, and P. Gerstoft, "Ship localization in santa barbara channel using machine learning classifiers," J. Acoust. Soc. Am. **142**(5), EL455–EL460 (2017).

[63] H. Niu, E. Reeves, and P. Gerstoft, "Source localization in an ocean waveguide using supervised machine learning," J. Acous. Soc. Am. **142**(3), 1176–1188 (2017).

[64] H. Niu, Z. Gong, E. Ozanich, P. Gerstoft, H. Wang, and Z. Li, "Deep-learning source localization using multi-frequency magnitude-only data," arXiv **211** (2019) doi: 10.1121/1.5116016.

[65] Z. Huang, J. Xu, Z. Gong, H. Wang, and Y. Yan, "Source localization using deep neural networks in a shallow water environment," J. Acoust. Soc. Am. **143**(5), 2922–2932 (2018).

[66] E. Ozanich, P. Gerstoft, and H. Niu, "A feedforward neural network for direction-of-arrival estimation," J. Acoust. Soc. Am. **147**(3), 2035–2048 (2020).

[67] J. Benson, N. R. Chapman, and A. Antoniou, "Geoacoustic model inversion with artificial neural networks," 1998 IEEE Symposium on Advances in Digital Filtering and Signal Processing, DFSP 1998 121–125 (1998) doi: 10.1109/ADFSP.1998.685708.

[68] J. Benson, N. R. Chapman, and A. Antoniou, "Geoacoustic model inversion using artificial neural networks," Inverse Problems **16**(6), 1627 (2000).

[69] A. Caiti and S. M. Jesus, "Acoustic estimation of seafloor parameters: A radial basis functions approach," The Journal of the Acoustical Society of America **100**, 1473–1481 (1996) doi: 10.1121/1.415994.

[70] J. Piccolo, G. Haramuniz, and Z.-H. Michalopoulou, "Geoacoustic inversion with generalized additive models," J. Acoust. Soc. Am. **145**(6), EL463–EL468 (2019).

[71] Y. Shen, X. Pan, Z. Zheng, and P. Gerstoft, "Matched-field geoacoustic inversion based on radial basis function neural network," The Journal of the Acoustical Society of America **148**, 3279–3290 (2020) doi: 10.1121/10.0002656.

[72] L. Mao, X. Pan, and Y. Shen, "Geoacoustic inversion based on neural network," 1–5 (2022) doi: 10.23919/oceans44145.2021.9705922.

[73] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," doi: 10.48550, original paper where ResNets were introduced, Reference for ResNet18.

[74] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," 7th International Conference on Learning Representations, ICLR 2019 (2019).