

Sub-Bottom Profiling Using an Autonomous Underwater Vehicle Equipped With a Sound Source and Towed Hydrophone Array

by

Paige Pfenninger

B.S., F.W. Olin College of Engineering, 2019

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY AND
WOODS HOLE OCEANOGRAPHIC INSTITUTION

February 2025

© 2025 Paige Pfenninger. All rights reserved.

The author hereby grants to MIT and WHOI a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

- Authored by: Paige Pfenninger
Department of Electrical Engineering and Computer Science, MIT
Applied Ocean Science and Engineering, WHOI
January 19, 2024
- Certified by: Ying-Tsong Lin
Professor and Victor C. Alderson Chair of Applied Ocean Science
Scripps Institution of Oceanography, UC San Diego
Thesis Supervisor
- Certified by: John Leonard
Samuel C. Collins Professor of Mechanical and Ocean Engineering, MIT
Thesis Supervisor
- Accepted by: Leslie A. Kolodziejcki
Professor of Electrical Engineering and Computer Science, MIT
Chair, Department Committee on Graduate Students
- Accepted by: Alexandra H. Techet
Professor of Mechanical & Ocean Engineering, MIT
Chair, Joint Committee of Applied Ocean Science and Engineering

Sub-Bottom Profiling Using an Autonomous Underwater Vehicle Equipped With a Sound Source and Towed Hydrophone Array

by

Paige Pfenninger

Submitted to the Department of Electrical Engineering and Computer Science
on January 19, 2024 in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

ABSTRACT

Sub-bottom profiling using an autonomous underwater vehicle equipped with a source and a towed array is an excellent method to finely survey large areas of the ocean bottom with minimal interference from the water column. This approach has the benefit of being able to determine the range dependence of the sub-bottom on a meter-by-meter scale rather than assuming constant sub-bottom properties over a large range. This thesis conducts theoretical and experimental studies to investigate the feasibility of using the arrival times of acoustic signals from an autonomous underwater vehicle source to a short, 16-element towed hydrophone array to determine the sound speed and layer thickness of the seabed through Bayesian geoacoustic inversion. This method provides range-dependent geoacoustic parameters with a resolution on the order of 10 meters. Numerical studies indicate that, for timing data with low variance, arrival times can be used to accurately estimate seabed properties. However, the performance of the Bayesian inversion model deteriorates as the variance of the timing data increases.

Experimental data were collected during the Seabed Characterization Experiment at the New England Mud Patch and the New England Shelf Break. This thesis attempts to improve the arrival times through the use of sub-array focusing but concludes that this method is not feasible due to the experimental data exhibiting a high level of variance in the sub-bottom timing returns, likely due to the presence of scatterers in the sediment layer. Therefore, the mean and variance of the direct path, bottom, and sub-bottom timing returns were calculated using Gaussian process regression. Furthermore, the results show that layer thickness and sound speeds are highly coupled, making it challenging to uniquely determine seabed properties.

Thesis supervisor: Ying-Tsong Lin

Title: Professor and Victor C. Alderson Chair of Applied Ocean Science
Scripps Institution of Oceanography, UC San Diego

Thesis supervisor: John Leonard

Title: Samuel C. Collins Professor of Mechanical and Ocean Engineering, MIT

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my advisors, YT Lin and John Leonard. I appreciate your mentorship throughout my first few years in graduate school and I thank you for your unwavering support and guidance throughout my studies. I look forward to your continued support as I pursue my PhD.

I would like to thank all of my lab mates in both APPG and MRG for welcoming me into your groups and providing great collaboration and encouragement. I would especially like to thank Will, Natalie, Brendan, and Billy for your support during research cruises. I would also like to thank Alan, JP, Amy, Ariel, Caroline, Camille, Kurran, and Jungseok for welcoming me into the MRG lab group and creating a supportive and enjoyable lab environment. A special thanks goes to Alan, whose invaluable help and unwavering support have been instrumental in completing this thesis. Your generosity with your time and knowledge has meant so much to me.

I am also thankful for my friends and family, who have been my constant sources of love, encouragement, and strength. Your belief in me has sustained me through the challenges of this journey.

Finally, I would like to thank my cat, Milo, for just being him.

This thesis would not have been possible without each and every one of you. Thank you.

This work was supported by funding from the Office of Naval Research Grant N00014-22-1-2013.

Contents

<i>List of Figures</i>	9
<i>List of Tables</i>	11
1 Introduction	13
1.1 Experimental Overview	15
1.1.1 SBCEX2017 and SBCEX2022	15
1.1.2 SBCEX AUV Setup	16
1.2 Problem Formulation	19
1.3 Summary	25
2 Literature Review	27
2.1 AUV Towed Array Processing	27
2.1.1 AUV Towed Array Beamforming	28
2.1.2 AUV Towed Array Synthetic Aperture Sonar	28
2.1.3 AUV Towed Array Near Field Beamforming	29
2.2 Geoacoustic Inversion	29
2.2.1 Bayesian Geoacoustic Inversion	30
2.2.2 Timing Based Geoacoustic Inversion	30
2.3 Uncertainty Estimation	31
3 Arrival Time Estimates with Sub-Array Beamforming	33
3.1 Linear Array Plane Wave Beamforming	34
3.2 Linear Array Approximation for the towed array	38
3.2.1 Near Field vs Far Field	41
3.3 Sub-Array Beamforming	42
3.4 Near Field Beamforming	43
3.5 Sub-Array Focusing	44
3.6 Timing Returns from Beamforming	46
3.7 Summary	50
4 Gaussian Process Regression	51
4.1 Estimating Arrival Times and Variance	51
4.2 Multi-Task Gaussian Process Regression	55
4.3 Range Independent Length	58
4.4 Summary	60

5	Geoacoustic Inversion	61
5.1	Bayesian Geoacoustic Inversion	61
5.2	Experimental Results	68
5.3	Summary	74
6	Conclusion	75
6.1	Thesis Contributions	76
6.1.1	Sub-Array Focusing	76
6.1.2	Scattering Estimation using Gaussian Process Regression	76
6.1.3	Timing Based Geoacoustic Inversion	77
6.2	Future Work	77
6.2.1	Multiple Sub-Bottom Layer Inversion	77
6.2.2	Alternative Bottom Sound Speed Profiles	78
6.2.3	Evaluating the Towed Array Geometry	78
6.3	Final Thoughts	79
A	Additional Signal Processing	81
A.1	Pulse Compression	81
A.2	Peak Detection	82
A.3	Sampling Rate Calibration	83
A.4	Hydrophone Distance Calibration	84
A.5	Rough Layer Detection	84
B	Bottom Interface Angle Derivation	87
	<i>References</i>	89

List of Figures

1.1	REMUS 600 and towed array used in both experiments. The towed array is 11.1 meters long and contains 16 elements.	17
1.2	Map of the New England Mud Patch experiment area with the AUV mission path shown	18
1.3	Altitude and depth measurements of the AUV mid-mission for the mud patch where the AUV was maintaining a constant altitude.	19
1.4	Spectrogram of AUV source during the mud patch survey	20
1.5	Map of the shelf break experiment area with the AUV mission path shown	20
1.6	Altitude and depth measurements of the AUV mid-mission for the shelf break experiment	21
1.7	A depiction of the AUV with the direct, bottom bounce, and sub-bottom bounce travel paths on two different towed array elements	22
3.1	Beam pattern of the array in the SBCEX22 experiment for different look directions. The look direction is shown in blue.	39
3.2	The hydrophone positions during a straight track (left) and during a turn (right)	40
3.3	Comparison between single element returns (left) and the result of sub-array beamforming (right) on the sub-bottom layer using data from the shelf break experiment	43
3.4	An illustration of sub-bottom as viewed by different sub-arrays during the AUV track	45
3.5	The sub-bottom processed with sub-array synthetic aperture beamforming compared to a single sub-array beamforming using data from the shelf break experiment	46
3.6	The timing returns calculated for the sub-bottom layer for four different sub-arrays using data from the mud patch experiment	47
3.7	The timing returns from all sub-arrays for the bottom (left) and sub-bottom (right). This data came from the mud patch experiment.	48
3.8	The amplitude and phase of approximately 100 waveforms on one channel with the direct path, bottom returns, and sub-bottom returns marked. This data came from the mud patch experiment.	49
4.1	The sub-bottom return on 16 channels with the estimates of the timing of the sub-bottom returns shown in red.	52
4.2	A histogram of the arrival times of the bottom and sub-bottom layers	53

4.3	The results of the multi-task Gaussian process regression on the timing data for the three different arrival paths. 40 pulses were used as the training data for the regression.	57
4.4	The results of the multi-task Gaussian process regression with the raw waveform indicating an accurate estimation of the mean and variance of the timing returns. The confidence interval is shown in green and the mean is shown in magenta.	58
4.5	The timing returns from the sub-bottom Gaussian process regression for four channels. The timing return data with segments from 20 to 50 pulses to estimate the mean arrival time	59
5.1	The Bayesian inversion results for test data with low sub-bottom variance in the timing data. The top figure shows the modeled timing returns for all three layers in red and the test timing returns in black. The bottom figures show the posterior distribution for the different parameters. The most likely estimate for each parameter is marked in red and the true value is shown in green. . .	64
5.2	The percent error between the most likely value of each parameter and the true parameter values as a function of the sub-bottom variance. The figure on the left shows all five parameters and the figure on the right is focused on the layer thickness and sound speed parameters. The shaded region indicates typical sub-bottom variance levels from the GP Regression	65
5.3	The posterior probabilities for each parameter of the test data as a function of the sub-bottom variance. The red line indicates the most likely value of that parameter and the black line indicates the true value. The white vertical line indicates the median sub-bottom variance of 20 minutes of mud patch data calculated using GP Regression	67
5.4	The Bayesian inversion results for experimental data. The top figure shows the modeled timing returns of the most likely estimated model for all three layers in red and the actual timing returns in black. The bottom figures show the posterior distribution for the different parameters. The most likely estimate for each parameter is marked in red.	69
5.5	The joint probability distributions between the parameters in the geoacoustic inversion.	70
5.6	The posterior probability distributions as a function of distance traveled for the five parameters in the geoacoustic inversion. The red line shows the most likely value at each distance	71
5.7	The posterior probability distributions as a function of distance traveled for the five parameters in the geoacoustic inversion. A Gaussian prior was placed on h_1 centered around 10 m with a standard deviation of 1 m. The red line shows the most likely value at each distance.	73
A.1	The shelf break data before and after pulse compression	82
A.2	The calibrated hydrophone positions over time with histograms showing the variation in position.	85

List of Tables

5.1 Results of other geoaoustic inversion results in the SBCEX2017 Mud Patch Experiment	72
---------------------------------------------------------------------------------------------------	----

Chapter 1

Introduction

The ocean covers more than 70% of the planet's surface, yet much about the ocean is still unknown. The seafloor plays a critical role in the ocean environment, and like the ocean, much of the seafloor is still poorly understood. Detailed mapping and study of the seafloor is essential for identifying habitats, understanding biodiversity, understanding geological processes, and uncovering resources. High-resolution mapping technologies, such as multi-beam sonar systems, are gradually improving knowledge of the seafloor. Still, as of 2024, only 26.1% of the ocean floor has been surveyed with this technology [1].

Due to the rapid attenuation of light and other electromagnetic waves in the ocean, acoustic waves are required to study the ocean and communicate underwater. Sound travels much further and more rapidly underwater than in the air. Acoustics are used underwater in numerous tasks such as navigation, object detection, wireless communication, and oceanographic observation. The propagation of sound in the ocean depends heavily on the ocean environment, and understanding that environment is critical to predicting how sound propagates through the ocean. When underwater sounds propagate at steep angles or in shallow water, the bottom interactions significantly impact sound propagation.

The ocean bottom is not just made up of one uniform layer. The bottom is commonly composed of many layers of sediment such as mud, silt, and limestone. Each of these

sediments has different geoacoustic properties and multiple sediment layers can impact sound propagation. Understanding the different sediment layers is crucial to understanding the oceans.

Studying the sea-bottom, however, is a challenging endeavor. High frequencies do not penetrate into the sub-bottom [2], and unknown water column interactions can make determining the seabed properties difficult. Autonomous underwater vehicles (AUV) are ideal for studying the ocean floor because they are capable of carrying sources that emit low to mid-range frequencies that can propagate into the sub-bottom, they can operate close to the ocean bottom to minimize water column interactions, and they can tow arrays of hydrophones that are capable of recording a wide range of grazing angles from the source. AUVs can also survey large regions of the seafloor to finely sample the sub-bottom properties.

This thesis aims to answer the following questions:

- How can the timing data from the sub-bottom be accurately estimated in a high-noise environment?
- How can timing data from different travel paths onto a towed array be used to determine the sub-bottom properties?
- How can scattering in the bottom and sub-bottom layers be quantified and included in the geoacoustic inversion scheme?

This thesis presents a method for using timing returns to estimate sub-bottom properties, discusses the use of sub-array beamforming to enhance the sub-bottom returns, shows the use of Gaussian process regression to estimate the variance of the sub-bottom, and implements a timing-based Bayesian geoacoustic inversion to estimate the bottom properties while including the variance seen in the sub-bottom returns.

1.1 Experimental Overview

1.1.1 SBCEX2017 and SBCEX2022

In 2017 and 2022, a series of seabed characterization experiments (SBCEX) was conducted at the New England Mud Patch (NEMP) and the New England Shelf Break off the coast of Massachusetts to study the geoacoustic properties of the region. Two of the primary objectives of this experiment were to characterize horizontal variations in the seabed and to quantify the uncertainties in the seabed parameter estimations over a broad range of frequencies. This region is of particular interest because the fine-grain sediments found here do not match the majority of the shelf sediments [3].

Before the 2017 and 2022 experiments began, a detailed sub-bottom survey was completed in 2015 and 2016 to measure the thickness of the sub-bottom and 53 core samples were taken to provide an understanding of the physical properties of the sediment layer. This survey found that the mud was up to 12 *m* thick in certain areas and showed that the mud was acoustically transparent. The results from this survey informed the locations of the later experiments [4].

The first portion of the main seabed characterization experiment was conducted at the mud patch in 2017 (SBCEX17). This portion of the experiment included 200 core samples and a variety of acoustic sources and measurement devices to quantify the acoustic properties of the seabed. The cores showed the NEMP sediment does not consist of pure mud, but contains silt, clay, and sand. High-frequency acoustic measurements (25 – 200 *kHz*) were also taken during the coring process to understand the sediment sound speed at those frequencies. Unfortunately, these sound speed measurements cannot be considered ground truth for lower frequency ranges due to the frequency dependence of the sediment sound speed. The experiment was conducted during the winter when the ocean temperature is more homogeneous to reduce the water column’s impact on sound propagation [5].

A follow-up experiment was conducted in the same region in 2022 (SBCEX22). The 2022 experiments included 36 additional gravity cores that were collected at both the NEMP and the shelf break [6]. The New England Shelf Break is geographically close to the NEMP. The shelf break has rapidly varying bottom topography with multiple distinct bottom layers. As part of these experiments, a sequence of AUV missions was conducted on the mud patch to study the spatial variability of the ocean bottom. A portion of the 2022 experiment also included an AUV sub-bottom survey of the New England Shelf Break. It is these AUV experiments that will be the focus of this thesis [7].

1.1.2 SBCEX AUV Setup

For the SBCEX22 experiments, a Remote Environmental Monitoring UnitS (REMUS) AUV 600 was equipped with a 3.5 kHz transducer, a conductivity, temperature, depth (CTD) sensor, an acoustic Doppler current profiler (ADCP) that measured altitude, latitude, and longitude, and a 16 hydrophone towed array. The AUV had bottom lock for the entire mission aside from the descent and ascent which means the estimated odometry should be accurate [8].

The first element of the towed array was located 20.77 m from the transducer, and the mean element spacing was 0.74 m . This gives a total array length of 11.1 m with the furthest element being 31.87 m from the source. Figure 1.1 shows the AUV and towed array used in this experiment.

For both missions, the AUV path was set as a lawn mower pattern as seen in Figure 1.2 and 1.5. The analysis portion of this thesis focuses on one of the long middle tracks where the AUV is operating in a straight line for an extended period.

New England Mud Patch

Figure 1.2 shows the mud patch experiment area with the mission path. The AUV was programmed to run at a constant altitude. Figure 1.3 shows the altitude and depth as



Figure 1.1: REMUS 600 and towed array used in both experiments. The towed array is 11.1 meters long and contains 16 elements.

measured by the AUV during one of the middle tracks for the mud patch survey. The bathymetry of the mud patch is slightly sloped. Over 10 minutes, the depth changes from 70.8m to 71.1m. During these 10 minutes, the AUV maintains a stable pitch of approximately -1.5° to maintain a constant altitude. Throughout the mud patch mission, the AUV maintained a mean altitude of 5.02 meters. These dimensions give a grazing angle range of the bottom reflection from $\arctan \frac{2 \cdot 5.02}{31.87} = 17.8^\circ$ to $\arctan \frac{2 \cdot 5.02}{20.77} = 25.8^\circ$ on the first and last elements of the towed array.

The transducer was programmed to emit a 50 ms chirp from 2.75 kHz to 4.25 kHz every second for nine seconds and then take a one-second pause. Figure 1.4 shows a spectrogram of the transducer as recorded on a single hydrophone on the array along with the power of the received signal in the frequency domain. The majority of the noise on the receiver is in the




Figure 1.2: Map of the New England Mud Patch experiment area with the AUV mission path shown

lower frequencies. This noise is created by AUV operation, but should not impede signal processing due to the frequency separation between the noise band and the signal band.

New England Shelf Break

The shelf break survey region with the AUV path is shown in Figure 1.5. As seen in Figure 1.6, the AUV was operating at a constant depth and the altitude changed from a minimum of 27.85 m to a maximum of 49.93 m . At the minimum altitude, the grazing angle range is from $\arctan \frac{2-27.8}{31.87} = 60.18^\circ$ to $\arctan \frac{2-27.8}{20.77} = 69.52^\circ$. At the maximum altitude, the grazing angle ranges from $\arctan \frac{2-49.8}{31.87} = 72.29^\circ$ to $\arctan \frac{2-49.8}{20.77} = 78.24^\circ$. During the shelf break survey, the transducer was programmed to emit a 100 ms chirp from 2.5 kHz to 4.5 kHz and also operated every second for nine seconds with a one-second pause.

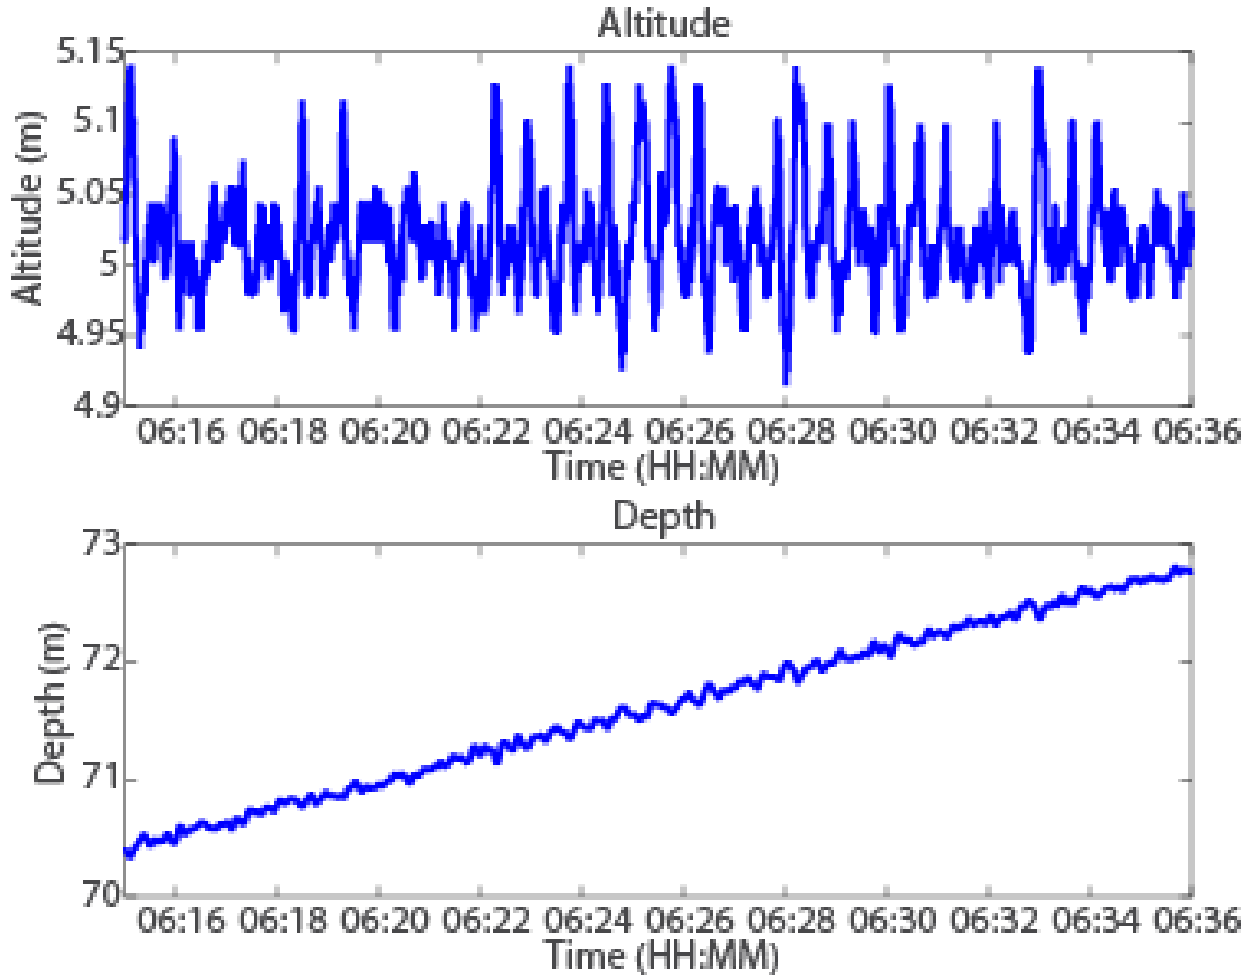


Figure 1.3: Altitude and depth measurements of the AUV mid-mission for the mud patch where the AUV was maintaining a constant altitude.

1.2 Problem Formulation

Each pulse emitted by the AUV source gives three different returns on the array. Figure 1.7 shows the array with the three different travel paths. The end-fire return travels the direct path between the source and the array, the bottom return reflects off the bottom before hitting the array, and the sub-bottom return transmits through the bottom and then reflects off of the sub-bottom before reaching the array. Each path provides different information about the environment that is essential for understanding the sub-bottom properties. Because the sub-bottom path travels through both the water column and the bottom layer, information

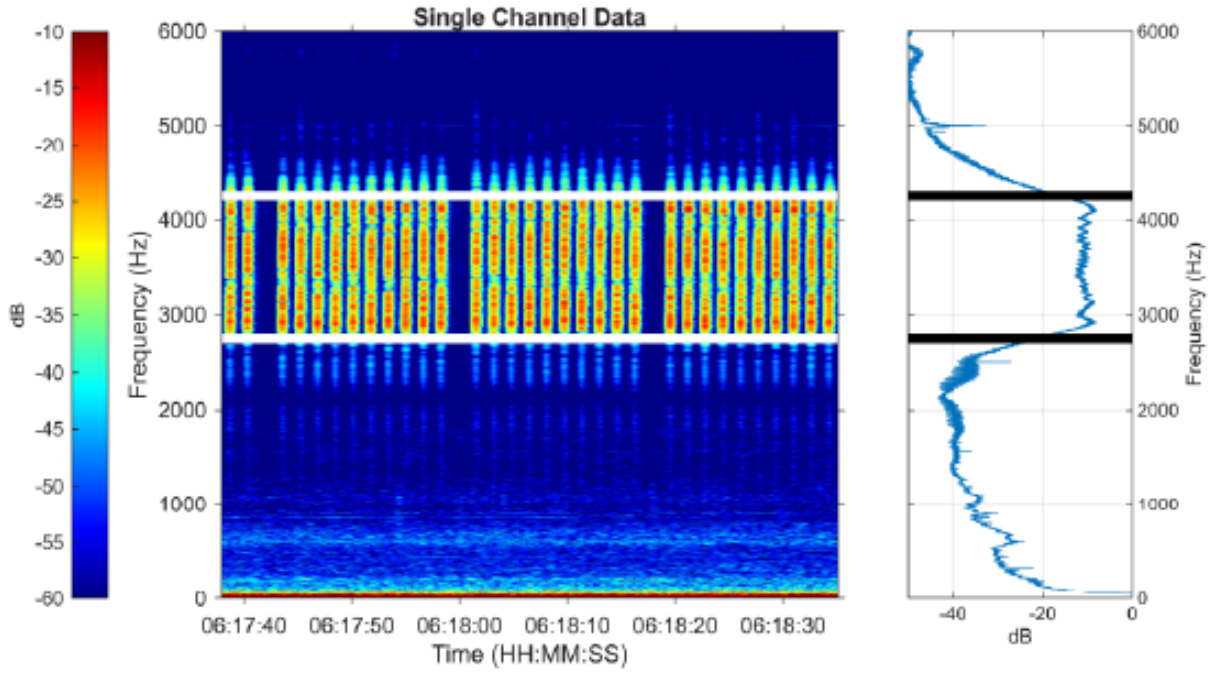


Figure 1.4: Spectrogram of AUV source during the mud patch survey




Figure 1.5: Map of the shelf break experiment area with the AUV mission path shown

about both the array element positions and the water column is needed to be able to determine the sub-bottom properties. The end-fire path is directly between the source and receiver,

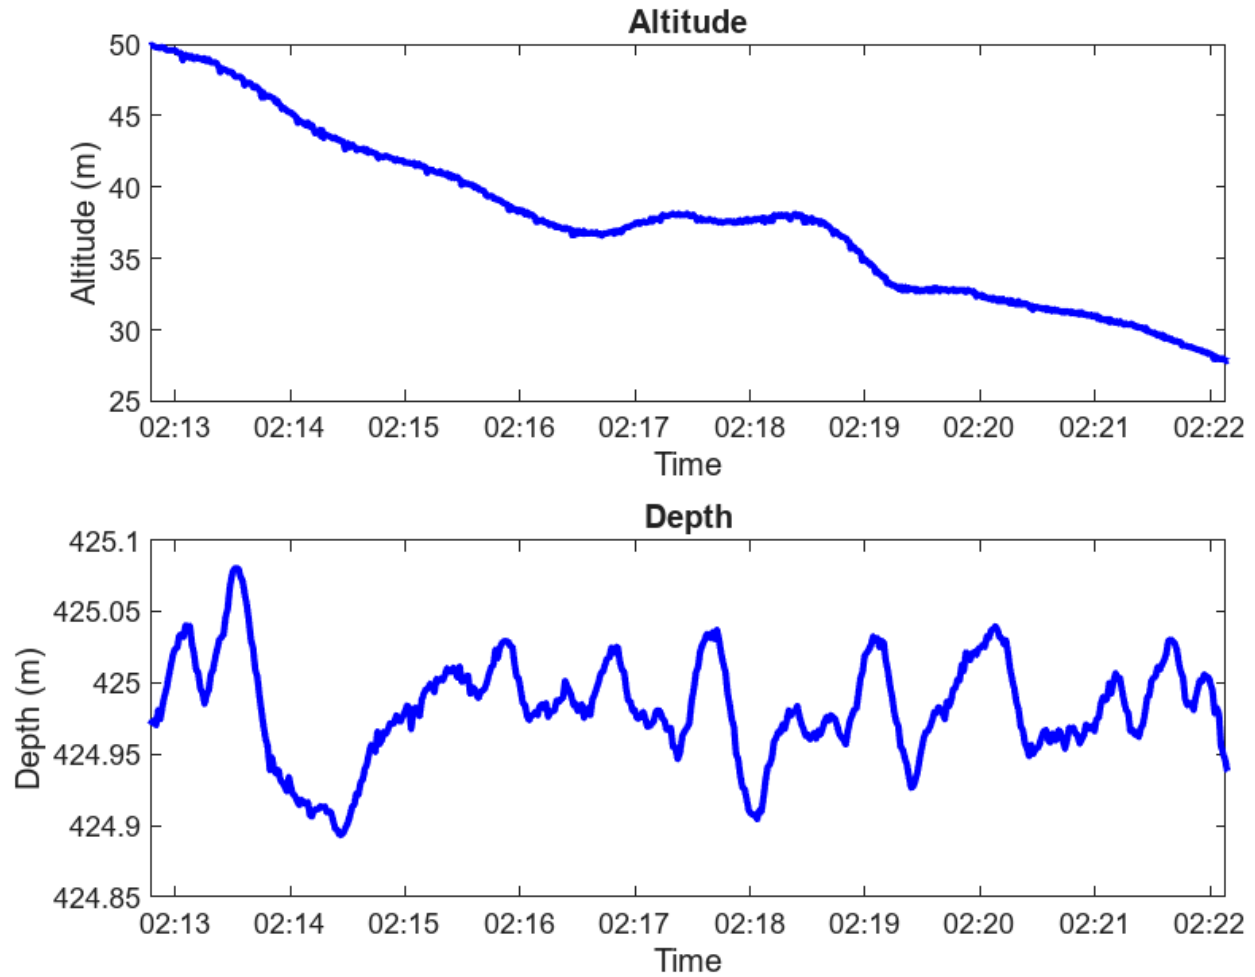


Figure 1.6: Altitude and depth measurements of the AUV mid-mission for the shelf break experiment

and the end-fire timing data provides information about the exact receiver positions. The bottom bounce path travels through the water column and provides information about water height and water sound speed. The goal of this thesis is to determine the sound speed and layer thickness of the sub-bottom using the timing measurement on the array of the different travel paths.

The AUV gives measurements of sound speed and altitude, and the rough location of the hydrophones is known. The system is defined with a series of equations. In these equations, the variables with the tilde are the measured variables, and the variables without the tilde are the actual values of these variables. For variables that have measurements, the actual

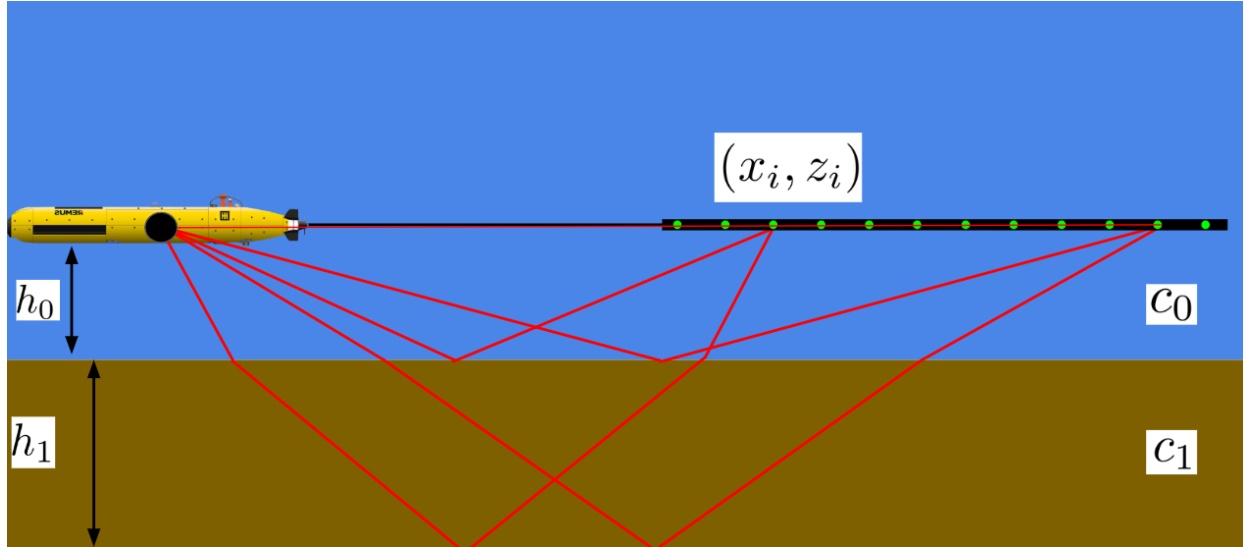


Figure 1.7: A depiction of the AUV with the direct, bottom bounce, and sub-bottom bounce travel paths on two different towed array elements

values are assumed to be close to the measured values with some amount of sensor noise.

$$\begin{aligned} x_i - \tilde{x}_i \\ z_i - \tilde{z}_i \end{aligned} \tag{1.1}$$

Equation 1.1 defines the x and z positions of hydrophone i where the origin $(0, 0)$ is defined as the AUV source and the array is assumed to be level with the source at $z = 0$. The positive x direction is towards the end of the array and the positive z direction is down. Due to the flexibility of the AUV towed array, the exact positions of the elements with respect to the source vary with the movement of the water and the AUV.

$$\begin{aligned} c_0 - \tilde{c}_0 \\ h_0 - \tilde{h}_0 \end{aligned} \tag{1.2}$$

Equations 1.2 define the difference between the measured water sound speed and the measured water height or AUV altitude. The sound speed is measured with a CTD on board the AUV, and the water height is measured with an altimeter.

$$\frac{\sqrt{x_i^2 + z_i^2}}{c_0} - (\tilde{t}_{end_i} - t_{offset}) \quad (1.3)$$

Equation 1.3 is the expected direct path arrival time on hydrophone i . All timing equations include a t_{offset} variable. The source emits a signal every second, but the exact time that the source is operating within that second period is unknown. t_{offset} is between $[0, 1]$ and is fundamentally when the source starts transmitting.

$$\frac{\sqrt{x_i^2 + (2h_0 - z_i)^2}}{c_0} - (\tilde{t}_{bottom_i} - t_{offset}) \quad (1.4)$$

Equation 1.4 is the expected bottom bounce arrival time on hydrophone i .

$$\frac{2h_0 - z_i}{c_0} \sin \theta_{sub-bottom_{0,i}} + \frac{2h_1}{c_1} \sin \theta_{sub-bottom_{1,i}} - (\tilde{t}_{sub-bottom_i} - t_{offset}) \quad (1.5)$$

Equation 1.5 is the expected arrival time of the sub-bottom bounce on hydrophone i . The angle $\theta_{sub-bottom_{0,i}}$ is the grazing angle of the sound hitting the bottom layer for hydrophone i . The sound then refracts through the layer. The angle $\theta_{sub-bottom_{1,i}}$ is the refraction angle of the sound as it travels through the bottom layer and reflects off of the sub-bottom. In this equation, c_1 is the sound speed of the sub-bottom layer, and h_1 is the sub-bottom layer thickness.

$$\frac{\cos \theta_{sub-bottom_{0,i}}}{c_0} - \frac{\cos \theta_{sub-bottom_{1,i}}}{c_1} \quad (1.6)$$

Equation 1.6 defines the relationship between $\theta_{sub-bottom_{0,i}}$ and $\theta_{sub-bottom_{1,i}}$ with Snell's Law [9]. Snell's law is a formula that defines the refraction of a wave as it passes between two layers of different sound speeds.

$$(2h_0 + z_i) \cot \theta_{sub-bottom_{0,i}} + h_1 \cot \theta_{sub-bottom_{1,i}} - x_i \quad (1.7)$$

Equation 1.7 relates the $\theta_{sub-bottom_{0,i}}$ and $\theta_{sub-bottom_{1,i}}$ angles to the receiver position.

Equations 1.6 and 1.7 are used together to uniquely determine the $\theta_{sub-bottom_{0,i}}$ and $\theta_{sub-bottom_{1,i}}$ angles. The derivation is done in Appendix B

When working with timing data, h_1 and c_1 are inherently linked because it takes the same amount of time for sound to travel through a thick layer with a fast sound speed as it does for sound to travel through a thinner layer with a slower sound speed. The assumption is that the timing data from different sound paths between the source and the various elements of the towed array will provide enough information to uniquely determine h_1 and c_1 .

There were a few modeling decisions that were made for this problem:

- The water and sub-bottom have a constant sound speed. For the dataset used, the AUV altitude is anywhere between 5 – 50 m . The highest order depth term in the simplified Mackenzie equation for sound speed is 10^{-2} [10]. This indicates that over the distance of 5 – 50 m , the sound speed is only likely to change on the order of 10^{-1} m/s due to the change in depth which is not significant enough to create issues. Equation 1.2 also accounts for any potential differences in the measured and actual sound speed.
- The array is straight with some amount of tilt. This assumption is valid in areas where the AUV is moving in a straight line and will be discussed further in Section 3.2.
- The bottom and sub-bottom are range-independent along the length of the array. This assumption means that the bottom properties are constant over some horizontal distance. Most geoacoustic inversion schemes assume a range-independent bottom over $O(10^3)$ meters [11, 12].
- Only the initial reflections are taken into account. Any surface-bottom reflections are ignored as are reflections between the bottom layers. This is known as the Born approximation where multiple reflections between layers are considered to be too low in amplitude [13]. The surface-bottom reflections and bottom-surface reflections are visible on the acoustic data, but because the AUV is in a deep water environment, it is relatively simple to remove these reflections via time gating.

The entire processing scheme has the following set of steps:

1. Sampling rate calibration
2. Hydrophone distance calibration
3. Pulse compression
4. Peak detection
5. Rough layer detection
6. Arrival time estimation
7. Multi-task Gaussian process regression
8. Bayesian geoacoustic inversion

The sampling rate calibration, hydrophone distance calibration, pulse compression, and rough layer detection steps are discussed in [Appendix A](#).

1.3 Summary

This chapter presents a solution to geoacoustic inversion using timing returns from an AUV equipped with a source to a towed array. The mathematical formulation was outlined and the experiments used in this thesis were discussed. The structure of the remainder of this thesis is as follows:

Chapter [2](#) discusses prior work in AUV tow array processing, Gaussian process regression, and geoacoustic inversion. Chapter [3](#) explores the arrival time estimation of the bottom and sub-bottom layers using sub-array beamforming and sub-array focusing. Chapter [4](#) focuses on estimating the scattering seen in the timing returns of the bottom and sub-bottom layers using Gaussian process regression. Chapter [5](#) uses the results from the Gaussian process regression to determine the sub-bottom properties using timing-based Bayesian geoacoustic

inversion. Chapter 6 summarizes the contributions of this thesis and discusses potential future research directions.

Chapter 2

Literature Review

2.1 AUV Towed Array Processing

Towed hydrophone arrays or towed arrays are long linear strings of hydrophones that are pulled by an ocean vessel. Towed arrays have been used since the early 1900s as a way to improve received signals underwater [14]. They are currently used in a wide variety of applications such as mine detection, seabed mapping, and ocean monitoring [15]. Typically, long towed arrays are pulled by surface vessels. AUV towed arrays are ideal for geoacoustic studies because they can operate close to the seabed which reduces the effects of the water column on sound propagation, they generate less noise than a surface vessel, and they operate autonomously for long periods which means that large regions can be surveyed without dedicated ship time [16]. The downside of using AUVs is that they are only capable of towing short arrays due to their smaller size and carrying capacity. In the experiments discussed in this thesis, the hydrophone portion of the array is 11.1 meters long with a total length of just over 30 meters. Towed arrays pulled by surface ships can be over 1000 meters long with multiple acoustic sections that can be up to 50 meters in length [17].

2.1.1 AUV Towed Array Beamforming

Beamforming using an AUV towed array is a challenging problem because of the vehicle's motion and the fact that AUVs are only capable of towing short arrays which reduces the spatial aperture and resolution. Beamforming using AUV towed arrays has been well studied. Lynch et al. [18] showed that the signal could be improved using a conventional beamformer on a short AUV towed array. Lynch et al. and Pallayi et al. [16, 18] also showed that the array was stable when the AUV was operating in a straight line. Holmes et al. [19] demonstrated the ability to use a time delay beamformer on an array towed by a REMUS AUV.

2.1.2 AUV Towed Array Synthetic Aperture Sonar

When an array is on a moving object such as an AUV, the array can be artificially lengthened using a technique known as synthetic aperture sonar (SAS) to improve bearing resolution [20]. This makes SAS particularly appealing for AUV towed arrays due to their inability to pull long arrays. Work has been done to show that passive synthetic aperture array processing is possible on an AUV towed array [20]. SAS has successfully been used to track fish and bubble clouds [21], map the seafloor [22], and inspect underwater pipelines [23].

SAS generally requires good knowledge of hydrophone location and feature identification [24]. This is tricky for towed arrays doing seabed surveys where the sensor locations can move and the seabed does not contain strong features that can be matched. There have been many bodies of work [15, 25, 26] on estimating the shape of the towed array as it moves through the water. Hansen et al. [22] discusses the challenges in SAS which include vehicle instability, navigation errors, sound velocity errors, issues related to unknown topography, and the presence of multi-path in shallow water environments.

SAS typically uses higher frequency multi-beam or side scan sonars, however, these frequencies are unable to penetrate the bottom and therefore cannot provide any geoacoustic information. Steele et al. and Steele et al. [27, 28] show that low-frequency SAS processing

can be used for geoacoustic studies.

2.1.3 AUV Towed Array Near Field Beamforming

Near-field beamforming is required when a source is in the near field of an array and the signal waveform is propagating spherically instead of in a plane wave [29]. Near field beamforming is a large research topic in wireless communications where arrays are becoming larger and waveforms are becoming shorter to accommodate more users and speed up data rates [30]. Despite the fact that AUV towed arrays are typically short due to the limited towing abilities of an AUV, and many ship-towed arrays are not close enough to signal sources to be in the near field, there have been several bodies of research focusing on near-field beamforming using towed arrays.

A large issue when using towed arrays is that the vehicle towing the array is commonly in the near field of the array. To locate far-field targets, suppressing the near-field interferer is required. Lee et al. [31] uses sub-arrays to minimize near-field interferers and focus on sources in the far field. The sub-arrays allow for the near-field interferers to be considered to be in the far field of the sub-array due to the shorter aperture. Wei et al. [32] also introduces an alternative algorithm to suppress near-field interferers. Sun et al. [29] compares far-field and near-field beamforming on an optical towed array and concludes that near-field beamforming outperforms far-field beamforming when objects are located in the near field. Sun et al. [29] includes tests with a very large aperture (308m) towed array where the far field includes a large range.

2.2 Geoacoustic Inversion

Geoacoustic inversion describes methods used to determine the properties of the seabed using acoustic data. There have recently been several studies using AUV towed arrays systems for geoacoustic surveys [12, 18, 33, 34]. Current geoacoustic methods with towed arrays typically

use the array to get a wide variety of reflection angles from the bottom and measure the amplitude of the reflection. A mapping between the angle and the return amplitude is then created. Different sediment types have different angle characteristics [35]. The returns are then compared to models of different sediment types and classify the sediment [12, 33, 34]. The downsides of these methods are that they require pre-computed models of sediment materials, they deal with a significant amount of noise between similar sediment materials, and they struggle when the bottom is composed of several different types of sediment.

2.2.1 Bayesian Geoacoustic Inversion

Bayesian geoacoustic inversion is a widely applicable inversion method that was first introduced for its use in geoacoustic inversion by Dosso et al. [36]. Bayesian geoacoustic inversion can apply to any dataset and model, and the method characterizes the most likely fit between the model and the dataset [37]. There have been many iterations of Bayesian geoacoustic inversion that use different datasets and models such as broadband geoacoustic inversion [38], passive geoacoustic inversion using ship noise [39], trans-dimensional geoacoustic inversion with a vertical array and impulse source [40], and image source inversion [41]. The Bayesian inversion model is widely applicable to many geoacoustic measurement techniques.

2.2.2 Timing Based Geoacoustic Inversion

The goal of this thesis is to do geoacoustic inversion based on travel time estimates. Travel time geoacoustic inversion has been a topic of interest since the 1950s [42]. Early attempts used sources and receivers at the top of the ocean and were sensitive to the sound speed profile of the water [43] and the layers had to be at least 1/15th of the water thickness to be able to resolve them [44]. These older works also usually assume a sound speed in the bottom layer to calculate the thickness.

More recent works have used the Bayesian inversion model to calculate the most likely geoacoustic model. Dettmer et al. [45] shows that travel time data can be used to invert

geoacoustic parameters at a high resolution. It also demonstrates that neglecting the error covariances of the measured data leads to inaccurate and misleading inversion results that have an unrealistically high confidence.

Pinson et al. [13] introduces the image source method as a way to sequentially estimate the geoacoustic parameters of multiple sediment layers. The image source method is a widely used method in ocean acoustics for multi-interface interactions [9]. Pinson et al. [13] uses both the angle of arrival and the travel times of the waves to estimate the geoacoustic properties by minimizing the difference between the modeled image source location and the actual image source location. Drira et al. [46] then improves upon this method by using the Teager-Kaiser energy operator to more accurately estimate the arrival times. The image source method is advantageous because it has a low complexity and the calculation time is low. Pinson et al. [47] shows the effectiveness of the image source method on experimental data using an AUV with a source and a towed hydrophone array.

2.3 Uncertainty Estimation

Understanding uncertainty is important for Geoacoustic inversion as seen in Dosso et al. [36] which suggests using a Bayesian formulation to estimate the uncertainty of the estimation based on the variance of the data. This work was then followed up with trans-dimensional geoacoustic inversion that uses a data-error model to account for unknown errors in the data [40].

In the Geoacoustic field, one of the main goals of the SBCEX experiments was to quantify the uncertainties of the seabed [5, 7]. Many different methods including source modeling [48], direct measurements [49], short-range remote sensing [50], and long-range remote sensing [51] were used to measure the frequency-dependent sound speed in this area which all yielded slightly different results, but showed significant overlap when the uncertainty of the results was considered.

Gaussian process regression is a non-parametric machine learning technique that is used to estimate regressions and quantify their uncertainties. Gaussian process models have successfully been used to fit and to understand the uncertainty of acoustic data to do inversion [52], robust localization [53], direction of arrival estimation, [54], and sound field reconstruction [55]. These results show that using Gaussian process regression yields an improvement over traditional methods.

Chapter 3

Arrival Time Estimates with Sub-Array Beamforming

The ocean is an acoustically noisy environment. There are many sources of noise including geological processes waves such as wind, currents, and seismic activity, anthropogenic sources like commercial shipping, and biological sources such as shrimp, fish, and whales. These sources of noise can make transmitting information in the ocean challenging [9]. The signal-to-noise ratio (SNR) is a comparison between the level of the desired signal and the level of the background noise (often expressed in decibels). Numerous signal processing and array processing techniques are used to improve the SNR by either reducing the noise level or improving the signal. The easiest method for improving the SNR is to increase the signal strength of the transmitted signal. In ocean environments, however, the maximum signal strength is frequently limited to protect marine life [39]. Pulse compression is one such method as described in Appendix A that is commonly used to increase the signal level of known signals. The goal of this chapter is to determine whether an array processing technique known as beamforming can improve the SNR and allow for accurate timing estimations.

3.1 Linear Array Plane Wave Beamforming

Arrays are commonly used in signal-processing contexts to increase the SNR via various filtering methods [56]. In situations where a signal is originating from a specific direction, spatial filters can be used to increase the SNR by increasing the gain in the direction of the signal while canceling noise that is originating from other directions. Two commonly used spatial filtering methods are beamforming and, in situations where an array is moving, synthetic aperture array processing. Beamforming is a technique of combining the inputs of several elements in an array coherently to increase the energy in a specific look direction and minimize interference from other directions [9]. Conventional beamforming methods such as linear beamforming or Bartlett beamforming adjust the phase of the signal for each element for a given specified source angle [9].

Bartlett beamforming aims to maximize the received signal power from a desired direction by steering the array response. Below is the mathematical formulation.

Consider a linear array of N uniformly spaced receivers. The signal received at the array can be expressed as:

$$\mathbf{y}(t) = \mathbf{A}(\theta)\mathbf{s}(t) + \mathbf{n}(t),$$

where:

- $\mathbf{y}(t)$: $N \times 1$ is the vector of received signals,
- $\mathbf{A}(\theta)$: $N \times 1$ is the steering vector for a signal arriving from direction θ ,
- $\mathbf{s}(t)$ is the desired signal,
- $\mathbf{n}(t)$ is the noise vector ($N \times 1$),

The steering vector $\mathbf{A}(\theta)$ for a uniform linear array with elements spaced d apart is:

$$\mathbf{A}(\theta) = \begin{bmatrix} 1 \\ e^{j2\pi d \frac{f}{c} \sin \theta} \\ \vdots \\ e^{j2\pi(N-1)d \frac{f}{c} \sin \theta} \end{bmatrix},$$

where f is the signal frequency and c is the sound speed of the medium.

The output of the beamformer for a specific direction θ_0 is:

$$z(\theta_0) = \mathbf{w}^H \mathbf{y}(t),$$

where:

- \mathbf{w} is the beamforming weight vector,
- \mathbf{w}^H is the Hermitian transpose or complex conjugate of \mathbf{w} .

In Bartlett beamforming, the weights are set to match the steering vector for the desired direction θ_0 :

$$\mathbf{w} = \mathbf{A}(\theta_0).$$

Thus, the beamformer output becomes:

$$z(\theta_0) = \mathbf{A}^H(\theta_0) \mathbf{y}(t).$$

The power of the beamformer output is given by the square of the magnitude of $z(\theta_0)$:

$$P(\theta_0) = |z(\theta_0)|^2 = |\mathbf{A}^H(\theta_0) \mathbf{y}(t)|^2.$$

In practice, the signal $\mathbf{y}(t)$ is represented by its cross-spectral density matrix \mathbf{R} to average out noise and analyze steady-state performance. The cross-spectral density matrix of the

received signal is:

$$\mathbf{R} = \mathbf{y}(t)\mathbf{y}(t)^H.$$

The power for Bartlett beamforming is then computed as:

$$P(\theta_0) = \mathbf{A}(\theta_0)^H \mathbf{R} \mathbf{A}(\theta_0).$$

The Bartlett beamformer scans across directions θ by computing the output power for each angle. The resulting power $P(\theta)$ as a function of θ forms the beam pattern:

$$P(\theta) = \mathbf{A}(\theta)^H \mathbf{R} \mathbf{A}(\theta).$$

Assuming the source is the strongest signal, the direction of the source can be found by determining the angle that maximizes the beamformed power output [9]. The beam pattern is the response of the array over a range of angles. It is typically defined by a main lobe in the look direction which is the distance between the nulls on either side of the look direction as well as side lobes outside of the look direction that amplify signals coming from unwanted directions [56]. Generally, good beam patterns have a narrow main lobe and low amplitude side lobes. The beam pattern $B(\theta)$ of an N element line array with uniform spacing is the following:

$$B(\theta) = \left| \sum_{n=0}^{N-1} e^{j2\pi n d \frac{f}{c} \sin \theta} \right|^2.$$

This simplifies to:

$$B(\theta) = \left| \frac{\sin \left(N \pi d \frac{f}{c} \sin \theta \right)}{\sin \left(\pi d \frac{f}{c} \sin \theta \right)} \right|^2. \quad (3.1)$$

The main lobe width depends on the length of the array, the frequency of the signal, the element spacing, and the look direction relative to broadside [56]. For a uniform linear array and a transmitted signal at a single frequency, the main lobe is the distance between the first

two zeros of Equation 3.1. The first nulls of the beam pattern occur where the numerator of $B(\theta)$ equals zero:

$$N\pi d \frac{f}{c} \sin \theta = m\pi, \quad m \neq 0, \pm 1, \pm 2, \dots$$

Solving for θ gives:

$$\sin \theta = \frac{mc}{Ndf} = \frac{m\lambda}{Nd},$$

where the wavelength $\lambda = \frac{c}{f}$. Therefore the first nulls on either side of the main lobe are:

$$\sin \theta_{\text{null}} = \pm \frac{\lambda}{Nd}.$$

The main lobe width or beamwidth is the angular is given by:

$$\Delta\theta = \theta_{\text{null},2} - \theta_{\text{null},1}.$$

For small angles where $\sin \theta \approx \theta$ the angular distance between the first nulls is:

$$\Delta\theta \approx 2 \frac{\lambda}{Nd}. \tag{3.2}$$

In Equation 3.2, Nd is another way to define the length of the array $L = Nd$. When the look direction is far from the broadside of the array, the apparent length of the array shrinks which means the main lobe becomes wider. The length of the apparent array is:

$$L_{eff} = L \cdot \cos(\theta).$$

Grating lobes appear in the beam pattern when the array geometry fails to satisfy the spatial Nyquist sampling theory [56]. For the array to satisfy the spatial Nyquist sampling theory, the array elements must be closer than half a wavelength apart. When a plane wave with wavelength λ and angle θ reaches the array, the wave introduces a phase difference between adjacent elements. This phase difference is given by:

$$\Delta\phi = \frac{2\pi d}{\lambda} \sin \theta.$$

To avoid grating lobes, the phase difference between the elements must uniquely determine the direction of arrival θ without any ambiguity. When the element spacing d exceeds $\frac{\lambda}{2}$, the phase differences become ambiguous because sine waves are periodic over a period of 2π which leads to grating lobes. These are undesired peaks in the array's beam pattern.

The positions of grating lobes can be determined by the following condition:

$$\sin \theta_m = \sin \theta + \frac{m\lambda}{d}, \quad m = \pm 1, \pm 2, \dots,$$

where θ_m are the angles corresponding to the grating lobes, and m is the aliasing order. When $d > \lambda/2$, grating lobes appear within the possible range of angles ($-90^\circ \leq \theta \leq 90^\circ$), which degrades the beam pattern of the array by introducing peaks in unwanted directions [56].

The beam pattern of the array in the SBCEX22 experiment can be seen in Figure 3.1. The array spacing in this experiment is $0.74 m$, and the maximum frequency of the source is $4.25 kHz$. For a $4.25 kHz$ signal, the array elements must be closer than $\sim \frac{1500}{4250} = 0.35 m$ apart to satisfy the spatial Nyquist sampling theory. The beam patterns shown in Figure 3.1 have grating lobes because the array elements are further than $0.35 m$ apart. The main lobe also increases in size as the look direction gets closer to end-fire.

3.2 Linear Array Approximation for the towed array

To beamform an array, the positions of each element relative to the source must be known. Most commonly used arrays are rigid and the element locations can be accurately measured. Despite a towed array being a flexible line of hydrophones, towed arrays are frequently approximated as linear arrays as long as the platform dragging the towed array is operating in a straight line [25]. Both the shelf break and mud patch experiments have the AUV operating

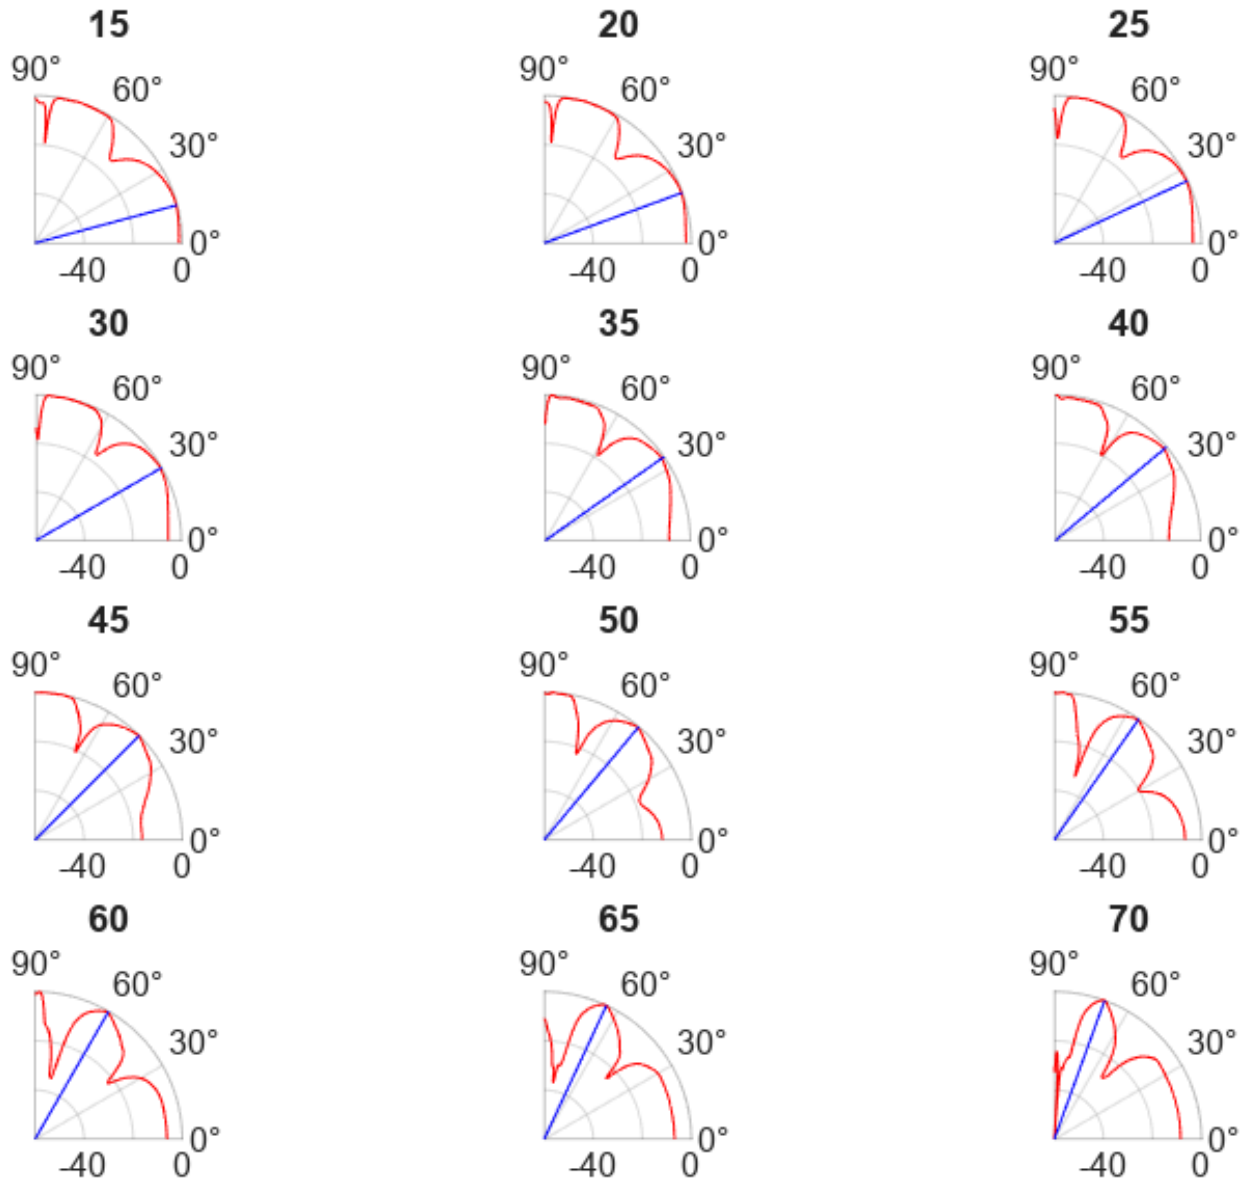


Figure 3.1: Beam pattern of the array in the SBCEX22 experiment for different look directions. The look direction is shown in blue.

in a lawn mower pattern where the AUV is operating in a straight line over a long duration as seen in Figure 1.2. For this thesis, only the long straight tracks or middle tracks were considered.

In the mud patch experiment, the AUV was operating at a constant altitude meaning the depth of the AUV changes as the bathymetry changes. The bathymetry in the mud patch is very slightly sloped, and therefore the AUV maintained a constant pitch of approximately

-1.5° during the analyzed segment. Because of the slight pitch of the AUV, the array is modeled as a linear array with some tilt. In the shelf break experiment, the AUV operated at a constant depth which means that the array is very likely to be level and straight during the middle tracks.




Figure 3.2: The hydrophone positions during a straight track (left) and during a turn (right)

The left image in Figure 3.2 shows the calculated distance of the hydrophones from the end fire returns during one of the straight middle paths whereas the right image in Figure 3.2 shows the calculated hydrophone distances during a time period where the AUV is turning. While the AUV is operating in a straight line, the calculated distance based on the direct path arrival time does not change and the calculated distances match the measured distances. In this situation, the array can be modeled as a straight linear array with known hydrophone positions. The calculated distances during the turn show that the array is bending and therefore cannot be assumed to be a straight linear array.

3.2.1 Near Field vs Far Field

To process the incoming field on the array with linear beamforming, the incoming signal must be able to be approximated as a plane wave which means that the array must be in the far field of the transducer [57]. If the transducer is approximated as a point source, the initial signal created by the transducer is a spherically spreading signal that takes the form:

$$p(r, t) = \frac{A}{r} \cos(\omega t - kr), \quad (3.3)$$

where:

- $p(r, t)$ is the acoustic pressure at a distance r from the source and time t ,
- A is the amplitude of the wave, proportional to the source strength,
- r is the radial distance from the point source,
- $\omega = 2\pi f$ is the angular frequency of the wave, with f being the frequency,
- $k = \frac{2\pi}{\lambda}$ is the wavenumber, with λ being the wavelength,
- t is the time,
- $\cos(\omega t - kr)$ represents the harmonic oscillation of the wave.

The $\frac{A}{r}$ term represents the spherical spreading. It is this term that plane wave beamforming does not account for.

As the wave travels, the sphere expands until at some point, the spherical wave starts to look locally linear and approximates a plane wave. The distance at which this happens depends on the wavelength of the signal and the aperture of the array. At this distance, the wave takes on the form of a plane wave [58]:

$$p(x, t) = P_0 \cos(\omega t - kx),$$

where P_0 is a constant amplitude of the pressure wave.

The distance at which a spherical waveform can be approximated as a plane wave is called the Fraunhofer distance R [58]. The Fraunhofer distance is defined as:

$$R = \frac{2D^2}{\lambda}, \quad (3.4)$$

where λ is the wavelength and D is the array aperture. For a linear array, the aperture is the array length. When in the near field, the wavefront cannot be approximated as locally linear, and so plane wave beamforming cannot be applied to the wavefront [57] [29]. In the far field, the wave can be approximated as a plane wave, and plane wave beamforming can be used to beam energy in a specific direction.

In the mud patch experiment, the maximum frequency is 4.25 kHz which corresponds to a wavelength of approximately $\frac{1500}{4250} = 0.35 \text{ m}$ and the array length is 11.1 m . The Fraunhofer distance for this geometry is $\frac{2 \cdot 11.1^2}{0.35} = 704 \text{ m}$ which means both the bottom and the sub-bottom are in the near field for the array. The same is true of the shelf break experiment where the maximum frequency is slightly higher at 4.5 kHz which corresponds to a wavelength of approximately $\frac{1500}{4500} = 0.33 \text{ m}$

3.3 Sub-Array Beamforming

The reason for splitting the array into several sub-arrays is twofold. First, at least two sub-bottom timing measurements are needed to uniquely determine both the sub-bottom sound speed and the sub-bottom layer thickness. Second, the array as a whole is in the near field of the bottom, and the only way to alter this is by shortening the array. The sub-arrays contain a subset of array elements and have a smaller aperture than the entire array. The bottom and sub-bottom returns can be approximated as in the far field for each sub-array which allows plane wave beamforming to be utilized to improve the SNR.

For the shelf break experiment, the mean altitude of the AUV is 20.36 m above the

seabed. The maximum array diameter where the seabed is in the far field is $\sqrt{\frac{1}{2}R\lambda} = \sqrt{\frac{1}{2} \cdot 20.36 \cdot 0.33} = 3.66 \text{ m}$. With an average array spacing of 0.74 m , a sub-array consisting of a maximum of 5 hydrophones is in the far field of the bottom returns.

The drawback of sub-array beamforming is by shortening the array, the beamwidth of the plane wave beamformer response is larger meaning it spans a wider range of angles. A smaller beamwidth means that the beamformer is more selective [9]. As the number of elements increases, the beamwidth decreases exponentially.




Figure 3.3: Comparison between single element returns (left) and the result of sub-array beamforming (right) on the sub-bottom layer using data from the shelf break experiment

Figure 3.3 shows the sub-bottom returns for a single element and the sub-bottom beamformed returns of a 4-element sub-array. Both plots have a decibel color range of 30 dB . The sub-bottom has a higher SNR in the beamformed sub-array compared to the single-element returns. Despite the wider beam width, the sub-arrays can still be used to increase the SNR.

3.4 Near Field Beamforming

Near-field beamforming is used when a wave is still propagating spherically and thus is in the near field of the array. Near-field beamforming has recently been a large area of research due to the increasing need to accommodate more and more users on wireless communication

arrays at faster data rates. Increasing the number of users and increasing the data rates requires larger arrays and smaller wavelengths that will force more and more communication devices into the near field of communication arrays [30]. Unlike linear beamforming, near field beamforming requires knowledge of both the angle to the signal source and the distance to the signal source [29]. Each element in the array must be pointed at a different angle so that all of the elements end up focusing on the same point.

3.5 Sub-Array Focusing

Because of the proximity of the array to the bottom, the length of the array, and the source frequency, the bottom and sub-bottom returns are no longer plane waves. Instead, they are spherical waves with curved wavefronts. Near-field beamforming requires knowledge of both the array element positions and the direction and distance to the signal source location. When using an AUV towed array on the sub-bottom, neither the location of the array elements is well known nor the location of the sub-bottom reflection point. To determine the sub-bottom reflection point and to average any noise in the vicinity of individual elements of the array, the array is subdivided into shorter overlapping sub-arrays of n hydrophones. Because the sub-arrays are shorter, they are in the far field of the bottom and sub-bottom, and therefore linear beamforming is possible.

In Figure 3.4, the process begins with the first sub-array, where the first group of elements in the array (from 1 to n) beamforms to a point on the sub-bottom, labeled s_0 . The height h between s_0 and the array is calculated based on the angle of maximum return found during the beamforming process. The sub-array then shifts down the array by $\frac{n}{2}$ elements or $\frac{n+1}{2}$ for n odd. The next sub-array consisting of elements $\frac{n}{2}$ to $\frac{n}{2} + n$ is beamformed to the same point by calculating the angle to s_0 using h and the position of the mid-point of the new sub-array. This cycle continues until there are no more elements left on the array. The equivalent equation for this process is:




Figure 3.4: An illustration of sub-bottom as viewed by different sub-arrays during the AUV track

$$\phi_0 = \max_{\phi} b_0(\phi)$$

$$h = \frac{1}{2} x_0 \tan(\phi_0)$$

$$B(s_0) = \sum_{i=1}^{N-n} b_i(\phi_i)$$

where

- $B(s_0)$: The point on the sub-bottom being beamformed.
- N : The number of hydrophones on the array.
- n : The number of hydrophones used in the sub-array beamforming
- ϕ_i : The phase adjustment for the i -th pulse to account for the movement of the AUV and ensure coherent addition. $\phi_i = \arctan\left(\frac{1}{2} \frac{h}{x_i}\right)$ where x_i is the center point of the sub-array
- $b_i(\phi_i)$: The beamformed signal from the i -th pulse to the angle ϕ_i .




Figure 3.5: The sub-bottom processed with sub-array synthetic aperture beamforming compared to a single sub-array beamforming using data from the shelf break experiment

Using this technique, the returns from the sub-arrays are enhanced as seen in Figure 3.5. All images have the same 30 *dB* range. The sub-bottom returns in Figure 3.5 have been greatly improved using the sub-array beamforming, and the two sub-bottoms are clearly visible in returns due to a higher SNR. Due to the wide beam width of the sub-arrays, the sub-array beamforming captures and enhances both sub-bottoms. This work is reproduced from [59], with the permission of AIP Publishing

3.6 Timing Returns from Beamforming

The primary goal of the sub-array beamforming is to improve the signal-to-noise ratio of the sub-bottom return to accurately estimate the timing returns of the sub-bottom. The time-based geoacoustic inversion method requires at least two accurate timing measurements to uniquely determine the sub-bottom layer thickness and sub-bottom sound speed.

The timing returns of the bottom and sub-bottom on each sub-array are calculated by beamforming to the bottom or sub-bottom and then finding the time of maximum amplitude of that layer.

Figure 3.6 shows the different sub-arrays beamformed to the sub-bottom and the calculated

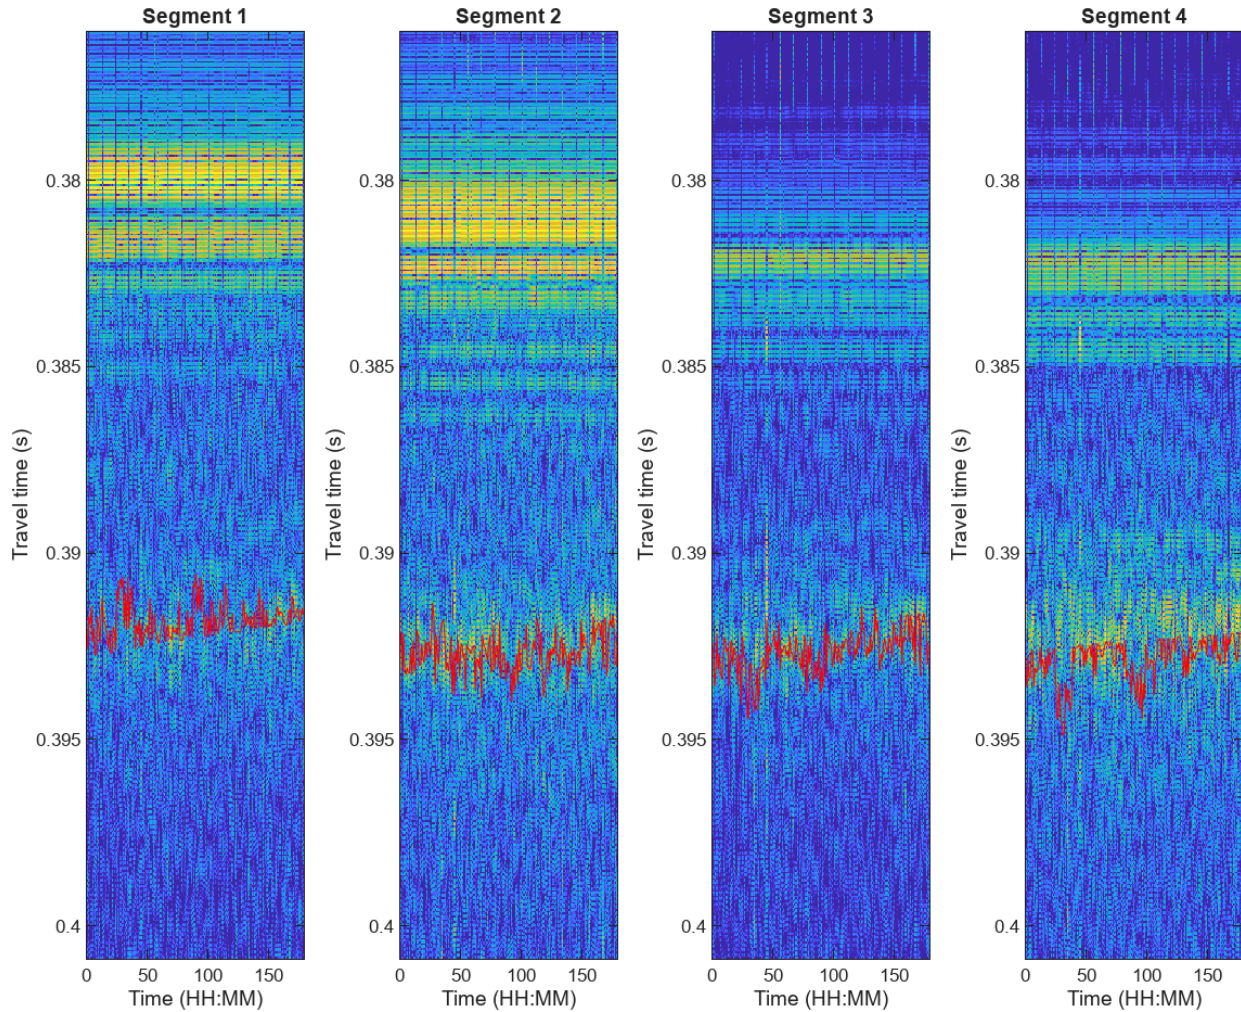


Figure 3.6: The timing returns calculated for the sub-bottom layer for four different sub-arrays using data from the mud patch experiment

timing returns overlaid on the beamforming results using data from the mud patch experiment. These figures show that the timing returns calculated match the incoming waveforms.

However, when the returns from each sub-array are plotted over one another as seen in Figure 3.7, an issue arises. The model assumes that the sound reaches the first sub-array before the second sub-array because the travel path between the source and the first sub-array is shorter than the travel path between the source and the second sub-array. As seen in the sub-bottom returns in Figure 3.7, this is not the case. There is a significant amount of overlap in the timing returns between the different sub-arrays. The data shown in Figure 3.6 shows that although the timing returns match the sub-bottom beamforming results, the

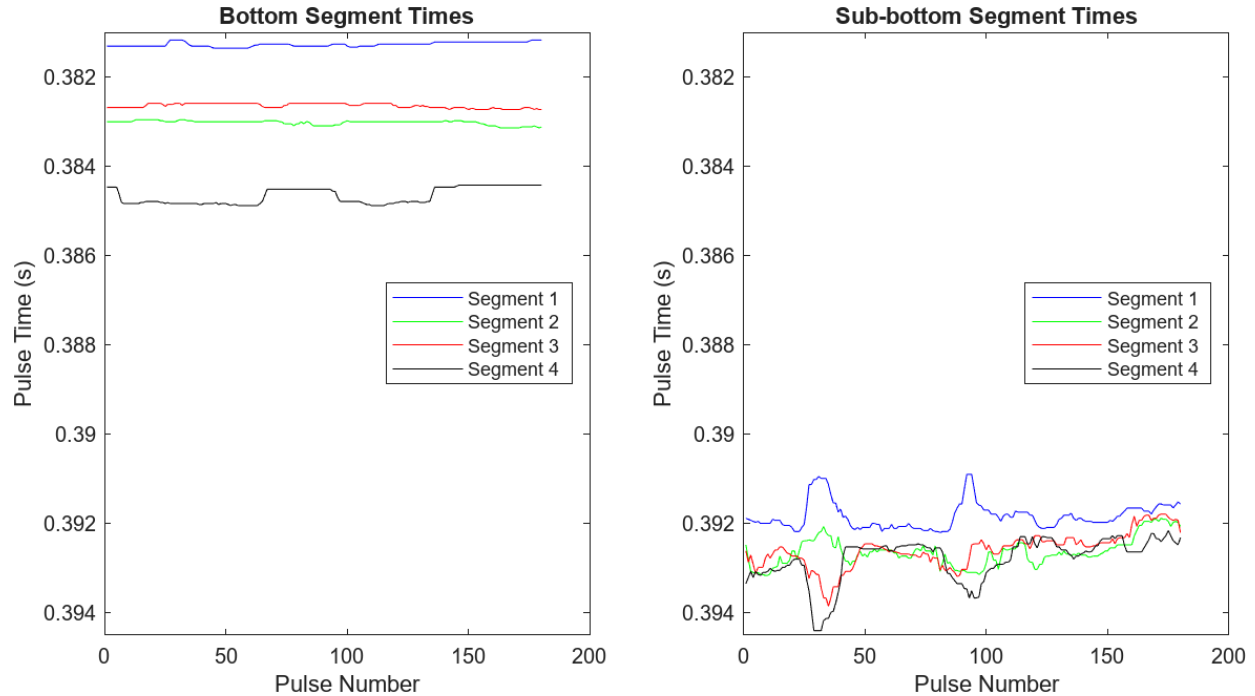


Figure 3.7: The timing returns from all sub-arrays for the bottom (left) and sub-bottom (right). This data came from the mud patch experiment.

sub-bottom returns are still very noisy which is causing the overlapping seen in Figure 3.7.

Beamforming assumes a correlation between the signal data from the different elements on the array [9]. If the noise on the array is uncorrelated, this allows the combination of the correlated signals on multiple array elements to create a signal with a higher SNR by coherently adding the signal and incoherently adding the noise.

Figure 3.8 shows the amplitude and phase of the complex envelope of the waveform of approximately 100 pulses from a single hydrophone. The rough locations of the end-fire, bottom, and sub-bottom returns are highlighted. The figure shows that the phase for all of the pulses align with one another for both the end-fire and bottom returns. This indicates that there is a strong correlation between pulses for both the end-fire and bottom returns. Each pulse is approximately $1.6m$ apart and each element is $0.74m$ apart, so the presence of a correlation of the received signal between pulses would indicate that there is also a correlation signal between hydrophones. However, there is no pattern in the phase of the sub-bottom arrivals. This indicates that there is a lack of correlation in the sub-bottom returns. One

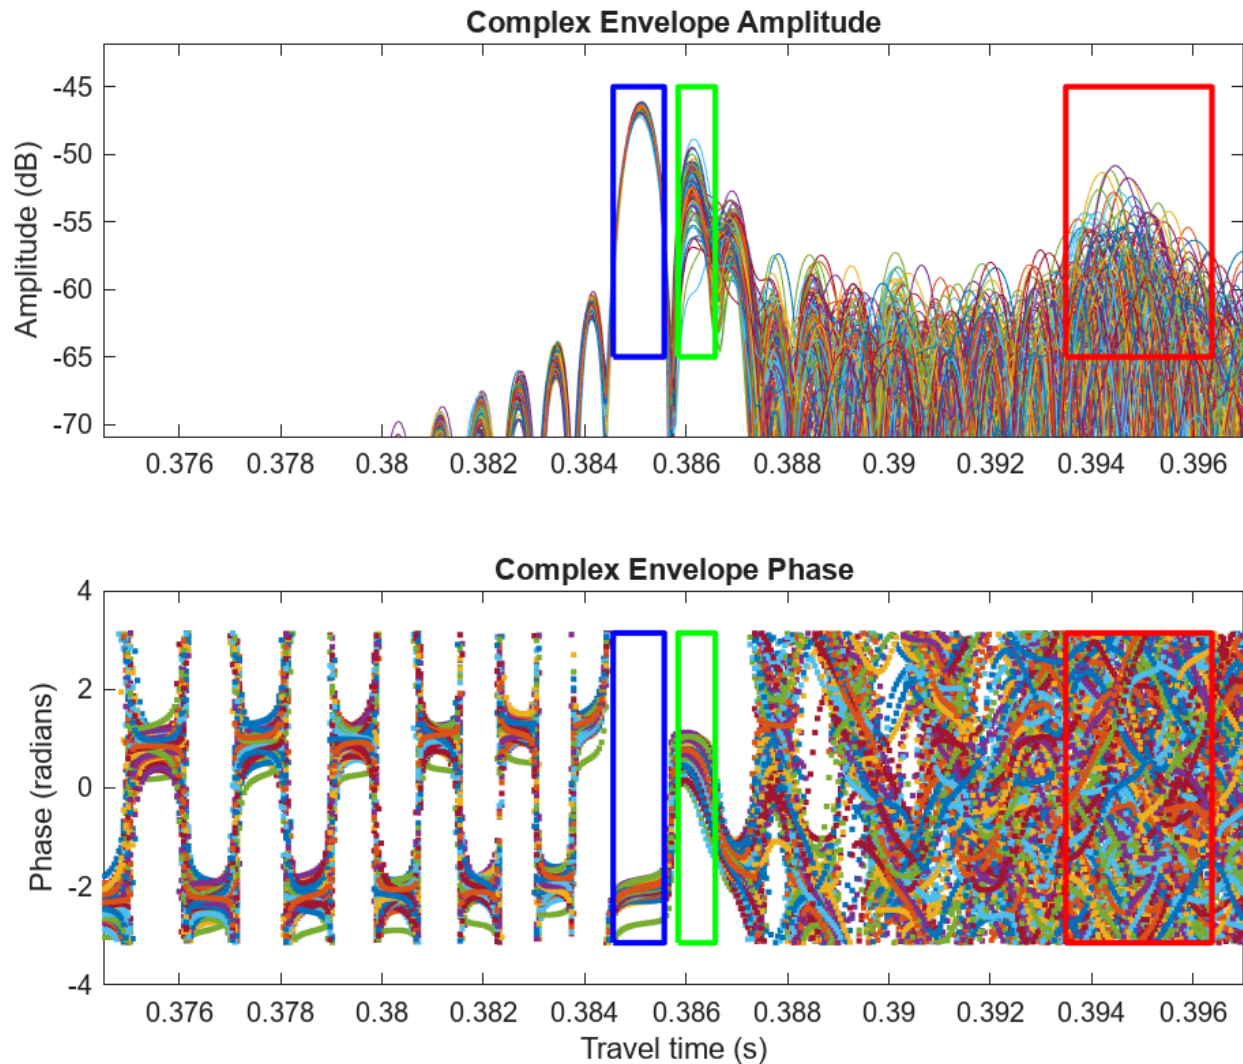


Figure 3.8: The amplitude and phase of approximately 100 waveforms on one channel with the direct path, bottom returns, and sub-bottom returns marked. This data came from the mud patch experiment.

potential cause of this is scattering in the sub-bottom layer is the presence of small-scale heterogeneities such as rocks [60]. The lack of coherence in the sub-bottom signals indicates that the sub-bottom scattering cannot be improved through beamforming. This scattering and lack of coherence in the sub-bottom returns are likely the reason that the sub-bottom arrival times are overlapping in Figure 3.7. Beamforming is able to improve the SNR of the sub-bottom layer, but because the scattering is in the sub-bottom layer, beamforming is unable to fix the overlap in the timing return estimates caused by the scattering. Therefore,

beamforming will not work as a method of accurately estimating the arrival times despite the SNR increases. To estimate the timing returns, the scattering in the sub-bottom layer needs to be understood and quantified.

3.7 Summary

This chapter demonstrated how the use of sub-array beamforming and sub-array focusing can be used to enhance bottom and sub-bottom reflections in the near field of a towed array. However, due to scattering in the sub-bottom layer, sub-array beamforming cannot be used to estimate the arrival times of the sub-bottom reflections on the array. The beamforming is unable to remove the effects of the scattering. The next chapter investigates using Gaussian process regression to quantify the scattering seen in the sub-bottom layer and estimate the timing returns of the sub-bottom layer.

Chapter 4

Gaussian Process Regression

4.1 Estimating Arrival Times and Variance

As seen in the previous chapter, the sub-bottom timing returns show a significant amount of variance due to scattering in that layer. Because beamforming did not work to remove the effects of the scattering in the sub-bottom timing returns, a different method is needed in order to estimate the timing returns of the bottom and sub-bottom layers. As seen in the previous section, there is likely scattering in the sub-bottom layer that causes the sub-bottom returns to arrive at inconsistent times across the array.

To estimate the sub-bottom returns, the scattering needs to be understood. To accomplish this, the sub-bottom layer is roughly marked by hand, and then the sub-bottom arrival time is defined as the point of maximum amplitude within the rough sub-bottom limits. The rough layer marking is described further in Appendix A. Figure 4.1 shows the received waveform and the calculated sub-bottom arrival times (shown in red) for each pulse for 16 channels. Figure 4.1 indicates that there is a trend between the channels where the sub-bottom arrival times are later for channels that are further from the source, however because of the scattering on the sub-bottom, the data is very noisy and the sub-bottom returns occur over a wide time span on each channel.




Figure 4.1: The sub-bottom return on 16 channels with the estimates of the timing of the sub-bottom returns shown in red.

Figure 4.2 shows histograms of the bottom and sub-bottom arrival times on the 16 channels of the array. The bottom arrival time histograms show that the signal is behaving as expected with the sound reaching closer hydrophones sooner and further hydrophones later. However the sub-bottom arrival time histograms show that there is significant overlap between the sub-bottom arrival times indicating that the variance in the sub-bottom layer is high. Generally, the sub-bottom arrival times do increase as the distance between the source and the hydrophone increases, but there is high variance in the data compared to the bottom arrival times. To calculate the Bayesian Inversion to determine the sub-bottom sound speed and layer thickness, estimating the mean and variance of the arrival times is required.

Gaussian process (GP) regression is a probabilistic machine learning method that defines a predictive function between a set of inputs and a set of noisy outputs as well as the uncertainty of the prediction [61]. A GP is determined by its mean function $m(\mathbf{x})$ and covariance function

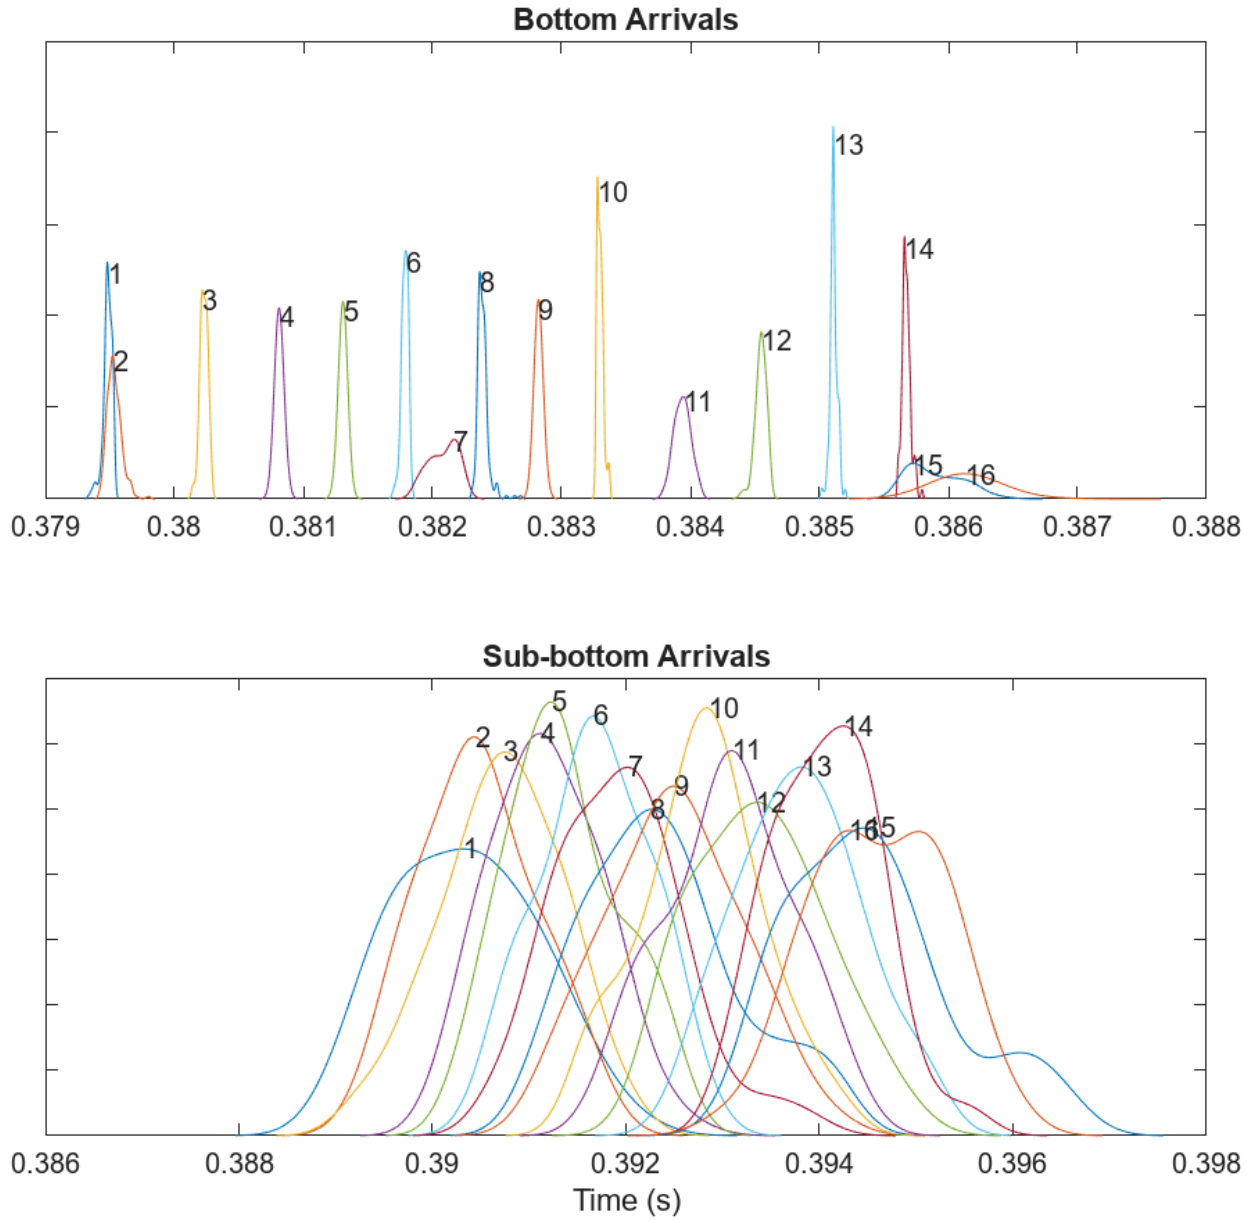


Figure 4.2: A histogram of the arrival times of the bottom and sub-bottom layers

$k(\mathbf{x}, \mathbf{x}')$ [62]:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \quad (4.1)$$

where

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})], \quad (4.2)$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))] \quad (4.3)$$

It is frequently assumed that $m(\mathbf{x}) = 0$ and datasets can be normalized to ensure a zero mean [62].

Given n training inputs $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]^T$ and corresponding outputs $\mathbf{y} = [y_1, y_2, \dots, y_n]^T$ where the outputs are assumed to be noisy observations of the true function $f(\mathbf{x})$:

$$y_i = f(\mathbf{x}_i) + \epsilon_i, \quad (4.4)$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma_n^2)$ is Gaussian noise, so the likelihood of $\mathbf{y} \mid \mathbf{f}$ is:

$$\mathbf{y} \mid \mathbf{f} \sim \mathcal{N}(\mathbf{f}, \sigma_n^2 \mathbf{I}). \quad (4.5)$$

Given some test input \mathbf{x}_* , let $f_* = f(\mathbf{x}_*)$ be the function of the true value at the testing point \mathbf{x}_* . Let $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^T$ be the values of the true function at the training points. The joint distribution of \mathbf{f} and f_* with the GP prior is:

$$\begin{bmatrix} \mathbf{f} \\ f_* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{0} \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{K} & \mathbf{k}_* \\ \mathbf{k}_*^T & k_{**} \end{bmatrix} \right), \quad (4.6)$$

where

$$\mathbf{K} = \{k(\mathbf{x}_i, \mathbf{x}_j)\}_{i,j=1}^n,$$

$$\mathbf{k}_* = [k(\mathbf{x}_*, \mathbf{x}_1), \dots, k(\mathbf{x}_*, \mathbf{x}_n)]^T,$$

$$k_{**} = k(\mathbf{x}_*, \mathbf{x}_*).$$

The posterior distribution of f_* given the observations \mathbf{y} is defined using a conditional probability on the data.

$$f_* | \mathbf{y} \sim \mathcal{N}(\mu_*, \sigma_*^2), \quad (4.7)$$

where

$$\begin{aligned} \mu_* &= \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}, \\ \sigma_*^2 &= k_{**} - \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_*. \end{aligned}$$

The mean μ_* represents the predicted value at the test point \mathbf{x}_* , while σ_*^2 quantifies the uncertainty of the prediction [62].

Given the end-fire, bottom, and sub-bottom timing data from a range of pulses. Gaussian process regression can be used to estimate the mean and variance of the timing data for all three layers.

4.2 Multi-Task Gaussian Process Regression

Multi-task Gaussian process regression extends traditional GP methods to model multiple related tasks simultaneously. Traditional GP regression works well for modeling complex relationships in data, but it assumes a single task where a task is a single relationship between inputs and outputs. Multi-task Gaussian process regression, on the other hand, utilizes correlations between multiple tasks to improve predictive performance for all outputs. Multi-task Gaussian process regression assumes one set of inputs for multiple sets of outputs. It is especially useful in scenarios where the different tasks are related or dependent in some way, like where tasks share the same physical system or in situations where certain tasks have smaller datasets [63].

The goal of multi-task Gaussian process regression is to model multiple output functions $f_1(x), f_2(x), \dots, f_N(x)$ where the outputs are correlated in some way. The first step of this

process is to define a joint Gaussian process for all of the tasks:

$$f(x) = [f_1(x), f_2(x), \dots, f_N(x)]^T \sim \mathcal{GP}(m(x), \mathbf{K}(x, x')),$$

where $m(x)$ is a vector and is the mean of the function for all tasks, and $\mathbf{K}(x, x')$ is a multi-task covariance function:

$$\mathbf{K}(x, x') = \mathbf{K}_{\mathbf{X}}(x, x') \otimes \mathbf{K}_{\mathbf{T}}. \tag{4.8}$$

The covariance function is made using the Kronecker product between the traditional GP covariance matrix $\mathbf{K}_{\mathbf{X}}(x, x')$ which models the correlations between inputs x and x' and the task covariance matrix $\mathbf{K}_{\mathbf{T}}$ that represents the correlations between different tasks. The covariance matrix in multi-task Gaussian process regression captures the inter and intra-task relationships which is beneficial in systems where there are dependencies between tasks. The combined covariance matrix leads to better performance modeling than running a GP on each task independently [63].

Because the end-fire, bottom, and sub-bottom arrival times are all part of the same physical system, multi-task GP regression can be used to estimate both the arrival times and the variance of the arrival times of all three layers simultaneously. To estimate the arrival times and variances, the data is segmented into sets of N timing measurements with a 50% overlap of timing measurements between sets. The bottom is assumed to be range-independent over these pulses. Multi-task GP regression is then performed on the arrival times of the three layers as three different tasks with the set of pulses as the training data points. The regression returns a mean time of arrival at each hydrophone for each layer as well as a variance of those times.

Figure 4.3 shows the multi-task Gaussian process regression for all three layers for 40 pulses. The regressions were calculated using GPyTorch [64]. The highlighted area is the confidence interval of the regression which is equivalent to two standard deviations. As

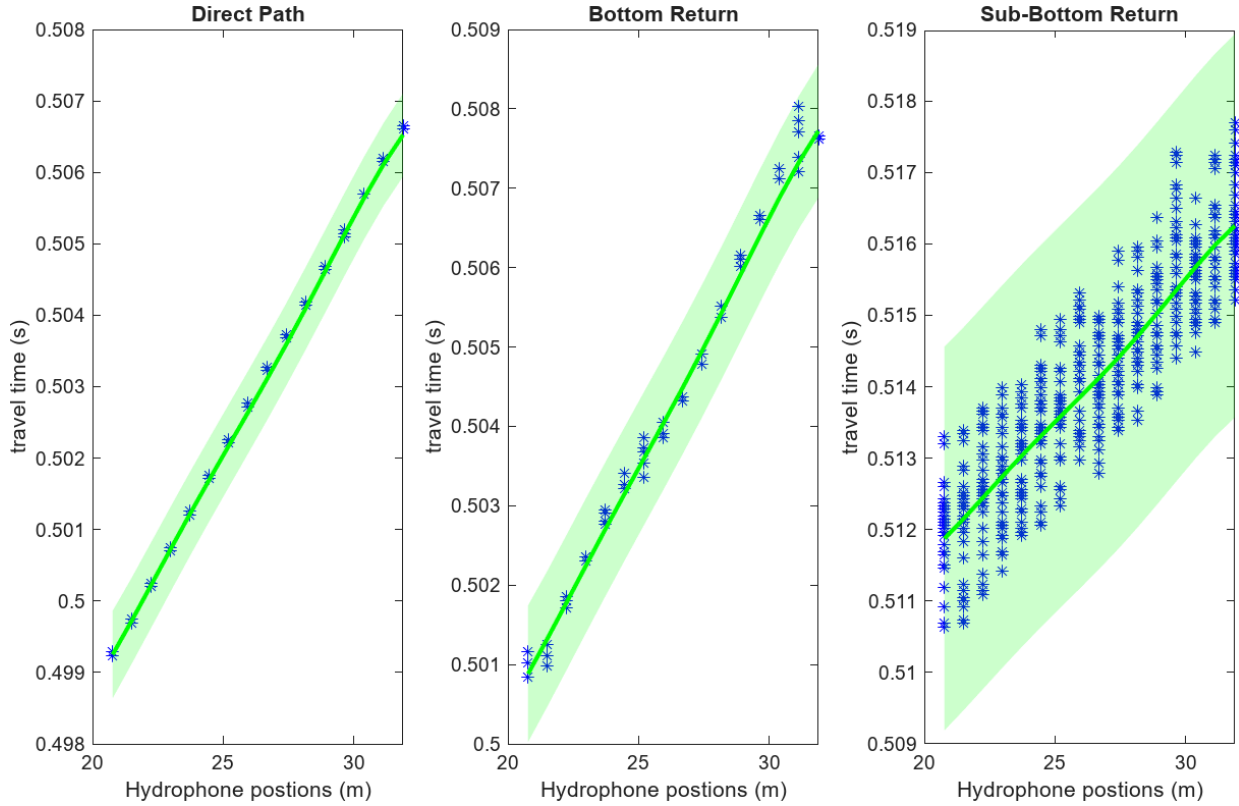


Figure 4.3: The results of the multi-task Gaussian process regression on the timing data for the three different arrival paths. 40 pulses were used as the training data for the regression.

expected from previous results, the confidence interval on the direct path returns is narrow which means the variance in those returns is small, the confidence interval in the bottom bounce returns is wider, and the confidence interval in the sub-bottom returns is large. This confidence interval defines the sub-bottom uncertainty that was seen in the previous section.

Figure 4.4 shows the timing results from the multi-task GP regression lining up with the raw waveform for three different channels. The confidence interval is outlined in green and the mean is the magenta lines. This figure shows that the confidence interval of the multi-task Gaussian process regression does account for the scattering and uncertainty in the timing returns of the different layers. These means and variances can then be used to calculate the timing base geoaoustic inversion.

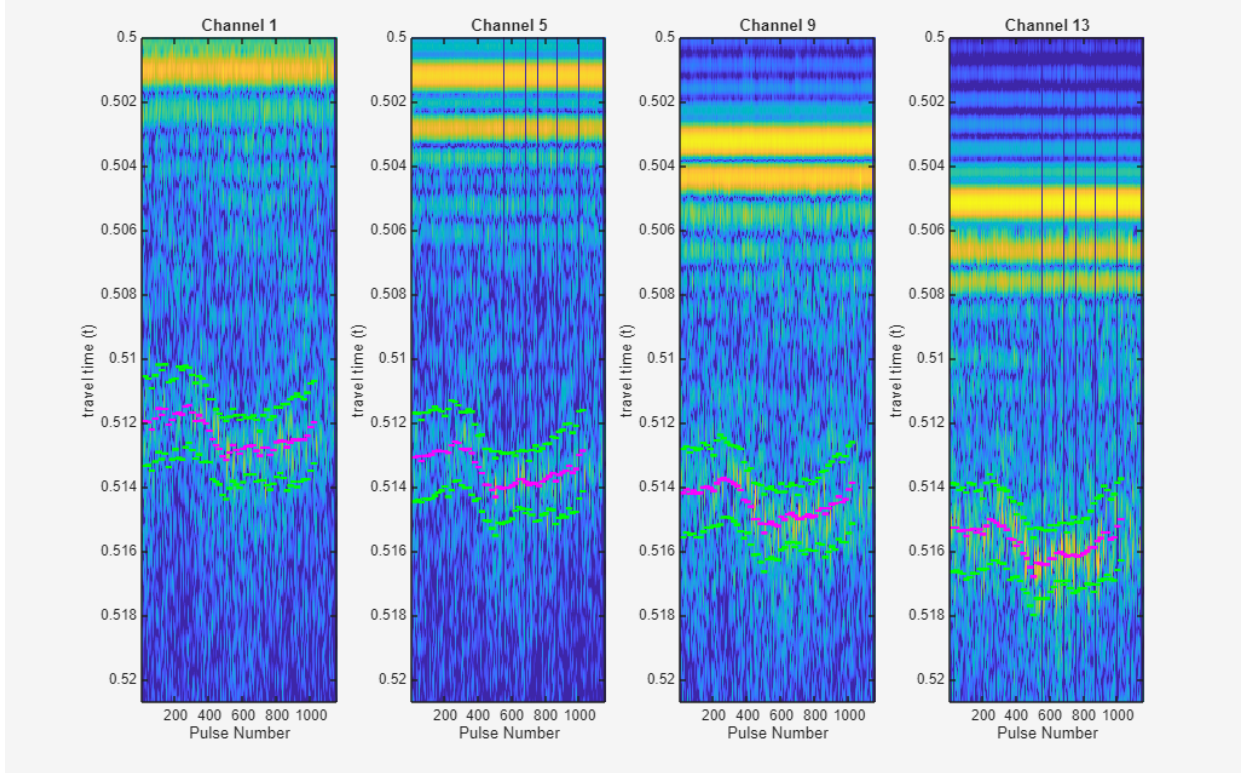


Figure 4.4: The results of the multi-task Gaussian process regression with the raw waveform indicating an accurate estimation of the mean and variance of the timing returns. The confidence interval is shown in green and the mean is shown in magenta.

4.3 Range Independent Length

One assumption that was made for the GP regression was that the bottom was range-independent over a series of N pulses. A range-independent bottom means that the bottom properties are constant over some distance. Most geoacoustic inversion methods assume range independence over distances from a few hundred meters to several kilometers [12].

The multi-task GP regression was run seven times with $N = [20, 25, 30, 35, 40, 45, 50]$ with a 50% overlap between the datasets used in each regression. The average AUV speed is 1.6 m/s , so $N = 20$ corresponds to a $\frac{20}{2} \times 1.6 = 16 \text{ m}$ distance between the different regressions and it assumes that the bottom is constant for 32 m . Figure 4.5 shows the results of the different GP regression runs for the different N values. From a high-level view, they all seem to agree with one another. The sub-bottom travel times remain fairly constant

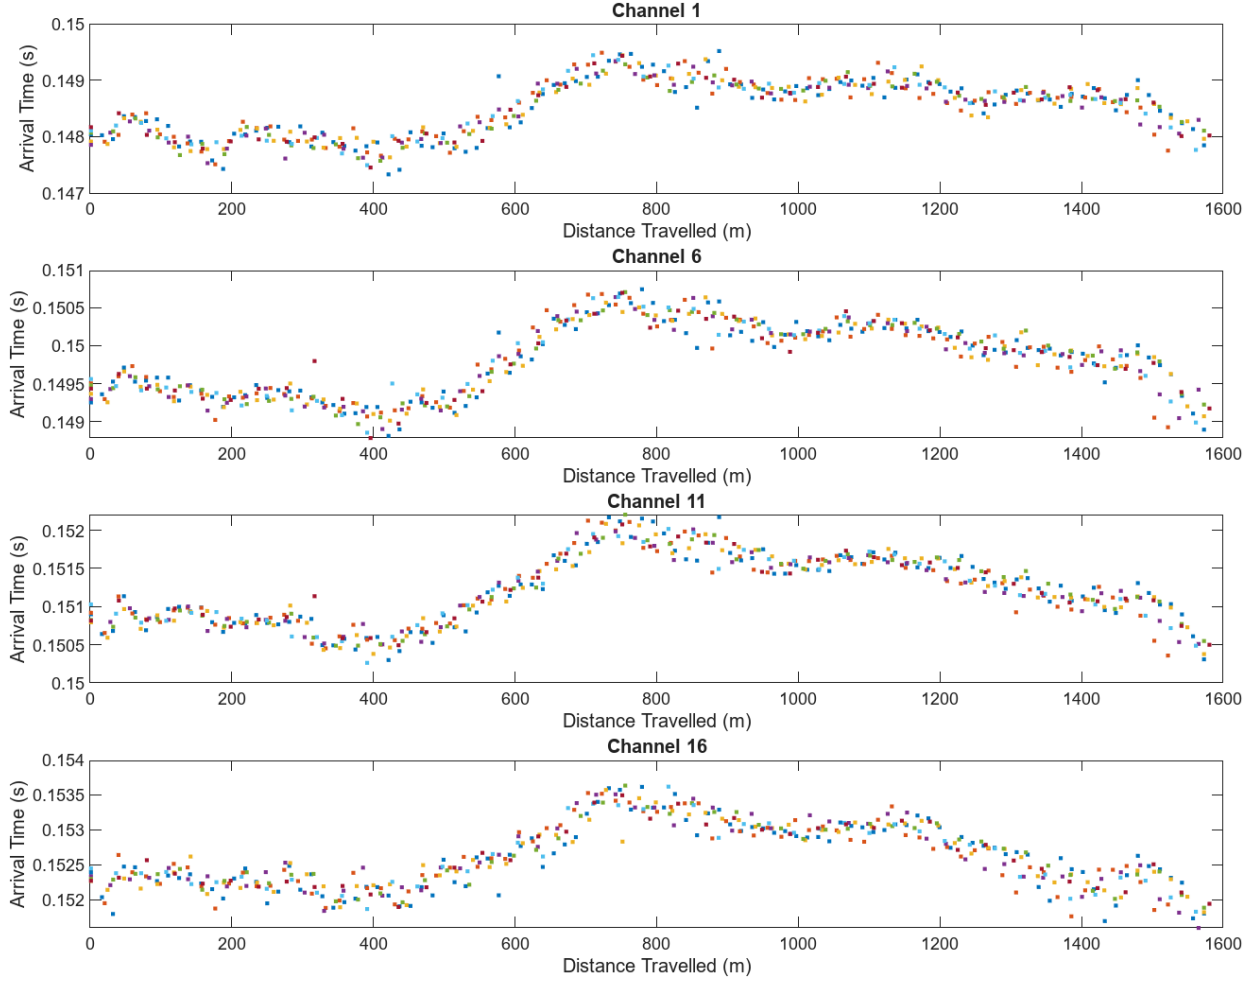


Figure 4.5: The timing returns from the sub-bottom Gaussian process regression for four channels. The timing return data with segments from 20 to 50 pulses to estimate the mean arrival time

from 0 – 200 *m* of travel distance, they linearly increase between 400 *m* – 800 *m* of travel distance, and then they linearly decrease between 800 *m* – 1600 *m* of travel distance. Between 400 *m* – 800 *m*, the mean slope of the arrival times is $4.5 \times 10^{-6} \text{ s/m}$. This means that over the duration of 20 pulses, the sub-bottom arrival times change approximately $9.1 \times 10^{-5} \text{ s}$. This is equivalent to the sub-bottom thickness changing by approximately 0.13 *m*. Over the duration of 50 pulses, the sub-bottom arrival times change approximately $2.2 \times 10^{-4} \text{ s}$. This is equivalent to the sub-bottom thickness changing by approximately 0.34 *m*. For this reason, the $N = 20$ data will be used for the geoacoustic inversion process. The results from Figure 4.5 indicate that the sub-bottom for the mud patch experiment can be approximated as range

independent over a range of at least 20 *m*, because 10 *cm* is within the range of uncertainty for the geoacoustic inversion.

4.4 Summary

This chapter demonstrated that multi-task Gaussian process regression can be used to estimate the variance in the bottom and sub-bottom timing returns. The results of the GP regression capture the scattering seen in the sub-bottom layer. The results also show that the sub-bottom layer changes slowly over the duration of the experiment, but that the sub-bottom layer can be considered to be range-independent over short intervals. In the next chapter, the mean and variance calculated using multi-task Gaussian process regression will be used to perform timing-based Bayesian geoacoustic inversion.

Chapter 5

Geoacoustic Inversion

5.1 Bayesian Geoacoustic Inversion

The goal of geoacoustic inversion is to find the geoacoustic model that is the most likely fit for the data and quantify the likelihood of the fit. This is commonly done through a matched-field method where a model m is created to represent the geoacoustic parameters. The error function $E(m)$ is used to quantify the mismatch between the model and the data. The best model is the one that minimizes the error. Bayesian inversion involves using a posterior probability density function to combine information about the model with observed data to characterize the solution [36]. Bayes' Theorem is used to determine the likelihood that a model fits the data. Bayes' theorem can be expressed as:

$$P(m|d) = \frac{P(d|m)P(m)}{P(d)}.$$

The prior is $P(m)$ which expresses the probability of the model. The likelihood function is $P(d|m)$ which expresses the probability of the data given the model and is a function of the model-data mismatch. The posterior is $P(m|d)$ which is the probability of the model given the data. The posterior contains information about the model fit and the uncertainty of the fit. $P(d)$ is the probability of the data. $P(d)$ is a normalizing constant that ensures that the

integral of the probability density function (PDF) is one.

The observed data d can be modeled as $d = g(m) + \epsilon$ where $g(m)$ is the model that predicts d and ϵ is the noise term. If the noise takes the form of a Gaussian distribution with zero mean and covariance Σ , $\epsilon \sim \mathcal{N}(0, \Sigma)$, then the probability density function of the noise term takes the form

$$P(\epsilon) = \frac{1}{\sqrt{2\pi\Sigma}} \exp\left(-\frac{1}{2}\epsilon^T \Sigma^{-1} \epsilon\right)$$

which is the likelihood function.

Since the noise term is equivalent to the difference between the model prediction and the observed data, $P(d|m)$ can be written as:

$$P(d|m) = \frac{1}{\sqrt{2\pi\Sigma}} \exp\left(-\frac{1}{2}(d - g(m))^T \Sigma^{-1} (d - g(m))\right)$$

The first term $\frac{1}{\sqrt{2\pi\Sigma}}$ does not depend on m , so $P(d|m)$ can be rewritten as

$$P(d|m) \propto \exp\left(-\frac{1}{2}(d - g(m))^T \Sigma^{-1} (d - g(m))\right)$$

The probability of the model or prior can then be added to get the posterior:

$$P(m|d) \propto \exp\left(-\frac{1}{2}(d - g(m))^T \Sigma^{-1} (d - g(m))\right) P(m)$$

This is derived in Dosso et al. [36] and Lin et al. [65].

Using the model defined in Section 1.2, along with the mean and variance of the timing data from the GP Regression, the posterior $P(m|d)$ can be calculated for a range of models to determine the most likely sub-bottom layer thickness and sound speed. The model for the timing-based Bayesian Inversion includes five variables: array tilt θ , water height h_0 , water sound speed c_0 , sediment or sub-bottom thickness h_1 , and sediment sound speed c_1 . The array tilt variable accounts for any difference between the measured array positions and

the actual array positions. The AUV provides measurements of the water height and water sound speed, however these measurements may still have some noise. The sediment thickness and sediment sound speed are completely unknown to the system.

The model of this system is relatively simple and easy to calculate as there are a small number of parameters. Because of this, a grid search over all parameters is computed and the likelihood of each point in the grid is calculated. In more complex systems, the parameter space is typically sampled using a sampling method such as a Gibbs sampler or a Monte Carlo method [36]. A grid search, while slower, provides a more complete picture of the posterior probability functions. To keep the computation time reasonable each parameter of the model is bounded. The array tilt calculation is bounded to $[-8^\circ, 5^\circ]$. This range was chosen based on the pitch of the AUV during the mud patch mission which is approximately -1.5° . A uniform distribution is assumed across the range of array tilts. The water sound speed is bounded between $\pm 3 \text{ m/s}$ from the measured sound speed. The water height is bounded to $\pm 0.75 \text{ m}$ from the measured AUV altitude. Both of these parameters assume a Gaussian distribution centered at the measured sound speed and water height. From previous studies of this area, the sediment layer thickness is known to be around 11 m [5], therefore the sediment layer thickness parameter is bounded between $[10 \text{ m} - 14 \text{ m}]$. The sediment sound speed is bounded between $[1425 \text{ m/s} - 1800 \text{ m/s}]$. A uniform prior is assumed for the sediment thickness and sound speed.

To test the model, test timing data was created with added noise. The variance of the noise for the direct path and bottom timing returns was set to the median variance from the GP regression results. The variance for the sub-bottom timing noise was varied. For low sub-bottom variance in the test data, the Bayesian geoacoustic inversion works well. The top graph in Figure 5.1 shows the modeled arrival times in red and the test arrival times in black. The lines match indicating that the model fits the test data well. The bottom graph in figure 5.1 shows the probability distribution functions for the individual parameters. The most likely value for each parameter is marked in red and the true value of each parameter is





Figure 5.1: The Bayesian inversion results for test data with low sub-bottom variance in the timing data. The top figure shows the modeled timing returns for all three layers in red and the test timing returns in black. The bottom figures show the posterior distribution for the different parameters. The most likely estimate for each parameter is marked in red and the true value is shown in green.

marked in green. The distributions where only a green line appears indicate a near perfect fit between the true and most likely values. With a low sub-bottom variance resulting in only a small amount of added noise, the timing based Bayesian inversion can correctly estimate the model parameters.




Figure 5.2: The percent error between the most likely value of each parameter and the true parameter values as a function of the sub-bottom variance. The figure on the left shows all five parameter values as a function of the sub-bottom variance. The figure on the right is focused on the layer thickness and sound speed parameters. The shaded region indicates typical sub-bottom variance levels from the GP Regression

The GP Regression results indicate that there is a high amount of uncertainty with the sub-bottom arrival times. By using test data, the variance of the data inputted into the Bayesian inversion model can be directly modified. Figure 5.2 shows the relationship between the variance level and percent error between the most likely calculated value of each

parameter and the actual values of each parameter. For each variance level, 100 different randomly initialized timing data test parameters were run through the model. The shaded region indicates typically sub-bottom variance levels from the GP Regression. These results show that the Bayesian Inversion is sensitive to the variance of the sub-bottom. As the variance of the sub-bottom increases, the inversion becomes less accurate.

The array tilt θ is the most sensitive variable. The approximately 60% error is only a difference of about 1° between the calculated and actual array tilt. A 1° difference in array tilt is only a 0.0011 m difference in the length of the direct path which would create approximately a $7 \times 10^{-7}\text{ s}$ travel time difference. With sub-bottom variances ranging from $1.5 \times 10^{-7}\text{ s}$ to $8 \times 10^{-7}\text{ s}$, this amount of array tilt error should be expected.

The water sound speed and AUV altitude are both resilient to sub-bottom variance. This is likely because these variables both have Gaussian priors centered around the measured values from the AUV, and they are both also dependent on the direct path and bottom bounce measurements which have a lower variance than the sub-bottom bounce measurements.

The sub-bottom layer thickness and the sub-bottom sound speed both see increasing error as the sub-bottom variance increases. These variables are only dependent on the sub-bottom measurements, and the posterior probability of these variables becomes more uncertain as the sub-bottom variance increases.

Another way of looking at the error between the calculated and the true value of the parameters is the data shown in Figure 5.3. This plot shows the posterior probability of all parameters as the sub-bottom variance increases. The red line indicates the most likely value for each parameter and the black line indicates the true value. The white line shows the median sub-bottom variance for the 20 minutes of mud patch data that was calculated using GP regression. A posterior probability for geoacoustic inversion that indicates a high likelihood that the maximum value is the correct value is a single narrow peak. A high narrow peak gives confidence that the true value is close to the most likely value. The relatively thin yellow regions on the water sound speed, the water height, and the array tilt graphs indicate

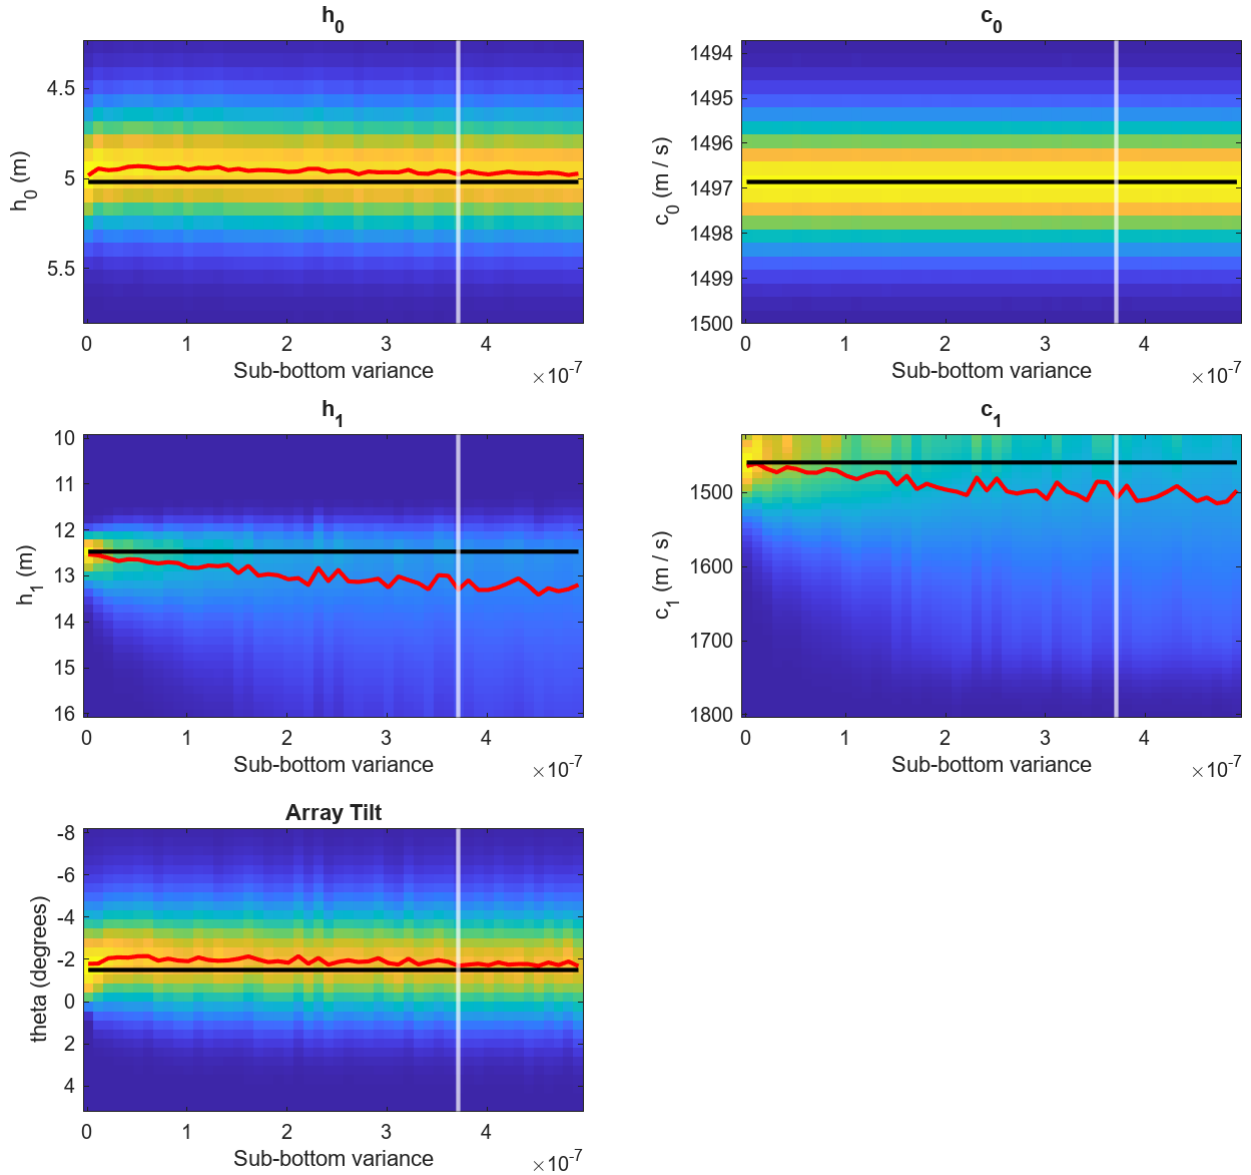


Figure 5.3: The posterior probabilities for each parameter of the test data as a function of the sub-bottom variance. The red line indicates the most likely value of that parameter and the black line indicates the true value. The white vertical line indicates the median sub-bottom variance of 20 minutes of mud patch data calculated using GP Regression

that posterior probability does not change as the sub-bottom variance increases, and the confidence level of these parameters is relatively high.

At very low sub-bottom variance levels, the posterior probabilities of the sub-bottom sound speed and sub-bottom layer thickness are both narrow and the maximum likely value is very close to the true value. However, as the sub-bottom variance increases, the

posterior probabilities quickly become wider and flatter indicating more uncertainty. The most interesting feature of these results is that as the sub-bottom uncertainty increases, the error between the most likely and true values of the sub-bottom layer thickness and sub-bottom sound speed increases in only one direction. The most likely layer thickness becomes higher than the true layer thickness and the most likely sound speed becomes faster than the true sound speed. As stated earlier, when working with timing data, these two values are inherently linked. A sound wave takes the same amount of time to travel through a thin layer at a slow sound speed as it does to travel through a thicker layer at a faster sound speed. These results indicate that this system is a biased estimator of the sub-bottom thickness and sound speed and the most likely results will generally show a faster sound speed and thicker sub-bottom layer as the sub-bottom variance increases.

5.2 Experimental Results

The same model that was run on the test data was then run on approximately 20 minutes of data from the mud patch experiment. To process a large middle track segment of data, the data was broken up into sections of 20 pulses with a 50% overlap between sections. The multi-task Gaussian process regression was run on each segment to get the means and variance of the timing returns. The means and variances were then inputted into the Bayesian Inversion model to get the likelihood distributions for the array tilt, water sound speed, water height, sub-bottom sound speed, and sub-bottom thickness.

The top graph in Figure 5.4 shows the arrival for the three layers over all 16 hydrophones for one section. The red dots are the modeled arrival times that correspond to the model with the maximum likelihood, and the black dots show the experimental data from the Gaussian process regression. This figure shows that the model does line up very well with the data indicating that there is a good match between the modeled system and the actual system.

The bottom graphs in Figure 5.4 show the likelihood of the array tilt, the water sound

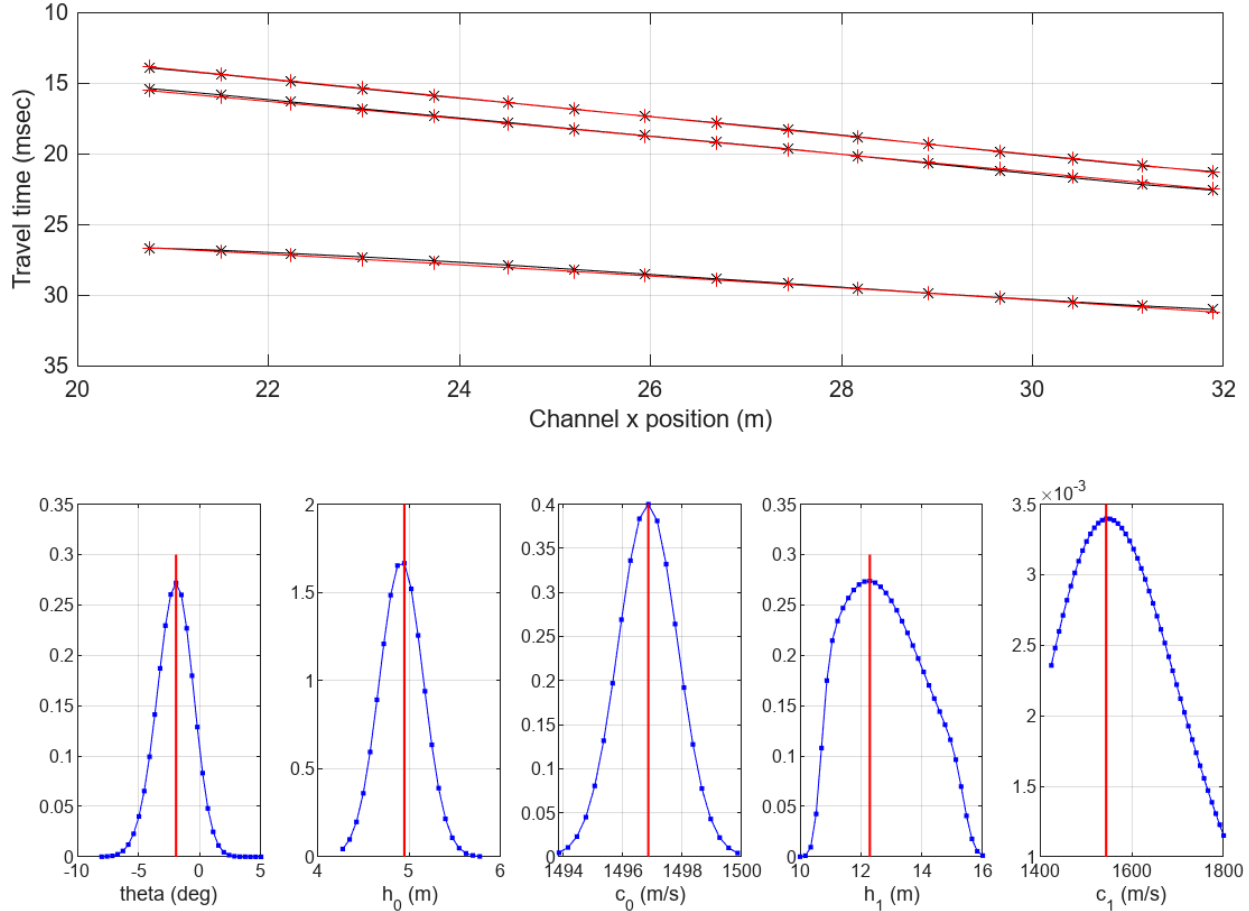


Figure 5.4: The Bayesian inversion results for experimental data. The top figure shows the modeled timing returns of the most likely estimated model for all three layers in red and the actual timing returns in black. The bottom figures show the posterior distribution for the different parameters. The most likely estimate for each parameter is marked in red.

speed, the water thickness, the sediment sound speed, and the sediment layer thickness for one segment. A wider distribution with a shorter peak indicates that the parameter estimation is more uncertain. As seen in Figure 5.4, the probability distributions for the water sound speed and water height are both very narrow which indicates that the maximum likely value is likely close to the actual value. However, the distributions for the sub-bottom sound speed and sub-bottom thickness are both wide and flat. This indicates that the most likely value is more uncertain and the true value could be very different from the maximum likely value.

Another way to look at the probability distributions is to look at the joint distributions between two variables. Figure 5.5 shows the joint distributions between different parameters.

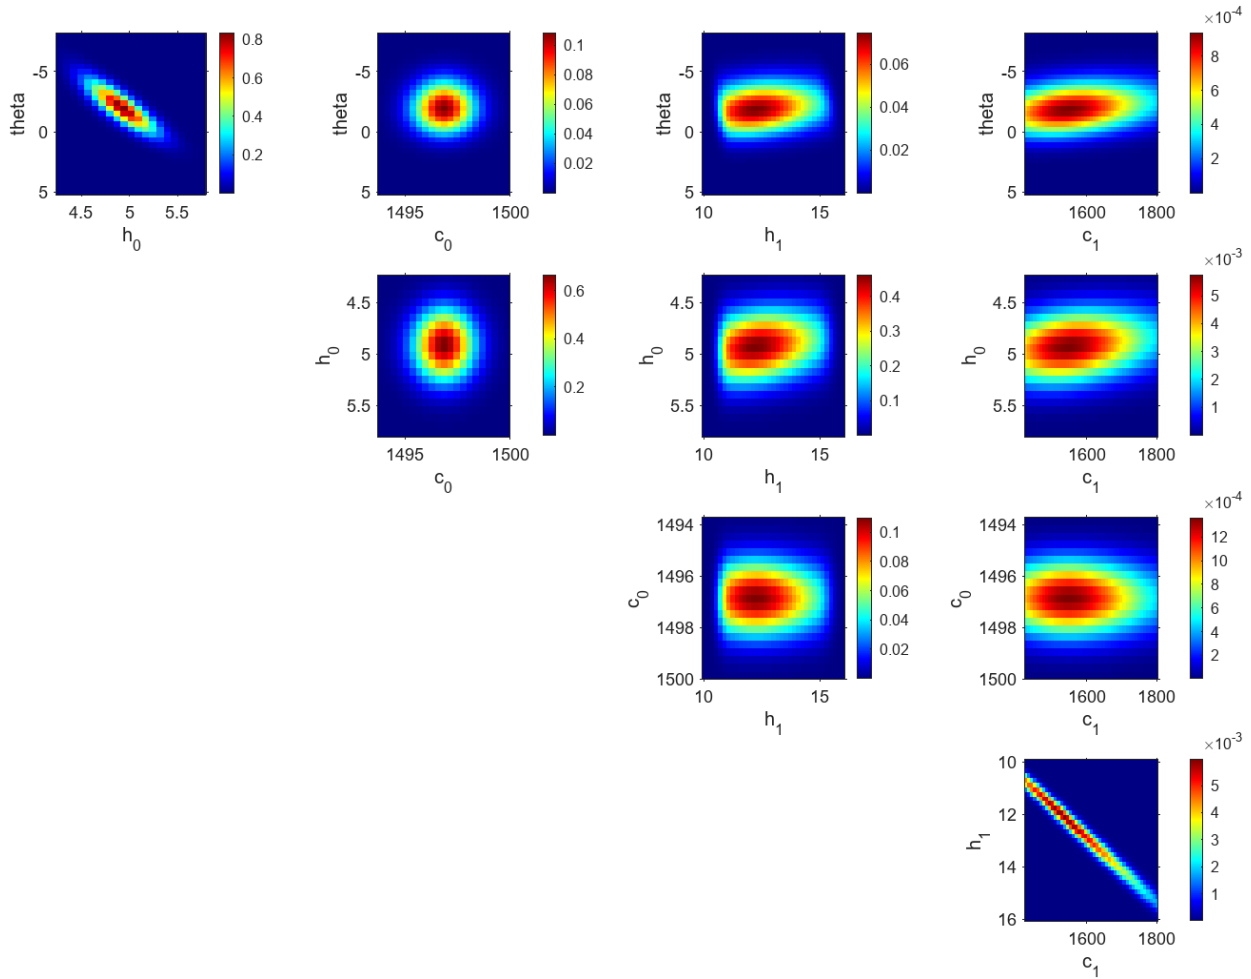


Figure 5.5: The joint probability distributions between the parameters in the geoacoustic inversion.

Most of the parameters have little to no correlation with other parameters. However, the h_1 and c_1 parameters which correspond to the layer thickness and layer sound speed show an extremely sloped joint distribution. This indicates that these two variables are highly coupled. The joint probability distribution shows that a shallower layer height with a slower sound speed is just as likely as a deeper layer height with a faster sound speed.

Figure 5.6 shows the likelihood distributions over the AUV path for the five parameters. The array tilt, water sound speed, and water height all have narrow maximum likelihood regions that indicate that the model can estimate these parameters accurately. The sub-bottom sound speed and sub-bottom layer thickness distributions however are far more

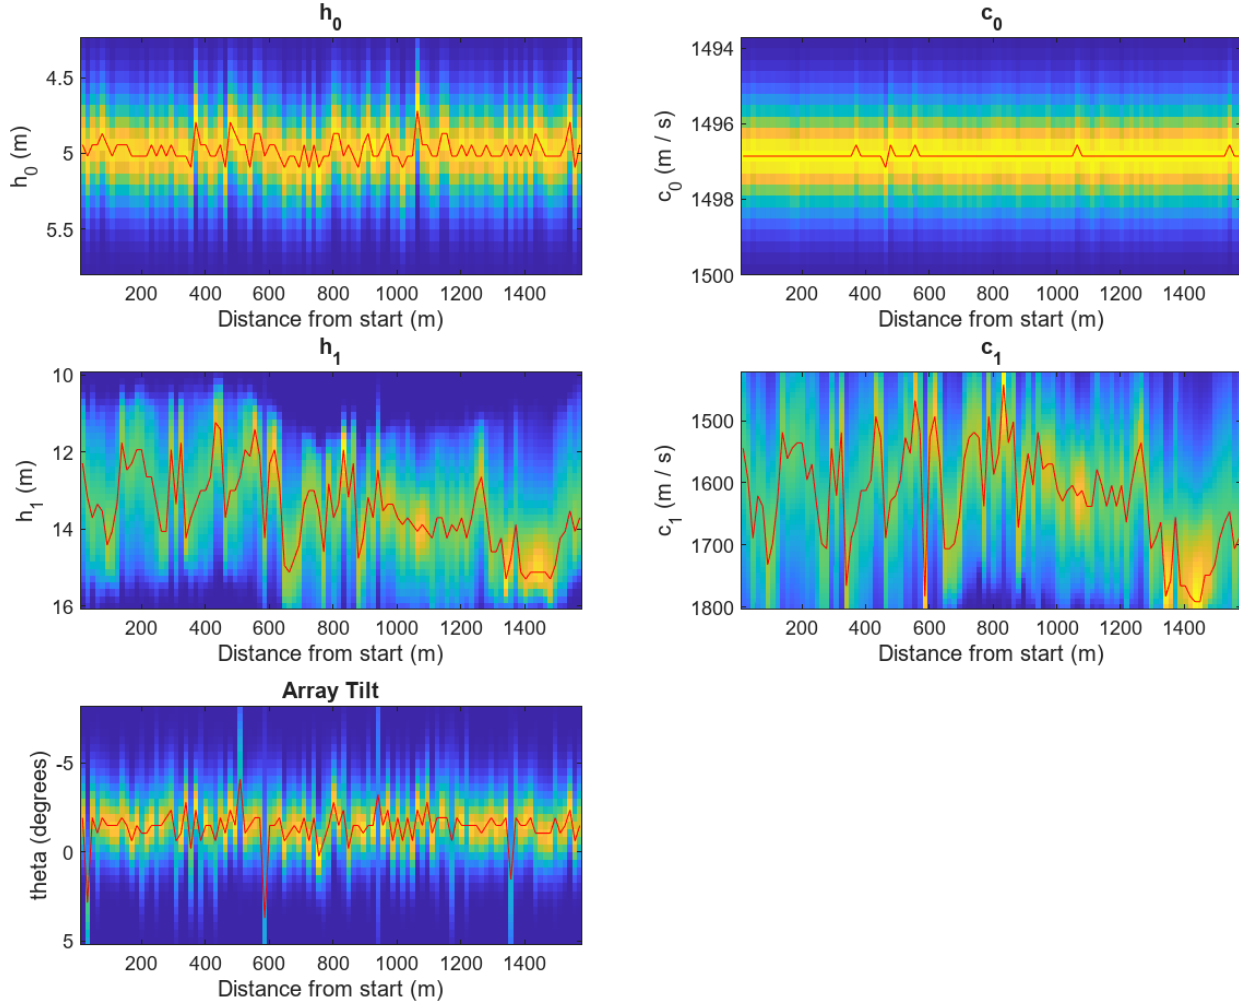


Figure 5.6: The posterior probability distributions as a function of distance traveled for the five parameters in the geoacoustic inversion. The red line shows the most likely value at each distance

ambiguous. The most likely cause of this is the large amount of scattering seen in the sub-bottom layer that is causing a high variance in the sub-bottom return times.

Studies vary on the accepted ground truth in the area. Direct measurements indicate a sediment sound speed to water sound speed ratio of between 0.96 and 1 and a sediment thickness of approximately 10 m [6]. The results from other geoacoustic inversions from additional experiments in the area are shown in Table 5.1. While the results vary, they generally conclude that the sound speed is less than the sound speed of the water column.

The most likely sub-bottom sound speeds and layer thickness are shown as a function

Reference	Mud Layer Compressional Wave Speed	Mud Layer Thickness
Bonnel et al. [66]	1464 – 1540 m/s (top to bottom)	min: 6.0 m mean: 8.0 m max: 11.1 m
Wan et al. [67]	$c_i = 1452 m/s$ (top) with gradient of $s_i = 9 s^{-1}$ with Airy phase $c_i = 1434 m/s$ (top) with gradient of $s_i = 2 s^{-1}$ without Airy phase	Thickness = $c_i \cdot \left(\frac{2000}{t_{i,j}} - \frac{s_i}{2} \right)^{-1}$ where $t_{i,j}$ is the two way travel time
Wan et al. [68]	$c_i = 1468 m/s$ (top) with a gradient of $s_i = 1.3 s^{-1}$	Same equation as above
Rajan and Lin [38]	Bayesian Inversion: 1375 m/s (top), 1595 m/s (bottom) Linear Inverse Case 2: 1394 m/s (top) \pm 5.57, 1578 m/s (bottom) Linear Inverse Case 3: 1411 m/s (top 5 meters) \pm 5.58, 1544.6 m/s (bottom 5 meters) \pm 5.07	Assumed 10 m
Timing Based Geoacoustic Inversion	1424 m/s to 1800 m/s	11.23 m - 15.29 m
Timing Based Geoacoustic Inversion with Gaussian Prior on Mud Thickness	1342 m/s to 1480 m/s	10.35 m - 12 m

Table 5.1: Results of other geoacoustic inversion results in the SBCEX2017 Mud Patch Experiment

of AUV travel distance as the red lines in Figure 5.6. As seen in Figure 5.6 the calculated sub-bottom sound speed ranges from 1424 m/s to 1800 m/s and the calculated sub-bottom thickness ranges from 11.23 m to 15.29 m . These results have a higher sound speed and a thicker sub-bottom than the results seen in Table 5.1. This is likely caused by the high variance seen in the sub-bottom returns. As seen in Figure 5.3, when the variance of the sub-bottom increases, the error in the timing-based inversion also increases. Specifically, the most likely sound speeds are higher than the true sound speed and the most likely layer thickness is thicker than the true layer thickness. As discussed previously the layer thickness and sound speed are highly coupled.

Specifically, the results seen in Figure 5.7 have a sound speed ranging from 1305 m/s to

1404 m/s in the 0 – 400 m of distance traveled which does align with some of the results seen in [38]. As seen previously in 4.5, the travel times do not change with range in the 0 – 400 m distance. The travel times changing could indicate a sloped sub-bottom layer which may cause the slower sound speeds seen in the later parts of Figure 5.7. Adding a sloped sub-bottom to the model could make the inversion results match more closely to previous work.

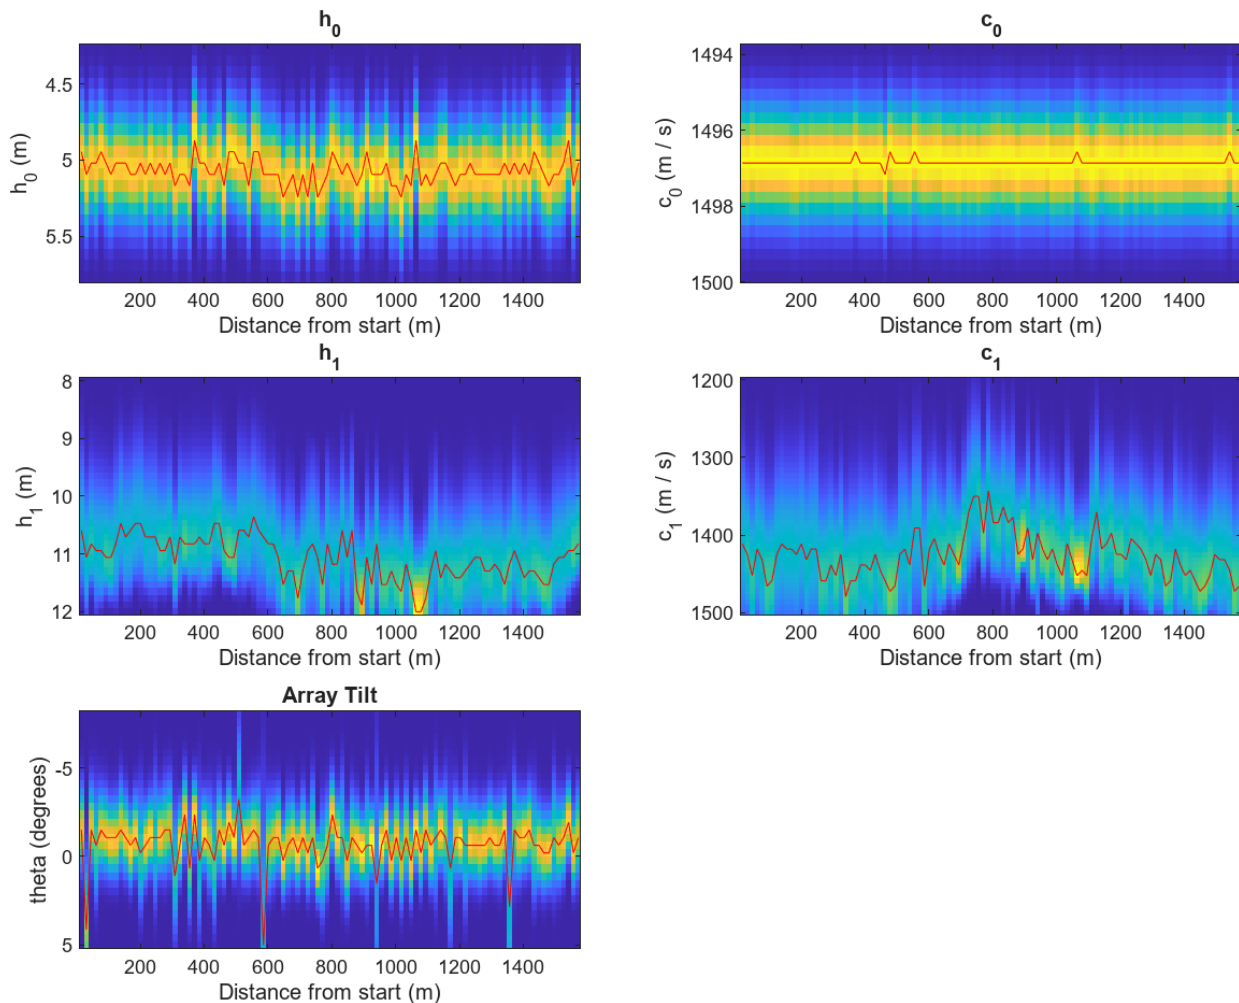


Figure 5.7: The posterior probability distributions as a function of distance traveled for the five parameters in the geoacoustic inversion. **A Gaussian prior was placed on h_1 centered around 10 m with a standard deviation of 1 m .** The red line shows the most likely value at each distance.

If, however, a Gaussian prior centered at 10 m with a standard deviation of 1 m is placed on the sub-bottom layer thickness, the timing-based Geoacoustic inversion results become

less uncertain and start to match the results seen in the other studies in Table 5.1. These results are shown in Figure 5.7. The sub-bottom thickness has been studied using both coring samples and two-way acoustic travel time which showed a thickness of around 10 *m* [5]. With the Gaussian prior, the calculated sub-bottom sound speed ranges from 1342 *m/s* to 1480 *m/s* and the calculated sub-bottom thickness ranges from 10.35 *m* to 12 *m*. These results more closely match the results seen in table 5.1. The posterior probability distributions seen in Figure 5.7 show far less uncertainty in the sub-bottom sound speed and sub-bottom layer thickness than the results seen in Figure 5.6 where has a uniform prior on the sub-bottom thickness.

5.3 Summary

This chapter presented the main results of this thesis, which consist of using the timing returns from an AUV equipped with a source and a towed hydrophone array. The timing returns from the sub-bottom layer exhibit high levels of scattering which was quantified in Chapter 4 using Gaussian process regression. With numerical simulations, this chapter showed that timing-based Bayesian geoacoustic inversion is capable of uniquely determining the sub-bottom layer thickness and sound speed when the variance in the sub-bottom timing returns is low. As the variance of the sub-bottom timing returns increases, the timing-based inversion becomes more uncertain. This chapter then applied the timing-based geoacoustic inversion to an experimental data collected at the New England Mud Patch. With a uniform prior on the layer thickness, the inversion results were highly uncertain and not consistent with prior studies of the area. When a Gaussian distribution was placed on the sub-bottom thickness, the inversion results matched prior studies and the inversion results were less uncertain. With the Gaussian prior, the timing-based inversion shows that the sub-bottom layer has range dependent sound speed and layer thickness. In the following chapter, the contributions of this thesis are summarized and potential future research topics are discussed.

Chapter 6

Conclusion

Sub-bottom profiling using AUVs equipped with sources and towed arrays has the potential to provide detailed ocean bottom information over wide regions. The ability of an AUV to measure the bottom acoustic response over a wide range of angles with minimal interference from the water column makes it a powerful tool for sub-bottom profiling. Many of the common geoacoustic inversion techniques require complex models and are computationally expensive. Timing-based geoacoustic inversion only requires a simple model that can be computed quickly which allows for faster processing on common computer hardware.

This thesis focused on using an AUV equipped with a source and a towed array to estimate the sound speed and layer thickness of the sub-bottom. Chapter 1 discussed the two experiments used in this thesis and sets up the problem formulation used in the remainder of this thesis. This thesis presented an approach of using timing returns from different travel paths to determine the sub-bottom properties. Chapter 3 focused on the use of beamforming and introduces near-field sub-array focusing to enhance the sub-bottom bounce returns. Chapter 4 explored scattering seen in the sub-bottom layer and the use of Gaussian Process Regression to determine the variance of the sub-bottom returns. Chapter 5 discussed timing-based Bayesian Geoacoustic inversion using the results from the Gaussian Process Regression.

6.1 Thesis Contributions

This thesis makes three different contributions: improving near-field bottom returns using sub-array focusing, estimating sub-bottom variation using Gaussian Process Regression, and determining sub-bottom properties using timing-based geoacoustic inversion. The specific elements of these contributions are reviewed here.

6.1.1 Sub-Array Focusing

This novel approach to beamforming allows for near-field bottom and sub-bottom beamforming using a towed hydrophone array. Operating an AUV in the near field of the bottom and sub-bottom layers minimizes the interference from the water column and allows for clearer surveying of the ocean floor. However, when the AUV is close to the bottom, the sub-bottom reflections are then frequently in the near field of the towed array. Reducing the size of the towed array by splitting it up into smaller sub-arrays allows plane wave beamforming to be performed on the sub-arrays. The beamforming returns from the different sub-arrays can then be combined to improve the SNR. In regions with high scattering in the sub-bottom, the beamforming results do not remove the scattering and therefore cannot be used to estimate the timing returns of the sub-bottom.

6.1.2 Scattering Estimation using Gaussian Process Regression

Estimating the variance in the data is important for the accurate estimation of uncertainty of geoacoustic inversion results [36]. This thesis shows that Gaussian Process Regression can be used to quantify the scattering and estimate the uncertainty in the timing returns of the sub-bottom and bottom paths.

The returns from the Gaussian Process Regression are then used to determine the range dependence of the timing returns for the mud patch data and conclude that while the timing returns from the mud patch do change with distance, the mud patch can be considered to be

range-independent over short distances within the margin of uncertainty of the geoacoustic inversion.

6.1.3 Timing Based Geoacoustic Inversion

This thesis develops a method of using the timing return data from different sound paths to determine the sub-bottom properties and the array element positions. The timing returns result in accurate sub-bottom property estimates when the variance of the sub-bottom timing data is low. As the variance of the sub-bottom timing data increases, the sub-bottom inversion results in sound speeds that are too high and sub-bottom thicknesses that are too thick. This thesis analyzes the timing-based inversion scheme using both simulated data and data from the mud patch experiment. This thesis also concludes that when a prior is added to the sub-bottom thickness, the uncertainty in the estimation decreases and the timing based inversion results agree with other inversion results in the same region.

6.2 Future Work

There are a few different directions that could extend the research presented in this thesis. These include modeling the sound speed profile as a linear profile instead of a constant sound speed in the sub-bottom layer, using timing-based geoacoustic inversions on sediments with more than one sub-bottom return, running further simulations to determine whether a towed array with different geometry would reduce the error in the geoacoustic inversion, and characterizing the ideal array length for different sub-bottom variance levels.

6.2.1 Multiple Sub-Bottom Layer Inversion

In theory, the timing-based geoacoustic inversion scheme can be extended to work with multiple sub-bottom layers. Each additional layer would give more information about the preceding layers which may reduce the uncertainty in the geoacoustic inversion in those layers.

The mud patch experiment only shows one sub-bottom return signal, however other data sets such as the shelf break experiment show multiple sub-bottom layers. Each additional layer would add two extra parameters to the inversion scheme (sound speed and layer thickness) which would increase computational complexity. However, the timing-based inversion scheme has the ability and potential to be extended to many layers.

6.2.2 Alternative Bottom Sound Speed Profiles

This thesis assumes that the sound speed in the sub-bottom layer is constant. Some studies suggest that the mud patch sub-bottom is composed of two layers: one mud layer with a constant sound speed, and a second layer with a linearly increasing sound speed at the mud base [69]. Other studies assume that there is a constant sound speed gradient in this layer [67]. Adding additional slope parameters to account for changes occurring within the sediment may increase the accuracy of the inversion and remove the inversion bias that causes sound speed estimates to be too high when the variance in the sub-bottom returns is high.

6.2.3 Evaluating the Towed Array Geometry

Timing-based inversion operates on the premise that there is enough angle variation between the first and last hydrophone to uniquely determine the sub-bottom thickness and sub-bottom sound speed. When working with timing data, these two values are inherently linked. A sound particle propagates through a thicker layer with a faster sound speed in the same amount of time that it propagates through a thinner layer with a slower sound speed.

Through simulation, the geometry of the towed array could be evaluated to determine how the geometry of the array influences the uncertainty in the geoacoustic inversion. These simulated results could then be used to inform the design of future AUV towed arrays for sub-bottom profiling and determine the minimum angle range required to reduce the uncertainty in the inversion estimation and remove the bias.

6.3 Final Thoughts

The study of the seafloor using acoustics has a rich history, with its roots tracing back to early observations about sound propagation in water. By the early 20th century, acoustic techniques began to find practical applications, such as depth sounding and submarine detection, marking the start of a transformative era in underwater exploration. Advances like multibeam sonar and sub-bottom profiling have since become essential tools for mapping and studying the seafloor. This thesis focuses on understanding the seafloor and how the seafloor changes using an AUV with a source and a towed array. By using advancements from other domains, such as Gaussian Process Regression from the machine learning community, acoustic seafloor mapping can evolve into an even more powerful tool for understanding Earth's underwater environments. These innovations promise to make significant contributions to environmental monitoring, resource management, and our broader understanding of the marine world.

Appendix A

Additional Signal Processing

A.1 Pulse Compression

Pulse compression is a type of matched filtering that is used to increase the SNR of known broadband signals and is commonly used in active sonar and radar systems [70]. In noisy ocean environments, pulse compression is often critical to raise the signal level sufficiently above the noise floor to do additional processing. Pulse compression raises the signal level without distorting the range measurements. In this system, the AUV is transmitting a linear chirp with bandwidth Δ and duration T centered at frequency f_0 . The envelope of the output response is a sinc with height \sqrt{D} and half pulse width of $1/\Delta$.

The half pulse width is $1/\Delta$ and the amplitude increase is \sqrt{D} where $D = T\Delta$.

The first step of pulse compression is to create a replica signal of the transmitted chip. The second step is to convolve the replica signal with the data. This functions as a type of matched filter.

In practice, the convolution step is completed in the frequency domain to decrease the computational complexity. The Fourier transform is taken of both the replica and the received signal. Convolution in the time domain is equivalent to multiplication in the frequency domain. In the frequency domain, the received signal is then multiplied by the complex conjugate of

the replica. The inverse Fourier transform is then taken of the resulting signal [71].

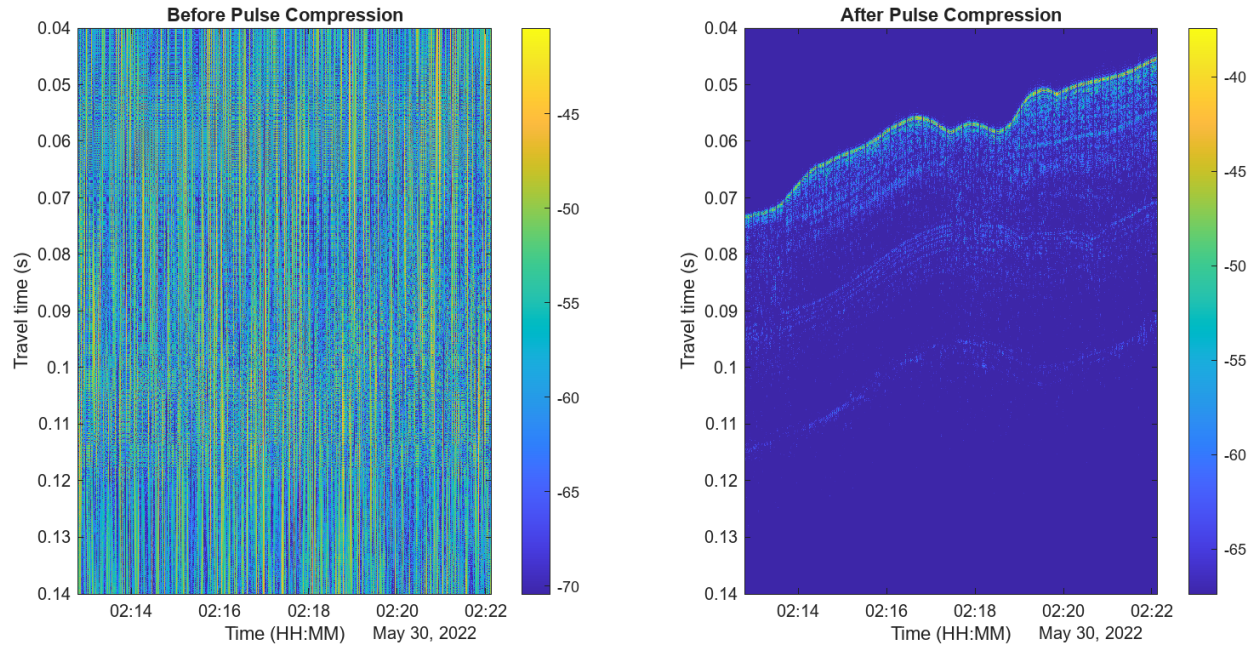


Figure A.1: The shelf break data before and after pulse compression

Figure A.1 shows the waveform before and after pulse compression. The data before pulse compression is not visible because of the noise level.

A.2 Peak Detection

Peak detection is the process of determining when the direct path, or end-fire, signal arrives at each hydrophone. This is used later to indicate the start of each pulse as well as to do both sampling rate calibration and hydrophone position calibration. After pulse compression, the direct path signals have a large SNR because they are not encountering any scatterers or boundary layers that would distort or weaken the signal.

The peak detection is a relatively simple process that involves separating the data into segments of length T where T is the repetitive period of the source and then finding the location of the maximum signal level in each of those segments. When the AUV is operating in a straight line at a constant depth, the direct path signals arrive at a constant rate that

varies only with the current sound speed of the water and slight movements of the hydrophone positions. Outliers where the amplitude of the peak is below a threshold value are then removed. The outlier detection removes the peaks in segments where the source was not transmitting.

A.3 Sampling Rate Calibration

The sampling rate of the hydrophone array needs to be calibrated before any processing can occur. This step is necessary because the clock of the computer recording the array data drifts between missions. The nominal sampling rate is around 20kHz, however in practice there is a $\pm 10\%$ difference between the nominal sampling rate and the actual sampling rate. The source on board the AUV runs off of a chip-scale atomic clock (CSAC) which has a very low drift rate and can be a trusted source of timing measurements.

To calibrate the sampling rate, a section of data is needed where the AUV is operating in a straight line at a constant depth. In those operating conditions, the towed array can be assumed to be a stable linear array and the propagation time of the direct path signal should be relatively constant. An initial sampling rate of 20 kHz is assumed, and the data is broken into segments of length T based on the initial sampling rate where T is the repetitive period of the source. Peak detection is then run on these segments. If the sampling rate is correct, all of the peaks will be in the same position in the repetitive period. If the sampling rate is incorrect, there will be a linear slope of the peak location in the repetitive period over the duration of the data section. This slope can be found with simple linear regression, and the sampling rate is updated based on the slope of the peaks using gradient descent. This process is then repeated until the slope of the peaks is zero. This process only needs to be completed once per mission. Once the sampling rate calibration is completed, the pulse compression and peak detection processes are redone with the corrected sampling rates.

A.4 Hydrophone Distance Calibration

The positions of the hydrophones also need to be calibrated. Before deployment, the positions of the hydrophones and the distance from the source to the first hydrophone are measured using a measuring tape. There is the potential for human error during this measurement process, so the distances are checked using the peak arrival times. To complete this process, a segment of the mission where the AUV is moving in a straight line at a stable depth is used. Here the towed array is assumed to be running in a straight line behind the AUV. Using the end-fire timing equation from Chapter 1 the positions of the hydrophones are calculated and then averaged over the duration of the straight segment. The speed of the AUV is negligible compared to the speed of sound. The speed of the AUV is approximately 0.1% the speed of sound.

Figure A.2 shows the calculated hydrophone positions with a histogram of those positions changing over time. The hydrophone positions are relatively stable, and the hydrophones further from the source have more variation in position than the ones closer to the source.

A.5 Rough Layer Detection

Once the signals have been improved through pulse compression the data is roughly labeled by hand to indicate the start and end times of each layer return for each channel. This is a relatively quick process and does not require high precision. However, due to the closeness of some of the layer returns, each individual layer was denoted by hand before proceeding to the fine-tuned detection. This process is not the main part of my thesis contribution and was not automated due to time constraints. However, in the future, this could become an automatic process with image processing techniques.

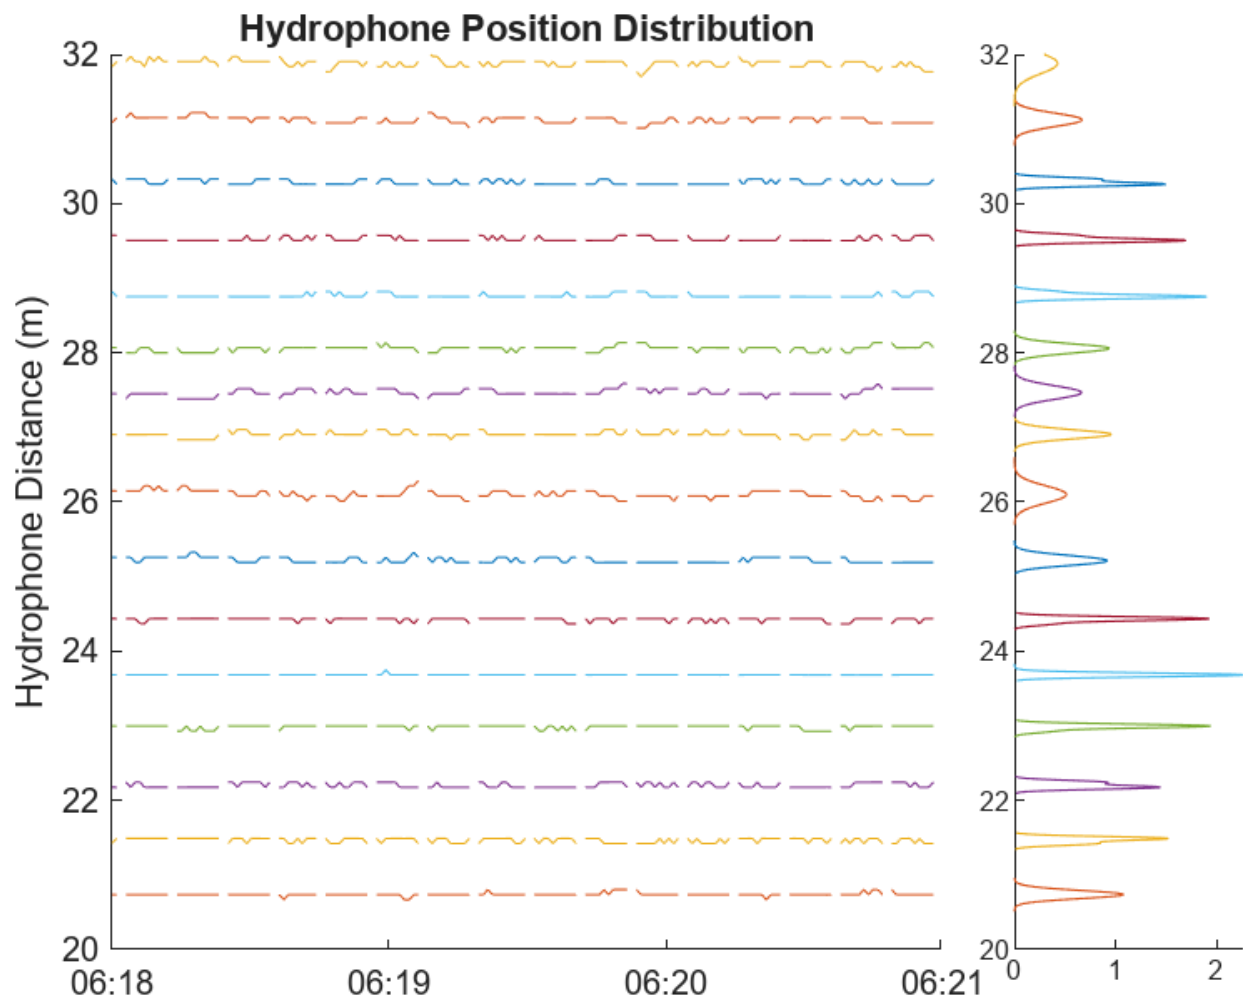


Figure A.2: The calibrated hydrophone positions over time with histograms showing the variation in position.

Appendix B

Bottom Interface Angle Derivation

The goal of this derivation is to determine the crossing point γ between the water and sediment along the x axis. The total travel time of the sub-bottom path is:

$$t_{total} = t_{water} + t_{sediment},$$

where t_{water} is the travel time in the water

$$t_{water} = \sqrt{\frac{(2h_0 - z)^2 + (2\gamma)^2}{c_0^2}},$$

and where $t_{sediment}$ is the travel time in the sediment

$$t_{sediment} = \sqrt{\frac{(2h_1)^2 + (x_i - 2\gamma)^2}{c_1^2}}.$$

Snell's Law is based on the principle that sound always takes the fastest path [9], so γ corresponds to the point where t_{total} is at a minimum. In order to determine this, the derivative of t_{total} with respect to γ is taken and set equal to zero:

$$\frac{\delta t_{total}}{\delta \gamma} = \frac{2\gamma}{c_0^2 \sqrt{\frac{(2h_0 - z)^2 + (2\gamma)^2}{c_0^2}}} - \frac{x_i - 2\gamma}{c_1^2 \sqrt{\frac{(2h_1)^2 + (x_i - 2\gamma)^2}{c_1^2}}} = 0.$$

The zero crossing of the above equation is γ that corresponds to the lowest travel time and therefore is the boundary crossing between the water and sediment layers. Once γ has been found, finding $\theta_{sub-bottom_{0,i}}$ and $\theta_{sub-bottom_{1,i}}$ are trivial.

$$\theta_{sub-bottom_{0,i}} = \arctan\left(\frac{h_0}{\gamma}\right)$$

$$\theta_{sub-bottom_{1,i}} = \arctan\left(\frac{h_1}{\frac{x}{2} - \gamma}\right)$$

References

- [1] US Department of Commerce and National Oceanic. *Seafloor mapping*. <https://oceanexplorer.noaa.gov/explainers/mapping.html>. Accessed: 2025-1-15.
- [2] B. T. Hefner, D. R. Jackson, K. L. Williams, and E. I. Thorsos. “Mid- to high-frequency acoustic penetration and propagation measurements in a sandy sediment”. In: *IEEE J. Ocean. Eng.* 34.4 (Oct. 2009), pp. 372–387.
- [3] J. A. Goff, A. H. Reed, G. Gawarkiewicz, P. S. Wilson, and D. P. Knobles. “Stratigraphic analysis of a sediment pond within the New England Mud Patch: New constraints from high-resolution chirp acoustic reflection data”. In: *Mar. Geol.* 412 (June 2019), pp. 81–94.
- [4] J. D. Chaytor, M. S. Ballard, B. J. Buczkowski, J. A. Goff, K. M. Lee, A. H. Reed, and A. A. Boggess. “Measurements of geologic characteristics and geophysical properties of sediments from the New England mud patch”. In: *IEEE J. Ocean. Eng.* 47.3 (July 2022), pp. 503–530.
- [5] P. S. Wilson, D. P. Knobles, and T. B. Neilsen. “Guest editorial an overview of the seabed characterization experiment”. In: *IEEE J. Ocean. Eng.* 45.1 (Jan. 2020), pp. 1–13.
- [6] M. S. Ballard, D. D. Garcia, K. M. Lee, G. R. Venegas, A. R. McNeese, P. S. Wilson, and J. D. Chaytor. “Direct measurements of sediment geoaoustic properties in the

- New England Mud Patch and shelf breaka)”. In: *J. Acoust. Soc. Am.* 156.4 (Oct. 2024), pp. 2889–2902.
- [7] P. S. Wilson, D. P. Knobles, and T. B. Neilsen. “Guest editorial: Continued exploration of fine-grained sediments from SBCEX2017”. In: *IEEE J. Ocean. Eng.* 47.3 (July 2022), pp. 497–502.
- [8] L. Paull, S. Saeedi, M. Seto, and H. Li. “AUV Navigation and Localization: A Review”. In: *IEEE J. Ocean. Eng.* 39.1 (Jan. 2014), pp. 131–149.
- [9] F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt. *Computational Ocean Acoustics*. Springer New York, 2011.
- [10] *Technical Guides - Speed of sound in sea water - Underlying Physics*. <http://resource.npl.co.uk/acoustics/techguides/soundseawater/underlying-phys.html>. Accessed: 2025-1-9.
- [11] A.-L. S. Lapinskia and D. M. F. Chapman. “The effects of ignored seabed variability in geoacoustic inversion”. In: *J. Acoust. Soc. Am.* 117.6 (June 2005), pp. 3524–3538.
- [12] C. W. Holland, P. L. Nielsen, J. Dettmer, and S. Dosso. “Resolving meso-scale seabed variability using reflection measurements from an autonomous underwater vehicle”. In: *J. Acoust. Soc. Am.* 131.2 (Feb. 2012), pp. 1066–1078.
- [13] S. Pinson, L. Guillon, and C. W. Holland. “Range dependent sediment sound speed profile measurements using the image source method”. In: *J. Acoust. Soc. Am.* 134.1 (July 2013), pp. 156–165.
- [14] S. G. Lemon. “Towed-array history, 1917–2003”. In: *IEEE J. Ocean. Eng.* 29.2 (Apr. 2004), pp. 365–373.
- [15] D. Zhang, B. Zhao, K. Zhu, and H. Jiang. “Dynamic analysis of towed cable with variable length during turning maneuvers”. In: *Sci. Rep.* 13.1 (Mar. 2023), p. 3525.

- [16] V. Pallayil, Y. T. Lin, E. Fischell, A. L. Kukulya, et al. “Integration and Testing of a Low-Profile Hydrophone Array with REMUS 100 AUV for Seabed Characterization and Marine Mammal Detection Application”. In: *2019 IEEE* (2019).
- [17] A. Barbagelata, P. Guerrini, and L. Troiano. “Thirty years of towed arrays at NURC”. In: *Oceanography (Wash. D. C.)* 21.2 (June 2008), pp. 24–33.
- [18] J. F. Lynch, D. Chu, T. C. Austin, W. M. Carey, A. D. Pierce, and J. D. Holmes. “Detection and classification of buried targets and sub-bottom geoacoustic inversion with an AUV carried low frequency acoustic source and a towed array”. In: *OCEANS 2006*. IEEE, Sept. 2006, pp. 1–5.
- [19] J. D. Holmes, W. M. Carey, J. F. Lynch, A. E. Newhall, and A. Kukulya. “An autonomous underwater vehicle towed array for ocean acoustic measurements and inversions”. In: *Europe Oceans 2005*. IEEE, 2005, pp. 1058–1061.
- [20] J. D. Holmes, E. J. Sullivan, and W. M. Carey. “Passive synthetic aperture processing with an autonomous underwater vehicle towed hydrophone array”. In: *OCEANS 2006*. IEEE, Sept. 2006, pp. 1–5.
- [21] T. M. Marston, B. R. Hall, C. Bassett, D. S. Plotnick, and A. N. Kidwell. “Motion tracking of fish and bubble clouds in synthetic aperture sonar data”. In: *J. Acoust. Soc. Am.* 155.3 (Mar. 2024), pp. 2181–2191.
- [22] R. E. Hansen, H. J. Callow, T. O. Sabo, and S. A. V. Synnes. “Challenges in Seafloor Imaging and Mapping With Synthetic Aperture Sonar”. In: *IEEE Trans. Geosci. Remote Sens.* 49.10 (Oct. 2011), pp. 3677–3687.
- [23] J. Carballini and F. Viana. “Using synthetic aperture sonar as an effective tool for pipeline inspection survey projects”. In: *2015 IEEE/OES Acoustics in Underwater Geosciences Symposium (RIO Acoustics)*. IEEE, July 2015, pp. 1–5.

- [24] D. Marx, M. Nelson, E. Chang, W. Gillespie, A. Putney, and K. Warman. “An introduction to synthetic aperture sonar”. In: *Proceedings of the Tenth IEEE Workshop on Statistical Signal and Array Processing (Cat. No.00TH8496)*. IEEE, 2000, pp. 717–721.
- [25] J. L. Odom and J. L. Krolik. “Passive towed array shape estimation using heading and acoustic data”. In: *IEEE J. Ocean. Eng.* 40.2 (Apr. 2015), pp. 465–474.
- [26] A. Xenaki, B. Gips, and Y. Pailhas. “Unsupervised learning of platform motion in synthetic aperture sonar”. In: *J. Acoust. Soc. Am.* 151.2 (Feb. 2022), p. 1104.
- [27] S.-M. Steele and A. P. Lyons. “An experimental test of endfire synthetic aperture sonar for sediment characterisation”. In: *IET Radar Sonar Navig.* 18.11 (Nov. 2024), pp. 2057–2065.
- [28] S.-M. Steele and A. P. Lyons. “Development and experimental validation of endfire synthetic aperture sonar for sediment acoustics studies”. In: *IEEE J. Ocean. Eng.* 47.2 (Apr. 2022), pp. 472–482.
- [29] C. Sun, H. Shen, X. Qiu, and Y. Liu. “Application of near-field beam-forming in large-aperture optical towed linear array”. In: *2016 IEEE/OES China Ocean Acoustics (COA)*. IEEE, Jan. 2016, pp. 1–5.
- [30] H. Lu et al. “A tutorial on near-field XL-MIMO communications towards 6G”. In: *IEEE Commun. Surv. Tutor.* 26.4 (2024), pp. 2213–2257.
- [31] H. Lee, J. Ahn, Y. Kim, and J. Chung. “Direction-of-arrival estimation of far-field sources under near-field interferences in passive sonar array”. In: *IEEE Access* 9 (2021), pp. 28413–28420.
- [32] K. Wei, S. Fang, and J. Tao. “Enhanced near-field interference suppression scheme for the non-cooperative underwater acoustic pulse detection of the towed linear array”. In: *J. Mar. Sci. Eng.* 10.2 (Feb. 2022), p. 250.

- [33] N. P. Chotiros and V. Pallayil. “Seabed Characterization Using Acoustic Communication Signals on an Autonomous Underwater Vehicle With a Thin-Line Towed Array”. In: *IEEE J. Oceanic Eng.* 38.3 (July 2013), pp. 410–418.
- [34] Q. Zhou, X. Li, B. Huang, L. Liu, S. Gao, H. Zhou, J. Liu, B. Liu, and C. Zhang. “Inversion of the physical properties of seafloor surface sediments based on AUV sub-bottom profile data in the northern slope of the South China Sea”. In: *Sci. Rep.* 11.1 (Mar. 2021), p. 6539.
- [35] G. V. Frisk. *Ocean and Seabed Acoustics: A Theory of Wave Propagation*. Pearson Education, Jan. 1994.
- [36] S. E. Dosso. “Quantifying uncertainty in geoacoustic inversion. I. A fast Gibbs sampler approach”. In: *J. Acoust. Soc. Am.* 111.1 Pt 1 (Jan. 2002), pp. 129–142.
- [37] N. R. Chapman and E. C. Shang. “Review of Geoacoustic Inversion in Underwater Acoustics”. In: *J. Theor. Comp. Acoust.* 29.03 (Sept. 2021), p. 2130004.
- [38] S. D. Rajan and Y. T. Lin. “Broadband Geoacoustic Inversions for Seabed Characterization of the New England Mud Patch”. In: *IEEE J. Oceanic Eng.* 48.2 (2022), pp. 264–276.
- [39] C. Gervaise, B. G. Kinda, J. Bonnel, Y. Stéphan, and S. Vallez. “Passive geoacoustic inversion with a single hydrophone using broadband ship noise”. In: *J. Acoust. Soc. Am.* 131.3 (Mar. 2012), pp. 1999–2010.
- [40] J. Dettmer, S. E. Dosso, and C. W. Holland. “Trans-dimensional geoacoustic inversion”. In: *J. Acoust. Soc. Am.* 128.6 (Dec. 2010), pp. 3393–3405.
- [41] L. Guillon, S. E. Dosso, N. Ross Chapman, and A. Drira. “Bayesian Geoacoustic Inversion With the Image Source Method”. In: *IEEE J. Oceanic Eng.* 41.4 (Oct. 2016), pp. 1035–1044.
- [42] C. H. Dix. “Seismic velocities from surface measurements”. In: *Geophysics* 20.1 (Jan. 1955), pp. 68–86.

- [43] G. M. Bryan. “Sonobuoy measurements in thin layers”. In: *Physics of Sound in Marine Sediments*. Boston, MA: Springer US, 1974, pp. 119–130.
- [44] G. M. Bryan. “The hydrophone–pinger experiment”. In: *J. Acoust. Soc. Am.* 68.5 (Nov. 1980), pp. 1403–1408.
- [45] J. Dettmer, S. E. Dosso, and C. W. Holland. “Uncertainty estimation in seismo-acoustic reflection travel time inversion”. In: *J. Acoust. Soc. Am.* 122.1 (July 2007), pp. 161–176.
- [46] A. Drira, L. Guillon, and A.-O. Boudraa. “Image source detection for geoacoustic inversion by the Teager-Kaiser energy operator”. In: *J. Acoust. Soc. Am.* 135.6 (June 2014), EL258–64.
- [47] S. Pinson, L. Guillon, and C. W. Holland. “The image-source method: a tool for geoacoustic inversion”. In: *Near Surf. Geophys.* 18.1 (Feb. 2020), pp. 39–48.
- [48] P. S. Wilson, D. P. Knobles, P. H. Dahl, A. R. McNeese, and M. C. Zeh. “Short-range signatures of explosive sounds in Shallow Water used for seabed characterization”. In: *IEEE J. Ocean. Eng.* 45.1 (Jan. 2020), pp. 14–25.
- [49] 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”. In: *IEEE J. Ocean. Eng.* 45.1 (Jan. 2020), pp. 26–38.
- [50] D. R. Barclay, D. A. Bevans, and M. J. Buckingham. “Estimation of the geoacoustic properties of the New England mud patch from the vertical coherence of the ambient noise in the water column”. In: *IEEE J. Ocean. Eng.* 45.1 (Jan. 2020), pp. 51–59.
- [51] J. Bonnel, S. E. Dosso, D. Eleftherakis, and N. R. Chapman. “Trans-dimensional inversion of modal dispersion data on the New England mud patch”. In: *IEEE J. Ocean. Eng.* 45.1 (Jan. 2020), pp. 116–130.
- [52] Z.-H. Michalopoulou and P. Gerstoft. “Inversion in an uncertain ocean using Gaussian processes”. In: *J. Acoust. Soc. Am.* 153.3 (Mar. 2023), pp. 1600–1611.

- [53] W. F. Jenkins and P. Gerstoft. “Bayesian optimization with Gaussian processes for robust localization”. In: *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, Apr. 2024, pp. 6010–6014.
- [54] I. D. Khurjekar, P. Gerstoft, C. F. Mecklenbräuker, and Z.-H. Michalopoulou. “Direction-of-arrival estimation using Gaussian process interpolation”. In: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, June 2023, pp. 1–5.
- [55] D. Caviedes-Nozal, N. A. B. Riis, F. M. Heuchel, J. Brunskog, P. Gerstoft, and E. Fernandez-Grande. “Gaussian processes for sound field reconstruction”. In: *J. Acoust. Soc. Am.* 149.2 (Feb. 2021), pp. 1107–1119.
- [56] H. L. Van Trees. *Optimum array processing: Part IV of detection, estimation, and modulation theory*. Detection, Estimation, and Modulation Theory. Nashville, TN: John Wiley & Sons, Mar. 2002.
- [57] E. Bjornson, O. T. Demir, and L. Sanguinetti. “A primer on near-field beamforming for arrays and reconfigurable intelligent surfaces”. In: *2021 55th Asilomar Conference on Signals, Systems, and Computers*. IEEE, Oct. 2021, pp. 105–112.
- [58] M. Franceschetti. *Wave theory of information*. Cambridge, England: Cambridge University Press, Nov. 2017.
- [59] K. Sabra, P. Pfenninger, and Y.-T. Lin. “Through-the-sensor sub-bottom imaging using the self-noise of an autonomous underwater vehicle”. In: *J. Acoust. Soc. Am.* In Press, JASA–EL–02695R1.
- [60] C. W. Holland, C. M. Smith, Z. Lowe, and J. Dorminy. “Seabed observations at the New England mud patch: Reflection and scattering measurements and direct geoacoustic information”. In: *IEEE J. Ocean. Eng.* 47.3 (July 2022), pp. 578–593.
- [61] J. Wang. “An intuitive tutorial to Gaussian process regression”. In: *Comput. Sci. Eng.* 25.4 (2023), pp. 4–11.

- [62] C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. Adaptive Computation and Machine Learning series. London, England: MIT Press, Nov. 2005.
- [63] E. V. Bonilla, K. M. Chai, and C. K. I. Williams. “Multi-task Gaussian Process Prediction”. In: *Neural Inf Process Syst 20* (Dec. 2007). Ed. by J. Platt, D. Koller, Y. Singer, and S. Roweis, pp. 153–160.
- [64] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson. “GPYtorch: Blackbox Matrix-matrix Gaussian process inference with GPU acceleration”. In: *CoRR* abs/1809.11165 (Oct. 2018), p. 1809.11165.
- [65] Y.-T. Lin, J. Bonnel, D. P. Knobles, and P. S. Wilson. “Broadband waveform geoacoustic inversions with absolute travel time”. In: *IEEE J. Ocean. Eng.* 45.1 (Jan. 2020), pp. 174–188.
- [66] J. Bonnel, Y.-T. Lin, D. Eleftherakis, J. A. Goff, S. Dosso, R. Chapman, J. H. Miller, and G. R. Potty. “Geoacoustic inversion on the New England Mud Patch using warping and dispersion curves of high-order modes”. In: *J. Acoust. Soc. Am.* 143.5 (May 2018), EL405–EL411.
- [67] L. Wan, M. Badiy, D. P. Knobles, and P. S. Wilson. “The Airy phase of explosive sounds in shallow water”. In: *J. Acoust. Soc. Am.* 143.3 (Mar. 2018), EL199–EL205.
- [68] L. Wan, M. Badiy, 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”. In: *IEEE J. Ocean. Eng.* 45.1 (Jan. 2020), pp. 201–211.
- [69] J. Bonnel, S. E. Dosso, J. A. Goff, Y.-T. Lin, J. H. Miller, G. R. Potty, P. S. Wilson, and D. P. Knobles. “Transdimensional geoacoustic inversion using prior information on range-dependent seabed layering”. In: *IEEE J. Ocean. Eng.* 47.3 (July 2022), pp. 594–606.

- [70] J. R. Klauder, A. C. Price, S. Darlington, and W. J. Albersheim. “The theory and design of chirp radars”. In: *Bell Syst. Tech. J.* 39.4 (July 1960), pp. 745–808.
- [71] S.A. Taylor, , J. L. MacArthur. “Digital Pulse Compression Radar Receiver”. In: *Johns Hopkins APL Tech. Dig.* 6.4 (Mar. 1967), pp. 2–10.