1-1-2010

Extraction of Small Boat Harmonic Signatures From Passive Sonar

George Lloyd Ogden
Portland State University

Extraction of Small Boat Harmonic Signatures From Passive Sonar

by

George Lloyd Ogden

A thesis submitted in partial fulfillment of the
requirements for the degree of

Master of Science
in
Electrical and Computer Engineering

Thesis Committee:
Lisa M. Zurk, Chair
Richard L. Campbell
Mark D. Sytsma

Portland State University
2010

Abstract

This thesis investigates the detection and classification of small boats using a passive sonar system. Noise radiated from a small boats consists of broadband noise and harmonically related tones that correspond to parameters in the boats engine and propeller. A novel signal processing method for detection and discrimination of noise radiated from small boats has been developed. There are two main components to the algorithm. The first component detects the presence of small boats by the harmonic tonals radiated from the boat propeller and engine. The second component was designed to extract the a signature from passive sonar data.

The Harmonic Extraction and Analysis Tool (HEAT) was designed to estimate the fundamental frequency of the harmonic tones, track the fundamental frequency using a Kalman filter, and automatically extract the amplitudes of the harmonic tonals to generate a harmonic signature for the boat. The algorithm is shown to accurately extract theses signatures, and results show that the signatures are unique enough that the same boat passing by the hydrophone multiple times can be recognized.

Acknowledgments

I would like to thank my advisor Dr. Lisa Zurk for her guidance during this research. Also I would like to thank Dr. Martin Siderius and the students in the NEAR-Lab for their help and support during this research. In addition, I would like to thank Pacific Northwest National Laboratory for presenting this project to the NEAR-Lab and funding this work. I am also grateful to Dr. Richard Campbell and Dr. Mark Sytsma for serving on my committee. Finally I would like to thank my wife and daughter for their patience and support during this research.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Executive Summary

The automated detection and classification of maritime traffic is a very challenging problem as well as one of great importance to many organizations. For marine protected areas (MPAs), an automated boat detection system could alert authorities of vessel traffic in the area. However, in some MPA's, like in Molokini off the coast of Maui, commercial snorkeling and diving boats are authorized where fishing vessels are not, so a classification or identification system is also needed to discriminate from these different types of boats. The need for similar systems arises in the monitoring of harbor traffic for national security. There are many different methods of detecting boats including radar [1], electro-optic (EO) and infrared (IR) cameras [2], and sonar - both active and passive. However, many of these methods provide little additional information beyond detection. Radar and optical methods are limited by line of sight for detection, and optical systems can be obscured by rain, fog, or may require daylight. Active sonar can be used for detection of quiet targets, but the high level of reverberation in shallow water environments often results in many false detections, limiting its utility. As an alternative, passive sonar has been proven to be an efficient tool for the detection

and identification of self-emitting targets [3, 4].

There has been limited research on the detection and classification of small boats using passive sonar. In [5], passive sonar was used to detect SCUBA divers by the peaks in the frequency energy distribution due to the divers breaths. The breathing rate and the spectrum intensity give information of the range of the diver. In [6], the same research group uses passive sonar to record the spectrum of small boats and investigated the effects of boat noise on the detection range of divers. However, this work is mainly focused on the detection of targets using passive sonar. There has yet to be any significant work on classification of small boats in the literature.

This thesis is focused on the detection and classification of small boats using passive sonar systems. Passive spectra of boats include broadband noise as well as tonals due to the harmonics of the engine speed and shaft/propeller rotation [7]. Using the above features, a novel method for detection and discrimination of boat noise has been developed. The algorithm has two major components; the first component detects the presence of small boats using the harmonic tonals radiated from the boat propeller and engine. The second component extracts the harmonic features and facilitates the exploration of the relationship between these features and the identification of specific boats. These features consist of harmonic amplitudes, SNRs, and the fundamental frequencies of the boat noise.

A Harmonic Extraction and Analysis Tool (HEAT) has been designed to estimate the fundamental frequency of the harmonic content generated by the engine and propeller of small boats. A discrete Kalman filter is applied to refine the estimated fundamental frequency and create a track through time. Harmonics of the fundamental frequency are extracted, and their amplitudes are used as signatures of the boat noise. The algorithm is shown to accurately and automatically extract these harmonic signatures for later use in classification.

The rest of the thesis is organized as follows. Chapter 1 provides a background on passive sonar and reviews the classification of large ships from radiated noise. Previous methods on estimating the fundamental frequency of a set of harmonics are introduced. Also included is as an outline of the proposed method of detection and estimation of small boat signatures. Chapter 2 introduces the acoustic model of the sound radiated from small boats. It also illustrates the pre-processing step which prepares the raw data for the detection algorithms. Chapter 3 details the boat detection algorithm and its performance on a data set collected by the Pacific Northwest National Laboratory (PNNL) in Sequim, WA. Chapter 4 introduces the HEAT algorithm for harmonic feature extraction. The data analysis results are shown in Chapter 5. Chapter 6 summarizes this research and suggests potential areas of future research.

1.2 Passive Sonar

SONAR (SOund Navigation And Ranging) it is a technique that uses acoustic signals for navigation, detection and communication. There are two main methods of sonar, active and passive. Typically in an active sonar system, a short pulse signal is transmitted through a medium and the echo is received by a hydrophone to determine range of a target. In comparison, passive sonar uses a hydrophone to record the sound generated by self-emitting sources.

In this work, passive sonar is used to record sound emitted from small boats due to engine noise and propeller movement. The acoustic signatures are extracted from the received signals to identify the source type. Passive sonar has been chosen for this research because we are recording the sound from moving boats, which are self-emitting sources. Passive sonar recording devices can also be relatively cheap to construct, very simple to deploy, and do not adversely affect the surrounding environment.

A major challenge in the algorithm development is to discriminate boat noise from interfering background noise sources including shipping, environmental and biological noises. The signal processing methods developed in this research are designed to be robust against constant noise sources as well as loud transient events. The hydrophone used to collect data by the PNNL is located at the mouth of a bay, with very strong currents during the changing tides. Objects

would regularly hit the hydrophone, causing loud, impulsive noise. Also located near the hydrophone was an underwater pump, which is a loud source of constant noise. In Chapter 3, a couple different techniques are introduced to help mitigate the effects of these noise sources.

1.3 Noise Radiation From Ships

Many researchers have studied radiated noise from large ships, both modeling and measurement. In the 1970's, Gray and Greely [8] developed a model to predict source level and frequency of the acoustic energy generated by propeller cavitations. In the 1990's Arveson and Vendettis [9] conducted a series of measurements of the noise radiated from the M/V Overseas Harriette and found agreement with Gray and Greeley's model. These references as well as many others characterize the radiated noise from large ships quite well. However, much less work has been done to characterize the radiated noise from small vessels.

Ross[3] and Urick[4] have given an excellent description of radiated noise of large surface ships and submarines. It has been shown that the radiated noise from a ship is a combination of broadband noise and sinusoidal tonal signals. The broadband noise is generated by many sources including propeller cavitations, and impulsive events in the engine such as the impact of a piston against the cylinder wall. This broadband noise propagates through the water and when received on a hydrophone, generates the classical bathtub pattern that is often associated

Table 1.1: Fundamental frequencies from the engine and propeller.

Engine Rates	**Propeller Rates**
Cylinder Firing Rate $f_{CF} = f_{CR}/2$	Shaft Rotation Rate $f_{SR} = f_{CR}/\Lambda_g$ Λ_g = Gear Ratio
Crankshaft Rotation Rate $f_{CR} = RPM/60$ RPM = Engine Speed	Blade Rotation Rate $f_{BR} = N_b f_{SR}$ N_b = Number of Blades
Engine Firing Rate $f_{EF} = N_c f_{CF}$ N_c = Number of Cylinders	

with passive acoustic signatures. This bathtub pattern is do to all the different multi-path arrivals of the noise adding up in and out of phase. The sinusoidal tonal signals can be related to details about the ships engine and propeller, and are the fundamental components of a harmonic set. Table 1.1 shows the major contributions to the tonals from the ships engine and propeller. The model of radiated ship noise represented as a sum of broadband noise and tonal frequencies will be used to describe the noise radiated from small boats.

1.4 Frequency Estimation Methods

Frequency estimation is a topic that spans many disciplines including speech recognition, musical pitch estimation, and biomedical signal processing to name a few. The thing in common with all of these different disciplines is the nature of

the signals they are trying to estimate; multi-harmonic periodic, or even quasi-periodic, signals. The fundamental frequency of these signals contains useful information. Also important is the progression of that fundamental frequency through time.

Frequency estimation and tracking is a complex problem as the frequency of a signal is inherently a non-linear parameter. Tracking non-linear parameters cannot be done with a conventional Kalman filter. In [10], a marginalized particle filter is used to track the instantaneous frequency of two biomedical signals: electrocardiogram and arterial blood pressure. A marginalized particle filter was also used in [11] to track the fundamental frequency, and in this case, multiple fundamental frequencies, of musical signals.

The method used in this thesis is a frequency domain method similar to the maximum likelihood method described in [12]. In their research, a time domain method for estimating the pitch period of voiced speech based on a maximum likelihood formulation. The frequency domain analog to that method is briefly described as matching a comb-like filter to the autocorrelation of the periodic signal. This frequency domain method is also similar to [13] and [14] where the goal is to minimize the difference between a comb filter and the signal itself. In this thesis, the Fourier transform of a signal is correlated to a comb filter to get an estimate of the fundamental frequency. This puts the frequency in a linear space, so a Kalman filter can now be used to track the fundamental frequency

through time.

1.5 Overview of detection and estimation algorithm

The overall goals of this research are to develop a boat detection algorithm that can be implemented real-time on passive acoustic systems, and to extract harmonic-related signatures that can be used to discriminate the boat types. Figure 1.1 shows the structure of the developed approach.

The signal is recorded on a hydrophone and sampled at rate f_s. The sampled signal is converted from time domain to a time-frequency domain (Module A) and passed to the detection algorithm (Module B). Once a ship signal has been identified, the data is passed to the HEAT algorithm which extracts a harmonic signature. The HEAT algorithm estimates the fundamental frequency (Module C) of all the harmonic tonals, then tracks the fundamental frequency f_0 though time using a Kalman filter (Module D). The harmonic signature is then extracted from the data by projecting the track of f_0 onto all the harmonics and estimating the amplitude of all the harmonics (Module E). These harmonic signatures can be used to build up a signal database for use in classification of small vessels.

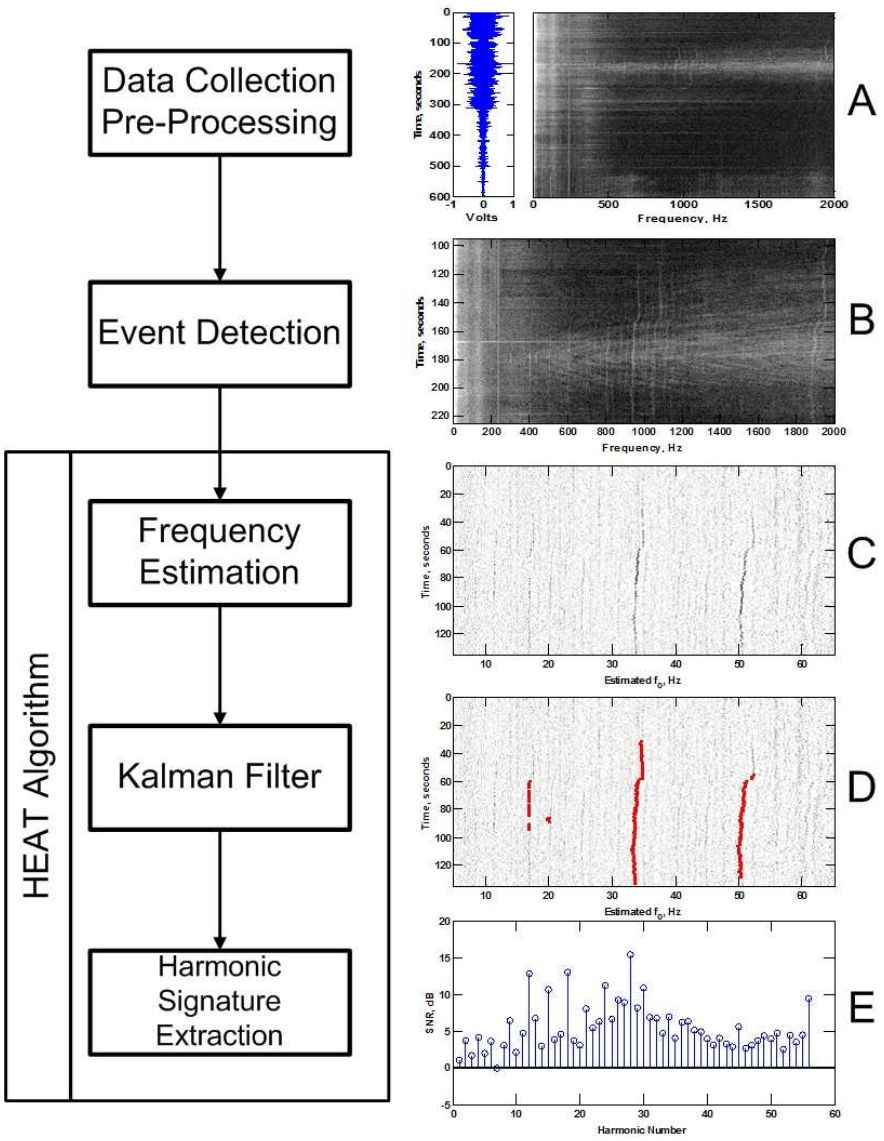

Figure 1.1: General overview of detection and estimation algorithm.

Chapter 2

Signal Model and Pre-Processing

In this chapter the model used to represent the signal received on the hydrophone is introduced. Also described is the pre-processing that was used to transform the data into a format that can be used by both the detection and estimation algorithms.

2.1 Signal Model

Consider a sum of many periodic sinusoidal signals whose frequencies are all harmonically related, being integer multiples of a fundamental frequency. This signal can be written as follows:

$$s(t) = \sum_{h=1}^{H} A_h \cos(2\pi h\gamma t + \phi_h),$$

where h is the harmonic number, A_h and ϕ_h are the amplitude and phase of the h^{th} harmonic component, and γ is the fundamental frequency. Assuming this signal is of infinite length, performing a Fourier transform on $s(t)$ will result in a series of delta functions with even spacing of γ.

Now consider the noise radiated from a ship as a combination of broadband noise as well as harmonically related sinusoidal tonal signals. This can similarly

be written as:

$$\begin{aligned} r(t) &= s(t) + n(t) \\ &= \sum_{h=1}^{H} A_h \cos(2\pi h\gamma t + \phi_h) + n(t), \end{aligned} \tag{2.1}$$

where once again, h is the harmonic number, A_h and ϕ_h are the amplitude and phase of the h^{th} harmonic component, and γ is the fundamental frequency, only now assume that the fundamental frequency and harmonic amplitudes are unknown. This fundamental frequency γ is related to the engine speed and other parameters by Table 1.1. We can rewrite $r(t)$ as $r(t, \theta)$, there he value θ represents a set of estimation parameters that consists of the fundamental frequency, γ, as well as the amplitude of all the harmonics, A_h, or $\theta = \{\gamma, A_h\}$.

The signal received on the hydrophone is different from the signal radiated from the boat for a number of reasons including changes in engine speed, Doppler shift, and other propagation effects. The signal can still be modeled as a sum of sinusoidal signals as in (2.1), but now the fundamental frequency γ is no longer constant. Instead, γ is written as $\gamma(t) = f_o + \Delta f(t)$ where f_o is the fundamental frequency, and $\Delta f(t)$ is the change in fundamental frequency over time, and $r(t, \theta)$ can finally be written as:

$$\begin{aligned} r(t, \theta) &= s(t, \theta) + n(t) \\ &= \sum_{h=1}^{H} A_h \cos(2\pi h\gamma(t) t + \phi_h) + n(t), \end{aligned} \tag{2.2}$$

Because the frequency content is constantly changing, a time-frequency representation is best to show the evolution of the frequency content.

2.2 Pre-Processing

Before processing the signal $r(t,\theta)$, it must be in a format that can properly represent the constantly changing frequency content. The time-frequency representation used here is the Short-time Fourier Transform (STFT). The STFT is computed by moving a short window along the data (creating a "snapshot") and computing the Fourier transform of the data along that window. The window length is assumed to be short enough that the change in fundamental frequency within the window is negligible, i.e., $\Delta f(t) \sim \Delta f(t_k)$ within the window. Now $\gamma_k = f_o + \Delta f(t_k)$, where t_k is the center time for the kth snapshot.

The signal in frequency domain for each window is now the convolution of the Fourier transformed data with the Fourier transform of the time-domain window, $S(f,\theta) = \mathscr{F}\{s(t,\theta)\} * \mathscr{F}\{rectwin(t)\}$, which results in $S(f,\theta)$ being a summation of weighted sinc (sinc $x \equiv \frac{\sin x}{x}$) functions:

$$S(f,\theta_k) = \sum_{h=1}^{H} \frac{A_h}{2} \text{sinc}\left[\pi\left(f - h\gamma_k\right)\right]. \tag{2.3}$$

Now consider $r(t,\theta)$ to be of finite length T_r. The signal is sampled at frequency f_s, sampling period $\Delta t = 1/f_s$, and can be written as $r(j\Delta t,\theta)$ where

$j = \{0, 1, \cdots, N_r - 1\}$ and $N_r = T_r/\Delta t$. Then the signal is partitioned into K overlapping segments, or snapshots, where each snapshot is of length T seconds, or N_s samples where $N_s = T/\Delta t$. The number of samples to overlap each snapshot is determined by the desired percent overlap, e.g, for 50% overlap, the number of samples to overlap is $50/100 * N_s = N_s/2$. The notation $r_k(n\Delta t, \theta_k)$ is used to represent the nth sample of the kth snapshot, where $n = \{0, 1, \cdots, Ns - 1\}$ and $k = \{1, 2, \cdots, K\}$. The received data, $r(n\Delta t, \theta)$, is now of dimension $[N_s \times K]$. Each snapshot is then transformed to the frequency domain by computing the N_s-point DFT using an FFT algorithm:

$$\begin{aligned} R_{ko}(m\Delta f, \theta_k) &= \tfrac{2}{N_s} \sum_{n=-N_s/2}^{N_s/2-1} r_k(t_k + n\Delta t, \theta_k) e^{-jm2\pi\Delta f n\Delta t} \\ & \qquad m = -N_s/2, \cdots, -1, 0, 1, \cdots, N_s/2 - 1, \end{aligned} \tag{2.4}$$

where t_k is the center time for the kth snapshot and $R_{ko}(m\Delta f, \theta_k)$ is the DFT coefficient at the frequency bin $m\Delta f$. The frequency resolution, or width of the bin Δf, is determined by the length of the snapshot window by $\Delta f = 1/T$. For simplicity $f_m \equiv m\Delta f$, where $f_m = \Delta f \{-N_s/2, \cdots, -1, 0, 1, \cdots, N_s/2 - 1\}$, so $R_{ko}(m\Delta f, \theta_k) = R_{ko}(f_m, \theta_k)$, which is of dimension $[N_s \times K]$.

Chapter 3

Detection Algorithm

In this chapter the details on a detection algorithm are presented. The algorithm was designed with the goal of eventually running on a real-time system to pick out only sections of data where a boat is present to pass on to an estimation algorithm. The performance of the detection algorithm will be measured by the number of false alarms and missed detections over a period of time.

3.1 Detection Algorithm

A detection algorithm was developed to search through data recorded from a single hydrophone and pull out sections of data with a ship signal present to be analyzed. The time-series, $r(t)$, is received on a single omni-directional hydrophone. H_0 is the hypothesis that the received time-series is only noise, and H_1 is the hypothesis that the received time-series is signal and noise.

$$\begin{aligned} H_0 &: r(t) = n(t), \text{ only noise} \\ H_1 &: r(t) = s(t) + n(t), \text{ signal and noise,} \end{aligned} \tag{3.1}$$

where $s(t)$ are the tonals generated by the boats engine and propeller after they propagate through the water, and $n(t)$ is the broadband noise and any environmental noise. An outline of the inputs to this algorithm are included in Appendix

B. The detection algorithm consists of 4 steps.

Step 1. Normalize the signal

Once the STFT is computed, each snapshot is normalized along frequency. The signal that is received is a combination of the tonals from the engine and propeller, broadband noise, and any environmental noise in the area. Since this noise is generally non-Gaussian, the received signal is normalized using a moving window of length W along frequency.

$$R_k(f_m, \theta_k) = \frac{R_{ko}(f_m, \theta_k) - \mu_{R,W}}{\sigma_{R,W}}, \tag{3.2}$$

where $\mu_{R,W}$ and $\sigma_{R,W}$ are the mean and standard deviation of $R_{ko}(f_m, \theta_k)$ in the window W.

Step 2. Initialize frequency detection

Once the signal is normalized, a threshold, λ_f, is applied to each snapshot. The threshold is a multiplier of the standard deviation, σ, and since the data is zero mean and unit variance, $\lambda_f \sigma = \lambda_f$. This means any frequency bin f_m that has normalized amplitude $R_k(f_m, \theta_k)$ greater than λ_f is considered a detection,

or

$$f_{m,det} : R_k(f_m, \theta_k) > \lambda_f.$$

Step 3. Cluster detected frequency bins

For each snapshot there are several detected frequency bins. Not all of these frequency bins are associated with tonals. In this step the detection are clustered in both time and frequency. This clustering is done to validate or invalidate any detection.

First, the detected frequency bins are clustered in frequency. The frequency clustering is done via a local max search along frequency for each snapshot. This is done in case there are multiple detected frequency bins associated with a single peak.

Then the frequency bins are clustered over time. This is done under the assumption that the tonals are relatively stationary not only within each snapshot, but also over N snapshots. It also assumes that the noise is not stationary over N snapshots. If a frequency bin, f_m, is detected for M out of N consecutive snapshots (from time $k - N$ to k) that detected frequency bin is considered valid at time k. This does introduce an M snapshot delay in detection.

It is at this point in the detection algorithm that certain frequencies can be excluded. If a frequency bin, f_m, is detected that is associated with a known

noise source, that detected frequency bin is considered invalid and is ignored.

Step 4. Event detection

Up to this point in the detection algorithm the evidence to support either H_0, no ship is present, or H_1, a ship is present has been built by validating detected frequency bins for each snapshot. In this step the evidence, or detected frequency bins in each snapshot, is added up and ran through a sequence of tests.

- The snapshots are clustered together in groups by searching for the first snapshot where the number of detections goes above a threshold, N_d, and the next snapshot in which the number of detections goes below that threshold, N_d. This is continued until the end of the recorded data is reached.
- If the time in between two clusters of snapshots is less than T_g seconds, the two clusters are grouped.
- If the total length of a group is greater than T_l seconds, the group is considered to satisfy H_1.

This series of time gating test works well to mitigate the effects of transient noises on the hydrophone.

Once a segment of data is identified as having a boat signature present, H_1, the data is run through an algorithm to extract the harmonic signature of the boat. All of these steps could easily be implemented on a real-time system.

Instead of calculating the STFT for the entire data set, the FFT of one snapshot at a time could be computed and the four steps presented above could operate at one snapshot at a time.

3.2 Performance of Detection Algorithm

To test the performance of the detection algorithm the algorithm was run on four hours of 23 consecutive days. The hours chosen were from 12:00 AM to 2:00 AM and from 8:00 AM to 10:00 AM. The early morning time was chosen because it is less likely that boats will be running at that time of day, so any detections during that time would likely be false alarms. The mid morning time was chosen as a time that boats would likely be passing by to test how well the detection algorithm picked them up. All the data used for this test was provided by PNNL from their Sequim, WA campus. More information on this data is provided in section 5.1.1.

The parameters used for the test of the algorithm are outlined in Table 3.1. Not included in Table 3.1 is the list of frequencies the detection algorithm filtered out as being sources on known noise; the list is too long for the table. The location that this data was collected is directly next to a pump that constantly refreshes salt water tanks located back on shore. There are a number of constant frequency lines that are associated with that pump that have to be ignored.

Table 3.1: Parameters used in application of the detection algorithm.

Variable	**Value**
Sample Rate	8000 Hz
Snapshot Window	1 sec
Snapshot Overlap	50%
Frequency Limits	0 to 2000 Hz
Normalizing Window	25 bins
λ_f	2σ
NM	3 snapshots
N_d	5 detections
T_g	5 seconds
T_l	30 seconds

Figure 3.1 shows the background noise, with arrows indicating the specific frequencies that are ignored by the detection algorithm. The algorithm also ignores any detections below 40 Hz, as the band from 0-40Hz is predominantly flow noise on the hydrophone.

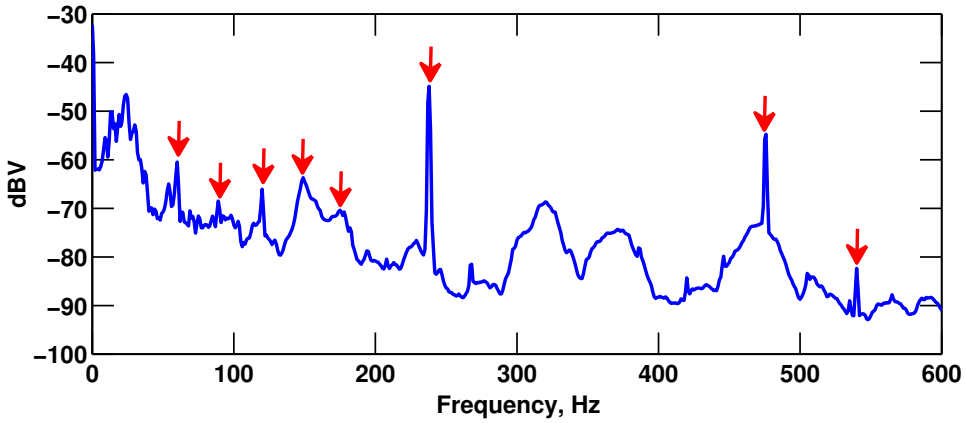

Figure 3.1: Average background noise from PNNL water pump. Arrows indicate the specific frequencies ignored by detection algorithm.

The total length of the time the detection algorithm evaluated was 92 hours

over 23 days. Over those 92 hours, only 12 false alarms occurred, and 1 missed detection, with 147 correct detections. Those detections that were false alarms were due to transient events, which are most likely objects hitting the hydrophone that were picked up with the current. The one missed detection was during a similar time where there was very strong broadband transient events that covered up the signature from the boat. During times of slack tide, when there was little to no impulsive noise on the hydrophone, there were no missed detections or false alarms. Figure 3.2 shows an example of a false alarm, along with a boat being picked up despite the transient events, plotted in a spectrogram. The green lines indicate the start time of the event, and the red lines indicate the end of the event.

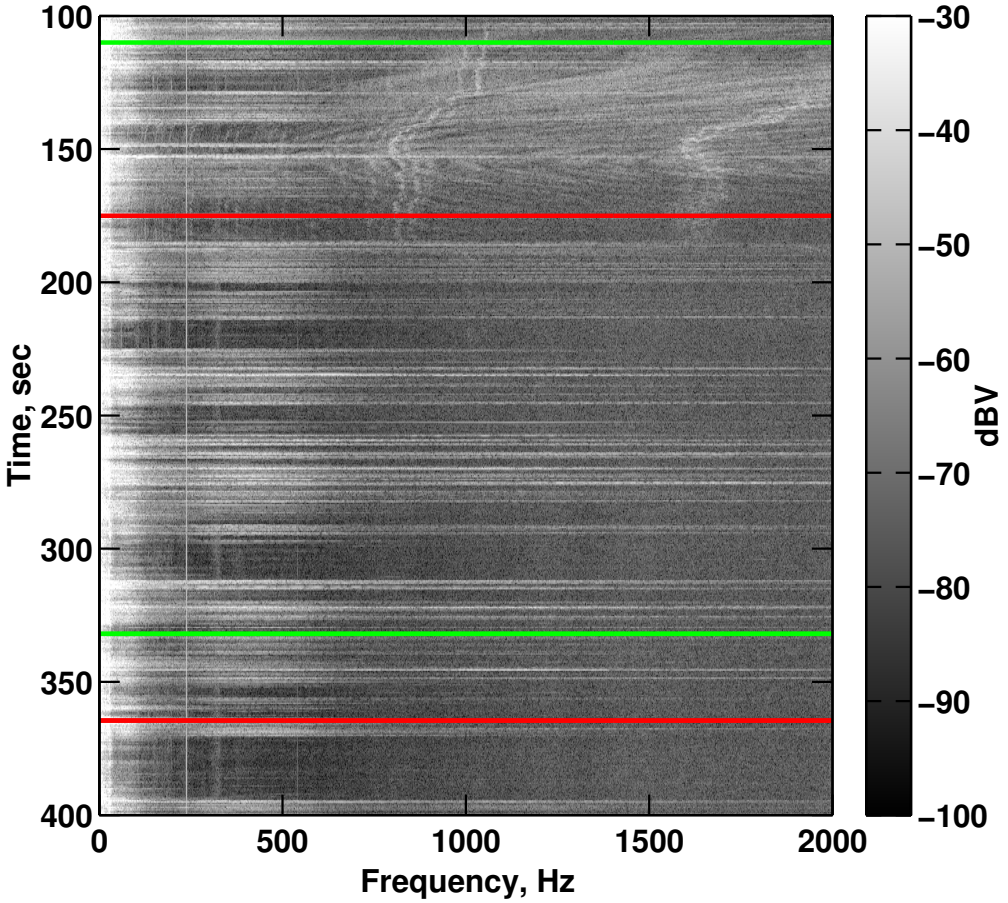

Figure 3.2: Example of two detected events. The green line indicates the start of an event and the red line indicates the end of an event. First event is a correct detection. Second event is a false alarm.

Chapter 4

Harmonic Extraction and Analysis Tool (HEAT)

The previous chapter discussed a detection algorithm used to pick out segments of acoustic data where a ship signature is present. Once a boat signature has been detected certain parameters about the boat need to be estimated in order to give some kind of information that will help in identifying the type of boat. An algorithm has been designed to extract important information from the data, and in this chapter, the details of the Harmonic Extraction and Analysis Tool (HEAT) are presented.

There are three main parts to the HEAT algorithm. First is the estimation of the fundamental frequency for the harmonic content from a boat. This is done by correlating the received signal with a hypothesized model with a known fundamental frequency at each snapshot. Second is refining the estimate by tracking the fundamental frequency with a Kalman filter to take advantage of the slow change in the fundamental frequency over time. Lastly is the extraction of the harmonic signature, which is an acoustic fingerprint of a boat. The fundamental frequency estimate from the Kalman filter is used as a basis for extracting the amplitude of the harmonic tonals from the data. These harmonic amplitudes are what make up the harmonic signature.

4.1 Fundamental Frequency Estimation

The noise radiated from small boats is modeled as a periodic, multi-harmonic signal, as was described in Section 2.1. In this section a method of estimating the fundamental frequency of a harmonic set is described. The Pearson Product-Moment Correlation Coefficient (PMCC) is used to estimate the correlation of the unknown harmonic set to a hypothesized model with known fundamental frequency.

4.1.1 Pearson Product-Moment Correlation Coefficient (PMCC)

The Pearson Product-Moment Correlation Coefficient (PMCC) is a measure of linear association between two random variables[15]. Consider the ordered pair of random variables X and Y with mean μ_x and μ_y, standard deviation σ_x and σ_y, and covariance σ_{xy}. The correlation coefficient between X and Y is defined as

$$\rho_{xy} = \frac{\sigma_{xy}}{\sigma_x \sigma_y},$$

where ρ is bounded between -1 and 1. This can be estimated from a sample of X and Y by

$$\hat{\rho}_{xy} = \frac{\sum_i (x_i - \mu_x)(y_i - \mu_y)}{\sqrt{\sum_i (x_i - \mu_x)^2}\sqrt{\sum_i (y_i - \mu_y)^2}}, \tag{4.1}$$

and is usually denoted as r, or Pearson's r. The notation $\hat{\rho}$ will be used as an estimate of the correlation coefficient, since r is used to to denote the received signal. The PMCC is used in this algorithm to measure the similarity of the measured signal $R_k(f_m, \theta_k)$ to a hypothesized model.

4.1.2 Hypothesized Model

In (2.3) the tonals recorded on the hydrophone were modeled as a sum of weighted, harmonically related sinc functions, with fundamental frequency γ_k. This same model is used as the hypothesized model, only here it is assumed that all the harmonics have equal amplitude:

$$\hat{S}(f_m, \bar{\gamma}) = \sum_{h=1}^{H} \operatorname{sinc}\left[\pi\left(f_m - h\bar{\gamma}\right)\right], \tag{4.2}$$

where $\bar{\gamma}$ is now a vector of fundamental frequencies. The vector $\bar{\gamma}$ is essentially a search window of fundamental frequencies bounded by γ_{min} and γ_{max}, with step size $\Delta\gamma$. Now $\bar{\gamma}$ can be written as $\bar{\gamma} = \gamma_{min} + p\Delta\gamma$, where $p = \{0, 1, \cdots, N_\gamma - 1\}$, and $N_\gamma = \left(\gamma_{max} - \gamma_{min}\right)/\Delta\gamma + 1$, which is of dimension $[1 \times N_\gamma]$. The model $\hat{S}(f_m, \bar{\gamma})$ is a matrix of size $[N_\gamma \times N_s]$, where N_s is the dimension of the frequency vector from the STFT.

4.1.3 Fundamental Frequency Estimation using PMCC

To estimate the fundamental frequency of the received signal $R_k(f_m, \theta_k)$, the signal is compared against a hypothesized model (4.2) using the PMCC (4.1):

$$\hat{\rho}_k(\bar{\gamma}) = \frac{\sum_{f_m} \left(\hat{S}(f_m, \bar{\gamma}) - \mu_S\right) R_k(f_m, \theta_k)}{\sqrt{\sum_{f_m} \left(\hat{S}(f_m, \bar{\gamma}) - \mu_S\right)^2 \sum_{f_m} \sqrt{R_k(f_m, \theta_k)^2}}}. \tag{4.3}$$

Notice that the mean for $R_k(f_m, \theta_k)$ does not appear in (4.3); it is already zero mean by equation (3.2). The vector $\hat{\rho}_k(\bar{\gamma})$ denotes the correlation coefficient for the kth snapshot of $R_k(f_m, \theta_k)$ for all values $\bar{\gamma}$. This is done for each snapshot which makes $\hat{\rho}(\bar{\gamma})$ dimension $[N_\gamma \times K]$. To distinguish this domain apart from the time-frequency domain such as $R_k(f_m, \theta_k)$, we will refer to $\hat{\rho}(\bar{\gamma})$ as being in the time-fundamental frequency domain.

When the fundamental frequency in $\hat{S}(f_m, \bar{\gamma})$ is equal the fundamental frequency in $R_k(f_m, \theta_k)$, i.e. $\bar{\gamma}_p = \gamma_k$, $\hat{\rho}_k(\bar{\gamma})$ will result in a high correlation and a peak at that frequency. Also, if the spacing in the hypothesized model is twice the fundamental frequency of the received signal, there will be another peak in $\hat{\rho}_k(\bar{\gamma})$ at that frequency. In fact, there are many peaks that show up in $\hat{\rho}_k(\bar{\gamma})$ as a result of this correlation analysis having to do with partial matches with multiples of the fundamental frequency. This is demonstrated in Figure 4.1, which shows the result the PMCC analysis of $\hat{S}(f_m, \bar{\gamma})$ to $\hat{S}(f_m, 20)$. The peak at 20 Hz shows the perfect correlation of the hypothesized signal with fundamental frequency of 20

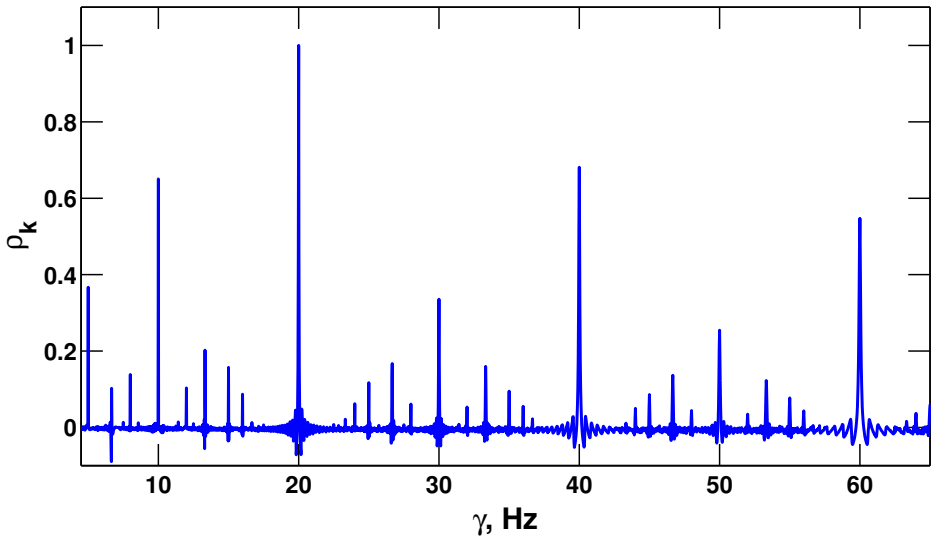

Figure 4.1: Example of the fundamental frequency correlation for the hypothesized signal with $\bar{\gamma}_p = 20$ Hz, $\hat{S}(f_m, 20)$.

Hz, $\hat{S}(f_m, 20)$, to itself. There are many other peaks that give high correlation due to only a fraction of the peaks lining up with the model. It is for this reason, as well as the desire to track the changing frequency content through time, that a Kalman filter is applied to $\hat{\rho}_k(\bar{\gamma})$.

4.2 Kalman Filter

The fundamental frequency estimator in (4.3) results in a correlation of the measured signal $R_k(f_m)$ to a hypothesized model $\hat{S}(f_m, \bar{\gamma})$ for all possible fundamental frequencies $\bar{\gamma}$ and for every snapshot k. A simple way to estimate of the fundamental frequency at time k, γ_k, would be to find the value of $\bar{\gamma}$ where $\hat{\rho}_k(\bar{\gamma})$ is maximum $\hat{\gamma}_k = \arg\max_{\bar{\gamma}} \hat{\rho}_k(\bar{\gamma})$. However, for this method it is not enough since

there can be multiple signals of interest (boats) present with different fundamental frequencies. Also, if the gear ratio in the engine is not an integer number, the fundamental frequency for the engine harmonics will be different from the fundamental frequency of the shaft and propeller harmonics. For that reason, a Kalman filter has been implemented in the correlation domain that will follow a time evolving estimate of the fundamental frequency for each of the peaks present.

4.2.1 Introduction of Kalman Filter

The Kalman filter is a widely used recursive algorithm used to estimate the mean and error covariance of a state through time given a series of noisy observations of the state [16]. The mathematical description of how the state propagates through time is given by a process model. The relationship between the state and the observations is given by a measurement model. The discrete Kalman filter is limited to linear process and measurement models with additive white noise. Modifications to the Kalman filter have been made to allow for non-linear models. Some methods are the extended Kalman filter which uses a first order Taylor series approximation, and the unscented Kalman filter which uses a deterministic sampling technique to pick a set of sigma points which are used to estimate the mean and error covariance of the state. Non-linear modifications may offer more flexibility, but come at a cost of complexity, and also have poorer performance

on linear models. This design fits into a linear model, so a conventional Kalman filter was a natural choice.

All of these methods work on the same underlying principles. The mean $\hat{\Gamma}_k^-$ and error covariance P_k^- of the state are projected forward, or predicted (4.4), from time $k-1$ to time k from initial estimates supplied to the filter, $\hat{\Gamma}_{k-1}$ and P_{k-1}.

$$\begin{aligned} &\underline{\text{Predict}} \\ \hat{\Gamma}_k^- &= F\hat{\Gamma}_{k-1} \\ P_k^- &= FP_{k-1}F^T + Q \end{aligned} \tag{4.4}$$

The estimate of the mean $\hat{\Gamma}_k$ and error covariance P_k are then corrected, or updated (4.5), based on the measurement of the state z_k at time k.

$$\begin{aligned} &\underline{\text{Update}} \\ K_k &= P_k^- H^T \left(HP_k^- H^T + R\right)^{-1} \\ \hat{\Gamma}_k &= \hat{\Gamma}_k^- + K_k\left(z_k - H\hat{\Gamma}_k^-\right) \\ P_k &= \left(I - K_k H\right) P_k^- \end{aligned} \tag{4.5}$$

At the next time step the corrected estimates are fed back into the prediction equations and the process repeats. This gives a filtered estimate of the mean and error covariance of the state for all times. That filtered estimate over time

is called a track.

4.2.2 Fundamental Frequency Tracking using Kalman Filter

A discrete Kalman filter has been implemented as a peak follower on $\hat{\rho}_k(\bar{\gamma})$. This stage of the algorithm follows many of the same steps as the detection algorithm discussed in Chapter 3. Step 1 is skipped because there is no need to normalize $\hat{\rho}_k(\bar{\gamma})$. Step 2 in the detection algorithm was to threshold the normalized signal to get an initial set of frequency detections. In the tracking algorithm, $\hat{\rho}_k(\bar{\gamma})$ is thresholded for each snapshot across fundamental frequency with threshold λ_γ. Then in Step 3, the detections are clustered with a local max search. This now gives for each snapshot of $\hat{\rho}_k(\bar{\gamma})$, a set of detections Z_k which are the inputs to the Kalman filter.

The tracker adapts a near constant velocity process model to tracking frequency, given by

$$
\begin{aligned}
\Gamma_k &= F\Gamma_{k-1} + w_k, \\
\Gamma_k &= \begin{bmatrix} \gamma_k \\ \dot{\gamma_k} \\ \rho_k \end{bmatrix}, \\
F &= \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \\
w_k &\sim \mathcal{N}(0, Q),
\end{aligned}
\tag{4.6}
$$

where Γ_k is the state vector for a single track, γ_k and $\dot{\gamma_k}$ represent the frequency and change in frequency (frequency velocity) from time $k-1$ to k respectively, ρ_k is the correlation value for frequency γ_k, F is the state transition matrix, and w_k is assumed to be a zero-mean white Gaussian process with variance Q. For a Kalman filter tracking an object' s position, the near constant velocity model would assume that from time $k-1$ to k, the change in velocity of the object is negligible. The position can be predicted as the previous position plus some change in position due to the object moving at some velocity over some time step. The adaptation of this model to tracking frequency assumes that the change in frequency from time $k-1$ to time k is negligible, so the frequency at time k, γ_k,

can be predicted as $\gamma_{k-1} + \dot{\gamma}_k$. Actual changes to the frequency are accounted for through the process noise term.

The measurement model directly relates the measurement z_k to the state at time k Γ_k as follows:

$$\begin{aligned} z_k &= H\Gamma_k + v_k, \\ z_k &= \begin{bmatrix} \gamma \\ \rho \end{bmatrix}, \\ H &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \\ v_k &\sim \mathcal{N}(0, R), \end{aligned} \tag{4.7}$$

where z_k is the measurement with frequency γ and correlation value ρ, H relates the state to the measurement, and v_k is assumed to be a zero-mean white Gaussian process with variance R.

The process noise covariance Q and the measurement noise covariance R are

both matrices defined as follows:

$$
\begin{aligned}
Q &= \begin{bmatrix} Q_\gamma & 0 & 0 \\ 0 & Q_{\dot{\gamma}} & 0 \\ 0 & 0 & Q_\rho \end{bmatrix}, \\
R &= \begin{bmatrix} R_\gamma & 0 \\ 0 & R_\rho \end{bmatrix}.
\end{aligned}
\tag{4.8}
$$

These parameters are the main tuning parameters of the Kalman filter. The amount of noise injected into the process or measurement model indicates the certainty with which the models are trusted to accurately estimate the state.

4.2.3 Tracker Logic

The Kalman filter requires an initial estimate of the state, so the tracking algorithm includes logic-based track initiation and termination [17]. There are several possible fundamental frequencies that fit the data, which are obtained by applying a threshold, λ_γ, to the fundamental frequency estimate, $\hat{\gamma}_k$, for each snapshot. These frequencies are presented to the tracker as a set of detections or observations for time k in the set Z_k. These detected frequencies are used to initiate tracks, as well as observations for the Kalman filter. There are three states which a track can be in:

- Initiated: If a fundamental frequency is detected in M out of N consecutive snapshots, a track is created.

- Flagged: If a track has no associated observations in the set Z_k, the track is flagged for termination.

- Terminated: If a track is flagged for N_F consecutive snapshots, the track is terminated.

At each time step, a set of observations need to be paired to the tracks. For an observation to be associated with a track the observation must satisfy the following threshold condition:

$$\left(z_k - H\hat{\Gamma}_k^-\right)\left(HP_k^-H^T + R\right)^{-1}\left(z_k - H\hat{\Gamma}_k^-\right)^T < \chi^2. \tag{4.9}$$

In the case where multiple observations satisfy this condition, the best match is the one with the smallest χ^2 value.

Before a track is initiated, the algorithm searches within a small window, W_γ, around the proposed frequency for any existing tracks. If a track already exists within that window, the new track is immediately terminated, as it is assumed to belong to the already existing track.

4.2.4 Harmonic Content Parameter

The Kalman filter outputs multiple tracks, each track being a time-evolving estimate of a fundamental frequency for the harmonic content in the signal $R_k(f_m)$. To determine which track best fits the data, a parameter Ψ is calculated from the estimate of the correlation, which is the third component of the state vector in Eq. 4.6, as follows:

$$\Psi\left[\kappa\right] = \sqrt{\frac{1}{L}\sum_{k=1}^{L}\left|\hat{\rho}_k\left[\kappa\right]\right|^2}, \tag{4.10}$$

where κ is the index of the track, $\hat{\rho}_k\left[\kappa\right]$ is the estimate of the amplitude from the Kalman filter for track κ, and L is the length of the track in snapshots. The track with the highest Ψ value is chosen as the best fit to the data and is deemed the best estimate of the fundamental frequency.

The value $\hat{\rho}_k$ is a measure of how well the data (3.2) fits the hypothesized model (4.2). The value $\hat{\rho}_k$ can also be described as a measure of how much harmonic content is present in the signal. A value of Ψ equal to one would mean that the data perfectly matches the model. Since the hypothesized model assumes that all the harmonics are present and equal in amplitude, and also that there is no noise present, this perfect match of measurement to model is not achievable. But this parameter can give an intuition on what kind of signal is there. Since the parameter $\Psi\left[\kappa\right]$ is the average of $\hat{\rho}_k$ over the length of the track, the higher

$\Psi[\kappa]$ is, the more harmonic content will be visible in the spectrogram over the length of the track. The opposite case is also true in that the lower $\Psi[\kappa]$ is, the less harmonic content will be visible in the spectrogram.

4.3 Harmonic Signature Extraction

The harmonic signature is thought to be a kind of acoustic fingerprint of a boat. We know from Table 1.1 that a motor will give rise to several different frequencies based on the number of cylinders, gear ratio, number of blades, etc. Those frequencies are fundamental frequencies of the tonals described in equation 2.2. In the previous sections the fundamental frequency $\gamma(t)$ was estimated and tracked through time using a Kalman filter. This estimate of $\gamma(t)$ is the first part of the estimation parameter θ. The second part is the amplitude of all the harmonics, A_h. These harmonic amplitudes are what make up the harmonic signature.

To obtain a harmonic signature, the fundamental frequency track of best fit is projected onto the spectrogram for all harmonics of that track. This is shown in figure 4.2, where 4.2(a) is the spectrogram ($R_k(f_m)$) and 4.2(b) is the harmonics of the fundamental frequency track projected onto the spectrogram. Once the track is projected onto the data, the amplitude for each of the harmonics is found by searching for a peak in a small window around the projected frequencies. The local noise of each harmonic is also estimated by averaging the spectrogram in a small window on each side of the peak.

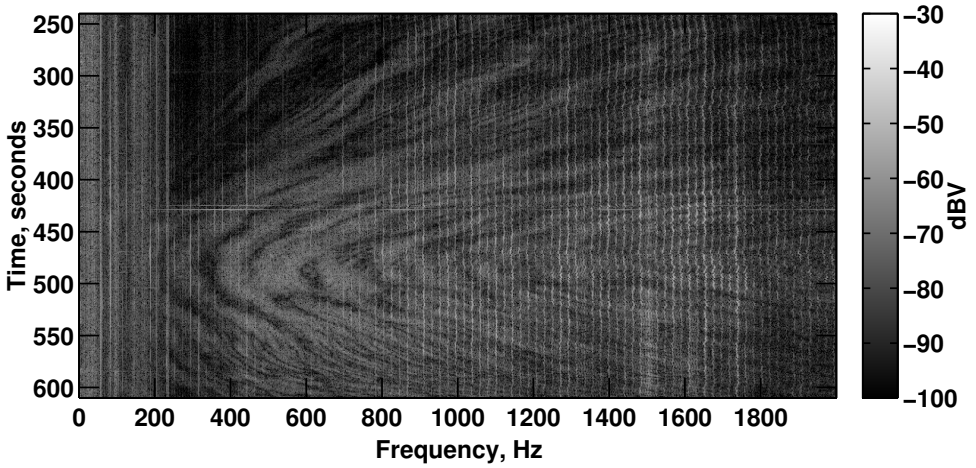

(a) Spectrogram of target.

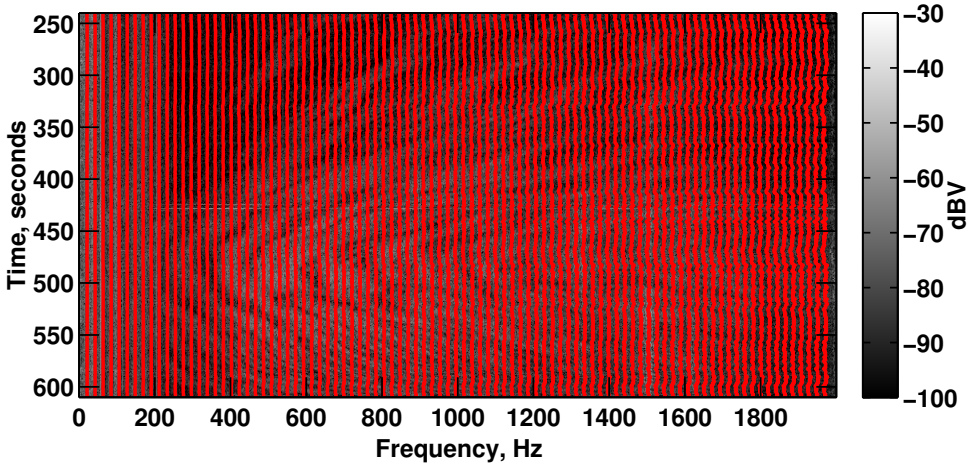

(b) Spectrogram of target fundamental frequency estimate projected out to all the harmonics.

Figure 4.2: Spectrogram with example of the fundamental frequency projected onto all possible harmonics.

The amplitude and noise are then averaged over the length of the track to give the final product of the HEAT algorithm, shown in Figure 4.3. This shows the harmonic amplitudes (stem plot) and the local noise around each peak (solid line). The amplitudes are in dB relative to the weakest harmonic. This is done since depending on the distance of the boat from the hydrophone, or the speed of the boat, the absolute amplitudes can widely vary, but the relative amplitudes of all the harmonics should stay the same regardless. An alternative signature

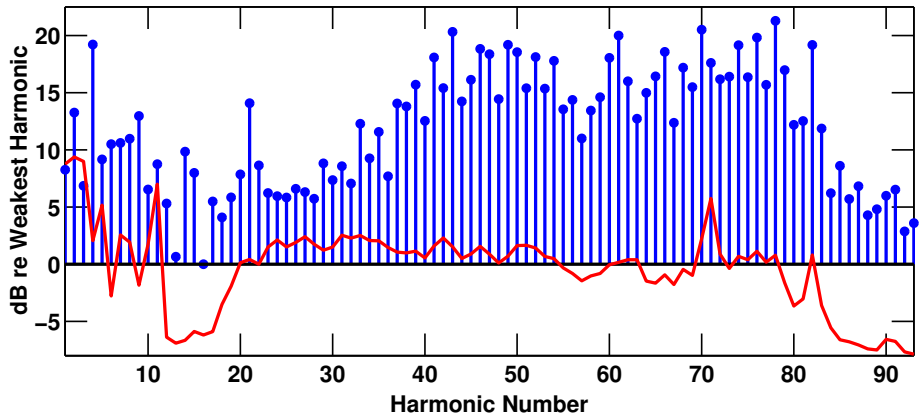

Figure 4.3: Harmonic Amplitude Signature (stem plot) extracted by HEAT algorithm with background noise (solid line)

could be shown by dividing the amplitude by the noise and then convert to dB to give the signal to noise ratio (SNR), but by putting the noise curve on the amplitude signature plot, both are visible.

4.4 HEAT Viewer GUI

The output of the HEAT algorithm is a structure, whose fields are described in Appendix D. A graphical user interface (GUI) was created to aid in the manual analysis of the HEAT algorithm, named HEAT Viewer (HEAT-V). Figure 4.4 shows the GUI with a target file loaded. This GUI loads a target file (A) and displays the spectrogram and the fundamental frequency correlation. It gives the user the option to over-plot the frequency tracks (B) on the correlation plot (C), as well as over-plot all the harmonics of that frequency track the spectrogram (D). Once the track that fits the data the best is selected, the user can generate the Harmonic Signature (E) for that boat.

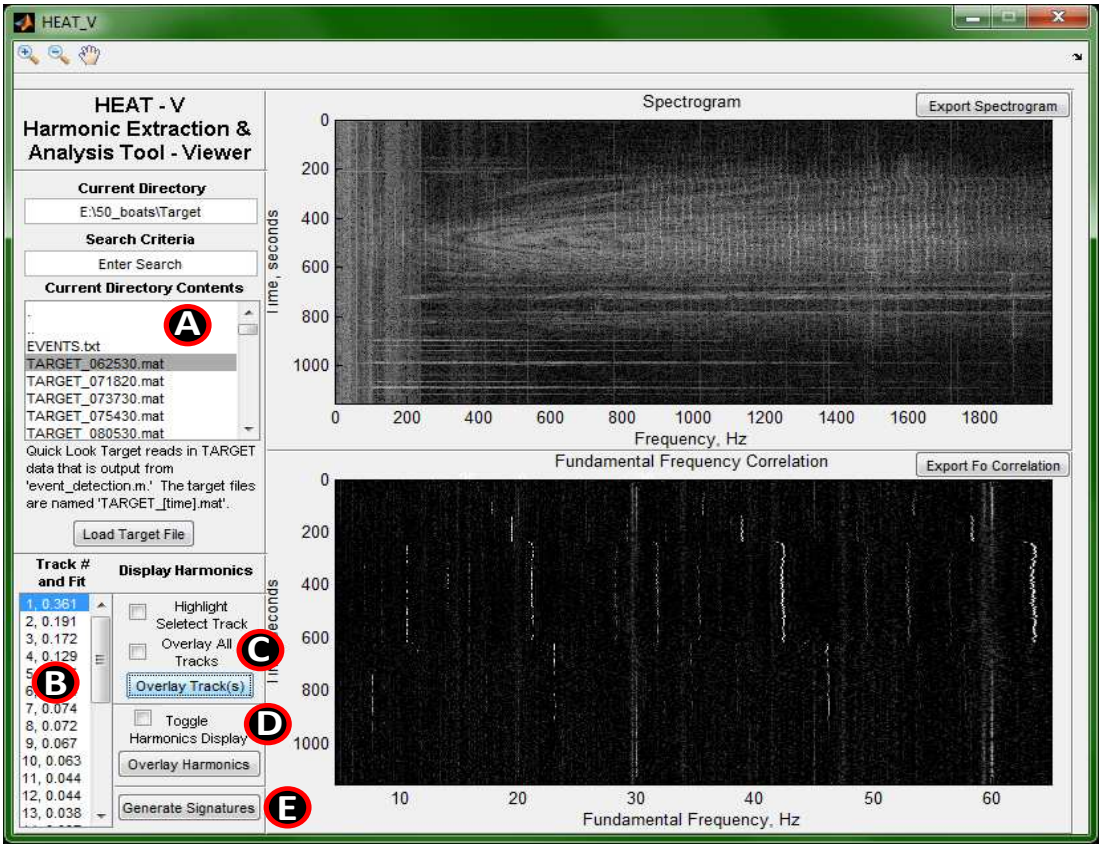

Figure 4.4: Screen capture of the HEAT Viewer GUI

Chapter 5

Results of Data Analysis

5.1 Data Description

Performance of the HEAT algorithm has been evaluated on boat noise collected at three locations under different environmental conditions. These are (1) multi-sensor data collected by PNNL in Sequim Bay, WA, (2) acoustic data collected using a laptop system on the Willamette River in Portland, OR and (3) acoustic data collected in Ahihi-Kinau Natural Area Reserve, HI.

5.1.1 Sequim Bay Data

The Sequim campus of PNNL is located at the mouth of Sequim Bay on the north part of the Olympic Peninsula of Washington. The John Wayne Marina located inside the bay allows for a very diverse population of boat traffic coming in and out of the area. Thus it is an excellent location for collecting test data from different types of small to mid scale boats. PNNL has been continuously monitoring boat traffic at this site for almost two years, and provided a data set including passes of 50 boats for testing the HEAT algorithm. Many of the 50 boats are duplicate passes by the same or a similar type of boat. This allows the evaluation of HEAT algorithm on repeated harmonic signatures.

The data set included both acoustic and non-acoustic information. The acoustic data was collected using a hydrophone mounted approximately 20 feet from the dock on the ocean floor (approximately 30 ft deep). The hydrophone used to record the boat noise was cabled back to the PNNL dock. At the dock the data was pre-processed using an anti-aliasing filter with cut-off frequency of 2.5kHz. The data was then sampled at 8kHz with 16-bit resolution.

Among the non-acoustic monitoring methods, there was a radar system providing an estimate of boat velocity, an electro-optic and infrared (EO/IR) camera recording a video of each boat pass, and a number of environmental sensors giving measurements of water temperature, current, etc. PNNL also provided records of each boat such as the hull material, engine type (e.g. inboard or outboard), and approximated length.

5.1.2 Willamette River Data

Data was also collected at the Riverplace Marina on the Willamette River in Portland, OR. The system used to collect this data was a laptop based system designed by the author. It is similar to the hydrophone system used in PNNL, but the additional mobility allows easy data collection at various locations. The data was recorded on a hydrophone pre-processed by a anti-aliasing filter with a cut-off frequency of 2.5kHz. It was sampled at 8kHz with 12-bit resolution using a Measurement Computing Corporation USB-1208FS data acquisition module.

In later versions, the data acquisition was done through the line-in audio port on the laptop which gives 16-bit resolution like the PNNL system.

5.1.3 Ahihi-Kinau Data

The last data set was collected in the Ahihi-Kinau Natural Area Reserve in Maui, HI. This data was recorded using the NEAR-Lab's SOREN system [18], which is a completely autonomous passive recording device. The battery power of the latest SOREN systems lasts approximately four weeks at a sampling rate of 44.1 kHz with 16-bit resolution. Three SOREN's were deployed in this area in March of 2010. They recorded more than 100 hours of environmental data including boat noise, whale noise, and snapping shrimp noise. There was no other surveillance systems available for this site, so there are no records of boat types or images.

5.2 Application of HEAT to data

Feature extraction results are shown for six boats including four boats from PNNL data, one boat from Willamette data and one boat from Hawaii data. Table 5.1 gives a summary of these boats. This table shows the boat identification letter (A-F), location where the data was collected, the hull material and engine type (if known) and the harmonic content parameter Ψ from the HEAT algorithm. The four boats chosen from the PNNL data were all 6 meter inboard boats. The objective for selecting these four boats was to discriminate the harmonic

signatures between inboard engine boats of approximately the same length with different hull materials.

Table 5.1: Summary of all the boats used to test the HEAT algorithm.

Boat ID	Location Collected	Approximate Length	Engine Type	Hull Material	Ψ
A	Sequim Bay, WA	6 m	Inboard	Fiberglass	0.468
B	Sequim Bay, WA	6 m	Inboard	Fiberglass	0.400
C	Sequim Bay, WA	6 m	Inboard	Unknown	0.195
D	Sequim Bay, WA	6 m	Inboard	Aluminum	0.138
E	Willamette River, OR	5 m	Outboard	Aluminum	0.280
F	Ahihi-Kinau, HI	Unknown	Unknown	Unknown	0.184

Table 5.2 gives the parameters used for the HEAT algorithm. The same parameters were used to process all the data from the three different locations. The data from Sequim and from the Willamette both had a low pass filter with cutoff at 2.5kHz. For this reason the analysis was only performed for frequencies up to 2kHz. This greatly reduces the dimension of the data to be analyzed which allows for faster processing.

Pictures of boats A through E can be seen in Appendix A. There was no photo record of boat F since it was recorded during an overnight deployment without human monitoring. Figures 5.1 through 5.12 show the spectrogram of each boat and their harmonic amplitude signatures extracted by HEAT. Boats A, B and E (figures 5.1, 5.3 and 5.9) show the case where the harmonics are clearly visible in the spectrogram throughout the entire frequency band. Boats C

and D (figures 5.5 and 5.7) show the opposite case where there are limited visible harmonics on the spectrogram within the frequency band. Boat F (figure 5.11) has clearly visible harmonics up to about 700 Hz, with another strong harmonic around 850Hz.

Table 5.2: Parameters used in application of the HEAT algorithm.

Variable	**Value**
Snapshot Window	1 sec
Snapshot Overlap	50%
Frequency Limits	0 to 2000 Hz
Normalizing Window	25 bins
$[\gamma_{min}\ \gamma_{max}]$	[4.5 65] Hz
$\Delta\gamma$	0.025 Hz
λ_γ	0.09
NM	3 snapshots
χ^2	3
N_F	3 snapshots
W_γ	0.5 Hz
Process Noise Covariances	
Q_γ	$(2\Delta\gamma)^2$
$Q_{\dot{\gamma}}$	$(2\Delta\gamma)^2/10$
Q_ρ	0.02
Measurement Noise Covariances	
R_γ	$(5\Delta\gamma)^2$
R_ρ	0.03

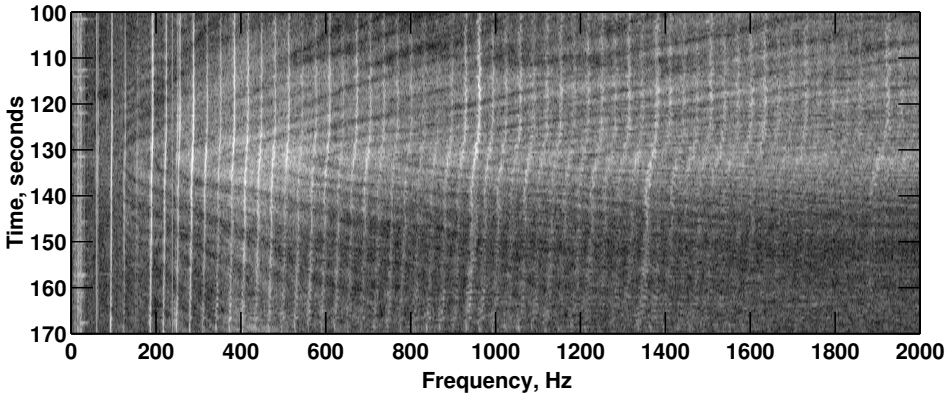

Figure 5.1: Spectrogram for boat A

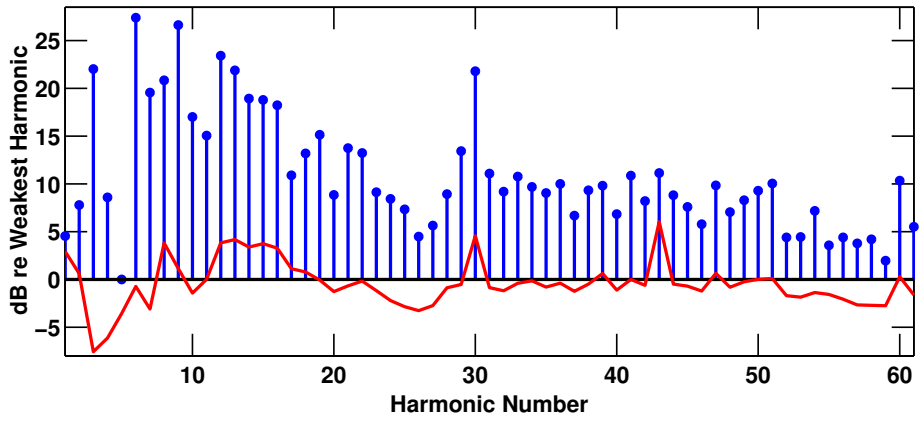

Figure 5.2: Harmonic Amplitude Signature (stem plot) with background noise (solid line) for boat A

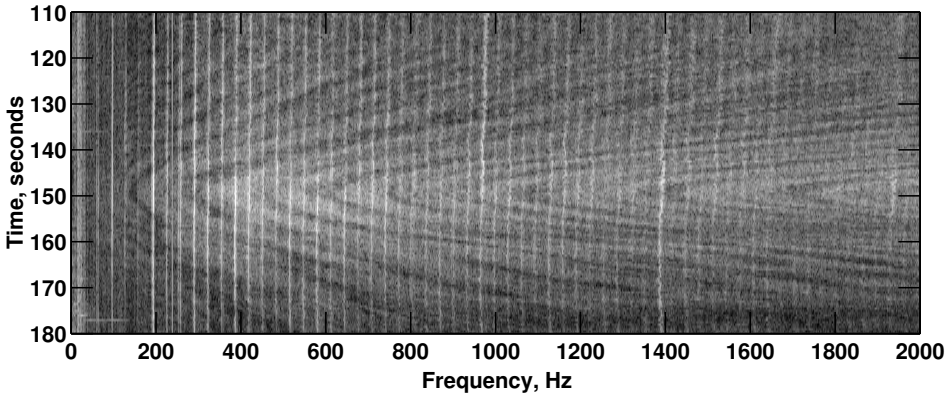

Figure 5.3: Spectrogram for boat B

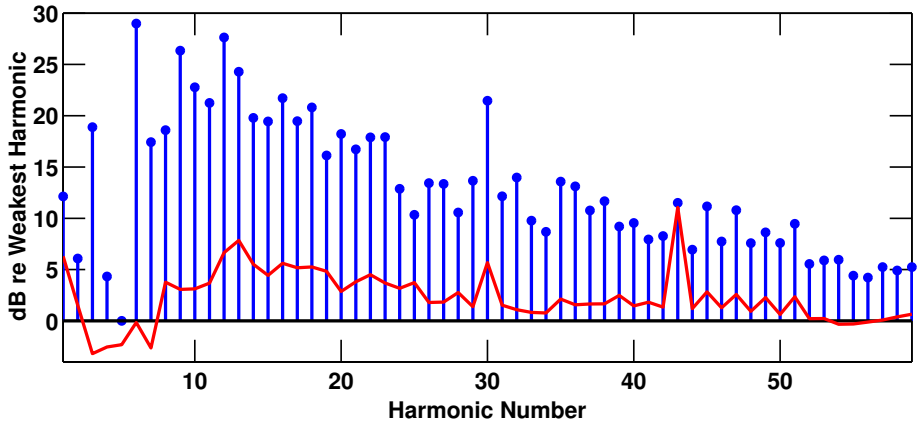

Figure 5.4: Harmonic Amplitude Signature (stem plot) with background noise (solid line) for boat B

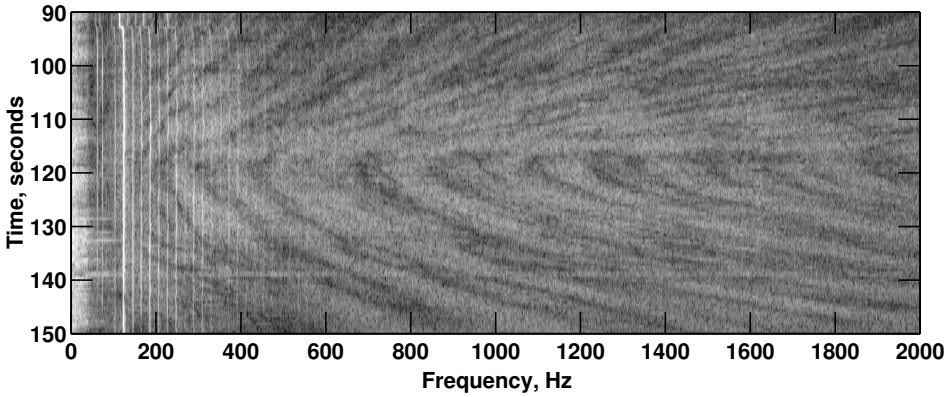

Figure 5.5: Spectrogram for boat C

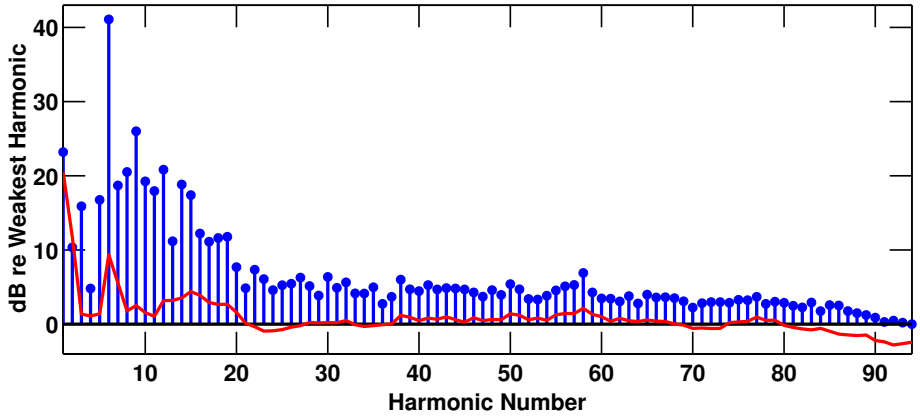

Figure 5.6: Harmonic Amplitude Signature (stem plot) with background noise (solid line) for boat C

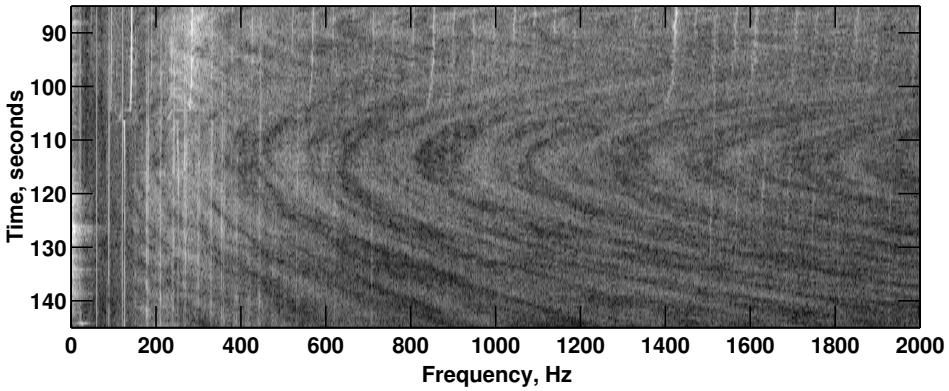

Figure 5.7: Spectrogram for boat D

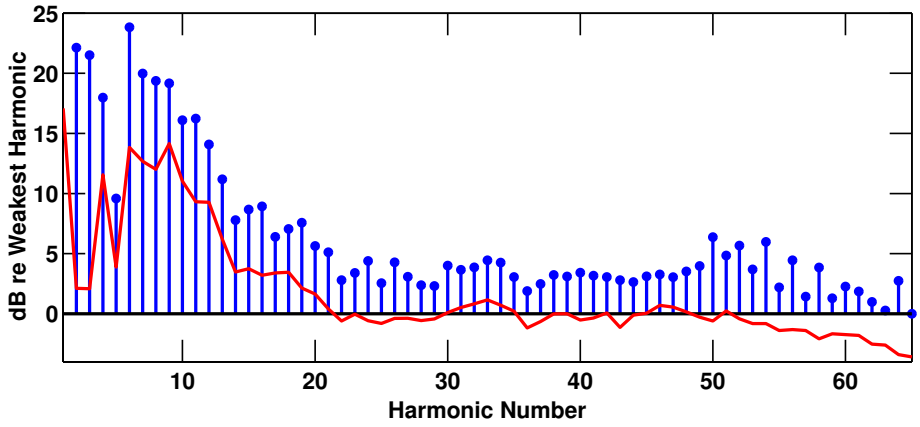

Figure 5.8: Harmonic Amplitude Signature (stem plot) with background noise (solid line) for boat D

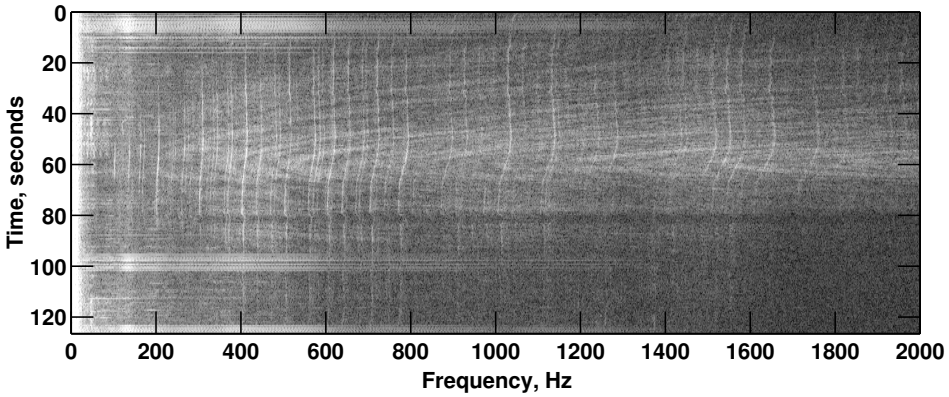

Figure 5.9: Spectrogram for boat E

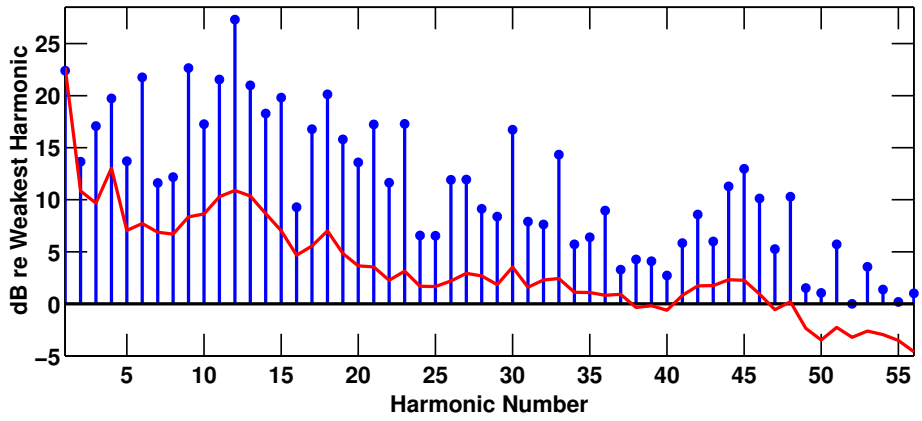

Figure 5.10: Harmonic Amplitude Signature (stem plot) with background noise (solid line) for boat E

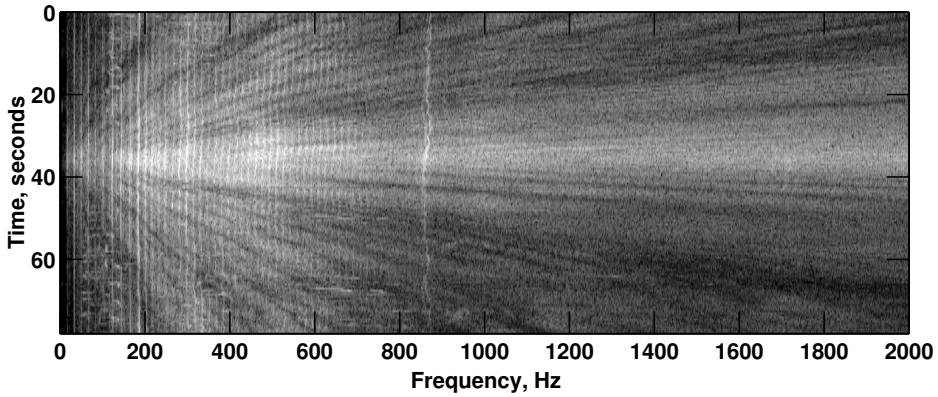

Figure 5.11: Spectrogram for boat F

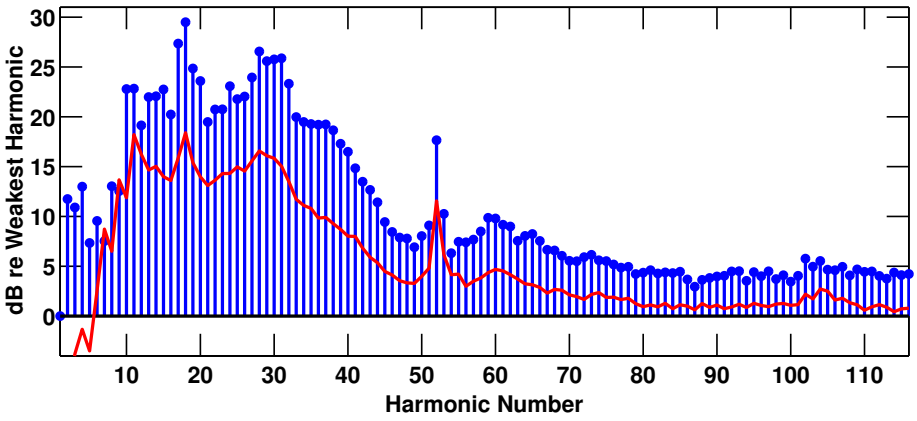

Figure 5.12: Harmonic Amplitude Signature (stem plot) with background noise (solid line) for boat F

Comparing the spectrograms of all the boats to the harmonic amplitude signatures, and specifically observing the intensity of the harmonic lines relative to the background noise, there is relatively good qualitative agreement between what can be seen in the spectrogram and what the HEAT algorithm extracted. Recall that the Ψ parameter was defined in Section 4.2.4 as a measure of how much harmonic content is visible in the spectrogram. In Table 5.1 boats A, B and E have higher Ψ values than boats C, D and F. This shows that HEAT algorithm is correctly extracting the harmonic amplitude signatures of the boats. These results also suggest that the performance of HEAT is not location dependent. The next step is to use the signatures for boat identification.

5.3 Harmonic Signature Correlation

As described in section 4.1.1, the product moment correlation coefficient (PMCC) is a measure of linear association between two ordered sets of data. The PMCC is used here to compare the harmonic amplitude signature from one boat to another. This is done by taking two amplitude signatures, cropping them to the length of the shortest signature, and calculating the correlation coefficient as in equation 4.1. Table 5.3 gives the results of this harmonic amplitude signature correlation for each of the four boats recorded in Sequim Bay, WA.

Table 5.3 shows a strong correlation of 0.87 between the harmonic amplitude signatures of boats A and B. Recall that in Table 5.1, boats A and B have

Table 5.3: Harmonic amplitude signature correlation results for the four boats from Sequim Bay, WA.

Boat	A	B	C	D
A	1.00	0.87	0.68	0.60
B	0.87	1.00	0.66	0.47
C	0.68	0.66	1.00	0.82
D	0.60	0.47	0.82	1.00

the same hull material and engine type. Boats A and B are in fact the same boat passing by the hydrophone at two separate times on two different days. High correlation between these two signatures shows that the extraction of the signatures is repeatable with enough accuracy to recognize the same boat.

Table 5.3 also that the second highest correlation (0.82) is found between boats C and D. According to the pictures of boats C and D in Appendix A, they are not the same boat. Comparing the harmonic amplitude signatures in Figures 5.6 and 5.8, the amplitudes follow a similar trend, but the local noise for the two boats is very different, especially at the lower frequencies. This observation suggests that the harmonic signature correlation on the amplitude signatures is not enough to distinguish between these two boats. Instead, the harmonic SNR signatures, as shown in Figures 5.13(a) and 5.13(b), are used in the harmonic signature correlation. Using the SNR measurement, boats C and D only have a harmonic signature correlation of 0.32.

Table 5.4 shows the correlation results for the harmonic SNR signatures of all

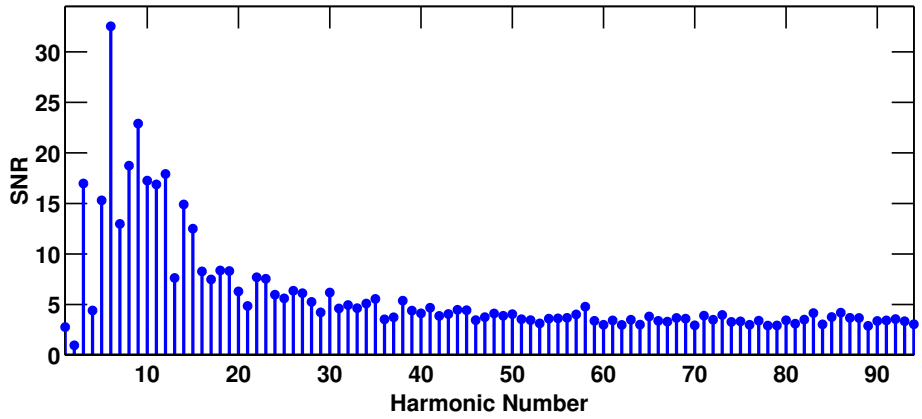

(a) Harmonic SNR Signature for boat C

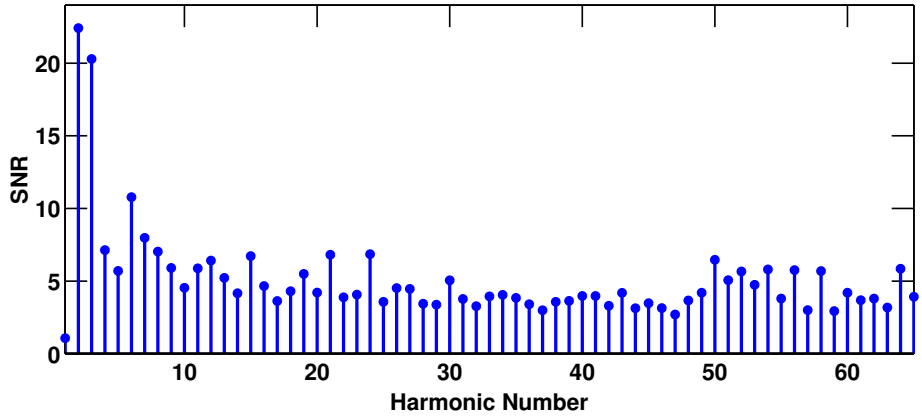

(b) Harmonic SNR Signature for boat D

Figure 5.13: The Harmonic SNR Signatures for boats C and D show a difference in the harmonic structure of boats C and D.

six boats. This table show that by correlating the harmonic SNR signature, boats A and B still have a high signature correlation, and now C and D have a low correlation, as to be expected. However, now boats A, B correlated to boat C have signature correlation values of 0.78. This does not give any conclusive evidence of using the harmonic SNR signatures for the correlation analysis. During this

Table 5.4: Harmonic SNR signature correlation results for the four boats from Sequim Bay, WA.

Boat	A	B	C	D
A	1.00	0.89	0.78	0.53
B	0.89	1.00	0.78	0.29
C	0.78	0.78	1.00	0.32
D	0.53	0.29	0.32	1.00

analysis the Ψ parameter has not yet been considered.

Boats A and B have high Ψ values, whereas boat C has a low value. This means that the majority of the harmonics extracted in the harmonic signatures of A and B are the actual tonals and not an estimation of the noise, where the opposite is true for boat C. By this measurement, the correlation of 0.78 between A, B and C can be held with less weight.

The harmonic signature correlation analysis was then applied on all of the 50 boats provided by PNNL. Of those 50 boats, 12 pairs of boats had a harmonic amplitude signature correlation of 0.7 or higher. Of those 12 pairs, only one pair was not a duplicate pass from the same boat. This means, given a data set collected at the same location with the same sound propagation paths, a harmonic amplitude signature correlation of above 0.7 between two boats strongly suggests that the two signatures are from duplicate passes of the same boat. This ability to match a signature to a specific boat could be used in MPA's like in Molokini where regular commercial snorkeling and diving boats are authorized and fishing

vessels are not.

Chapter 6

Conclusions and Future Work

Detection and classification of small boats is a relatively unexplored research area. There has been a large amount of work on classification of large ships in the past, but for small boats there are very few efforts in the literature. The goal of this thesis was to develop signal processing methods to detect and discriminate different types of boat noise recorded using passive acoustic sensors. There are many applications to this work including national security in littoral zones and surveillance in marine protected areas.

Two new algorithms were developed during the course of this research. First, a detection algorithm was designed to search through a large amount of hydrophone data for sections where a boat is passing. This algorithm searches for frequency tonals which arise from moving parts in a boats engine and propeller. The detection algorithm was tested on a data set recorded for a duration of 92 hours, in which there were 12 reported false alarms and 1 missed detection. Results of this analysis suggest that the detection performance is dependent on the tides. During slack tide there were zero missed detection and zero false alarms. The false alarms and missed detection occurred during the changing tides when objects get picked up in the strong current and hit the hydrophone, causing loud impulsive noises. A further advancement to this algorithm is to make it more

robust against those impulsive noises.

Another signal processing strategy developed in this thesis is the HEAT algorithm. This algorithm estimates the fundamental frequency of the harmonically related tonals and extracts the amplitude of the harmonics. The relative amplitudes of all the harmonics constitute a harmonic signature that can be used in identification of the boats. It was shown that the signatures can be extracted using the HEAT algorithm. Also, results show that the signatures are unique enough that the same boat passing by the hydrophone multiple times can be recognized.

The signatures extracted by the HEAT algorithm were used in a correlation analysis. It was found that a correlation value greater than 0.7 strongly suggests that the two signatures are from duplicate passes of the same boat. Further advancements of using the harmonic correlation in conjunction with the harmonic content parameter for boat identification could be made to ensure robustness of correlation analysis.

Another field for future research is to include an estimate of the fundamental frequency representing the engine and propeller harmonic sets at the same time. This could be done by adding a second fundamental frequency in the hypothesized model that is related by some non-integer gear ratio. The extracted fundamental frequencies could also be related back to the engine parameters to estimate the speed of the engine as well the number of cylinders in the engine and number of

blades on the propeller.

References

[1] M. McDonald and S. Lycett. Fast versus slow scan radar operation for coherent small target detection in sea clutter. *IEE Proceedings: Radar, Sonar and Navigation*, 152(6):429 – 435, 2005.

[2] Sebastiaan P. Van Den Broek, Henri Bouma, and Marianne A.C. Degache. Discriminating small extended targets at sea from clutter and other classes of boats in infrared and visual light imagery. volume 6969, Orlando, FL, United states, 2008. SPIE.

[3] Donald Ross. *Mechanics of Underwater Noise*. Pergamon Press, 1976.

[4] Robert J. Urick. *Principles of Underwater Sound*. McGraw-Hill, 3rd edition, 1983.

[5] Rustam Stolkin, Alexander Sutin, Sreeram Radhakrishnan, Michael Bruno, Brian Fullerton, Alexander Ekimov, and Michael Raftery. Feature based passive acoustic detection of underwater threats. volume 6204, Kissimmee, FL, United states, 2006. SPIE.

[6] Brian Borowski, Alexander Sutin, Heui-Seol Roh, and Barry Bunin. Passive acoustic threat detection in estuarine environments. volume 6945, Orlando, FL, United states, 2008. SPIE.

[7] George Ogden, Lisa Zurk, Martin Siderius, Eric Sorensen, Josh Meyers, Shari Matzner, and Mark Jones. Frequency domain tracking of passive vessel harmonics. *The Journal of the Acoustical Society of America*, 126(4):2249–2249, 2009.

[8] Leslie M. Gray and David S. Greeley. Source level model for propeller blade rate radiation for the world's merchant fleet. *Journal of the Acoustical Society of America*, 67:516 – 522, 1980.

[9] Paul T. Arveson and David J. Vendittis. Radiated noise characteristics of a modern cargo ship. *Journal of the Acoustical Society of America*, 107(1):118 – 129, 2000.

[10] Sunghan Kim, Lars Holmstrom, and James McNames. Multiharmonic tracking using marginalized particle filters. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, pages 29 –33, 20-25 2008.

[11] Corentin Dubois and Manuel Davy. Joint detection and tracking of time-varying harmonic components: A flexible bayesian approach. *IEEE Transactions on Audio, Speech and Language Processing*, 15(4):1283 – 1295, 2007.

[12] James D. Wise, James R. Caprio, and Thomas W. Parks. Maximum likelihood pitch estimation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-24:418 – 423, 1976.

[13] M.J. Ross, H.L. Shaffer, and A. Cohen. Average magnitude difference function pitch extractor. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP22(5):353 – 362, 1974.

[14] J.A. Moorer. The optimum comb method of pitch period analysis of continuous digitized speech. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-22:330 – 338, 1974.

[15] Theodore W. Anderson. *An Introduction to Multivariate Statistical Analysis*. Wiley-Interscience, 2nd edition, September 1984.

[16] Dan Simon. *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*. Wiley & Sons, 1. edition, August 2006.

[17] Stefano P. Coraluppi and Doug Grimmett. Multistatic sonar tracking. volume 5096, pages 399–410. SPIE, 2003.

[18] Eric Sorensen, Helen H. Ou, Lisa M. Zurk, and Martin Siderius. Passive acoustic sensing for detection of small vessels. submitted to IEEE Oceans, Seattle, 2010.

Appendix A

Boat Photos

(a) Boat A

(b) Boat B

(c) Boat C

(d) Boat D

(e) Boat E

Figure A.1: Camera images of five of the six boats under test. There is no photo for Boat F.

Appendix B

Detection Algorithm Input Parameters

The input parameters for the detection algorithm are put in fields of a data structure. The fields of the structure are as follows:

snapshot_window
* T, length of the snapshot window in seconds

overlap
* percent for snapshots to overlap

frequency_limits
* frequency limits of spectrogram, $[f_{min} f_{max}]$ in Hertz

normalizing_window
* W, number of bins to normalize the STFT

peak_det_thr
* λ_f, normalized data detection threshold

NM
* number of snapshots required for a detection, $NM - 1$ out of NM

filter_freqs
* frequencies associated with known noise sources

boat_thr
* N_d, minimum number of detected frequency bins per snapshot

clust_space
* T_g, maximum spacing for grouping detection clusters in seconds

clust_length
* T_l, minimum total length of detection clusters in seconds

Appendix C

HEAT Algorithm Input Parameters

The input parameters for the HEAT algorithm are put in fields of a data structure. The fields of the structure are as follows:

snapshot_window
* T, length of the snapshot window in seconds

overlap
* percent for snapshots to overlap

frequency_limits
* frequency limits of spectrogram, $[f_{min} f_{max}]$ in Hertz

normalizing_window
* W, number of bins to normalize the STFT

NM
* number of snapshots required for a detection, $NM - 1$ out of NM

fomin, fomax
* γ_{min}, γ_{max}, limits for fundamental frequency search window in Hertz

delta_fo
* $\Delta\gamma$, spacing of fundamental frequency search window in Hertz

corr_det_thr
* λ_γ, detection threshold for fundamental frequency correlation

track_window
* W_γ, frequency blanking window for track creation in Hertz

Chi_thres
* χ^2 threshold

R_amplitude, R_frequency
* R_ρ, R_γ, ρ and γ measurement noise covariance

Q_amplitude, Q_frequency
* Q_ρ, Q_γ, ρ and γ process noise covariance

stop_track_thr
* N_F, threshold for track termination in number of snapshots

Appendix D

HEAT Algorithm Output File Format

The HEAT algorithm outputs a data structure and saves the file in a folder with the following format: `..\yyyymmdd\HHMMSS\TARGET_HHMMSS.mat` where `yyyymmdd` is the date the detection occurred on and `HHMMSS` is the start time of the detection. The fields of the structure are as follows:

Event_ID
* start date and time of the event, yyyymmdd_HHMMSS

Event_Duration
* length of the event, in seconds.

parameters
* structure containing all the input parameters of the HEAT algorithm.

stft
* short-time Fourier transform of the data, size size $[N_s \times K]$.

time
* time vector for the target, size $[1 \times K]$.

frequencies
* frequency vector, size $[N_s \times 1]$.

fos
* fundamental frequency vector, size $[N_\gamma \times 1]$.

fo_correlation
* fundamental frequency correlation for the data, size $[N_\gamma \times K]$.

tracks
* all the fundamental frequency tracks for the target, size $[N_r \times K]$.

track_fit
* harmonic content parameter for all the tracks, size $[N_r \times 1]$.

harmonic_amplitude_signature
* harmonic amplitude signature for all the tracks, size $[N_r \times H]$.

harmonic_noise_background
* average noise around each harmonic for all the tracks, size $[N_r \times H]$.

harmonic_snr_signature
* harmonic SNR signature for all the tracks, size $[N_r \times H]$.