

A Comparison of Detection and Classification Techniques of Cuvier's Beaked Whales in Passive Acoustic Monitoring Data

Ellen Jacobs, University of California, San Diego

Mentors: Danelle Cline and John Ryan

Summer 2016

Keywords: Passive Acoustic Monitoring, Wavelets, PAMGuard, Cuvier's Beaked Whales

ABSTRACT

Passive acoustic monitoring is an emerging field in the realm of marine mammal research that provides unique opportunities to observe cetaceans in their own environment. Increased volumes of data and advances in acoustic technology have both necessitated and facilitated the development of a variety of automated signal detection methods. However, the efficacy of these methods is not always easily known. Our study sought to analyze and compare two different methods of cetacean signal detection and species classification, PAMGuard software and wavelet analysis, on their ability to detect and classify Cuvier's beaked whale echolocation pulses. Both methods were applied to two different files of MARS hydrophone data to determine their accuracy in detecting and classifying expert-annotated ground truth clicks and their agreement with each other. To determine pure classification accuracy, they were also applied to a file of concatenated ground truth clicks. Finally, to determine the effect of a persistent 50 kHz tone from the MARS power supply, the concatenated file was filtered to remove the tone and two noise overlays were applied before being run through PAMGuard. Our results indicate that PAMGuard has a strong click detector, whereas the wavelets method has a more accurate classifier. Both methods, however, are shown to be strongly affected by the ambient acoustic environment. Information from this analysis can be used to better inform future efforts in automating cetacean acoustic signal detection and classification.

INTRODUCTION

Passive acoustic monitoring is a field of study with great potential for application in cetacean research. Cetaceans are traditionally observed with visual surveys, but this limits detection to when they choose to come to the surface, making the surveys unsuitable for deepdiving species. In addition, the acoustic detection range in the ocean is much larger than the visual detection range, as sound propagates much further than light through the water (Marques *et al.* 2013). Passive acoustic monitoring is also not limited by weather, as visual surveys are, and fewer human labor hours are required to collect data as hydrophones can be left for long periods of time using battery packs or connected to power sources via cabled observatories (Marques *et al.* 2013). The Monterey Accelerated Research System (MARS), with a hydrophone attached to one of its eight nodes, provides continuous passive acoustic data to the Monterey Bay Aquarium Research Institute, creating a much larger and more comprehensive dataset than is possible with typical short-term deployments. Recent advances in hydrophone and signal processing technology have allowed for more and more advances in detection and classification of cetacean acoustic signals, opening up many new possibilities for marine mammal research.

One cetacean particularly well suited to passive acoustic monitoring is the Cuvier's

beaked whale (*Ziphius Cavirostris*). A member of the family Ziphiidae, Cuvier's beaked whales are the most common of the beaked whale species and are found in tropical to temperate offshore waters globally (Allen *et al.* 2011). They eat cephalopods or mesopelagic fish, and tend to measure between 5 and 7 meters (Tyack *et al.* 2006). Cuvier's beaked whales hold

the world record for longest and deepest dive by a mammal, with a record-breaking



Figure 1: A Cuvier's beaked whale, courtesy of the Cetacean Research & Rescue Unit, Banff, Scotland.

2992 m and 137.5 minutes (Schorr *et al.* 2014). This deep-diving behavior makes Cuvier's beaked whales difficult to observe with traditional visual surveys, as they have short surface intervals between their long dives (Schorr *et al.* 2014). In addition to their deep diving behavior,

their tendency to echolocate consistently through the day and night and active avoidance of boats also make them good candidates for passive acoustic monitoring (Baumann-Pickering *et al.* 2014). They are known to be strongly affected by Mid-Frequency Active Sonar, frequently used in the Navy's Anti-Submarine Warfare training, which has greatly increased public interest in the role of sound in the whales' lives (Tyack *et al.* 2011).

Cuvier's beaked whales use upswept frequency-modulated pulses for echolocation. These pulses are species-specific, meaning passive acoustic data can be used to detect their presence without the accompaniment of visual detections (Baumann-Pickering *et al.* 2014). Their interclick intervals tend to be around 0.4 ms, longer than most other beaked whale species (Zimmer *et al.* 2005). The pulses generally occur between 15 and 70 kHz, with peak frequencies generally around 40 kHz (Baumann-Pickering *et al.* 2014). The clicks exhibit strong directionality, meaning that the probability of detection is greatly increased when the whales are directly facing the hydrophone but strongly decreased when they are off-axis (Zimmer *et al.* 2005).

Many passive acoustic surveys of Cuvier's beaked whale population density have been attempted in recent years. A key limitation on these studies, however, is the actual detection and classification of clicks once the data has been collected. Our study sought to compare two methods for click detection and classification, an open-source software known as PAMGuard, and a method of signal comparison known as wavelet analysis. PAMGuard analyzes data by detecting high amplitude signals matching a set of user-specified parameters, whereas wavelet analysis creates an image of a click based on its similarity to a known signal that is then compared to known Cuvier's beaked whale clicks to find a percent similarity.

MARS data was analyzed using MATLAB's signal processing toolkit to find optimal parameters for PAMGuard classification of Cuvier's beaked whales, determining strengths and weaknesses of the user-friendly software. After identifying probable click events from expert annotations, files containing the events were run through PAMGuard. This data was analyzed against data from wavelet analysis provided by a collaborator to identify as many Cuvier's beaked whale clicks as possible. The output from each detection method was then analyzed in comparison to expert annotations functioning as ground truth to determine accuracy. PAMGuard was found to have the most powerful click detector, while wavelet analysis was found to be the most accurate method of classification. In addition, the accuracy of detection and classification of both methods was greatly affected by both the background acoustic environment and the

presence of a 50 kHz tone from the MARS cabled observatory power supply. This information will hopefully aid in attempts to automate marine mammal acoustic detections in MARS data.

MATERIALS AND METHODS

MARS DATA



Figure 2: A diagram of the Monterey Accelerated Research System.

Acoustic data was recorded with a digital broadband hydrophone connected to the Monterey Accelerated Research System (MARS). Data is streamed in ten minute WAV files to a computer at MBARI, where it is ready for analysis. The two files used in this analysis were recorded on 28 October, 2015, from 13:29-13:39 and 19:09-19:19. The file beginning at 13:29 was selected for this analysis because of its high signal to noise ratio, and the file beginning at 19:09

was selected for this analysis because of existing expert annotations for the data.

GROUND TRUTH

Ground truth for the low signal to noise file was determined from annotations by Tetyana Margolina, of the Naval Post-Graduate School in Monterey, California. The file was opened using the Triton Software Package for MATLAB and the file's spectrogram was analyzed at a fine scale for beaked whales as well as other cetacean echolocation pulses. The time of each echolocation pulse was recorded by hand in a Microsoft Excel spreadsheet. As the duration of the detected pulse was not recorded in the original annotation, for this analysis a padding of 100 samples on each side of the annotation was included to comprise the duration of the pulse.

PAMGUARD

PAMGuard is an open-source passive acoustic monitoring software developed in collaboration with the University of St Andrews' Sea Mammal Research Unit. PAMGuard detects cetacean clicks by first passing the raw acoustic data through a bandpass filter, which

removes all the data outside a user-specified frequency range that can be tailored to the species of interest. Individual events are identified with a minimum-amplitude trigger, then verified with other user-specified parameters. Once clicks have been detected, they are passed through species classifiers. The classifiers look at parameters such as click length, specifying the duration of the click, energy bands, specifying the proportion of energy present in different frequency bands, peak and mean frequency, width of the peak frequency, the number of zero crossings, and the presence and magnitude of frequency upsweep. A classifier was specifically tuned to pick up Cuvier's beaked whale clicks in MARS hydrophone data.

WAVELET TRANSFORMS

Wavelet analysis is a method of representing an acoustic signal in the time and frequency domain. Because it is impossible to know the exact frequency at an exact point in time of a signal, signals must be cut into parts to analyze the frequencies in windows of time. Wavelet transforms provide a way to look at a signal with a resolution matching the scale of the part by comparing the signal at different scales to a known signal, in this analysis a Daubechies wavelet. The ability to change the scale of the frequency allows for greater time and frequency resolution than traditional Fourier transforms. This analysis against a known signal creates an image of the click, which can be compared to an image of an ideal beaked whale click that has been created based on expert-specified parameters. Because of this methodology, wavelet analysis does not have the same differentiation between detection and classification, as the steps are carried out simultaneously. Other versions of wavelet analysis can compare each signal to expert-identified beaked whale clicks rather than the artificially created click used for this analysis. The amount of similarity between the two click images informs whether a classification is made for each signal.

COMPARISON

The time of peak amplitude was recorded for each of the positively identified clicks for the ground truth data from the October 28th 19:09 recording, the PAMGuard data from both October 28th recordings, and the wavelets data from both recordings. For the first analysis, PAMGuard data and wavelets data were analyzed in terms of the ground truth. For the second analysis, PAMGuard and wavelets data were directly compared for the 19:09 recording and the

13:29 recording. Finally, the two methods were tested on a file containing only the clicks identified by the ground truth annotations with a padding of 300 samples on either side.

For the PAMGuard comparison to the ground truth, a nearest neighbor search was performed on both the PAMGuard data and the ground truth to determine whether each identified click had been identified in the other dataset as well. Distance to the nearest click was obtained for each click in the dataset, and clicks with a neighbor within a distance of 500 were counted as having matches. From there, a new dataset with all the unique click detections by both methods was compiled from the nearest neighbor search, with detections made only by PAMGuard, only by the ground truth, or by both, noted as different categories. The number of clicks in each category was then divided by the total number of unique clicks detected in the clip to get the percentages classified by each method. The same protocol was also used to compare wavelets to the ground truth.

For the comparison of the detectors' performance in the low signal to noise environment to the high signal to noise environment, a nearest neighbor search was performed on the wavelets data in comparison to the PAMGuard data to determine whether individual clicks had been identified in both datasets. From there, the total number of unique clicks detected in each file, both by the methods separately and together, was calculated. As ground truth data was not available for the high signal to noise file, accuracy measurements could not be calculated, and instead percent agreement between methods was obtained.

The analysis of the methods on the concatenated ground truth clicks was done by pulling clicks from the annotated file and putting them in a single file that was run through the two detection methods. The number of classifications in this data from each method was used to obtain pure classification evaluations, as unaffected as possible by differences in their detection power.

FILTERING

The effects of the 50 kHz hum present in all MARS hydrophone data were determined by filtering out the 50 kHz hum using a low-pass filter that attenuated all frequencies over 49.5 kHz. This filter was applied to the concatenated ground truth click file and then run through PAMGuard to see the effect of a file without that noise. Then, to see if the differences between the filtered and unfiltered data were due to the presence of a noise specifically at 50 kHz or due

to the presence of any noise, a layer of white Gaussian noise of the same amplitude (2 dB) was applied to the filtered data. To see the effects of a higher amount of noise, a layer of white Gaussian noise at a much higher amplitude (6 dB) was added to the filtered data. All four files were then run through PAMGuard.

To determine whether the classified clicks were actual classifications of the ground truth clicks or random noise, a k-nearest neighbor search was performed on the filtered, 2 dB, and 6 dB files to compare distance to the unfiltered ground truth clicks. A distance of fewer than 500 units was determined to be a match. This search was used in conjunction with a visual inspection of the clicks to determine accuracy of the classifications.

RESULTS

| | Ground Truth vs PAMGuard | Ground Truth vs Wavelets |
|---------------------|--------------------------|--------------------------|
| Total Unique Clicks | 76 | 47 |
| Method Only | 43 | 14 |
| Ground Truth Only | 13 | 8 |
| Shared Detections | 20 | 25 |
| Accuracy (%) | 26.31 | 53.19 |

Table 1: PAMGuard vs Wavelets



Wavelets performed much better in terms of matching the ground truth in the 19:09 file, with roughly twice the accuracy of PAMGuard, sharing 25 of 47 detections to PAMGuard's 20 of 76. This resulted in an accuracy of 26.31% for PAMGuard and 53.19% for wavelets. For both methods, clicks detected only

Figure 3: The number of total classifications for each method. Yellow represents the number of clicks picked up by both methods, while purple represents the number of clicks picked up by either the method or the ground truth only.

by the ground truth represented the smallest proportion of clicks, with 13 and 8 for PAMGuard and wavelets, respectively.

PAMGuard also detected about three times as many unique clicks as wavelets, with 43 clicks to wavelets' 14. The rate of wavelets detecting clicks rather than noise was determined to be 100%, but in PAMGuard the rate was estimated to be around 95%.



Figure 4: The number of ground truth classifications detected by each method. Orange represents the number of clicks classified by the method, while blue represents the number of ground truth clicks the method missed.

| Table 2: Low Signal to Noise vs High Signal to Noise | | | | |
|--|---------------------|----------------------|--|--|
| | Low Signal to Noise | High Signal to Noise | | |
| Total Unique Clicks | 74 | 2665 | | |
| PAMGuard Only | 34 | 1331 | | |
| Wavelets Only | 11 | 1199 | | |
| Shared Detections | 29 | 135 | | |
| Percent Match (%) | 39.19 | 5.07 | | |

Table 2: Low Signal to Noise vs High Signal to Noise



Figure 5: Number of detections by both methods combined, for the low signal to noise file and the high signal to noise file.

The analysis of the file with a low signal to noise ratio versus a high signal to noise ratio revealed significantly higher agreement between PAMGuard and Wavelets for the low signal to noise file, with 39% agreement in the low and 5% in the high signal to noise file. Without ground truth for the file, however, accuracy is not possible to determine. The high signal to noise file had over 30 times more detections than the low signal to noise file, though they shared only four times as many classifications in the high signal to noise file as in the low signal to noise ratio. The error rate for PAMGuard in the high signal to noise ratio was lower than the low signal to noise ratio, about 67%



Figure 6: The percent of total unique detections classified by each of the methods. Yellow represents the percent classified by both, wavelets represents the percent classified only by wavelets, and purple represents those only classified by PAMGuard.



Figure 7: Total number of detections of ground truth clicks. Purple represents those classified by the method, while yellow represents those missed by the method.

The comparison of the two methods on the file of only the 33 concatenated ground truth clicks demonstrated that PAMGuard's detector has no effect on its classifier, as PAMGuard detected 20 of the ground truth clicks in both the full file and the concatenated ground truth click file. Wavelets, however, found 8 fewer ground truth clicks in the concatenated file than it did in the full file, for a total of 17.

| | Detections | Classifications |
|-----------------|------------|-----------------|
| Unfiltered Data | 33 | 20 |
| Filtered Data | 61 | 28 |
| 2 dB of Noise | 33 | 21 |
| 6 dB of Noise | 33 | 20 |

The filtered data came up with roughly twice the number of detections as the unfiltered and noisy data, with 61 to the other files' 33. There were 33 clicks in the expert ground truth

accuracy.

annotations and concatenated ground truth file. The filtered data found 28 classifications, more than the 20, 20 and 21 of the unfiltered, 2dB, and 6 dB files, respectively. All 20 of the classifications made in the unfiltered file were true detections, and 20 of the 28 detections in the filtered data were true detections. Because the process of adding noise alters the waveform of the clicks, a nearest neighbor test was performed on the 2 dB and 6 dB files to determine whether the

21 and 20 detections were the same clicks as detected in the unfiltered data.

The number of these classified clicks that were actually detections of the ground truth clicks varied depending on whether it was determined by visual inspection or the nearest neighbor search. For the filtered data, the nearest



Figure 8: Total number of detected ground truth clicks by each file. Red represents unclassified detections, while green represents detections classified as beaked whales.

neighbor search determined 23 of the 28 classifications were close enough to count as matches, whereas the visual inspection determined that only 20 classifications were true detections. For the 2 dB file, 19 matches were determined by the nearest neighbor search and 19 were determined by visual inspection. For the 6 dB file, 17 matches were determined by the nearest neighbor search and 17 were confirmed by visual inspection.

DISCUSSION

PAMGUARD VS WAVELETS

The results of the comparison of PAMGuard and wavelets to ground truth suggest that the wavelet method is more accurate than PAMGuard because it detected a larger number of the ground truth clicks than PAMGuard did. This difference could be due to the wavelet transform retaining more information about an individual click than PAMGuard looks at when determining whether the click falls within the specified parameters. However, the margin between the two methods was very similar, as only 5 clicks separated the two methods. In order for the difference between the two methods to be definitely significant, a larger sample size of clicks to compare would be necessary. In addition, it was revealed in further analyses that the addition of 100 samples on either side of the expert annotation marking was not sufficient to fully capture every annotated click. As the clicks were compared based on peak amplitude, which generally occurs at the middle of the signal, this is unlikely to have a significant effect on this analysis. There is, however, the possibility that the peak amplitude of the signal was located in a part of the click not captured by the analysis, so it remains worth noting.

The larger number of clicks detected by PAMGuard suggests that PAMGuard may have a more powerful detector than the wavelets method. From an analysis of those classifications, the rate of classifying random noise appears low, meaning that PAMGuard uncovered a larger number of potential beaked whale clicks than wavelets. Not all of the clicks classified by PAMGuard are likely to belong to beaked whales, but a strong majority was likely to belong to some sort of cetacean and was worth further classification analysis. A similar analysis of the wavelets classifications revealed that the clicks detected by wavelets, although fewer in number than those detected by PAMGuard, were all likely beaked whale clicks. This provides further support for the idea that the wavelets classifier is more powerful than that of PAMGuard.

LOW VS HIGH SIGNAL TO NOISE RATIO

No accuracy judgment is possible in regards to the results of the low and high signal to noise comparisons, due to the lack of ground truth for the high signal to noise file. However, a larger proportion of shared detections in the low signal to noise file suggests that the accuracy of both detectors was likely higher, as the probability of the detectors agreeing on the detection of a real click is likely higher than the probability of both agreeing on the same piece of random noise. From an analysis of the rate at which PAMGuard detected random noise instead of actual potential clicks, PAMGuard appeared to be performing worse in the high signal to noise environment. The lack of ground truth was very limiting for this particular analysis.

Although true accuracy measurements cannot be determined from this analysis, information can still be obtained about the performance of the detectors in different sound environments due to the considerable difference in the amount of agreement between the methods. A decrease in agreement from over 40% down to 5% means a significantly different performance by the methods, showing a weakness in the consistency of the two methods. For the

ultimate implementation of a method as a means for automation of detection and classification, a consistent performance through any acoustic environment is optimal, and PAMGuard and wavelets are not consistent.

CONCATENATED CLICK ANALYSIS

To test the pure classification from each of the two methods without the effect of the detectors, a file of only the annotated ground truth clicks concatenated, normalized and separated by padding on either side, was created and run through the two methods. However, neither method picked up more ground truth clicks in the concatenated file than it did in the full clip, which instead provides information on the effect of the click detection on the later classification. PAMGuard classified the same number of clicks in the full as the concatenated file, so PAMGuard's detection process is shown to not affect its later classification process. As wavelets' performance classifier is shown to be dependent on having detected the click on its own, likely in part because of the calculation of the mean energy, which would change depending on where the start and end points of the click are defined.

An important consideration in this analysis, however, is the process of concatenation. All the clicks had to be normalized in order to be evenly padded as separation, which fundamentally altered the waveform and mean amplitude of the signals. This did not appear to alter PAMGuard's performance, although it lowered that of wavelets. To try to correct for some of this effect, the threshold for matching the ideal beaked whale click was lowered, decreasing the accuracy of the classifier. The effect that detection has on wavelets' classifier accuracy potentially casts doubts on whether the wavelets classifier could be used in a joint application with the PAMGuard detector.

GROUND TRUTH

A fundamental issue with this analysis can be found in the lack of accurate ground truth. For most passive acoustic monitoring studies, human expert annotations are used as ground truth. The first issue with this arises from the shortage of annotated data. To do a real evaluation of any automated detection, annotated data is necessary, but this annotation can take many times longer than real time and is thus at a premium. A second issue arises with using human expert

annotations as ground truth because it assumes that the human annotator does not miss any signals and correctly identifies everything, which is a very challenging task in itself. In some situations, a method of automation might correctly detect or classify clicks better than a human annotator can because the clicks are difficult to see on a spectrogram, meaning there is no truly accurate ground truth against which data can be compared.

Another issue with human annotated ground truth is the lack of consensus within the beaked whale community about where exactly the limits of what counts as a Cuvier's beaked whale lies. Many experts believe that a distinctive notch in the waveform is necessary for a positive identification, whereas others do not. The human annotator who provided ground truth for this analysis believes the notch, in addition to a strong frequency upsweep, is necessary for a positive identification. Although the creator of PAMGuard agrees that a strong frequency upsweep is a necessary part of classification, another collaborator on the project believes that not all Cuvier's beaked whale calls are frequency upsweep as a parameter for detection.

Because of these issues with the expert annotations' accuracy and consistency with other experts, it is difficult to use it as ground truth. A judgment on the agreement between the three sources of data, the two methods and the ground truth, is not totally possible without a verification that all three sources are looking for the same thing.

FILTER ANALYSIS

The results of the filter analysis were surprising, as the hypothesis had been that removing the 50 kHz noise would decrease the number of noise detections. However, similar to how the low signal to noise file had fewer detections than the high signal to noise file with a lower error rate, the unfiltered data had half the number of detections of the filtered data. As the unfiltered and noise-overlaid files detected 33 clicks and there were only 33 clicks in the concatenated file to begin with, detecting twice as many clicks means the error rate on the filtered data is higher than it is in the other files. An inspection of the 28 clicks detected on the filtered file revealed that 20 of the classifications, all of which were true classifications, were in the same amplitude range as those classified in the unfiltered file, and the 8 extras were random noise detected in a slightly lower amplitude range. This means that the ground truth click true classification rate is the same between the filtered and unfiltered data. The fact that the number

of clicks detected went back down to 33 from the doubled rate when Gaussian white noise at either amplitude was added is evidence that it is likely the presence of noise, rather than the frequency of the noise, that affects the performance of PAMGuard.

The differences between the nearest neighbor test and the visual inspection in estimating the number of true detections for the filtered data is likely due to the nature of the random noise classifications. As there is no random noise in between the clicks in the concatenated ground truth click file, and random noise detections would have to be very close to the clicks in the added padding. For the three clicks that were determined by the nearest neighbor test to be true but by the visual inspection to be false, it is likely that the noise that was detected happened to be close enough to the click itself that its distance fell within the required threshold and erroneously counted.

With the rates of error included into the analysis of the classification accuracy for these files, it is evident that the detections and classifications in the unfiltered data are the most accurate, with the filtered file, 2 dB file, 6 dB file proving to have higher rates of false detection and classification or lower accuracy. This is consistent with the analysis of PAMGuard's performance in the low and high signal to noise files. Although intuitively a click detection and classification tool might be expected to function better when there is less background noise, PAMGuard appears to function better in noisy environments. This may be intentional, as PAMGuard was designed to filter raw data from the ocean, which is already known to be noisy. This may, however, be an artifact of the parameter-based detection. When there is not a loud tone to drown out the random noise, it is possible that more random pieces of noise slip through the parameters.

These filtered files were not run through wavelets for analysis because of the poor performance of wavelets on the unfiltered concatenated file. PAMGuard's performance on the concatenated file proved the same as on the full file, so the effects of concatenation could be ignored.

OVERLAP ANALYSIS

In order to determine the difference in the marked start and end times of the clicks detected by each method, an attempt was made at finding the number of samples during the 10 minute analysis period in which two or more of PAMGuard, wavelets, or the ground truth

detected a click. However, due to the large number of data points produced by analyzing every sample of a ten-minute file at a sample rate of 256,000 samples per second, this analysis was abandoned as MATLAN crashed numerous times and could not display or export variables containing all the data. A method of determining overlap that does not involve creating a data point for each sample for each detection method could be applied to acquire this information. The overlap could be used to not only evaluate the consistency of start and end times between the two methods, but also determine an appropriate allowance for distance between clicks to be considered a match in the Nearest Neighbor test. However, until a different method of calculating the overlap is proposed, one including fewer than one data point per sample over the ten-minute file, the analysis will have to go unfinished.

CONCLUSIONS

Although PAMGuard and wavelets are both powerful tools, both have significant barriers before it would be efficient to automate them. PAMGuard's detector appears to be very effective but its classifier is not very accurate, in terms of replicating the results of an expert annotator. Wavelets, on the other hand, does not have as strong a detector, but has a slightly more accurate classifier. Both, however, are less than 60% accurate, which brings into question whether the process of automation is worth it for these two methods.

Another consideration is processing speed. On an adequately powered computer, PAMGuard computes fairly quickly, whereas wavelets take longer to run. In the context of running individual files this difference is not significant, but in the case of automation for running on data that is being streamed in constantly, even a slightly slower processing speed could lead to significant backup in data processing.

A recent development in the wavelets classifier is the use of real Cuvier's beaked whale clicks as comparisons for potential detections, rather than the artificially generated Cuvier's beaked whale click that was used in this analysis. This advance, while not improving the processing speed of wavelets, likely improves the accuracy of the wavelets method. Further analysis would be required to say whether the improvement affects wavelets' detection and classification enough to be definitively more effective than PAMGuard.

A future implementation of this analysis could be a workflow in which clicks are first detected by PAMGuard detector, then classified by wavelets. This joint method would combine

the stronger parts of both individual methods, creating a more powerful tool. This is complicated by the results of the methods working on the file of concatenated ground truth clicks, however, so a different method of compiling PAMGuard's clicks than what was used to compile the ground truth annotated clicks would have to be employed to maintain wavelets' higher classification accuracy.

The future of automated detection of cetacean acoustic signals likely lies in neural networks. Using a neural network would allow the detection to be dynamic through varying oceanographic conditions, which has been shown to strongly affect the performance of PAMGuard and wavelets. It would also eliminate the need for user identification of pertinent parameters, creating the most effective set of parameters instead with machine learning. A drawback of neural networks, however, is the need for a large training set. There is only just over a year of data from the MARS hydrophone currently, and based on whale watching data Cuvier's beaked whales do not appear to spend very long in the bay. However, a potential way around this need for data exists with the joint tool approach. The tool, half PAMGuard and half wavelets, could be used to extract all the potential beaked whale clicks from the year of data. This would provide a larger amount of training data for the neural network than could be acquired by the alternative of human annotation.

Once an accurate method of detecting and classifying cetacean echolocation clicks has been developed, the possibilities for cetacean research will expand considerably. Insight into their communication will likely have implications for research in population density, behavior, social dynamics, and many other fields of study. With more knowledge about these vital members of the ecosystem, more can be done to protect and conserve the ocean's cetaceans.

ACKNOWLEDGEMENTS

Thank you to my mentors, Danelle Cline and John Ryan, for their counsel and support, to my collaborators, Mark Fischer, Marjolaine Calliat, and Tetyana Margolina, for their data and guidance, to Nicholas Raymond for his help with MATLAB programming and signal processing, and George Matsumoto and Linda Kuhnz for facilitating the internship program.

References:

- Allen, B.M., Brownell, R.L., and Mead, J.G. (2011). Species Review of Cuvier's beaked whale, *Ziphius Cavirostris*.
- Baumann-Pickering S, Roch MA, Brownell Jr RL, Simonis AE, McDonald MA, Solsona-Berga A, Oleson EM, Wiggins SM, Hildebrand JA. (2014). Spatio-temporal patterns of beaked whale echolocation signals in the North Pacific. PLoS ONE. 9(1):e86072. 10.1371/journal.pone.0086072
- Marques, T. A., Thomas, L., Martin, S. W., Mellinger, D. K., Ward, J. A., Moretti, D. J., Harris, D. and Tyack, P. L. (2013). Estimating animal population density using passive acoustics. Biological Reviews, 88: 287–309. doi: 10.1111/brv.12001
- Schorr GS, Falcone EA, Moretti DJ, Andrews RD. (2014). First Long-Term Behavioral Records from Cuvier's Beaked Whales (*Ziphius cavirostris*) Reveal Record-Breaking Dives. PLoS ONE 9(3): e92633. doi:10.1371/journal.pone.0092633
- Tyack, P. L., M. Johnson, N. A. Soto, A. Sturlese and P. T. Madsen. (2006). Extreme diving of beaked whales. Journal of Experimental Biology 209:4238-4253.
- Tyack, P. L., W. M. X. Zimmer, D. Moretti, et al. (2011). Beaked whales respond to simulated and actual navy sonar. PLoS One 6.
- Zimmer, W. M. X., M. P. Johnson, P. T. Madsen and P. L. Tyack. (2005). Echolocation clicks of free-ranging Cuvier's beaked whales (*Ziphius cavirostris*). Journal of the Acoustical Society of America 117:3919-3927.