Dual Peak Detection and its Functions in Thoracic Identification

Paper and Research by Dylan Tokotch

Supervising Professor: Dr. Eric McGregor

Introduction

Historically, humans have identified other humans by observing one another's physiological and behavioral biometric traits. Examples of these traits include facial structure, body size and shape, skin pigmentation, hair color, and gait. Regardless of how reliable this process is for humans to identify other humans, this process can be rendered less successful by a variety of countering activities, such as applying disguises. However, there are several biometric traits common among humans that are not easily disguisable. In [1-10], the authors analyze the biometric trait that is present in every living human being and yet difficult to mimic or disguise: the human heartbeat.

In [1], the authors collected three types of data from 36 subjects via an experiment. The data included phonocardiogram (PCG) data, carotid pulse (CP) data, and electrocardiogram (ECG) data. The carotid data, the electrocardiogram, and the phonocardiogram were collected by a Laser Doppler Vibrometer, electrodes placed on the subjects' chests, and digital stethoscope, respectively. In addition, they developed the Thoracic Identification System (THIS), which performs biometric identification using the three previously stated signal types. Prior to this research, the Thoracic Identification System had a satisfactory accuracy when trained on a set of ECG signals and tested against ECG signals collected eight minutes later. But the identification rate when using carotid pulse data stood at about 61%. A large portion of this error was believed to be due to the inability of the classifier algorithm to properly and accurately account for the changes in speed, amplitude, and structure in the heartbeat that occur between training and testing data.

To mitigate these differences, Chen [9] experimented with Dynamic Time Warping and Optimal Subsequence Bijection. Chen's research, however, showed that implementation of both of these methods, whether on their own or combined together, showed no statistically significant improvement in the accuracy of their system.

The problem that both Dynamic Time Warping and Optimal Subsequence Bijection sought to solve was a desynchronization problem between heartbeats. Ideally, in order to compare two heartbeats, the signals would be of the same length with their distinguishing features being lined up at the same discrete sample/time placement in the signal. However, this is not usually true, even for different heartbeats of the same person. DTW and OSB are both algorithms that can be implemented in order to match the key features of one heartbeat to the respective matching features of another. But in [9] , the authors reasoned that these algorithms did not have statistically insignificant increases in accuracy because they altered the original signals too much and removed critical data from the signal, thus removing any accuracy improvement the system might have gained from the now easier to compare signals. Instead, the authors suggest detecting both the primary and secondary peaks of each heartbeat and then using the primary and secondary peaks of each signal as the key features for the alignment process.

In this work we implement a dual-peak detection algorithm and create a new feature selection algorithm following [9].

The Peak Detection Algorithm

The new peak detection algorithm is given below.

1. Detect_Primary_And_Secondary_Peaks(signal)
    a. Primary_Peaks = Detect_Primary_Peaks(signal)
        i. Iterate through a range of window sizes, 500 to 900 with an increment of 50, to find an optimal window size that minimizes the standard deviation between the distances of each detected maximum for each window.
        ii. Once the ideal window size is determined, detect the peaks by finding the maximums in each window.
    b. Secondary_Peaks = Detect_Secondary_Peaks(signal, Primary_Peaks)
        i. Compute the primary peak width for each primary peak to the left and to the right of each peak.
            1. Iterate to the left and to the right of each peak maximum by steps of 50 until there are two consecutive windows that are on average, not increasing (for the left peak width) or decreasing (for the right peak width).
        ii. Partition the signal into segments that are the primary peak and that are not the primary peak.
        iii. For each non-primary peak segment, find the local maxima and identify this as a secondary peak for the signal.
    c. Return the locations of the primary and secondary peaks.

Methodology of THIS

In order to understand the need for and application of our algorithm that we implemented, one must first understand the THIS system and its algorithms. The program, which is written in Matlab, reads in the signal data from preexisting files based on the options that the user has selected. These options include which data sets to train the classifier on and test the classifier against, as well as the kind of signal, the rate of down sampling, and the format of the data file. After these choices are made by the user, the data files are read into the THIS program. These signals, however, come into the program as a single 1 to 4 minute continuous signal, which could contain an average of anywhere from 50 to 300 heartbeats, depending on the heart rate of the subject at that point of recording.

The software then partitions the signal into individual heartbeats so that they can be properly analyzed by the classifier. The software uses a sliding window to mark the peak maximum for each window for a given window size chosen by the user through the GUI. At this point, the THIS program has a list of maximum peaks for the signal which coincide with the primary peaks of heartbeats.

Once the software has this series of maximum peaks, it takes a certain number of samples both before and after the primary peak for each window. Each window's set of samples is a single heartbeat for that subject. The parts of the signal that were not included in these heartbeats were discarded.

After the heartbeats are segmented for each subject, spectrograms were created for each heartbeat. A spectrogram is a visual representation of the frequencies in a signal; in this context, the spectrograms are visual representations of the heartbeats. In order to create a spectrogram from a given heartbeat, a discrete Fourier transform is performed on the signal, which partitions the signal into a number of time-variant bins, each of which is given an assigned value based on the frequencies of the signal in that interval of time.

After the spectrograms are created, two types of average, or mean spectrograms, are created. For each subject, a subject nominal spectrogram is created such that the spectrogram is the average of all of their spectrograms. For the population (all of the subjects), a single overall nominal spectrogram is created such that the spectrogram is the average of all of the subject's nominal spectrograms.

Among the user inputs into the THIS program are the number of spectrograms for both training and testing that are used in order to identify an individual. For each subject, the THIS program selects the n most similar spectrograms to the subject's nominal spectrogram, where n is the given parameter for number of spectrograms for training. Also for each subject, the THIS program selects the m most similar spectrograms to the subject's nominal spectrograms, where m is the given parameter for number of spectrograms for testing. These are defined as the subject's maximum likeliness spectrograms for training and testing, respectively.

Each of these spectrograms, of course, still have a one to one relation to their respective heartbeat signal from which they were created. These can be plotted to visualize the selection of the spectrograms (and by extension, heartbeats) that most accurately represent the average of the individual's heartbeats as shown in figure 1:
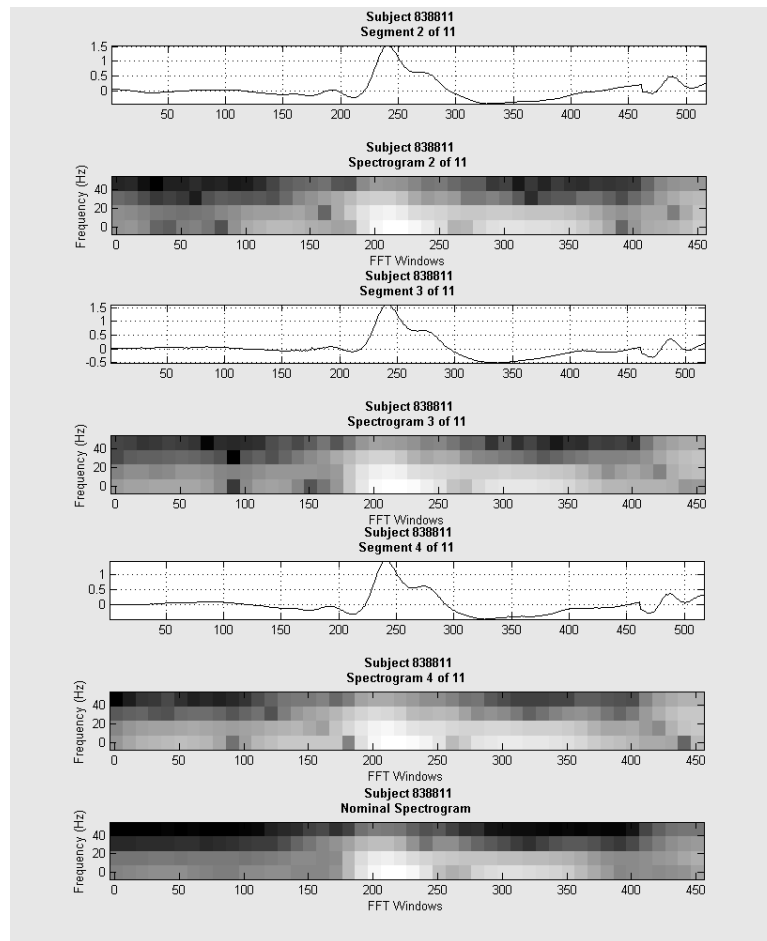
Figure 1 - Above is a screenshot of three heartbeats and their respective derived spectrograms from a particular testing subject. Also displayed is the nominal spectrogram for that subject.

From each subject's respective maximum likeliness spectrograms, Symmetric Relative Entropy Scores are calculated for each bin. A Symmetric Relative Entropy Score has two equally weighted parts. In order to uniquely identify each subject by their spectrogram information, we wish to discard the parts of their spectrogram (bins) that are common across the population, but keep the bins that are common across each subject. In other words, we wish to create a scoring system which gives more weight to bins that have high variance compared to the population nominal model than bins that have low variance compared to the population nominal model, as well as weigh bins that have low variance compared to the subject's nominal model higher than bins that have high variance compared to the subject's nominal model.

Once the Symmetric Relative Entropy Scores are calculated for each bin, the software creates a feature set consisting of a matrix containing the k highest scoring bins for each subject, where k is another parameter specified by the user. At this point, the classifier, whether testing or training, has n feature sets to train on and m feature sets to test with for each subject/signal that most accurately identify that subject uniquely.

Once both the training and testing features are calculated, the next step is for the classifier to actually classify. With the aforementioned feature sets for each subject, the

classifier creates a model for both testing and training. It then compares these models probabilistically and computes a probability for each testing signal to belong to each training subject. It then checks its own accuracy by comparing its most likely candidates to the actual source of each feature set and outputs its own accuracy, for the purposes of data collection and improvement.

Changes to Methodology of THIS

Firstly, before we began the changes that would become the dual peak detection algorithm in the software, we attempted many implementations of algorithms that did not improve the accuracy of the software. As described previously, the algorithms of Optimal Subsequence Bijection and Dynamic Time Warping were thought to be solutions for solving the problem of desynchronized key features of heartbeats that was believed to be the source of the software's inaccuracy.
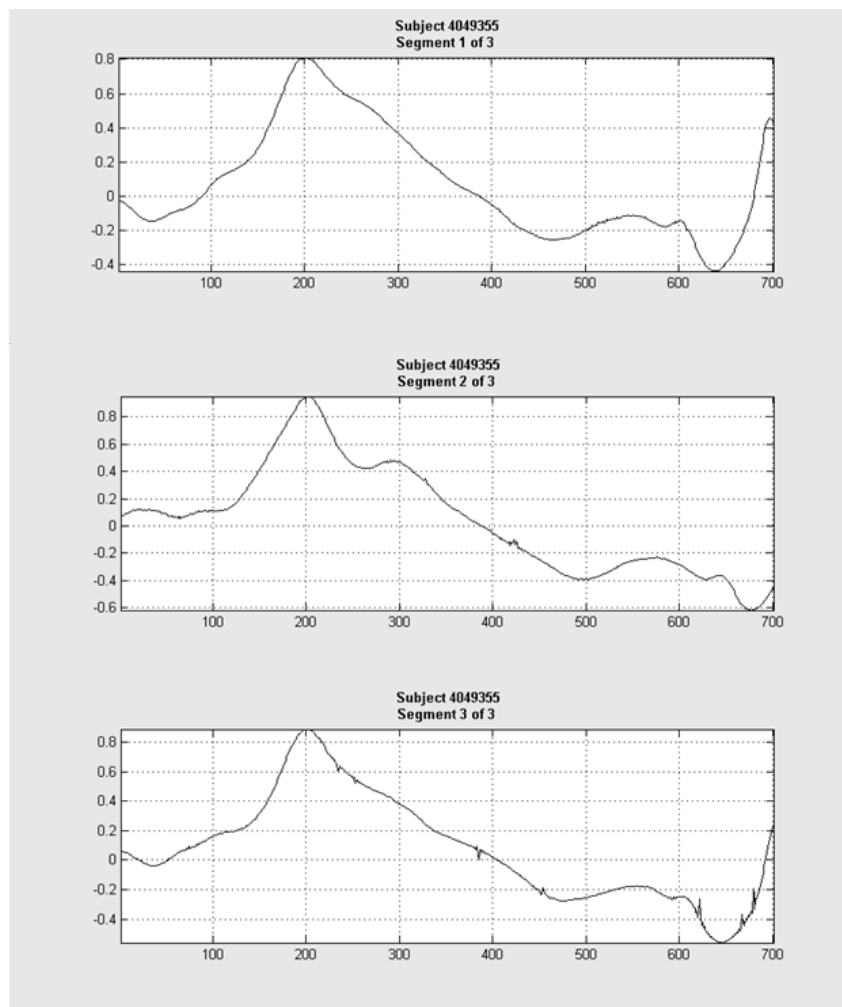

Figure 2 – Above are three heartbeat segments taken from the same subject.

Notice above that there is a two-fold problem here. Firstly, since the samples are determined solely by their relation to the primary peak key feature, only the primary peak key

feature is aligned between heartbeats. In this example, primary peaks will always occur at position 201 across all segments across all subjects. However, this means that the secondary peak is not bound to any particular point. This means that the secondary peak locations for multiple heartbeat segments of the same subject do not necessarily have to match. In addition to the variability of the location of secondary peaks, the figure above also shows an uncommon but problematic issue: the secondary peak can sometimes not even exist in the heartbeat segment because it is too far from the primary peak to be sampled. Since spectrogram bins are created from the data present at that location, this presents an information loss problem.

Initial delves into the implementation of OSB and DTW proved difficult, however, and further research into similar applications revealed that they did not statistically significantly improve accuracy. It became evident that DTW and OSB both remove too much of the information in LDV signals because of excessive changes to the original signal in order to make the key features of the signal match.

It was determined, however, that a less extreme variation of the signal key feature matching could likely be used without removing important information. And so, after an extensive amount of research, consideration, and deliberation, the algorithm that we decided to implement was a dual peak detection algorithm and segmentation algorithm. The purpose of the dual peak detection algorithm is to align all primary and secondary peaks; the purpose of the new segmentation algorithm is to concatenate segments surrounding the primary and secondary peaks of each heartbeat during the data selection process.   Our solution to the information loss problem was to create a hybrid concatenated signal that contained a segment around the primary peak and a segment around the secondary peak of each heartbeat. This effectively aligned the locations of the most information critical areas, that is, the primary and secondary peaks of each heartbeat, without skewing any of the other data.

We began by implementing a dual peak detection algorithm by Jose Corona and Dr. Verne Leininger. Firstly, we implemented their variation of primary peak detection. The advantage of their primary peak detection over our previous implementation was that, given a signal, the user wasn't required to input a window size in order to detect peaks. A poorly chosen user's window size produces many false positives and false negatives. Secondly, after implementing their algorithm for primary peak detection, we wrote and implemented another method for secondary peak detection based on conversations with Dr. Leininger.
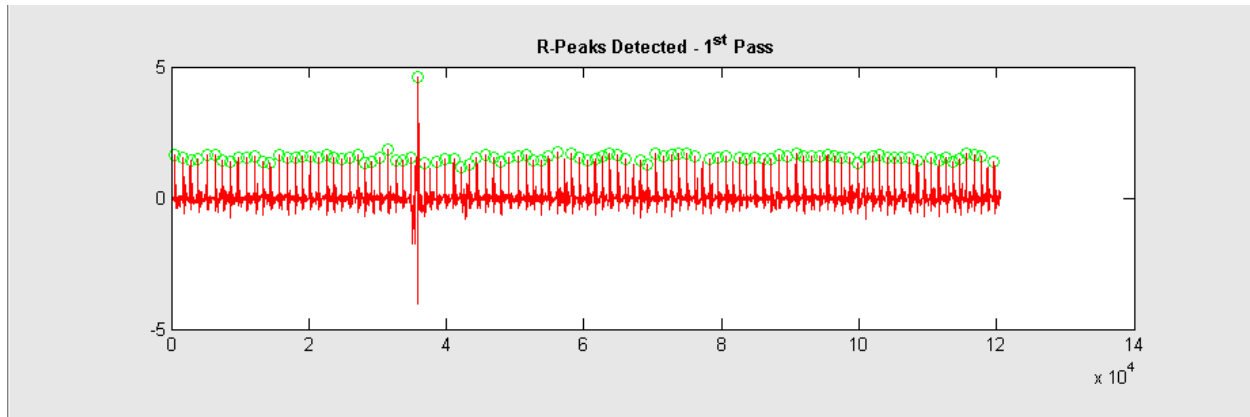
Figure 3 - Above is the prior singular peak detection algorithm with fitted window size parameter. Primary peaks are shown in green.
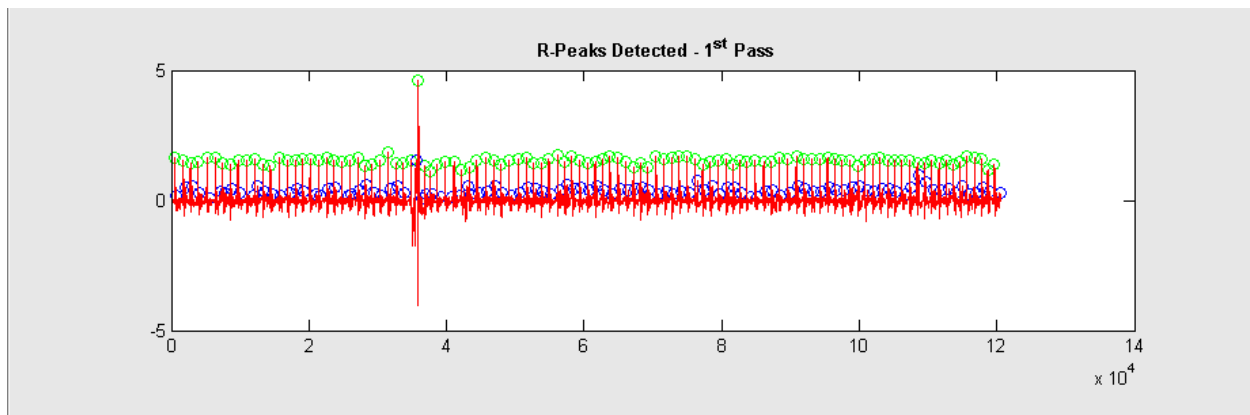

Figure 4 - Above is dual peak detection algorithm on same signal with primary peaks in green and secondary peaks in blue

Next, having implemented primary and secondary peak detection, we implemented the necessary changes to the software so that it would be able to segment a region around the secondary peaks of each heartbeat as well as around the primary peaks and concatenate them together to make a spliced signal before spectrogram creation.
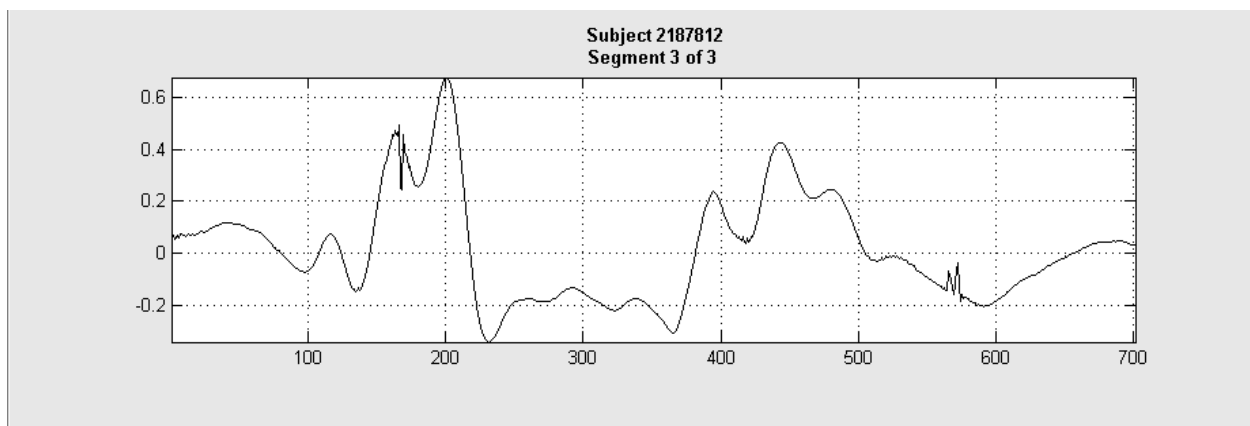

Figure 5 - Pictured above is the original implementation of segment selection. Parameters include 200 samples before and 500 samples after the primary peak, which consequentially places our primary peak at the 201st sample.
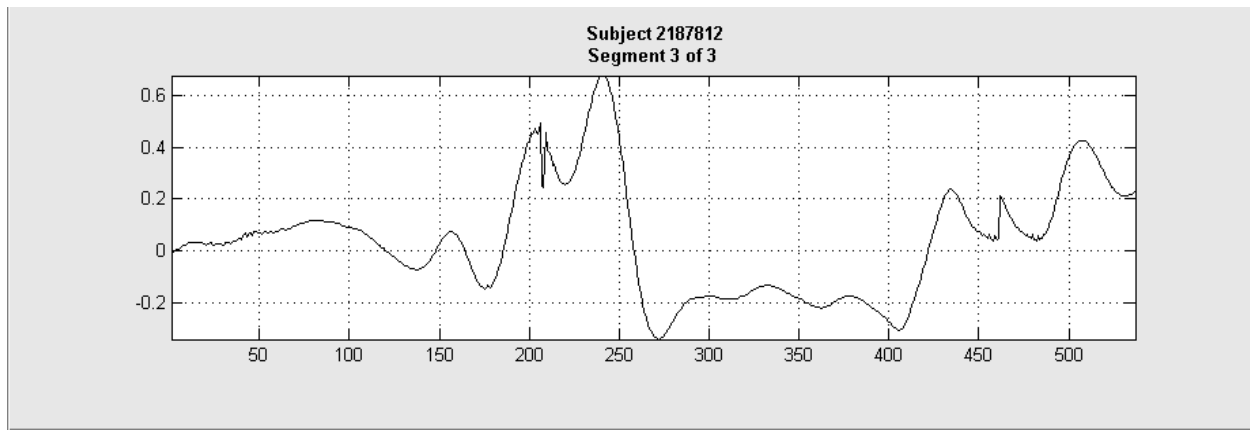
Figure 6 - Pictured above is the same heartbeat with new segmentation. Parameters include 240 samples taken before the primary peak and 220 after; it also contains 45 samples before the secondary peak and 30 after. The splice occurs at sample 462. The primary and secondary peak locations are located at the 241st and 506th samples, respectively.

Of course, the benefit of this implementation really shines when comparing heartbeat segments of different subjects:
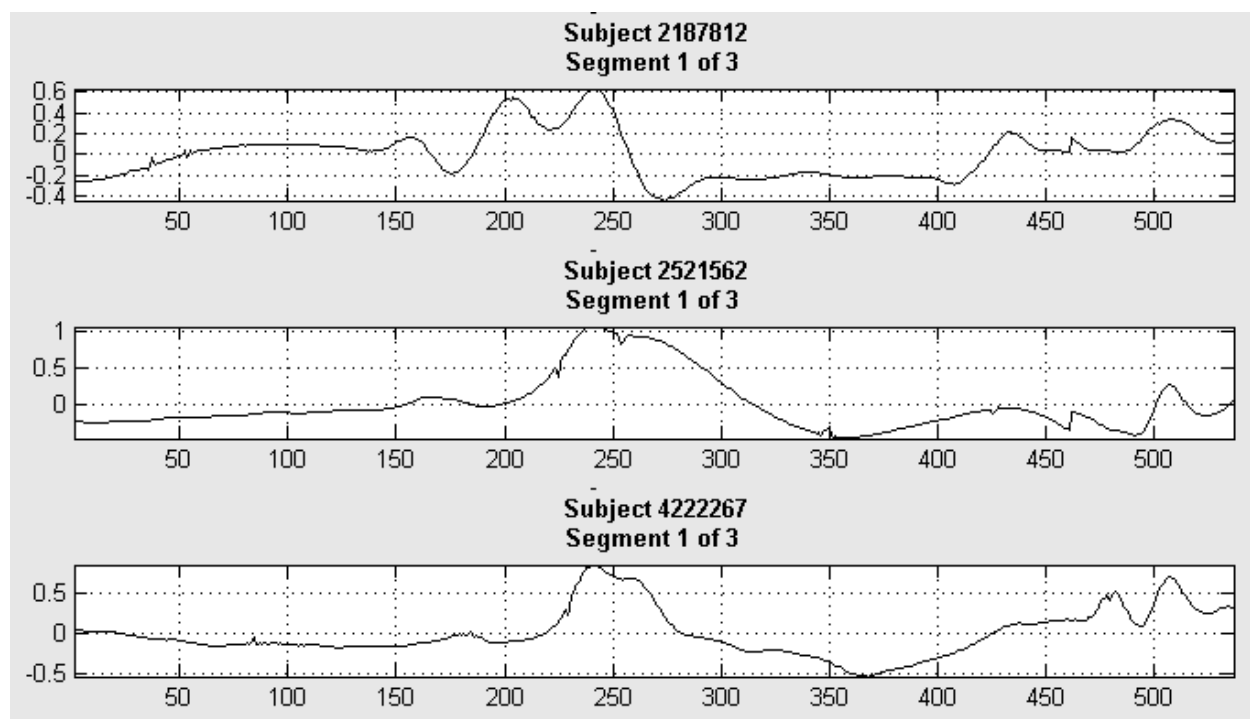
Figure 7 – Different subjects have the same locations for the primary peak and secondary peaks.

While the secondary peak detection algorithm aimed to increase the accuracy and decrease the loss of information, it does so at a cost: time. Since the secondary peak detection process was so computationally intensive, we also implemented a caching, saving, and loading system in order to speed up the aforementioned process. Caching allows the user to save as much as 15 minutes for 32 subjects.

In addition, numerous edits to the software's GUI were made in order to make the dialogues for additional input parameters, such as: added number of samples to take before and after the secondary peak, added option to use new peak detection algorithm, and removal of some of the unneeded parameters.

Lastly, we tested to the software by adjusting seven target parameters, namely: samples around the primary and secondary peaks, the number of bins selected for each model, and the number of spectrograms to use for training and for testing.

Testing Results

The following chart is a summary of the data created from the optimization of THIS program's parameters over 87 iterations.

Trial #	Peak Detection	# Train Segs	# Test Segs	# Bins	Before Primary	After Primary	Before Secondary	After Secondary	1st Rank Match Rate
0	O'Brien	3	13	27	200	500	N/A	N/A	61%
1	Tokotch	↑	13	27	200	200	200	200	66%
2	↑	↑	↑	↑	↑	↑	↑	100	↑
3	↑	↑	↑	↑	↑	↑	50	↑	69%
4	↑	↑	↑	↑	↑	220	↑	↑	↑
5	↑	↑	↑	↑	210	↑	↑	↑	↑
6	↑	↑	↑	↑	↑	↑	↑	↑	↑
7	↑	↑	↑	25	↑	↑	↑	↑	↑
8	↑	↑	15	↑	↑	↑	↑	↑	↑
9	↑	15	↑	↑	↑	↑	↑	↑	↑
10	↑	↑	↑	↑	↑	↑	↑	30	↑
11	↑	↑	↑	↑	↑	↑	25	↑	↑
12	↑	↑	↑	↑	↑	220	↑	↑	↑
13	↑	↑	↑	↑	240	↑	↑	↑	↑
14	↑	11	↑	↑	↑	↑	↑	↑	↑

Figure 8 – A table showing the streamlined optimization process for the Tokotch peak detection parameters.

Our final best 1st rank match rating maxed at 69%, only marginally higher than the 61% 1st rank match rating that was achieved through the prior single peak detection algorithm. However, this dual peak detection appears to have a much higher robustness rating in regards to its parameters. The O'Brien peak detection algorithm's parameters are very volatile and

require extremely fine tuning in order to achieve its 1st rank match rating of 61% whereas this algorithm showed 69% 1st rank match accuracy over almost half of the total 87 trials of changing its parameters. Not directly evident in the data is the dual peak detection's increase of the 2nd, 3rd, 4th, and other higher match rankings of the software. In fact, our final optimized parameters had a continual increase in higher order match rankings, with 72%, 78%, and 81%, respectively.

Conclusion

In conclusion, the implementation of the dual peak detection algorithm into the THIS program has raised both the 1st rank match accuracy and stability with regards to parameter changes. In addition, the dual peak detection algorithm has also increased the percentages of lower rank match accuracies. Some common trends in the data, including noisy signals and a sizeable subset of the subjects that the classifier often took at least 24 out of 32 tries to match suggest that there could be other factors at work affecting the accuracy of the THIS program. For example, better data could be trained and tested on, whether that data was better filtered, recorded more accurately, or even averaged with another kind of signal. Another explanation might be that our current system of creating and analyzing spectrograms of heartbeats as read from a Laser Doppler Vibrometer is only applicable to a sizeable majority of humans and has outliers in a given population. Further studies could include looking into algorithms that dynamically adjust target parameters in order to achieve more accurate results, or even algorithms that might replace those target parameters entirely.

Bibliography

1. R.E.McGregor; S.Schuckers; J.Skufca, "Thoracic Biometrics - Investigations of the Human Heartbeat as a Biometric: Heart Sounds, Electrocardiogram, and Vibrometry," CIA Report (not published), 2013.
2. I. Odinaka; P. H. Lai; A. Kaplan; J. O'Sullivan; E. Sirevaag; S. Kristjansson; A. Sheffield; J. Rohrbaugh, "ECG Biometrics: A Robust Short-Time Frequency Analysis," IEEE Int. Workshop on Information Forensics and Security, pp. 1-6, 2010.
3. M. Chen, "Laser Doppler Vibrometry as Biometric", M.S. Thesis, Washington University, 2007.
4. P.H. Lai; J.A. O'Sullivan; M. Chen; E.J. Sirevaag; A.D. Kaplan; J.W. Rohrbaugh, "A Robust Feature Selection Method for Noncontact Biometrics Based on Laser Doppler Vibrometry," Biometrics Symposium, 2008, pp.65-70, 2008.
5. M. Chen, "Laser Doppler Vibrometry Measures of Physiological Function: Evaluation of Biometric Capabilities", Ph.D. Thesis, Washington University, 2009.
6. M. Chen; J.A. O'Sullivan; A.D. Kaplan; P.H. Lai; E.J. Sirevaag; J.W. Rohrbaugh, "Biometrics With Physical Exercise Using Laser Doppler Vibrometry Measurements of the Carotid Pulse," 2009 Int. Conf. on Biometrics, Identity and Security, pp. 1-6, 2009.
7. A.D. Kaplan; J.A. O'Sullivan; E.J. Sirevaag; J. W. Rohrbaugh, "Laser Doppler Vibrometry Measurements of the Carotid Pulse: Biometrics Using Hidden Markov Models", Proc. SPIE, (2009).
8. A.D. Kaplan; J.A. O'Sullivan; E.J. Sirevaag; S.D. Kristjansson; P.H.Lai; J.W. Rohrbaugh, "Hidden State Dynamics in Laser Doppler Vibrometery Measurements of the Carotid Pulse Under Resting Conditions," 2010 Annual Int. Conf of the IEEE Engineering in Medicine and Biology Society, pp. 5273-5276, 2010.
9. M. Chen; J.A. O'Sullivan; N. Singla; E.J. Sirevaag; S.D. Kristjansson; P.H. Lai; A.D. Kaplan; J.W. Rohrbaugh; , "Laser Doppler Vibrometry Measures of Physiological Function: Evaluation of Biometric Capabilities," IEEE Transactions on Information Forensics and Security, v..5, no.3, pp.449-460, 2010.
10. A.D.Kaplan, "Information Processing for Biological Signals: Application to Laser Doppler Vibrometry", Ph.D. Thesis, Washington University, 2011.
11. Y. Michalevsky; A. Schulman; G.A Veerapandian; D. Boneh; , "PowerSpy: Location Tracking using Mobile Device Power Analysis," 24[th] USENIX Security Symposium, 2015.
12. Muller, M., Information Retrieval for Music and Motion, Ch. 4. New York: Springer, 2007. http://www.springer.com/cda/content/document/cda_downloaddocument/9783540740476-1.pdf?SGWID=0-0-45-452103-p173751818