Thoracic Biometrics – Investigations of the Human Heartbeat as a Biometric: Heart Sounds, Electrocardiogram, and Vibrometry

Report Update

January 20, 2015

2011-2013 Intelligence Community Postdoctoral Fellow: R. Eric McGregor, Ph.D.
Principal Investigators: Stephanie Schuckers, Ph.D. and Joseph Skufca, Ph.D.
Project Manager: Kelly Faddis

Introduction

For the project *Thoracic Biometrics –Investigations of the Human Heartbeat as a Biometric: Heart Sounds, Electrocardiogram, and Vibrometry,* the authors investigated the use of traits from the human heart as a biometric. An experiment was conducted that collected three types of cardiac data: the carotid pulse, the electrical activity and the sound from 36 subjects. The carotid pulse (CP) was measured using a Laser Doppler Vibrometer, the electrical signal displayed as an electrocardiogram (ECG) was measured using electrodes placed on the subjects' thoracic region and the cardiac sound displayed as a phonocardiogram (PCG) was measured using a digital stethoscope.

For each subject, 12 recordings of each of the 3 signals were taken. The first 4 recordings (set 1) were taken when the subject was at rest, the second set of 4 (set 2) were taken after the subject had pedaled a recumbent bicycle until their heart rate was elevated, and the third set of 4 recordings (set 3) were taken at rest approximately 5 minutes after exercising. Within each set of recordings the experimental conditions remained unchanged for the electrical and acoustic signals, however blurring of the laser was used to distort the carotid pulse signals in 3 of the 4 recordings. For each set, the first recording had a finely focused laser and for the other 3 recording the laser focusing apparatus was rotated 90 degrees.

To analyze the data the authors developed biometric identification and authentication software named Thoracic Identification System (TIS) written in MatLab. TIS utilizes a probabilistic classifier and a support vector machine (SVM) for identification or verification. Using the probabilistic classifier the authors reported good performance on ECG data but rather poor performance on the CP data, contrary to previously published reports [2]. It should be noted that the authors of [2] tested and trained on rested heartbeats whereas [1] trained on rested heartbeats and tested on elevated heartbeats.

This report details the modifications to TIS that were made by McGregor and undergraduate student Tayseer O'Brien. These modifications include a new heart beat detection function and modifications to the algorithm which chooses which beats to use for training and testing the classifier. We also present test results that show a significant increase in performance for CP identification when using heartbeats with different heart rates. Specifically, when using the older version of TIS and training the classifier on recording 1 of the CP signal and testing on recording 9 of the CP signal for 33 of the 36 subjects the identification rate was only 27.27%. When using our modified software on the same subjects, again training the classifier on recording 1 of the CP signal and testing on recording 9 of the CP signal, the identification rate is 60.61%. When training on recordings 1 and 5 of the CP signal and testing with recording 9 of the CP signal we obtain a 69.7% identification rate.

Algorithm Modifications

The original TIS software was based on the algorithms in [2]. The algorithm first segments individual heartbeats from the recordings and then using a sliding short-time Fourier transform constructs a spectrogram for each heartbeat. The spectrogram is a matrix of floating-point numbers, each representing the power of the signal over a window of time and a range of frequencies. Each cell in the matrix is referred to as a bin. A predetermined number of bins are chosen as identifying features for each individual. The feature selection algorithm picks bins for an individual based on how well the values of those bins remain constant over time among all the individual's own heart beats and how well they differ from the rest of the population. Once the feature sets are calculated, classification is performed.

Our revised software does two things differently. First, we rewrote the beat detection algorithm and second, we modified the algorithm that chooses which heartbeats to use for training and testing. The remaining parts of the software were left unchanged.

In the original software we used an open source beat detection package written by Sergey Chernenko, downloadable from http://www.librow.com. The beat detection software worked well on ECG signals but did not detect carotid pulses consistently. This resulted in a set of heartbeats that were not aligned with one another. This is important because the feature selection algorithm computes nominal heart beats that individual heartbeats are later compared against. When a nominal heartbeat is computed using unaligned beats the result does not represent an *average* heartbeat. To address this problem, we rewrote the beat detection algorithm from scratch. The new algorithm segments the LDV signals into beats that are 100% aligned for segments of the signal that are void of EMI or movement artifacts.

The second modification that we made to the software was to add a new algorithm to determine which heartbeats to use for each individual. The original software had two modes: *first consecutive* and *random*. Using the first consecutive mode, the software would choose the first k consecutive beats and when using the random mode the software would pick k random beats. The new *smarter* algorithm picks k beats that are closest to the individual's nominal heartbeat. The new heartbeat selection algorithm increases performance because the chosen heartbeats have similar heart rates and the heartbeats that have noise due to movement or EMI are less likely to be selected.

This aids feature selection where we pick a set of bins for an individual. To select the bins, we construct spectrograms for each of the individual's selected heartbeats, construct a nominal heartbeat for each individual and construct a nominal spectrogram of the population using the spectrograms of all individuals. An individual's feature set is then chosen by looking for the bins that vary little among the individual's spectrograms, yet different greatly from the population's nominal spectrogram. By selecting a set of heartbeats for an individual, where each is *near* their own nominal heartbeat, we are able to choose features that better represent the individual.

Data Used for Testing

In the tests we performed for this work, we used data from 33 of the 36 subjects. Subject 526928's data folder was missing all data files, subject 8579635 was missing only trial 9 data (which is used in our tests) and subject 653146 had no CP data in their data files. In our previous work [1] we also omitted subject data that contained obvious EMI. In that work we only used data from 24 of the 36 subjects. Here we do not omit any data containing EMI.

Since we only have one session of data with 12 separate recordings for each of our subjects we train and test our classifier on different recording. For each subject, recordings 1, 5, and 9 have the laser acutely focused on the subject's carotid artery. Recording 1 is taken when the subject is at rest, recording 5 is taken after the subject pedals a recumbent bicycle, and recording 9 is taken approximately 5 minutes after exercise[1]. All of our tests utilize recordings 1 for training and recordings 9 for testing because both have a sharply focused laser and the subject's heart rate in each recording is less elevated than recording 5.

Test Results

In Table 1 we show the results of 7 tests on TIS using the data set described above. The identification rate is in the bottom row. The values for each of the software parameters are given for Test 1. Fields highlighted in blue indicate settings that were changed compared to the previous test. Fields containing ← indicate the test used the same setting as the previous test.

Test 1 shows the results (27.27% identification rate) of using the original software and settings from [1]. For Test 2 we use the original software and only increase the window size used in the beat selection algorithm from 400 samples to 770 samples. This produces a modest increase in performance (+3.03%). In Test 3 we drop the maximum frequency when selecting features. This produces a smaller set of bins to choose features from, later illustrating (in Test 6) that the best identifying features are below 50 Hz. Again Test 3 has a modest increase in performance (+6.06%). Test 4 uses only 3 heartbeats to test the classifier and 13 heartbeats to train the classifier. Though the performance is poor (6.06%), it is shown for comparison with Test 7 that uses the same parameters.

Test 5 uses our modified software, which does not choose heartbeats randomly, but rather uses our new algorithm. All other parameters are the same as Test 1. Here we see a measurable increase in performance compared to the original software with an identification rate of 42.42% (+15.15%). By limiting the size of the spectrograms to 50 Hz, in Test 6 we see a substantial increase in performance with an identification rate of 57.58% (+30.31%). Finally, in Test 7, by limiting the number of heartbeats used for training and testing we get an identification rate of 60.61% (+33.34).

---

[1] Though recording 9 is taken approximately 5 minutes after the subject has exercised, their heart rate is still elevated in most cases.

Table 1: Test Parameters and Results

Test #	1	2	3	4	5	6	7
Peak Detection Options							
Algorithm Implementation	Chernenko	←	←	←	O'Brien	←	←
Window Size	400	770	←	←	←	←	←
Modeling Options							
Algorithm Type	Homogenous	←	←	←	←	←	←
System Mode	Identification	←	←	←	←	←	←
Classifier	Probabilistic	←	←	←	←	←	←
LDV Preprocessing Options							
Downsample Rate	1000	←	←	←	←	←	←
Apply High Pass Filter	Yes	←	←	←	←	←	←
Remove Dropouts	No	←	←	←	←	←	←
Use Displacement Signal	No	←	←	←	←	←	←
Segmentation Options							
Segments Per Beat	1	←	←	←	←	←	←
Segment Boundaries	Fixed Length Relative to Peaks	←	←	←	←	←	←
Before Primary Peak Length (samples)	200	←	←	←	←	←	←
After Primary Peak Length (samples)	500	←	←	←	←	←	←
Train & Test Data Options							
Segment Selection Method:	Automatic	←	←	←	←	←	←
Randomly Select Training and Testing Segments	Yes	←	←	←	No	←	←
Number of Segments Per Training Subject	12	←	←	3	12	←	3
Number of Segments Per Test Subject	7	←	←	13	7	←	13
Number of Segments Per Test	1	←	←	←	←	←	←
Spectrogram Options							
Window Length (samples)	64	←	←	←	←	←	←
Window Step (samples)	10	←	←	←	←	←	←
Window Type	Hamming	←	←	←	←	←	←
Feature Selection Options							
Max Frequency Bin (Hz)	250	←	50	←	250	50	←
Number of Selected Bins	27	←	←	←	←	←	←
Identification Rate (%)	**27.27**	**30.30**	**33.33**	**6.06**	**42.42**	**57.58**	**60.61**

Additional tests were run using multiple recordings of the CP signal to train the classifier. For example, when training on recordings 1 and 5 and testing on 9 we get an identification rate of 69.70%.

Future Work

As mentioned above, we are continuing to investigate ways to mitigate the differences in heart rate. The problem here is that when a subject's heart rate increases, our algorithms only align the primary CP peak, not the rest of the signal. We are looking for ways to *normalize* each heartbeat regardless of heart rate. Another area we are currently working on is using Gaussian and other filters to remove the noise (or blur) from the CP signals that were recorded with an unfocused laser. This would prove useful in the field where a subject is moving and the laser becomes out of focus. We are also looking for additional data sets and the means to collect them ourselves.

Conclusion

In this work we showed a significant increase in the performance of TIS in identifying subjects based on their carotid pulse. In [1] the authors removed from testing subjects who had signals containing EMI. In this work we only omit subjects that have missing CP signals. The original software from [1] produced an identification rate of 27.27% using the parameters cited in [1] while training on rested heartbeats from recording 1 and testing on elevated heartbeats in recording 9. By modifying only the window size used in the beat detection algorithm and the maximum frequency used in feature selection we were able to increase its performance to 33.33%. By rewriting the beat detection code and changing the maximum frequency used in feature selection along with limiting the number of beats used for training and testing we were able to increase the performance to 60.61%. When training the classifier with heartbeats having both rested and elevated heart rates and testing the classifier on elevated heart rates we obtain an identification rate of 69.70%.

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. 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.