<3 ID
Non-Intrusive Cardiac Rhythm-Based User Identification using an Accelerometer

Yoolhee Kim   Vasudevan Nambeesan   Anand Krishnan Prakash
{yoolheek, vnambees, anandkrp}@andrew.cmu.edu
Carnegie Mellon University

ABSTRACT

There has been significant work in recent times on the use of human cardiac rhythm as a tool for identification. While most work in this area explores the use of electrocardiogram signals for analysis, this paper investigates a novel concept of the use of accelerometer signals from heart-induced vibrations in an automobile seat, for the specific application of user identification. Experiments are performed on eight individuals with accelerometers placed on the neck for proof of concept, and underneath a chair where the user sits to replicate a real world scenario. Filtering was performed on the data to eliminate noise, and features that could be used to identify the user were extracted from the signals. Decision Tree, Random Forest and Gaussian Naive Bayes classifiers were used for classification of the dataset. We have found that the Random Forest model worked best for this work: the random forest classifier model was tested using randomized 20-fold cross validation and yielded an accuracy of 86%.

1. INTRODUCTION

This project is motivated by the growing significance of security and user identification. User identification is required for authorization (to verify if a user has access to perform a particular action), and for authentication (to prove if the user is who they claim to be). Occupant identity information in a smart home setting or an work environment is useful for enabling personalized environment settings like specific temperature and lighting settings, which also improves energy efficiency. In this project, we look at the specific application of driver identification in automobiles. We will try to identify the person occupying the driver's seat for authorization of starting the vehicle. This system could stop the car in the case of detecting a user that is not in the list of authorized users, and alert the owner of the vehicle.

While there exist more popular forms of biometric identification, we believe this non-intrusive method of user identification using cardiac rhythm has advantages over these existing techniques. Retina scanning is highly invasive and many users feel uncomfortable while scanning. Fingerprint scanning is another form of identification, but users' fingerprint can be subject to some form of wear and tear and hence this is not always permanent. The closest form of identification using cardiac rhythm is user identification using electrocardiogram (ECG), but it is not practical to use record readings using ECG each time you need to identify a user. While heartbeat rate is also subject to changes due to aging [14], the cardiac rhythm between pulses still remains the same - it is only the frequency of the pulses that change.

In this project we propose a non-intrusive method for user identification using accelerometers, based on heart-induced vibration in chairs. The accelerometer picks up vibration data from the chair caused by the cardiac rhythm of the user seated in the chair, and that data is used to identify the individual.

In addition to accelerometer data from the chair, we have taken data from an accelerometer attached to the neck. Because the heartbeat signal is much stronger and suffers much less from noise, from the neck, we hope that this will offer a proof-of-concept for the possibility of user identification through heartbeat signals.

The shape of the heartbeat is quantified by the following parameters (as can be seen in Figure 1): P is the first peak, Q is the next trough, R is the next largest peak, S is the next trough, and T is the final peak in a heartbeat pattern.

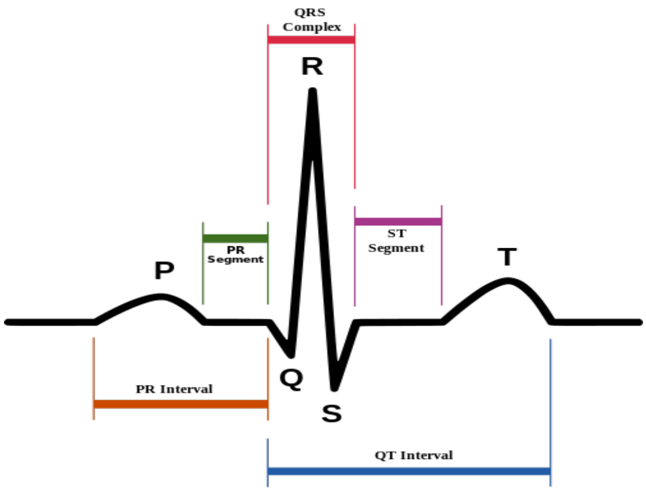

Figure 1: PQRST form of the Heartbeat [4]

This method works based on the fact that unique characteristics in the heart beat pattern correspond to unique individuals. Various features, such as the distances between the peaks of the heartbeat pulse, are extracted from the accelerometer and used to identify the user. These features have been used in the past in order to identify users based on heartbeats from ECG data.

2. RELATED WORK

Recent studies in the field of biometrics have shown that human electrocardiogram (ECG) signals can serve as a biometric identification tool [5]. As the muscles in the heart expand and contract, electrochemical signals are generated near the muscles which cause spikes in the ECG signal. Use of this signal for human identification have been investigated and proven with accuracy of over 90% by various authors. The unique features of this signal are the relative distances between the P, Q, R, S, T points in the heartbeat wave, which was used by Israel et al in 2005 [9]. Wubbeler et al in 2007 verified the long term stability of individual ECG signals and achieved an equal error rate smaller than 3% [16].

A combination of processing techniques was used by Irvine et al in 2008 for human identification with ECG signals [8]. They achieved a near 100% enrollment rate. Fang et al in 2009 developed an unsupervised identification method by measuring similarity in ECG phase space portraits [7]. The use of ECG signal from the fingers was investigated by Lourenco et al in 2011 [3].

While the research discussed above proved the concept of biometric identification as well as developed techniques to improve accuracy, the signal source was an ECG electrode from either the chest or fingers and therefore are a partially intrusive method of detection. In our project we attempt to use an accelerometer for sensing the heartbeat from underneath the chair seat, with no direct contact to the user. This is a novel method compared to previous work because 1) it uses an accelerometer for identification and 2) it is an extremely non-intrusive method, as the accelerometer is not in direct contact with the user.

Additionally, there has been some work on the use of accelerometer for heart rate measurement. Phan et al in 2008 used an accelerometer mounted on the chest to measure heart rate and respiratory function [10]. This preliminary study verifies that cardiac patterns can be detected and observed with an accelerometer. The accelerometer picks up changes in acceleration caused by the flow of blood vessels against the skin [13].

A recent paper from MIT shows the use of an accelerometer from an iPhone for detection of heartbeats. This preliminary study verifies that cardiac patterns can be detected and observed with an accelerometer. The accelerometer picks up changes in acceleration caused by the flow of blood vessels against the skin [13]

3. CHALLENGES

We have faced some research challenges as this project proves a concept that has never been attempted in the past. Since the use of accelerometer for identification has not been done before, identifying the right features is a challenge since there are no existing proven methods. In addition, the physical phenomenon that connects the heart's pumping cannot be easily associated with peaks that occur in the accelerometer signal (as in the case of an ECG). Therefore choosing the appropriate features to identify the user, and filtering parameters to isolate the peaks, needs to be chosen through a brute force method.

Once we have extracted the appropriate features from the accelerometer signal, there is a challenge in fully understanding the nature of these extracted features. These features may be correlated, independent, or have some type of probabilistic function, but these characteristics are not well understood. This makes it difficult to decide the appropriate classifier to use for identification of the individual based on these features. As a result, we have tested a few simple classifiers and will compare the results between these classifiers.

4. SYSTEM ARCHITECTURE

There are three major sources of data for this project. The first is from an accelerometer placed on the neck. The second is taken from an accelerometer mounted on a chair beneath the surface where the user sits. The third is from a Fitbit band placed on the user's wrist, which serves as the ground truth for the experiment. The heartrate BPM information from the Fitbit is matched with the accelerometer data to ensure that the peaks from the accelerometer match up with the expected heartrate. Ten 30-second runs of data were taken from ten individuals. The data from the accelerometer was sent through a National Instruments Data Acquisition system which is then passed to the computer. Figure 2 shows the data acquisition tools we used in this work. The position of the two accelerometers from which data was collected can be seen in Figure 3, which shows data collection session in progress.

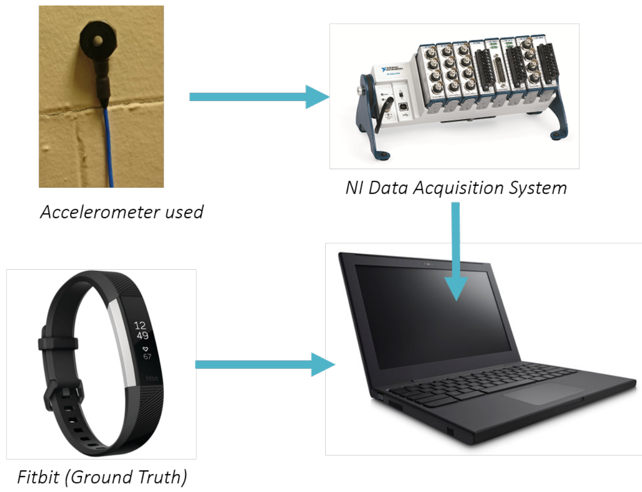

Figure 2: Schematic of the Experimental Setup

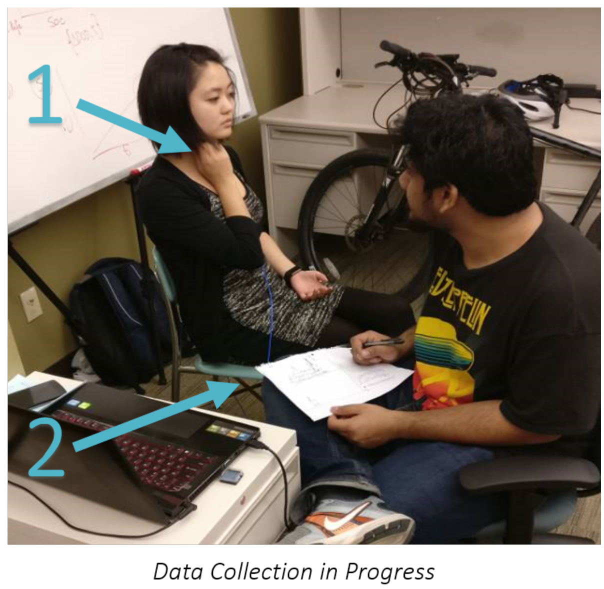

Figure 3: Data Collection

5. SIGNAL PROCESSING AND FEATURE EXTRACTION

5.1 Filtering the Data

The heartbeat pattern is a continuous analog signal that has been discretized by the sensing system at a sampling frequency of 10.24 kHz. The heartbeat from the user is picked up by the accelerometer and processed. This raw data is sent to a Butterworth band pass filter with cut off frequencies of 10 and 90 Hz and order of 3 to remove ambient noise. After applying the Butterworth band pass filter with these parameters on the raw signals shown in Figures 4(a) and 4(c), we obtained filtered data samples presented in Figures 4(b) and 4(d) respectively. Two sharp peaks and two sharp troughs can be observed for each pulse of the individual from the neck data and one peak and one trough for the chair data.

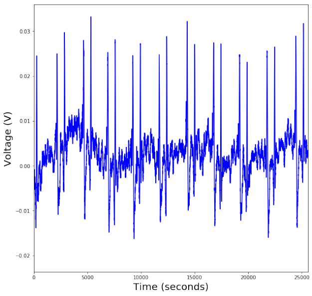

(a) Unfiltered Neck

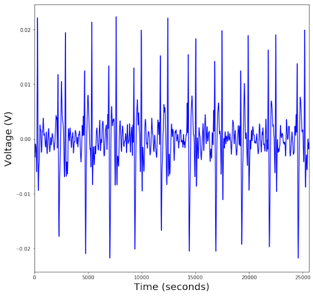

(b) Filtered Neck

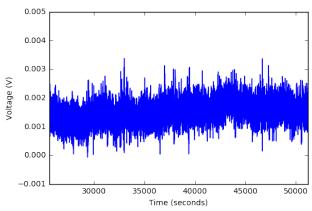

(c) Unfiltered Chair

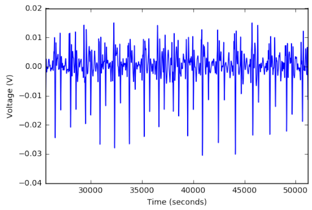

(d) Filtered Chair

Figure 4: Filtering Data

5.2 Extracting Features

Since the properties of these peaks and troughs are the features for the project, a peak detection algorithm is used to identify both the amplitude and the time of the peaks and trough (peaks from chair data can be seen in Figure 5(b)). The peaks are distinguished as first and second peaks of a pulse for the neck data only (Figure 5(a)). The features used for this project are the various distances between the peaks and troughs of each pulse. The absolute values of amplitudes of the peaks and troughs were also used as features. Finally, a Fast Fourier Transform was performed on parts of the signal and the frequency with the maximum amplitude in the FFT is used as a feature. For each peak detected in the earlier steps, we perform FFT on a set of 1000 points before the peak, the peak and the set of 1000 points after the peak (around 0.2 seconds of data per peak). The frequency at which we obtain a maximum amplitude is used as a feature for that peak. Figure 5(c) shows an example of the FFT output performed on a set of points around a neck-peak.

The combination of the features described above contains information that is unique to an individual. Since the objective of the project is to use the features listed above to identify individuals, the classification labels for this project are the identity of the individuals observed.

6. MACHINE LEARNING METHODS

For our classification of individuals based on accelerometer data, we will be using a few machine learning methods: Decision Tree Classifier, Random Forest, and Gaussian Naive Bayes. Because we do not fully understand the nature of the features we extract, we have used a few simple classification methods to try to determine which classifier performs best. Through further analysis of the features as well as the results of the classifiers, we hope to gain a better understanding of the features we have chosen to use for user identification, and the appropriate type of classifiers to use for these datasets.

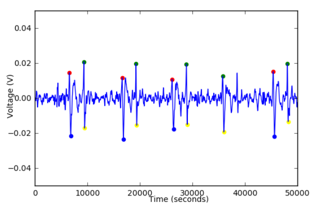

(a)

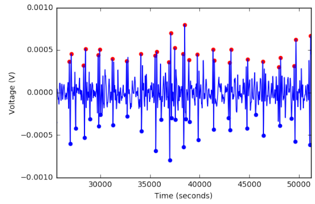

(b)

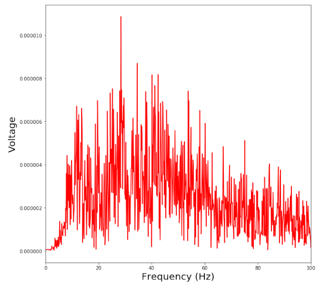

(c)

Figure 5: Peaks and troughs from 5(a) Neck and 5(b) Chair and 5(c) Output of FFT done on one neck peak

6.1 Decision Tree Classifier

The decision tree classifier is a machine learning technique that creates a set of sequential optimal splits in the feature space such that the variance within each leaf is minimized. The splits are made based on decisions pertaining to the features [15]. Thus the final decision tree contains leaves which each have a set of unique feature criteria. One or more leaves can be associated to a label. Decision trees are one of the most visually simple classifiers.

6.1.1 *Justification for Use*

The decision tree classifier was chosen in this project because it makes no assumption about the nature of the features, and simply splits the features to create leaves that minimizes the variance of the labels within each leaf. This seemed ideal considering we do not have a great understanding of the nature of the features we have extracted from the accelerometer heartbeat data. Additionally, the decision tree classifier gives an idea of which features are the strongest to obtain the desired labels [1]. Since the idea of using accelerometer data to identify individuals is novel and has never been done before, this information will be of significant importance.

6.1.2 *Parameter Optimization*

Various parameters of the classifier have been optimized for each dataset to give us the best accuracy of user identification. These parameters include the depth of the tree, number of features, criterion, and splitter.

The depth of the tree (i.e. the number of levels of leaves) can be controlled. Increasing the depth can refine the nuances in the feature space but a very high depth can cause an over-fit. The number of features considered when choosing for the best feature split can also be determined. Unless it is specified in the function, it automatically considers all the features for each split.

The criterion parameter determines the function to measure the quality of a feature split, including gini impurity and entropy. Gini impurity is measured by the probability of incorrect label selection if labels are randomly distributed in a leaf, while entropy defines the criterion by amount of information gain given a split. The splitter parameter determines the strategy used to choose the feature splits at each node [1].

6.2 Random Forest Classifier

Random decision forests are an ensemble learning method for classification, and operates by constructing many decision trees at training time. The classifier fits a number of decision trees on various sub-samples of the dataset, and uses averaging to improve accuracy by controlling over-fitting of the dataset /citerandom-forest.

6.2.1 Justification for Use

Like the decision tree classifier, random forests do not make any assumption about the nature of the features, and simply splits the features to create leaves that minimizes the variance of the labels within each leaf. Random forests run efficiently on large data bases, and can handle thousands of features without any deletion of the features. The random forest also has an effective method for guessing missing data, and is able to maintain accuracy when a large portion of the data is missing [6].

Additionally, similar to the decision tree classifier, this classifier gives an estimates of which features are most important in the classification. Since we do not know which features are most important and useful in identifying individuals, this can be a very useful tool for boosting our classification accuracy [2].

6.2.2 Parameter Optimization

Various parameters of the classifier have been optimized for each dataset to give us the best accuracy of user identification. These parameters include the depth of the tree, number of features, number of estimators, and criterion.

The depth of the tree (i.e. the number of levels of leaves) can be controlled. Increasing the depth can refine the nuances in the feature space but a very high depth can cause an over-fit. The number of features considered when choosing for the best feature split can also be determined. Unless it is specified in the function, it automatically considers all the features for each split [2].

The number of estimators decides how many trees will be considered in the forest. The default value is 10 trees if a specific value is not set. The criterion parameter determines the function to measure the quality of a feature split, including gini impurity and entropy. Gini impurity is measured by the probability of incorrect label selection if labels are randomly distributed in a leaf, while entropy defines the criterion by amount of information gain given a split [2].

6.3 Gaussian Naive Bayes Classifier

Gaussian Naive Bayes classifiers are simple probabilistic classifiers based on applying Bayes theorem. The classifier assumes strong independence amongst the features; in other words, it assumes that the presence of a particular feature in a class is unrelated to the presence of any other feature [12]. For this report, we will assume independence and a Gaussian distribution of our features, meaning that our features are unrelated and follow a normal distribution. However, this is not definite and we will have to interpret our features further to determine that this is the case in order to justify use of this classifier. The Gaussian Naive Bayes Classifier in the python scikit-learn library does not have any parameters for optimization.

6.3.1 Justification for Use

Naive Bayes model is easy to build and useful for large data sets. This classifier is simple and known to outperform highly sophisticated classification methods. The Naive Bayes classifier performs best when the assumption of independence of the features holds true. It is useful because less training data is required. The use of the Naive Bayes may be an issue with this application due to the assumption of independent predictors [12]. This is almost never true in real life, and in our case, many of our features are likely to be correlated, such as the distances between peaks of the heartbeat pulse and the maximum FFT frequency.

7. RESULTS

We analyze the performance of our classifier models through k-fold cross validation, confusion matrices, and visualization of the decision boundaries. These results reflect the classification of 8 different people, for whom we have taken heartbeat accelerometer data over a 5-minute period. Some sections of this data had to be removed due to noise issues.

We test the classifiers on 6 datasets - Chair_All, Chair_Mean, Chair_Median, Neck_All, Neck_Mean, and Neck_Median. The "Mean" and "Median" data are the datasets with features averaged over a run, meaning each person will have 10 samples per feature since each person had 10 experiment runs. In contrast, the "All" data corresponds to the datasets where each heartbeat pulse is a sample, corresponding to around 500 samples per person, which is the number of heartbeat pulses over a 5 minute period.

7.1 K-Fold Cross Validation

Results from a 20-fold cross validation on the classifier data using Random Forest, Decision Tree and Gaussian Naive Bayes classifiers is tabulated in Table 1. Different sets of data, as mentioned above, has also been used. It can be seen that Chair Mean data performs exceptionally well irrespective of the classifier used. Classification on the Chair data has produced more accurate classifications than the Neck data. Random Forest has the overall best accuracy of the three classifiers used. It is important to note that running Random Forest Classifier took more time than the other two classifiers due to the computational complexity.

7.2 Confusion Matrix

Figures 6, 7 and 8 show the confusion matrix for Decision Tree, Gaussian Naive Bayes, and Random Forest Classifiers respectively with the ChairMean data. Each cell in the matrix can be is the intersection of row i and column j and it represents the number of user samples which are actually user i and has been classified as user j by our algorithm. Sum of all cells in a row i gives the total number of samples for that user i. It can be seen that the Decision Tree and Random Forest classifiers performs better than Gaussian Naive Bayes classifier as they do not have any incorrect classifications. The poor accuracy given by Gaussian Naive Bayes could be due to the collinearity in the features, since Naive Bayes assumes feature independence as well as a Gaussian distribution. It should also be noted that the confusion matrices for Decision Tree and the Random Forest classifiers are perfect, not even a single wrong classification. This may suggests over-fitting of the data. K-Fold cross validation (Section 7.1) reveals that the classifier results are not perfect as shown in Table 1.

7.3 Decision Boundaries

Table 1: Results of K-Fold Cross Validation with k = 20

Classifier	Chair_All	Chair_Mean	Chair_Median	Neck_All	Neck_Mean	Neck_Median	Classifier Average
Random Forest	0.70	0.86	0.79	0.85	0.73	0.73	0.78
Decision Tree	0.66	0.86	0.67	0.75	0.69	0.67	0.72
Gaussian Naive Bayes	0.45	0.88	0.65	0.49	0.68	0.55	0.62
Average	0.60	0.87	0.70	0.70	0.70	0.65	

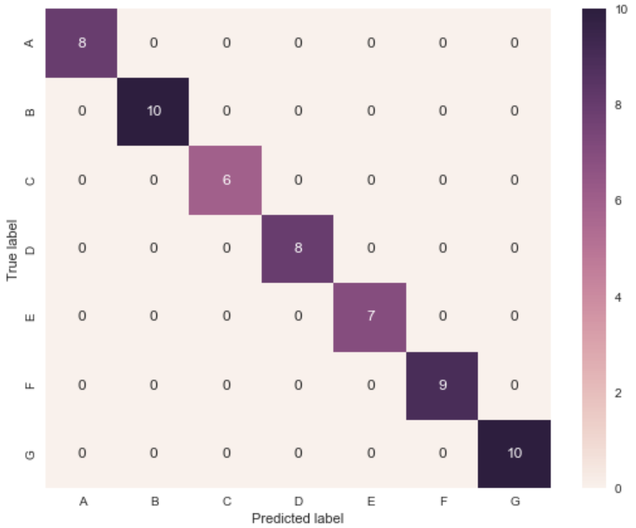

Figure 6: Confusion Matrix for Decision Tree for Chair Mean Data

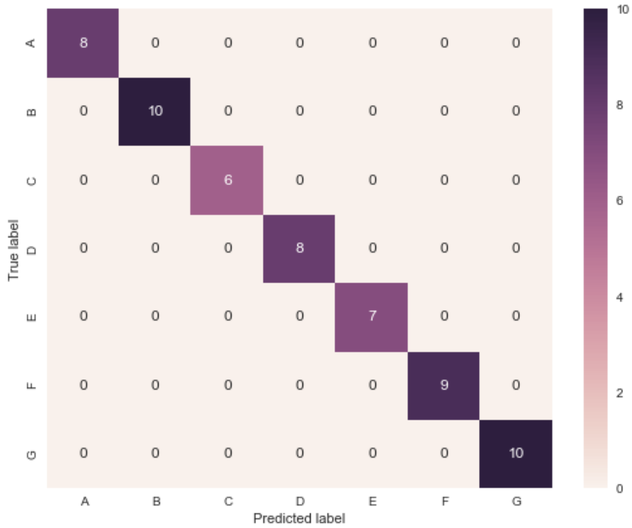

Figure 8: Confusion Matrix for Random Forest for Chair Mean Data

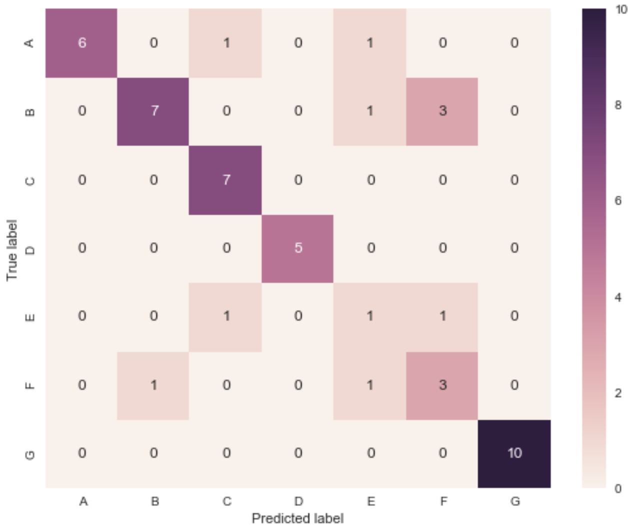

Figure 7: Confusion Matrix for Gaussian Naive Bayes for Chair Mean Data

A classifier divides the feature space into multiple decision regions and all points in one region will have the same output label. Decision Boundaries are formed at the intersection of these decision regions [11] and classification of any point lying on the decision boundary is always ambiguous. Figures 9, 10 and 11 show the decision regions and the decision boundaries of the classification when we use Decision Tree, Gaussian Naive Bayes and Random Forest Classifiers on the Chair Mean data respectively (for the two more important features for each classifier). Both the Decision tree and the Random Forest classifiers make sequential optimal splits on the data, one feature at a time and hence it divides the feature space into multiple rectangular regions. Additionally, these plots show that the features show some relationship to each other, as the data points are not randomly or evenly scattered.

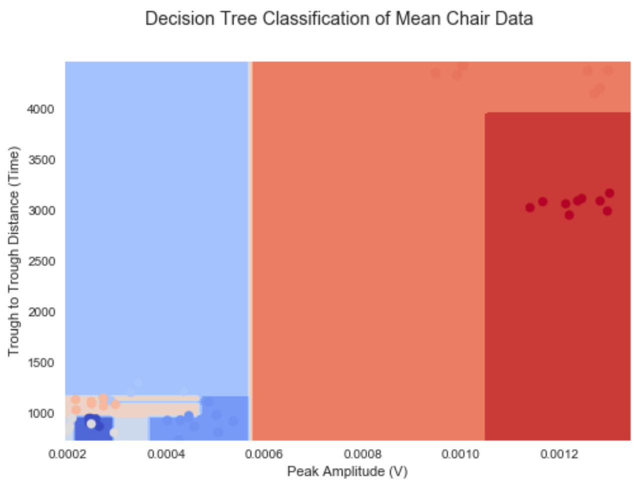

Figure 9: Decision Boundary using Decision Tree for Chair Mean Data

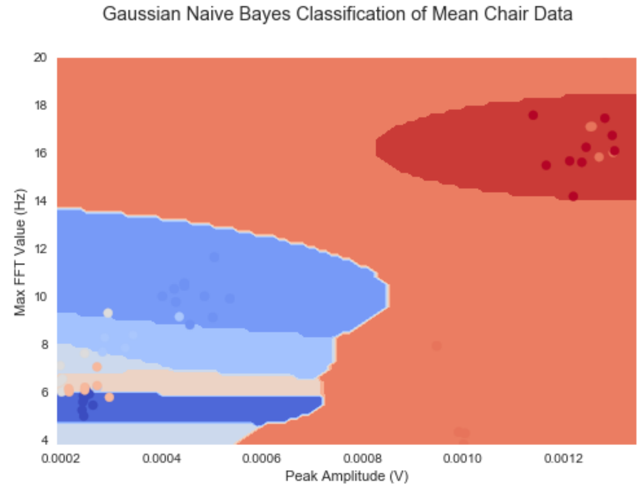

Figure 10: Decision Boundary using Gaussian Naive Bayes for Chair Mean Data

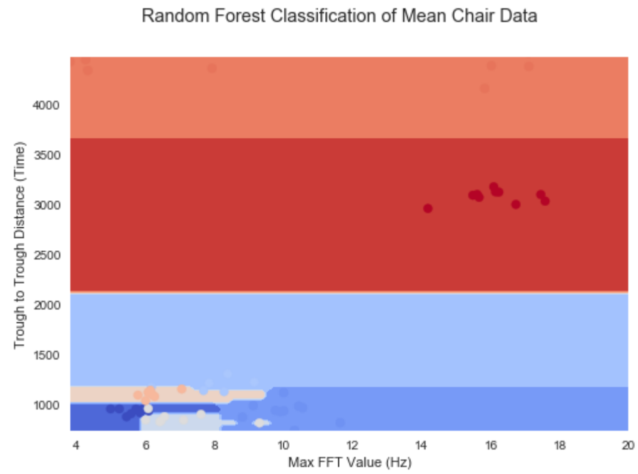

Figure 11: Decision Boundary using Random Forest for Chair Mean Data

7.4 Analysis

Random Forest clearly performs the best of all the classifiers. This is not surprising, since random forest is more sophisticated than decision trees, and unlike the Gaussian Naive Bayes classifier, does not assume anything about the nature of the features. Random forest usually does best since the use of multitudes of trees and averaging helps to prevent over-fitting of the data. Additionally, we did not do significant research on the nature of our features, which may be detrimental for some classifiers such as the Naive Bayes Classifier. However, random forest classifiers typically perform better with more features fed to the classifier, regardless of what characteristics these features may have.

Gaussian Naive Bayes performs decently, but it assumes independence amongst the features, and also assumes a normal distribution. Based on the visualization of the decision boundaries, it seems that some of our features may be correlated to each other, weakening the Gaussian Naive Bayes classification.

Surprisingly, the chair data performs better than the neck data, and of the chair data, the mean data performs the best. We believe that the chair data performs better than the neck data due to the fact that the neck data was not collected in a standardized way; each person would hold the accelerometer to their neck. This introduces the potential for much experimental error, as the user can hold the accelerometer in a different manner throughout the experimental runs. This inconsistency can cause errors in the signal, and cause complications in the classification of the individual from the extracted features.

In addition, because the Chair_Mean data performs better than the Chair_All data, this would imply that although there may be some outliers in the feature samples, 30-second average of the feature samples gives a good consistent features for identification of the individuals.

8. CONCLUSION

This paper demonstrates that the use human cardiac rhythm measured using an accelerometer-based sensing system can be used as a tool for unique user identification. For the eight individuals tested in this paper, it was observed that the use of appropriate filtering and feature extraction methods from chair signals generated successful results. The randomized 20 fold validation was performed on the classifier models to evaluate their performance. Of the different classifiers used, the random forest method yielded the best results and maximum accuracy of 86% was obtained when Chair Mean data was used.

9. FUTURE WORK

Because we were able to prove good initial accuracy in identification of users through heart-induced chair vibrations picked up by the accelerometer, we would like to continue making improvements in our model to continue to boost our accuracy.

First, we would like to take cleaner data. From the chair, we would like to take data in a less noisy environment such that we have less disruptions from external vibrations other than the heartbeat. Additionally, we would like to take data from the neck in a more standardized manner, with either a flexible band or collar around the neck attached to the accelerometer. This would allow for consistent data acquisition of the heartbeat pulse from the neck, and hopefully a higher classification accuracy from the neck data.

Additionally, we would like to take a deeper dive into interpretation of the features we have extracted. We would like to analyze our features and look at the nature of the features and whether they occur in clusters, have a linear relationship, etc. We would like to determine the collinearity of our features by looking at partial dependence plots. Based on a deeper understanding of the characteristics and nature of these features, we can determine the most appropriate classifier to use for our dataset, and to improve our classification performance.

Finally, we would like to data from people under various conditions. In order to test the robustness of the model and to ensure the appropriate use of features, the data shall be collected at different heart rates of the same individual. This can be achieved by taking data under different conditions such as rest, running etc. A potential way to overcome this would be to use dynamic time warping to standardize the heartbeat rate for every individual.

10. REFERENCES

[1] Decision tree classifier: python-sklearn. `http://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeClassifier.html`.

[2] Random forest classifier: python-sklearn. http://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html.

[3] D. P. S. André Lourenço, Hugo Silva and A. L. N. Fred. Towards a finger based ecg biometric system. *International Conference on Bio-inspired Systems and Signal Processing - BIOSIGNALS 2011*, pages 348–353, 2011.

[4] A. Atkielski. Schematic diagram of normal sinus rhythm for a human heart as seen on ecg (with english labels). https://upload.wikimedia.org/wikipedia/commons/9/9e/SinusRhythmLabels.svg, January 2007.

[5] L. Biel, O. Pettersson, L. Philipson, and P. Wide. Ecg analysis: a new approach in human identification. In *IMTC/99. Proceedings of the 16th IEEE Instrumentation and Measurement Technology Conference (Cat. No.99CH36309)*, volume 1, pages 557–561 vol.1, 1999.

[6] L. Breiman and A. Cutler. Random forests. https://www.stat.berkeley.edu/~breiman/RandomForests/cc_home.htm.

[7] S.-C. Fang and H.-L. Chan. Human identification by quantifying similarity and dissimilarity in electrocardiogram phase space. *Pattern Recognition*, 42(9):1824 – 1831, 2009.

[8] J. M. Irvine, S. A. Israel, W. T. Scruggs, and W. J. Worek. eigenpulse: Robust human identification from cardiovascular function. *Pattern Recognition*, 41(11):3427 – 3435, 2008.

[9] S. A. Israel, J. M. Irvine, A. Cheng, M. D. Wiederhold, and B. K. Wiederhold. {ECG} to identify individuals. *Pattern Recognition*, 38(1):133 – 142, 2005.

[10] D. H. Phan, S. Bonnet, R. Guillemaud, E. Castelli, and N. Y. P. Thi. Estimation of respiratory waveform and heart rate using an accelerometer. In *2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 4916–4919, Aug 2008.

[11] Princeton. Decision boundaries. https://www.cs.princeton.edu/courses/archive/fall08/cos436/Duda/PR_simp/bndrys.htm, 2008.

[12] S. Ray. 6 easy steps to learn naive bayes algorithm (with code in python). https://www.analyticsvidhya.com/blog/2015/09/naive-bayes-explained/, September 2015.

[13] G. C. K. P. S Kwon, J Lee. Validation of heart rate extraction through an iphone accelerometer. https://www.ncbi.nlm.nih.gov/pubmed/22255524/, September 2011.

[14] J. B. Schwartz, W. J. Gibb, and T. Tran. Aging effects on heart rate variation. *Journal of Gerontology*, 46(3):M99, 1991.

[15] P.-N. Tan, M. Steinbach, and V. Kumar. Data mining cluster analysis: basic concepts and algorithms. *Introduction to data mining*, 2013.

[16] G. Wubbeler, M. Stavridis, D. Kreiseler, R.-D. Bousseljot, and C. Elster. Verification of humans using the electrocardiogram. *Pattern Recognition Letters*, 28(10):1172 – 1175, 2007.