ECG-Based Human Identification

Whitepaper by OMsignal
(Dated: December 13, 2017)

***Abstract*—In this work, a fast, reliable and accurate ECG-based human identification strategy is presented. A neural network is trained to identify individuals from a pool of 33 participants, given a window of 5 heart beats. The participants are drawn from the OMsignal MyHeart project. The windows are extracted from ECG recordings captured by OMsignal apparel while the participants go about their daily activities. The neural network is augmented with an uncertainty measure based on Monte Carlo dropout, allowing predictions to be made only on high quality data. On a testing dataset created from different recordings of the same participants and unobserved until development had finished, a window accuracy of 99.7% is achieved with a window rejection rate of 18.4%. Using majority voting classification across all collected windows, 31 out of 33 of the participants are correctly identified. The individuals who are not correctly identified belong to a subset of users who had less than 5 recordings on 5 separate days. This is hypothesized to be the lower limit on the amount of data necessary to observe enough variation to make accurate identifications of an individual.**

Identifying an individual from their ECG requires the presence and measurement of subtle features in the signal. The performance of this system provides encouraging evidence that OMsignal apparel, combined with machine learning algorithms may be able to find subtle patterns associated with developing medical conditions.

I. Introduction

The ECG signal has an individual-specific morphology due to the differences in position, size, and anatomy of the heart, age, sex, relative body weight, chest configuration, and various other factors.

Generally, the currently available ECG-based identification systems can be categorized as either fiducial points based methods (i.e. the PQRST signature along with their onsets and offsets) dependent or fiducial points independent algorithms. Usually the fiducial points detection is an error prone strategy that can be affected by changes of the signal slopes, inverted or abnormal waves, noise and artifacts. In scientific literature, the fiducial points independent approaches can be divided into three groups: auto-correlation, phase space, and frequency based analyses. Moreover, different pattern classification approaches, including minimum distance classifiers, support vector machines and neural networks are used for the final identification. The proposed ECG-based human identification strategy is depicted in Fig. 1.

II. Data

The data for this work comes from a subset of the data from the OMsignal MyHeart project. This study involved collecting data from several dozen women as they went about their daily lives with both the OMbra and the OMsignal comfort bra. No restrictions were placed on their activity or how they were to record the data. The raw ECG signal from these individuals were remotely pushed to the OMsignal servers.

At the time the data was frozen for this work, 2,328 hours of data had been recorded, encompassing 9,877,496 registered heart beats.

III. Signature Window Extraction

This section describes how the signature windows are extracted from a continuous ECG recording. This involves noise reduction, segmentation and normalization.

A. Noise Reduction

Baseline wander is removed from the signal using a smoothed median filter with an order corresponding to one second. High frequency noise is also attenuated using a wavelet filter.

B. Window Normalization

Before beat segmentation, each detected heart beat is normalized by:

- Vertical scaling: The amplitudes of signature window samples are scaled to be limited between zero and one.
- Horizontal scaling: Each beat is stretched or compressed to take up exactly 100 samples.

C. Beat Segmentation

Beats are segmented out of admissible blocks. A block is a ten minute segment of ECG. An admissible block must not belong to the first 10 minutes of a new recording, since this is approximately how long it takes the electrode-skin impedance to reach an acceptable level. Any admissible block is then processed through beat segmentation.

The R peaks of an admissible block are detected and all non overlapping data windows of 10 consecutive heart beats are then tested for window admissibility. An admissible window is any selection of five beats out of ten having an average heart rate below 90 BPM and where the Pearson product correlation coefficient of each one of the five beats with the mean of all ten beats is greater than 0.98. The signature window is constructed by choosing the five beats with the

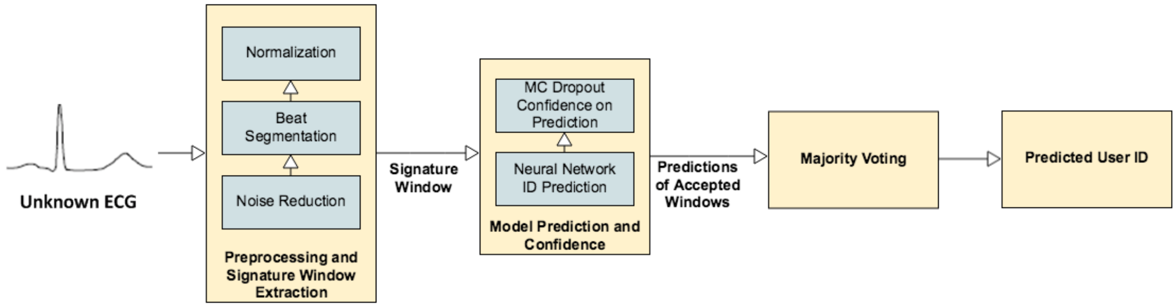

Fig. 1: The overall structure of the ECG-Based Human Identification System

highest correlation coefficient with the mean, provided these five beats form an admissible window.

An example of a signature window after all prepossessing steps is given in Fig. 2.

IV. Model Structure and Training

A neural network comprised of nine hidden layers with a softmax output layer is trained to identify an individual among 33 participants. Each input beat is normalized to a mean of zero and a standard deviation of 1. Batch normalization is used at the input and every hidden layer. Dropout regularization is used at every hidden layer and the output, and Gaussian noise is added to the input. This is in order to both prevent over-fitting, and to enable monte-carlo dropout uncertainty to be used at test time. Each beat is analyzed separately by a network that shares its weights across all the beats, to produce a feature vector and contribution strength for each feature. These are combined into a weighted average before the output layer. The intuition is to allow the network to learn to extract useful features for identification from each beat as well as to assess the quality of features per beat. Training is performed using a batch size of 256 signature windows, with a cross-entropy loss function optimized with an adaptive learning rate method: Adadelta. The network structure is depicted in Fig. 4.

V. Monte Carlo Dropout Confidence Measure

To ensure robust performance of the identification system, a confidence measure is necessary to discard those signature windows for which the decision made by the model is not sufficiently reliable. In recent years, Monte Carlo dropout has become a popular technique for quantifying epistemic uncertainty in neural networks. The following steps show how a confidence measure and the window rejection criterion is computed using Monte Carlo dropout samples.

- Run network *without* dropout to compute the most likely class given the signature window,
- Sample 500 times from network *with* dropout, and collect output samples only for the most likely class,
- Define the confidence value as the 5th percentile of the collected samples,
- If the confidence value is above 0.9 the classification is accepted, otherwise it is rejected.

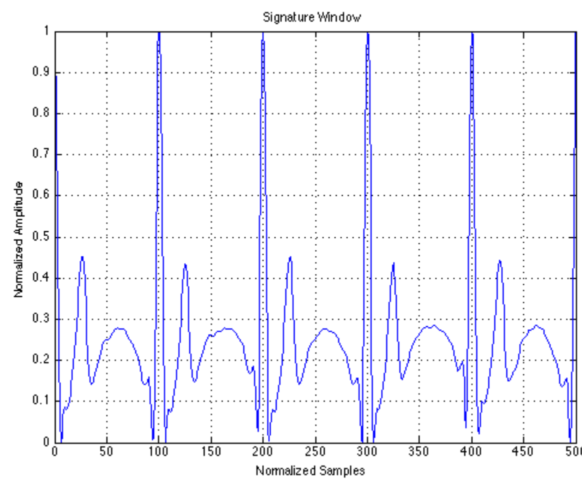

Fig. 2: Example of Signature Window after all prepossessing steps that is given to neural network as input. The QRS complex, T and P waves are all clearly identifiable.

Fig. 3 shows the distribution of confidence values across the signature windows of the validation set. The chosen acceptance threshold of 0.9 is justified on the basis that it is right before the density of confidence values begins to increase rapidly. Choosing a threshold between 0.1 and 0.9 would have little effect compared with choosing a threshold at one of the boundaries. The more specific of the two is chosen.

VI. Results

The overall results for each data set are given in TABLE I. The per user performances on both the validation and testing sets are given in TABLE III. The overall results are broken down into all users, users that have less than 5 recordings on 5 distinct days that are eligible for inclusion in this analysis, and users that have 5 or more days of recordings. It can be seen that there is a significant improvement in accuracy due to the proposed confidence-based rejection mechanism as well as requiring users to provide at least 5 days of recordings.

TABLE II also highlights the problems that are seen in some users on either the validation or testing set performance:

	Overall Accuracy	Overall Rejected Fraction	Overall Accuracy of Accepted Windows	Overall Accuracy of Rejected Windows	Mean Per User Accuracy	Mean Per User Rejected Fraction	Mean Per User Accuracy of Accepted Windows	Mean Per User Accuracy of Rejected Windows
	Training Set							
All Data	99.9%	5.0%	100.0%	97.3%	99.8%	7.2%	100.0%	98.1%
>= 5 Days	99.9%	4.9%	100.0%	97.3%	99.8%	6.6%	100.0%	98.0%
<5 Days	99.8%	6.6%	100.0%	97.2%	99.8%	8.7%	100.0%	98.3%
	Validation Set							
All Data	95.4%	18.7%	99.8%	76.1%	88.9%	30.3%	92.2%	78.4%
>= 5 Days	96.9%	15.9%	99.9%	80.8%	95.2%	22.8%	99.5%	83.9%
<5 Days	80.1%	46.8%	97.9%	59.8%	72.3%	50.2%	72.8%	63.9%
	Test Set							
All Data	95.3%	18.4%	99.7%	75.7%	87.9%	34.2%	92.4%	81.0%
>= 5 Days	96.3%	16.2%	99.8%	78.3%	94.2%	23.8%	98.6%	86.3%
<5 Days	79.8%	52.6%	97.7%	63.6%	71.0%	62.0%	75.8%	66.7%

TABLE I: Summary of per signature window accuracy of the system for each data set divided into all data, users with 5 or more days of data and users with less than 5 days of data. The results are presented in overall accuracy per window and mean of the accuracy per user. These differ since different users have different quantities of recorded data.

Issues Breakdown	<5 Days	>= 5 Days	Total
Accepted Window Accuracy Below 85%	5	0	5
Fraction Rejected Above 50%	6	4	10

TABLE II: The number of users that have each issue among those who have less than 5 days of data and those with 5 or more days or data.

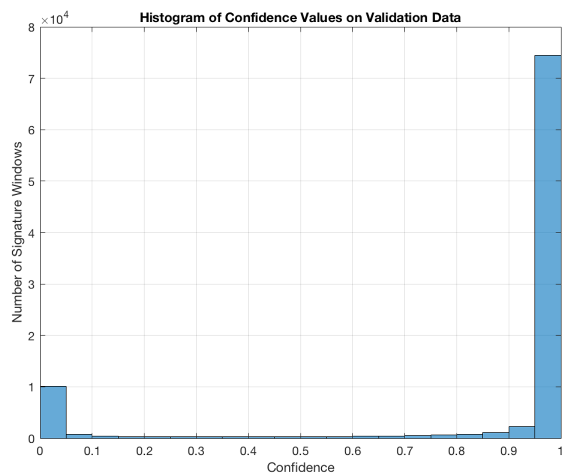

Fig. 3: Distribution of Confidence Values across all the signature windows on the validation set. Most of the confidences are above 0.9 or below 0.1. 0.9 is the chosen acceptance threshold.

some users have very low accuracy even after rejecting due to low confidence and some users have very high rejection rates. It can be seen that most of the users with issues, including all users with low accuracy issues, are among those with less than 5 days of data.

VII. Discussion

Preliminary work for this study focused on a group of 23 who were required to give at least 5 recordings on 5 separate days at different times of day each time. The results of this analysis were quite good. The stark difference in performance between those who have 5 or more days and those who have less in the current data is hypothesized to result from a lack of variability observed in the training data, leading to a un-representative sample. The pre-study work requirements appear to be a lower bound on what is necessary to observe sufficient variability in an individuals signature.

The results show that the system achieves almost perfect accuracy at identifying individuals, provided they meet the minimum data requirements discussed above, a feat that requires the model to have learned features that generalize across different recordings. The confidence measure also adds significant robustness to the system. The model is able to positively identify when it will be accurate, without rejecting too much data, as evidenced by the increase in mean accuracy of 94.2% to 98.6% on the testing set, without and with rejection respectively. The benefit of the confidence based rejection is further evidenced by user 7, who has an accuracy without rejection of 34.4% on the testing set, which improves to 86.2% after rejecting 86.8% of their data.

Preliminary work for this study investigated the performance of several other techniques, including K-nearest neighbour classifiers and random forest classifiers. Other network structures including either less or more number of hidden layers with various numbers of neurons were also evaluated.

Achieving a high accuracy with a low rejection rate demonstrates that OMsignal apparel provides a sufficiently high quality signal for this purpose in a lifestyle context. Identifying an individual from their ECG requires the presence and measurement of subtle features in the signal. The promising performance of this ECG-based identification system provides encouraging evidence that OMsignal apparel, combined with machine learning algorithms will enable finding subtle patterns associated with developing medical conditions. Furthermore, this system demonstrates the capability to manage the occasional periods of variable or corrupted signal that are unavoidable in lifestyle data.

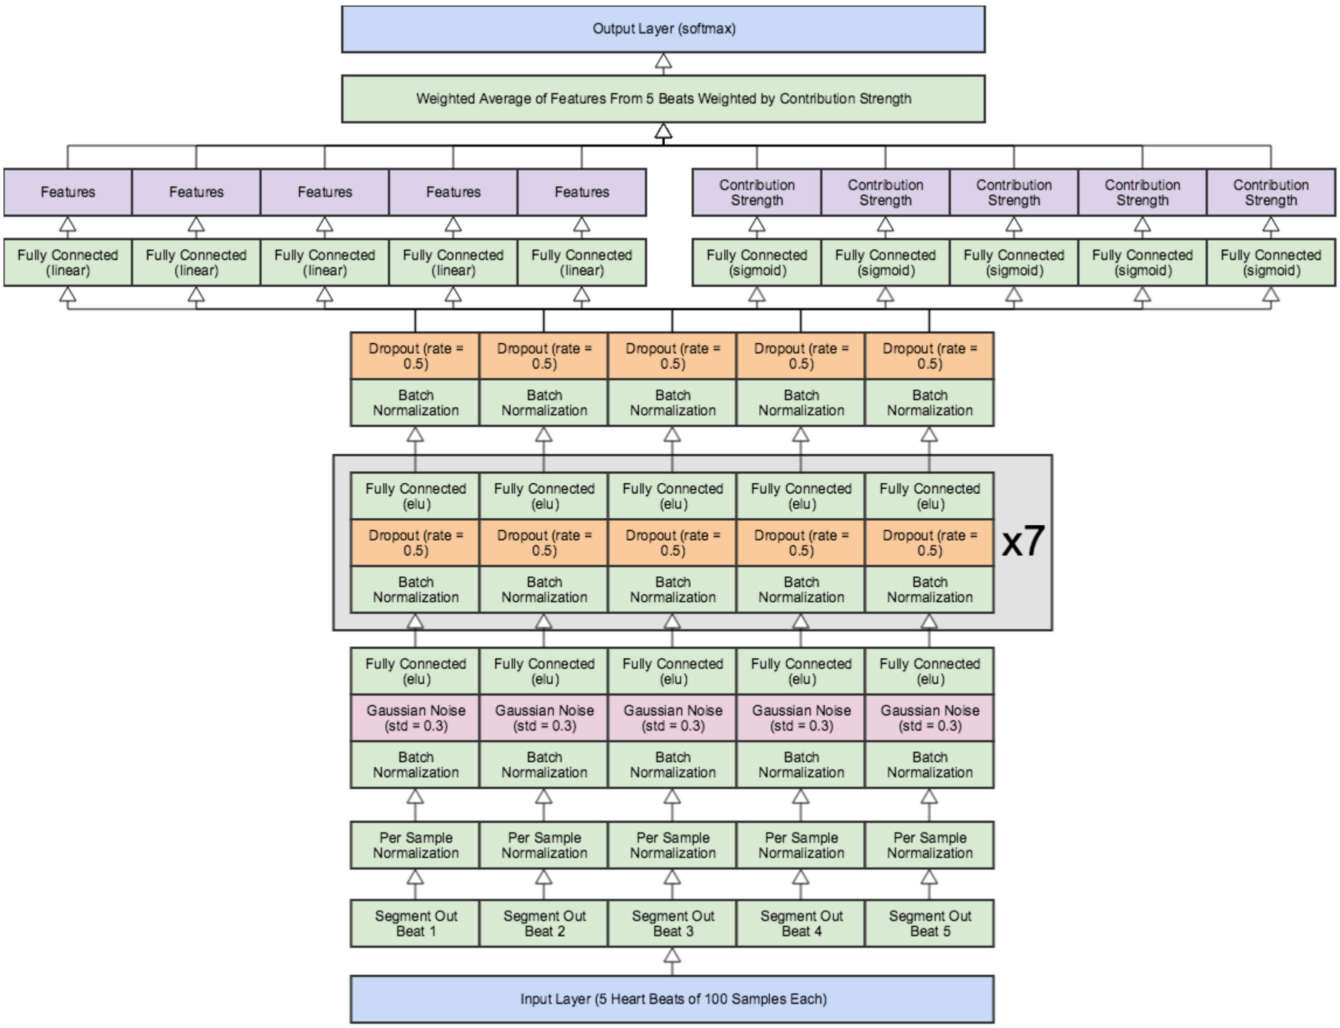

Fig. 4: Structure of the neural network. The weights of the layers are shared across each of the 5 beats. Each beat is individually processed through the network and produces a 300 dimensional feature vector along with each features associated contribution strength to the weighted average at the combination layer. The 5 beats are combined by taking a weighted average of the feature vectors from the 5 beats.

	Validation Data		Testing Data	
User ID	**Rejected Fraction**	**Accuracy of Accepted Windows**	**Rejected Fraction**	**Accuracy of Accepted Windows**
3	16.6%	100.0%	22.1%	99.5%
4	8.5%	100.0%	9.4%	100.0%
5	28.3%	100.0%	8.5%	100.0%
6	27.9%	100.0%	26.3%	99.7%
7	64.2%	98.9%	86.8%	86.2%
8	59.2%	100.0%	77.0%	88.4%
9	0.5%	100.0%	0.4%	100.0%
10	20.4%	99.8%	10.2%	100.0%
12	1.3%	100.0%	1.5%	100.0%
13	47.3%	98.2%	29.8%	100.0%
14	6.6%	100.0%	9.5%	100.0%
15	16.2%	93.5%	33.3%	100.0%
16	0.9%	100.0%	1.6%	100.0%
17	48.2%	99.8%	51.1%	94.2%
18	54.4%	98.6%	40.2%	100.0%
19	18.9%	99.9%	9.7%	100.0%
21	6.9%	100.0%	4.3%	100.0%
22	31.8%	100.0%	15.3%	100.0%
25	12.5%	100.0%	23.4%	99.9%
27	14.7%	100.0%	9.9%	99.9%
28	14.3%	99.9%	27.6%	99.9%
29	20.7%	100.0%	35.4%	100.0%
31	11.6%	100.0%	11.8%	100.0%
33	15.1%	99.9%	26.5%	99.9%
1	36.2%	98.0%	84.6%	18.2%
2	18.7%	100.0%	98.1%	50.0%
11	20.9%	99.0%	27.8%	100.0%
20	90.5%	9.1%	97.2%	16.7%
23	2.9%	100.0%	47.4%	97.1%
24	8.2%	100.0%	40.5%	100.0%
26	85.6%	98.5%	86.5%	100.0%
30	94.9%	0.0%	48.1%	100.0%
32	94.2%	50.8%	27.9%	100.0%

TABLE III: Fractions of signature windows rejected and identification accuracy rates of accepted windows for each user. The users beginning with 1 below the double horizontal bar are the users with less than 5 days of valid data when the results were computed.