# Ballistocardiography-based Authentication using Convolutional Neural Networks

by

Josh Hebert

A Thesis

Submitted to the Faculty

of the

WORCESTER POLYTECHNIC INSTITUTE

In partial fulfillment of the requirements for the

Degree of Master of Science

in

Computer Science

by

May 2018

APPROVED:

Professor Krishna Venkatasubramanian, Thesis Advisor

Professor Robert Walls, Thesis Reader

Professor Craig E. Wills, Head of Department

# Abstract

This work demonstrates the viability of the ballistocardiogram (BCG) signal derived from a head-worn device as a biometric modality for authentication. The BCG signal is the measure of an individual's body acceleration as a result of the heart's ejection of blood. It is a characterization of an individual's cardiac cycle and can be derived non-invasively from the measurement of subtle movements of a person's extremities. Through the use of accelerometer and gyroscope sensors on a Smart Eyewear (SEW) device, derived BCG signals are used to train a convolutional neural network (CNN) as an authentication model, which is personalized for each wearer. This system is evaluated using data from 12 subjects, showing that this approach has an equal error rate of 3.5% immediately after training, and only marginally degrades to 13% after about 2 months, in the worst case. We also explore the use of our authentication approach for individuals with severe motor disabilities, and observe that the results fall only slightly short of those of the larger population, with immediate EER values at 11.2% before rising to 21.6%, again in the worst case.. Overall, we demonstrate that this model presents a longitudinally-viable authentication solution for passive biometric authentication.

## Acknowledgments

I would like to thank the many people, that, with their support, made this project possible:
My advisor, Professor Krishna Venkatasubramanian, for providing the continued direction and apt criticism that made this project a success.
Professor Robert Walls, who graciously agreed to be the reader of this paper.
Both Professor Jeanine Skorinko and Kelly Charlebois and all of the amazing staff of TechAccess, who coordinated the data collection process from the many individuals sampled in this report.
Finally, my fellow graduate students of B17, who provided constant inspiration and encouragement throughout my time at WPI.

# Contents

# List of Figures

# List of Tables

# 1 Introduction

In recent years, there has been an explosion in the proliferation of head-mounted wearable devices, with products such as smart eyewear (SEW) devices [1, 2, 3, 6], smart ear-buds [20], and smart headwear [4]. Most of these devices come with in-built accelerometer and gyroscope sensors which can measure a variety of the wearer's head and body movements. One of the potential uses of such head-mounted wearable devices is to provide new sources of biometrics for authentication [24].

Our principal **goal** in this thesis is to develop a passive authentication approach for head-mounted wearable devices (i.e., it works without requiring explicit action or gestures of any kind from the wearer). Passive authentication approaches are more difficult for adversaries to copy and spoof. Physiological signals (e.g., electrocardiogram (ECG) and electroencephalogram (EEG)) make good biometrics for enabling passive authentication due to their inherent limited observability for adversaries. However, typical head-mounted wearable devices do not come with physiological sensors. Hence, our strategy is to use measurements from the two *movement sensors* (i.e., accelerometer and gyroscope), which are commonly available on such devices, and use them to derive something called a **ballistocardiogram (BCG)**. BCGs represent the body's motion as the blood flows through it, in response to the beating of the heart, and thus capture the characteristics of the cardiac cycle [11]. We derive the BCGs used for authentication by asking each subject to sit still and then capturing and filtering the subtle and involuntary[1] head movements, which their body makes as a result of the pumping action of the heart. In fact we divide the sensor measurements into fixed-duration *segments* and derive six versions of BCGs (which we refer to as *BCG waveforms*), one for each axis of accelerometer and gyroscope, from each segment.

The BCG waveforms from several segments are then used to train a Convolutional Neural Network (CNN) classifier which acts as the *authentication model*). Each authentication model is subject-specific (i.e., personalized), therefore different subjects have their own models. Once the models are trained, *authentication* is done by collecting one (or more) segment-long accelerometer and gyroscope measurements from an *unknown individual* wearing the head-mounted wearable device. The sensor measurements in the segment are filtered to derive six BCG waveforms, which

[1] Here, we use the term involuntary in medical sense to mean not under the conscious control of a person [5].

are then fed into a subject-specific model (belonging to the subject in our system the unknown individual claims to be). This model determines if these newly derived BCG waveforms are similar to the waveforms already seen from the subject earlier during model training. If deemed sufficiently similar, the unknown individual is then successfully authenticated. We evaluate our authentication models using data from 12 subjects. Our approach demonstrates an equal error rate of 3.5% immediately after training and 13% after about 2 months, in the worst case. Using a separate dataset of 6 subjects with motor disabilities, our approach demonstrates a worst-case equal error rate of 11.2% after training and 21.6% after about 2 months.

In this thesis, we specifically use a Smart Eyewear (SEW) device (Google Glass [2]) as our head-mounted wearable device. We chose to use an SEW device because of its relative ubiquity and easy programmability compared to other head-mounted wearable devices. Note that our work is generalizable to any head-mounted wearable device and is in no way limited to SEW devices or to Google Glass.

The **contributions** of this thesis are then three-fold. (a) A passive authentication approach that uses BCGs derived from accelerometer and gyroscope measurements of subtle and involuntary head movements. (b) Demonstration of the viability of the authentication approach longitudinally, using accelerometer and gyroscope measurements collected from 12 subjects over the course of approximately 2 months. (c) Demonstration of the promise of our authentication approach for people with motor disabilities.

## 1.1 System and Threat Models

The principal problem that we address in this thesis is user authentication. That is *to determine if the current wearer of an SEW device (or more generally, a head-mounted wearable device) is the person that they claim to be.* In order to achieve this, our system is comprised of two distinct components: the user-worn SEW, and the remote authentication server. In our system, raw sensor measurements are recorded by the SEW and sent to the authentication server. Once received, these readings are refined into BCGs, and used to either train a user-specific authentication model, or are presented to an already trained model for an authentication decision. Summarily, in our system, all processing is performed remotely, leaving the SEW as simply a data collection device. This has the advantage of requiring very little in the ways of hardware in the SEW, furthering our narrative that such a system can be implemented for existing hardware without any specific modifications required.

We assume that the **threat** to our authentication approach comes from adversaries trying to declare themselves to be a particular subject (i.e., the victim) and try to mimic their head movements to try to authenticate successfully. For the purposes of this thesis, we assume that adversaries: (1) do not have access to the model used for authentication, (2) cannot pollute the model during the training stage, and (3) do not have access to any form of cardiac signal from the victim's past or present.

## 1.2 Thesis Organization

For the remainder of this thesis, Section 2 discuss concepts integral to understanding this thesis, as well as similar work done in this field. Later, Section 4 illustrates the process used to extract data from a head-worn SEW, along with the pre-processing and transformation techniques used to convert them into BCG waveforms. Section 5 describe the layout of the authentication system itself, as well as the methodology with which the parameters were chosen. Subsequently, Section 6 demonstrates that the work presented here presents a high level of accuracy through the use of our collected dataset. Additionally, in Section 7, we have explored the viability of this authentication system in the context of individuals with severe motor disabilities. Finally, in Section 8, we summarize the contributions of this thesis and explore possible avenues of future work.

# 2 Background

In order to better describe the work of this thesis, we discuss two integral concepts used at length throughout the course of this work. These concepts, *Ballistocardiography* and *Convolutional Neural Networks* form the backbone of this paper, and are discussed below.

## 2.1 Ballistocardiography

Ballistocardiography is a non-invasive method of characterizing a person's cardiac cycle, based on the measurement of the body's motions as the blood flows through the body in response to the two phases of the beating of the heart [11]: the diastolic phase, when the heart is resting and the chambers are simply being filled with blood, and the systolic or ejection phase, when the heart contracts to pump the blood throughout the body. The force needed to pump the blood travels throughout the body resulting in a subtle, but observable, movement.

Therefore, if we can measure the movement of a person's body part, we can use ballistocardiography to derive a measurement of their cardiac cycle [12]. Ballistocardiography works by normalizing and filtering the movement data to remove artifacts and thereby derive cardiac signals. It has been shown that using ballistocardiography, accelerometer and gyroscope measurements can be converted into several physiological signals, such as blood volume pulse (BVP), respiration [11], and *BCG* [12]. A sample BCG can be seen in Figure 1, with major points labeled.

## 2.2 Convolutional Neural Networks

Much like traditional neural networks, convolutional neural networks (CNN) allow the results of predictions to be backpropagated throughout the model, such that the resultant network will classify new data more correctly going forward. However, unlike traditional neural networks, CNNs add some number of convolutional layers, each of which has a number of filters to be tuned by backpropagation. As a result, CNNs are able to learn filtering operations to apply to data to make relevant features much more prevalent, such that when they are fed into a densely connected neural network, the accuracy is greatly improved.

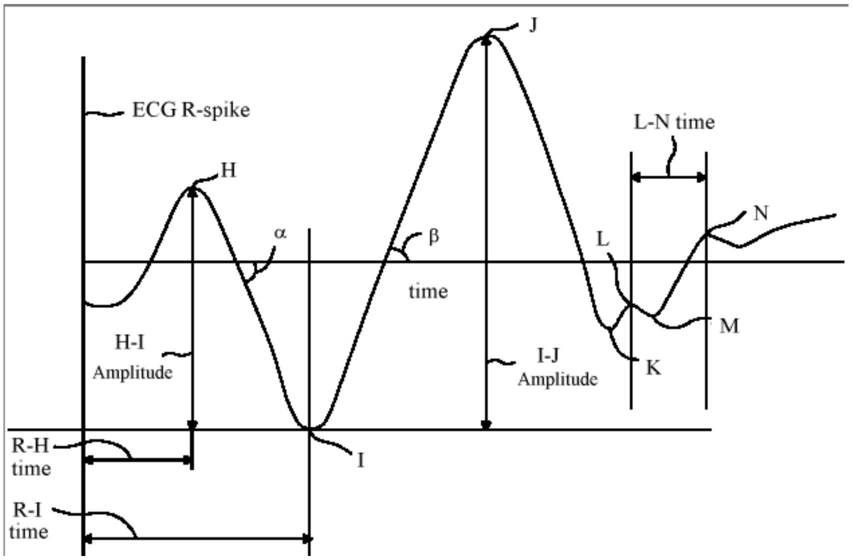

Figure 1: **Typical BCG waveform, with major spikes at H and J, representing the two major contractions of the heart during the systolic phase.**

With regards to classification, often, in the case of time series data such as the cardiac data used here, a recurrent neural network (RNN) will be used to make predictions based not only on previously seen features, but also transient memory stored within the network itself[8]. Graves et al. [9] demonstrate that such a network can be used to recognize speech patterns with a high degree of accuracy. However, in our particular case, we would prefer that our network not maintain an internal state, such that each data point be evaluated independently of its neighbors. This can be though of as analogous to fingerprint-based biometrics; fingerprint based authentication only evaluates the print at a single point in time, not its state over a period of time. In our case, a single waveform can be thought of as a user's cardiac fingerprint. Conversely, CNNs have been demonstrated to be able to classify discrete samples, such as fingerprints[7] and even time-series based electrocardiogram (ECG) readings. The later application is demonstrated by Zhang et al. [25], where it is shown that across a wide population of individuals, a CNN can be trained to identify a user based on readings of ECG bio-potentials, which, like BCGs, are based upon cardiorhythms. In this way, we believe it to be possible to substitute a BCG in place of the ECG and still obtain comparable results.

# 3 Related Work

Biometric, and more explicitly, cardiac-based identification of users is hardly a new problem. In [17], it is shown that by aggregating a variety of measurements collected passively from a user, it is possible for a support vector machine (SVM) to differentiate users with a reasonable degree of accuracy, managing to correctly identify a user 92% of the time in the best case.

Similarly, and perhaps more closely related to the work proposed here, ballistocardiography has been tried for authentication purposes [10, 12, 23]. In [10, 23], ballistocardiography was used on movement data collected from the person's torso, which makes it substantially easier to accurately detect the person's cardiac properties. However, the nature of ballistocardiography precludes us from using these results because as we move away from the heart, the noise in the derived BCG waveform increases dramatically [12]. This is derived from the fact that the original attempt to identify a subject based BCG signals measured using hand-worn movement sensors, produced only a 66% accuracy rate [12].

Authentication have been previously explored in the context of SEW devices as well. In [16], the authors use the notion of head movement in response to a specific song as a signature for authentication. The head movements used by this system are, however, very simple and can be easily spoofed by an adversary. In [21], the authors induce white noise into the subject's skull, the response to which is then picked up to determine who is wearing the device. This is a better solution in terms of spoofing resistance, but it requires the use of bone conductance speakers, which are something that not all SEW devices possess. This approach requires around 23 seconds of inducing noise before it can identify the subject, which is too slow to be practical. Further, the use of white-noise for authentication has been found to be uncomfortable to some subjects, as noted by the authors. In [18], the authors present an approach for user identification in SEW devices, which uses the blinking and head movement pattern of the subject while they watch a short video on the device's screen. However, this approach requires 34 seconds to identify the user, presenting an obvious temporal barrier to usefulness.

We address the major limitation present in all aforementioned papers, which is that not a single one covers longitudinal accuracy; each relies solely on data collected in a single sitting. As the ultimate goal of this research is to create a system capable of authenticating a user, it is essential that it be resilient to possible changes over time in the biorhythms of the users. As such, a corollary goal of this research is to demonstrate that the designed solution is capable of still reliably authenticating users days or even weeks later.

# 4 Approach

Our authentication approach has five stages. *Data collection*, and subsequently, *data preprocessing*, which describes our process for gathering of the accelerometer and gyroscope measurements from the SEW device followed by a preprocess step to synchronize them. *BCG extraction*, which shows the derivation of BCG from the head movement data. *Model training*, which uses the derived BCG to train a convolutional neural network (CNN) classifier as the authentication model. Finally, the *authentication process* itself using the trained CNN classifier to authenticate someone.

## 4.1 Data Collection Protocol

The first stage in our authentication approach is to collect accelerometer and gyroscope measurements from subjects wearing the SEW device (a Google Glass in our case). We *standardize* the data collection protocol across all the subjects to minimize motion artifacts and to ensure reproducibility of results. The data collection protocol used is as follows: In particular, we ask the subjects to sit comfortably, upright, and still. Their hands are resting, such that their elbows form a 90-degree angle. We then situate an SEW device on their faces such that the upper edge of the device aligns with their brow. We make sure that the device fits comfortably (not pinching or sitting unevenly) on the subjects' heads. Subjects wearing prescription glasses are asked to remove them to minimize SEW fitting issues. An example of this posture is shown in Figure 2.

## 4.2 Data Preprocessing

During the data collection, the accelerometer and gyroscope sensors in the SEW device were set to sample at 50 Hz. The SEW device relayed measurements from the accelerometer/gyroscope sensors wirelessly to a nearby laptop, where the measurements were stored. As a result of the data collection process, we obtained six discrete, *raw sensor-streams*: the three axes of the accelerometer and three axes of the gyroscope measurements. Sensor data from any Android device, like Google Glass, is not guaranteed to align exactly to a particular sampling rate. Therefore, a sampling rate of 50Hz resulted in an inter-sample-interval of anywhere from 5

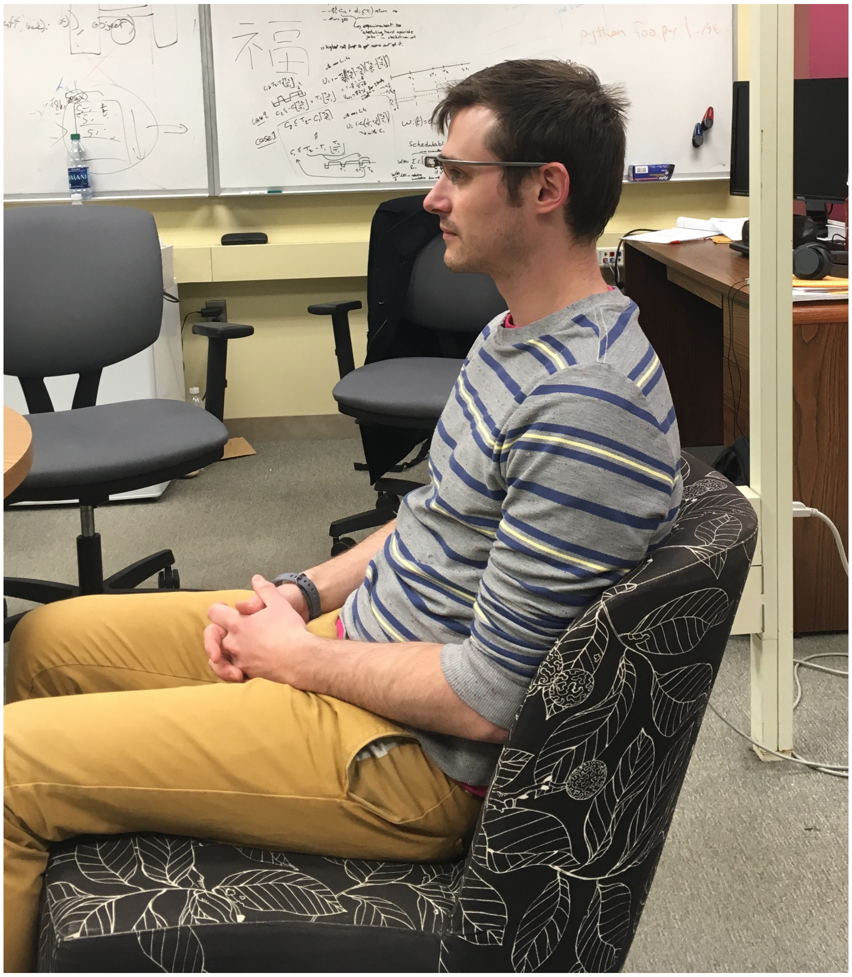

Figure 2: **Example positioning and posture of a subject during data collection.**

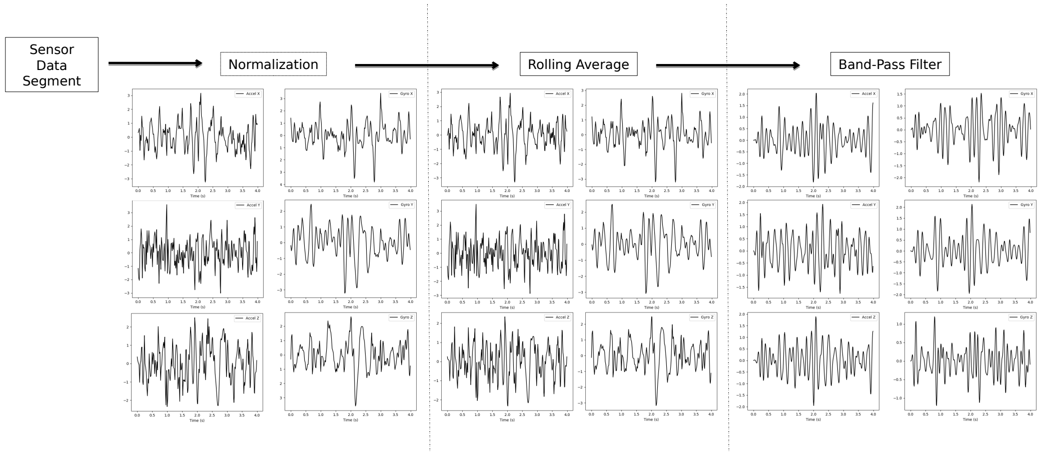

Figure 3: **The process of generating a BCG waveform from the segmented data of the three axes of accelerometer and gyroscope sensors.**

to 20ms. Additionally, there exists no guarantee that the samples recorded from the gyroscope and accelerometer measurements are synchronized or aligned in any way. In order to address these two concerns, we preprocess the sensor-streams. In this regard, once the raw sensor-streams are collected, we truncate the beginnings and endings of both the gyroscope and accelerometer measurements such that the time-stamp of the first and last samples of both sensor measurements are as close as possible. We then interpolate the data and align the samples with one another. The first sample of the gyroscope and the accelerometer now share they same timestamp, as do all the subsequent points in the sensor-stream. All subsequent analysis use these *preprocessed sensor-streams* instead of the raw sensor-streams.

## 4.3 Deriving the BCG Waveform

Once we obtain the preprocessed sensor-streams of accelerometer and gyroscope measurements from individual subjects, the next step in our approach is to derive the BCG. In this regard, we first divide each preprocessed sensor-stream into overlapped *segments* of size $w$ seconds. Between two sequential segments, there is a $w - 1$ seconds overlap. Hence, two segments with $w = 3$ seconds would share 2 seconds of data. Then, inspired by [12], we perform a three-step BCG derivation process. (1) *Normalization:* We normalize each of the six sensor-streams to have zero mean and unit variance within each segment. (2) *Rolling Average Filter:* We then subtract a rolling-average filter of 35 samples from each sensor-stream to cor-

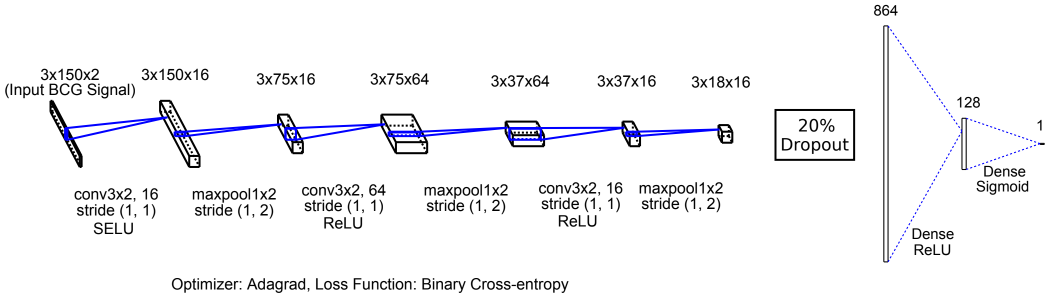

Figure 4: **Final CNN topology and the chosen parameters used by our subject-specific authentication model.**

rect for large motions, as well as gyroscope and accelerometer drift. (3) *Band-Pass Filter:* Finally, we apply a 4th-order band-pass Butterworth filter with cutoff frequencies at 4 and 11Hz to each sensor-stream. Figure 3 shows the main stages of the BCG generation for one segmented of a sensor stream. In all, we derive six BCG signals (heretofore referred to as *BCG waveforms*); one per axis of accelerometer and gyroscope; *per segment*. These six BCG waveforms are then used as input for our authentication model.

## 4.4 Model Training and Authentication

We now construct an authentication model that learns the features of each subject's BCG waveforms and use their uniqueness to authenticate the subject at a later time. We use a convolutional neural network (CNN)-based classifier as our authentication model. We use CNNs due to their demonstrated ability to effectively classify time-series such as electrocardiogram (ECG), which, like BCG waveforms, are representations of the cardiac rhythm [25]. In particular, our authentication approach has two stages, the training stage and the authentication stage.

### 4.4.1 Model Training Stage

During the *training stage*, the goal is to enroll the subject into the authentication system by training a **subject-specific authentication model** for them. This requires the collection of accelerometer/gyroscope measurements from a subject wearing the

SEW device for $\Delta_E$ time-units. We use the 6 BCG waveforms from $\Delta/w$ segments (obtained from the subjects in our dataset (see Section 5.1)) to *train a CNN-based classifier* as an authentication model for a subject. During training, we label the BCG waveforms from the subject's own segments as belonging to the positive class, while we consider the BCG waveforms from other subjects in out dataset part of the negative class. As such we use a one-versus-all strategy that allows each subject to get a custom, subject-specific model.

In our CNN setup, the network has a particular emphasis on the relationships between the three axes of accelerometer and gyroscope derived BCG waveforms. To this end, we re-arrange each $w$ segment to form a $2 \times 3 \times (w \times 50)$ tensor to be used as input to the CNN. Here, the first dimension refers to measurement source (accelerometer or gyroscope), the second refers to axis, and the third refers to time (where 50 is the sampling frequency in Hz). In designing our CNN we only fixed the inputs and outputs. As an output, the final layer is expressed as a single neuron reporting a value between 0.0 and 1.0, indicating the *confidence of the CNN* that the given sample should be accepted. In order to determine the intermediate layers of this CNN and their parameters, we took a *genetic algorithm-based approach*. More details on the algorithm are given in Section 5.3. Figure 4 shows the final topology of the CNN we used. We initialize all weights in the network with random values and train the CNN over the course of 100 epochs, at which point the loss of the model stabilizes to minimum. At this point, the model has been trained and is ready to perform authentication.

#### 4.4.2 Authentication Stage

Once the model is trained, it is capable of performing *authentication*. To authenticate an unknown individual, we collect one (or more) $w$-second segment of raw accelerometer and gyroscope measurements while they (i.e., the unknown individual) are wearing the SEW device. We derive BCG waveforms for each segment. These BCG waveforms are fed into the CNN of the authentication model, which produces a confidence score pertaining to how strongly it believes that the unknown individual is the same as the subject for whom the model was trained. If this confidence score is greater than a *decision threshold* $T$, the authentication is complete, and the subject is deemed authenticated. $T$ refers to the minimum confidence of the CNN required to accept a given input tensor. A higher value of $T$ will make it more difficult to mistakenly accept negative class BCG waveforms, but will also hinder the acceptance of positive BCG waveforms.

The authentication step in our approach need not be a one-shot event, it can be repeated over several $w$-second segments. We use the variable $s$ to denote the number of segments (hence the number of *authentication attempts*) that were performed during authentication. In our approach, if any one of the $s$ sequential segments is accepted, the wearer is authenticated. Note that, since between any two $w$-second segments their is a $w-1$-second overlap, a sequence of $s$ segments will only require $s + w - 1$ seconds to measure.

# 5 Experimental Setup

We now describe the dataset and metrics that we used to validate our approach, along with our experiments that establish the various parameters used in our approach.

## 5.1 Dataset

In order to validate our approach we collected data from 12 volunteer subjects. We assembled these subjects from three distinct populations. After collecting data from fellow graduate students, we collected data from student participants who volunteered through WPI's SONA system, as well as from faculty from TechAccess, an organization aimed at assisting individuals with disabilities. By assembling our dataset in this way, we see a number of advantages. Namely, it allows us to create a dataset largely balanced in regard to gender, with a relatively large spread in age with high spread. This creates a more realistic dataset more representative of a real-world scenario. IRB approval was obtained for all data collection. Subjects were informed that no personally identifiable information would be collected from them, and asked to sign a consent form acknowledging the data that was to be collected.

From each subject we collected 10 minutes of accelerometer and gyroscope measurements using the Google Glass, as well as the age and gender of the subject. We additionally maintained internal records that were used to keep track of participants such that subsequent sessions were accurately matched with older participants. These internal records have been withheld and will not be published.

Our goal with this project was to measure the *longitudinal effectiveness* of our authentication approach. Consequently, we collected **three 10-minute sessions** from our subjects over approximately three months. During a session, sitting still for 10 minutes can be tedious, therefore, we broke up the 10-minute session into *five 2-minute intervals*. Between each interval subjects were given ample time to take a break and readjust themselves. During the 2-minute intervals, we asked the subjects to either focus on an eye-level post-it card stuck to the wall nearby, or close their eyes. Further, we specifically asked the subjects to not focus on the screen of the Google Glass display as it was found to be uncomfortable for anyone to do so for long periods of time. During any of the five 2-minute intervals, if the subjects moved in any way, the data collection for that interval was stopped and repeated.

Table 1: **Dataset demographics**

| Set | Avg. Age | Std. Dev. Age | ♂ | ♀ |
|---|---|---|---|---|
| Validation | 32.83 | 13.13 | 4 | 8 |
| External | 28.50 | 10.91 | 7 | 3 |
| **All** | **30.86** | **12.09** | **12** | **10** |

All but one of the 12 subjects provided us with three sessions of data. For every subject, the time-difference between any two session was at least 10 days. The span of days following training for each session falls roughly in the following bins: **Session 2:** 10-32 days, **Session 3:** 28-60 days. As such, the data collected covers a period of time approaching two months.

We refer to these 12 subjects as the *validation set*. In addition, we also collected data from 10 other subjects, which we call the *external set*. The data in the validation set is used to train the models for authentication. As we are building subject-specific models, we generate 12 models from the validation set.

In contrast, the people in the external set include subjects from whom we only obtained one session (due to scheduling reasons) or those subjects whose data was collected in the pilot phase of the data collection. We use the external set for the evaluation of the ability of the subject-specific authentication models to reject data from subjects that they have never seen before. Table 1 summarizes the demographics of our dataset.

## 5.2 Metrics

To be able to evaluate the efficacy of our approach we use the following core metrics: *false acceptance rate (FAR)*, *false rejection rate (FRR)*, and *equal error rate (EER)*. FAR is the fraction of negatively labeled test BCG waveforms that were misclassified as positive. FRR is the fraction of positively labeled BCG waveforms that were misclassified as negative. Similarly, we also use the complements of FAR and FRR, namely *true acceptance rate* (TAR) as the fraction of positively labeled BCG waveforms that were classified as positive, and *true rejection rate* (TRR) as the fraction of negatively labeled BCG waveforms that were classified as negative. The average of TAR and TRR provide us with *accuracy*. Finally, *equal error rate* (EER) is the rate at which FAR and FRR are equal. Even though we compute these metrics for every subject in our dataset, we present summary statistics of these met-

Table 2: **Accuracy for different segment lengths ($w$)**

| **segment length** | **1s** | **2s** | **3s** | **4s** | **5s** |
|---|---|---|---|---|---|
| Accuracy (%) | 89.79 | 94.81 | **97.55** | 96.51 | 98.53 |

rics as an average over all subjects.

## 5.3 Parameter Selection

The data we use for all our parameter selection comes from the first session of the subjects in the validation set. In order to tune our model parameters, we create 12 separate models, one for each of the subjects in the validation set. We then train the models using $\Delta_E = 8$ minutes of data, and tune using the remaining 2 minutes of data. Consequently, each model has 8 minutes of data from one subject in its positive class, and 88 minutes (8 minutes $\times$ 11 subjects) from other subjects in the validation set in the negative class. The remaining 24 minutes (2 minutes $\times$ 12 subjects) of data is used to test each model and tune its parameters based on the test results. All test are done with $s = 1$, i.e., using one-shot authentication using one segment of movement data.

Using this set up, we select two distinct characteristics: the length of the input segment $w$, and the topology and hyperparameters of the convolutional neural network itself.

### 5.3.1 Choosing Segment Length $w$

In order to determine the length of segments (i.e, $w$ seconds) of accelerometer and gyroscope measurement, from which we obtain the BCG waveforms, we evaluate values of $w$ from 1 second to 5 seconds in discrete steps using the data from the first session as described previously. For each model, we compute the accuracy ((TAR+TRR)/2), and then average the accuracy across all models. We found that the the accuracy increases until $w = 3$ seconds and then flattens, as shown in Table 2. We therefore choose our segment length as $w = 3$ seconds.

| Trait | Possible Values |
|---|---|
| Number of convolutional layers | 1, 2, 3, 4, 5 |
| Number of filters in each convolutional layer | 8, 16, 32, 64, 128, 256 |
| Activation function following each convolutional layer | ReLU, SELU, tanh, sigmoid |
| Dropout between the last convolutional layer and the first dense layer | 0%, 20%, 40%, 50% |
| Number of dense layers | 1, 2, 3 |
| Number of neurons in each dense layer | 16, 32, 64, 128, 256, 512 |
| Activation function following each dense layer | ReLU, SELU, tanh, sigmoid |
| Dropout following each dense layer | 0%, 20%, 40%, 50% |
| Optimizer | Adam[15], SGD, Adagrad[13] |
| Loss function | Binary Cross-entropy, Mean Squared Error, Mean Absolute Error, Mean Squared Logarithmic Error, Kullback-Leibler Divergence[14] |

Table 3: **CNN network genome traits and possible values.**

#### 5.3.2 Selecting Convolutional Neural Network Hyperparameters

As mentioned before, we took a *genetic algorithm-based method* to find the parameters of our CNN-based authentication model. We optimized over a total of 10 traits, shown in Table 3, that capture all the elements associated with convolutional and dense layers in the network.

Our algorithm first generates 20 CNNs, randomly selecting values from all possible traits and training them the first four 2-minute intervals (which forms $\Delta_E$) of the first session of data collection and evaluating their performance using the final 2 minutes of data in the first session. We use the scoring function $(FAR)^2 + (FRR)^2$ to evaluate each CNN. This is done to favor models that have a balanced error rate, as opposed to ones that vastly favor one metric.

In each generation we select the best 25% CNNs ($20 \times 0.25 = 5$) as the ”parents” for the next generation. Additionally, we select 3 CNNs at random from the bottom $75\%$ as parents. As such, each generation has 8 parents. In order to replenish the next generation back up to 20, 12 children are created. This is done by selecting two of the parents at random from the pool of 8, and for each trait, randomly selecting between the two values held by the parents. A 15% mutation rate is applied to the children in each generation, in that for each trait in the child, there is a 15% chance that, rather than inheriting from one of the two parents, the trait is selected at random from all possible values. We ran this algorithm for 10 generations and selected the CNN with the overall best score as our model. The generated topology and parameters are shown in Figure 4.

# 6 Evaluation

We now train a CNN classifier (parameterized with the values chosen above) as the authentication model for each of the 12 subjects in our validation set. Subsequently, we *evaluate the efficacy of these models*, in realistic settings, by using the yet unseen samples of our dataset. This simulates the actions of the primary adversary of our threat model; someone who views a subject (i.e., victim) authenticating and then tries to mimic their head movements to authenticate successfully. As the victim makes no explicit gestures and sits still during authentication, the adversary cannot observe any specific movements of the victim's head and is therefore reduced to using their own head movements to try to authenticate.

## 6.1 Evaluation Data Categories

We evaluate our approach based on three categories of data from the dataset. (1) *Positive Validation:* Here, we test each subject's model with BCG waveforms derived from the subject's own yet unseen (i.e., not used for training[2]) segments. These "positive" BCG waveforms compute the TAR (referred to as **Validation TAR**) of our authentication approach. (2) *Negative Validation :* Similarly, we also test the BCG waveforms from rest of the segments in the validation set, i.e., those belonging to other 11 users. The TRR (referred to as **Validation TRR**) thus derived demonstrates how well the model can prevent other subjects in the validation set from impersonating a particular subject. These BCG waveforms do not possess any temporal meaning with respect to the model being evaluated, and are treated as one large set. (3) *Negative External:* Finally the external set is used to generate a second TRR (referred to as **External TRR**), which demonstrates that our model has not overfit and is equally capable of denying entry to subjects whose data it has never seen in any form. Similar to (2), these BCG waveforms also do not possess any temporal meaning with respect to the model being evaluated.

Note that we have used the last 2-minutes of data from the subjects in the validation set to select the CNN parameters and choose the appropriate segment length. We maintain that this will not affect the correctness of our results because at no point was this last 2-minutes of data used in the actual training of models.

[2] This include segments from the final 2 minutes of data in session 1 along with segments from 10 minutes of data from sessions 2 and 3

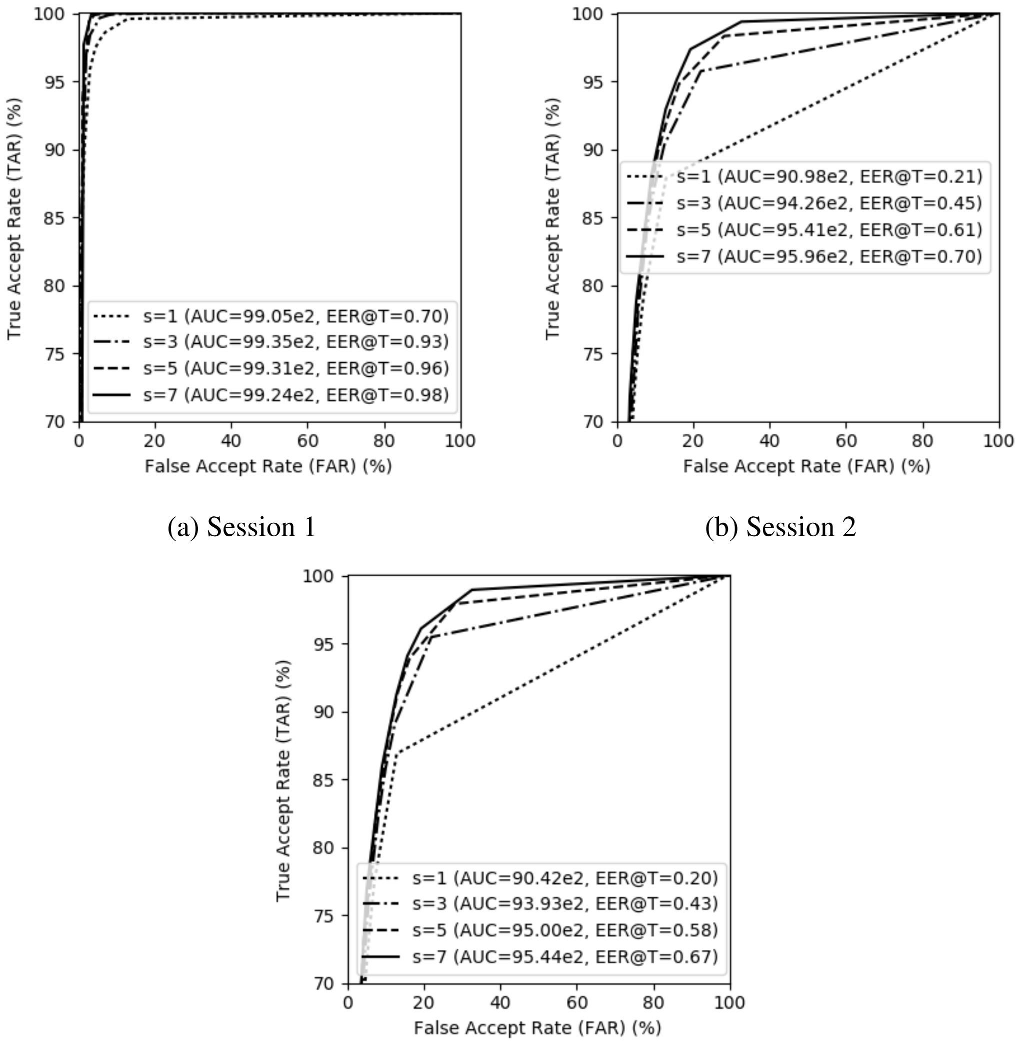

(a) Session 1

(b) Session 2

(c) Session 3

Figure 5: **For each session, we generate ROC curves for each value of $s$ (authentication attempts) by varying $T$ (decision threshold). Note, the y-axis of the graph does not start at the origin.**

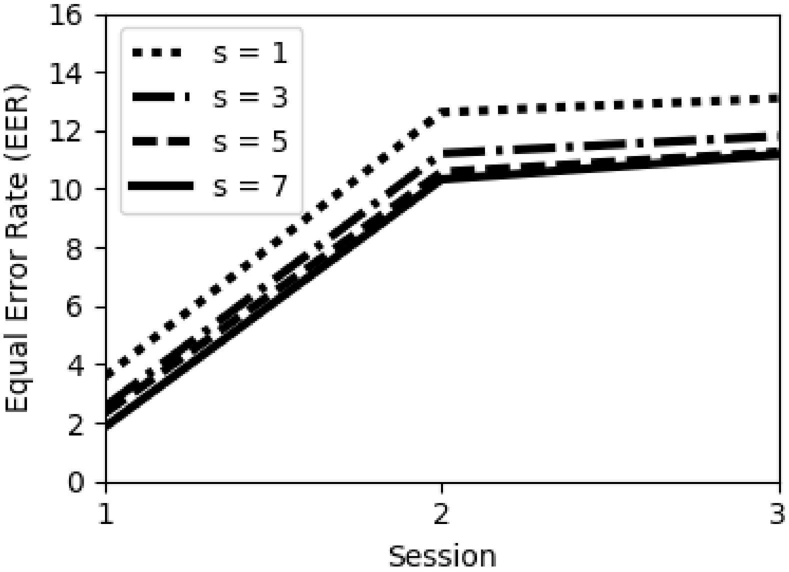

Figure 6: **For each ROC curve, we isolate the EER. As expected, we see a degradation in EER as time between the training and authentication increases.**

## 6.2 Performance Analysis

Figure 5 shows the ROC curves for the 12 trained models, for multiple values of $s$, i.e., authentication attempts. The curves are produced by the varying $T$, the decision threshold that provides the lower bound of the confidence level of the CNN. Here, $TAR$ = Validation TAR, and $FAR$ = $\{(1-$Validation TRR$) + (1-$External TRR$)\}/2$. We generate a total of three ROC graphs, one for each data collection session. It can be seen that overall our approach performs well. In session 1, the ROC curves show that the authentication models are very accurate with area under the curve (AUC) values greater than $0.99$ irrespective of the number of authentication attempts. It is easy to see that the greater the number of authentication attempts, the higher the overall accuracy (i.e., $s = 7$ always outperforms $s = 1$). The accuracy of the authentication models drop in sessions 2 and 3. The performance drop between sessions 2 and 3 is however minimal as denoted by the AUC values. Remember that as our segments are overlapped, consequently, when $s = 7$, it only requires $s + w - 1$ seconds (i.e., 9 seconds) of sensor measurements.

The point in Figure 5 where the ROC curve meet the top left and bottom right diagonal (not shown in the figures) of the graph is the EER value. Figure 6 shows the EER values and its evolution over time. The EER value occurs at different $T$ values depending on how many authentication attempts are allowed and which session's data is being used in the evaluation. The magnitude of the EER value again shows the efficacy of our approach, with EER values below $5\%$ in session 1, and then increasing to around $13\%$ in sessions 2 and 3 for $s = 1$, in the worst case. Again higher values of $s$ produce lower EER values. These results show that one of the future extensions of this approach could be in enabling continuous authentication of wearer, something we plan to explore in the future.

Figure 7 shows how the Validation TAR, Validation TRR, and External TRR vary for our 12 subjects over the three sessions when the best $T$ (in terms of the EER) is chosen. That is, as time goes long, we vary the value of $T$ to keep authentication model at corresponding EER (obtained from the ROC curves) at all times. It can be seen that for session 1, the box plots are very tight for all $s$. However, as we move to sessions 2 and 3, we vary the threshold $T$ such that we remain at the EER. This ensures that our TAR does not drop too precipitously, while still keeping a reasonably high TRR. The External TRR has a higher spread than Validation TRR, however, the medians of both are still very similar to each other.

## 6.3 Summary

These results show the viability of using BCG, collected using head-mounted wearable devices, for authentication. In this section, we have shown that, after training an authentication model with 8 minutes of data, we can authenticate an individual with a high degree of accuracy immediately afterwords using a process of multiple authentication chances. Furthermore, by adjusting the decision threshold as time since training grows, we demonstrate that this system can maintain a reasonably high EER even as extended periods of time has elapsed. We evaluate samples up to two months after training, but data here suggests that this performance would hold well into the future.

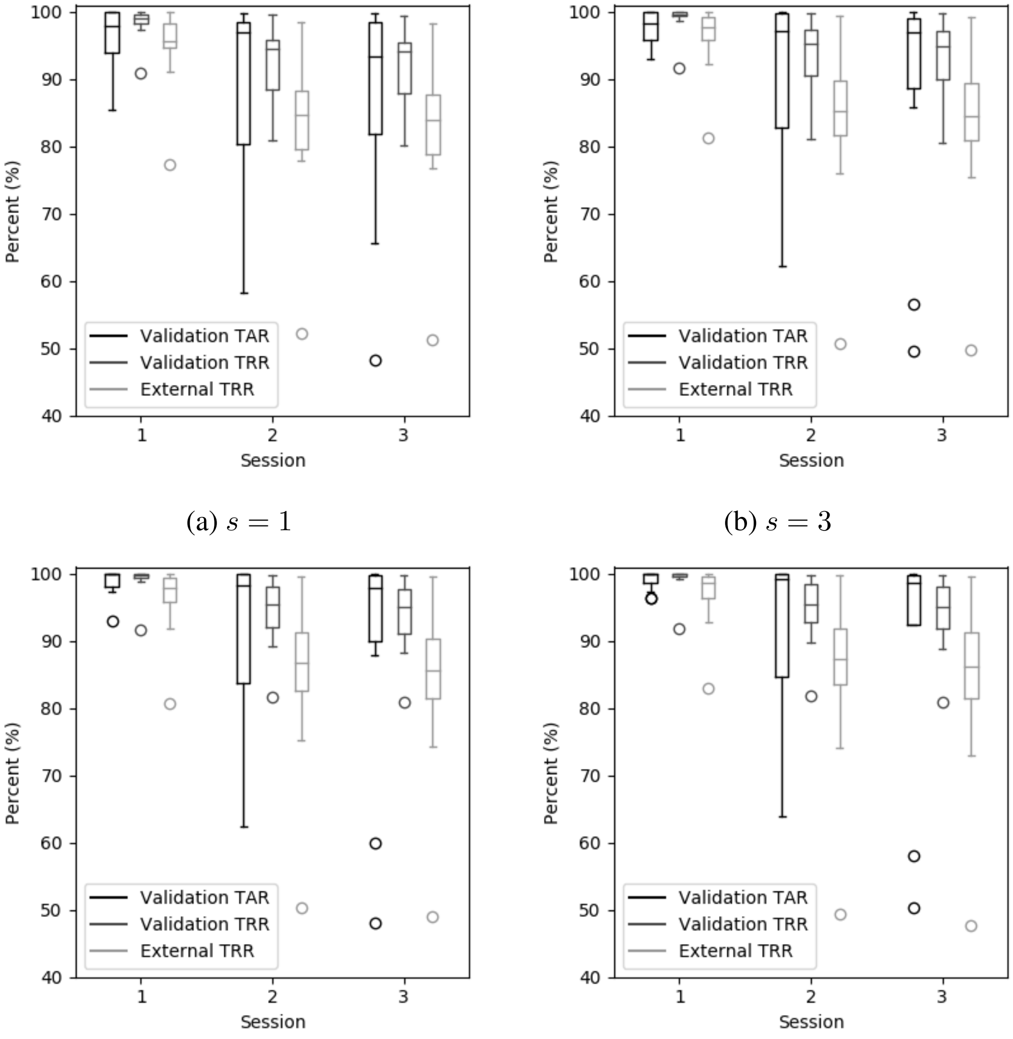

(a) $s = 1$

(b) $s = 3$

(c) $s = 5$

(d) $s = 7$

Figure 7: **Validation TAR, Validation TRR and External TRR for values of $s$ (authentication attempts) between 1 and 7, by using optimal values of $T$ over the three sessions. Note, the y-axis of the graph does not start at the origin.**

# 7 Performance for Individuals with Motor Disabilities

We also deployed our authentication approach on a dataset where the subjects had severe motor disabilities. Such individuals cannot independently authenticate to their computing devices because traditional authentication approaches, like passwords, were not designed for such populations, and therefore, are difficult to use [22]. The passive nature of our approach has the potential to allow individual with motor disabilities to authenticate to their devices with minimal effort and therefore increase their level of independence in using modern computing devices.

## 7.1 TechAccess of Rhode Island

As alluded to above, we worked closely with TechAccess of Rhode Island, a non-profit focused on improving the lives of individuals with disabilities through the use of assistive technology. By working to foster and explore collaborations such as this thesis, they help find solutions that provide functional outcomes for these individuals.

With regard to IRB approval, in the case of individuals from TechAccess that were unable to sign for themselves due to motor impairments, a guardian was permitted to sign on their behalf, with a witness standing by.

Table 4: **Demographics of smaller dataset comprised of subjects with cerebral palsy.**

| Set | Avg. Age | Std. Dev. Age | ♂ | ♀ |
|---|---|---|---|---|
| Subjects w/ Disabilities | 42 | 10.6 | 3 | 3 |

## 7.2 Data Collection

We collected a separate dataset with 6 subjects, shown in Table 4, with non-spastic cerebral palsy for this portion of the work. We obtained the requisite IRB approval and collaborated with a local non-profit to obtain the data. We collected *two sessions* of data from every subject. The second session data, due to scheduling rea-

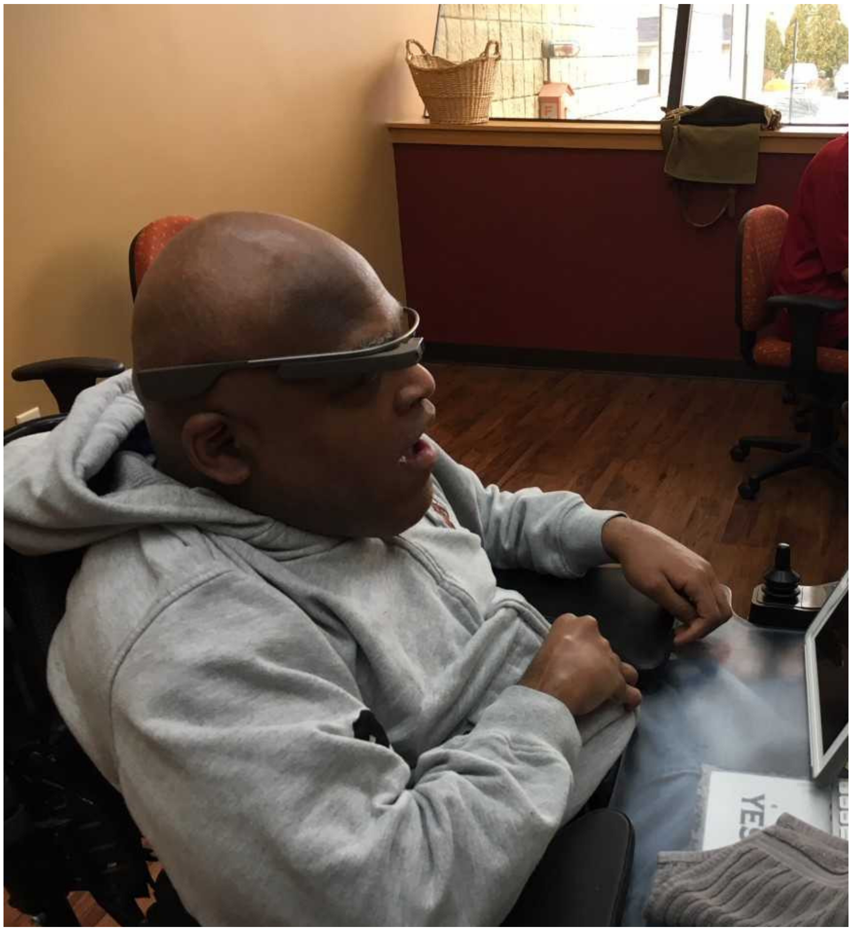

Figure 8: **One subject with motor disabilities providing SEW readings at TechAccess.**

sons, was collected after 15 days for two subjects and 57 days for the other three, after the initial session. Much like the prior section, data was collected from subjects in a controlled environment with the subject remaining as still as possible. However, these individuals were unable to sit in the same chair provided to other participants, due to being wheelchair-bound; this did not impede them from following procedure to the best of their abilities. One such subject is shown in Figure 8.

Once again, we created subject-specific authentication models and repeated the evaluation process described in Section 6 for each of these five subjects. All the parameters for our authentication model were identical to those described in Section 5. We tested these models with all their data collected other than that used for training, and data in the external set mentioned in Section 5.1.

## 7.3 Evaluation

Similar to Section 6, we generate EER values for the ROC curves shown in Figure 9 based on adjusting the value of $T$. Following respectable performance from the curves of the first sessions, in the second sessions, ROC curves with AUC values around $88\%$ show that the EER has degrade more than it had in the larger population, with immediate EER values at $11.2\%$ before rising to $21.6\%$ in the worst case. However, it is very likely that this comes about as a result of the extreme spacing in the sessions in this dataset. In the larger data set, even among third sessions, very few participants had such a large time elapsed, and those that did were balanced by those with more reasonable spacings. In this scenario, almost half of the test set had spacings at the upper extremity of 57 days. We plot the spread of the validation TAR, validation TRR and external TRR in Figure 11, and in these box plots, it can be seen that there is a much larger spread in the performance, with participants with a larger spacing tending towards the bottom of the plot. Additionally, we observe a much larger deviation between the validation TRR and external TRR than was seen in the external set. This is likely due to the fact that the external set comes from individuals not affected by motor disabilities, and as a result, are markedly different than the subjects in this pool.

Finally, we observe the EER values derived from the ROC curves in Figure 10, where we observe a degradation in performance compared to the larger dataset, particularly in the case of $s = 1$. We attribute this to a number of factors: (1) we observed that this group had increased difficulty in maintaining a motionless posture given that their condition often affects their head muscles as well, (2) many of them were assisted by medical devices that would on occasion generate vibrational noise. That being said, our disabilities dataset is too small to make definitive statements. Undoubtedly, if we possessed a much larger pool of subjects with similar disabilities with which to train our authentication models, we speculate that the performance, especially with regard to testing against pools such as the external set, would markedly improve.

However, given the poor state of authentication solutions for people with motor disability, we contend that any approach that increases the authentication independence of such individuals, even in small terms, has merit. In this regard, we view these results as a promising first step.

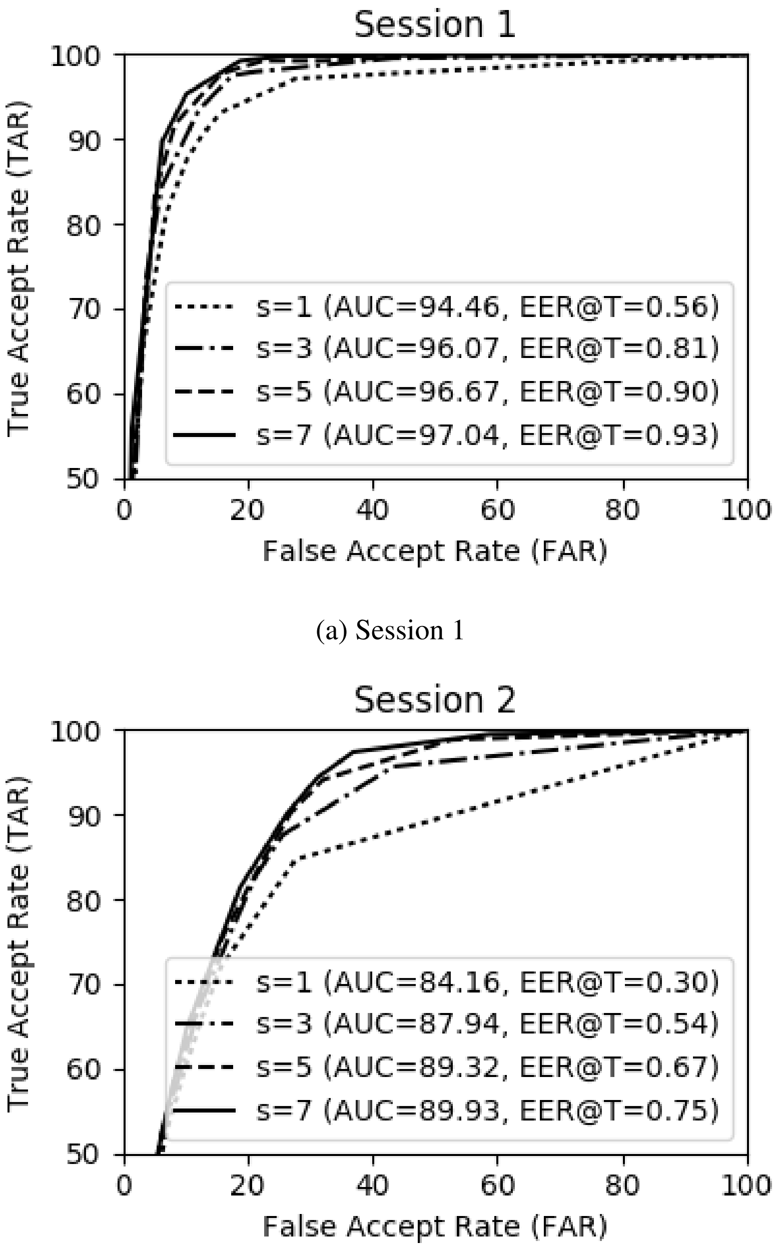

(a) Session 1

(b) Session 2

Figure 9: **For each session collected from individuals with motor disabilities, we generate ROC curves for each value of $s$ (authentication attempts) by varying $T$ (decision threshold). Note, the y-axis of the graph does not start at the origin.**

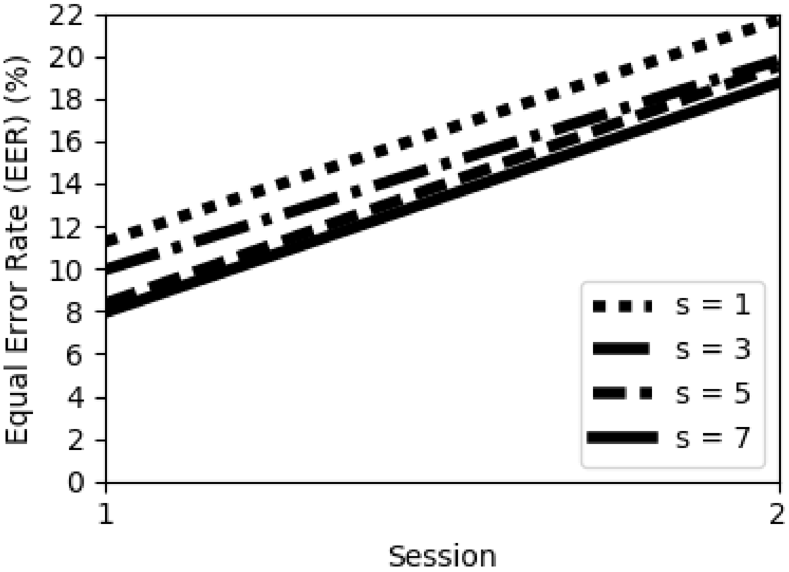

Figure 10: **EER calculated from models trained with data from subjects with motor disabilities.**

## 7.4 Summary

Ultimately, here we have shown the authentication system previously described can be applied to a population which currently has very few options with regard to secure authentication. Furthermore, we have shown that, even with a small population, and a population that often times has immense difficulty with holding perfectly still, we are able to achieve classification results better than any that have been seen in prior ballistocardiography-based work.

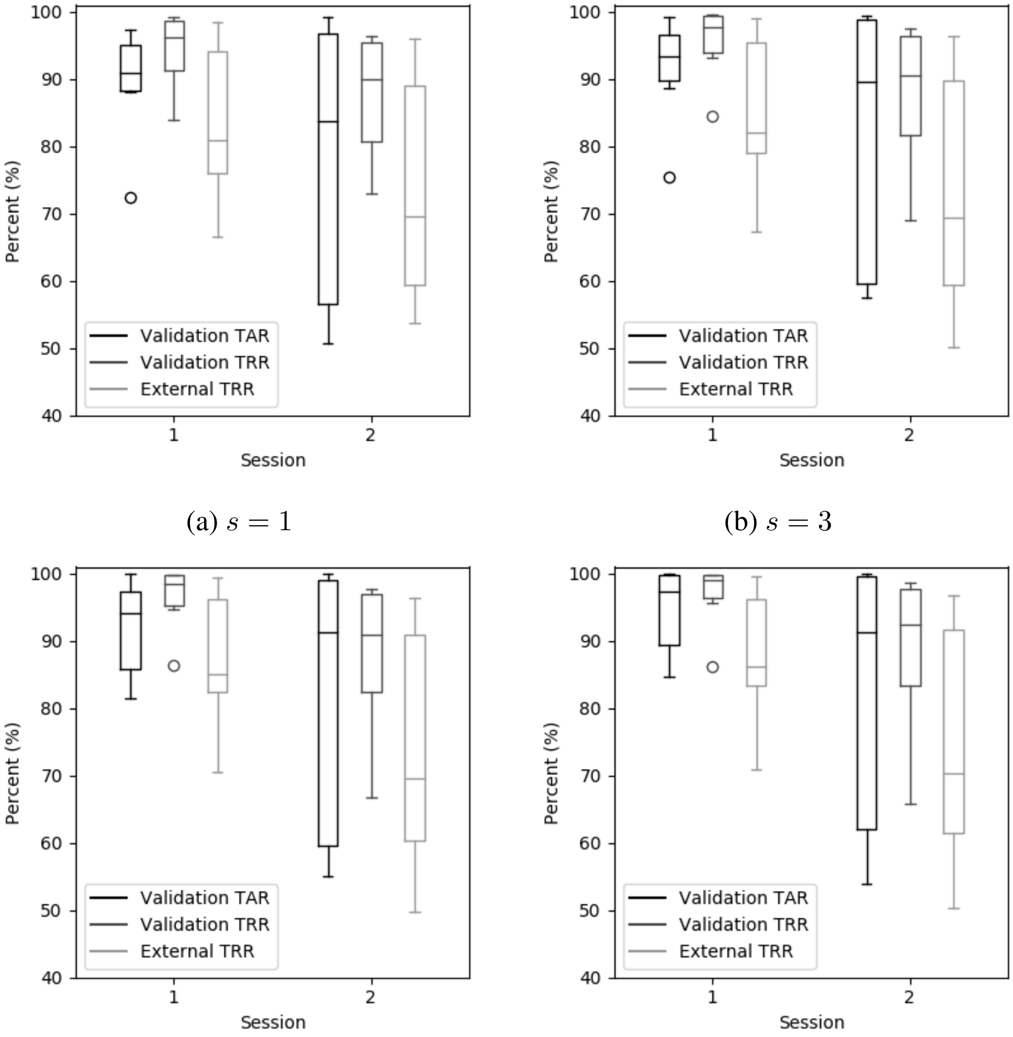

(a) $s = 1$

(b) $s = 3$

(c) $s = 5$

(d) $s = 7$

Figure 11: **In the context of data collected from subjects with severe motor disabilities: Validation TAR, Validation TRR and External TRR for values of $s$ (authentication attempts) between 1 and 7, by using optimal values of $T$ over the three sessions. Note, the y-axis of the graph does not start at the origin.**

# 8 Conclusions and Future Work

In this thesis, we have demonstrated a new authentication solution using BCGs collected from SEW devices. Our approach used ballistocardiogram (BCG) waveforms derived from accelerometer and gyroscope sensor measurements in a Smart Eyewear device to generate BCG waveforms. These waveforms were then used to train a user-specific CNN to authenticate the user. Subsequent samplings of the subject BCGs for the purpose of authentication achieved an EER of $\sim$4% immediately after training, and $\sim$13% after almost 60 days of training. Additionally, we demonstrated that our approach holds promise as a new authentication option for individuals with motor disabilities in increasing their level of independence in securely using modern computing devices. The results demonstrate the potential of our approach. However, there are many avenues through which this work can be improved:

**Compensating for various types of artifacts.** Motion artifacts are a big problem with extracting physiological responses from movement data [11]. Therefore, we need strategies to compensate for noise induced in the signal from the motion of the user. Particularly: (1) when used in locations with heavy foot traffic and (2) when used with the user performing different activities (e.g., standing, walking, etc.).

**Potential implementation issues.** This thesis only describes an analysis of our authentication approach, not its implementation. Ideally, we want the authentication decision to be made locally on the SEW device, in order to minimize any security risk with off-loading the authentication decision. Implementing our approach on a typical head-mounted wearable devices needs to be explored further.

**Effects of fatigue:** In our current dataset, the subjects were alert during all data collection sessions. However, it has been shown that factors such as fatigue effect an individual's movement, posture [19]. At this time it is not clear how well our approach works when authenticating individual is fatigued or under the weather in general.

**Model re-training schedule.** As physiological responses from a subjects change over time, we see that our authentication accuracy drops in sessions 2 and 3. Adjusting the decision threshold will help with reducing the authentication error to only some extent. At some point we have to re-train the models to capture the current physiology of the individual. Experiments are required to determine when to re-train the authentication models such that the overall drop in the authentication accuracy of the models is balanced with the inconvenience taking the system offline causes.

**Larger motor disabilities population.** As discussed prior, the small population of subjects with motor disabilities leads to a lower amount of confidence with regard to the accuracy of the results. Studying the use of ballistocardiogram for a larger population of people with motor disabilities could confirm the results demonstrated here, paving the way for a potentially usable system for such a population.

# References

[1] Epson Moverio. https://epson.com/moverio-augmented-reality.

[2] Google Glass. http://www.google.com/glass/start/.

[3] Microsoft Hololens. https://www.microsoft.com/microsoft-hololens/en-us.

[4] Smart Cap. http://www.smartcaptech.com/.

[5] TheFreeDictionary's Medical dictionary. http://medical-dictionary.thefreedictionary.com/.

[6] Vuzix. https://www.vuzix.com/.

[7] P. Baldi and Y. Chauvin. Neural networks for fingerprint recognition. *Neural Computation*, 5(3):402–418, 1993.

[8] G. Dorffner. Neural networks for time series processing. *Neural Network World*, 6:464–466, 1996.

[9] A. Graves, A. r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6645–6649, May 2013.

[10] H. Guo, X. Cao, J. Wu, and J. Tang. *Ballistocardiogram-based person identification using correlation analysis*, pages 570–573. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[11] J. Hernandez, Y. Li, J. M. Rehg, and R. W. Picard. Bioglass: Physiological parameter estimation using a head-mounted wearable device. In *Wireless Mobile Communication and Healthcare (Mobihealth), 2014 EAI 4th International Conference on*, pages 55–58, 2014.

[12] J. Hernandez, D. J. McDuff, and R. W. Picard. Bioinsights: Extracting personal data from still wearable motion sensors. In *2015 IEEE 12th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*, pages 1–6, 2015.

[13] Y. S. John Duchi, Elad Hazan. Adaptive subgradient methods for online learning and stochastic optimization. In *Journal of Machine Learning Research 12*, pages 2121–2159, 2011.

[14] J. M. Joyce. *Kullback-Leibler Divergence*, pages 720–722. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.

[16] S. Li, A. Ashok, Y. Zhang, C. Xu, J. Lindqvist, and M. Gruteser. Whose move is it anyway? authenticating smart wearable devices using unique head movement patterns. In *2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 1–9, 2016.

[17] A. Mosenia, S. Sur-Kolay, A. Raghunathan, and N. K. Jha. Caba: Continuous authentication based on bioaura. *IEEE Transactions on Computers*, 66(5):759–772, May 2017.

[18] C. E. Rogers, A. W. Witt, A. D. Solomon, and K. K. Venkatasubramanian. An approach for user identification for head-mounted displays. In *Proceedings of the The 19th International Symposium on Wearable Computers*, ISWC'15, 2015.

[19] S. S and T. D. Early detection of silent ischaemic heart disease by 24-hour electrocardiographic monitoring of active subjects. *British Heart Journal*, 36(5):481–486, 1974.

[20] Sarah Perez. Misfit Debuts Spector, Smart Headphones That Double As A Fitness Tracker. https://techcrunch.com/2016/01/06/misfit-debuts-spector-smart-headphones-that-double-as-a-fitness-tracker/, 2016.

[21] S. Schneegass, Y. Oualil, and A. Bulling. Skullconduct: Biometric user identification on eyewear computers using bone conduction through the skull. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 1379–1384, 2016.

[22] S. Singh, A. Cabraal, C. Demosthenous, G. Astbrink, and M. Furlong. Password sharing: Implications for security design based on social practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 895–904, 2007.

[23] E. Vural, S. Simske, and S. Schuckers. Verification of individuals from accelerometer measures of cardiac chest movements. In *2013 International Conference of the BIOSIG Special Interest Group (BIOSIG)*, pages 1–8, 2013.

[24] J. Wei. How wearables intersect with the cloud and the internet of things : Considerations for the developers of wearables. 3:53–56, 07 2014.

[25] Q. Zhang, D. Zhou, and X. Zeng. Heartid: A multiresolution convolutional neural network for ecg-based biometric human identification in smart health applications. *IEEE Access*, 5:11805–11816, may 2017.