

A Wearable System That Knows Who Wears It

Cory Cornelius
Security and Privacy
Research
Intel Labs
Hillsboro, OR

Ronald Peterson
Institute for Security,
Technology, and Society &
Dept. of Computer Science
Dartmouth College
Hanover, NH

Joseph Skinner
Thayer School of Engineering
Dartmouth College
Hanover, NH

Ryan Halter
Thayer School of Engineering
& Geisel School of Medicine
Dartmouth College
Hanover, NH

David Kotz
Institute for Security,
Technology, and Society &
Dept. of Computer Science
Dartmouth College
Hanover, NH

ABSTRACT

Body-area networks of pervasive wearable devices are increasingly used for health monitoring, personal assistance, entertainment, and home automation. In an ideal world, a user would simply wear their desired set of devices with no configuration necessary: the devices would discover each other, recognize that they are on the same person, construct a secure communications channel, and recognize the user to which they are attached. In this paper we address a portion of this vision by offering a wearable system that unobtrusively recognizes the person wearing it. Because it can recognize the user, our system can properly label sensor data or personalize interactions.

Our recognition method uses *bioimpedance*, a measurement of how tissue responds when exposed to an electrical current. By collecting bioimpedance samples using a small wearable device we designed, our system can determine that (a) the wearer is indeed the expected person and (b) the device is physically *on* the wearer's body. Our recognition method works with 98% balanced-accuracy under a cross-validation of a day's worth of bioimpedance samples from a cohort of 8 volunteer subjects. We also demonstrate that our system continues to recognize a subset of these subjects even several months later. Finally, we measure the energy requirements of our system as implemented on a Nexus S smart phone and custom-designed module for the Shimmer sensing platform.

1. MOTIVATION

We are entering a period of rapid expansion of wearable and pervasive computing due to the continuing advances in low-power electronics, including sensors and actuators. Today, it is not uncommon for people to carry multiple computing devices, such as smart phones, music players, and cameras. Increasingly, we carry, hold, or wear devices to measure our physical activity (e.g., Fitbit), to interact with our entertainment devices (e.g., the Xbox One), or to monitor our physiology (e.g., a cardiac patient concerned about

heart arrhythmia or a diabetic managing her blood glucose). These unobtrusive, wearable devices make it possible to continuously or periodically track many health- and lifestyle-related conditions at an unprecedented level of detail. Their use of wireless connectivity enables interaction with other devices nearby (e.g., entertainment systems, climate-control systems, or medical devices), and allows sensor data they collect to be automatically shared with a social-networking service, or (in the case of health applications) uploaded to an Electronic Medical Record system for review by a healthcare provider.

In this paper, we focus on a fundamental problem involving wearable devices: who is wearing the device? The ability to recognize who is interacting with a device is essential for many applications. For an entertainment device, it can recognize the user and load the correct profile. For a home climate control, it can adjust the environment to the wearer's preference. Most compellingly, for a health-monitoring device, it can label the sensor data with the correct identity so that it can be stored in the correct health record. (A mix-up of sensor data could lead to incorrect treatment or diagnosis, with serious harm to the patient.)

In our vision, a person should be able to simply attach the desired set of devices to their body – whether clipped on, strapped on, stuck on, slipped into a pocket, or even implanted or ingested, and have the devices *just work*. That is, without any other action on the part of the user, the devices would discover each other's presence, recognize that they are on the same body (as opposed to devices in radio range but attached to a different body nearby), develop shared secrets from which to derive encryption keys, and establish reliable and secure communications. Furthermore, for many of the interesting applications described above, the devices must also recognize *who* is wearing them.

We have previously developed a method for a networked set of devices to recognize that they are located on the same body; our approach uses correlations in acceleration signals for this purpose [9]. If even one device can recognize *which* body, then transitively the set of devices know who is wearing them. Indeed, it is unlikely that every device will have the technology, or suitable placement, to recognize the user; in our model, only one such device needs that capability.

One easy solution, common in many devices today, is for the device to be statically associated with a given user. This smart phone is *my* phone, whereas that fitness sensor is *your* fitness sensor. The device is assumed to be used by only that user and any data generated

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiSys'14, June 16–19, 2014, Bretton Woods, New Hampshire, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2793-0/14/06 ...\$15.00.

<http://dx.doi.org/10.1145/2594368.2594369>.



Figure 1: A subject wearing one of our bioimpedance sensors. The exposed wires connect the electrodes to our custom sensor module (which is connected to the Shimmer via the internal expansion port).

by a sensor is associated with that user. There are, however, many situations where this static association fails. In some households, a given device might be shared by many users (e.g., a blood-pressure cuff). In other settings, two people might accidentally wear the wrong sensor (e.g., a couple who go out for a run and accidentally wear the other’s fitness sensor). For some applications, a person may actively try to fool the system (e.g., a smoker who places his “smoking” sensor on a non-smoking friend in order to receive incentives for smoking cessation).

What we require is an unobtrusive, wearable device that leverages biometrics to recognize the wearer. It would then share that identity with the body-area network of other devices (earlier confirmed to be on the same body [9]). This device would be trained once for each user that might wear it, but thenceforth be completely automatic and unobtrusive.

Our approach is to use *bioimpedance*, which measures how a living tissue responds when exposed to an electrical current. Our wearable device, shown in Figure 1, places eight electrodes in contact with the wearer’s wrist. Using two of these electrodes, the device applies a small, harmless current to the wrist and measures bioimpedance. By comparing these bioimpedance samples with a model of a subject that was built earlier in a training phase, the device can both determine that some person is wearing it and verify the particular subject who is wearing it. If we train the device for a set of users, e.g., the members of a household, then the device can recognize which of those household members is wearing the device, or that none of them are wearing the device. We previously demonstrated the feasibility of this approach [10] on a set of nearly 50 subjects; although that study validated our method’s ability to distinguish subjects, all of our data was captured in the lab in a controlled setting using an *unwearable* device.

In this paper we demonstrate the feasibility of our approach by deploying a custom-designed wearable device outside of the lab to collect data over the course of a day for each subject. As in our prior work, we target a household-size cohort of subjects because we believe these kinds of wearable devices will typically be used by a few subjects over their useful lifetime. We also provide evidence that bioimpedance samples remain recognizable even several months later.

Our solution has many advantages. Not all wearable devices require the ability to recognize the user; only one device need do so, assuming it can communicate the identity to other devices proven to be on the same body. The devices may be smaller and simpler than a device like a smart phone since they need no interface for user recognition (or PIN or password for authentication). Use of a biometric

provides important security and privacy properties, preventing unauthorized users from either accessing sensitive data (e.g., in which an adversary Alice tricks Bob’s sensor into divulging his activity data to her smart phone), and preventing the mis-labeling of sensor data that might later be used for medically important decisions. Privacy is particularly important in health-related pervasive applications [5]. Furthermore, these methods can support personalization techniques so often envisioned in pervasive computing.

Contributions. First, we devised a wearable sensor system capable of sensing bioimpedance at the wrist. Second, we extend our previous method to include a different set of features that substantially increases recognition performance. Third, we deployed our system outside of the lab to 8 individuals who wore our system for a day. Finally, we demonstrate the feasibility of our device in its ability to recognize the wearer over time and to distinguish among several wearers in identification and verification settings under a naive imposter attack model.

2. WEARABLE SENSOR SYSTEM

We designed our wearable device to be a watch-like bracelet that contains small electrodes to measure bioimpedance. The form factor of a watch has several technical advantages. First, it is worn the same way each time, more or less; issues with placement of the electrodes are diminished because it can sense data from nearly the same location each time and in the same orientation. Second, a watch can be instrumented to detect when it has been placed on and taken off a person. Attachment can be detected, for example, by the ends of the watchband being clasped together or by detecting properties of the skin such as temperature or moisture. Because we require the electrodes to be in contact with the body and not all form factors will afford continuous contact, a mechanism to detect when the device is in contact with a body is necessary. Such simple detection mechanisms also allow us to conserve energy by only performing recognition when the device is actually in contact with a person. We describe such a method in Section 2.4.

2.1 Bioimpedance

Bioimpedance is a physiological property related to a tissue’s resistance to electrical current flow and its ability to store electrical charge. In *in vivo* human applications, it is typically measured through metallic electrodes placed on the skin and around an anatomic location of interest (e.g., the wrist). These electrical properties are predominantly a function of the underlying tissue, including the specific tissue types present (e.g., blood, adipose, muscle, bone), the anatomic configuration (i.e., bone or muscle orientation and quantity), and the state of the tissue (normal or osteoporotic bone, edematous versus normally hydrated tissue, and so forth). Significant impedance differences exist between the varying tissue types, anatomic configurations, and tissue states, each of which may provide a unique mechanism for distinguishing among people. Figure 2 shows the anatomy of a human wrist with several electrodes and example paths of where current can travel. For more details about bioimpedance, see Cornelius [8, Section 3.3].

2.2 Hardware

To capture a bioimpedance sample, we designed and manufactured a sensor designed to be worn on a person’s wrist. The wearable sensor was built on top of the Shimmer Platform [26] – an open-source, low-power wireless sensing platform – and uses a bioimpedance sensor module we designed. The Shimmer provides processing (via a MSP430 microcontroller), wireless communication (via Bluetooth or 802.15.4), storage capabilities (via SD card), and measurement of movement (via an accelerometer). It also pro-




Figure 2: The anatomy of a human wrist [13]. The bones (ulna and radius) are shown in gray, muscle in red, skin in brown. The gray boxes on the outside of the anatomy represent the placement of electrodes with corresponding labels. If electrode 0 were chosen as the current-applying electrode, the blue dashed line represent possible paths from it to electrode 1, 3, or 4. Notice how, depending upon the selected electrodes, the path may pass through a variety of tissues.

vides internal and external expansion connectors that allow it to interact with custom sensor modules (e.g., a gyroscope, magnetometer, ECG, EMG, or GSR). Our custom bioimpedance sensor module (Figure 3) used the internal expansion connector, and was enclosed by a custom-designed case. The bioimpedance sensor module included a receptacle that enables a series of electrodes to be connected to it. We designed and manufactured an elastic sleeve with 8 evenly spaced electrodes. The sleeve connected to the bioimpedance sensor module and included a pocket that holds the Shimmer. Both the sleeve and sensor module went through several iterations of design to ensure it would be comfortable for subjects to wear and hold up to daily use. Figures 1 and 4 show the final form factor.

We designed the bioimpedance sensor module (Figure 3) around the Analog Devices AD5933 Impedance Analyzer [3]. The AD5933 includes a frequency generator that allows excitation at a specified frequency between 1 kHz and 100 kHz with a resolution of 0.1 Hz and can measure impedances between 1 k Ω and 100 M Ω to within 0.5% total system accuracy. The Shimmer controlled the AD5933 via the I²C bus. Because we wanted to allow multiple electrode locations, the sensor module included two Analog Devices ADG1608 8-Channel Multiplexors [4]. Thus, the sensor module was capable of selecting 2 of 8 possible electrodes, called an *electrode pair*, for bi-polar sensing. These multiplexors were controlled by setting specific GPIO pins. The electrodes were connected to the sensor module via an 8-pin Hirose 3260-8S3(55) Connector [16] such that custom electrode configurations could be independently built and interfaced with the sensor module. The Shimmer provided regulated power to the sensor module, which was fed to an Analog Devices ADR433 ultra-low noise voltage reference. This ADR433 [2] provided a stable 3V supply voltage needed for the impedance analyzer while a Microchip MCP1252 charge pump [20] fed a 5V supply voltage to the multiplexors.

Internally, the AD5933 computes the Discrete Fourier Transform of 1024 analog-to-digital samples at each frequency. This computation yields the power of the signal in the form of real r and imaginary




Figure 3: Our custom designed sensor module with the major components labeled. It was approximately 45mm long by 19mm wide and fitted comfortably on top of a Shimmer with a custom-designed enclosure.

i components. We computed the magnitude of the impedance sample at frequency ω as $|Z(\omega)| = \sqrt{r_\omega^2 + i_\omega^2}$. We computed the phase of the impedance sample as $Z\emptyset(\omega) = \tan^{-1}\left(\frac{i_\omega}{r_\omega}\right)$, taking care to ensure $\frac{i_\omega}{r_\omega}$ was a positive quantity by rotating this angle until it fell into the appropriate quadrant in the complex plane. Given the magnitude and phase, we computed the resistive and reactive components of impedance, represented by $R(\omega)$ and $X(\omega)$ respectively, by projecting onto the Cartesian plane:

$$\begin{aligned} Z(\omega) &= R(\omega) + jX(\omega) \\ &= |Z| \cos(Z\emptyset(\omega)) + |Z| \sin(Z\emptyset(\omega)) \end{aligned}$$

where j is the imaginary number. We calibrated of these values to a known reference impedance as described in Appendix A.

2.3 Software

The Shimmer ran the TinyOS operating system [28]. We wrote custom software to communicate with the impedance analyzer and multiplexors. The software was divided into three major parts: a low-level driver, a high-level driver, and a logging application.

The low-level and high-level drivers allowed applications to communicate with our sensor module. The low-level driver was a bare-bones interface to the AD5933 that wrapped the I²C communications. The high-level driver implemented a state machine that allows more natural interaction with the sensor module. It also allowed applications to adjust the settling time (i.e., the amount of time between the stimulus and sampling) and handled failures gracefully. In total, the low-level driver compromised 381 lines of nesC code while the high-level driver compromised 363 lines of nesC code.

The logging application used the high-level driver to interact with the sensor module. Before taking bioimpedance samples, the application collected acceleration for 5 seconds at 50 Hz to classify the type of motion (i.e., low energy, medium energy, high energy) the subject’s wrist was experiencing. Next, two electrodes on the wrist strap were selected using the multiplexors. The impedance analyzer was then commanded to measure at 50 logarithmically spaced frequencies from 1 kHz to 100 kHz. Because of the logarithmic spacing of the frequencies, the impedance analyzer’s internal mechanism to step through frequencies linearly could not be used, so we




Figure 4: Our wearable bioimpedance sensor. The sleeve (top) is shown inside out to display the electrodes. The Shimmer (bottom) and wires were typically housed within the sleeve, but are shown exposed here. For reference, the Shimmer was approximately 53 mm long by 32 mm wide.

commanded each measurement individually. After one frequency sweep, we selected another set of stimulus and measurement electrodes and another frequency sweep was started. After a complete sample, we commanded the impedance analyzer to conserve power by sleeping. A timer was used to wake the impedance analyzer after a preset interval and start another sample. Sample data could either be stored to a microSD card or sent via the Bluetooth or 802.15.4 radio depending upon the application. In total, the logging application compromised 1570 lines of nesC code.

We also developed a smart phone application that allowed collection of bioimpedance samples via Bluetooth. A Nexus S smart phone [12] communicated with the Shimmer via Bluetooth using the Serial Port Profile. To ensure the confidentiality and integrity of the communications, we relied upon the underlying authentication and confidentiality mechanisms present in Bluetooth. The smart phone application was a 932 line Java Android 4.1 application. We used the OpenUAT toolkit [19] for some of the signal processing and the WEKA toolkit [14] for the machine-learning aspects of the application. We used these toolkits to implement the algorithms described in Section 3 below. We trained each user’s model offline using WEKA.

2.4 Detecting presence

Because our device required contact with a person’s wrist, we needed to detect when the device was in contact with a person. Fortunately, skin contact was easy to detect. Figure 5 shows some example bioimpedance samples, one of which shows the impedance sample when there was no contact with a human wrist. When there was no contact, the magnitude of the impedance for that sample was high across all frequencies. Thus, we used a simple threshold to detect presence.

This presence-detection scheme, however, required a lot of energy to execute. High security applications, for example, might require near-instant detection of wristband removal. Given that a full impedance sample took 98 mW to measure, this presence-detection method is not suitable for continuous use. It might be feasible to look at an individual electrode pair and a single frequency, but we




Figure 5: Example impedance samples from electrode pair 04 for 1) no contact, 2) bracelet in contact with a human wrist, and 3) bracelet attached to 4 kΩ resistor. In the case of no contact, impedance is effectively infinite. Note: impedance and frequency are plotted on logarithmic scales.

leave such optimizations for future work; alternate approaches for presence detection are described in Section 5.

3. RECOGNIZING WEARERS

Recall that the goal of our system is to recognize who is wearing it under a naive imposter attack model. The intuition is quite simple: gather bioimpedance samples from a subject and build a model that represents that subject’s bioimpedance samples; later, use that same model to see if a new and unknown bioimpedance sample matches the model. This implies there are two distinct phases: enrollment and recognition. Before we discuss the details of the enrollment and recognition phases, we first discuss the details of the bioimpedance samples and features we use in those phases.

3.1 Bioimpedance samples

It would be infeasible to measure bioimpedance from all pairs of the eight electrodes, so we carefully chose specific electrode pairs. We captured samples from two types of electrode pairs: those electrodes directly across from one another (e.g., electrodes 0 and 4 as shown in Figure 2), since they were the maximal distance away from each other and therefore provide more tissue for the current to travel through; and those electrodes that were exactly one electrode apart (e.g., 0 and 2), since the current would travel through the outer regions of the wrist. We did not measure bioimpedance at electrodes directly next to each other since the skin would be the primary tissue affecting the reading. The skin is subject to variability due to sweat and externally applied fluids (e.g., lotions, hand sanitizers, or topical medicines). Likewise, we did not measure bioimpedance for those electrodes spaced exactly two electrodes apart, because those samples exhibited characteristics similar to those collected from electrodes maximally apart. Figure 6 shows example bioimpedance samples from a single subject for the 12 different electrode pairs we used.

3.2 Feature extraction

Given a set of frequencies and their corresponding bioimpedance measurements, we extracted four features from each bioimpedance sample: two from the bioimpedance magnitude and two from the bioimpedance phase. The features were simple: we fitted one line to

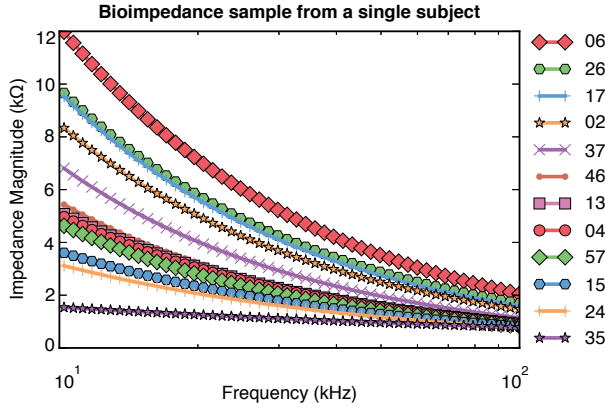


Figure 6: Example bioimpedance samples collected from a single subject for the 12 different electrode pairs after fitting to a line in log-log space. Note: frequency is plotted on a logarithmic scale.

the bioimpedance magnitude and another line to the bioimpedance phase, both in log-log space. The inspiration for this feature can be seen in Figure 5, which shows that a bioimpedance sample of a human wrist is mostly linear in log-log space. Because samples were inherently noisy, fitting a line to a bioimpedance sample smoothed it while preserving its general shape. Fitting a line yields four features because a line is succinctly described by a slope and intercept (a slope and intercept for each of the magnitude and phase components). We thereby reduced the dimensionality of the data from 100 to 4, which in turn lowered the computational and energy requirements. Although we explored other features, including using the raw data itself, we found these features yielded the best recognition performance.

Our final *feature vector* consisted of the concatenation of these features for each electrode pair. Since we took samples from 12 electrode pairs, this resulted in a feature vector of dimension $12 \times 4 = 48$, which was much smaller than the raw data (which is of dimension $12 \times 50 \times 2 = 1200$). This concatenation assumed all electrode pairs would provide some information about the identity of the wearer. (In Section 4 we explore different concatenations of electrode pairs.)

Because many of our studies took place outside of the lab, we had to discard some bioimpedance samples. We discarded those samples where the system failed to measure a complete bioimpedance sample. (In a production system, this could be detected in the moment and the sample could be retaken.) We also discarded all samples where the wrist was deemed to not be in contact with the device. We decided that a sample was not in contact when the maximum bioimpedance magnitude was greater than 10^3 kΩ, a value determined by the maximum non-contact bioimpedance magnitude shown in Figure 5. Finally, we discarded all samples where the sum of squared errors of the fitted line in log-log space was above an empirically determined threshold of 0.5. We discarded these samples because such a poor fit indicated a noisy sample, probably due to motion or other interference with the reading.

3.3 Enrollment

Before a subject could use the system, we trained the system to recognize his/her bioimpedance by putting the system into *enrollment mode*. In this mode the system captured bioimpedance samples for a designated time (12 bioimpedance samples took about 15 seconds). The system then used these training samples as inputs

to an *enrollment algorithm* that learns a model of the enrollee’s bioimpedance. (It was necessary to compute this model off the wearable device because of resource constraints.) The model was then loaded into the system for use.

Given a set of *training feature vectors* from a subject, we learned the model of their bioimpedance samples using the enrollment algorithm. There were two modes of operation for these algorithms: *identification* and *verification*. Identification is a many-to-one matching, while verification is a one-to-one matching. Identification is used to determine which person from a population is wearing the device. Verification, on the other hand, is used to confirm that a chosen person is wearing the device. In our experiments, our target population is the size of a household since we believe that over the lifetime of a typical wearable device it will only encounter a few people.

3.3.1 Identification

In identification mode, we used a discriminative algorithm to learn a model of each subject’s bioimpedance. Our method learns a classifier for each subject by using that subject’s feature vectors as positive examples (i.e., they are labeled positively) and all other subjects’ feature vectors as negative examples (i.e., they are labeled negatively). This is the one-versus-all strategy used for multi-class classification [24]. In this mode, each new subject requires retraining the models; however, we only need to collect samples from the new subject because we can reuse the already enrolled subjects’ training data.

We examined two algorithms for use during identification-mode enrollment. The first classifier, Naive Bayes (NB), independently models the mean and variance of each feature. A NB classifier is a relatively simple classifier to train because all it requires is computing the mean and variance of each feature for each label. The second classifier, a Support Vector Machine (SVM), finds the hyperplane that best separates the positive examples from the negative samples (i.e., the maximum-margin hyperplane). The intuition is that only those training examples near the hyperplane (i.e., the “support vectors”) are necessary to describe it and thus constitute the model. Even if the data is not linearly separable, one can use the so-called “kernel trick” to map the examples into a higher dimensional space where they might become linearly separable [6].

The SVM classifiers require us to choose some parameters. (The NB classifier has no such parameters.) To choose parameters, we ran a 10-fold cross-validation of a small subset of our dataset over the parameter space of the SVM classifier. In the SVM case, we looked at different kernels (linear, polynomial, and radial-basis function), soft-margin costs, and, in the case of polynomial and radial-basis function kernels, their respective kernel coefficient gamma [6]. Using a grid search, we found that a 3rd degree polynomial kernel with a cost of 32 and gamma of 0.03125 was optimal.

3.3.2 Verification

In verification mode, we used a generative algorithm to learn a model of each subject’s bioimpedance. This mode naturally supports multiple subjects because each subject’s model can be learned independently of any other subject. Thus, a new subject could simply be loaded into the system without regard to the other subjects already in the system.

We examined one algorithm for use during verification-mode enrollment. A *Gaussian Mixture Model* (GMM) modeled the bioimpedance samples using a weighted linear combination of Gaussian densities, where each Gaussian density was parameterized by a mean vector and covariance matrix. We chose initial Gaussian densities by clustering the set of feature vectors using *k*-means

clustering [17], where k was set to the desired number of Gaussian densities. We then iteratively refined these initial Gaussian densities using the expectation-maximization algorithm [11] until the maximum likelihood remained stable (i.e., the difference between successive iterations was less than 0.01) or after a maximum number of iterations (100). We modeled the full covariance since the dimensionality of a bioimpedance sample was relatively low. Because some values of the covariance matrix could become very small, as in the case of outliers, we enforced a variance floor of 0.001 on the covariance matrix. For our experiments, we found that 4 Gaussian densities best modeled a subject’s bioimpedance samples.

3.4 Recognition

Once a subject was enrolled, the system entered *recognition mode*. In recognition mode, the system periodically determined whether it was on a human body (Section 2.4), then collected bioimpedance samples. The system used a *recognition algorithm* to determine whether a bioimpedance sample matches the enrollee’s model. Like the enrollment algorithm, this recognition algorithm operated in one of two modes: identification or verification.

3.4.1 Identification

In identification mode, we used the enrolled models to choose which subject best matched a test feature vector from an unknown subject. A feature vector that was classified as positive for a given subject’s model was said to match that subject’s bioimpedance; otherwise, the test feature vector was classified as negative because it did not match that subject’s bioimpedance. Choosing the best match is left to each classifier.

Each classifier has a different mechanism for classifying test feature vectors. The NB classifier chooses the label (i.e., is subject or not) of the test feature vector with the maximum likelihood as the classification. The best match is chosen as the sample with the highest likelihood. That is, it computes the Gaussian probability density function for the test feature vector given each label, and the label with the maximum value is the classification. Since an SVM is a linear classifier, it simply computes the linear transformation of the test feature vector and returns the sign of that value as the predicted label. The best match is chosen by the classifier with the largest margin between the support vectors and the sample.

3.4.2 Verification

In verification mode, we used the enrolled models to decide whether a test feature vector came from a known subject. We did so by asking the generative model to tell us how likely the test feature vector matched the model. Thus, we queried the subject’s GMM for the likelihood of the new test feature vector given the model. Given some threshold τ , we accepted those bioimpedance samples came from that subject if the likelihood was greater than τ , and rejected otherwise. Because there is no good way of choosing the threshold *a priori*, we varied this threshold τ to determine how well our method performed.

3.5 Metrics

Consider a set of test feature vectors from a given subject and a set of test feature vectors from other subjects. We labeled the test vectors measured from the given subject as *positive* and all other test vectors as *negative*. We then used the model trained for that subject to classify all the test feature vectors, resulting in a positive or negative classification for each. Thus for each subject, every other subject acted as a naive imposter against that subject’s model. Ideally, the model would classify only those test feature vectors from that subject as positive and all other test feature vectors from

other subjects as negative. Given such classification results, we present our results using the following metrics. The *false accept rate* (FAR) is the fraction of negatively labeled feature vectors that were misclassified (i.e., they were classified as positive). A perfect classifier would perform at 0% FAR. The *false reject rate* (FRR) is the fraction of positively labeled feature vectors that were misclassified (i.e., they were classified as negative). A perfect classifier would perform at 0% FRR. The *balanced accuracy* (BAC) is the sum of half of the true accept rate (i.e., the fraction of positively labeled feature vectors that were correctly classified, or $1 - \text{FRR}$) and half of the true reject rate (the fraction of negatively labeled feature vectors that were correctly classified, or $1 - \text{FAR}$). This metric weighs the negative and positive examples equally, which is necessary because there were more negatively labeled feature vectors than positively labeled feature vectors (because each subject’s model is tested against everyone else’s). A perfect classifier would perform at 100% BAC. Finally, the *equal error rate* (EER) is the rate at which the FAR equals the FRR. A perfect classifier would perform at 0% EER.

Although we computed these metrics for every subject in our data sets, we present summary statistics of these metrics over all subjects. In what follows, then, any mentions of BAC, FAR, and FRR should be interpreted as the average BAC, average FAR, and average FRR over all subjects. Where possible, we also include the standard deviation of each metric. Note that because the number of positive samples for any given subject was smaller than the number of negative examples by a factor of $N - 1$, where N is the number of subjects, a classifier that always predicts the negative case will perform at a FAR of 0%, BAC of 50%, but FRR of 100%. For comparison’s sake, we also computed these metrics under a classifier that predicted randomly based on the frequency of labels present in the training set (i.e., it wholly ignored any feature vectors associated with the labels). This classifier served as a baseline performance measure to compare with our method.

4. EXPERIMENTAL RESULTS

Using the above recognition method, we conducted a series of experiments to evaluate the feasibility of our system. In our first experiment, we replicated our prior work to ensure our device was working properly. Our second and third experiments validate the feasibility of our system when worn in an out-of-lab setting for both identification and verification. Finally, we measured the energy requirements of an implementation of our system.

Recall that our goal is to recognize the wearer in a small cohort of subjects using a custom-designed wearable bioimpedance sensor. Our prior work showed promise that bioimpedance is sufficiently unique among individuals [10]. In this study we sought to validate the use of a wearable bioimpedance sensor by deploying our device to subjects in an out-of-lab study. Such an evaluation will necessarily have some variance in bioimpedance samples and provides some measure of our system’s performance in the real world. Furthermore, in this study we evaluate recognition in verification mode as well as identification mode.

4.1 Dataset

We collected bioimpedance samples from 8 people over a period of one day each. We informed participants of the risks involved in wearing the device (e.g., localized skin irritation at electrode site). If they agreed to enroll, we asked each subject to self-report their age and gender, and we measured the circumference of their wrist using a Health-O-Meter Digital Tape Measure at the location shown in Figure 7. We asked each subject to wear the sensor on their non-dominant wrist for as long as possible during the day and to return the device at the end of the day. We informed subjects that



Figure 7: The Health-O-Meter Digital Tape measure we used to measure each subject’s wrist circumference. It was placed just above the ulnar styloid process of the subject’s non-dominant wrist where the device was worn.

they could remove the device at their leisure and should certainly remove it if they believed the device would come in contact with water (e.g., before showering or swimming). Should they remove the device, we instructed subjects to put the device back on in the same orientation as it was previously worn. Subjects were paid \$8 for their participation and both our device and data-collection protocols were approved by our Institutional Review Board.

The enrolled participants (Table 1), 3 female and 5 male, had an average of 27 ± 8 years. The average wrist circumference was 18.0 ± 1.0 cm, meaning that the group of subjects selected for this study each had a similar wrist circumference. One could use the wrist circumference as a distinguishing factor, but such a device would need to accurately and precisely measure the wearer’s wrist circumference. Even with a moderate measurement error of 5 mm, distinguishing the individuals in our dataset would be difficult. On average, the device was in contact with a subject’s wrist for 9.3 hours resulting in 112 bioimpedance samples. Note that this group of 8 subjects was larger than the groups we examined in our prior work; we previously examined cohorts of 5 subjects representing a household or similar group that may share devices.

4.2 Replication

We first validated how well our methods performed under conditions similar to our prior work [10]. We ran a 10-fold cross-validation for each subject using only the electrode pairs that were maximally distant (i.e., 04, 15, 26, and 37) as in our prior work. Figure 8 shows the results of this validation for the best-performing combinations of the specified electrode pairs using the NB classifier and SVM classifier. This validation extended our prior work by also examining the use of an SVM classifier. The best performing single-electrode pair in this study (04) confirmed our prior finding that the best performing single-electrode pair was the pair with electrodes on the top and bottom of the wrist. In the SVM case, the best double-electrode pair combination was 04, 37. Due to the design of our new device, pair 04 did not sit exactly at the top of the wrist (as in our prior experiments), but rather it shared the top of the wrist with electrode pair 37. Overall, the SVM classifier achieved the best performance when all four electrode pairs were combined. It achieved a BAC of $95.0\% \pm 3.00\%$, a FAR of $1.40\% \pm 0.61\%$, and a FRR of $8.62\% \pm 5.74\%$. In contrast, the NB classifier (which was used in our prior study) achieved a BAC of $86.4\% \pm 3.82\%$,

Table 1: The total number of samples taken, the number of samples that were deemed not to be in contact with the subject’s wrist, the total time in hours the device was in contact with the subject’s wrist, and the subject’s wrist circumference in centimeters. We measured each subject’s wrist circumference because we believe it can be a useful feature in distinguishing subjects (although we did not use it in these studies). See Cornelius [8, Section 3.6.3 & 3.7.3] for how wrist circumference could be integrated as a feature.

Sub.	Samples	No Cont.	Cont. (h)	Wrist (cm)
1	111	6	8.8	18.2
2	96	8	7.3	16.2
3	107	22	7.1	18.4
4	262	27	19.6	17.0
5	134	12	10.2	19.4
6	123	38	7.1	18.4
7	134	59	6.2	17.6
8	111	13	8.2	19.0
Avg	135 ± 50	23 ± 17	9.3 ± 4.0	18.0 ± 1.0



a FAR of $15.4\% \pm 6.35\%$, and a FRR of $11.7\% \pm 3.45\%$, similar findings to our prior experiments. A random classifier achieved $50.2\% \pm 0.87\%$, a FAR of $13.0\% \pm 5.83$, and a FRR of $86.7\% \pm 7.12$. The SVM classifier outperformed our prior work while the NB classifier performed similarly.

4.3 Identification

To examine how well our system performed in the identification setting, we ran a cross-validation over the entire dataset. Recall that identification requires our algorithms to choose which subject is present in a cohort of candidate subjects. This validation assumes that we have nearly an entire day’s worth of training samples available, since we train on 90% of the data, leaving 10% for testing. This assumption means the training samples will necessarily capture the variability in a subject’s bioimpedance across different environments, motions, and orientations. For each subject we ran a 10-fold cross-validation over the set of feature vectors for each classifier defined in Section 3. That is, we trained our model using 90% of the samples, leaving 10% of the samples to be classified. We computed the FAR, FRR, and BAC for each subject, and we report the average and standard deviation of these measures over all subjects.

Figure 9 shows these results for the top-performing electrode-pair combinations using the NB and SVM classifiers respectively. The top-performing electrode-pair combination (02 04 06 13 15 17 26 37) achieved a BAC of $98.1\% \pm 0.98\%$, FAR of $0.96\% \pm 0.51\%$, and FRR of $2.83\% \pm 1.17\%$ for the SVM classifier. The best the electrode-pair combination (04 06 35 37 46 57) for the NB classifier achieved was a BAC of $87.9\% \pm 3.83\%$, FAR of $14.6\% \pm 8.14\%$, and FRR of $9.62\% \pm 3.44\%$. The SVM classifier benefited from more electrode pairs, while the performance of the NB classifier did not significantly increase except for reductions in FAR and FRR.


Figure 10 shows a visualization of these top-performing electrode-pair combinations for each classifier. In most of the top combinations for the SVM classifier, electrode pairs 04, 37, 02 were present. The electrode locations (see Figure 2) of the top-performing electrode-pair combination correspond to the ulnar side of wrist (or medial side of the forearm) for the SVM classifier, while the top-performing electrode-pair combinations for the NB classifier tend to encompass the whole wrist. In general, as the number of elec-




(a)

(b)

Figure 8: A replication of our prior work using a NB classifier over our dataset restricted to only those electrode pairs maximally distant. For a single electrode pair, 04 performed best, while the full set of electrode pairs does not perform significantly better. The SVM classifier, on the other hand, performed significantly better than the NB classifier when all four electrode pairs were included.




(a)



(b)

Figure 9: The top-performing electrode-pair combinations of a 10-fold cross-validation as classified by a NB classifier. The performance of the NB classifier was flat while the SVM classifier benefits from more electrode-pair combinations.



(a) NB classifier

(b) SVM classifier

Figure 10: Visualization of best performing combinations of electrode-pairs in cross validation for the NB classifier and SVM classifier. Notice how the top-performing electrode-pair combination for the NB classifier encompassed much of the wrist, while the top-performing electrode-pair combination for the SVM classifier encompassed the medial side of the wrist.




Figure 11: The recognition rates for each subject in a verification setting. The average per-subject EER was $13.10\% \pm 8.67\%$. A global threshold achieved an EER of 12.7% . The threshold, however, can be changed to suit the needs of the application (i.e., more false-negatives or more false-positives).

trode pairs increased, our device sensed more of the geometry of the wrist. This result shows that using electrode pairs that are not just maximally distant (as was done in our prior work) will benefit recognition performance.

4.4 Verification

In this experiment, we sought to understand how well our wearable device performs in a verification setting. Recall that verification is the process of verifying whether the wristband’s wearer is its owner, that is, the individual this wristband was trained to recognize. To validate the performance of our system, we used a hold-out validation where we held out the first 90% of a subject’s data and left the remaining 10% for testing. We computed the FAR, FRR, and EER for each subject, and we report the average and standard deviation of these measures over all subjects.

Figure 11 shows how well our method performs for a 90% hold-out validation. Recall that we can vary the threshold τ for each subject and compute the corresponding FAR and FRR for each threshold. At some threshold τ the FAR will equal the FRR, which the legend in Figure 11 also shows as the EER. The EER varied across subjects, but the average per-subject EER was $13.10\% \pm 8.67\%$. Rather than varying the threshold τ for each subject, we can also vary a global threshold τ over all subjects. That is, rather than computing subject-specific thresholds, we can also compute a threshold that works for any subject. The EER for such a global threshold was 12.7% . In verification mode, one can easily change the threshold to suit the needs of the application to account for fewer false-positives or false-negatives, as Figure 11 shows. Verification mode, however, tends perform worse than identification mode because identification mode incorporates information about other subjects.

4.5 Longitudinal verification

To understand the longitudinal recognition rates of bioimpedance, we collected 10 additional bioimpedance samples from three subjects (1, 4, and 5) 140 days after their initial enrollment in our wearable study. We ran a hold-out validation where the testing dataset was equal to these new bioimpedance samples for Subjects

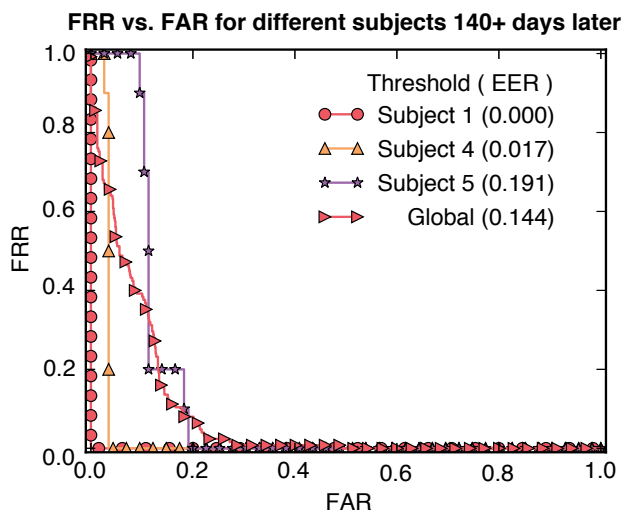


Figure 12: The longitudinal recognition rates of bioimpedance. We collected data from 3 subjects 140 days after their initial enrollment in our wearable study. The per-subject threshold EER was $6.93\% \pm 10.60\%$ while the global threshold EER was 14.4% , indicating results similar to our wearable verification evaluation.

1, 4, 5. We used the last 10% of the other subjects’ data as negative test data for these three subjects.

Figure 12 shows the results of this longitudinal verification. The average EER for a per-subject threshold was $6.93\% \pm 10.60\%$ while the average EER for a global threshold was 14.4% . The recognition rates of Subjects 1 and 5 were similar to their recognition rates in the initial verification evaluation, while Subject 4 performed better than its initial rate. Recall that these longitudinal samples were taken immediately after each other and thus would be similar enough that if one should match a subject’s model, then the majority of them would. Likewise, the per-subject threshold EER and global threshold EER did not significantly differ from the initial verification evaluation. These results suggest that a subject’s bioimpedance remains stable enough to be verified at least 4.5 months later.

4.6 Energy measurements

Wearable devices require careful design for energy conservation. Although our prototype was not optimized for low power, we report here on its power consumption as a worst-case analysis. To capture energy measurements, we used the Monsoon Power Monitor [21] connected to a Windows laptop. The Power Monitor acts like a battery and samples the current drawn every $200\mu\text{s}$. We down-sampled the current measurements to 100 ms intervals via averaging.

Figure 13 shows the energy measurement of a Shimmer sampling bioimpedance and sending these values to the smart phone. There are five distinct phases in the energy measurement. In the first phase, the Shimmer was idle. This consumed 6.42 mA on average. Next, the Shimmer turned on its Bluetooth radio (at the 3.0 s dashed line) and attempted to pair with the smart phone. On average, this phase consumed 9.33 mA . Again, the spikes in this phase correspond to the times when the Bluetooth radio was searching for the smart phone. Beginning at the 13.8 s dashed line, the Shimmer and smart phone established a connection. This phase consumed 27.0 mA on average. Next, the smart phone instructed the Shimmer (at the 16.3 s dashed line) to collect 12 bioimpedance samples and send them via Bluetooth. This phase consumed 52.5 mA on average. About 30%

of this current was due to the actual bioimpedance sensor board, while about 60 % was a result of the Bluetooth radio (the remaining 10 % was the overhead of just running the device). Next, the smart phone told the Shimmer to stop sampling and disconnected at the 35.0 s and 36.5 s dashed lines, respectively. The phase between stopping sampling and disconnecting consumed 26.8 mA on average. Once again, the dominating factor was the Bluetooth radio, although the bioimpedance sensor board required more energy than the accelerometer. A Shimmer with a 450 mAh battery could last more than a day with the Bluetooth radio on continuously and with hourly bioimpedance recognition. In a production system, the wearable device would also compute the classification results (training would probably still have to be done on another device), sending only these results via Bluetooth, thus eliminating the dominant energy cost. The wearable device could easily last for a week. Furthermore, our prototype system was built on top of a commercial off-the-shelf system and so was constrained by the underlying system itself. One could also minimize the energy profile and size by using a system-on-chip that combines the radio and micro-controller into a single chip [22].

Figure 14 shows the energy measurement of a smart phone engaged in bioimpedance recognition. (This figure is not aligned in time with Figure 13 because we could only measure energy from one device at a time.) Like the measurement in Figure 13, we assume the smart phone has already verified that the Shimmer is on the same body. There are four phases in the energy measurement. The first phase was a steady-state energy measurement of the smart phone. This phase consumed 54.5 mA on average. In the next phase, the smart phone turned on its Bluetooth radio at the 5.0 s dashed line and began connecting to the Shimmer. On average, this phase consumed 87.8 mA. At the 13.8 s dashed line, the smart phone was connected to the Shimmer. In this phase, the smart phone collected bioimpedance samples from the Shimmer. The dashed lines at 10.5 s, 12.2 s, 13.9 s, 15.4 s, 17.0 s, 18.8 s, 20.2 s, 21.8 s, 23.5 s, 24.9 s, 26.3 s and 28.1 s correspond to the times when the Shimmer started sensing a new electrode configuration. This phase consumed 67.9 mA on average. At the 29.6 s dashed line, the smart phone ran the recognition algorithm, told the Shimmer to stop sensing, and began disconnecting from the Shimmer. On average, this phase consumed 91.8 mA. Compared to the overhead of the Bluetooth radio and Android operating system, our bioimpedance recognition method did not significantly impact the current drawn; thus, the bioimpedance calculations causes negligible impact on the smart phone.

5. DISCUSSION AND FUTURE WORK

The usefulness of a biometric relates to its ability to recognize a person within some population. The target population is especially important in the forensic sciences. For a long time the Federal Bureau of Investigation believed fingerprints were unique until an innocent man was linked to the 2004 Madrid train bombings using fingerprint matching [30]. In this paper we emphasize that our target population size is that of a household; that is, we ought to be able to distinguish individuals in a household. While bioimpedance may be able to distinguish individuals in larger populations, such explorations remain future work. We believe that tetra-polar sensing combined with different electrode pair combinations will yield recognition rates on par with biometrics like ECG. One could also combine identification and verification to improve robustness. For example, we could use the identification algorithm to identify who is using the device (like a username) and then use the verification algorithm to decide if it is that person with sufficient probability (like a password).

One advantage of the wrist location is that the wristband is placed in about the same location and at about the same orientation every time it is worn. We experimented with changes in wristband orientation, and determined that it does have an effect on the bioimpedance samples, depending upon the amount of rotation about the wrist. A better physical design might reduce this problem by ensuring the proper band orientation on the wrist. If not, it may be possible to use kinematic sensors to determine the orientation of the band and compensate for different orientations. It might also be possible to compute rotation- and reflection-invariant features. The details of such computations are left for future work.

We did not explicitly consider variations in the bioimpedance due to changes in skin temperature (e.g., for a person with a fever, or who steps outside on a cold winter day), or due to changes in diet (e.g., level of hydration or blood sugar). These and other body conditions may have a measurable impact on bioimpedance that could make it more difficult to develop a robust model for each subject. It might be the case, for example, that a change in blood glucose alters bioimpedance samples measured at the wrist. To be truly confident in this method we need to explore the stability of bioimpedance over weeks or even months, to sample a larger number of subjects, and to explicitly and implicitly explore a broader range of environmental conditions than we captured in our day-long field experiment. We plan to perform such validations in the future.

Although we designed the bracelet for ourselves, a few subjects complained about the tightness of the bracelet. Future bracelets would be designed with different wrist sizes in mind and with better electrodes. Some subjects complained that the electrodes pulled the hair on their wrist. Other subjects mentioned that the device was too bulky to fit under a coat. Our reliance on the Shimmer platform is the source of much of the bulk. Future bracelets could incorporate their own storage, processing, communication, and power without relying on external sources. Custom silicon would also allow model training to be done on the wearable itself.

Our method will suffice for the purpose of identifying the bracelet's wearer, in many interesting applications. In some applications, however, there may be individuals with the motivation to fool the sensor into believing that the wearer is a different person – for example, if the bracelet is used as part of a biometric authentication system, or if the person wishes to have body-area sensor data collected under someone else's identity. We believe, however, that it would be exceptionally difficult to 'forge' another person's bioimpedance. In principle, an adversary could capture the desired person's bioimpedance (by hacking our bracelet to extract the data) and then construct a bracelet liner that 'replays' the impedance using fixed resistors, but this attack would be difficult to accomplish given the frequency-dependent nature of bioimpedance. More so, the threshold for verification algorithms can be chosen to produce more false accepts (i.e., affects security) versus false rejects (i.e., affects usability), and so is a policy decision.

Finally, there are several ways the current design could be optimized for lower cost or reduced energy consumption. The current wristband includes 8 electrodes yet we use only two at a time for measuring bioimpedance. It would be worthwhile exploring whether effective models can be built using only a single pair, or perhaps two pairs, of electrodes. Furthermore, we measured bioimpedance across a wide sweep of 50 frequencies; it may be possible to focus on a smaller number of frequencies, decreasing the energy and time needed for each measurement.

Rather than using bioimpedance itself as a presence detection mechanism (Section 2.4), we could integrate capacitive sensing technologies into our device. In particular, the SemTech SX9300 is an ultra low power, Specific Absorption Rate (SAR) controller




Figure 13: An energy measurement of a Shimmer when collecting bioimpedance data for recognition. The dashed line at 3.0 s is when the Shimmer turned on its Bluetooth radio. The dashed line at 13.8 s is when the smart phone connected to the Shimmer. The dashed line at 16.3 s is when the smart phone told the Shimmer to start sensing bioimpedance. The dashed line at 35.0 s is when the smart phone told the Shimmer to stop sensing bioimpedance. The dashed line at 36.5 s is when the smart phone disconnected from the Shimmer.




Figure 14: Energy measurement of a Nexus S smart phone when running bioimpedance recognition. The dashed line at 5 s is when the smart phone turned on its Bluetooth radio. The dashed line at 10.1 s is when the smart phone and Shimmer established a connection. The dashed lines at 10.5 s, 12.2 s, 13.9 s, 15.4 s, 17.0 s, 18.8 s, 20.2 s, 21.8 s, 23.5 s, 24.9 s, 26.3 s and 28.1 s are when the Shimmer started sensing a new electrode configuration. The dashed line at 29.6 s is when the smart phone classified the bioimpedance samples and disconnected. The dashed line at 32.9 s is when the smart phone turned off its Bluetooth radio.

that can discriminate between the human body and inanimate objects [25]. It distinguishes a human body from inanimate objects by measuring the permittivity (a measure of how freely charged particles can rotate and become polarized when subject to an electric field) of the space near small capacitive sensors (essentially small land areas on a PCB). The integrated circuit comes in a 3 mm x 3 mm x 0.6 mm QFN-20 package and consumes only 459 μW in active mode. It can generate an interrupt to wake a host microcontroller upon a “body close” or a “body far” event, allowing the controller to sleep until human presence is detected or to react to wristband removal. To reduce power costs further it has a doze mode that can scan for capacitive events at a programmable rate of 30 ms to 400 ms per scan while consuming just 48.6 μW .

Alternatively, we could integrate electric-field sensing technologies into our device. Cohn et al. describe such a low-power wake-up method using electric-field sensing technology [7]. Their sensor, like ours, requires contact with the skin, and using their low-power wake-up method only requires 9.3 μW total, which is three orders of magnitude lower power than required to operate our bioimpedance sensor. It would be easy to adopt this approach for use in our device.

Readers interested in more details about our device or its evaluation may wish to review Cornelius [8, Chapter 3].

6. RELATED WORK

There have been many biometrics proposed in the literature, however not all of them are suitable for a system such as we describe.

Most biometrics are unsuitable for our system because they cannot be captured continuously or they need to interrupt the user. Fingerprint recognition for example, requires the users to swipe or hold their fingers on a sensor. Electrocardiogram (ECG) recognition, while continuous, requires an electrical connection across the heart. Electrodes could be integrated into a shirt, however current form factors require the user to touch two electrodes using both hands.

Bioimpedance is often used to measure a person’s body-fat percentage, since they are proportional to each other; several bathroom scales measure both weight and body-fat percentage. There has only been limited use of bioimpedance in support of biometrics, however. Ailisto et al. used body fat (as measured by bioimpedance) and weight to reduce error rates of fingerprint biometrics from 3.9% to 1.5% [1]. Others have used bioimpedance to detect liveness in fingerprint biometrics – see, for example, Martinsen et al. [18] – since a fingerprint reader can be easily fooled; such techniques could be incorporated into our system as well.

We are the first to suggest bioimpedance itself as a biometric. Our prior experiments provided promising evidence that bioimpedance could be a viable method for distinguishing among individuals in a small cohort, such as the members of a household. Those experiments, however, were based on a few samples of each user, under controlled laboratory conditions, using a large bench-top prototype [10]. In this paper we report our success in building such a sensor in the form of a wearable wristband, and demonstrate its potential outside of the lab, with each subject wearing the wristband for the full day. Rasmussen et al. [23] report similar results to ours

but for measurements across the whole body (hand-to-hand) using a non-wearable device intended to be integrated to devices like a laptop or automated teller machine.

Others have used capacitive sensing to differentiate subjects using a capacitive touchscreen. Vu et al. [29] require the subject to wear a special ring that would inject a signal through the subject's finger and into the tablet screen while they are touching the screen. They could encode the ring wearer's identity into this signal, but the ring has no ability to biometrically identify its wearer. Indeed, this signal could be used to communicate anything to the tablet while the user touches the screen, although the data rate (4 bit/s to 5 bit/s) limits the amount of information that can be communicated. Harrison et al. [15] show how to differentiate between subjects using a capacitive touchscreen. Rather than identifying each subject, they focus on determining and tracking the number of users touching the screen. They accomplish this by modifying the touchscreen to measure the impedance between the user and ground across many frequencies. By doing this, they differentiate between subjects interacting with the touchscreen.

Finally, Srinivasan et al. [27] used height sensors to distinguish the subjects of a household. Although height might not be a distinguishing factor for large populations, they showed it is sufficiently distinct for a population the size of a household. Our cohort size was inspired by their household population approach. Our method, however, is suitable for wearable sensors that can be used anywhere, even outside of the home.

7. SUMMARY

In this paper we present a wearable system that can continuously recognize the person wearing our system. Our system is intended for applications that need to confirm that the wearer is indeed the device owner, or that need to distinguish among a small cohort such as the members of a household. To recognize people, our system uses a custom-designed bioimpedance sensor that is used for biometric identification. In contrast to our prior work, we show the effectiveness of our system by prototyping it in the form of a wearable, wrist-worn device with electrodes embedded on the inside of the wristband. This wristband was connected to a custom designed impedance-measuring module we built for the Shimmer research platform. We evaluated the ability of our system to correctly identify its wearer within a cohort of eight subjects (modeling a household of eight people). We found that the device was successful in recognizing its wearer almost 98% of the time using data collected outside of the lab. Furthermore, we show that our recognition method does not adversely affect the battery-life of a smart phone and that our wearable bioimpedance sensor could easily last for a day or longer. Finally, showing that a wearable bioimpedance system works provides the foundation for future studies of bioimpedance as a biometric.

Acknowledgments

We thank Rianna Starheim for her invaluable help with data collection. We also thank the anonymous reviewers for their comments, our fellow colleagues in the Dartmouth TISH group for their feedback, and Mary Baker for her feedback and guidance as our paper shepherd. This research results from a research program at the Institute for Security, Technology, and Society at Dartmouth College, supported by the National Science Foundation under Grant Award Number 0910842 and by the Department of Health and Human Services (SHARP program) under award number 90TR0003-01. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors.

8. REFERENCES

- [1] H. Ailisto, E. Vildjiounaite, M. Lindholm, S.-M. Mäkelä, and J. Peltola. Soft biometrics—combining body weight and fat measurements with fingerprint biometrics. *Pattern Recognition Letters*, 27(5):325–334, Apr. 2006. DOI 10.1016/j.patrec.2005.08.018.
- [2] Analog Devices. Analog Devices ADR433 Ultra-low Noise Voltage Reference. Online at <http://www.analog.com/en/special-linear-functions/voltage-references/adr433/products/product.html>, visited Feb. 2013.
- [3] Analog Devices AD5933 Impedance Analyzer. Online at <http://www.analog.com/en/rfif-components/direct-digital-synthesis-dds/ad5933/products/product.html>, visited Apr. 2013.
- [4] Analog Devices ADG1608 8-Channel Multiplexor. Online at <http://www.analog.com/en/switchesmultiplexers/multiplexers-muxes/adg1608/products/product.html>, visited Apr. 2013.
- [5] S. Avancha, A. Baxi, and D. Kotz. Privacy in mobile technology for personal healthcare. *ACM Computing Surveys*, 45(1), Nov. 2012. DOI 10.1145/2379776.2379779.
- [6] B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory (COLT)*, pages 144–152, 1992. DOI 10.1145/130385.130401.
- [7] G. Cohn, S. Gupta, T. J. Lee, D. Morris, J. R. Smith, M. S. Reynolds, D. S. Tan, and S. N. Patel. An ultra-low-power human body motion sensor using static electric field sensing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp)*, pages 99–102, 2012. DOI 10.1145/2370216.2370233.
- [8] C. Cornelius. *Usable Security for Wireless Body-Area Networks*. PhD thesis, Dartmouth College Computer Science, Hanover, NH, Sept. 2013. Available as Dartmouth Computer Science Technical Report TR2011-741, Online at <http://www.cs.dartmouth.edu/reports/TR2011-741.pdf>.
- [9] C. Cornelius and D. Kotz. Recognizing whether sensors are on the same body. *Journal of Pervasive and Mobile Computing*, 8(6):822–836, Dec. 2012. DOI 10.1016/j.pmcj.2012.06.005.
- [10] C. Cornelius, J. Sorber, R. Peterson, J. Skinner, R. Halter, and D. Kotz. Who wears me? Bioimpedance as a passive biometric. In *Proceedings of the USENIX Workshop on Health Security and Privacy (HealthSec)*, Aug. 2012. Online at <http://www.cs.dartmouth.edu/~dfk/papers/cornelius-impedance.pdf>.
- [11] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B (Methodological)*, 39(1):1–38, 1977. DOI 10.2307/2984875.
- [12] Google. Nexus S – Android Device Gallery. Online at <http://www.android.com/devices/detail/nexus-s>, visited Dec. 2013.
- [13] H. Gray and W. H. Lewis. *Anatomy of the human body*. Lea & Febiger, 20th edition, 1918. Online at http://en.wikipedia.org/wiki/File:Gray417_color.PNG.
- [14] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA data mining software: an update. *SIGKDD Explorations Newsletter*, 11(1):10–18, Nov. 2009. DOI 10.1145/1656274.1656278.
- [15] C. Harrison, M. Sato, and I. Poupyrev. Capacitive fingerprinting: exploring user differentiation by sensing

electrical properties of the human body. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST)*, pages 537–544. ACM, Oct. 2012. DOI 10.1145/2380116.2380183.

- [16] Hirose Electric Co., Ltd. Hirose 3260-8S3(55) Connector. Online at http://www.hirose.co.jp/cataloge_hp/e23200014.pdf, visited Mar. 2013.
- [17] T. Kanungo, D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu. An efficient k-means clustering algorithm: analysis and implementation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):881–892, July 2002. DOI 10.1109/tpami.2002.1017616.
- [18] Ø. G. Martinsen, S. Clausen, J. B. Nysæther, and S. Grimnes. Utilizing Characteristic Electrical Properties of the Epidermal Skin Layers to Detect Fake Fingers in Biometric Fingerprint Systems—A Pilot Study. *IEEE Transactions on Biomedical Engineering*, 54(5):891–894, May 2007. DOI 10.1109/tbme.2007.893472.
- [19] R. Mayrhofer. OpenUAT: The Open Source Ubiquitous Authentication Toolkit. Online at <http://www.openuat.org/>, visited Aug. 2013.
- [20] Microchip Technology Inc. Microchip MCP1252 Charge Pump. Online at <http://ww1.microchip.com/downloads/en/DeviceDoc/21752B.pdf>, visited Feb. 2013.
- [21] Monsoon Power Monitor. Online at <http://www.monsoon.com/LabEquipment/PowerMonitor/>, visited Nov. 2013.
- [22] Nordic Semiconductor. nRF51822. Online at <http://www.nordicsemi.com/eng/Products/Bluetooth-R-low-energy/nRF51822>, visited Nov. 2013.
- [23] K. B. Rasmussen, M. Roeschlin, I. Martinovic, and G. Tsudik. Authentication using pulse-response biometrics. In *Proceedings of the Annual Network and Distributed System Security Symposium (NDSS)*, Feb. 2014.
- [24] R. Rifkin and A. Klautau. In defense of one-vs-all classification. *Journal of Machine Learning Research*, 5:101–141, Dec. 2004. Online at <http://jmlr.org/papers/v5/rifkin04a.html>.
- [25] Semtech. SX9300 Ultra Low Power, Dual Channel, Smart Proximity SAR Compliant Solution. Online at <http://www.semtech.com/touch-interface/capacitive-touch-controllers/sx9300/>, visited Apr. 2013.
- [26] Shimmer-Research. Online at <http://www.shimmer-research.com/>, visited Apr. 2013.
- [27] V. Srinivasan, J. Stankovic, and K. Whitehouse. Using Height Sensors for Biometric Identification in Multi-resident Homes. In *Proceedings of the 8th International Conference on Pervasive Computing (PERVASIVE)*, volume 6030 of *LNCS*, pages 337–354. Springer, May 2010. DOI 10.1007/978-3-642-12654-3_20.
- [28] TinyOS. Online at <http://www.tinyos.net/>, visited Apr. 2013.
- [29] T. Vu, A. Baid, S. Gao, M. Gruteser, R. Howard, J. Lindqvist, P. Spasojevic, and J. Walling. Distinguishing users with capacitive touch communication. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (Mobicom)*, pages 197–208. ACM, 2012. DOI 10.1145/2348543.2348569.
- [30] Wikipedia. Brandon Mayfield — Wikipedia, The Free Encyclopedia. Online at http://en.wikipedia.org/wiki/Brandon_Mayfield, visited Aug. 2013.



Figure 15: The resistor array used to calibrate our system. The $1\text{ k}\Omega \pm 5\%$ resistors are connected in series. The reference resistance between a pair of electrodes will vary depending on the number of resistors between them. For example, the resistance between electrodes 0 4 is $4\text{ k}\Omega$.

APPENDIX

A. CALIBRATION

To compute impedance from the raw data, we calibrated the system using a reference impedance. Calibration is important because different sensor modules will exhibit different impedances due to different parasitic elements and/or artifacts from the amplifiers present in the hardware. We built three of these sensor modules, so we calibrated each device to be sure that our studies are not affected by the particular choice of hardware. In addition, our particular device senses from different electrode pairs via a pair of multiplexors which can themselves introduce additional parasitic elements. Thus, we also must calibrate for every electrode pair.

To calibrate our devices for an electrode pair, we used a known reference impedance value $|Z|_{ref}$ in the form of an array of $1\text{ k}\Omega \pm 5\%$ resistors (we used a digital multimeter to record the exact values between electrodes) as shown in Figure 15. To calibrate impedance magnitude, we first compute the uncalibrated impedance magnitude $|Z(\omega)|_{uncal}$ of the impedance as described in Section 2.2. Because we know the expected magnitude of the reference impedance $|Z|_{ref}$ (the impedance magnitude of a resistor is equal to its resistance), we computed the *gain factor* as:

$$\text{Gain Factor}(\omega) = \frac{1}{|Z|_{ref} \times |Z(\omega)|_{uncal}}$$

For some new uncalibrated impedance magnitude, we compute the calibrated impedance magnitude as:

$$|Z(\omega)| = \frac{1}{\text{Gain Factor}(\omega) \times |Z(\omega)|_{uncal}}$$

We followed a similar calibration procedure for the impedance phase. Since resistors exhibit no impedance phase shift (aside from some parasitic capacitance, which we ignore because these effects fall outside of the bandwidth of our system), we first compute the uncalibrated impedance phase $Z\emptyset(\omega)_{uncal}$ as above. This value is the *phase difference* that accounts for all phase shift attributed to the hardware itself. For some uncalibrated impedance phase, we computed the calibrated impedance phase as:

$$Z\emptyset(\omega) = Z\emptyset(\omega)_{uncal} - \text{Phase Difference}(\omega)$$

This calibration procedure yields a gain factor and phase difference for each frequency and electrode pair that are specific to that sensor module.

During calibration we noticed the lower frequencies exhibited random noise in the signal, of unknown source, so we discard frequencies below 10 kHz. This left 50 frequencies. The AD5933 evaluation boards exhibited similar noise, so we do not believe it is something inherent in the design of our own sensor module.