

Who wears me? Bioimpedance as a passive biometric

Cory Cornelius*, Jacob Sorber*, Ronald Peterson*, Joe Skinner†, Ryan Halter†§, David Kotz*

**Department of Computer Science*

† *Thayer School of Engineering*

§ *Geisel School of Medicine*

Institute for Security, Technology, and Society

Dartmouth College, Hanover, NH

Abstract

Mobile and wearable systems for monitoring health are becoming common. If such an mHealth system knows the identity of its wearer, the system can properly label and store data collected by the system. Existing recognition schemes for such mobile applications and pervasive devices are not particularly usable – they require *active* engagement with the person (e.g., the input of passwords), or they are too easy to fool (e.g., they depend on the presence of a device that is easily stolen or lost).

We present a wearable sensor to passively recognize people. Our sensor uses the unique electrical properties of a person’s body to recognize their identity. More specifically, the sensor uses *bioimpedance* – a measure of how the body’s tissues oppose a tiny applied alternating current – and learns how a person’s body uniquely responds to alternating current of different frequencies. In this paper we demonstrate the feasibility of our system by showing its effectiveness at accurately recognizing people in a household 90% of the time.

1 Introduction

Body-worn sensing systems and wearable devices are becoming more prevalent in our lives. Today, it is not uncommon for people to carry, hold, or wear devices that measure physical activity (e.g., Fitbit [6]), interact with entertainment devices (e.g., the Wii), or monitor their physiology (e.g., a cardiac patient concerned about heart arrhythmia or a diabetic managing her blood glucose). Many more have been proposed or developed as research prototypes. These unobtrusive wearable devices make it possible to continuously or periodically track many health- and lifestyle-related conditions at an unprecedented level of detail. Wireless connectivity allows interaction with other devices nearby (e.g., entertainment systems, climate-control systems, or medical devices). Sensor data may 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.

However, in spite of recent advances, significant challenges remain. Reliably interpreting data from a body-worn sensor often requires information about who is wearing the sensor as well as the current person’s environment, location, current activity, and social context. Techniques exist for collecting some of this information, but today’s body-worn sensors lack the ability to reliably determine who is wearing the device.

In this paper, we focus on a fundamental problem involving wearable devices: who is wearing the device? This problem is key to nearly any application. 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 mixup of sensor data could lead to incorrect treatment or diagnosis decisions, with serious harm to the patient.)

Today, these devices are usually statically associated with a particular person. This smartphone is *my* phone, whereas that fitness sensor is *your* fitness sensor. The device is assumed to be used by only that person; any data generated by a sensor is associated with that person. There are many situations where this model fails, however. In some households, a given device might be shared by many people (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). In some scenarios, 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). When a device does employ some type of authentication mechanism, the traditional solutions are manual, intrusive, and make no guarantees the device is sensing the authenticated user. They require the input of passwords or pin codes, or some kind of challenge-response.

We imagine a device that can be worn on the wrist

and unobtrusively recognize its wearer. Then, a person would be able to simply attach other 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 discover each other’s presence, recognize that they are on the same body (and transitively learn from the wrist device *whose* body), develop shared secrets from which to derive encryption keys, and establish reliable and secure communications. This ability to operate unobtrusively, collecting in situ data and seamlessly integrating computing into a person’s daily life, without interrupting it, is vital to the success of any wearable sensing system.

Not every device need have the capability to recognize who is wearing it. Some will be limited by size or power constraints; others will simply not have access to suitable biometric features of the person. It is sufficient for one device to recognize its wearer, as long as the other worn sensors can recognize that they are on the same body [4] and can communicate securely among themselves; transitively, they all learn the identity of the wearer.

Contributions

Our approach uses the electrical characteristics of biological tissues in a person’s body to recognize who is wearing a wearable device. More specifically, we measure the *bioimpedance* – a measure of how the body opposes a tiny applied alternating current – at the person’s wrist. We evaluate the viability of bioimpedance as a novel, passive biometric in the context of common scenarios for wearable pervasive computing, specifically mHealth. In our experiments using 46 subjects, we were able to reliably recognize the correct person in a hypothetical household 90% of the time.

2 Design goals

Attaching an identity to sensor data requires some method of recognizing whom the device is sensing. One approach, biometric recognition, uses some tell-tale characteristic of the person to determine whether that same person is present at some later time [3]. Biometrics leverage physiological or behavioral characteristics of a person to accomplish recognition. Physiological characteristics range from non-invasive characteristics like facial features and hand geometry to more invasive characteristics like the impression of a finger, the structure of the iris, or the makeup of DNA. Behavioral characteristics include things like the dynamics of using a keyboard, the acoustic patterns of the voice, the mechanics of locomotion, and how one signs a signature. To qualify as a biometric, the chosen characteristic must have the following properties: universality, uniqueness, and permanence. A *universal* characteristic is

one that every person (or most people) possess. Although everyone may possess such a characteristic, the characteristic must also be individually *unique* within a given population. Lastly, the characteristic must have some *permanence* such that it does not vary over the relevant time scale. These properties, with their stated assumptions, are necessary but not sufficient for a biometric that we desire.

In the context of personal health sensors and other pervasive applications, a biometric needs to also be *unobtrusively measured* yet difficult to circumvent. The ability to unobtrusively measure a biometric stems from our desire to provide usable security for personal health sensing systems. Likewise, a biometric needs to be *difficult to circumvent* because there are incentives for people to circumvent them. For example, a person might want to game their insurance provider or fool a physician into believing they have a certain ailment for prescription fraud. Thus, a sufficient biometric will be universal, unique, permanent, unobtrusively measurable, and difficult to circumvent.

Unfortunately, the above-mentioned biometrics are all ill-suited for use with wearable sensing systems. The makeup of DNA, the structure of the iris, and the impression of a finger may be difficult, if not impossible, to forge; however, they are also difficult to unobtrusively measure. Recognition requires the user to interrupt what they are doing to measure the biometric. The behavioral characteristics mentioned above can be measured unobtrusively as the person goes about their day, but they may be easier to forge since they can be easily measured. A microphone can capture a person’s voice, a camera can observe a user’s gait, or a malicious application could learn one’s typing rhythm [9]. In contrast, recognition for wearable sensing applications demands a biometric that is simultaneously difficult to circumvent and easy to measure continuously. We propose to use bioimpedance.

3 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 (transducers) 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 being gauged, including the specific tissue types present (blood, adipose, muscle, bone, etc.), the anatomic configuration (i.e., bone or muscle orientation and quantity), and the state of the tissue (normal or osteoporotic bone, edematous vs. normally hydrated tissue, etc). Significant impedance differences exist between the varying tissue types, anatomic configurations, and tissue state, each of which may provide a unique mechanism for distinguishing between people.

Bioimpedance can be measured by applying a small sinusoidal current between a pair of electrodes attached to the skin. The injected current establishes an electrical field within the tissue and results in a measurable voltage difference between the two electrodes. Thus, potential voltage difference is a function of the underlying tissue impedance. Specifically, the alternating current version of Ohm’s law, $V = IZ$, can be used to relate the voltage V and current I to the bioimpedance Z of the tissue sample. Many tissues exhibit dispersive characteristics, meaning that their electrical properties are dependent on the frequency at which they are measured. Typically, the frequency of the alternating current is swept over a specific band and enables so-called electrical impedance spectroscopy. As a result, complex bioimpedance, $Z(\omega)$, combines resistive and reactive components, $Z(\omega) = R(\omega) + jX(\omega)$, where R is the frequency dependent tissue resistance, X is the frequency dependent tissue reactance, ω is the signal frequency, and j represents the imaginary quantity $\sqrt{-1}$.

Resistance and reactance are dependent on the tissue being measured and the configuration and geometry of the impedance-measuring probe (i.e., electrode size and electrode spacing). In terms of its dependence on the tissue, resistance is primarily associated with the ability of a tissue to carry charge (i.e., current flow through ionic solutions, both intra- and extra-cellular), and reactance is associated with the ability of a tissue to store charge (i.e., the capacitive nature of a cell’s double membrane).

The anatomy of the forearm proximal to the wrist include skeletal bones (radius and ulna), arteries, veins, nerves, muscles, adipose, skin, and interstitial fluids. Over the frequency range of 10 kHz to 10MHz reported values of bone conductivity and adipose conductivity are relatively stable. In muscle, skin, and blood, however, the conductivity monotonically increases with frequency [7].

Person-to-person differences at the wrist include: size, skin thickness, skin water content, bony anatomy (bone sizes), vascular branch size and locations, sub-dermal water content, and adipose/muscle/bone/vasculature content within the sensing region. All of these parameters will have an impact on the actual impedance measured at the wrist. For example, difference in wrist size would represent a change in electrode location and difference in the content, size, and distribution of the underlying tissue types would represent a person-specific conductivity.

4 Wearable device

We imagine the device to be a piece of jewelry, not unlike a watch, that would contain 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




Figure 1: Our bench-top system. The bracelet has a hook-and-loop fastener to hold it in place during bioimpedance measurements. Here, the bracelet is attached to a resistor array we use for calibration.

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 watch 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 (but outside the scope of this paper). Such simple detection mechanisms also allow us to conserve energy by only performing recognition when the device is actually in contact with a person.

Indeed, we have proposed the concept of a wearable device, in a wristwatch form factor, that would coordinate a person’s body-area network of sensors, providing a root of trust; we call it Amulet [12]. Such a device also provides a perfect platform for implementing a biometric recognition mechanism. We expect that the necessary electronics and skin-contact sensors for bioimpedance could easily be integrated into an Amulet-like device.

Experimental apparatus

To collect a large dataset over many frequencies and electrode patterns, we used a custom-designed impedance analyzer constructed specifically for the purpose of recording *in vivo* bioimpedance measurements [8]. Figure 1 shows this bench-top system.

The impedance analyzer is interfaced to a bracelet with eight electrodes (visible at lower right) through an 8-to-1 multiplexer and digital input/output control module [5] that connects current and voltage channels to individually chosen electrodes. This permits impedance measurements to be recorded between any pair of electrodes and thus across almost any part of the wrist. Depending on the selected electrodes, there are two types of mea-

surements that can be taken by the system. A *bi-polar measurement* (i.e., a measurement using two electrodes) occurs when the voltage-measuring electrodes also apply current. Conversely, a *tetra-polar measurement* (i.e., a measurement using four electrodes) occurs when the two voltage measuring electrodes are distinct from the current-applying electrodes. Tetra-polar measurements do not suffer the high contact impedances occurring at the electrode-tissue interface of current-applying electrodes, while bi-polar measurements do. However, as a result, tetra-polar measurements require more complex hardware for making accurate impedance readings. Because impedance can be measured across different pairs of electrodes, the system is (in effect) sensing different parts of the anatomy. For example, impedances recorded between adjacent electrodes are a function primarily of skin and peripheral structures, while impedances recorded between opposing electrodes sense more internal structures. By switching through multiple pairings of electrodes, a list of bioimpedance measurements associated with an individual’s wrist can be recorded and ultimately used for recognizing an individual from within a group of individuals. Finally, we use a resistor array to calibrate the system, and a laptop computer to communicate with the analyzer and multiplexer through a USB-based serial communication protocol.

We wrote custom software to control the system. For each electrode pattern, we run the following sequence. First, the multiplexer is instructed to select the correct electrode pattern. Next, the network analyzer is instructed to sweep through the desired frequency range. Once complete, bioimpedance is computed from the returned data for each frequency and subsequently saved to a file. We repeat this sequence to acquire five measurements per electrode pattern per subject.

5 Method

Before a person can use the device on a daily basis, they must train it to recognize their bioimpedance. They put the device into *enrollment mode*, during which the device captures five bioimpedance measurements from the user in under a minute. The device uses these training measurements as inputs to an *enrollment algorithm* that learns a model of the enrollee’s bioimpedance. (It might be necessary to compute this model off the device because of resource constraints.) Once a model of the enrollee’s bioimpedance is trained, it is loaded into the device for use. (For a device being used by multiple people, it may be loaded with multiple models, but in this paper we limit our analysis to a single enrolled user.)

Once a user is enrolled, the device enters *recognition mode*. In recognition mode, the device periodically determines whether it is on a person’s body (using the mecha-

nisms described above), then collects bioimpedance measurements. The device uses a *recognition algorithm* to determine whether the enrollee’s model matches the measured bioimpedance.

We next describe how our device measures bioimpedance and extracts features for analysis.

Bioimpedance measurements

As shown in Figure 1, the bracelet we designed has 8 electrodes from which to measure bioimpedance. Because it would be infeasible to measure bioimpedance from all combinations of these electrodes, we carefully chose specific electrode patterns. For bi-polar measurements, we only chose those electrodes directly across from one another (e.g., 1 and 5) since they are the maximal distance away from each other and therefore provide more tissue for the current to travel through. Similarly for tetra-polar measurements, we chose to apply current between those electrodes directly across from each other and measure from the other electrodes that are directly across from each other (i.e., apply 1 and 5, measure between 2 and 6, 3 and 7, and 4 and 8). We represent these pairs as a compact list where the first two elements are the electrodes applying current and the last two elements are the electrodes measuring bioimpedance (e.g., 1515 for a bi-polar measurement, and 1526, 1537, or 1548 for a tetra-polar measurement).

Figure 2 shows an example of five bioimpedance measurements from a single subject. Because impedance is a complex value, the plot shows both the resistance (the real part) and reactance (the imaginary part) individually, along with a combined plot showing them both simultaneously. Although there is some inter-measurement variation between the resistance and reactance components of the measurements, the magnitude of the impedance measurement exhibits a linear relationship with the 1 kHz to 100 kHz frequencies, after which there is more variability in the higher frequencies. The resistance and reactance components do exhibit a linear relationship with the middle frequencies (between 10 kHz and 100 kHz).

Feature extraction

Given a set of frequencies and their corresponding bioimpedance measurements, we extract 7 features from each bioimpedance measurement to form a *feature vector*. We extract these features because individual bioimpedance measurements are inherently noisy, and also to reduce the dimensionality of the data and hence the computational and energy overhead.

The first feature we extract is the *maximum magnitude* of all the bioimpedance measurements. We chose this feature because subjects tended to have different maxima according to our initial analysis.

The other six features capture the shape of the bioimpedance measurements as a whole. We fit a line to




Figure 2: Five example bioimpedance measurements collected from a single subject. The electrode pattern was 3751.

the bioimpedance measurements in log-log space, which smooths over the measurements themselves while also preserving the general shape of the curve formed by the measurements. We fit a line for the resistance part, the reactance part, and the magnitude of the bioimpedance measurements. Because each line is succinctly described by a slope and intercept there are six such features (two for each of the three parts described).

Cohort of subjects

Because we cannot know *a priori* the population of subjects who will be using the device, we necessarily need to choose a *cohort of subjects* as an example population. We believe the archetypal cohort is a family since that is the target population. This means that the the cohort of subjects in which the device will be used is relatively small (typically 2 to 5 subjects).

For a large population and large cohort size, it would be infeasible to evaluate all combinations of subjects for the specified cohort size. Thus we randomly selected a subset of all possible cohort combinations such that we have a 95% confidence (with a 1% margin of error) that our sample size is representative.

Enrollment algorithm

Given a set of *training feature vectors* from a cohort of subjects, we learn a model for each user using an *enrollment algorithm*. To do so, for each subject we learn a binary classifier using that subject’s feature vectors as positive examples (i.e., they are labeled positively) and all other subject’s feature vectors as negative examples (i.e., they are labeled negatively).

We examined four different classifiers. The first classifier, *k*-Nearest Neighbors, requires no computation in enrollment mode. Rather, the algorithm simply stores all training examples and their respective labels for use in the recognition algorithm. The second classifier, Naive Bayes, independently models the mean and variance of each feature assuming a Gaussian distribution. A Naive Bayes classifier is a relatively simple classifier to learn because all it requires is computing the mean and variance of each feature for each label. The third classifier, Logistic Regression, learns the parameters of a sigmoid function that best fits the training feature vectors. To avoid

over-fitting, the parameters are L2 regularized. The final classifier, Linear Discriminant Analysis, learns the linear combination of features that best separate the training feature vectors into their respective labels.

Recognition algorithm

Given a set of *test feature vectors* from a cohort of subjects, we can use the enrolled subject’s trained model to classify (i.e., choose a label) whether a particular test feature vector came from that subject. A feature vector that is classified as positive for a given subject’s model is said to match that subject’s bioimpedance; otherwise, the test feature vector is classified as negative because it does not match that subject’s bioimpedance.

Each classifier has different mechanism for classifying test feature vectors. The *k*-Nearest Neighbors classifier chooses the label of the *k*-nearest training feature vectors as the classification, where the nearest training feature vector is defined to be the training feature vector with the smallest Euclidean distance to the test feature vector. For $k > 1$, a majority vote over the *k* labels of the training feature vectors is used to determine the classification. The Naive Bayes classifier chooses the label of the test feature vector with the maximum likelihood as the classification. That is the Gaussian probability density function is computed for the test feature vector given each label, and the label with the maximum value is the classification. Both the Logistic Regression and Linear Discriminant Analysis classifiers computing the dot product of the test feature vector and the learned parameters of the classifier. The classification, then, is the *sign* (i.e., positive or negative) of this dot product.

Parameters

Aside from empirically choosing the classifier that most accurately recognizes subjects, we examined two other parameters to determine their affect on recognition rates.

The first parameter is the size of the cohort. In general, the smaller the size of the cohort, the better recognition rates we expect; with fewer subjects being considered, the easier it should be to distinguish between subjects. We evaluated cohort sizes of 2, 3, 4, 5. We also evaluated a cohort size of 46 to see how well our method works in the presence of all subjects.



Figure 3: The AccuFitness MyoTape Body Tape Measure we used to measure the circumference of each subject’s wrist just below the subject’s ulnar styloid process.

The second parameter is the electrode pattern, which affects bioimpedance measurements (and therefore recognition rates) since the current travels through different tissue for different electrode patterns. Because there are many combinations of electrode, we examined a subset of them to determine which ones yielded the best recognition rates. For bi-polar measurements, the electrodes applying current are also sensing, so there were only four possible patterns: 1515, 2626, 3737, and 4848. In the case of tetra-polar measurements, the sensing electrodes can be chosen independently of the applying electrodes; we examined electrode patterns 1526, 1537, 1548, 2637, 2648, 2651, 3748, 3751, 3762, 4851, 4862, and 4873.

6 Dataset

We collected data from human subjects using a protocol and device approved by our Institutional Review Board. After obtaining informed consent, we instructed users to fill out a questionnaire to collect their age and gender. We used an AccuFitness MyoTape Body Tape Measure [11] (as shown in Figure 3, at right) to measure the circumference of their left wrist, to millimeter precision. We measured the circumference at the location just below the ulnar styloid process as shown in Figure 3 at left. Once enrolled, we placed the electrode bracelet on the subject’s left wrist and they were instructed to keep their wrist still until data collection finished. The data collection sequence took roughly 12 minutes per subject. After completion, the subject was compensated for their time.

We collected bioimpedance measurements from 46 subjects, 22 males and 24 females. The average subject age was 21 years ($\sigma = 3$); all subjects were 18 years or older. In total, we collected 80 measurements from each subject (5 measurements for each electrode pattern), resulting in 3680 total bioimpedance measurements.

Figure 4 shows a histogram of wrist circumferences by gender. The average subject wrist circumference was




Figure 4: A histogram of wrist circumferences by gender. Males tend to exhibit larger wrist circumferences.

16.0 cm ($\sigma = 1.33$). For females, the average wrist circumference was 15.1 cm ($\sigma = 0.55$), and males 17.0 cm ($\sigma = 1.21$).

7 Evaluation

To evaluate the efficacy of our approach, our experiments focus on the ability of bioimpedance measurements to distinguish among people. We sought to determine how well our method performs in the context of a family where all the users are known and the goal is to recognize whether the device is worn by its intended user.

Metrics

Consider a given subject S , and a set of test feature vectors from a larger cohort of subjects. We label the test vectors measured from subject S as *positive* and all other test vectors as *negative*. Now consider the model trained on subject S and apply it to all the test feature vectors, resulting in a positive or negative classification for each; ideally, the model classifies only those test feature vectors from S as positive. We define the *false accept rate* (FAR) as the fraction of negatively labeled feature vectors that were misclassified (i.e., they were classified as positive). Similarly, we define the *false reject rate* (FRR) as the fraction of positively labeled feature vectors that were misclassified (i.e., they were classified as negative). We also define the *balanced accuracy* (BAC) as 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 $1 - \text{FAR}$). Balanced accuracy weights the negative and positive examples equally, since for some sizes of cohorts there are more negatively labeled feature vectors than positively labeled feature vectors.

For each subject we ran a leave-one-out cross-validation over the set of feature vectors in a cohort of subjects according to the algorithms specified in Section 5. We computed the FAR, FRR, and BAC for each subject,

and we report the average and standard deviation of these measures over all subjects in the cohort.

Recall that our method can be parameterized by cohort size and by electrode pattern. We explored the parameter space to find an optimal setting that maximizes recognition rates across all subjects. For each experiment we show the top three performing electrode patterns.

Single-pattern bioimpedance recognition

In this experiment, we tested all combinations of electrode patterns according to the method described in Section 5. The purpose of this experiment is to understand how accurately a single electrode pattern can recognize a subject.

Figures 5 and 6 show the results of this experiment for bi-polar and tetra-polar electrode patterns respectively. As one would expect, recognition rates decrease with the cohort size. With limited space we can present only the results from Naive Bayes, which yielded the best overall recognition rates for both bi-polar and tetra-polar electrode patterns. The other classifiers performed particularly poorly in the 46-subject cohort.

Among bi-polar electrode patterns, there was a clear winner: the 4848 electrode pattern yielded the best recognition rates for three classifiers and came in a close second for the fourth classifier. The location of these electrodes on the bracelet roughly correspond to the bottom and top of the wrist, which is the smallest distance between any pair of electrodes we examined.

There was little difference in terms of recognition rates between bi-polar and tetra-polar measurements. Thus, either mode of measurement would suffice when using a single pattern for bioimpedance recognition.

Multi-pattern bioimpedance recognition

Since we are not limited to just one particular electrode pattern, we hypothesized it might boost recognition rates to concatenate feature vectors from multiple electrode patterns into a multi-pattern feature vector. For example, in the bi-polar case we could incorporate feature vectors from the 1515 and 2626 electrode patterns by concatenating them together. This approach might boost recognition rates because the applied current takes different paths through the subject's wrist for different electrode patterns.

In the bi-polar case, we explored all combinations of bi-polar electrode patterns (e.g., 1515 2626, 1515 3737, ..., 1515 2626 3737 4848); there are 11 such combinations. In the tetra-polar case, we explored all combinations of tetra-polar electrode patterns such that the electrodes supplying current are distinct (e.g., 1515 2626, but not 1515 1526); there are 243 such combinations.

Figures 7 and 8 show the results of this experiment. In comparison to the single electrode feature vector, there was only a modest boost in balanced accuracy and de-

crease in its variance for both bi-polar and tetra-polar electrode patterns. This result implies that a two-electrode device may be sufficient for most purposes. Notice, however, that the false accept rates decreased at the expense of a higher false reject rate. Thus, if an application requires fewer false accepts, then we recommend using multiple-pattern bioimpedance recognition. Four electrodes in a bi-polar configuration should be sufficient.

Wrist circumference recognition

From the data in Figure 4, we hypothesize that a subject's wrist circumference might serve as a good feature for recognizing subjects since many of the subjects fall into their own bin on the histogram. We did not take multiple measurements of each subject's wrist circumference, because the error distribution experienced by a real wearable device would depend on the measurement characteristics of that device. Instead, we simulated taking multiple measurements by assuming some measurement error. For example, if a subject's wrist circumference was measured to be 15 cm and the device has a measurement error of 5 mm, then we compute 10 linearly spaced measurements from 14.5 cm to 15.5 cm. In this experiment, we examined how measurement errors of 0.1 cm, 0.5 cm, and 1.0 cm affect recognition rates where the feature vector is the measurement itself.

Figure 9 shows the result of this experiment. With no measurement error, we can accurately recognize users 98% of the time regardless of the cohort size. (In a typical household, with a broader cohort age diversity than in our subject population, the accuracy should be even better.) As measurement error increased, however, the recognition rates fell: with 1 mm measurement error, recognition rates fell to 54% for the full cohort of subjects. For smaller cohort sizes, recognition rates remain above 90% when the measurement error is 1 mm or less. This result implies that wrist-size is a good biometric for small cohorts; can we do better with a hybrid approach?

Combining bioimpedance with wrist circumference

Since wrist circumference appears to be a good indicator of identity, we added a wrist-size feature to our bioimpedance models. Figure 10 shows the results of this experiment for a cohort of 5 subjects, for the top-performing electrode patterns (for both single and multi pattern). In contrast to using wrist circumference alone, combining bioimpedance and wrist circumference dramatically lower the false reject and false accept rates for both larger cohort sizes and larger measurement errors. If a device can be built to measure both wrist circumference and bioimpedance, this device would be ideal for mobile health applications, especially large cohorts, because of the lower false accept rates.

The analysis above assumes there is only one enrolled user per device. One could extend our method to multiple




Figure 5: The three bi-polar electrode patterns that yielded the best recognition rates for various cohort sizes using a Naive Bayes classifier.




Figure 6: The three tetra-polar electrode patterns that yielded the best recognition rates for various cohort sizes using a Naive Bayes classifier.




Figure 7: The three bi-polar electrode multi-patterns that yielded the best recognition rates for various cohort sizes using a Naive Bayes classifier.




Figure 8: The three tetra-polar electrode multi-patterns that yielded the best recognition rates for various cohort sizes using a Naive Bayes classifier.




Figure 9: Wrist circumference recognition rates for various cohort sizes and measurement errors using Naive Bayes.




Figure 10: Bioimpedance with wrist-circumference models: recognition rates for a 5-subject cohort and various measurement errors, using a Naive Bayes classifier.

enrolled users by changing the recognition algorithm to choose the model that classified the test feature vector positively, or choose no model in case that no model, or more than one model, positively classified the test feature vector.

8 Impedance as a passive biometric

Recall that a passive biometric has the following characteristics: universality, uniqueness, permanence, unobtrusively measurable, and difficult to circumvent. As with any system, there is a set of assumptions we made about the threats and adversaries the system is designed to handle with regard to these characteristics.

Universality: We assume that every person has a wrist where we can measure bioimpedance. Of course this is true for most people; nonetheless, our technique could be used at other locations on the body.

Uniqueness: The core threat we address is an attacker simply wearing the device to confuse it into thinking the enrolled subject is wearing it. We assume that a device knows the population of attackers, that is, the population of potential impostors who may attempt to wear the intended subject’s device. In the household context, the population is the set of household members. Given this population, the device should be able to determine whether it is on the body of the legitimate user, based on the size and

bioimpedance measurements of the wearer’s wrist. Our analysis simulated ‘attackers’ by testing all other subjects in the population (i.e., the set of negatively labeled test feature vectors). The proportion of times these simulated attackers were successful is represented by the FAR. For a cohort of five, and wrist-size measurement error 1mm, an attacker would be successful only 4% of the time; see the combined bioimpedance-circumference experiment in Figure 10. If many bioimpedance measurements are taken over the course of a day, while the user wears the device, such an attack can be mitigated: a non-legitimate wearer would have to repeatedly succeed to be recognized – highly unlikely with only 4% chance of success each time. Thus, we believe bioimpedance measurements are individually unique for a household population.

Permanence: Although we expect the bioimpedance of a person to be reasonably permanent, changing over long time scales because the size and shape of our wrist changes as we age, we need to explore short- and medium-term variations due to diet or physical activity. We are currently constructing a wearable prototype using a chip [1] capable of bi-polar measurements, and we plan to use this wearable prototype in a longitudinal study.

Unobtrusive: We chose the wrist location to allow unobtrusive measurements, as many people already wear watches or bracelets and our method could easily be integrated into such a form factor. We plan to study how

comfort versus the contact and movement of the electrodes affects classification rates.

Circumvention: Although we did not experimentally explore methods to actively circumvent our approach, we believe the bar is high enough to make such attacks infeasible. An attacker would have to model the physiology of an enrolled user’s wrist in order to succeed. We plan to follow up with a study of the feasibility of various attack scenarios.

9 Related work

To our knowledge, no one has used bioimpedance itself as a biometric. However, it has been used to measure a person’s body fat percentage since they are proportional to each other. Ailisto et al. [2] used bioimpedance and body weight to reduce error rates of fingerprint biometrics from 3.9% to 1.5%. We, on the other hand, use bioimpedance itself as the biometric and combine it with wrist size, two measurements that can be realized in a wearable device. Others have used bioimpedance to detect liveness in the case of fingerprint biometrics, since a fingerprint reader can be easily fooled. Martinsen et al. [10] present such a system to detect liveness. Such techniques could be incorporated into our system as well. Finally, Srinivasan et al. [13] 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.

10 Conclusion

In this paper we describe a method for wearable sensors to recognize a person using bioimpedance, with the goal of supporting body-area mHealth sensors by giving them the ability to recognize wearer identity. We studied this approach in the context of a household population, and experimentally show that our method has a balanced accuracy of 85%. When combined with wrist circumference measurements with 1 mm measurement error, our method is 90% accurate at recognizing users. Our next step is to construct and evaluate a wearable prototype.

Acknowledgments

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 pro-

gram) 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.

We also thank Rianna Starheim for her invaluable help with collecting data from subjects and the anonymous reviewers and our fellow colleagues in the Dartmouth TISH group for their valuable feedback.

References

- [1] Analog Devices AD5933. Online at <http://www.analog.com/en/rfif-components/direct-digital-synthesis-dds/ad5933/products/product.html>, visited Apr. 2012.
- [2] 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.
- [3] R. M. Bolle, J. H. Connell, S. Pankanti, N. K. Ratha, and A. W. Senior. *Guide to Biometrics*. Springer Professional Computing, 2004.
- [4] C. Cornelius and D. Kotz. Recognizing whether sensors are on the same body. In *Proceedings of the 9th International Conference on Pervasive Computing (Pervasive)*, volume 6696 of *LNCS*, pages 332–349. Springer, June 2011. DOI 10.1007/978-3-642-21726-5_21.
- [5] DT9835. Online at <http://www.datatranslation.com/products/dataacquisition/usb/DT9835/>, visited Apr. 2012.
- [6] Fitbit. Online at <http://www.fitbit.com/>, visited 2010.
- [7] S. Gabriel, R. W. Lau, and C. Gabriel. The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz. *Physics in Medicine and Biology*, 41(11):2251–2269, Nov. 1996. Online at <http://view.ncbi.nlm.nih.gov/pubmed/8938025>.
- [8] R. J. Halter, A. Hartov, and K. D. Paulsen. A broadband high-frequency electrical impedance tomography system for breast imaging. *IEEE Transactions on Biomedical Engineering*, 55(2 Pt 1):650–659, Feb. 2008. DOI 10.1109/TBME.2007.903516.
- [9] P. Marquardt, A. Verma, H. Carter, and P. Traynor. (sp)iPhone: Decoding Vibrations From Nearby Keyboards Using Mobile Phone Accelerometers. In *Proceedings of 18th ACM Conference on Computer and Communications Security (CCS)*, pages 551–562. ACM, Oct. 2011. DOI 10.1145/2046707.2046771.
- [10] 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.
- [11] AccuFitness MyoTape Body Tape Measure. Online at <http://www.accufitness.com/index.php/body-tape-measure-myotape>, visited Apr. 2012.
- [12] J. Sorber, M. Shin, R. Peterson, C. Cornelius, S. Mare, A. Prasad, Z. Marois, E. Smithayer, and D. Kotz. An Amulet for trustworthy wearable mHealth. In *Workshop on Mobile Computing Systems and Applications (HotMobile)*, Feb. 2012. Online at <http://www.cs.dartmouth.edu/~dfk/papers/sorber-amulet.pdf>.
- [13] 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.