9-1-2013

Usable Security for Wireless Body-Area Networks

Cory Cornelius
Dartmouth College

Cory Cornelius

Usable Security for Wireless Body-Area Networks

Dartmouth Computer Science
Technical Report TR2013-741
Submitted September 2013

USABLE SECURITY FOR WIRELESS BODY-AREA NETWORKS

A Thesis

Submitted to the Faculty

in partial fulfillment of the requirements for the

degree of

Doctor of Philosophy

in

Computer Science

by

Cory Cornelius

DARTMOUTH COLLEGE

Hanover, New Hampshire

September 2013

Examining Committee:

David Kotz, Chair

Sean Smith

Andrew Campbell

Ryan Halter

Stephen Intille

F. Jon Kull
Dean of Graduate Studies

Abstract

We expect wireless body-area networks of pervasive wearable devices will enable *in situ* health monitoring, personal assistance, entertainment personalization, and home automation. As these devices become ubiquitous, we also expect them to interoperate. That is, instead of closed, end-to-end body-worn sensing systems, we envision standardized sensors that wirelessly communicate their data to a device many people already carry today, the smart phone. However, this ubiquity of wireless sensors combined with the characteristics they sense present many security and privacy problems.

In this thesis we describe solutions to two of these problems. First, we evaluate the use of bioimpedance for recognizing who is wearing these wireless sensors and show that bioimpedance is a feasible biometric. Second, we investigate the use of accelerometers for verifying whether two of these wireless sensors are on the same person and show that our method is successful as distinguishing between sensors on the same body and on different bodies. We stress that any solution to these problems must be usable, meaning the user should not have to do anything but attach the sensor to their body and have them *just work*.

These methods solve interesting problems in their own right, but it is the combination of these methods that shows their true power. Combined together they

allow a network of wireless sensors to cooperate and determine whom they are sensing even though only one of the wireless sensors might be able to determine this fact. If all the wireless sensors know they are on the same body as each other and one of them knows which person it is on, then they can each exploit the transitive relationship to know that they must all be on that person's body. We show how these methods can work together in a prototype system. This ability to operate unobtrusively, collecting *in situ* data and labeling it properly without interrupting the wearer's activities of daily life, will be vital to the success of these wireless sensors.

Greetz

Foremost, none of this would have been possible without the guidance and support of my advisor David Kotz. I could not think of a better advisor. I have also been lucky to work with Apu Kapadia, Minho Shin, and Jacob Sorber, who have shaped my thinking in so many ways. I am grateful to Ron Peterson, who has been an invaluable source of technical expertise. To my committee, Sean Smith, Andrew Campbell, Ryan Halter, and Stephen Intille, I am indebted for their patience and feedback. To my labmates, Shrirang Mare and Aarathi Prasad, who kept me energized on those late nights before deadlines. I am thankful to the interns, Janet Kim, Zachary Marois, Rihanna Starheim, who I have had the pleasure of working with. To my friends, Dan Peebles and Steve Gomez, who have kept me curious. To my family, who have been supportive throughout. I am grateful to my Brother, who has been a source of inspiration and levelheadedness, and to my Mother, who has showed me nothing is impossible and has kept me going. Finally, to my Grandmother, who gave me my first computer and taught me to stand up for what I believe in.

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 under award number 90TR0003-01.

Contents

1

Introduction

Wearable devices are becoming more prevalent in our lives. Today, it is not uncommon for people to carry multiple computing devices, such as smart phones, music players, and cameras; increasingly, they also carry, hold, or wear devices to measure physical activity (e.g., Fitbit [50] or BodyBugg [12]), to interact with entertainment devices (e.g., the Xbox 360 + Kinect [111]), or to monitor their physiology (e.g., a cardiac patient concerned about heart arrhythmia or a diabetic managing her blood glucose). Many more such devices and systems have been proposed or developed as research prototypes [80, 81]. 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 enables interaction with other nearby devices (e.g., entertainment systems, climate-control systems, or medical devices), the automatic sharing of sensor data with a social-networking service or (in the case of health applications) a user's Personal Health Record (PHR) system or an Electronic Health Record (EHR) system for review by a health care provider.

More so, as the population of the world grows, it will become increasingly necessary to use these technologies to monitor, diagnose, and treat individuals. These wearable devices will enable the collection of longitudinal health-related data, and physicians will use this data to make better medical decisions. However, these devices present new and unique security and privacy challenges. If we expect them to become a solution for monitoring the world's aging population, then it is imperative that we be mindful of these challenges. Indeed, the kind (medical) and magnitude (continuous) of data they collect in a mobile health application can be considered protected health information [71]. More generally, many researchers have noted that these wearable devices are energy scarce, and, in doing so, have made significant steps towards developing energy-aware security solutions such that the overhead of security is minimal if not nil [66, 139]. Finally, we must also remind ourselves that any security and privacy solutions ought to be usable, because they will only be adopted if they are both useful and usable.

1.1 Motivation

Suppose Alice and Fred, a health-conscious couple living together, each decide to buy a fitness-monitoring sensor. The instructions indicate that they should each "pair" their respective sensor with their own smart phone. Pairing ensures, through cryptographic means, that a sensor is only able to communicate with a specific smart phone. One day, when Alice and Fred go for a run, Alice unknowingly wears Fred's

sensor, and Fred wears Alice's sensor. As they run together, and therefore remain in communication range, Fred's smart phone would be collecting data about Alice, but labeling the data as Fred's and placing it in Fred's health record. Alice's smart phone would be similarly swapping labels and data. This problem, a result of the one-to-one pairing model, is even more likely as the number of sensors grows. Pairing makes the implicit assumption that the sensor paired with the smart phone will not be used by anyone else but the user of the smart phone. Depending on the usage of the data, this can be a security problem.

Today, many of the commercial versions of these devices one can buy are highly specialized, end-to-end solutions that use protocols (e.g., Bluetooth [23], ANT [11], or ZigBee [163]) with little to no security. It has been shown, for example, that there is very little security built into the Fitbit [122]. When a commercial device does use some type of security, the solutions are usually manual and intrusive. They often require the input of passwords or personal identification numbers or require some kind of challenge-response scheme. Additionally, as the example with Alice and Fred shows, these devices are also usually statically associated with a particular person. That is, this smart phone is *my* phone, whereas that fitness sensor is *your* fitness sensor. This assumption implies that any data generated by a sensor must also be associated with that person. However, there are many situations where this assumption fails. In some settings, 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 (as in the example). Finally, a person might actively try to fool the device (e.g., a smoker who places his "smoking" sensor on a non-smoking friend to receive incentives for smoking cessation).

Our goal is to make life easier for people like Alice and Fred. Although Alice and Fred buy identical sensor devices, Alice should be able to put on either device and have her smart phone recognize which device is attached *to her*, automatically

creating the phone-device association without an explicit pairing step. Similarly, if Alice and Fred jointly own another sensor device, either may use the sensor at any time, and again the correct smart phone should detect which body is wearing the sensor and receive the data into the correct person's health record.

1.2 Challenges

Given this vision, there are significant security, privacy, and usability challenges. We have outlined elsewhere many of these problems [134, 36]. For example: How does each device authenticate the other? How do devices discover each other without exposing their own presence? How does a user pair devices together? How can users trust that their devices are not leaking the data they are collecting? How do we balance privacy, usability, and utility of a network of these devices? We have proposed a device and architecture, called *Amulet*, that could act as solution to some of these problems [134].

This thesis focuses on three specific challenges whose solutions would contribute to the realization of an Amulet:

1. How can we recognize who is wearing a device?
2. How can we verify whether a suite of devices are attached to same person?
3. How do the solutions to these challenges fit into an system?

These problems matter for at least one important reason. If we expect physicians to make decisions about health data collected from wearable devices, then they will need some confidence that all the data they are examining has come, at the very least, from the same body. It is risky to make medical decisions without said confidence, and until such measures are in place, we should be wary of using data from these devices.

1.3 Contributions

This thesis contributes two methods as solutions to challenges described above and integrates them into a working system.

The first method explores the use of *bioimpedance* to recognize wearers. We describe a sensor that is the first wrist-worn sensor capable of sensing bioimpedance. We also describe a dataset we collected from individuals that wore our sensor. Additionally, we empirically show the feasibility of this sensor in its ability to recognize the wearer over time and to distinguish among several wearers. Finally, we examine the feasibility of this sensor in different environmental conditions and motion artifacts.

The second method explores the use of *acceleration* and *orientation* to verify whether sensors are on the same body. We sense these quantities using an accelerometer and gyroscope placed at different locations on the body – the wrists, the ankles – and describe a method to determine whether these sensors are on the same body as an accelerometer and gyroscope placed at a person's waist. Using a dataset of subjects wearing these sensors, we provide empirical results that our method works and is robust to cases even when subjects are walking together.

Finally, we integrate these methods into a network of wearable devices. We describe methods for detecting when these sensors are attached to a person. Additionally, we show how to pair these devices to bootstrap a privacy-preserving protocol. We also show that our system is feasible by analyzing the number of walking intervals per subject. Finally, we provide an analysis the amount of energy our system consumes.

1.4 Outline

This thesis is organized according to each method. Chapter 2 establishes a common language and some background about each method. Chapter 3 motivates and evaluates the use of bioimpedance as a biometric. In Chapter 4, we describe how to use acceleration and orientation to determine whether two sensors are on the same body and empirically evaluate this method. Chapter 5 neatly integrates these two methods into a usably secure wireless body-area network. Lastly, Chapter 6 describes some open problems with this work and Chapter 7 concludes.

2

Background

The solutions to these problems were designed with two considerations in mind. First, that a network of these wearable devices exists (i.e., there is more than just one device present), and second, that usability was of paramount concern. To better understand the context in which these solutions were developed, this chapter defines what we mean by these two considerations and formalizes the particular kind of adversary and threats we intend to mitigate. It also introduces some terminology regarding the machine-learning techniques we use.

2.1 Wireless body-area networks

We believe the sensors in these wearable devices will become commodities. They will interoperate with each other and other devices that people carry with them, like a smart phone. In the future, we imagine smart phones will be replaced with "an mHealth architecture that provides strong security and privacy guarantees" we call Amulet [134]. A person will wear several sensors of varying types (e.g., blood-pressure monitor, pulse oximeter, pedometer, blood-glucose meter) and due to physiological requirements, or comfort, these sensors will necessarily be attached at different locations on the body. These sensors will primarily communicate with each other and other devices over some wireless medium. The central hub of this network will store and aggregate all data coming from the sensors. In fact, constructing such a system is feasible today as there are commercially available medical and fitness sensors capable of communicating with smart phones via Bluetooth.

This network of wearable devices is called a Wireless Body-Area Network (WBAN). Figure 2.1.1 depicts such a network. In a WBAN, a person wears one or more Sensor Nodes (SNs), like a blood-glucose monitor, pedometer, or electrocardiography sensor, on their body and carries a single Mobile Node (MN), like a smart phone or Amulet, at their waist. The MN and SNs communicate wirelessly at a range near the body (typically, less than 2 m) in a star network topology. If desired, this WBAN can communicate with some back-end using the MN's Internet connectivity. However, we are more concerned with the WBAN itself rather than the Internet connectivity of the MN.

In a WBAN, we assume the MN is capable of communicating with some back-end service system (e.g., an EHR or PHR). The SNs, on the other hand, have limited computational and energy resources, and, as such, are only capable of communicating with the MN. Additionally, we assume SNs have the ability to detect when

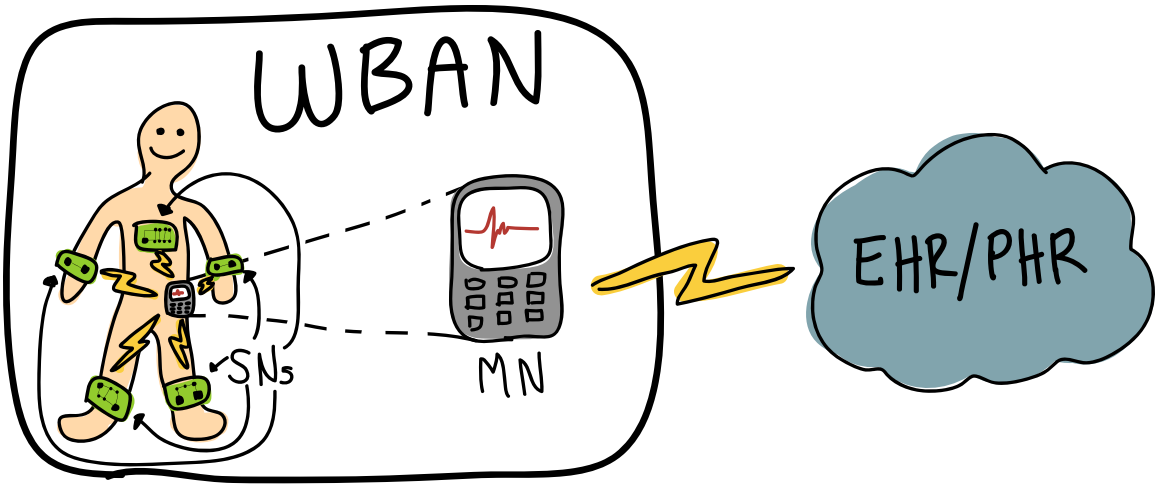

Figure 2.1.1: The components of a Wireless Body-Area Network. Many SNs communicate wirelessly with a central MN. The MN is capable of forwarding information from the SNs to an application in the cloud (such as an EHR or PHR).

they are attached to a person (although they might not know to whom). An SN, for example, might contain a circuit that is completed. For example, when a person straps an SN onto their body, the two ends of a necklace or wrist-strap come into contact and complete the circuit. How this might be accomplished for other locations is explained in more detail in Chapter 5.

Sensor Nodes can be further divided into one of two categories: personal and shared. A *personal* SN is one that is exclusively used by one person (e.g., a wearable blood-glucose monitor). A *shared* SN is one that can be used by several people (e.g., a fitness monitor). Sensor Nodes can also be classified by their intended usage. For example, an *ephemeral* SN is one that a person sporadically uses (e.g., a wireless digital scale). These types of SNs tend to be embedded in the environment, typically allowing multiple people to use them (although there can be personal, ephemeral SNs like a blood-glucose monitor). On the other hand, *continuous* SNs are those that a person wears continuously (e.g., a fitness monitor). This thesis is primarily concerned with these kind of SNs because they are continuously worn by people.

We assume the MN and SNs interact with the following entities: a manufacturer, a user, and a service provider. A *manufacturer* is an entity that produces the sensors,

and distributes them to the service provider and the users. A *user* is a person that uses the MN and SNs to get information about their own health. Typically, users obtain an MN or an SN either directly from a manufacturer or from a service provider. Users can also choose to forward this information to a service provider for consultation or keep the data local for self-monitoring. A *service provider* is an entity that provides services such as configuring sensors, analyzing sensed data, and providing consultation to users. For example, a service provider can be a hospital, PHR provider, a home-health organization, or some other business or service provider.

Wireless Body-Area Networks have long been studied by researchers, often under different names (e.g., body-area sensor networks and personal area networks) [142, 66]. The Institute of Electrical and Electronics Engineers (IEEE) 802.15 task group six formally defines a Body-Area Network (BAN) as "a communication standard optimized for low power devices and operation on, in or around the human body (but not limited to humans) to serve a variety of applications including medical, consumer electronics / personal entertainment and others" [73]. However, while the underlying physical, medium access control, and network layers are interesting, we take them for granted since we are primarily concerned with providing usable security mechanisms for the wireless variant of BANs.

2.1.1 Locations of interest

There are many locations on your body where you could wear an SN. Figure 2.1.2 graphically depicts several likely locations. The highlighted locations – Left Hand (HL), Right Hand (HR), Left Ankle (AL), Right Ankle (AR), Right Waist (WR) – indicate those locations we examine in this thesis. Notably, we did not investigate locations located on the upper body. Your head, for example, is one location where you might wear an SN (e.g., Google Glass [57]). Likewise, you might also wear an SN on your Left Chest (CL) if you wanted to monitor your heart [129] or respiratory

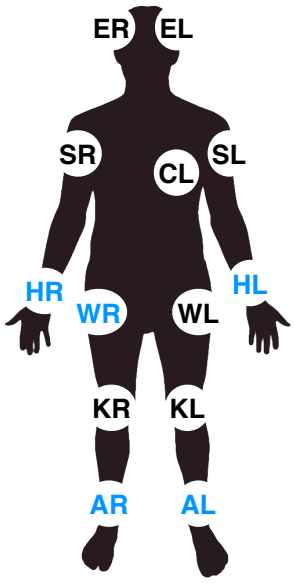

Figure 2.1.2: Some locations where a person could wear an SN. The highlighted ones – HL, HR, WR, AR, and AL – are the locations we examined in this thesis.

rate [115]. Additionally, some sensors are designed to be worn on the upper arm (e.g., BodyBugg [12]). We leave these other locations for future work.

2.2 Usability

One can imagine all kinds of solutions to the challenges we outlined, but we emphasize that any solution to these challenges must be usable. Often the target population for these kind of systems is the elderly, very young, ill, or disabled, and so usability is important. More so, the more usable a device is the more likely people are to adopt it. We define a usable solution to mean a person should only need to attach an SN to their body – whether clipped on, strapped on, stuck on, slipped into a pocket, or even implanted or ingested – and have them *just work*. That is, without any other action by the user, the MN and SNs would discover each other's presence, recognize that they are on the same body, transitively learn *whose* body they are on, independently compute shared secrets from which to derive encryption keys,

and finally establish reliable and secure communications. Some solutions to these problems might require a training period (e.g., to learn some characteristic about the user that can be later used to recognize them), which should also be accomplished in a manner that minimize the amount of work the user has to do. This ability to operate unobtrusively, collecting *in situ* data without interrupting the wearer's activities of daily life, is vital to the success of WBANs.

The idea that security mechanisms ought to be usable is by no means a new idea. Venkatasubramanian et al. define usable security mechanisms "as those which activate on deployment, in a plug-n-play manner, with minimal (ideally none) initialization procedures" [147]. We adopt this same definition of "usable" and use it to guide our own usable security mechanisms. The idea that there ought to be minimal initialization is important, because some solutions to these challenges will necessarily require some initialization on part of the user. For example, to identify who is wearing a device, the device needs to have some data about that user. In the ideal case, the device would detect that a new user is wearing it and prompt them (via the MN, for example). Alternatively, the MN need not prompt the user, rather, it could create a unique but temporary label until a more permanent label can be supplied by the user. Such a scheme would not interrupt the user at all.

2.3 Security model

In any system, designers make a set of assumptions about the system and the entities that interact with it. For example, designers make assumptions about the capabilities of an adversary and the kind of threats adversaries might try to commit against the system. We state those assumptions here and define those security goals our system achieves under these assumptions.

2.3.1 Hardware capabilities

We make the following assumptions about the kind of hardware that is necessary to support our system. The MN and SNs must have wireless capabilities (e.g., 802.15.4, 802.11, or Bluetooth) and support cryptographic primitives like encryption and authentication.[1] We also assume the MN and SNs can securely pair; that is, they can authenticate each other and share keys using a pairing method that makes use of the cryptographic primitives.[2] We also assume the MN and SNs have some common sensor or sensors, like an accelerometer or gyroscope.

2.3.2 Trust assumptions

We define *trust* to mean those assumptions that one entity has about another entity. We make the following assumptions about how the three entities – manufacturer, service provider, and users – trust each other in our system. First, a user and the service provider trust the manufacturer to produce calibrated SNs that operate correctly. This assumption allows users and service providers to trust that the SNs are providing correct data. A user also trusts the service provider to not disclose the sensor data it has received from the MN. The service provider trusts the users to not tamper with the hardware or software of the MN or the SN. The manufacturer assumes nothing about any of the other entities.

2.3.3 Adversary model

Adversaries want to thwart the security of the system. As such, we assume they have the following capabilities that would allow them to attack the system. We assume the adversary has the ability to observe all (encrypted) messages at all times in

[1] The TI CC2420 [30], for example, supports 802.15.4 wireless communications and provide a hardware-based implementation of Advanced Encryption Standard (AES) with support for both encryption and authentication.

[2] For a good overview of pairing methods, see Kumar et al. [87].

the wireless medium. That is, they can observe those messages being sent between the MN and SNs. An adversary can use this capability, for example, to identify the types of SNs and MN in the system. We also assume the adversary has the ability to arbitrarily inject, modify, and discard messages in the wireless medium. Thus, by disrupting a message and later re-injecting that message, they can also arbitrarily delay or replay the message. An adversary can use these capabilities to track users while they use the system or to tamper with the data in the system.

We also assume the following limitations about any adversary. First, we assume an adversary is computationally bounded, meaning that it cannot break the cryptographic primitives without using a brute-force attack. Second, we assume an adversary will not access or modify the hardware nor modify the software of the MN or SN in our system. Third, we assume an adversary will not be able to acquire a subject's biometric. Like a password, each subject has an incentive to make sure their biometrics will not be divulged to an adversary. Finally, we assume an adversary will not jam the wireless medium; that is, an adversary will not try to disrupt the usage of the system by performing denial-of-service attacks.

2.3.4 Security goals

The goal of our system is to allow users to wear an MN and SNs that will collect data about their self and their environment and forward this data to a service provider for analysis according to the trust assumptions. An adversary seeks to disrupt this process and obtain sensitive data according to the threat model outlined above. Despite such an adversary, our system should still achieve the following goals:

G1 Our system should preserve the confidentiality of the sensed data and meta-data. That is, the sensed data and its meta-data should not be revealed to anyone but the user and service provider.

G2 Our system should preserve the integrity of the sensed data and meta-data,

meaning that the service provider and user should be able to trust that the sensed data or meta-data has not been tampered with.

G3 Our system should preserve the authenticity of the sensed data and meta-data. That is, the service provider should be able to trust that the sensed data and meta-data was acquired from the specified user and the specified SN.

G4 Our system should preserve the obscurity of the sensors, meaning that the type of SN and MN the user is wearing should not be revealed to anyone but the user and the service provider.

We provide an analysis of how our system achieves these goals in Chapter 5.

2.4 Machine learning

Machine learning is the broad field of using computers to learn representations of data. It encompasses many techniques and continues to be an open area of research. Additionally, it has proven successful at solving many real-world problems, including natural language processing, speech recognition, handwriting recognition, and others.

In this thesis, we formulate many of the challenges as pattern-recognition problems, and, in particular, we construct them as classification problems. *Classification* is the process of determining which class a piece of data belongs to. That is, we want to assign a label to a piece of data. For example, given biometric data, we would like to know from which user that biometric data was collected; in this case the set of labels (classes) is the set of user identities (users).

Because we assign these labels, this task is called *supervised learning*. The process works by first collecting some raw data, and then assigning labels to them to create a *dataset*. Instead of using raw data itself, one usually transforms the raw data into a *feature vector*. A feature vector represents *features* of the raw data (like

the mean). Given such a dataset, we can train a *classifier*. The classifier's job is to generate a *model* that is capable of assigning the correct label to each feature vector in the dataset. One can view this task as learning a function that outputs some value (i.e., label) for some input (i.e., feature vector). This process is called *training* and when a dataset is used for training we call it a *training dataset*. Once a classifier learns a model, then we can use the model to classify some new feature vector (computed from some new sample) into one of the labels in the training dataset the model was trained on. This process is called *testing* and a dataset used for testing is called a *testing dataset*.

2.4.1 Classifiers

There are many different classifiers one can use to solve supervised learning tasks. In this thesis we use four popular classifiers that span the gamut of possible functions one can learn.

k-Nearest Neighbors

The first classifier, and perhaps the simplest to understand, is the k-Nearest Neighbors (KNN) classifier. To train a k-Nearest Neighbors (KNN) classifier, no computation is necessary. Instead, the training dataset is itself the model. At test time, the k-nearest data points to the new sample are found (typically using some specified metric like Euclidean distance) and the model outputs the label equal to the most frequent label (i.e., majority vote) of those k neighbors. It is also possible to apply other statistics than majority vote. If, for example, the distribution of labels is skewed in the training dataset, then voting can be weighted according to the distribution of labels.

Naive Bayes

The second classifier, Naive Bayes (NB) [161], takes a probabilistic view of classification by assuming that a label C can be conditionally modeled by the features F_i of the training feature vector:

$$p(C|F_0, \ldots, F_n)$$

Given this probability, we can label some new feature vector f by computing the posterior probability $p(C = c|F_i = f_i)$ for each label c. The model chooses the label c that maximizes the posterior probability:

$$\arg\max_c p(C = c|F_i = f_i)$$

The remaining question is: how does one compute the posterior probability $p(C|F_0, \ldots, F_n)$? According to Bayes' rule, the posterior probability is equivalent to:

$$p(C|F_0, \ldots, F_n) = \frac{p(F_0, \ldots, F_n|C)p(C)}{p(F_0, \ldots, F_n)}$$

We can compute $p(C)$ easily because we know the distribution of each label in the training dataset. We also know that $p(F_0, \ldots, F_n)$ is a constant since the feature vector is also known. However, it is not clear how to compute $p(F_0, \ldots, F_n|C)$. To do so, we make the simplifying (naive) assumption that features of the feature vector are independent:

$$p(F_0, \ldots, F_n|C) = p(F_0|C) * p(F_1|C) * \ldots * p(F_n|C)$$

Under this assumption, it is now easy to compute $p(C|F_0, \ldots, F_N)$.

The assumption is that given the label c, each feature F_i of the feature vector can be modeled by some distribution $p(F_i|C = c)$. For continuously valued features,

the chosen distribution is often the Gaussian distribution since it is easy to compute and many values tend to be normally distributed. Since a Gaussian distribution is described by a mean and variance only, the NB classifier computes these values for each feature across all training feature vectors. In addition to the mean and variance of each feature, the NB classifiers also computes the prior probability of each label in the training dataset. These values are the model. As test time, the NB classifier classifies a test feature vector as the label that maximizes the posterior probability of the model.

Random Forest

A Random Forest (FOREST) [27] is a collection of Decision Trees (hence, forest) trained on different subsets of the dataset. For classification, a Random Forest outputs the mode of the outputs of the decision trees. The use of many classifiers (or, more generally, models) is called *ensemble learning* and they are the state-of-the-art method for many tasks [39].

A Decision Tree can be thought of as tree of simple `if-then-else` statements. That is, starting at the root, each node of the tree uses one particular feature of the input and decides to traverse the left or right child node based on the value of that feature. The leaves of the Decision Tree represent outputs (or classifications, in the case of classification trees) corresponding to those labels. Thus, a path from the root node to a leaf can be thought of as a conjunction that represents one of the labels. Other paths represent the same or different label. Given a test data, each sample evaluates the root node and recursively evaluates the appropriate child node until a leaf is reached. The output of the classification is the label of the leaf node.

Support Vector Machine

A Support Vector Machine (SVM) [38] is linear, binary classifier. An SVM classifies

samples by finding a hyper-plane with the largest separation between the set of positive training feature vectors and negative training feature vectors (i.e., the maximum-margin hyper-plane). Mathematically, this hyper-plane can be described using the vectors that lie exactly on the margin of the hyper-plane (i.e., those vectors that "support" the hyper-plane). Thus, these support vectors are the model. Given this model, an SVM can classify a test feature vector by determining on which side of the hyper-plane the test feature vector lies on and outputting the associated label.

Even if the training feature vectors are not linearly separable, we can use the "kernel trick" to map the examples into a higher dimensional space where they might become linearly separable [26]. An oft-used example of such a kernel is the Gaussian radial-basis function.

Random

Finally, we compare each classifier to the Random (RAND) classifier as a baseline to see how well the other classifiers are performing. It classifies test feature vectors by randomly choosing a label weighted according to the occurrence of that label in the training dataset. Thus, this serves merely as a baseline classifier, and, ideally, all other classifiers should do at least as well as the RAND classifier.

2.4.2 Statistical measures

The problems we examine can be formulated as binary classification problems. This means, that there are exactly two labels the classifier can output: positive or negative (the semantics of the label are left to the specific problem). As such, we define the statistical measures by which we will evaluate these classifiers according to these labels.

We define a *true feature vector* as a feature vector in the training set that has a positive label, and a *false feature vector* as a feature vector in the training set that has

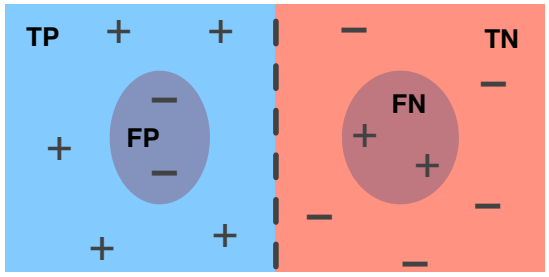

Figure 2.4.1: The definition of different classification results: TP, FP, TN, FN. Each sample is represented by a $+$ and $-$. The left side represent a positive classification, while the right side represents a negative classification. Thus, $-$ sample on the positive side represents a false positive, while a $+$ on the negative side represents a false negative.

a negative label. A *positive classification* means the classifier assigned a positive label to a test feature vector, while a *negative classification* means the classifier assigned a negative label to a test feature vector. It follows, then, that a True Positive (TP) occurs when a true feature vector is classified as positive, and a True Negative (TN) occurs when a false feature vector is classified as a negative. A False Positive (FP) occurs when a false feature vector is classified as positive, and a False Negative (FN) occurs when a true feature vector is classified as negative. Figure 2.4.1 shows this graphically.

Given such classification results, there are many statistical measures one can compute. Table 2.4.1 summarizes the most common types of statistical measures one can compute. Since different fields (e.g., biometrics, information retrieval, medicine) tend to use different nomenclature, several measures go by multiple different names.

We use the following statistical measures to evaluate the performance of each classifier:

Balanced Accuracy (BAC) tells us how well the classifier is performing by weighting the negative and positive examples equally, since sometimes there are more negatively labeled feature vectors than positively labeled feature

Measure	Formula	Also Known As
True Positive Rate	TP/(TP+FN)	recall, sensitivity, true accept rate
True Negative Rate	TN/(TN+FP)	specificity, true reject rate
Pos. Predictive Value	TP/(TP+FP)	precision
Neg. Predictive Value	TN/(TN+FN)	
False Negative Rate	FN/(TP+FN)	false reject rate, false non-match rate
False Positive Rate	FP/(TN+FP)	false accept rate, false match rate
False Discovery Rate	FP/(TP+FP)	
False Omission Rate	FN/(TN+FN)	
Accuracy	(TP+TN)/(TP+TN+FP+FN)	
Balanced Accuracy	1/2*(TP/(TP+FN) + TN/(TN+FP))	

Table 2.4.1: List of common binary classification statistical measures, their formulas, and what they are also known as.

vectors in our datasets. A perfect classifier performs at 100 %.

Precision tells us what portion of positively predicted feature vectors were correctly classified. A perfect classifier performs at 100 %.

Recall tells us what portion of positively labeled feature vectors does the classifier classify correctly. A perfect classifier performs at 100 %.

False Accept Rate (FAR) tells us what portion of negatively labeled feature vectors does the classifier mis-classify. A perfect classifier performs at 0 %.

False Reject Rate (FRR) tells us what portion of positively labeled feature vectors does the classifier mis-classify. A perfect classifier performs at 0 %.

2.4.3 Validations

Given a dataset, it is useful to validate how well it performs under the given metrics. We performed three types of validations to evaluate the performance of our method.

k-fold cross-validation

A *k-fold cross-validation* divides the dataset into k subsets where each subset contains the same proportion of samples according to the distribution labels for the entire dataset. We then train a classifier over $k-1$ of the subsets (the training set) and classify the remaining subset (the testing set). We repeat this procedure for each subset. This type of cross-validation tells us how well our classifier generally performs since it will classify every sample in the dataset.

Leave-one-out cross-validation

Rather than split the dataset into random subsets, a *leave-one-out cross-validation* splits the dataset according to some rule. For example, a *leave-one-sample-out cross-validation* would divide the dataset into n subsets, where n is the number of samples in the dataset. That is, we train a classifier using $n-1$ samples, and test it on the left-out sample. This would be repeated for each sample. This is equivalent to an n-fold cross-validation. There are other rules one could use to split the dataset.

A *leave-one-subject-out cross-validation* leaves an entire subject's samples out as the testing set and trains the classifier using the remaining samples. We then test the classifier using the left-out subject's samples, repeating this procedure for each subject. This type of cross-validation will tell us how general our classifier is. Ideally our classifier would not be subject-specific yet do well in the case of a never-before-seen subject.

A *leave-one-location-out cross-validation* leaves out all of the samples from a

specific location as the testing set, and trains the classifier using the remaining samples. We then test the classifier using the left-out samples from the chosen location, repeating this procedure for each location. Like the leave-one-subject-out cross-validation, this cross-validation tells us how sensitive our classifier is to location. In the ideal case, our classifier would not be location-specific yet do well for each of the locations we specify.

Hold-out validation

A *hold-out validation*, holds out a proportion from the dataset as the testing dataset. That is, we train a classifier on the first N samples and leave the remaining samples for testing. As the number of training samples increases, we expect the performance of the classifier to improve since the classifier has more training samples to learn from.

One can think of this kind of validation as a 2-fold cross-validation, however it is somewhat different. First, the portion of training data versus testing data can be different and not just half and half. More over, the training datasets and testing datasets are not randomly chosen; rather, the first N samples are chosen for training and the remaining are left for testing. This is important when time can be a factor (e.g., the data is a time series) because it means we have no information from the future, so to speak. As such, a hold-out validation more accurately simulates the performance of a method when time is a factor.

3

Recognizing wearers using bioimpedance

We recognize people by learning a tell-tale characteristic about them. At some later time, we can use this characteristic to determine whether that same person is present. We call such characteristics *biometrics* [24]. Today, many systems use a wide variety of biometrics to recognize individuals. A biometric, for example, can be used as an authentication mechanism to control access to a particular resource.

3.1 Introduction

In this chapter, we focus on a fundamental problem involving WBANs: who is wearing the MN and SNs? The ability to recognize who is interacting with a device is fundamental to many applications. For an entertainment device, it can recognize the user and load the right game profile or music play list. 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 SN's data with the correct identity so that it can be stored in the correct health record. (A mix-up of SN's data could lead to incorrect treatment or diagnosis, with serious harm to the patient.)

This problem is a form of the *one-body verification problem* [36], which asks: how can one ensure that the SNs in a WBAN are collecting data about one individual and not several individuals? This chapter deals with the *strong* version of this problem, which requires identifying the specific individual these SNs are attached to. Chapter 4 deals with the *weak* version of this problem.

To solve this problem, we present a biometric-based solution. In a biometric-based solution, an SN or an MN collects some identifying information about the wearer. It will then learn a model of that biometric so that it can use the model to recognize that user at some later time. What we want is a simple, wearable device that uses biometric techniques to recognize the user, then share that identity with the WBAN. This device should be trained once, for each user that might wear it, but thenceforth be completely automatic and unobtrusive.

3.2 Biometrics

Attaching an identity to sensor data requires some method of recognizing who the device is sensing. To accomplish recognition, we use a biometric. In this section we

describe what is a biometric, the kinds of biometrics that exist, and how they can be used to recognize people.

3.2.1 Characteristics

According to Jain et al. [78], to qualify as a biometric, the chosen characteristic must have the following properties: universality, measurability, uniqueness, permanence, performance, circumvention, and acceptability. We adopt this characterization. A characteristic must be *measurable*, meaning that there is some sensor capable of capturing the characteristic for processing. Next, a characteristic must be *universal*, meaning that most people have it. Likewise, although everyone may have such a characteristic, the characteristic must also be individually *unique* within a given population. For personal health sensors and other pervasive applications, *performance* is the ability to unobtrusively measure the characteristic while maintaining good measures of success. This ability to unobtrusively measure a biometric stems from our desire to provide usable security for personal health sensing systems. The characteristic must also have some *permanence* such that it does not vary over the relevant time scale. Likewise, a biometric needs to be *difficult to circumvent* because there are incentives for people to circumvent them. Finally, the biometric should be *acceptable*, meaning that the target population is both willing to use the sensor necessary to acquire the characteristic and allow that characteristic to be used for recognition purposes.

3.2.2 Categorization

Physiological biometrics use some characteristic of your physiology to identify you. These tend to be the biometrics people are most familiar with: fingerprint, hand geometry, facial recognition, iris or retina 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 makeup of DNA, or the structure of the iris. These type of biometrics typically require a sensor attached to the subject or require the subject to place a limb on a sensor.

On the other hand, *behavioral* biometrics use some characteristic of your behavior to identify you. 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. In contrast to a physiological biometric, behavioral biometrics can exhibit wide within-subject variation since they are sensitive to things like mood. Likewise, they also tend to be easier to collect since they generally do not require the subject to be interrupted. For an overview of behavioral biometrics, see Yampolskiy et al. [159].

These categories are not mutually exclusive. For example, a voice-recognition biometric has both physiological and behavioral aspects. Your voice is shaped by your vocal tract (the larynx, pharynx, oral and nasal cavities), however your current behavior can also affect your voice. For example, your current state of mind (e.g., being excited or nervous) can alter your vocal tract and therefore your voice.

Unfortunately, many of the 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 subjects 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 subject's gait, or a malicious application could learn one's typing rhythm [98]. In contrast, recognition for wearable sensing applications demands a biometric that is simultaneously difficult to circumvent and easy to measure continuously. We call such a biometric a *passive* biometric.

3.2.3 Recognition

The usefulness of biometrics relates to their ability to be used to recognize a person for some population.[1] The size of the population is important since some sensors, while seemingly unfit for distinguishing large populations, maybe be able to distinguish smaller populations. For example, Srinivasan et al. [136] 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. Given a population, biometrics can be used in one of two ways: *identification* and *verification*.

Identification

Identification is a one-to-many matching. First, the system collects many biometric samples from a population. The combination of the biometric samples and the subject's identity forms a *biometric template* for each subject. The system can then use these biometric templates to determine the identity of any subject in that population. That is, some unknown subject would present themselves to the system, and it is the system's job to determine which subject that is from the population. A biometric system accomplishes by first measuring the unknown subject's biometric. It then examines its database of biometric templates, and finds the biometric template that best matches that subject's biometric sample. The system then asserts that the identity of the unknown subject is the identity of the biometric template that best matches that unknown subject's biometric sample.

[1] The target population is especially important in forensic science [35]. 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 [152].

Verification

Verification, on the other hand, is a one-to-one matching. That is, a subject presents an identity (e.g., a name), and it is the job of the biometric system to verify that identity. First, the system collects a biometric template from the subject and stores it. At some later time, a subject presents themselves to the system and asserts their identity. The system would then retrieve from the database the biometric template associated with that identity. It would then measure that subject's biometric, and try to match that subject's biometric sample with the biometric template. If the biometric sample matches the biometric template (for some matching metric), then the asserted identity is verified. For example, one might use a generative model to learn the distribution of a subject's biometric and select a threshold to accept a new biometric sample according to likelihood. Notice that there is no assumed population; rather, it should reject everyone else in the world.

3.3 Bioimpedance

Our approach is to use *bioimpedance*, which measures the response of a living tissue when exposed to an electrical current. In our wearable device, as seen in Figure 3.3.1, eight electrodes are placed in contact with the wrist. Using two of these electrodes, the device applies a small current to the body and then measures the impedance of the tissue. Comparing these bioimpedance measurements with a model, built earlier during a training phase, allows the device to determine that the person being measured is indeed the person for whom we trained the model. If we train the device for a set of users, e.g., the members of a household, then the device should be able to recognize which of those people is wearing the device, or that none of them are wearing the device.

Figure 3.3.1: 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).

Such solutions have many advantages. Not all wearable devices need have 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 personal identification number or password for authentication). The use of a biometric provides important security and privacy properties. They can prevent unauthorized users from accessing sensitive data (e.g., in which an adversary Alice tricks Bob's sensor into divulging his activity data to her smart phone), or prevent the mis-labeling of sensor data that might later be used for medically important decisions. Privacy is particularly important in health-related pervasive applications [13]. Furthermore, these methods can support personalization techniques so often envisioned in pervasive computing.

3.3.1 Theory

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 being measured, 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 state, each of which may provide a unique mechanism for distinguishing among people.

We measure bioimpedance by applying a small sinusoidal current between a pair of electrodes attached to the skin. The injected current creates 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 impedance Z of the tissue sample. Many tissues show 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 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[2] $\sqrt{-1}$. Because two electrodes apply current and two electrodes measure voltage, there are two modes of sensing. The first mode, called *bi-polar sensing*, occurs when the voltage-measuring electrodes also apply current, which only requires two electrodes. The second mode, called *tetra-polar sensing*, occurs when the two voltage measuring

[2]We use j, instead of i, so as to not confuse it with I for current.

electrodes are distinct from the current-applying electrodes, which requires four 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 [126]. However, as a result, tetra-polar measurements require complex hardware for making accurate impedance readings.

Both tissue (physiologic) conditions (i.e., pH, temperature, and fluid flow) and geometry will influence impedance measurements. Bioimpedance is dependent upon the tissue being measured along with the configuration and geometry of the impedance-measuring probe (i.e., electrode size and electrode spacing). 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). Geometrically, in the most generic case of a parallel plate geometry in which a tissue sample is placed between electrodes of area A and spaced a distance d apart, the resistance is proportional to $d/(\sigma A)$ where σ is a geometry-independent tissue conductivity.

3.3.2 Location of interest

To explore the possibility of using bioimpedance as a biometric device, we identified the wrist as the most viable anatomic location for which a wearable device could be realized. We constructed a wrist wearable device with eight Ag/AgCl electrodes (8 mm diameter) embedded within an elastic strap. The eight electrodes are interfaced through a set of 8×1 multiplexors to a custom-built impedance analyzer [63] such that current source and sink and voltage sensing pins can be connected to any of the 8 electrodes. This configuration permits impedance measurements to be recorded between any pair of electrodes as seen in Figure 3.3.2. The anatomy of the forearm proximal to the wrist includes skeletal bones (radius and ulna), arteries,

veins, nerves, muscles, adipose, skin, and interstitial fluids. Over the frequency range of 10 kHz to 10 MHz reported values of bone conductivity and adipose conductivity are relatively stable at 0.1 Siemens per meter (S/m) and 0.01 (S/m). In muscle, skin, and blood, however, the conductivity monotonically increases with frequency with reported values ranging from 0.3 S/m to 0.5 S/m, 0.001 S/m to 0.5 S/m, and 0.7 S/m to 1.0 S/m [55]. While these numbers represent the values from a single sample, "biological tissues are inhomogeneous and show considerable variability in structure or composition and hence in dielectric properties" [55].

We can sense these different parts of the wrist anatomy by using multiple electrode locations because different combinations of electrode locations will measure impedance across different pathways. 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 configurations of electrodes, a list of impedances associated with an individuals' wrists can be recorded and ultimately used for recognizing an individual within a group of individuals. Figure 3.3.2 shows the anatomy of a human wrist with several electrodes and paths where the current can travel.

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 content, muscle volume, bone size, and vasculature within the sensing region. All of these parameters affect the impedance measured at the wrist. For example, a difference in wrist size would represent a change in electrode location, while a difference in the content and distribution of the underlying tissue types would represent a person-specific conductivity.

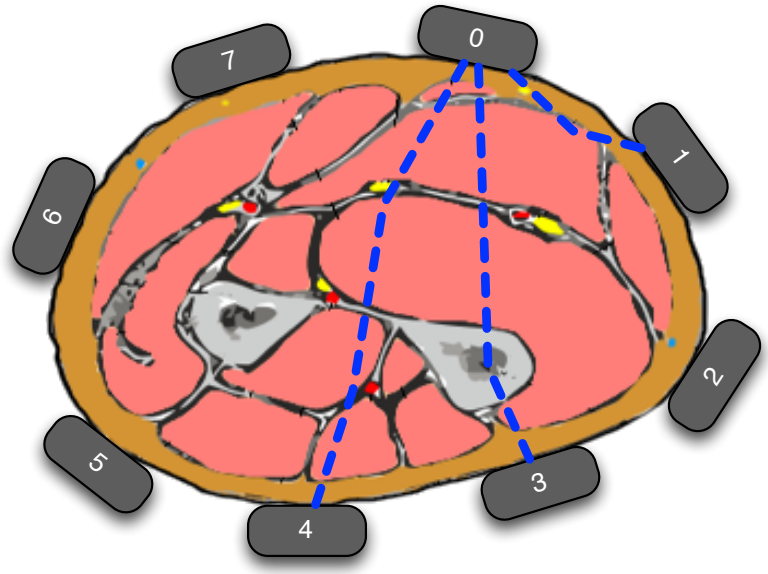

Figure 3.3.2: The anatomy of a human wrist [60]. We colored the bones (ulna and radius) in gray, the muscle in red, and the skin in brown. The gray boxes on the outside of the anatomy represent the placement of electrodes with corresponding labels. If electrode `0` were chosen as the current applying 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 structures.

3.4 Recognizing wearers

Recall that the goal of our device is to recognize who is wearing it. The intuition is quite simple: gather bioimpedance measurements from a person and build a model to represent that person's bioimpedance signature; later, use that model to determine whether a new and unknown bioimpedance measurements match. Thus, there are two specific phases: enrollment and recognition. Before we discuss the details of the enrollment and recognition phases, we first discuss the details of the bioimpedance measurements and features we use in those phases.

3.4.1 Bioimpedance measurements

Due to collection time, it would be infeasible to measure bioimpedance from all combinations of the eight electrodes. Instead, we carefully chose specific electrode configurations. We captured measurements from two types of electrode configurations: those electrodes directly across from one another (e.g., electrodes `0` and `4` as shown in Figure 3.3.2), since they are the maximal distance away from each other and therefore provide more tissue for the current to travel through; and those electrodes that are exactly one electrode apart (e.g., electrodes `0` and `2`), since the current will 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, which 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 in our preliminary tests those measurements exhibited characteristics similar to those collected from electrodes maximally apart. Figure 3.4.1 shows an example bioimpedance measurement from a single subject for the different electrode

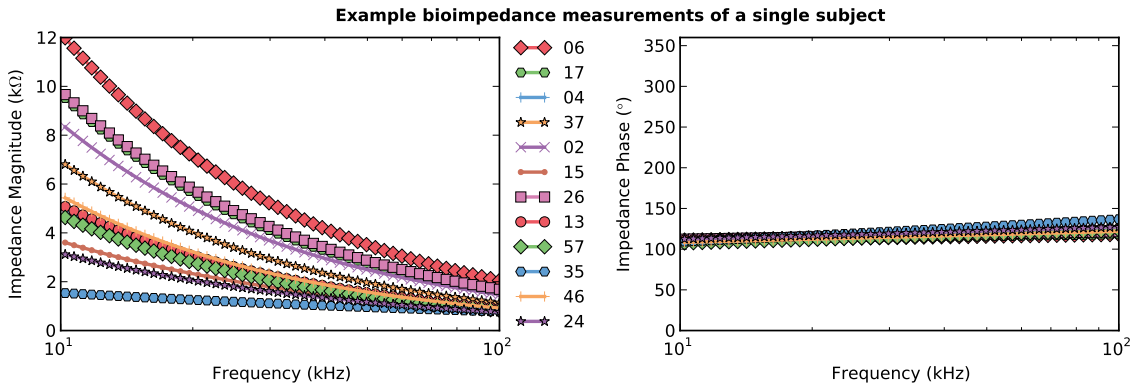

Figure 3.4.1: Example bioimpedance measurements collected from a single subject for the 12 different electrode configurations. Note: frequency is plotted on a logarithmic scale.

configurations.

3.4.2 Feature extraction

Given a set of frequencies and their corresponding bioimpedance measurements, we extract four features from each bioimpedance measurement: two from the impedance magnitude and two from the impedance phase. The features are simple: we fit one line to the impedance magnitude and another line to the impedance phase both in log-log space. The inspiration for this feature can be seen in Figure 3.4.1, which shows the impedance magnitude and phase of a human wrist. In log-log space, they exhibit a linear relationship. Because measurements are inherently noisy, a line smooths over the individual measurements while also preserving the general shape of the curve formed by the measurements. Because each line is succinctly described by a slope and intercept there are four such features (two for each of the magnitude and phase components). We found these properties to be mostly unique among individuals. Additionally, these features reduce the dimensionality of the data from 100 (i.e., the number of frequencies) to 4, which in turns lowers the computational and energy overhead. Although we explored other features, including using the raw data itself, we found these features performed the best according to our metrics.

Our final feature vector consists of the concatenation of these features for each electrode configuration. Since we take measurements from 12 electrode configurations, this results in a feature vector of dimension $12 \times 4 = 48$, which is much smaller than the raw data (which is of dimension $12 \times 50 \times 2 = 1200$). This concatenation assumes all electrode configurations will provide some information about the identity of the wearer; however, this may not always be the case. We explore different combinations of electrode configurations that perform best according to our metrics.

Because many of our studies took place outside of the lab, we had to discard some bioimpedance samples. We discarded all samples where we did not collect a full bioimpedance measurement at each electrode configuration and frequency. There were only 5 such instances of incomplete bioimpedance measurements and the reason for this is unknown. In a real system, incomplete bioimpedance measurements could be detected and the measurement could be retaken. We also discarded all samples where the wrist was deemed to not be in contact with the device. We deemed a sample to be non-contact if the maximum impedance magnitude was greater than an empirically determined threshold of $10 \times 10^6\,\Omega$ as the non-contact impedance magnitude shows in Figure 3.4.2. 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, due to motion or other interference with the reading.

3.4.3 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 for many applications of interest. This means that the cohort of subjects in which the

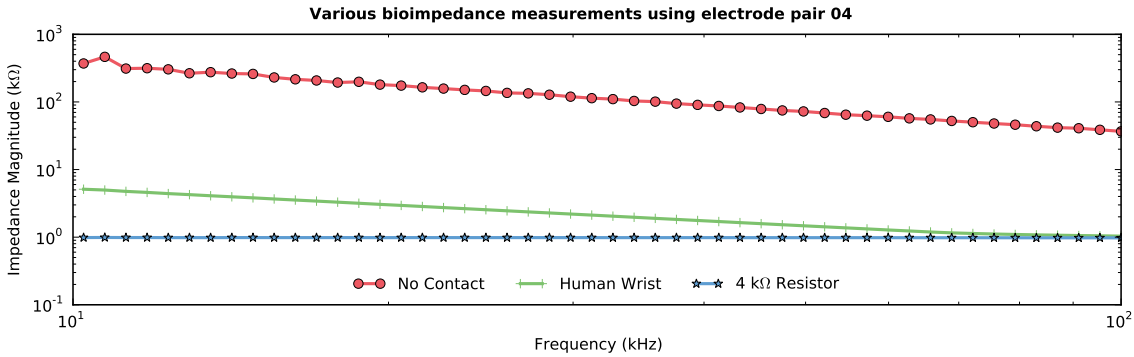

Figure 3.4.2: Example impedance measurements from electrode configuration `04` for 1) no contact, 2) bracelet in contact with a human wrist, and 3) bracelet attached to 4 kΩ resistor. In the case of no contact, impedance is effectively infinite. Note: impedance and frequency are plotted on logarithmic scales.

device will be used is relatively small (typically 2 to 8 subjects).

3.4.4 Enrollment mode

Before a person can use the device daily, they must train it to recognize their bioimpedance by putting the device into *enrollment mode*. In this mode the device captures bioimpedance measurements for some designated time (a single bioimpedance measurement takes about 15 seconds). 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.

Given a set of *training feature vectors* from a subject, we learn a model of their bioimpedance measurements using the 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 subjects' feature vectors as negative examples (i.e., they are labeled negatively). For a device being used by multiple people, we load it with multiple models and use a one-versus-all strategy for multi-class classification [124], but we limit our analysis to a single enrolled user.

We explore how the amount of training data affects our performance metrics, since typically more samples means better performance.

We examined three different classifiers for use as an enrollment algorithm: NB, SVM, and FOREST. See Section 2.4 for details about these classifiers.

3.4.5 Recognition mode

Once a user is enrolled, the device enters *recognition mode*. In recognition mode, the device periodically collects bioimpedance measurements. The device uses an *identification algorithm* or a *verification algorithm* to determine whether the enrollee's model matches the measured bioimpedance.

Given a set of *test feature vectors* from an unspecified subject, we can use the enrolled subject's trained model to determine 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.

3.4.6 Evaluation parameters

The FOREST and 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 each classifier. In the case of the FOREST classifier, we looked at different settings of the number of trees in the forest and found that 10 worked best. 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 [26]. We used a grid-search to search this parameter space and found that a 3rd degree polynomial kernel with a cost of 32 and gamma of

0.03125 worked best. During this parameter search, the performance of the SVM classifier tended to be better than the KNN classifier, so we excluded the KNN from any further consideration.

3.4.7 Evaluation metrics

Consider a set of test feature vectors from a given subject and a set of test feature vectors from other subjects. We label the test vectors measured from that subject as *positive* and all other test vectors as *negative*. We then use the model trained for that subject to classify all the test feature vectors, resulting in a positive or negative classification for each. Ideally, the model classifies only those test feature vectors from the subject as positive and all others as negative. Given such classification results, we evaluated the performance of these classifiers using the BAC, FAR, and FRR metrics described in Section 2.4.

Although we computed these metrics for every subject in our datasets, we present summary statistics of these metrics over all subjects. Thus any mentions of BAC, FAR, and FRR should be interpreted as the average BAC, average FAR, and average FRR over all subjects. Note that because the number of positive samples for any given subject is 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 using the RAND classifier. This classifier serves as a baseline performance measure to compare with our methods that we present.

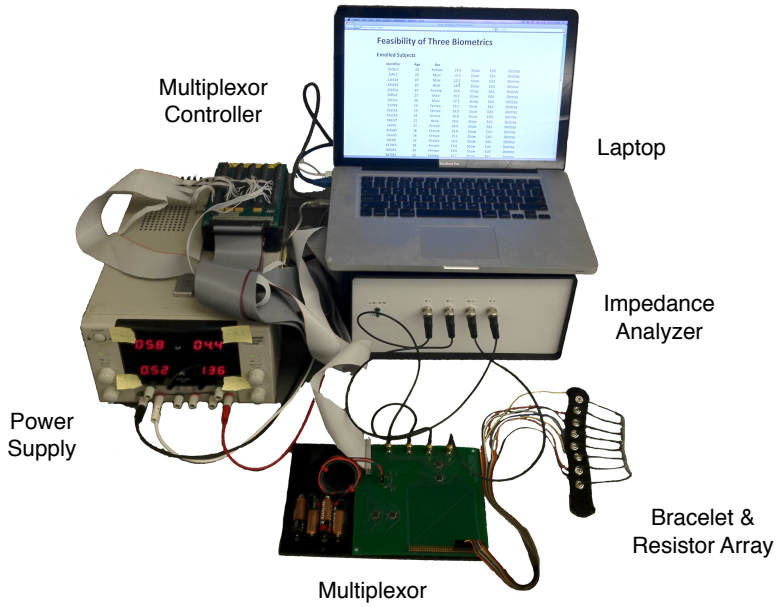

Figure 3.5.1: Our bench-top bioimpedance system. The bracelet has a Velcro fastener to hold it in place during bioimpedance measurements. Here, the bracelet is attached to a resistor array we use for calibration.

3.5 Measurability

Recall that a passive biometric has the following characteristics: universality, uniqueness, permanence, unobtrusively measurable, and difficult to circumvent. In this section, we describe the hardware used to measure bioimpedance.

3.5.1 Unwearable device

To collect a large dataset over many frequencies and electrode configurations, we used a custom-designed impedance analyzer constructed specifically to record *in vivo* bioimpedance measurements [64]. Figure 3.5.1 shows this bench-top system.

Hardware

The impedance analyzer is designed around a 32-bit ADSP-21065L digital signal processor [7] and a Spartan XC2S30 field programmable gate array [158]. These digital processing devices interface with a wide-band (up to 300 MHz with 1 μHz resolution) AD9852 direct digital synthesizer [5] used to drive digital signals through a 14-bit AD9754 digital-to-analog converter [4] and generate sinusoidal voltages over wide range of frequencies (10 kHz to 10 MHz). A front-end analogue amplification stage drives these voltages through a 100 Ω current sensing resistor to a channel output (and ultimately to an electrode in contact with the skin) and 16-bit AD7677 analogue-to-voltage converters [3] are used to gauge both voltages and currents from the system through which impedance can be calculated. This impedance analyzer generates output voltage signals up to 2 V peak-to-peak and maximum currents of 10 mA, has a bandwidth of 10 MHz, a signal-to-noise ratio greater than 94 dB up to 2 MHz, 90 dB up to 7 MHz, and 65 dB at 10 MHz, an accuracy of 99.7 %. It is able to record both bi-polar and tetra-polar impedance sweeps from 20 logarithmically spaced frequencies ranging from 10 kHz to 10 MHz in 10 seconds. A Topward 6303D digital power supply [1] provides power to the impedance analyzer.

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 [45] 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, our system can take two kinds of measurements. 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. By switching through multiple configurations 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.

Software

We wrote custom software in Ruby [127] to control the system. For each desired electrode configuration, we ran the following sequence. First, the multiplexer is instructed to select the correct electrode configuration. Next, the impedance analyzer is instructed to sweep through the desired frequency range. Once complete, we compute the impedance from the returned data over all frequencies and save the impedance to a file. We repeat this sequence to acquire five measurements per electrode configuration per subject.

3.5.2 Wearable device

We imagine the device to be a piece of jewelry, not unlike a watch or bracelet, 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 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. Such simple detection mechanisms also allow us to conserve energy by only performing recognition when the device is in contact with a person.

To explore this vision, we designed and manufactured a wearable sensor to measure the bioimpedance of a person's wrist. The wearable sensor is built on top of the Shimmer Platform [132] and uses a custom bioimpedance sensor module we designed. Shimmer is an open-source, low-power, wireless sensing platform. It provides processing (via a MSP430 microcontroller unit), wireless communication (via Bluetooth or 802.15.4), and storage (via a Secure Digital (SD) card) capabilities and includes a simple sensor in the form of an accelerometer. It also provides an internal and external expansion connector that allows it to interact with custom sensor modules (e.g., a gyroscope, magnetometer, Electrocardiography (ECG), Electromyography (EMG), or Electro-dermal Response (EDR)). Our custom bioimpedance sensor module (Figure 3.5.2) uses this internal expansion connector, both enclosed by a custom-designed case. The bioimpedance sensor module includes a receptacle that enables a series of electrodes to be connected to it. As such, we designed and manufactured an elastic sleeve with eight evenly spaced electrodes. The sleeve connects to the bioimpedance sensor module and includes a pocket that holds the Shimmer. Figures 3.3.1 and 3.5.3 show the final form factor.

The bioimpedance sensor module (Figure 3.5.2) is designed around the Analog Devices AD5933 Impedance Analyzer [2], which allows us to do bi-polar bioimpedance sensing. The AD5933 includes a frequency generator that allows

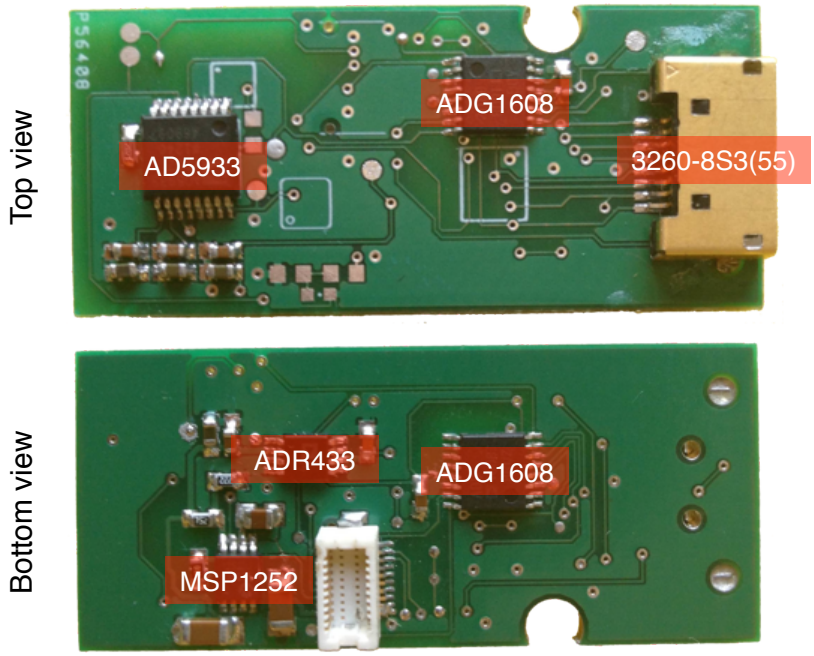

Figure 3.5.2: Our custom designed bioimpedance sensor module with the major components labeled. It is approximately $45\,\mathrm{mm}$ long by $19\,\mathrm{mm}$ wide and fits comfortably on top of a Shimmer with a custom-designed enclosure.

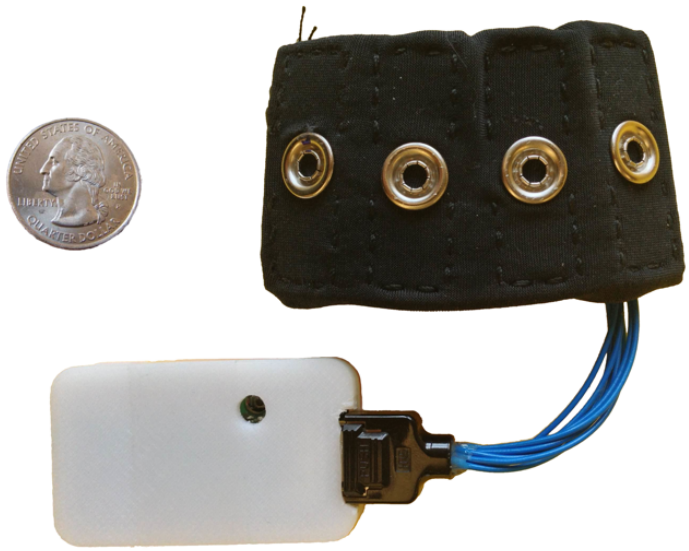

Figure 3.5.3: Our wearable bioimpedance sensor. The sleeve (top) is inside out to display the electrodes. The Shimmer (bottom) and wires are typically housed within the sleeve, but are shown exposed here. For reference, the Shimmer is approximately $53\,\mathrm{mm}$ long by $32\,\mathrm{mm}$ wide.

the excitation at a specified frequency between 1 kHz and 100 kHz with a resolution of 0.1 Hz, and it can measure impedances between 1 kΩ and 100 MΩ to within 0.5 % total system accuracy. It also includes an internal temperature sensor capable of sensing between −40 °C to 125 °C (±2.0 °C), since error in impedance measurements can be on the order of 30 ppm/°C. The Shimmer controls the AD5933 via the I^2C bus. Because we want to allow multiple electrode locations, the sensor module includes two Analog Devices ADG1608 8-Channel Multiplexors [6]. Thus, the sensor module is capable of selecting 2 of 8 possible electrodes for bi-polar sensing. We call such a selection an electrode configuration. These multiplexors are controlled by setting specific general purpose input/output pins. Electrodes are connected to the sensor module via an 8-pin Hirose 3260-8S3(55) Connector [70] such that custom electrode configurations can be independently built and interfaced with the sensor module. The Shimmer provides regulated power to the sensor module, which is fed to an Analog Devices ADR433 ultra-low noise voltage reference. This ADR433 [10] provides a stable 3 V supply voltage needed for the impedance analyzer while a Microchip MCP1252 charge pump [110] feeds a 5 V supply voltage to the multiplexors.

Software

The Shimmers run the TinyOS operating system [140]. We wrote custom software to communicate with the impedance analyzer and multiplexors. The software is divided into three major parts: a low-level driver, a high-level driver, and a logging application.

The low-level and high-level drivers allow applications to communicate with our sensor module. The low-level driver is a barebones interface to the AD5933 that wraps the I^2C communications. The high-level driver implements a state machine that allows more natural interaction with the sensor module. It also allows applications

to adjust the settling time (i.e., the duration between the stimulus and measurement) and handles failures gracefully. In total, the low-level driver compromises 381 lines of nesC code while the high-level driver compromises 363 lines of nesC code.

The logging application uses the high-level driver to interact with the sensor module. Before taking bioimpedance measurements, it samples the accelerometer for 5 seconds at 50 Hz to classify the type of motion (i.e., low energy, medium energy, high energy) the subject's wrist is experiencing. Afterwards, the impedance analyzer is commanded to measure temperature using its internal temperature sensor. This measurement allows us to account for variable temperatures since the input amplifier gain varies with temperature. Next, two electrodes on the wrist strap are selected using the multiplexors, both to be used as stimulus and measurement electrodes. We then tell the impedance analyzer to take measurements at 50 logarithmically spaced frequencies from 1 kHz to 100 kHz. Due to the logarithmic spacing of the frequencies, the impedance analyzer's internal mechanism to step through frequencies linearly can not be used, so each measurement must be commanded individually. After one frequency sweep, another set of stimulus and measurement electrodes is selected and another frequency sweep is started. After a complete measurement, the impedance analyzer is sent to sleep to conserve power. We use a timer to wake the impedance analyzer after a preset interval and another measurement is started. Measurement data can either be stored to a MicroSD card or sent via the Bluetooth or 802.15.4 radio. In total, the logging application compromises 512 lines of nesC code.

In addition to the logging application, we also developed an application that allows collection of bioimpedance samples via Bluetooth. This proved useful for doing many of the in-the-lab studies. It is also possible to use this application to communicate with a smart phone or other Bluetooth-enabled mobile device, thereby making our system a truly mobile system.

Energy Measurements

Figure 3.5.4 shows how much energy our wearable device uses over the course of one full sample. When idle, the Shimmer draws 1.0 mA on average. There are two phases in the measurement. During the first phase, we sample the accelerometer, which draws an average current of 6.8 mA, lasts 6 seconds, and consumes 11 μA h of charge. During the second phase, we sample bioimpedance at 50 different frequencies for 12 different electrode configurations (i.e., 600 frequencies in total). On average it draws 32 mA of current, lasts 16 seconds, and consumes 140 μA h of charge. The spikes present in the second and final phases are a result of writing data to the MicroSD card. The full measurement consumes 151 μA h. Given that the Shimmer has a 450 mA h battery and we sample every 5 minutes, our wearable sensor can log bioimpedance measurements for about 1.5 days. This energy measurement represents our logging application. In a mobile system, the device would communicate wirelessly with some other mobile device like a smart phone. According to the Shimmer specifications, the Bluetooth radio consumes 20 mA "once paired regardless of data payload" [132]. This 20 mA is an order of magnitude larger than the average cost of writing to the MicroSD card (quoted as 3 mA), but we only need to send one message per reading, and we can optimize the size of the data payload.

Calibration

To compute bioimpedance from the raw data, the system first needs to be properly calibrated. The AD5933 internally computes the Discrete Fourier Transform of 1024 ADC samples at each frequency. This gives the power of the signal in the form of real r and imaginary i components. The magnitude of the impedance measurement at frequency ω is computed as $|Z(\omega)| = \sqrt{r_\omega^2 + i_\omega^2}$. The phase of

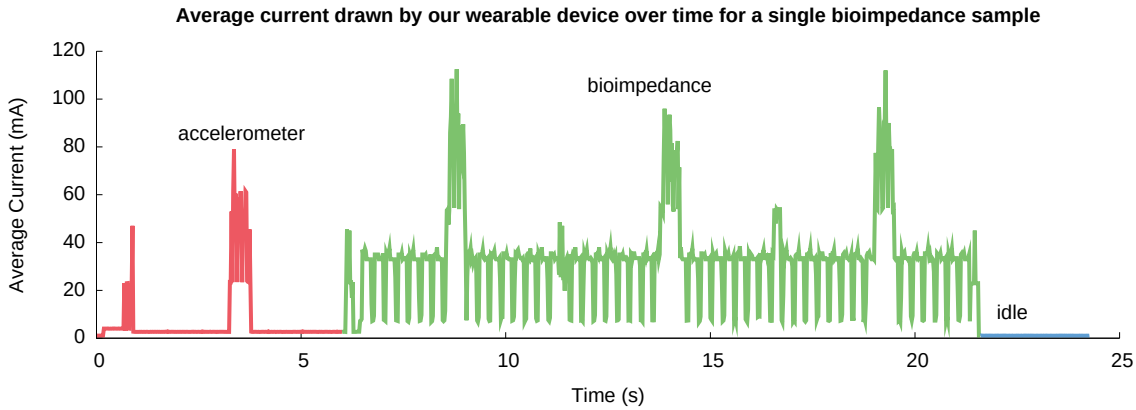

Figure 3.5.4: A representative energy measurement of our bioimpedance sensor. Notice the two phases present in the measurement: the accelerometer sampling phase, and the bioimpedance sampling phase. The large spikes in each phase correspond to writing to the SD memory card.

the impedance measurement is computed as $Z\emptyset(\omega) = \tan^{-1}\left(\frac{i_\omega}{r_\omega}\right)$, taking care to ensure $\frac{i_\omega}{r_\omega}$ is positive quantity and rotating this phase angle such that it falls in the appropriate quadrant in the complex plane depending on the actual signs of the real and imaginary components. Given the magnitude and phase, we compute the resistive and reactive components of impedance, represented by $R(\omega)$ and $X(\omega)$, by projection 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.

However, to get these impedance measurements, the device must be calibrated using a known impedance value. Calibration is important because different sensor modules will have different uncalibrated impedances due different parasitic elements and/or artifacts from the amplifiers present in the hardware. We built three of these sensor modules, so we must calibrate 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 configurations via a pair of multiplexors which can themselves introduce additional parasitic elements. Thus, we also must calibrate for every electrode configuration.

To calibrate our devices for a specific electrode configuration, we used a known reference impedance value $|Z|_{ref}$ in the form of an array of $1\,\mathrm{k}\Omega \pm 5\,\%$ resistors (shown in Figure 3.5.5; we used a digital multimeter to record the exact values between electrodes). To calibrate impedance magnitude, we first compute the uncalibrated impedance magnitude $|Z(\omega)|_{uncal}$ of the measured impedance as described above. Because we know the expected magnitude of the reference impedance $|Z|_{ref}$ (the impedance magnitude of a resistor is equal to its resistance), we compute a *gain factor* as:

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

This gain factor allows us to compute the actual impedance magnitude for some new uncalibrated impedance magnitude as:

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

A similar calibration procedure can be followed for the impedance phase. Since resistors have no impedance phase shift (aside from some parasitic capacitance, which we can ignore because these effects fall outside of the 100 kHz bandwidth of our system), we first compute the uncalibrated impedance phase $Z\emptyset(\omega)_{uncal}$ as specified above. This value is the *phase difference* that accounts for all phase shift attributed to the hardware itself since the impedance phase of a resistor should be zero. Given this phase difference, we compute the actual impedance phase for some new uncalibrated impedance phase as:

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

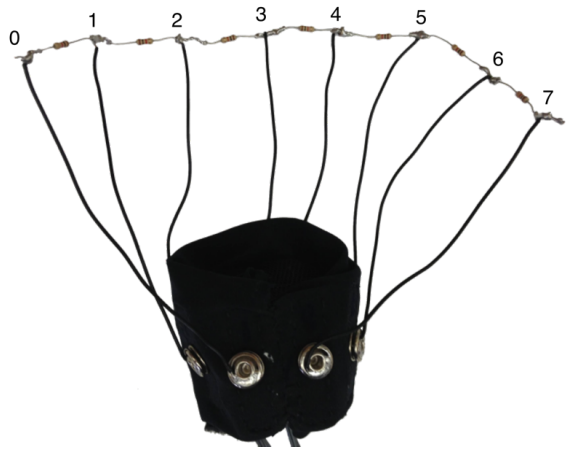

Figure 3.5.5: The resistor array used to calibrate our bioimpedance device. The $1\,\mathrm{k}\Omega \pm 5\,\%$ resistors are connected in series and act like a pseudo-wrist that can be electrically connected to the electrodes to calibrate the device. Since we measure bioimpedance at multiple electrode configurations, the reference resistance between a pair of electrodes will vary depending on the number of resistors between them. For example, the resistance between electrodes `04` is $4\,\mathrm{k}\Omega$.

As mentioned above, these calibration procedures are performed for each electrode configuration, which results in a gain factor and phase difference for each frequency and electrode configuration that is specific to that sensor module. More information on system calibration can be found in the AD5933 data-sheet [2].

The gain factor can also be computed for different temperatures, however, the influence of temperature variation on the AD5933 is linear (30 ppm/°C) and can easily be accounted for. Moreover, the temperature variation in our datasets showed little variation so we did not do any such compensation.

During calibration we noticed many of the lower frequencies exhibited random noise in the signal. We therefore discard frequencies below 10 kHz, resulting in 50 total frequencies. Although the source of this noise is unknown, development boards provided with the AD5933 exhibited similar noise so we do not believe it is an inherent limitation in the design of our own sensor module. Henceforth, all data reported represents measurements that have been calibrated in this way.

3.5.3 Electrode configurations

Recall that there are two modes of measuring bioimpedance: tetra-polar and bi-polar sensing. Tetra-polar sensing uses four electrodes, while bi-polar uses two as described in Section 3.3. Additionally, recall that we sense from eight locations on a person's wrist according to Figure 3.3.2. We compactly represent the choice of electrode configurations as such. For tetra-polar sensing, `1234` means that electrodes `1` and `2` were used to apply the current while electrodes `3` and `4` were used to sense the voltage change. For bi-polar sensing, `12` means that electrodes `1` and `2` were used to both apply the current and sense the voltage change.

3.6 Uniqueness

Uniqueness is an import characteristic of biometrics. However, it must be qualified by stating the size of the population for which the biometric is presumed to be unique. Over the lifetime of a personal device, it might only be used by a very small set of people (e.g., those living in a household). Thus, the purpose of this section is to understand how unique bioimpedance is for small populations.

3.6.1 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 [113] (as shown in Figure 3.6.1, 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 3.6.1 at left.

Figure 3.6.1: 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.

Once enrolled, we used the bench-top system (Figure 3.5.1) to collect data from subjects. We placed the electrode bracelet on the subject's left wrist and we instructed them 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 ± 3 years; all subjects were 18 years or older. In total, we collected 80 measurements from each subject (5 measurements for each electrode configuration), resulting in 3680 total bioimpedance measurements. Figure 3.6.2 shows a histogram of wrist circumferences by gender. The average subject wrist circumference was $16.00\,\text{cm} \pm 1.33\,\text{cm}$. For females, the average wrist circumference was $15.10\,\text{cm} \pm 0.55\,\text{cm}$, and males $17.00\,\text{cm} \pm 1.21\,\text{cm}$.

We collected bioimpedance samples from 16 different electrode configurations, 4 bi-polar and 12 tetra-polar electrode configurations. The bi-polar electrode configurations we collected were: `04`, `15`, `26`, and `37`. The tetra-polar electrode configurations we collected were: `0415`, `0426`, `0437`, `1526`, `1537`, `1540`, `2637`, `2640`, `2651`, `3740`, `3751`, and `3762`. The electrode configurations correspond to those electrodes that are maximally distant apart.

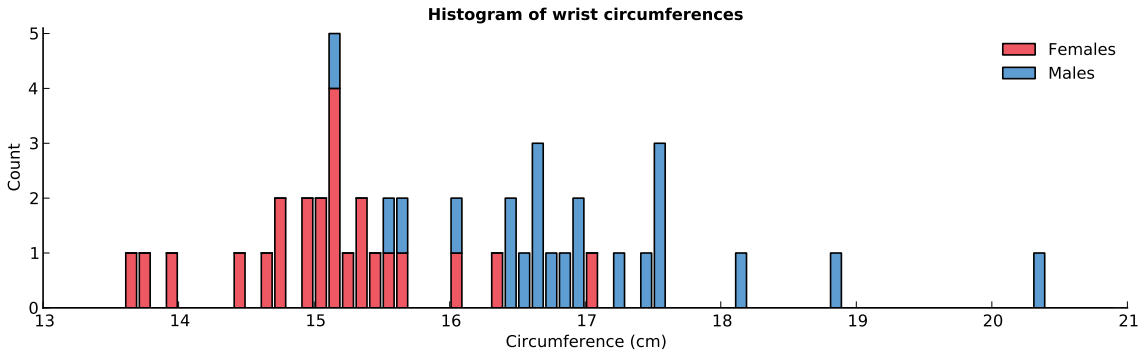

Figure 3.6.2: A histogram of wrist circumferences by gender. Males tend to have larger wrist circumferences.

3.6.2 Parameters

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

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 3.4. We computed the FAR, FRR, BAC for each subject, and we report the average of these measures over all subjects in the cohort. See Chapter 2 for the definition of these statistical measures.

Because of the large size of our dataset, we need to subdivide it into cohorts. It is, however, infeasible to evaluate all combinations of subjects for the specified cohort size (e.g., there are more than 260000000 combinations of subjects for a cohort size of 8.) Thus, we randomly sampled all possible cohort combinations such that we have a 95 % confidence (with a 1 % margin of error) that our sample size is representative. In our experiments, we examined cohort sizes from 2 to 8 subjects. We also present a 46-subject cohort for the sake of comparison.

3.6.3 Evaluation

To evaluate the uniqueness of bioimpedance, this experiment focused on the ability of our method to distinguish among people, also known as identification. We sought to determine how well our method performs for small population sizes where all the users are known and the goal is to recognize whether the device is worn by subjects in that population.

Single electrode configurations

In this experiment, we tested all combinations of electrode configurations according to the method described in Section 3.4. The purpose of this experiment was to understand how well a single electrode configuration can recognize a subject. Figures 3.6.3 and 3.6.4 show the results of this experiment for a single bi-polar and tetra-polar electrode configuration and for different classifiers.

Recognition rates decreased as the cohort size increased, as one would expect. The NB and FOREST classifiers perform better than the SVM classifier. The FOREST classifier had lower FAR than the NB classifier, but at the expense of a higher FRR. If we extrapolate out to the 46-subject cohort size, the NB classifier asymptotes around 80 % BAC, while the FOREST classifier decreases linearly as the cohort size increases. The SVM classifier becomes useless at such a large cohort size. Both the SVM and FOREST classifiers were more sensitive to the unbalanced dataset than the NB classifier.

There was little difference between the top performing single tetra-polar `2637` and bi-polar electrode configuration `37`. Both the NB and FOREST classifiers performed similarly. The SVM classifier, on the other hand, performed better slightly better when it uses the tetra-polar electrode configuration. For a single electrode configuration it makes more sense to use bi-polar sensing because the overhead of

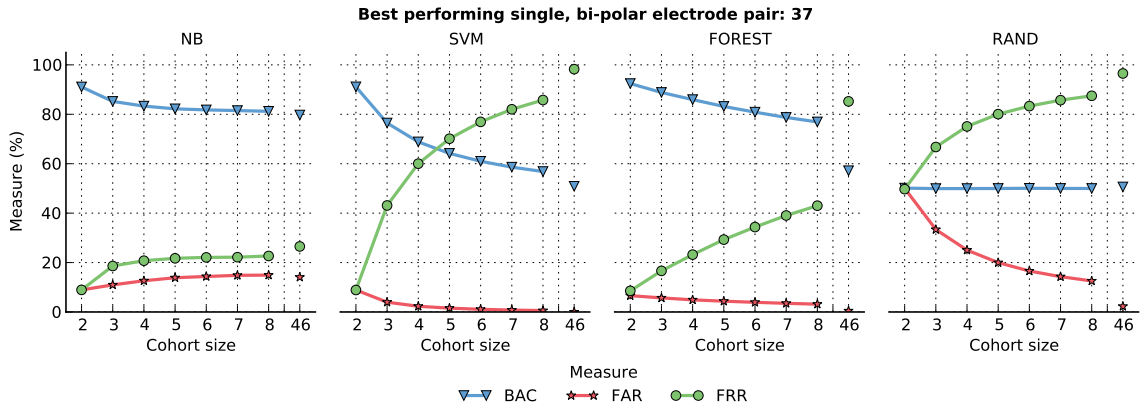

Figure 3.6.3: Electrode configuration `37` was the best-performing single, bi-polar electrode configuration. The NB and FOREST classifiers performed better than the SVM classifier.

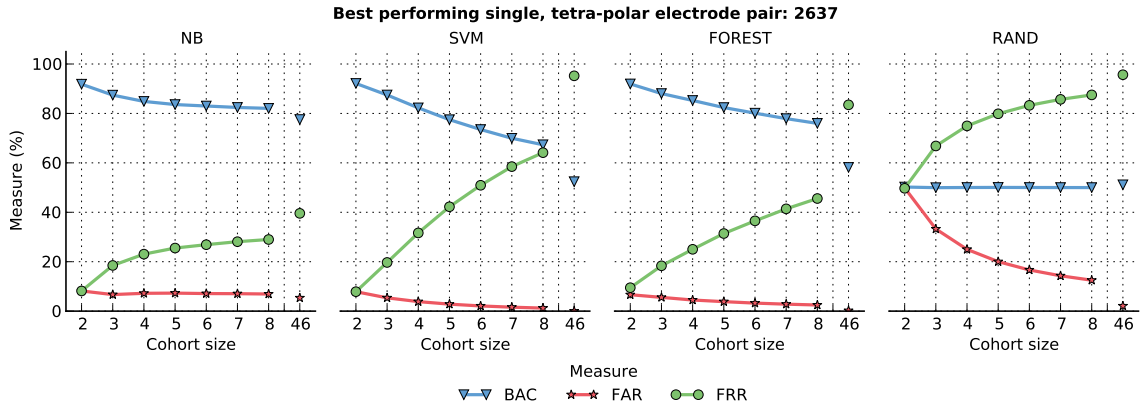

Figure 3.6.4: electrode configuration `2637` was the best-performing single, tetra-polar electrode configuration. These results are similar to the bi-polar electrode configuration case `37` shown in Figure 3.6.3.

tetra-polar sensing does not compensate for increased recognition performance.

Electrode configuration `37` clearly played a part in both the top performing bi-polar and tetra-polar electrode configurations. In fact, the next best-performing tetra-polar electrode configuration was `1537`. There was power in sensing at electrode configuration `37`. The location of these electrodes on the bracelet corresponded to the medial dorsal and medial palmar sides of the wrist. This was one of the smallest distances between any pair of electrodes we examined because of the oval shape of people's wrists.

Combined electrode configurations

Since we are not limited to just one particular electrode configuration, we hypothesized that it might boost recognition rates to concatenate feature vectors from multiple electrode configurations into a combined electrode configuration feature vector. For example, in the bi-polar case we can incorporate feature vectors from electrode configurations `15` and `26` by concatenating them together. This approach could boost recognition rates because the applied current takes different paths through the subject's wrist for different electrode configurations.

In the bi-polar case, we explored all combinations of bi-polar electrode configurations (e.g., `15 26`, `15 37`, ..., `15 26 37 48`); there are 11 such combinations. In the tetra-polar case, we explored all combinations of tetra-polar electrode configurations such that the electrodes supplying current are distinct (e.g., `1526 2637`, but not `1526 1537`); there are 243 such combinations.

Figures 3.6.5 and 3.6.6 show the results of this experiment. In comparison to the single electrode configuration, the difference in performance was minimal for both the NB and FOREST classifiers. For all classifiers, however, using multiple tetra-polar electrode configurations reduced the FAR, but typically at the expense of the FRR in the case of the NB classifier. However, the performance of the SVM classifier increased dramatically with the use of multiple electrode configurations, so much that it becomes the best-performing classifier with a BAC of 88.0 % for a cohort of eight subjects. Overall, both the NB and SVM classifiers benefited from multiple electrode configurations.

Once again, the electrode configuration `37` made an appearance in both experiments. The top 6 bi-polar combined electrode configurations included `37`, while at least the top 10 tetra-polar combined electrode configurations do as well. Thus, we recommend that a bracelet needs to include electrodes that have contact with the

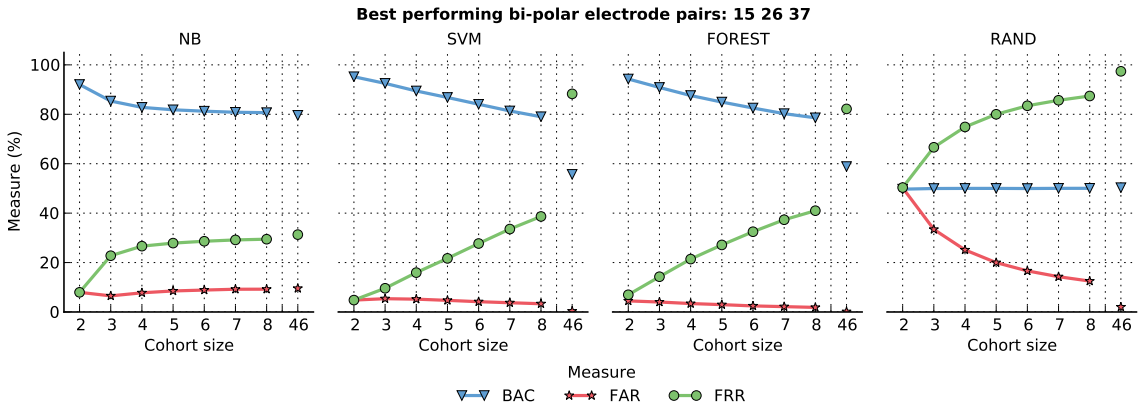

Figure 3.6.5: electrode configurations `15 26 37` were the top performing bi-polar combined electrode configuration. The usage of multiple electrode configurations increase performance slightly.

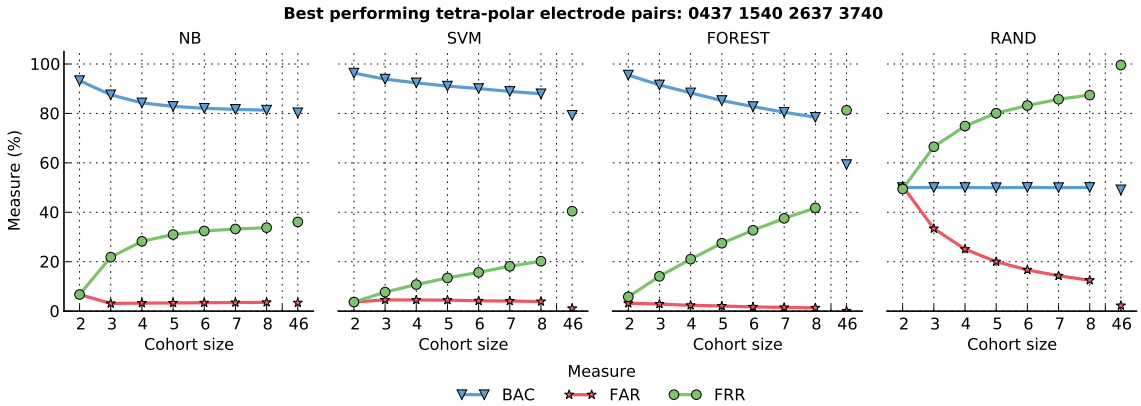

Figure 3.6.6: electrode configurations `0437 1537 2637 3740` were the best-performing tetra-polar combined electrode configurations. These results follow similar trends to previous experiments in Figure 3.6.4 with a slight performance increase.

medial palmar and medial distal sides of the wrist to achieve the best recognition performance.

Notice that the FAR decreased at the expense of a higher FRR. Thus, if an application requires fewer false accepts, then we recommend using multiple combined electrode configurations. The best-performing bi-polar combined electrode configurations included almost all of the electrode configurations we used to take bi-polar measurements. This suggests that we should have added other bi-polar electrode configurations to our collection process. For example, we could have sensed

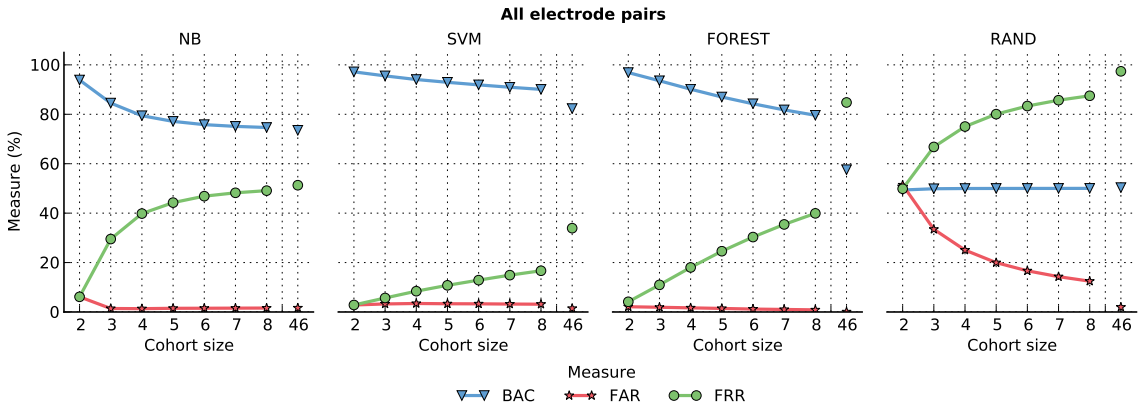

Figure 3.6.7: Recognition performance for all electrode configurations. The SVM classifier benefits most from all electrode configurations while the NB and FOREST do not benefit.

from electrode configurations `02`, `03`, etc. Even though the best-performing bi-polar combined electrode configurations included most of the electrodes, electrode configurations `26 37` performed similarly. Thus at least four electrodes are necessary for bi-polar measurements. For tetra-polar measurements, we recommend at least six to get recognition on par with the top performing electrode configurations. The best-performing six-electrode tetra-polar combined configuration was `0437 2637`.

All electrode configurations

Finally, we tried combining tetra-polar and bi-polar measurements into one feature vector. That is, what happens when we give all the data we collected to each classifier and evaluation how each performs? They should at least perform as well as the top performing tetra-polar combined electrode configurations.

Figure 3.6.7 shows the results of this experiment. Given all electrode configurations, the SVM classifier performed the best once again. Since it takes around 2 minutes to collect all of these electrode configurations, we do not recommend this approach. It does, however, give us a bound on how well our method performs when given all the information in our dataset.

Wrist circumference recognition

From the data in Figure 3.6.2, we hypothesized that a subject's wrist circumference might serve as a good feature for recognizing subjects because many of the subjects fall into their own bin on the histogram. We did not, however, take multiple measurements of each subject's wrist circumference, and so the distribution of measurement errors is unknown. Knowing this distribution is important because a hypothetical device that measures wrist circumference will have some measurement error that would depend on the measurement characteristics of that device. We can, however, simulate taking multiple measurements by assuming some distribution of 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 can randomly sample a normal distribution with mean 15 cm and standard deviation of 5 mm. That is, we assume the measurement error follows a normal distribution. 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 3.6.8 shows the result of this experiment. With no measurement error, we recognized subjects 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 performance should be even better because wrist sizes tend to correlate well with age.) As measurement error increased, however, the recognition rates fell: with 1 mm measurement error, recognition rates fell to 54 % BAC for the full cohort of subjects. For smaller cohort sizes, recognition rates remained above 80 % when the measurement error was 1 mm or less. Recognition rates were even worse for 5 mm and 1 cm measurement errors. This result implies that wrist size is a good biometric for small cohorts when there is little measurement error; but can we do better with a hybrid approach?

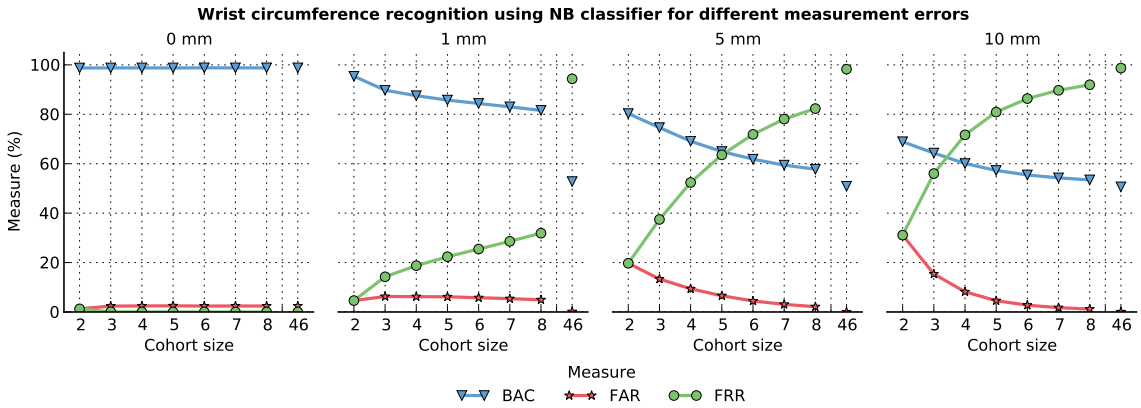

Figure 3.6.8: Wrist circumference recognition rates for various cohort sizes and measurement errors using the NB classifier.

To get a sense of how accurately an electronic device could measure a person's wrist circumference, we used the digital tape measure shown in Figure 3.7.1. We took 10 consecutive measurements from a subject's wrist, resetting the digital tape measure each time. The standard deviation of these measurements was 0.199 cm. The precision of the digital tape measure is only to the nearest 0.2 cm. Thus, we expect a measurement error between 1 mm to 5 mm.

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 3.6.9 shows the results of this experiment for different cohort sizes using all electrode configurations and the SVM classifier. In contrast to using wrist circumference alone, combining bioimpedance and wrist circumference dramatically lowers the FRR and FAR 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 applications that require security, especially in large cohorts, because of the lower FARs.

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

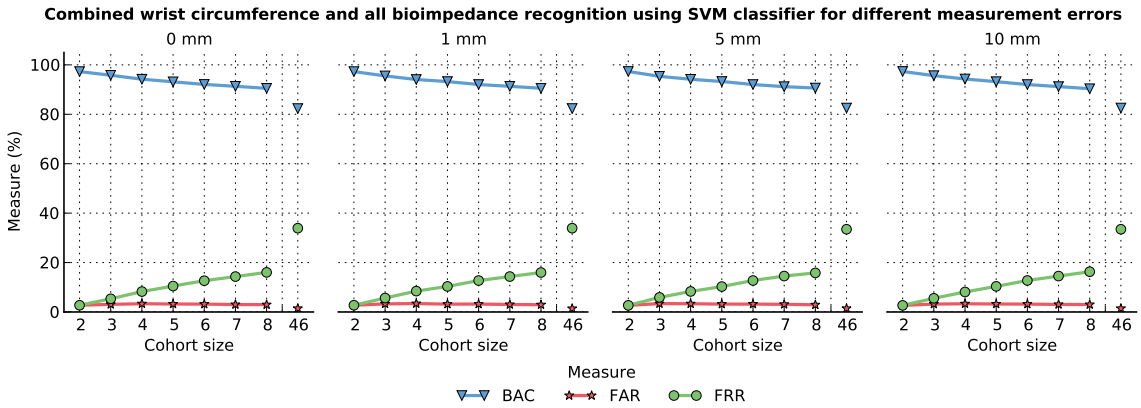

Figure 3.6.9: Bioimpedance with wrist-circumference models: recognition rates for a 5-subject cohort and various measurement errors, using theSVM classifier.

algorithm to choose the model that classified the test feature vector positively, or choose no model if no model or more than one model positively classified the test feature vector. For a cohort of five subjects and wrist-size measurement error 1 mm, our device would confuse subjects only 0.5 % of the time (2.0 % of the time for a 5 mm measurement error). Given such rates, we believe bioimpedance measurements are individually unique for a household population.

3.7 Performance

The third characteristic we examine is the performance of the biometric. To determine the performance of bioimpedance as biometric, we manufactured a wearable bioimpedance bracelet device. We then gave this device to subjects to wear for a day to understand how well a wearable version of our bench-top system would perform.

3.7.1 Dataset

We collected bioimpedance data using our wearable sensor (Figure 3.5.3) from eight people over a period of one day each. Before enrollment, participants were told the reason we were collecting data and informed of the risks involved in wearing the

Figure 3.7.1: 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 bracelet was worn.

device (e.g., possible inflammation at electrode site). Once they agreed to enroll, we gathered three pieces of information about them: age, sex, and wrist circumference. We measured the subject's wrist using a Health-O-Meter Digital Tape Measure [69] at the location where the bracelet would be worn as shown in Figure 3.7.1. Age and sex were self reported. We then introduced each subject to the device and asked them to wear the sensor on their non-dominant wrist for as long as possible during the day and return it to us after at least eight hours. We told subjects that they could remove the device at their leisure but should remove the device if it could 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 put on. Subjects were paid $8 for their participation. Our device and data-collection protocols were approved by our Institutional Review Board.

Of the participants enrolled (Table 3.7.1), 3 were female, 5 were male, with an average age of 27 years $\pm$ 8 years. The average wrist circumference was 18.0 cm $\pm$ 1.0 cm meaning that the group of subjects selected for this study each had a similar wrist circumference. Table 3.7.1 shows, among other statistics, the duration the bracelet was in contact for each subject. On average, it was on a subject's wrist for 9.3 hours resulting in 112 bioimpedance samples. This group of 8 subjects is larger than our

Subject	Total	No Contact	Contact	Contact Time (h)	Wrist (cm)
1	111	6	105	8.8	18.2
2	96	8	88	7.3	16.2
3	107	22	85	7.1	18.4
4	262	27	235	19.6	17.0
5	134	12	112	10.2	19.4
6	123	38	85	7.1	18.4
7	134	59	75	6.2	17.6
8	111	13	98	8.2	19.0
Avg	135 ± 50	23 ± 17	110 ± 52	9.3 ± 4.0	18.0 ± 1.0

Table 3.7.1: The total number of samples taken, the number of samples that were deemed not to be in contact with the subject's wrist, and the total time the bracelet was in contact with the subject's wrist.

uniqueness experiment, since we previously examined group sizes of up to 5 subjects.

Figure 3.7.2 shows how well an SVM classifier classified subjects for different circumference error measurements. For no measurement error, using wrist circumference alone we could recognize subjects with 93 % BAC, 1.4 % FAR, and 12 % FRR. However, when the measurement error was 5 mm, the BAC fell to 56 %, FAR to 0.4 %, and the FRR to 87 %. Because the FRR increased while the FAR remained stable, the SVM over-fit each subject and thus rejected more and more of each subject's genuine samples. Thus, using wrist circumference alone would not be feasible for this dataset.

3.7.2 Parameters

We chose parameters similar to those in our uniqueness study (Section 3.6). However, since our wearable device does not support tetra-polar measurements, we only sampled bi-polar measurements. To compensate for the lack of tetra-polar measurements, in addition to sensing at the maximally distant electrodes as in our uniqueness study, we also measured bioimpedance at those electrodes that were at least one electrode

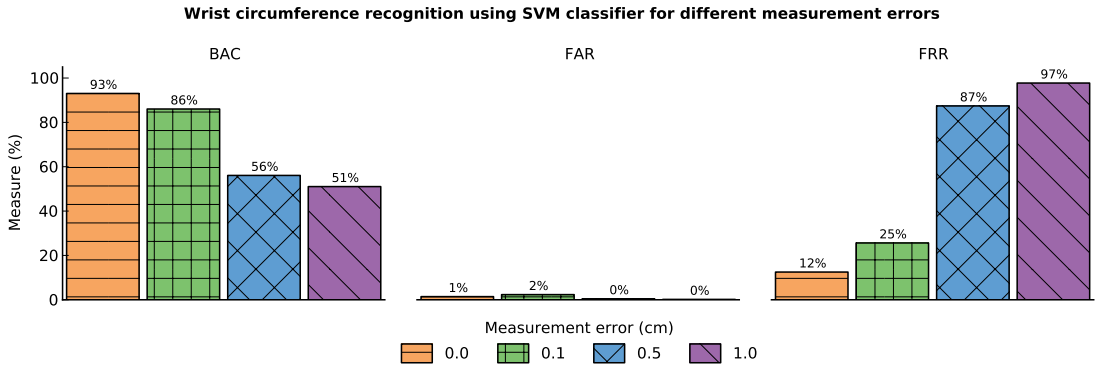

Figure 3.7.2: Wrist circumference recognition rates for different circumference measurement errors using an SVM classifier on our wearable dataset. As the measurement error increases, the SVM classifier over-fits on each subject's genuine examples more.

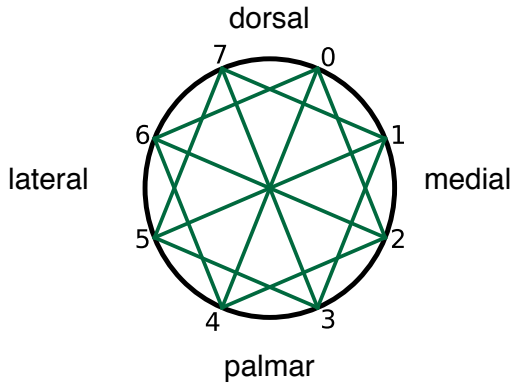

Figure 3.7.3: The electrode configurations our wearable bioimpedance sensor used for bi-polar measurements.

apart. Figure 3.7.3 graphically depicts these electrode configurations. According to our uniqueness study, the combination of all of the maximally distant electrode configurations performed the best of the bi-polar electrode configurations. Thus, we hypothesized that adding other electrode configurations would only increase performance because they give us considerably more coverage of the anatomy and geometry of the wrist. These twelve electrode configurations gave four times more coverage than those electrode configurations in our uniqueness study.

3.7.3 Identification evaluation

Using this new wearable device, we sought to understand its performance. In contrast to the setting in our previous study on uniqueness, the purpose of this study was to understand the performance of our wearable bioimpedance sensor outside of the lab and in the field. That is, we gave subjects our wearable bioimpedance sensor to wear for a day while they went about their usual activities. In this set of experiments, we sought to understand the performance of our wearable device in an identification setting outside of the lab. Recall that identification is the process of identifying which subject was wearing the device from a population of known subjects.

Uniqueness replication

We first validated the performance of our methods under conditions similar to our uniqueness study. We did a 10-fold cross validation for each subject using only the electrode configurations that were maximally distant (i.e., `04`, `15`, `26`, and `37`).

Figure 3.7.4 shows the results of this validation for the best-performing combinations of the specified electrode configurations for the NB, SVM, and FOREST classifiers. The best-performing single electrode configuration in this study was `26` for the NB classifier and `04` for the SVM and FOREST classifiers. That the `04` electrode configuration performed well confirms our previous finding that the best-performing single electrode configuration corresponded to those electrodes on the dorsal and palmar sides of the wrist.

For the SVM and FOREST classifiers, the best combined electrode configuration was `04 37`. Because of the design of our bracelet, electrode configuration `04` did not sit exactly at the top of the wrist, as in our uniqueness study, but shared the top of the wrist with electrode configuration `37`. Contrast this result to the best electrode configuration for the NB classifier, where `15 26` correspond to the medial

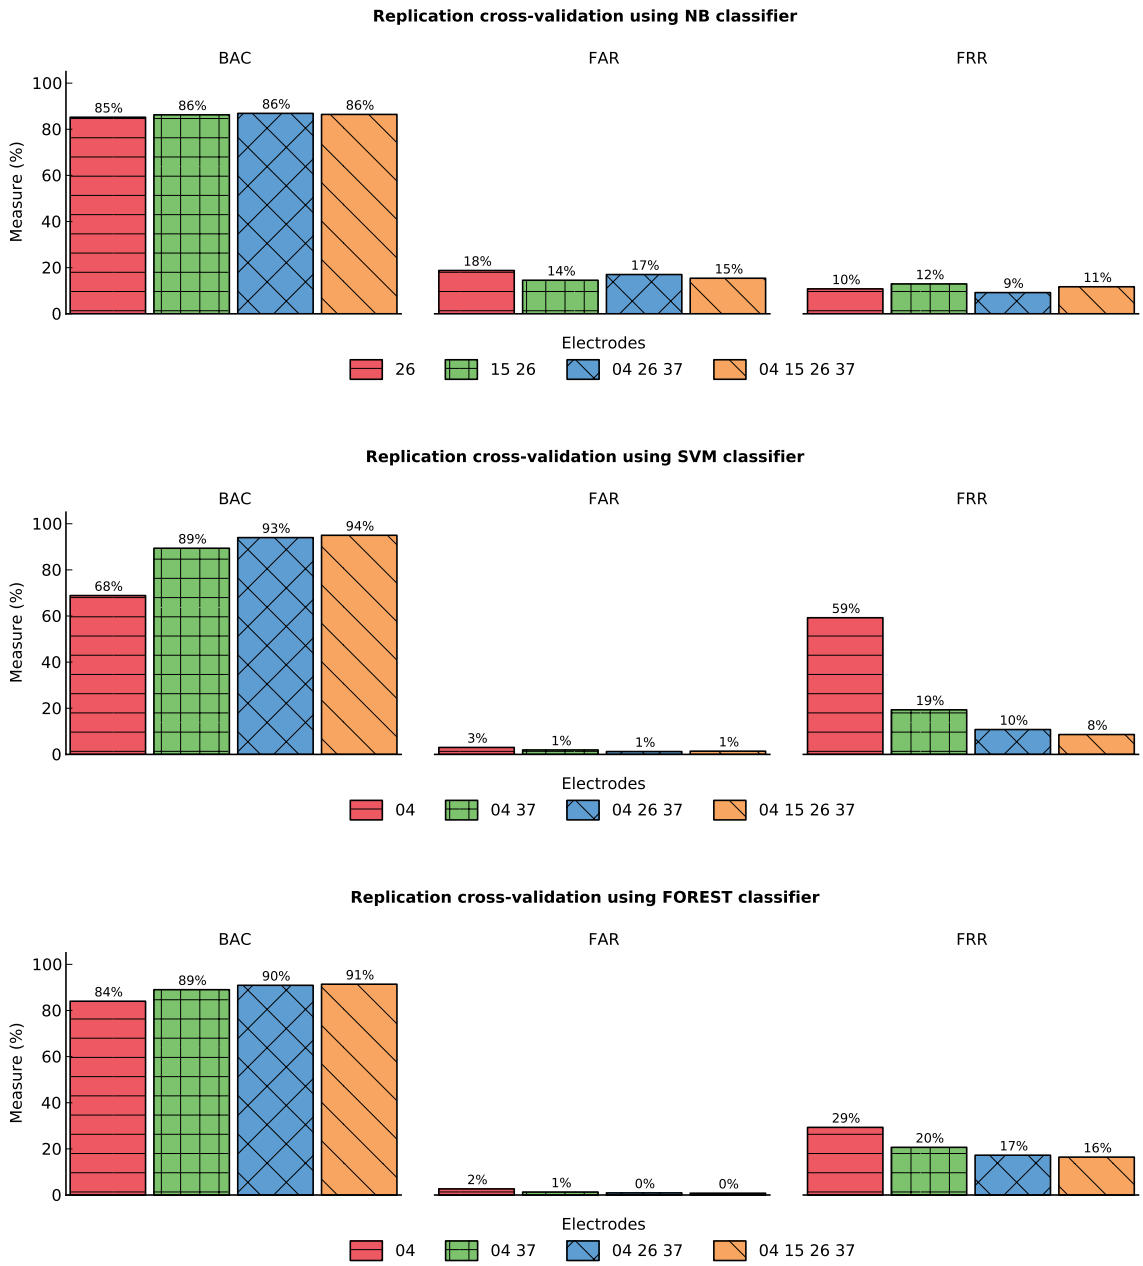

Figure 3.7.4: A replication of uniqueness study using the NB, SVM, and FOREST classifiers over our wearable dataset restricted to only those electrode configurations maximally distant. For a single electrode configuration, `04` and `26` does best. The SVM and FOREST classifiers are significantly better than the NB classifier when using the all electrode configurations.

and lateral sides of the wrist.

The SVM classifier achieved the best performance when all electrode configurations are combined. It achieved a BAC of 95.0 %, a FAR of 1.40 %, and a FRR of 8.62 %. The performance of the FOREST classifier was similar to the SVM classifier with a BAC of 91.4 %, a FAR of 7.83 %, and a FRR of 16.4 %. In contrast, the NB classifier only achieved a BAC of 86.4 %, a FAR of 15.4 %, and a FRR of 11.7 %, similar to the findings in our uniqueness study. For comparison, the RAND classifier empirically achieved 50.2 %, a FAR of 12.9 %, and a FRR of 86.7 %. In comparison to our uniqueness study, our wearable device performed significantly better because we have more data from each subject. The best-performing classifier in the uniqueness study, the SVM classifier, only achieved a BAC of 88.4 % for a cohort of eight subjects using all electrode configurations.

Cross-validation

Next, we did a cross-validation experiment over all combinations of electrode configuration that were used to collect bioimpedance samples. This validation makes the assumption that we have almost an entire day of training samples available. It also 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 according to the algorithms specified in Section 3.4. That is, we trained our model using 90 % of the samples randomly chosen, leaving 10 % of the samples to be classified. We computed the FAR, FRR, and BAC for each subject, and we report these measures over all subjects.

Figure 3.7.5 shows these results for the top performing multi-electrode configurations using the NB, SVM, and FOREST classifiers. The top performing multi-electrode configurations achieved a BAC of 98.1 %, FAR of 0.70 %, and FRR of 2.83 %

using the SVM classifier and electrode configurations `02 04 06 13 15 17 26 37`. The NB classifier, on the other hand, achieved a BAC of 87.9 %, FAR of 3.83 %, and FRR of 9.62 %. The FOREST classifier performed better than the NB classifier but worse than the SVM classifier with a BAC of 93.6 %, a FAR of 0.53 %, and a FRR of 12.2 %. The SVM classifier benefited from more electrode configurations, while the performance of the NB classifier did not significantly increase except for reductions in FAR and FRR.

Figure 3.7.6 shows a visualization of these top performing multi-electrode configurations for each classifier. The top performing multi-electrode configurations for the NB classifier was `04 06 26 35 37`, for the SVM classifier was `02 04 06 13 15 17 26 37`, and for the FOREST classifier was `02 04 06 13 15 17 26 35 37`. In these top multi-electrode configurations for the SVM classifier, electrode configurations `04`, `06`, `26`, and `37` were present in each. Two of these electrode locations of top performing multi-electrode configurations, `04` and `37`, corresponded to the ulnar side of wrist or medial palmar and medial dorsal (imagine your hand sticking out of the page towards yourself, palm down). The top performing multi-electrode configurations for the FOREST classifier tended to encompass the whole wrist. As the number of electrode configurations increased, more of the geometry of the wrist was being sensed by our device. This result shows that using electrode configurations that are not just maximally distant increase recognition performance. We recommend sensing from electrode configurations that capture the ulnar side of the wrist.

Cross-validation with wrist circumference

In our uniqueness study, we considered adding wrist circumference as a distinguishing feature, which decreased the FAR. We performed this same experiment using our

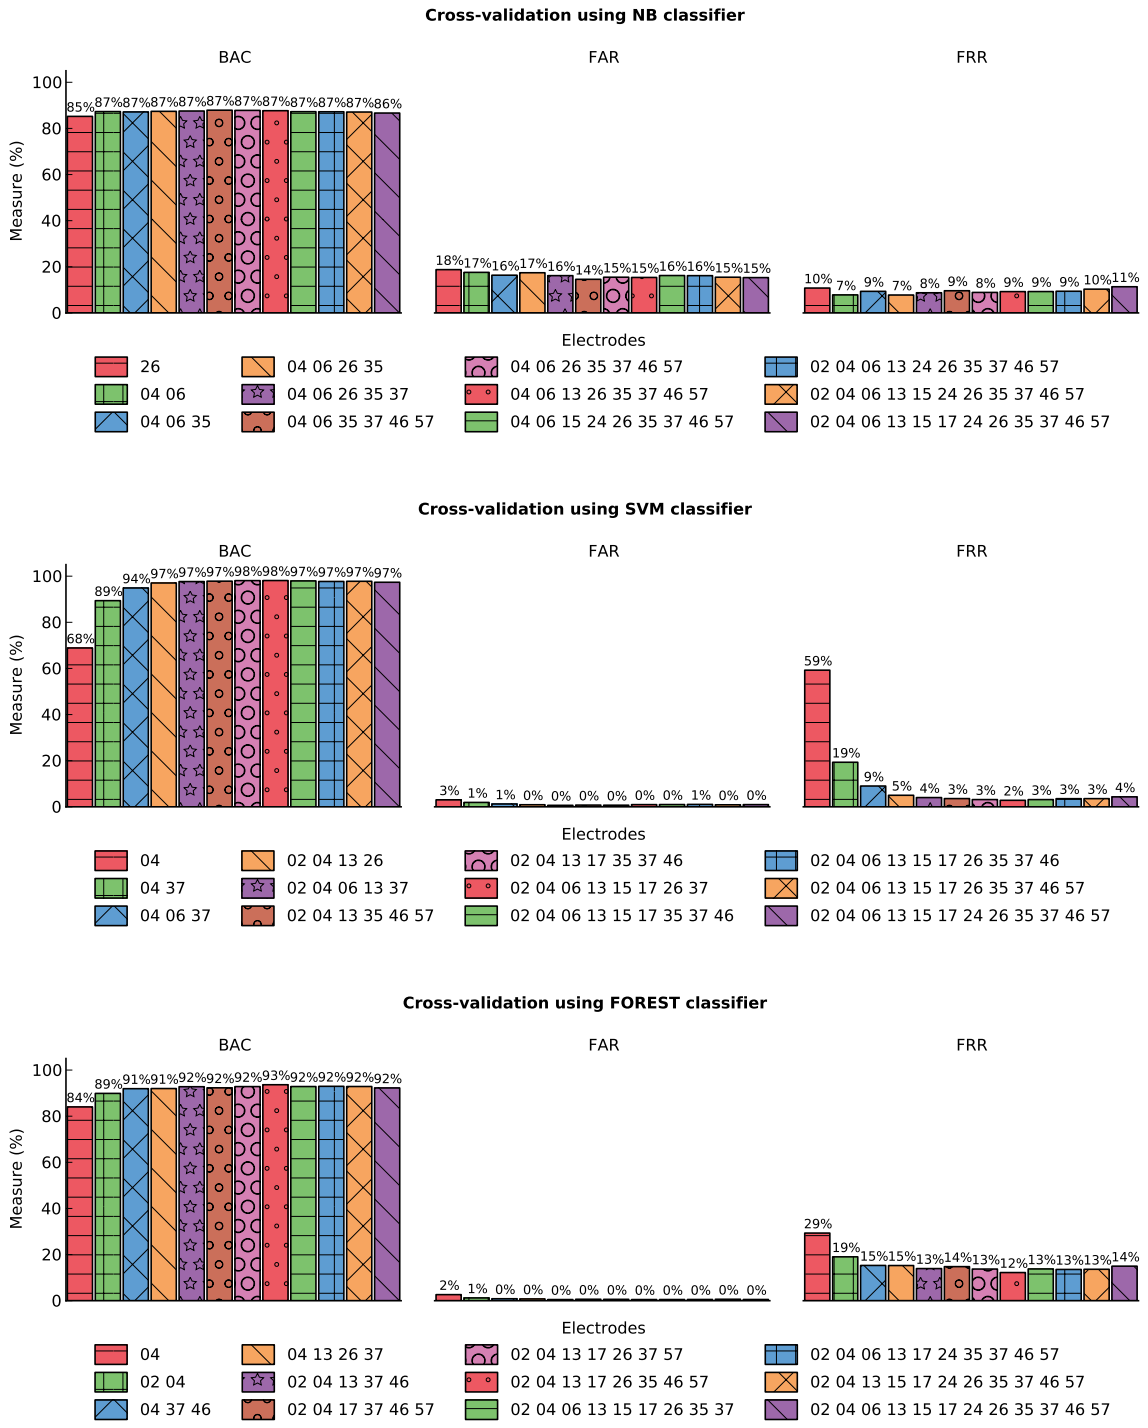

Figure 3.7.5: Top performing electrode configurations for a 10-fold cross-validation as classified by the NB, SVM and FOREST classifiers. The performance of the NB classifier was flat while the SVM and FOREST classifiers benefited from more electrode configurations.

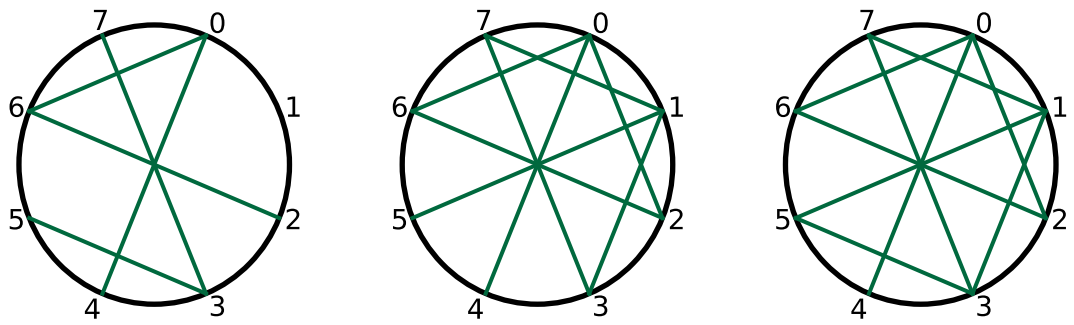

(a) NB **(b)** SVM **(c)** FOREST

Figure 3.7.6: Visualization of best-performing combinations of electrode configurations for the NB, SVM, and FOREST classifiers. Notice how the top performing multi-electrode configurations for the SVM and FOREST classifiers encompasses the medial side of the wrist.

wearable device. We assumed we can measure the subject's wrist circumference with measurement errors 0 cm, 0.1 cm, 0.5 cm and 1.0 cm.

Figure 3.7.7 shows the results of this experiment for the SVM classifier using all electrode configurations. Notice how these results do not significantly differ from the results shown in Figure 3.7.6. This means that wrist circumference provided no additional recognition power. While this claim might seem contradictory to our claim in the uniqueness study, it is not. Rather, we measured from more electrode configurations which allowed the enrollment algorithm to learn more about the wrist. Indeed, one of the defining features of a wrist is its circumference, and by measuring bioimpedance at many locations on the wrist, they can serve as a surrogate measure of the wrist circumference. Thus, such a device need only use bioimpedance to achieve the best recognition rates.

Hold-out validation

Next, we performed a hold-out validation to see how well bioimpedance performs over time. This kind of validation tells us how much training data is necessary to achieve acceptable recognition performance. In this validation, for each classifier

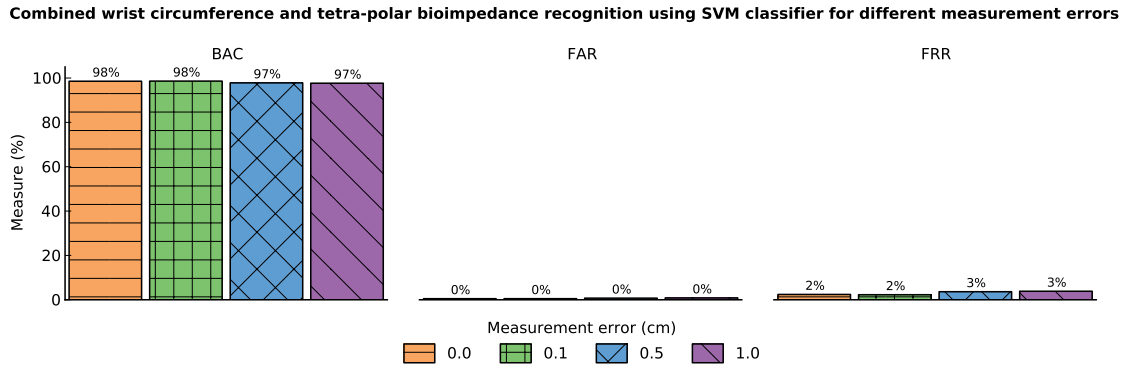

Figure 3.7.7: Combined bioimpedance and wrist circumference recognition rates for different circumference measurement errors using an SVM classifier on our wearable dataset.

we choose the best-performing electrode configuration combination reported in the cross-validation as shown in Figure 3.7.6.

Figures 3.7.8 shows how well the NB, SVM, and FOREST classifiers performed as the duration of training data increased. As expected, additional training data caused the BAC to increase while the FAR and FRR decreased in the case of the SVM classifier. However, in the case of the NB and FOREST classifiers, the BAC remained relatively flat due to the FAR increasing while the FRR decreased. This suggests that the NB and FOREST classifiers are unsuitable for recognition when there is limited training data. Even when only 10 % of a subject's bioimpedance samples were used for training (which was about 10 training samples), the SVM classifier achieved a BAC of 81.2 %, FAR of 5.62 %, and FRR of 32.0 %. This is on par with uniqueness study despite the increased cohort size and number of bioimpedance samples.

3.7.4 Verification evaluation

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 a subject is the person they claim to be. Thus, the general task is to train a classifier

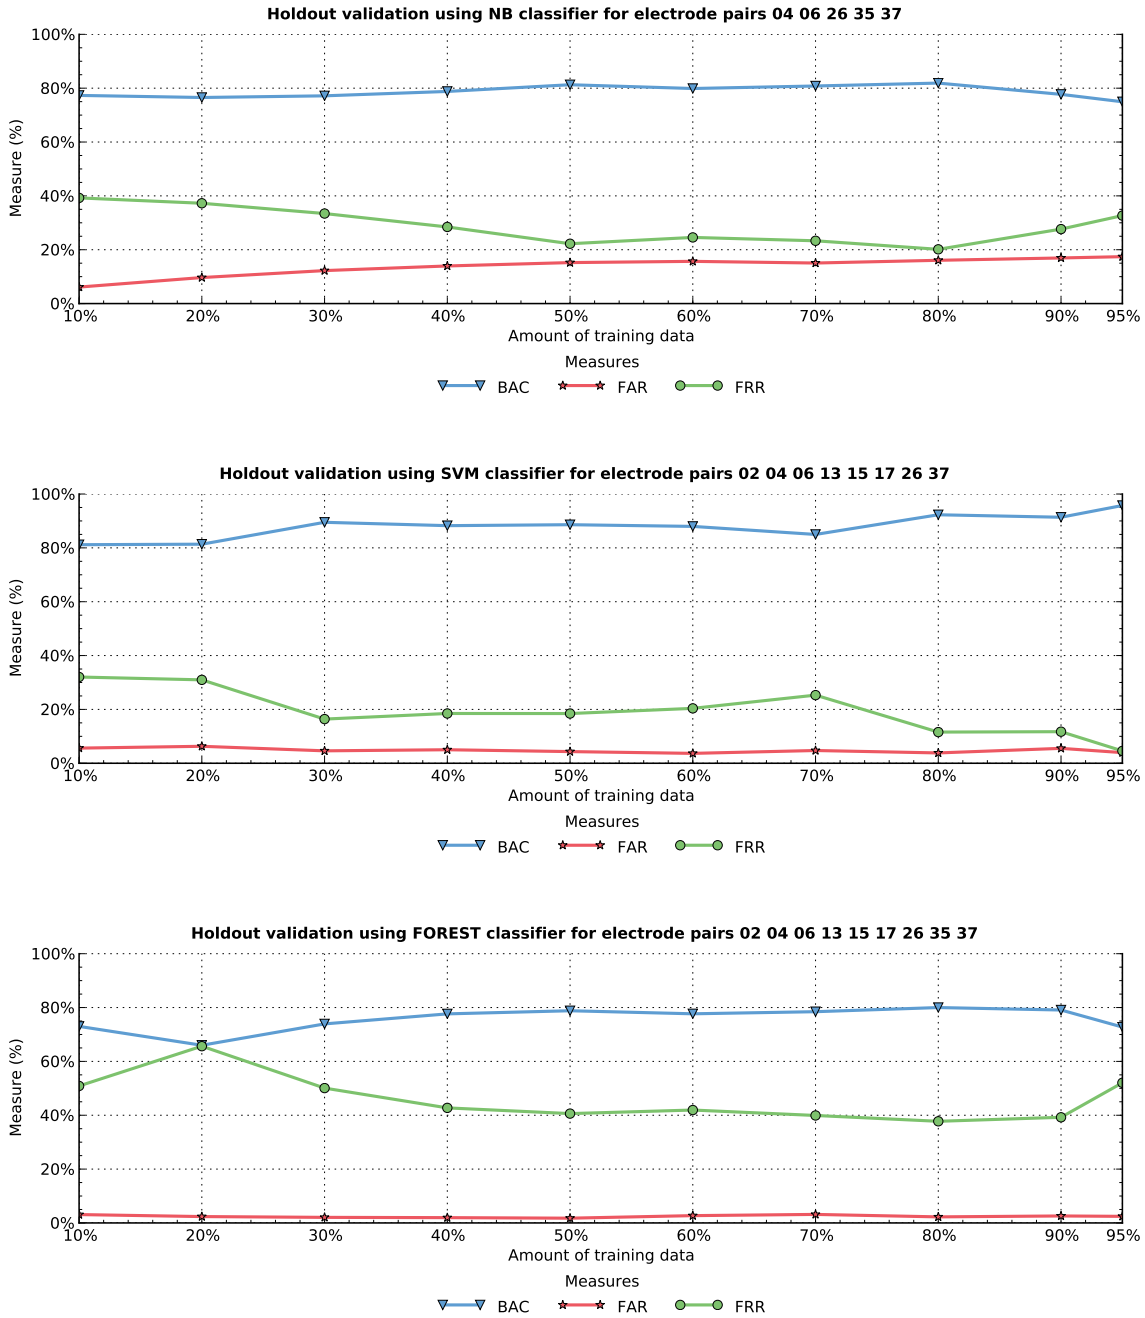

Figure 3.7.8: Hold-out validation as classified by NB, SVM, and FOREST classifiers for different amounts of training data for the best-performing electrode configuration. As the amount of training data increases, so did the BAC while the FAR and FRR decrease in the case of the SVM classifier. The performance of the NB and FOREST classifiers remains flat due to an increasing FAR.

using a single subject's data. In the testing phase, we see how well that classifier performs for that particular subject's data in contrast to all the other subjects' data. Notice, however, that our standard classifiers will not work for this procedure since they require both negative and positive examples. In the verification setting, we only have positive examples because the model is trained on only one subject.

Because we only have positive examples, we need to use a *generative* model that learns the distribution of a subject's bioimpedance samples. One such generative model is a Gaussian Mixture Model (GMM). A GMM models the distribution of bioimpedance samples using a weighted linear combination of Gaussian densities. Each Gaussian density is parameterized by a mean vector and covariance matrix. To learn the underlying distribution of a subject's bioimpedance, we use the Expectation-Maximization (EM) algorithm [40] to iteratively refine the mixture of Gaussian densities until the maximum likelihood remains stable (i.e., the difference between successive iterations is less than 0.01) or after a maximum number of iterations (100). We choose initial Gaussian densities by clustering the set of feature vectors using k-means clustering [84], where k is set to the desired number of Gaussian densities. We then iteratively fit these initial Gaussian densities using the EM algorithm. We modeled the full covariance matrix since the size of a bioimpedance sample is relatively small. Because some values of the covariance matrix can become very small, as in the case of outliers, we enforce a variance floor of 0.001 on the covariance matrix. For our experiments, we found that 4 Gaussian densities best modeled the distribution of a subject's bioimpedance samples.

In this experiment we used a hold-out validation to validate the performance of our method. However, it is not clear how to test a given bioimpedance sample since a GMM has no intrinsic notation of a label like the classification algorithms do. A GMM can tell us the likelihood of a bioimpedance sample given the parameters θ of the GMM (i.e., $p(x|\theta)$). Given some threshold τ, we can accept that bioimpedance

sample came from that subject if $p(x|\theta) > \tau$ and reject it otherwise. However, there is no good way of choosing the threshold τ *a priori*. Instead of choosing, we can vary the threshold τ to see how well our method performs.

Figure 3.7.9 shows how well our method performs for a 90 % hold-out validation. We varied the threshold for each subject to see how the FAR and FRR vary. In the ideal situation, both the FAR and FRR would remain near 0. In the legend of Figure 3.7.9, we also show the Equal Error Rate (EER). This was the rate at which the FAR equaled the FRR. There was a significant difference between EERs of subjects. The average per-subject EER was 12.90 % $\pm$ 8.15 %. Rather than a subject-specific threshold, we also used a global threshold over all subjects. The EER for a global threshold was 12.9 %. This EER resembles the result in our identification hold-out validation using an SVM classifier as shown in Figure 3.7.8. In a verification-based method one can easily change the threshold to suit the needs of the application to account for less false-positives or false-negatives as Figure 3.7.9 shows.

3.8 Permanance

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. Likewise, we want to understand how motion, orientation and the environment affect bioimpedance measurements. In this section, we explore these factors.

3.8.1 Motion effects

To study the effect of motion, we collected a dataset from a single subject performing three tasks: stationary, walking, and running. In the stationary task, the subject

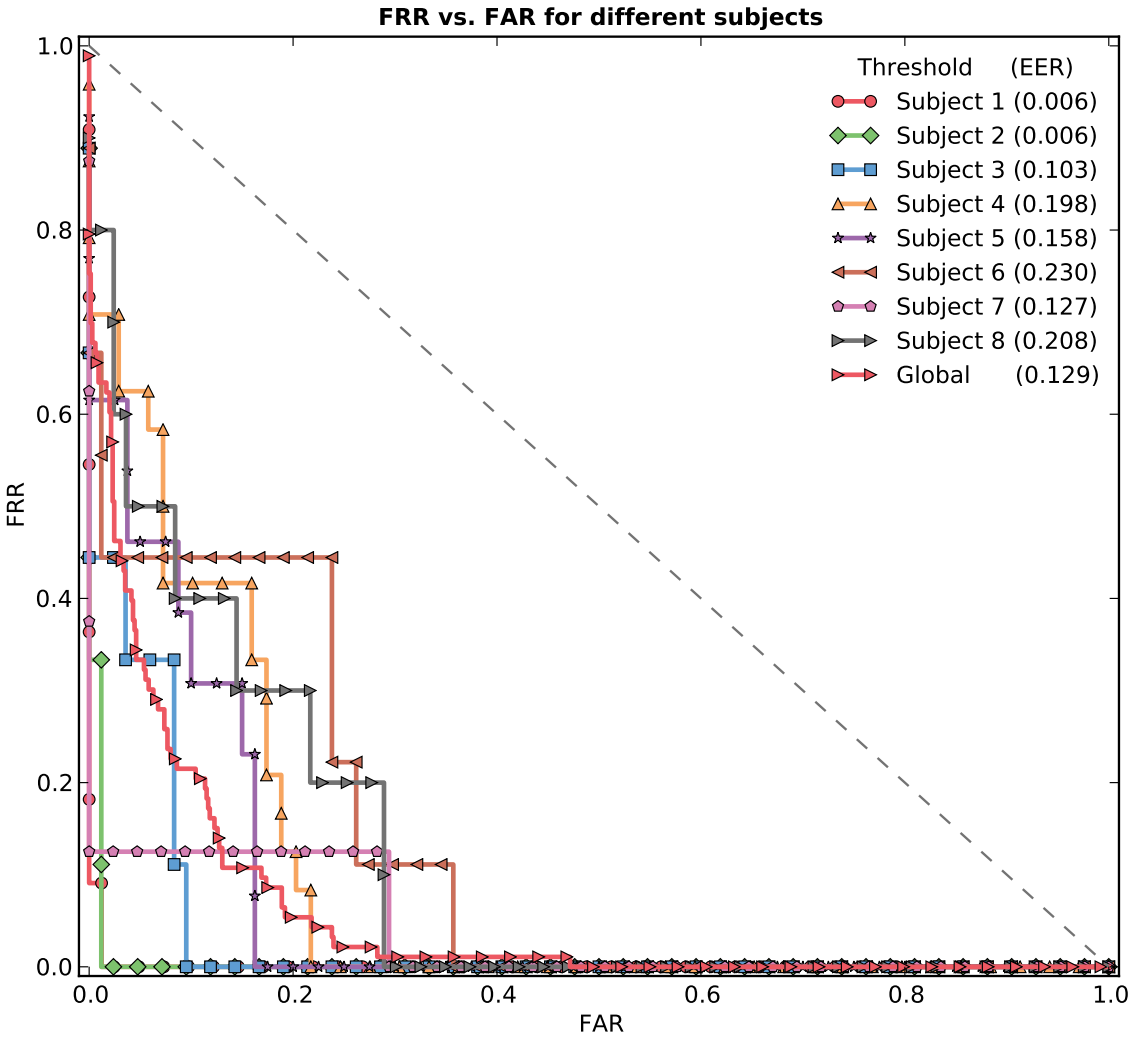

Figure 3.7.9: The recognition rates for each subject in a verification setting. The average per-subject EER was $12.90\,\% \pm 8.15\,\%$. A global threshold achieves a EER of $12.9\,\%$. The threshold, however, can be changed to suit the needs of the application (i.e., more false-negatives or more false-positives).

remained still with their wrist supported on a desk. For the walking task, the subject was asked to walk from when collection started to when it ceased as they normally would. And in the running task, the subject was asked to run in place with exaggerated arm movements. We captured five samples from each task, and averaged over the samples.

Figure 3.8.1 shows the results of this experiment as impedance magnitude plots for each task. While the samples in the stationary and walking scenarios have similar shapes, they have different offsets. The relative ordering of the bioimpedance samples in terms of magnitude is mostly preserved. The running task, on the other hand, varies significantly from the stationary task. Because vigorous motion affects bioimpedance measurements in our device (perhaps due to electrode contact since our bracelet is top-heavy due to the Shimmer being firmly attached near the electrodes), we can account for this effect by ensuring bioimpedance measurements are only taken when the subject is mostly stationary.

3.8.2 Orientation

The orientation of the wristband matters because the absolute position of electrodes will affect the pathway of the stimulus current. For this reason, the trained model will be sensitive to the orientation of the device. We measured this change by rotating the sleeve around a user's wrist and capturing bioimpedance measurements after each rotation. We rotated the device 45° around the wrist fully five times to record five samples and finally averaged over the samples.

Figure 3.8.2 shows the results of this experiment in the form of impedance magnitude plots for each rotation from 0° to 315°. While the general shape of the bioimpedance measurements is preserved, the offset of the measurement changes with each rotation as expected. For rotations that are exactly 180° apart, the curves

generally mirror each other. We expected this result because the magnitude of impedance has no polarity. While our measurements are sensitive to orientation about the wrist, we believe they can be compensated for by either detecting orientation or by using different features that are invariant to rotation.

3.8.3 Environmental effects

We hypothesize that some environments might have effects on bioimpedance. Humidity, for example, may affect the sensor due to the decreased resistance at the contact site between the skin and electrode. To simulate this condition, we collected bioimpedance measurements from three types of environments: a low humidity control, high humidity, and a water-soaked bracelet. For each environment, we captured five samples and averaged over them.

Figure 3.8.3 shows the results of this experiment for the different environments. Humidity has a tendency to compress (i.e., their offsets are clumped together) and reduce the magnitude of bioimpedance measurements overall. The relative ordering is not preserved except for the highest and lower bioimpedance measurements. In the case of a wet bracelet, bioimpedance measurements are useless. We need to account for environmental effects due to humidity and moisture. It would be easy to add a humidity sensor to the Shimmer, and use humidity as a feature in a training model. We could easily detect a wet bracelet (or other short between electrodes), and exclude those measurements. Such methods for compensating for these conditions remain future work.

3.8.4 Longitudinal effects

To understand the longitudinal recognition rates of bioimpedance, we collected 10 additional bioimpedance samples from Subjects 1, 4, and 5 140 days after their

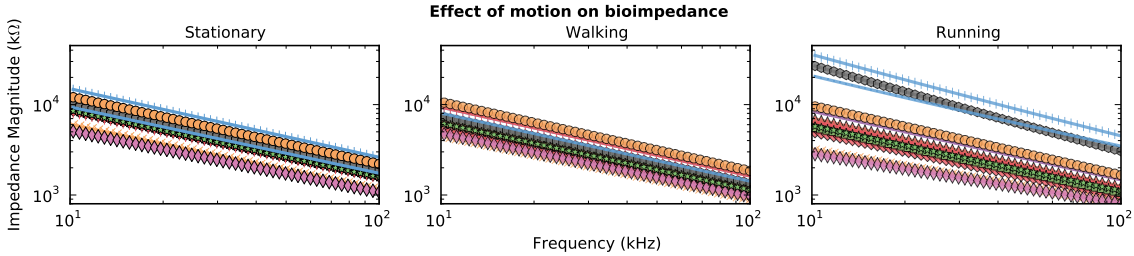

Figure 3.8.1: Bioimpedance plots for different types of motion: stationary, walking, running. Although the shape of the curves remain similar, the offsets change as the amount of motion increases.

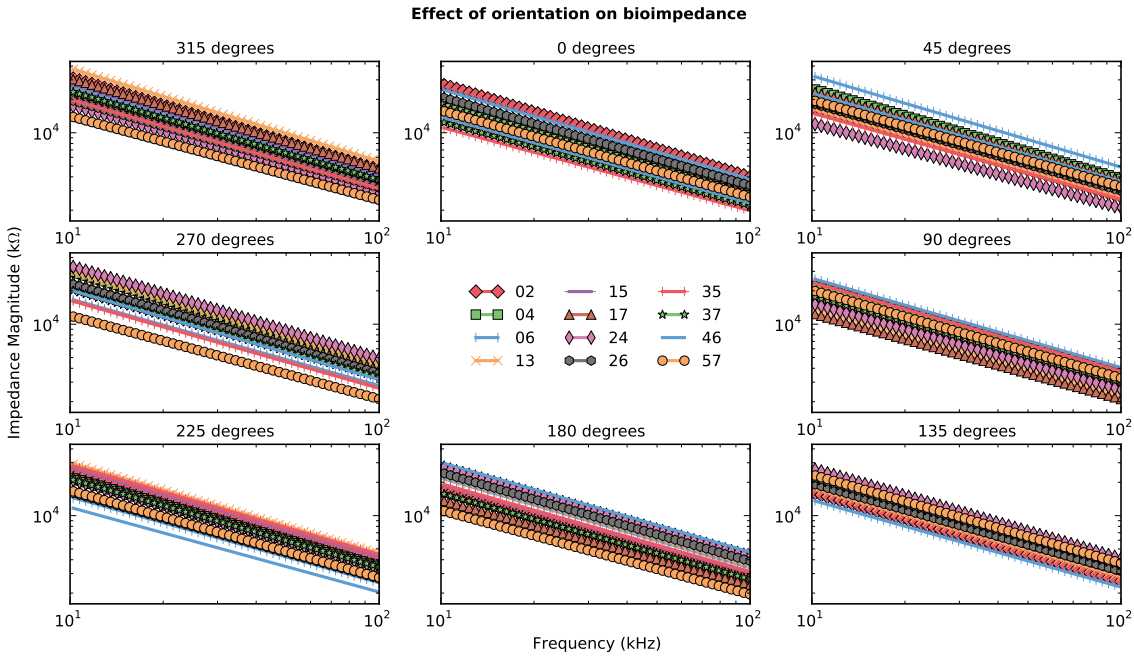

Figure 3.8.2: Bioimpedance plots for different device orientations rotated medially about the wrist every 45°.

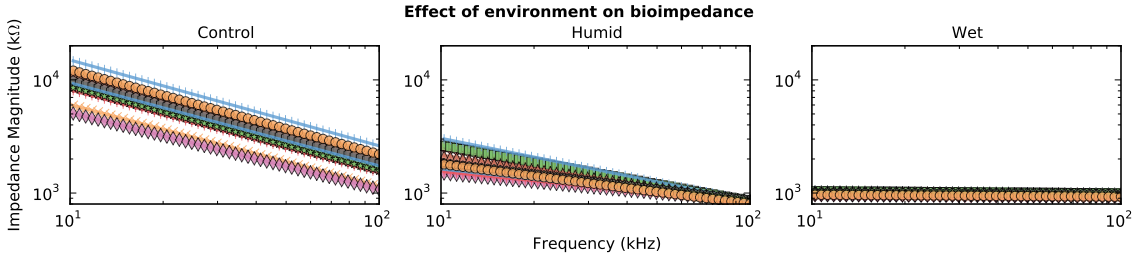

Figure 3.8.3: Bioimpedance plots for different simulated environments: control, humid, and a wet bracelet. Bioimpedance measurements are very sensitive to humidity and moisture.

initial enrollment in our wearable study. We ran a hold-out validation where the testing dataset is equal to these new bioimpedance samples for Subjects 1, 4, 5 or the last 10 % for every other subject. We trained a GMM for each subject similar to our verification evaluation.

Figure 3.8.4 shows the results of this longitudinal verification validation. The per-subject threshold EER was 9.57 % $\pm$ 8.12 % while the global threshold EER was 13.3 %. The recognition rates of Subjects 1 and 5 are similar to their recognition rates in the initial verification evaluation, while Subject 4 performed better than their initial rates. 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 indicate that a subject's bioimpedance remains relatively stable enough to be verified at least 4.5 months later.

3.9 Universality, acceptability, and circumvention

Finally, we come to the universality, acceptability and circumvention of bioimpedance. For these characteristics, we offer qualitative arguments for why we believe bioimpedance has these characteristics. More formal studies of these characteristics are left for future work.

3.9.1 Universality

We assume that every person has a wrist where we can measure bioimpedance. Of course this is true for most people. Our technique could be used at other locations on the body, but is unclear how one could create an unobtrusive device that can measure bioimpedance at, say, the chest. Ideally, the current should pass through as

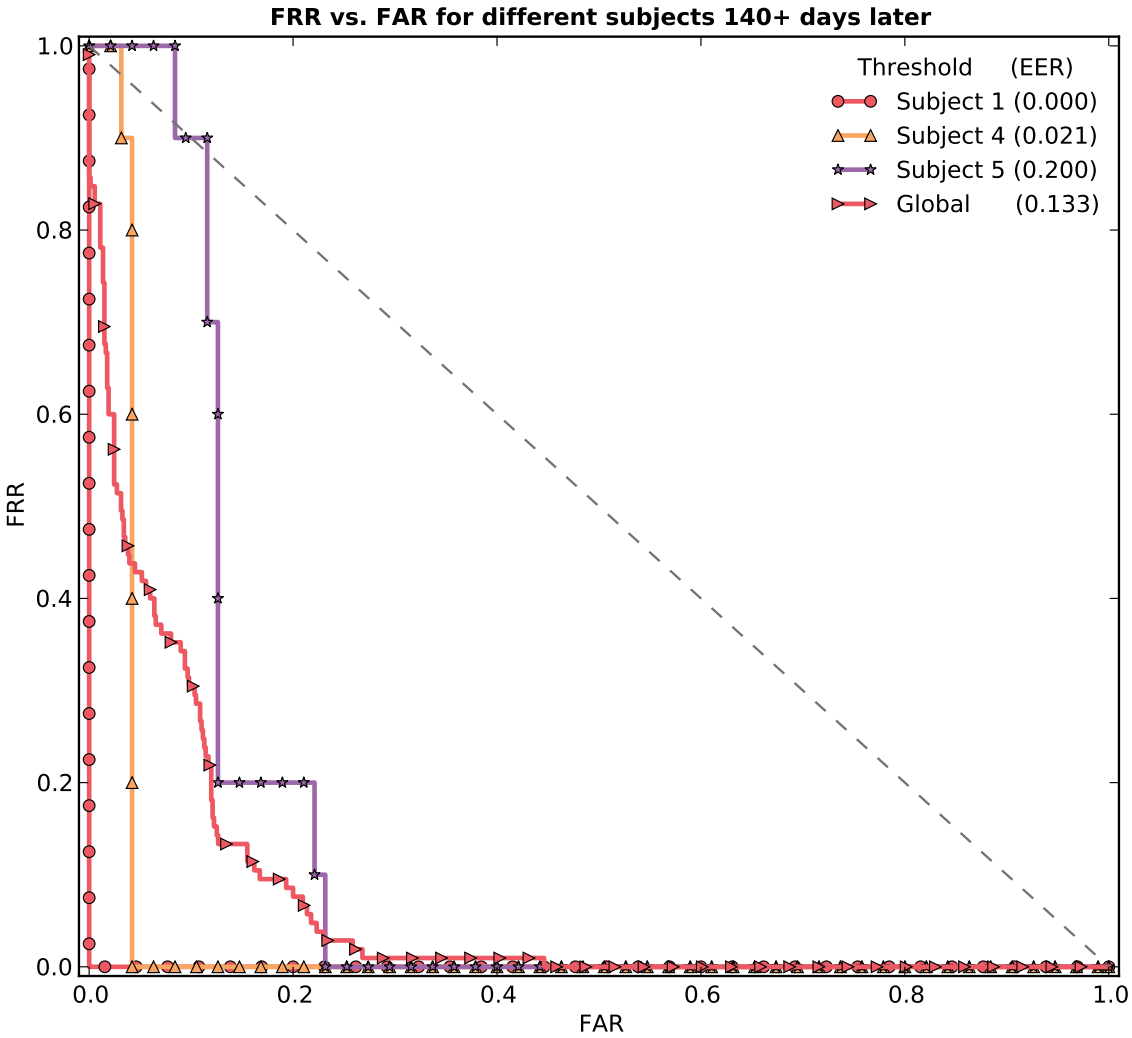

Figure 3.8.4: The longitudinal recognition rates of bioimpedance. We collected data from 3 subjects 140 days are their initial enrollment in our wearable study. The per-subject threshold EER was $9.57\,\% \pm 8.12\,\%$ while the global threshold EER was $13.3\,\%$, indicating results similar to our wearable verification evaluation.

much tissue as possible, and, in the case of the chest, this should be from the anterior to posterior part of the body. Having electrodes on the chest and back would require some sort of wire, which could be intrusive. For some locations, like the ankle, it is trivial to use our wearable device in those locations.

3.9.2 Acceptability

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. Anecdotally, many of our subjects found the sensor to be comfortable yet a little constricting. While comfort might affect the adoption of bioimpedance, we believe this obstacle can be overcome with further engineering of the form factor.

Our current form factor is a custom-sewn Lycra sleeve. Manufacturing these sleeves is not ideal. The electrodes were manually sewn into the fabric and much pain was taken to ensure the wires to the electrodes could withstand any torsion and tension.

With 3D printing becoming more prevalent, it is now possible to custom-print 3D models. By using this technology to print a custom-designed wristband for the specific geometry of the wearer, electrodes can always have good contact with the skin while being minimally constricting. Mayton et al. [108], for example, describe such a 3D printed wrist-worn device.

3.9.3 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 active attacker would have to model the physiology of an enrolled user's wrist to succeed. To model a user's physiology, they would need to either get access to a subject's bioimpedance data or measure their bioimpedance data. Like anything

that can be used for authentication purposes (e.g., a password), that piece of data should be kept private. The same holds true for bioimpedance data. However, using bioimpedance also raises the bar for an active attacker by forcing them to measure a subject's bioimpedance. To do so, they would need to attach a bioimpedance sensor to a person's wrist. With a little education, users can learn to not allow such attacks to happen.

3.10 Related work

This work is the first research to use bioimpedance as a biometric. It has long been used to measure a person's body fat percentage since they are proportional to each other [51]. Ailisto et al. [8] used bioimpedance and body weight to reduce error rates of fingerprint biometrics from 3.9 % to 1.5 %. Others have used bioimpedance to detect liveness in the case of fingerprint biometrics, since a fingerprint reader can be easily fooled; Martinsen et al. [99] present such a system to detect liveness. Such techniques could be incorporated into our system as well.

There are many other biometrics people have used to recognize users [78]. Face recognition [92] uses a camera and locates common features on a person's face or determining how a person's face relates to other faces. Fingerprint recognition [95] uses a sensor and extracts features based on the ridges and valleys of a person's finger. It is probably the most common biometric used today. Gait recognition [114] uses a camera and extracts features from a sequence of images of a person walking. Iris recognition [29] uses a camera and extracts features from the appearance of a person's iris. Vein pattern recognition [155] uses an infrared sensor and extracts features from an image of veins from some part of a person's body. Finally, speaker recognition [19] uses a microphone and extracts features from a person's voice. Without liveness tests, most of these biometrics can be fooled, while others can be

intrusive to measure.

Srinivasan et al. [136] 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.

Sriram et al. [137] provide a method to identify patients using ECG and acceleration data for remote health monitoring. While ECG has proven to be useful for authentication, they observe that these methods do not perform well in the real world because physical activity perturbs the ECG data. By employing an accelerometer to differentiate physical activities, they can use ECG data from those physical activities to authenticate patients. We both make the observation that "the monitoring system needs to make sure that the data is coming from the right person before any medical or financial decisions are made based on the data" [137]. Sriram's framework for using other sensors to compensate for certain affects could be similarly applied to our method.

Finally, others have used capacitive sensing to differentiate subjects using a capacitive touchscreen. Vu et al. [149] 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, for example, encode the ring wearer's identity into this signal. 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) limit the amount of information that can be communicated. Harrison et al. [67] show how to differentiate between subjects using a capacitive touchscreen. Rather then 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.

3.11 Limitations

We indicated above that 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 affect bioimpedance measurements, 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, we could use kinematic sensors to determine the orientation of the band and compensate for different orientations. For example, the accelerometer can be used to determine the orientation of the device relative to gravity. This would not, however, detect the case that the device is put on backwards (i.e., electrodes are swapped in direction) or on the wrong wrist. In the general case, it might be possible to engineer rotation- and reflection-invariant features. For example, an electrode configuration could be assumed to be rotated by duplicating the measured impedance, while reflection could be handled by swapping measurements at different electrode configurations. The details of such a super feature vector are left for future work.

We did not explicitly consider potential 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 tight model for each subject. It might be the case, for example, that a change in blood glucose adversely affects bioimpedance at the wrist.

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.

We evaluated the feasibility of this method as tested on eight subjects each wearing the device for an average of 10 hours. Although we achieve good performance according to our metrics, this only gives us confidence that bioimpedance is stable over the time periods for which we collected data. To be truly confident in this method, however, we need to explore the stability of bioimpedance over years, to sample a larger number of subjects, and to explicitly and implicitly explore a broader range of environmental conditions.

Our method will suffice for 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, perhaps the bracelet is used as part of a biometric authentication system, or the person wishes to have data collected under someone else's identity. We believe, however, that it would be difficult to forge another person's bioimpedance. In principle, an adversary could measure 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. See Chapter 6 for more discussion about this issue.

Finally, there are several ways the current design could be optimized for lower cost or reduced energy consumption. The current wristband includes 8 electrodes, but

we can remove some electrodes at the cost of reduced recognition rates. Furthermore, we measure bioimpedance across a wide sweep of 50 frequencies, but we could optimize the number of frequencies necessary and decrease both the energy and time needed for each measurement. A low-power human presence detection method would also allow us to sense only when a wearer is detected.

3.12 Summary

We present a method for continuous identity verification of a bracelet's wearer, using bioimpedance as a biometric. We constructed a prototype device with electrodes embedded on the inside of the wristband, connected to a custom impedance-measuring device built into a Shimmer research platform. We evaluated the ability of this device to correctly identify its wearer, within a cohort of 8 subjects. We evaluated the stability of this biometric under various environmental conditions, and over an average of 10 hours per subject. We found that, depending upon the amount of training data, the device was successful in recognizing its wearer with 98 % BAC in a cross-validation study. In a hold-out study, the device was successful at recognizing its wearer with 81 % BAC when trained with at least 10 bioimpedance samples. In a verification study with data collected from subjects more than 4 months later, we achieved an EER of 13.3 %. For small cohorts of subjects, we believe bioimpedance can be used as a passive biometric, but it remains to be seen how well bioimpedance performs for larger cohorts. Likewise, there are significant barriers to overcome in compensating for environmental and longitudinal physiological changes.

4

Verifying whether sensors are on the same body

In WBANs there are typically many SNs transmitting sensed data to a specific MN. It's easy to imagine a cryptographic scheme that would ensure the confidentiality and integrity of all communications. But even with such a security mechanism, the assumption that the SNs are all attached to the same human body as the MN remains. This chapter focuses on this problem.

This problem matters for at least one important reason. If we expect physicians to make decisions about health data collected from WBANs, then they will need

some confidence that all the data they are examining has come, at the very least, from the same body. It is risky to make medical decisions without said confidence, and until such measures are in place, we should be wary of using data from WBANs.

4.1 Introduction

In this chapter, we focus on the *weak* version of the one-body verification problem. The weak version, unlike the strong version, requires determining whether the sensors are on the same body. This might seem like an easier problem than the strong version, but it is more difficult. Rather than finding a unique characteristic of a specific individual (as in a biometric), we need to find something that is not subject specific yet distinguishes sensors on one body from those on another body. For example, we could compute the distance between the MN and each SN using time-of-flight algorithm. However, for the distances we are interested in (less than a meter), the accuracy of these kind of localization methods is just not good enough (error rates of 0.5 m according Hamida et al. [65]). But there are at least two others ways to solve this problem: using biometrics or using statistics.

Our method described in Chapter 3 trivially solves this problem. However, not all types of SNs will produce data that is suitable for biometric recognition. It is unclear, for example, how an SN could recognize its wearer when it only has a temperature sensor. Even supposing that an SN did have a suitable biometric sensor (like the one discussed in Chapter 3), its not clear that biometric could work for most locations on a person's body. It is unknown, for example, whether bioimpedance is a suitable biometric when measured at a person's waist. More so, an SN designed to be carried in a person's pocket might never come in contact with the wearer's skin, a requirement for the use of bioimpedance. Other biometrics have similar limitations. Finally, the use of a biometric requires non-trivial hardware and software

requirements that might not fit the form-factor of the SN or even make economic sense. Bioimpedance, for example, requires an impedance analyzer and electrodes. Thus, using a biometric-based solution to determine whether an SN is on the same body as an MN will not work for all use cases.

We offer a statistics-based solution. The idea is that if the MN and SN can measure some quantity at the same time, then we can use statistics to determine whether there is a relationship between these two quantities, that is, some correlation between these two measured quantities over time. Others have shown how to accomplish this goal using physiological-based quantities via photoplethysmograph [144] and electrocardiography [146]. This kind of solution does not work in all cases because they typically require skin contact to measure the physiological values.[1]

Instead, we take a biomechanical approach and measure the acceleration and orientation of the subject while they are walking. In the ideal circumstance, we would find some way of measuring acceleration and orientation using the existing sensors of the SN. Depending upon the sensors present on the SN, this may not be possible. Thus, we propose the following compromise: every sensor device will include an accelerometer and gyroscope in addition to its primary sensor (ECG, blood pressure, etc.). Since accelerometers and gyroscope are tiny, cheap, and require little energy to operate, this is a reasonable assumption. Thus, rather than taking advantage of different relationships for each kind of sensor data, placement, and usage conditions, we only need to find a relationship for the acceleration and orientation data that answers the question: are these devices attached to the same body?

[1] There are non-contact methods of measuring some physiological values [120, 118, 119, 157, 148, 56, 49], however they require a camera (visible or infrared) pointed at a non-moving subject. Such a requirement is both resource intensive and unrealistic for our use cases.

4.2 Sensors

To maximize the generality of our solution, we require each SN to have an accompanying accelerometer and gyroscope. Today, many SNs already include an accelerometer because it is often useful to know how a device is oriented relative to gravity.

4.2.1 Accelerometers

Accelerometers measure acceleration using the principle that any mass m will experience a force F proportional to the amount of acceleration a is it experiencing according to Newton's second law of motion (under the assumption of non-relativistic accelerations, which humans will probably never experience):

$$F = ma$$

Instead of measuring the conventional rate of change of velocity, accelerometers measure force per mass, also known as specific weight when the mass is constant. That is, they measure how much force is applied to a known mass. Although acceleration is typically measured in meters per second (m/s^2), we deviate from the International System (SI) of units to distinguish what is being measured. Since the typical force we experience is gravity, we measure acceleration in gravities, or g-force.[2] All stationary masses on the earth experience 1 g-force, while a mass in free-fall would experience no g-force (until terminal velocity is reached).

[2] The use of g should not be confused with the SI unit g for grams. The use of grams to describe mass makes no appearance in this thesis.

Hardware

The most common type of accelerometer today is the capacitive, micro-machined accelerometer. In our experiments we used two of these types of accelerometers: the Freescale MMA7260Q [52] and MMA7361L [53]. This type of accelerometer operates by detecting the deflection of a small proof mass located between two fixed beams. Since the proof mass is attached to a small spring, we can equate Hooke's law, $F = -kx$ (where k is a constant factor related to the spring and x is the displacement of the spring), to Newton's second law, $F = ma$. When the proof mass experiences an acceleration, a proportional force causes it to deflect in the direction opposite of the direction of the acceleration. To measure this deflection, a simple circuit is used. By design, each beam forms a capacitor with the proof mass, and when the proof mass moves it induces a change in capacitance:

$$C = \epsilon \frac{A}{d}$$

Since the dielectric constant ϵ and area A are a constant quantities, capacitance is inversely proportional to distance d between the proof mass and the fixed beam. Thus, this change in capacitance results in a change in voltage which is related to the amount of displacement of the proof mass experiencing a force. Such an accelerometer is limited to operation in the axis it is oriented. To allow sensing in all 3 dimensions, three such accelerometers must be machined orthogonal to each other.

Calibration

Since these accelerometers output a voltage change corresponding to acceleration, it is necessary to calibrate the output. Even if we used only one kind of accelerometer, it would still be necessary to calibrate each accelerometer due to variances in

manufacturing. Additionally, calibration also allows us to compare accelerations from both types of accelerometers.

To calibrate a specific axis x, the accelerometer is placed such that the axis aligns with direction of gravity and allowed to remain stationary (i.e., the proof mass is allowed to experience 1 g-force). The voltage output is then recorded. The accelerometer is then flipped 180° along that axis, allowed to remain stationary again, and the voltage is recorded. These voltages correspond to a minimum V_{min} and maximum V_{max} voltage experienced under gravity alone. Because these kind of accelerometers exhibit linear voltage change per g-force, we therefore have a system of linear equations that equates voltages V_{min} and V_{max} to g-forces -1 and 1, respectively:

$$1 = m_x V_{max} + b_x$$

$$-1 = m_x V_{min} + b_x$$

Solving this system of linear equations for coefficients m_x and b_x yields:

$$m_x = \frac{2}{V_{max} - V_{min}}$$
$$b_x = \frac{-V_{max} - V_{min}}{V_{max} - V_{min}}$$

After calibration, the corresponding g-force a_x for some new voltage V_x can be computed as:

$$a_x = V_x m_x + b_x$$

This procedure is repeated for each axis of the accelerometer.

Cost

The selection of these specific accelerometers was driven more by availability and ease of use than unit cost. For example, it possible to purchase a Freescale MMA845xQ 3-axis accelerometer for $0.56 (in quantities of 1000) [44]. This particular model consumes 1.4 μA in standby mode, as low as 6.6 μA in low power mode, and 27 μA in normal mode [54]. So while our choice in accelerometers was not optimized for energy-usage profile or cost, it is certainly possible to do so.

4.2.2 Gyroscopes

Vibrating structure gyroscopes measure angular velocity using the principle that a mass m moving at velocity v will experience an orthogonal inertial force F_C proportional to the amount of angular velocity ω it is experiencing according to the Coriolis effect:

$$F_C = -2m\omega \times v$$

By vibrating the proof mass m at a fixed velocity v, the gyroscope can measure this force F_C using the same principle an accelerometer uses. The SI unit for angular velocity is radians per second (rad/s).

Hardware

In our experiments we used two gyroscope: the InvenSense ISZ-500 [76] and IDG-500 [75]. The ISZ-500 is a single gyroscope oriented along the z axis, while the IDG-500 is a dual gyroscope oriented along the x and y axes. Thus, when combined they allow us to capture orientation in all 3 dimensions. Like the accelerometer, angular velocity induces a change in capacitance which results in a change in voltage.

Calibration

Since these gyroscopes output a voltage change corresponding to angular velocity, it is necessary to calibrate the output. Even if we used only one 3-axis gyroscope, we would still need to calibrate each axis of the gyroscope due to variances in manufacturing. Additionally, calibration allows us to compare angular velocities from any other kind of gyroscope.

To calibrate a specific axis x, the gyroscope is placed such that the axis is allowed to remain stationary (i.e., the proof mass is allowed to experience no angular velocity). Then, the output voltage $V(t)$ is continuously recorded for all times $0 < t < T$ while the gyroscope is manually rotated ϕ radians about the axis. The integral of the output voltage is then proportionally related to rotation angle ϕ via:

$$m_x = \frac{1}{T}\sum_{0}^{T} V(t)$$
$$\phi = s_x \int_0^T (V(t) - m_x)\, dt$$

Solving for the scaling factor s_x is simple:

$$s_x = \frac{\phi}{\int_0^T (V(t) - m_x)\, dt}$$

We approximate this definite integral using the trapezoidal rule. In our calibration routines, we rotated the device exactly 360°. Thus, given the calculated coefficients m_x and s_x and a new voltage V_x from the gyroscope, the corresponding angular velocity ω_x is computed as:

$$\omega_x = s_x(V_x - m_x)$$

This procedure is repeated for each axis of the gyroscope.

Cost

Again, the selection of these specific gyroscopes was driven more by availability and ease of use than unit cost. However, one can purchase an ST L3GD20 3-axis gyroscope for $2.99 (in quantities of 15000) [43]. It consumes 2 mA in sleep mode, and 6.1 mA in normal mode [138]. It is also possible to purchase a combined Bosch BMI055 6 degree-of-freedom inertial measurement unit that contains a 3-axis accelerometer and 3-axis gyroscope that costs $3.04 (in quantities of 10000) [42] and consumes 25 μA in suspended mode, and 5 mA in normal mode [25].

4.2.3 Usage

To verify whether SNs are on the same body as the MN, our approach is to find a relationship in the acceleration and orientations that the MN and SN experience for a given time period. Our intuition is that if an SN is on the same body as the MN, then (at a coarse level) the accelerometers and gyroscopes on both the SN and MN must experience similar accelerations and changes in orientation. If a person is motionless (and hence no acceleration except gravity or change in orientation), then there is no information we can extract from the accelerometer or gyroscope to determine whether those SNs are on the same body.[3] Thus, we must require that the person do some kind of activity.[4] Luckily, there are a several of activities a person can perform that cause changes in acceleration and orientation.

Not all activities are permissible for our approach. First, it must be something that humans are capable of doing with their body. This rules out activities like flying

[3] This is not theoretically true. A person on top of a mountain above sea level experiences less gravitational acceleration than a person at sea level. This relationship is the well known physical quantity $F = G\frac{m_1 m_2}{r^2}$ where G is the universal gravitational constant, m_1 is the mass of the earth, m_2 is the proof mass in the accelerometer, and r is the distance between them. So one could, in theory, distinguish between sensors on bodies at different elevations using an accelerometer.

[4] This might be a limiting factor of our method for some kinds of subjects. For example, an elderly patient confined to a bed would experience few accelerations or changes in orientation.

for obvious reasons. First, the person's whole body must experience the acceleration and change in orientation. This rules out activities like using a computer because typically only the upper body is moving. Second, the activity must be something a person does periodically throughout their day. This rules out activities like jumping or running. Third, it must be something that is suitably unique to any individual. This rules out activities like driving a car, since any person in that car will be experiencing the same acceleration and change in orientation. Finally, it must be an activity that is computationally fast and easy to detect using an accelerometer and gyroscope. The reason for this constraint will become obvious soon.

The obvious type of activity that fulfills these constraints is walking. When walking, a human body is largely rigid in the vertical direction. Although our limbs do bend, we hypothesize that the vertical acceleration (i.e., the acceleration relative to gravity) experienced by sensors placed anywhere on a walking body should correlate well. As one foot falls, that side of the body experiences a downward acceleration due to gravity, followed by an abrupt deceleration when the foot contacts the ground. Sensors on one side of the body should experience a similar vertical acceleration, while sensors on the other side of the body will experience the opposite. We should expect positive correlation for one side of the body, and an inverse correlation on the other side. This observation is complicated by the fact that it is difficult to extract the vertical acceleration component without knowing the orientation of the sensor. Furthermore, although the signal can be very noisy, the accelerations due to walking are likely to dominate the accelerations due to intra-body motion (such as arm swings or head turns) and we should be able to reliably make a determination that the supposed suite of sensors are on the same body. Finally, there is an existing body of literature that shows how to do activity recognition given user-annotated data [16], and even on a mobile-phone-class device [28]; these techniques are particularly good at detecting when a user is walking. Our approach, therefore, is to

detect periods when a user is walking by monitoring the MN's accelerometer data periodically; when the data indicates the user is walking, we then use our method.

4.3 Method

As stated previously, we assume each SN has an accompanying accelerometer and gyroscope, which we will use in our method. In this section we describe how our method works for acceleration only, however the same procedure can be applied to the gyroscope data as well. That is, one can substitute "gyroscope" wherever there is a mention of "accelerometer" since the signal-processing techniques are agnostic to the type of signal.

Consider a signal s sampled at some frequency such that:

$$s = \{(x_1, y_1, z_1), (x_2, y_2, z_2), \ldots\}$$

where x_i, y_i, and z_i are the three axes of the instantaneous acceleration, relative to gravity, at time i. Because sensors might be mounted in different orientations, or might be worn in different orientations each time they are worn, we discount orientation by using the *magnitude* of the acceleration. Figure 4.3.1 shows that the magnitude exposes the overall walking motion well. Thus, we compute the magnitude of all three axes for all samples in s:

$$m_i = \sqrt{x_i^2 + y_i^2 + z_i^2}$$

This measure gives us the rate of change of speed over time for that particular SN.

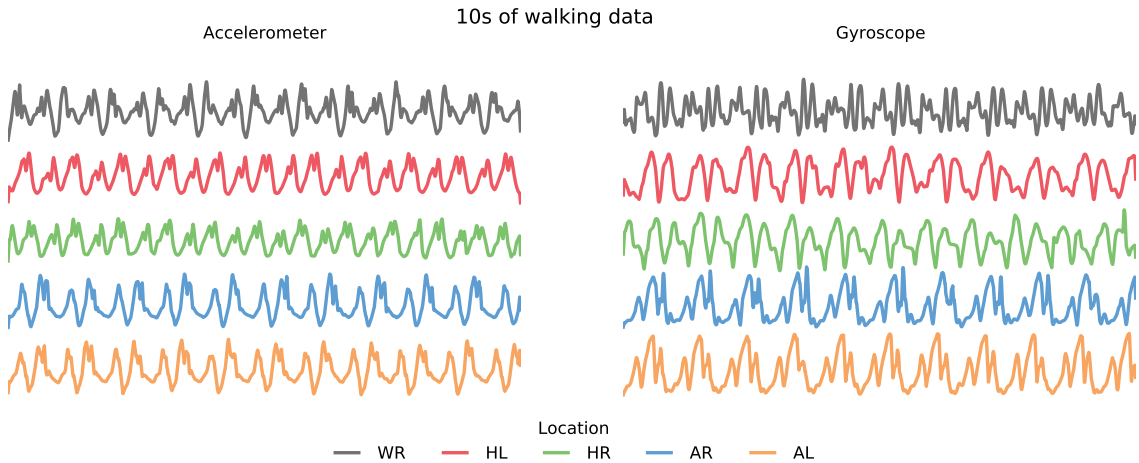

Figure 4.3.1: Ten seconds of accelerometer and gyroscope magnitude data for each position on the body for one subject. This subject took about 20 steps.

4.3.1 Feature extraction

We partition this orientation-agnostic signal $\{m_1, \ldots, \}$ into feature window lengths w with some desired step size s. The step size allows overlapping windows. For each feature window $j = 0, 1, \ldots$ comprising $\{m_{js+1}, \ldots, m_{js+w}\}$, we extract seven common features:

mean $\frac{\sum_{k=0}^{w} m_{js+k}}{w}$

variance $\frac{\sum_{k=0}^{w} (m_{js+k} - \textbf{mean})^2}{w}$

standard deviation $\sqrt{\frac{\sum_{k=0}^{w} (m_{js+k} - \textbf{mean})^2}{w}}$

mean absolute deviation $\frac{\sum_{k=0}^{w} |m_{js+k} - \textbf{mean}|}{w}$

inter-quartile range $\mathrm{P}_{75}(\{m_{js+1}, \ldots, m_{js+w}\}) - \mathrm{P}_{25}(\{m_{js+1}, \ldots, m_{js+w}\})$ where $\mathrm{P}_n(S)$ is the nth percentile of S.

power $\frac{\sum_{k=0}^{w} m_{js+k}^2}{w}$

energy $\sum_{k=0}^{w} m_{js+k}^2$

Collectively, these seven values form the *feature vector* $F_j = (f_j^1, f_j^2, \ldots, f_j^7)$. We chose these features primarily because others [103, 123] have used these features successfully to detect physical activities, and we hypothesize they would similarly be useful for our problem. We also explore subsets of these features. If they can capture the physical activity of walking and we examine the correlation of these features, we should expect them to correlate if and only if they are attached to the same body.

4.3.2 Coherence

Coherence is a measure of how well two signals correlate in the frequency domain. More precisely, it is the cross-spectral density of two signals divided by the auto-spectral density of each individual signal. Like Lester et al. [91], we approximate coherence by using the magnitude-squared coherence:

$$C_{xy}(\phi) = \frac{|S_{xy}(\phi)|^2}{S_{xx}(\phi)S_{yy}(\phi)}$$

In the equation, x and y are the signals, S_{xy} is the cross-spectral density between signals x and y, S_{xx} is the auto-spectral density of signal x, and ϕ is the desired frequency. Cross-spectral density is calculated by the Fourier transform of the cross-correlation function. If x and y are well correlated at some frequency ϕ, then $C_{xy}(\phi)$ should be close to 1.

Because we are interested in many frequencies, we compute the normalized magnitude-squared coherence up to some frequency $\phi_{\max}$:

$$N(x, y) = \frac{1}{\phi_{\max}} \int_0^{\phi_{\max}} C_{xy}(\phi) d\phi$$

We chose $\phi_{\max} = 10$ because, as Lester et al. notes, "human motion rests below the 10Hz range" [91].

In addition, to compute the cross-spectral density over different frequencies, it

is necessary to window the signals x and y. We choose a window of length equal to one-half of the feature window length with no overlap between adjacent windows.

4.3.3 Feature coherence

Given two sets of feature matrices $A = (F_1, F_2, \ldots)$ and $B = (F_1, F_2, \ldots)$ with entries F_j as described above, we want to determine how well A and B are correlated. Here, A and B represent the *feature matrices* extracted from the accelerometer data of the MN and SN.

We apply coherence to the feature matrices in the following manner. For some window length c (the feature coherence window) with some step size to allow overlap, we compute the normalized coherence of A and B as such:

$$N_k^{AB} = \left\{ N(A^1_{k\ldots k+c}, B^1_{k\ldots k+c}), N(A^2_{k\ldots k+c}, B^2_{k\ldots k+c}), \ldots, N(A^7_{k\ldots k+c}, B^7_{k\ldots k+c}) \right\}$$

where $A^1_{k\ldots k+c} = \{f^1_n \in A : k \leq n < k + c\}$ is the window of c samples from the first feature of A. That is, we take each feature (i.e., a column of the matrix) of A and the corresponding feature of B, and compute the normalized coherence using c samples (i.e., the rows of the matrix). At this stage, we are left with a matrix of normalized coherences for each feature and window k.

Because we want to capture how the two signals are related over time, the coherence window c should be sufficiently large to capture periodicity in the features. Because the typical walk cycle is on the order of seconds, it is advisable to chose a coherence window on the order of several seconds.

As mentioned, we can also apply this same method to the data from the gyroscope. Rather than having one feature coherence matrix, we would have two such matrices: one computed from the accelerometer data and the other computed from the gyroscope data. These two matrices can then be combined together.

4.3.4 Verification

To account for the many positions an SN might be placed on the body, we collect data from several locations. In our method, we compare the MN's accelerometer data to each other SN's accelerometer data. That is, the MN acts as a reference accelerometer, to which every other SN must correlate using the method described above. For a given set of locations and one reference location, we compute the feature coherence of each location A relative to the reference location B. In our experiments, we compute the coherence of the HR and WR; HL and WR; AL and WR; and AR and WR as shown Figure 2.1.2. When we compute the coherence for one user, this method yields feature coherences of the sensor on the same body, and we can label them as such. To yield feature coherences of sensors on different bodies, we take pairs of users and mix their locations. For example, at the waist and left hand there are two possible ways to mix up the sensors: Alice's waist and Fred's left hand, Fred's waist and Alice's left hand. By mixing locations for any pair of users, we can compute an equal number of feature coherences that are and are not on the same body, labeling them as such.

Given a set of feature coherences and their respective labels, we can train a classifier to learn a model that maps a feature coherence to a label. We examine several classifiers and many different parameters in our evaluations.

4.3.5 Smoothing

The classification method described above makes an instantaneous classification of a feature coherence for that particular coherence window. It is, however, possible to boost the classification rates by examining a series of classifications over time. For example, if over the course of three classifications, two classifications are positive and the third one is negative, we can use a simple voting scheme to smooth over these

mis-classifications. In the example, because most of the classifications are classified as on the same body, we assume the SN is on the same body for that classification window. We can empirically determine the best smoothing window size by varying the it and choosing the one that yields the best classification rates.

4.4 Exploratory evaluation

In this exploratory study, we sought to understand how well the accelerometer and gyroscope performed when used independently and together. This study is useful for several reasons. First, it tells us which sensor performs the best at determining whether two SNs are on the same body. If, for example, one of the sensors does not help determine this fact, then we could ignore that sensor completely. Doing so would allow us to reduce the total physical size of the required components, thereby reducing energy use, component cost, and code footprint.

4.4.1 Dataset

For our exploratory study, we collected data from two subjects at the five highlighted locations shown in Figure 2.1.2. We instructed each subject to walk a flat course of length approximately 200 m. The course consisted of three straight segements connected at right angles as shown in Figure 4.4.1a. We used the Shimmer platform [132] with an attached 9 Degree of Freedom sensing module capable of sensing acceleration (using the Freescale MMA7361L [53]), angular velocity (using the InvenSense ISZ-500 [76] and IDG-500 [75]), and magnetic field (using the Honeywell HMC5843 [72]), each in three dimensions. We did not use the magnetometer in our experiments, however. We wrote custom software to synchronize the Shimmers to a global clock with millisecond precision (see Chapter 5 for details) and collect data from the accelerometer and gyroscope at 250 Hz. Our device and data-collection

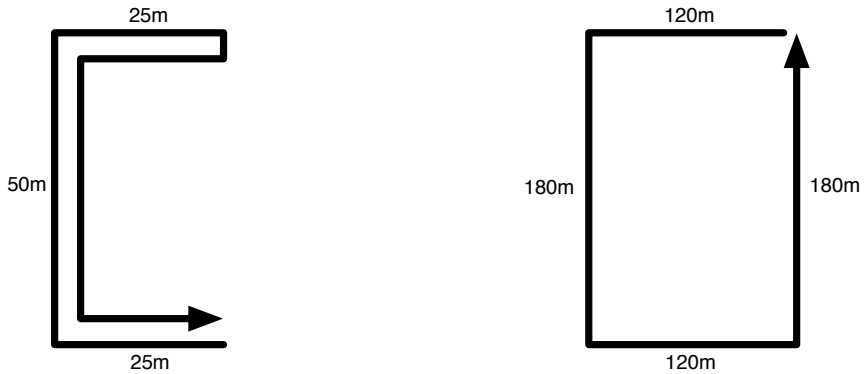

(a) Exploratory study **(b)** Dual user study

Figure 4.4.1: Abstract and idealized top-down view of the courses we asked subjects to walk. These are not to any scale.

protocols were approved by our Institutional Review Board. On average, it took the user about 2 minutes to walk the full course. While this time is short, this study was intended to be a controlled study on one aspect our of method.

Recall that our method works by comparing the accelerometer and gyroscope data from the SN with the MN's own accelerometer and gyroscope data. That is, we compare Subject 1's WR with their own AL, AR, HL and HR. Since we know these SNs are all on Subject 1's body, we label them positively. We collected data from a single user at a time in this exploratory study, thus we only have positive examples. However, we can create negative examples by mixing the two subjects' data together. That is, we can compare Subject 2's WR with Subject 1's AL, AR, HL and HR, and vice-versa. Such a scheme yields exactly 4 sets of positively labeled samples and 4 sets of negatively labeled samples per subject. Thus, this particular dataset is balanced.

4.4.2 Parameters

Our method requires us to choose a feature window length and step, feature coherence window length and step, a smoothing window length, and a classifier. For

simplicity, we represent these parameters as a tuple in the form `(feature window length, step; coherence feature window length, step; smoothing window length)`. For some figures, the smoothing window length is left out of the tuple. The meaning of these parameters is described below.

Feature window length and step

The length of the feature windows is specified as the number of samples. Ideally this length should be smaller than the sampling rate to capture features of a person's gait, but not so small as to miss these features. Anywhere from 100 ms to 500 ms is ideal since humans tend to take about 1 to 2 steps per second [156].

Because we sampled at 250 Hz, we chose feature window lengths of 25, 50 and 125. These were specifically chosen because they fall within the 100 ms to 500 ms range (i.e., they correspond to 100 ms, 200 ms, 500 ms), and they evenly divide the sample rate.

The feature window step is specified as a percentage of the feature window length. For simplicity, we examined steps of 50 % (i.e., half overlap) and 100 % (i.e., no overlap).

Feature coherence window length and step

The length of the feature coherence window is specified in seconds. This length should be chosen such that periodicity in a person's gait will be captured by the feature coherence. Since humans tend to take about 1 to 2 steps per second, this value should be on the order of several seconds. Since our sampling rate was 250 Hz, we chose feature coherence window lengths of 4 s, 8 s and 16 s. Again, for simplicity's sake, we chose feature coherence window steps of 50 % (i.e., half overlap) and 100 % (i.e., no overlap).

Smoothing window length

The smoothing window length is specified as a number of classifications. Ideally a classifier would make the correct classification each time. However, this is not always possible in practice. Depending upon how mis-classifications are distributed, a smoothing window can improve classification rates by smoothing over mis-classifications. We examined lengths of 1, 3, 5, 7 and 9. Because our smoothing procedure employs majority voting, we chose odd lengths to avoid ties.

Classifiers

Finally, we examined three different classifiers – SVM, NB, FOREST – as described in Section 2.4. We did not specifically tune any of these algorithms. In addition to these three classifiers we also present results for the RAND classifier as a baseline classification rate. The RAND classifier should perform around 50 % BAC on average.

4.4.3 Accelerometer-only validation

Recall that we collected both accelerometer and gyroscope data. For this first validation, we only examined the accelerometer data. That is, we fed the method described in Section 4.3 acceleration data only. Figure 4.4.2 shows results of a 10-fold cross-validation for a selected subset of the parameters using the described dataset. Each graph represents one classifier, while each line in the graph represents one particular setting of the feature window length and step, and the coherence feature window length and step. We varied the smoothing window and computed the BAC for each window. For all classifiers but the RAND classifier, as the smoothing window increased so does the BAC until it decreases once again due to over-fitting. Some parameter settings yielded poorer performance than other settings.

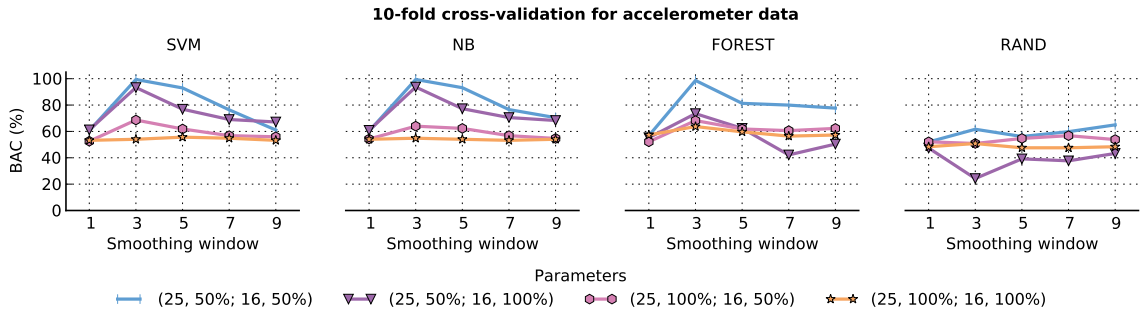

Figure 4.4.2: Ten-fold cross-validation using accelerometer data for different feature coherence parameters, classifiers, and smoothing windows. The NB classifier performs the best with a smoothing window of three and small windows and steps.

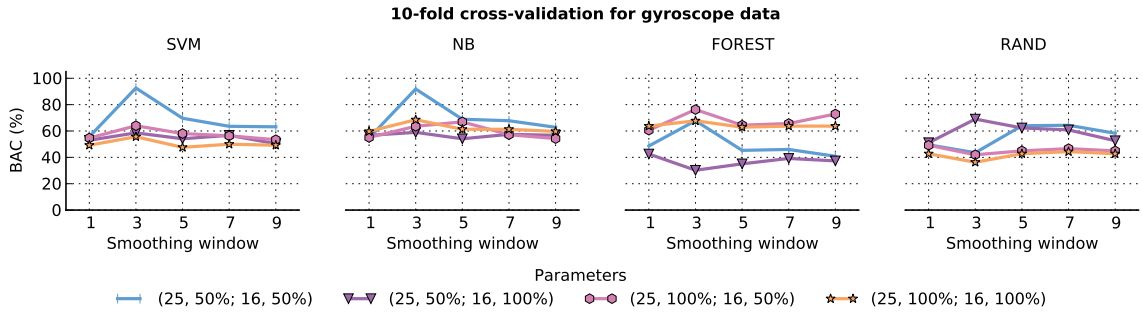

Figure 4.4.3: Ten-fold cross-validation using gyroscope data for different feature coherence parameters, classifiers, and smoothing windows. The gyroscope data alone performs worse than accelerometer data alone.

4.4.4 Gyroscope-only validation

For this next validation, we used only gyroscope data. Figure 4.4.3 shows the results of a 10-fold cross-validation for the same parameters chosen in Figure 4.4.2. Overall, using gyroscope data alone performed much worse than just accelerometer data. Since gyroscopes measure angular velocity, as a person walks, their limbs are most likely experiencing much different angular velocities as compared to accelerations. Additionally, the preliminary course subjects walked did not include many turns and hence no significant full-body angular velocities; it was primarily straight with a few ninety-degree turns. The lack of such turns made the gyroscope data not very useful.

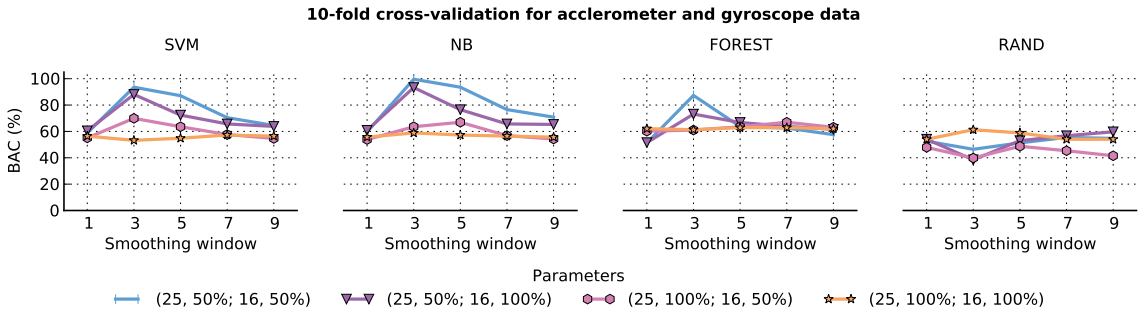

Figure 4.4.4: Ten-fold cross-validation using accelerometer and gyroscope data for different feature coherence parameters, classifiers and smoothing windows. We recommend only using the accelerometer since these results match the results of Figure 4.4.2.

4.4.5 Combined accelerometer and gyroscope validation

Finally, we examine the use of the combined accelerometer and gyroscope data. Figure 4.4.4 shows the results of a 10-fold cross-validation for the same parameters used previously. Overall, these results are similar to the results in Figure 4.4.2. This is expected since each classifier should perform at least as well as the accelerometer-only validation since they have the same information. This suggests that the overhead of a gyroscope does not improve classification rates enough to warrant inclusion.

4.5 Single-subject evaluation

The purpose of this study was to understand how our method performs over a larger subject pool and longer walking times. This study should confirm whether our method generalizes to many subjects and does not require training from any specific subject.

4.5.1 Dataset

We collected data from 7 subjects with accelerometers attached to each subject at the five highlighted locations specified in Figure 2.1.2. The accelerometer we used in

	Location						
Subject	**AL**	**AR**	**HL**	**HR**	**WR**	**Total**	**Average**
1	18:49.5	18:49.5	18:49.5	18:49.5	18:49.5	01:34:07.4	18:49.5
2	30:14.4	30:04.0	29:58.2	30:05.7	30:04.1	02:30:26.3	30:05.3
3	21:06.5	21:06.5	21:06.5	21:06.5	21:06.5	01:45:32.6	21:06.5
4	19:31.5	19:30.7	19:17.7	19:17.9	19:34.7	01:37:12.5	19:26.5
5	20:34.8	20:33.2	20:25.0	20:18.2	20:28.3	01:42:19.5	20:27.9
6	28:46.8	28:49.6	28:39.6	28:36.7	28:39.5	02:23:32.2	28:42.4
7	19:10.2	19:07.3	18:58.8	18:57.9	19:05.0	01:35:19.1	19:03.8
Average	22:36.2	22:34.4	22:27.9	22:27.5	22:32.5	01:52:38.5	22:31.7

Table 4.5.1: Statistics of our 7 subject single-subject dataset. The average subject walked for 22 min along a course of their choosing.

this experiment was the WiTilt [135], which contains the Freescale MMA7260Q [52] 3-axis accelerometer. The WiTilt allowed us to collect data wirelessly via Bluetooth to a computer that we carried with the subject as they walked. The WiTilt sampled the accelerometer at 255 Hz. Table 4.5.1 provides some summary statistics for each subject. On average, each subject walked for approximately 22 minutes. Subjects were not instructed to walk any particular course; rather, we required that they walk for roughly 20 minutes while we followed them. Our device and data-collection protocols were approved by our Institutional Review Board.

Again, we only have positive examples of SNs on the same person, so we used the same strategy described in Section 4.4 to create negative examples. However, for this dataset, this approach creates many more negative examples than positive examples since there are six other subjects to compare against for each subject. As such, this dataset is unbalanced and heavily skewed towards negative examples.

4.5.2 Parameters

We explored most of the same parameters as specified in Section 4.4. However, because the sample rate of the accelerometer used in this study differed from the previous study (255 Hz versus 250 Hz), we choose different feature window lengths that were a factor of 255 Hz. The feature window lengths we examined in this study were 17, 51 and 85. The other parameters remained the same.

4.5.3 Feasibility

The purpose of this section is to understand the feasibility of our method. That is, we first make some assumptions about the location of the SN and subject wearing the SN. Then, we progressively relax these assumptions until we make no assumptions about the location of the SN or who is wearing it.

For the studies in this section, we only present results from the best choice of parameters for each classifier, based on our preliminary experiments, which were a feature window length of 17 with a 100 % step, a feature coherence window length of 16 with a 50 % step, and a smoothing window of 3. This combination corresponds to 32 seconds of walking data.

Given known subject and location

Our first experiment aimed to determine how well our method works when we know who the subject is and where the subject was wearing the SN. That is, how well does our method perform when we train a classifier using only a specific subject's data at a specific location? This assumption might be appropriate for devices designed to be worn only at a certain location, such as a wristband. We simulate these assumptions by training a classifier for each subject-location combination. This method is feasible because we could use the bioimpedance sensor described in Section 3 to identify the

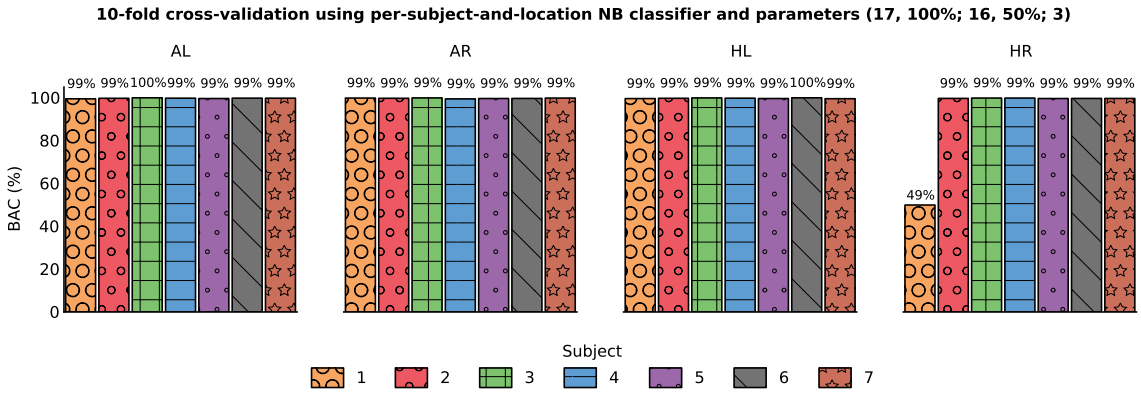

Figure 4.5.1: BAC of a 10-fold cross-validation for each location and subject using the NB classifier. It performed perfectly except for Subject 1's right hand.

subject and leverage work that describes a method to determine where on the body a sensor is located [89, 9]. However, the drawback is that this method would require subject-specific training data.

Figure 4.5.1 shows the results for this experiment for the NB classifier. As one would expect, our method performed nearly perfectly, except for Subject 1's right hand. (We are uncertain why that one case failed.) This experiment tells us that there are some unique characteristics about each location and subject that can easily be discriminated. Indeed, much of the gait recognition literature backs up this claim [77, 117].

Given known location but unknown subjects

In the next experiment we removed the assumption that we know which subject is wearing the SN. That is, we do not know who the subject is, but we do know where on the body the subject is wearing the SN. We simulate this assumption by training a classifier for each location, across all subjects' data. Figure 4.5.2 shows the results for this experiment for each type of classifier. As expected, our method performed worse than the case when we know which subject is wearing the SN. Using the NB classifier, our method could tell with 89 % BAC on average whether the SN and MN

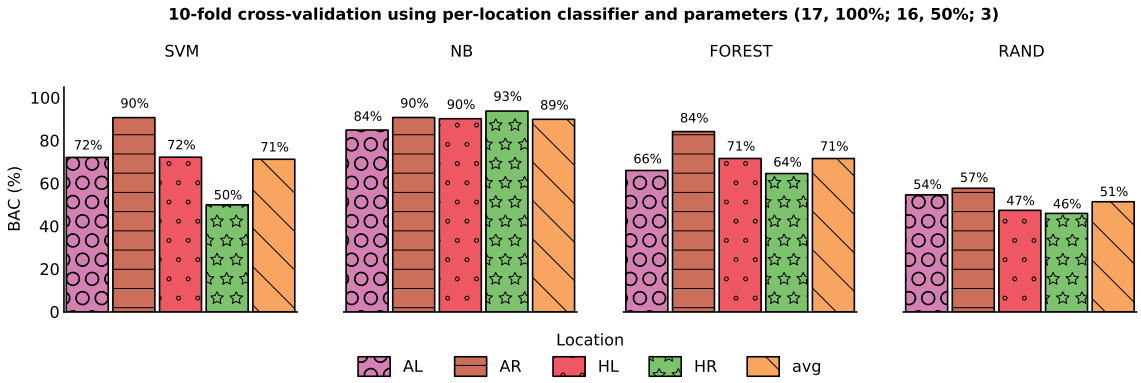

Figure 4.5.2: BAC of a 10-fold cross-validation for each location (but subject unknown) using different classifiers. The NB classifier performed the best overall.

are on the same body under the assumption that we know the location of the SN *a priori*.

Given known subject but unknown locations

In this experiment, we remove the assumption that we know the location of the SN on the subject's body. That is, we know who the subject is, but we do not know where on their body they are wearing the SN. We simulate this assumption by training a classifier for each subject, across all locations' data. Figure 4.5.3 shows the results of this study for each type of classifier. The NB classifier achieved a 89 % per-subject average BAC, which is on par with the previous experiment. As expected, our method performed worse than the case when we know the identity of the subject and the location of the SN.

Given unknown subjects and locations

Finally, in this experiment we make no assumptions about the subject or the location of the SN on their body. This is the ideal circumstance since having a global classifier, as opposed to a per-subject or per-location classifier, means fewer models to train and test. In this case, we train one classifier across data from all subjects and locations.

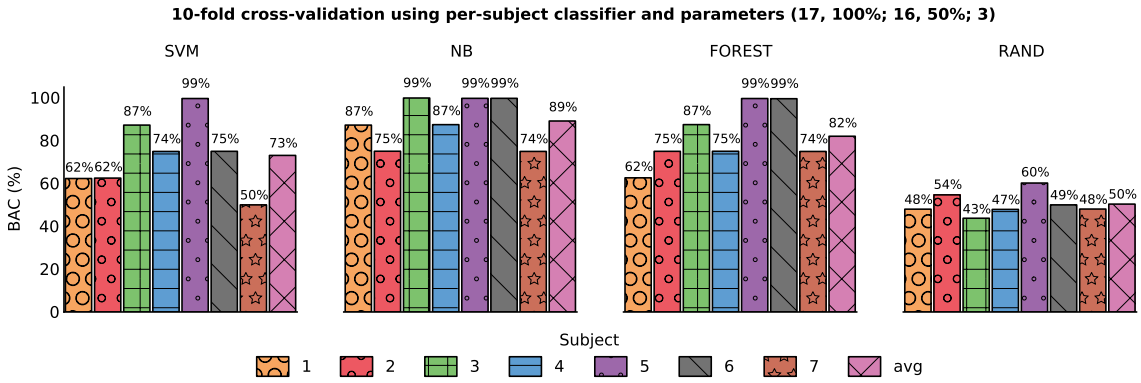

Figure 4.5.3: BAC of a 10-fold cross-validation for each subject (but location unknown) using different classifiers. Again, the NB classifier performed best overall.

Figure 4.5.4 shows the results of this experiment for each type of classifier. In addition to BAC, we also show the precision and recall of each classifier. The NB classifier performed the best over all the statistical measures. Notice that the precision is near perfect, meaning that the number of false positives was very low compared to the number of true positives. However, recall is not so good. For a security-related application this situation is optimal, since false positives are more detrimental to the system than false negatives. More false negatives just means we need more time to verify whether these sensors are on the same body. It is possible, for example, to tune a classifier to produce more false positives or more false negatives depending upon the application. For our purposes, these classifiers work fine without much tuning. Thus we conclude that using a NB classifier we can achieve 86 % BAC with few false-positives when given 32 s of acceleration data from a subject who is walking.

4.5.4 Generalizability

The purpose of these experiments was to understand the generality of our method. Ideally our method would require neither subject-specific nor location-specific training to achieve good performance. To test this goal further, we ran a leave-one-*subject*-

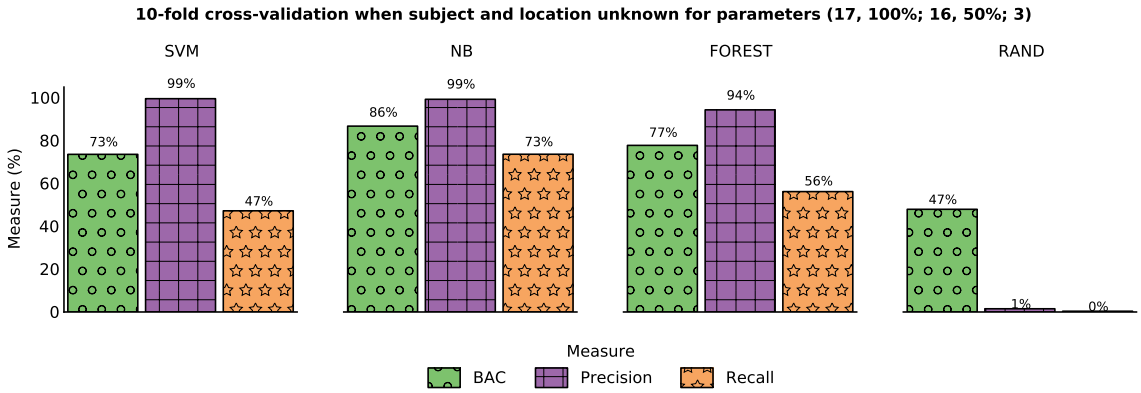

Figure 4.5.4: BAC, precision, and recall for a 10-fold cross-validation for different classifiers when both subject and location are unknown. The NB classifier performed best for this case.

out and leave-one-*location*-out cross-validation. Recall that a leave-one-subject-out cross-validation means that we take one subject's data and set it aside as the test dataset. We then use the remaining subjects' data to train a classifier. Finally, we use that left-out subject's data to test how well the model performs. Because a subject's data is left out in the training phase, then the classifier will have no knowledge about the distribution of that particular subject's data.

Location left out

Ideally, our method would not be sensitive to any particular location. Figure 4.5.5 shows that this is the case, because the average location BAC for this validation is nearly identical to the BAC in Figure 4.5.4. Notice how the ankle locations performed better than the hand locations in all classifiers. We believe this is the case because a person's hands tend to move more sporadically while walking than do a person's feet. A person's feet need to keep them moving forward while their arms do not. Additionally, the right ankle tended to perform the best because it was on the same side of the body as the MN in our experiments. Likewise, the right hand tended to perform the worst because, while it is on the same side of the body as the MN, a person's hand tends to swing opposite of their forward leg while they are walking.

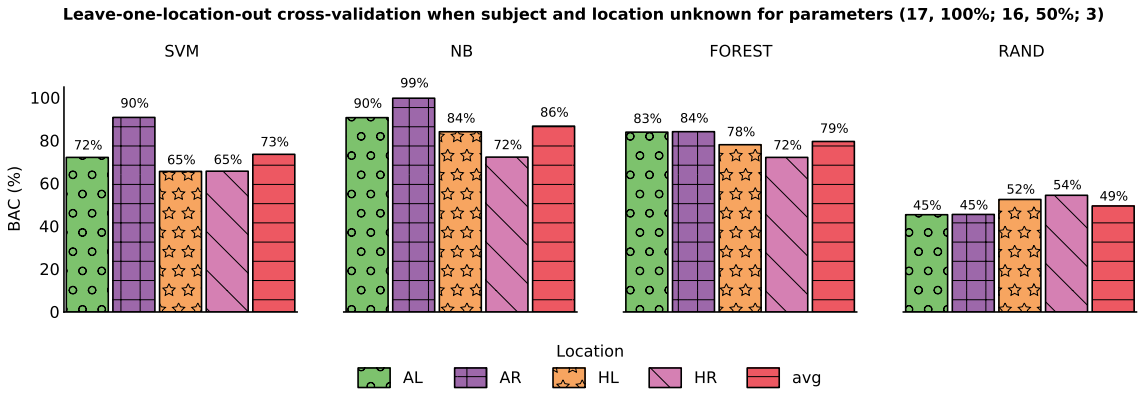

Figure 4.5.5: A leave-one-location-out cross-validation for different classifiers when both subject and location are unknown. The similarity of the average location BAC in this figure to Figure 4.5.4 means that our method is not sensitive to any particular location for these parameters.

That is, a person's right hand moves similarly to their left ankle. Because our method performed just as well on average as the feasibility experiment, we do not need to train a per-location classifier.

Subject left out

Ideally, our method would not be sensitive to any particular subject. Figure 4.5.6 shows that this is the case, because the average subject BAC for this validation is similar to the BAC in Figure 4.5.4. Notice how the BAC for each subject had more variability in the NB classifier as compared to the SVM classifier. Thus, the SVM classifier generalized better while the NB classifier was more sensitive to subject-specific data. However, the NB classifier performed better than the SVM classifier. Because the SVM and NB classifiers exhibit similar performance to our cross-validation, our method must not be sensitive to any subject.

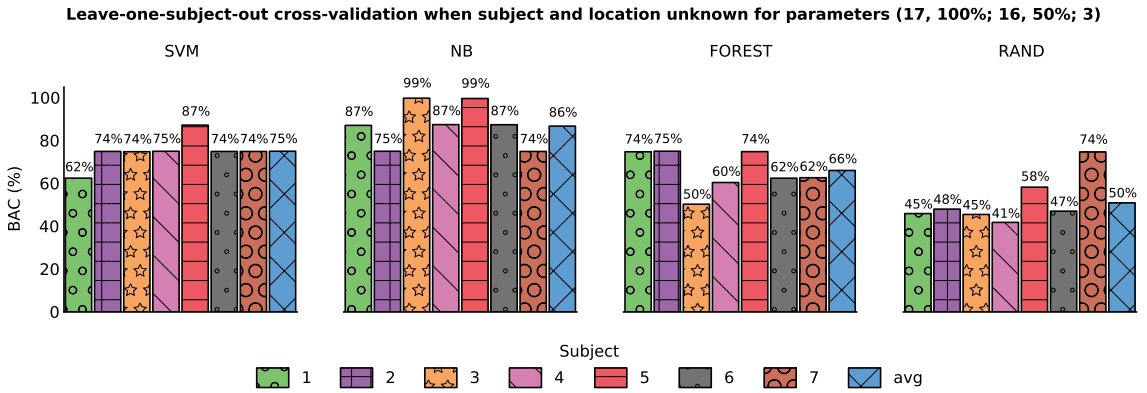

Figure 4.5.6: A leave-one-subject-out cross-validation for different classifiers when both subject and location are unknown. The similarity of the average subject BAC in this figure with Figure 4.5.4 means that our method is not sensitive to any particular subject for these parameters.

4.5.5 Feature Analysis

Although the features are not necessarily expensive to compute, minimizing the amount of data that needs to be communicated reduces energy. On typical sensor platforms, wireless communication can be an order of magnitude more expensive than computation on the same number of bits. Thus we can select a subset of features that provide the highest classification rates and ignore those features that contribute little to the classification performance. For example, we compute both the standard deviation and variance, but because standard deviation is the square root of the variance we would expect one of these computed features should be eliminated.

To accomplish this feature selection, we employed a standard correlation-based feature selection with a greedy hill-climbing algorithm to select candidate subsets. We evaluated each subset using a 10-fold cross-validation over all locations and subjects for the best selection of parameters. The feature selection algorithm indicated that standard deviation, power and energy provide the bulk of the classification performance. The other features contribute little or nothing to the classification performance. Figure 4.5.7 shows the same experiment in Figure 4.5.4 but with the

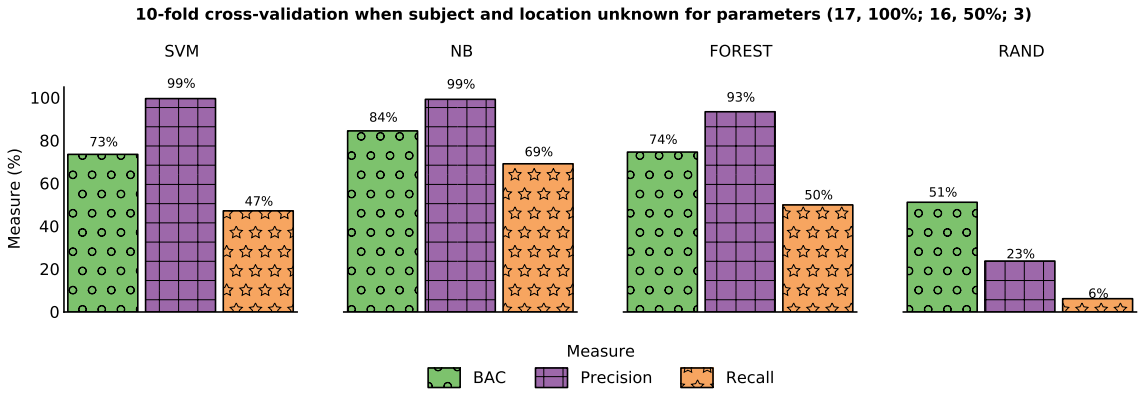

Figure 4.5.7: Verification performance for a truncated set of features. We only used the coherence of standard deviation, power and energy.

truncated set of features. The SVM and FOREST classifier performed just as well, while the recall of the NB classifiers suffers slightly.

4.6 Dual-subject evaluation

In this evaluation, we sought to understand how well our method works when two subjects wore accelerometers at the same time. That is, we gave accelerometers to two subjects and had them walk together in-step. The goal is to determine whether our method can successfully distinguish which SN was on which body, even when both SNs stayed in radio range of both bodies' MNs, and both subjects were walking in stride.

4.6.1 Dataset

To collect this data, we instrumented the Shimmer platforms to collect acceleration data at 250 Hz. To keep the individual Shimmers synchronized in time, we developed an application that uses the Flooding Time Synchronization Protocol [97] to provide μs precision time synchronization. One Shimmer was designated as the root node that was separate from the Shimmers collecting data. The acceleration data is stored

to the SD card on each Shimmer for later retrieval.

Due to our limited supply of Shimmers, each subject could only wear 4 Shimmers at a time. For the first trial, both subjects carried a Shimmer at their WR, AR, AL, and HL, and on the second trial they switched the Shimmer at HL to the HR.

4.6.2 Walking in step

For this experiment, both subjects were instructed to walk in-step around the course in Figure 4.4.1b. That is, both their legs and arms followed the same motions to the best of the subjects' abilities. The purpose of this study was to understand how well our method performs even when two people try to fool our method. Figure 4.6.1 shows the performance of our method for this dataset. The BAC of our method dropped significantly due to a decrease in recall. The precision, however, remained the same. This result tells us that for people purposely walking in-step our method is more likely to determine that an MN and SN are not on the same body when they actually are. The consequence of this is that it will take longer to verify whether two devices are on the same body. However, because the precision remains the same, it is difficult to fool our method.

4.7 Related work

There are other methods of determining whether sensors are on the same body. One could localize the wireless sensors to ensure they are within some bodily distance of each other. This trivially fails when users are close together because of the limited granularity of wireless localization [65]. Likewise, depending upon the frequency at which the wireless sensors communicate, the body may block all or some of the wireless signal necessary for the localization scheme.

A simple solution is to put labels on the wireless sensors to indicate with which

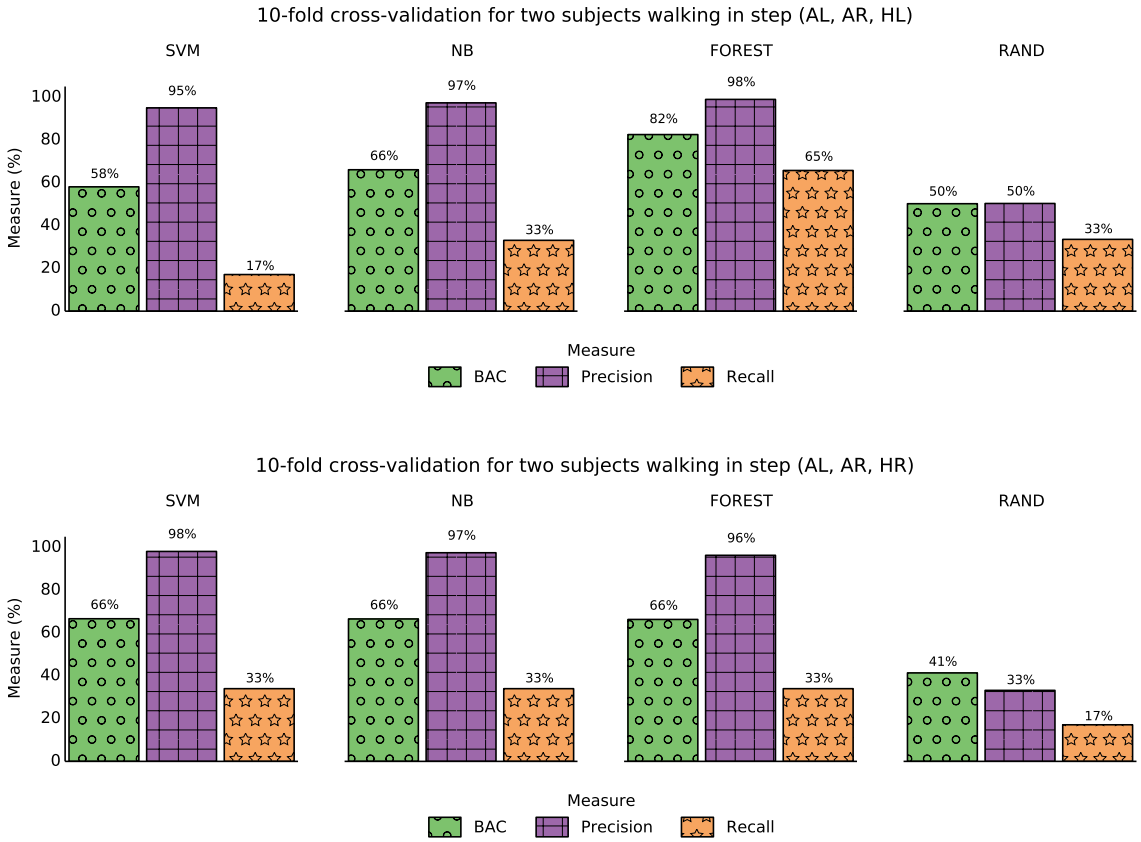

Figure 4.6.1: The performance of our same-body verification method when two people deliberately walk in step for two different trials. The top plot includes locations AL, AR, and HL while the bottom plot includes locations AL, AR, and HR. For both trials, the recall drops while precision remains the same. This result means that our method is producing more false negatives.

user they are logically paired. While this does provide some confidence to the user that they are wearing the correct wireless sensors, it does not provide a third party, like a physician, with the same confidence. It also does not take into account that some sensors might be intentionally shared. Thus, we require a kind of proof that confirms the wireless sensors were attached to the same body.

Lester et al. [91] provide a solution for the one-body verification problem, but only for sensors that are carried in the same location on each body. They also propose using accelerometers attached to each sensor and measure the *coherence* of the accelerometer data. "Coherence measures the extent to which two signals are linearly related at each frequency, with 1 indicating that two signals are highly correlated at a given frequency and 0 indicating that two signals are uncorrelated at that frequency" [91]. By looking at the coherence at the 1 Hz to 10 Hz frequencies (the frequency range of human motion), they can experimentally determine a threshold (e.g., coherence > 0.9) at which it is appropriate to deem two sensors as located on the same body.

We extend Lester et al. [91] to sensors carried at different locations on the body – wrist, ankle, and waist – by using features often used for activity recognition. We then extract the pairwise coherence of features for the sensors on the same body. Given these coherences, we can train a classifier and use it to determine whether the alleged set of sensors are on the same body. We train our classifier to be as general as possible by using data collected from several individuals; the same model can then be used by all users for all sensor devices. Lester et al. used a predefined threshold.

Prior work has explored the use of body-coupled communication [102, 109, 18, 14, 116, 162, 130, 48], which is a means of transmitting data by way of the physical human body. Although this is a promising solution, the security of these types of communications is wholly unexplored. What happens, for example, when two users touch each other? The security properties of these kinds of communication

channels are not well understood, and until then we should be wary of their use. Most compelling, many sensors will not have direct contact with the body, and in any case such a scheme would require additional hardware.

Mayrhofer et al. [106] describe a solution to exchange a cryptographic key between two devices by manually shaking the two devices together. They use the method described in Lester et al. [91] to determine whether two devices are being shaken together. But, as they notice, coherence "does not lend itself to directly creating cryptographic key material out of its results" [106]. To extract key material they extract quantized Fast Fourier Transform (FFT) coefficients from the accelerometer data to use as entropy for generating a key. One could imagine applying this technique to our acceleration data, but it is unclear whether it would still work. Our problem is made difficult by the fact that the accelerometers are not being shaken together but are attached to a body and will therefore experience less-correlated accelerations. These methods run over our datasets did not work.

Kunze et al. [89] describe a method for using accelerometers to determine where on a body a particular sensor is located. They detect when a user is walking regardless of the location of a sensor, and by training classifiers on a variety of features (RMS, frequency range power, frequency entropy, and the sum of the power of the discrete wavelet transform at different levels) on different positions on the body they can use the classifier to determine where on the body the sensor is located. We seek to provide a method that determines whether a suite of sensors is located on the same body without having to use multiple classifiers for different body locations. Although, we could use their method to boost our own results, this approach would require per-location classifiers.

Kunze et al. [88] also describe a similar method to account for sensor displacement on a particular body part. This problem is difficult primarily because "acceleration due to rotation is sensitive to sensor displacement within a single body

part" [88]. To alleviate this problem, the authors observe that "combining a gyroscope with an accelerometer and having the accelerometer ignore all signal frames dominated by rotation can remove placement sensitivity while retaining most of the relevant information" [88]. This solution could be used to extend our work to many other positions on the limbs we examined.

Finally, others have examined a similar problem, determining where people are seated a car. Wang et al. [151] describe a method for determining whether a mobile phone is on the driver side or passenger side of the vehicle. We could, in theory, use this work to extend our method to include times when two people are in a car; SNs on one side of the car will experience different centripetal accelerations than SNs on the other side. If we could factor this knowledge into our model, we may be able to use our method even though the person is not walking.

4.8 Limitations

Our method relies on the assumption that a user is capable of walking, which may not be true for some users. It remains as future work to determine whether we can extend the method for a person who is confined to a wheelchair, for example. Even for a user who is able to walk, there may be an extended period of time after binding an SN and before the user walks. The MN could alert the user that they should walk around so that verification can be performed. As future work, we may explore other acceleration and orientation-changing events; for example, to ask the user for clap their hands, shake the devices together, or perform some unique movement. Likewise, we could extend our method to other activities such as driving a car or riding a bicycle.

Ideally the algorithm should be tuned to produce more false negatives (i.e., the algorithm determined the SNs to be on different bodies when they really were on

the same body) than false positives (i.e., the algorithm determined the SNs to be on the same body when they were not) because the consequences of a false positive (recording the wrong person's data in someone's health record) are more severe than the consequences of a false negative (losing data). It is possible to 'bias' the classifier toward false negatives, if desired. For example, in an SVM classifier, one can weight the classes such that the model will output more false negatives than false positives.

Although we do not discuss encryption mechanisms, ensuring data confidentiality is paramount in any health-related scenario. If one were to optimize the verification phase by simultaneously verifying all bound SNs, it might be necessary to encrypt the acceleration data to avoid replay attacks (in which the adversary replays one SN's acceleration data in hopes that its rogue SN will be verified as being on the same body as the victim). Even if such an attack is discounted, the accelerometer data itself might be privacy sensitive because accelerometer data may be used to recognize a victim's activity. Some activities are clearly privacy sensitive, and some of those sensitive activities might be detected from accelerometer data alone.

One could also imagine using motion-capture technology to determine the acceleration an SN would be experiencing, and feed the MN this data to verify an SN that is not on a body. This technology typically requires special hardware that would make it infeasible today, but it might be feasible with new technology in the future.

In a practical system, one must consider energy and computational costs. In our model, the SN sends raw acceleration data to the MN. If this proves to be too expensive, then the SN could compute features from a window of acceleration and communicate those features instead. We leave exploring this delicate balance between extendability (allowing use of other features in the future), computability (due to limited computational capabilities on an SN), and energy requirements (with trade-offs specific to the technology in an SN) as future work. We assume the processor on the MN will be more than capable of computing correlations, but the

energy cost of these functions is unknown and requires more careful analysis. Should the computation prove to be too expensive or time consuming, then one may need to explore optimizations or approximations, or the assistance of a back-end server, with due consideration to the trade-off of computational overhead, performance, and privacy. Indeed, many of these features need to be empirically determined by a manufacturer.

In addition, a practical system must be aware of the locations a sensor might be used. Although we only explore the wrist and ankles for our experiments, we expect our method to work for other locations since these are the extremes of the body. However, this remains unverified and left for future work. Adding more locations would increase the time it takes to train the classifiers, but training can be accomplished offline. We also expect sensor manufacturers would know the general location where the sensor would be placed on the body and thus could train an SVM for that particular location.

4.9 Summary

In this chapter we described a method for determining whether two sensors are on the same body. We exploit the fact that the human body experiences similar accelerations and orientation changes as a person walks. Sensors on other people will experience much different accelerations, a fact we can exploit to determine whether two sensors are on the same body. In essence, our method generically verifies that all the SNs bound to an MN are the same body. We show that our method can achieve a balanced accuracy of 86 % with few false positives using 32 seconds of accelerometer data from different locations on the body. Our method can be generically applied regardless of the sensor type and without subject-specific or location-specific training data. In summary, we make the following contributions:

- We provide a solution to the weak version of the one-body verification problem, complementing our bioimpedance-based biometric described in Chapter 3 to address the strong version of the problem.
- We extend Lester et al. [91] to sensors carried at different locations on the body – wrist, ankle, and waist – by extracting features used for activity recognition.
- We provide empirical results to show that our method works using a dataset of seven users walking for 22 minutes.

5

Putting it all together

Now we come to the point where the two methods we previously described come together. In this chapter, we describe how our method for recognizing wearers and our method for verifying whether sensors are on the same body can be used together in a system. Although these methods can be used independently of each other, they were designed to be complementary. The power of these methods is really shown when a WBAN can use them to determine which person is wearing its MN and SNs.

We imagine a world where personal health sensors are ubiquitous and wirelessly communicate with a person's Amulet [134] or smart phone. More generally, these WBAN systems have many potential uses. Recall there are two components in

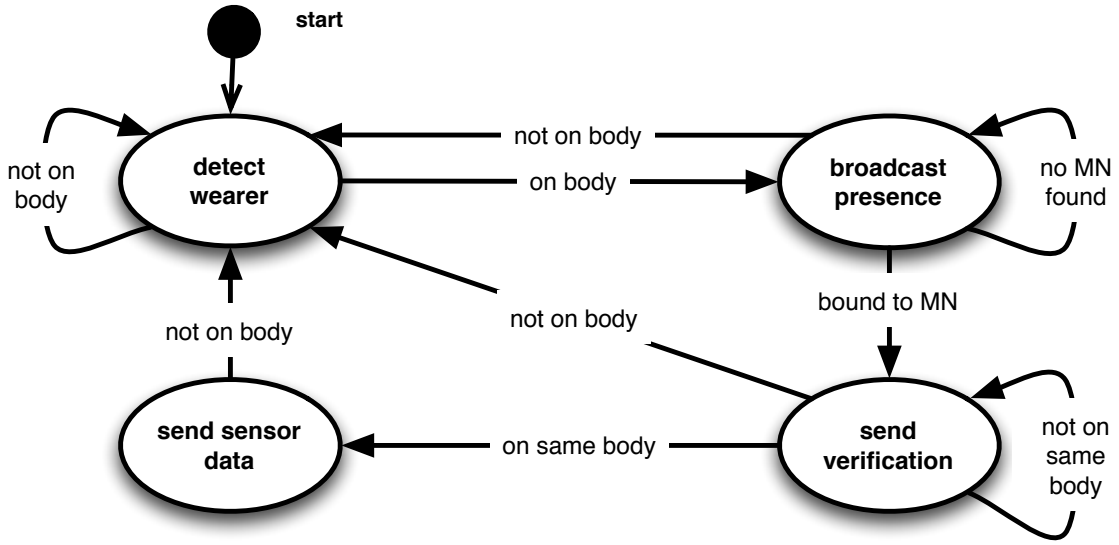

Figure 5.1.1: The protocol SNs follow. Notice that in any given state, the SN is always detecting whether it is attached to a person, because it can be removed at any time.

such a system: an MN and many SNs. Because these two components communicate wirelessly, without our methods they would not know whether they are sensing the same person, let alone which person.

5.1 Implementation

The goal of our implementation is to show how our methods could be used in a personal health sensing system. To show how this is possible, we outline the protocols the MN and SNs follow, and the methods that are necessary to realize this protocol.

5.1.1 Protocol

The purpose of each SN is to communicate its sensed data to the MN. However, in addition to the data it is sensing, it also needs to communicate data about the data (i.e., meta-data) it is sensing. This includes things like the time of the sensed data as well as who it is sensing. To accomplish this, each SN follows the protocol outlined in Figure 5.1.1.

When a user wishes to use an SN the following happens:

1. The user attaches an SN to their body.
2. The SN detects that someone is wearing it, turns on and begins broadcasting its presence.
3. An MN hears this broadcast and initiates same-body verification with the SN.
4. Once verified, the MN can request bioimpedance samples (if the SN supports bioimpedance recognition) or other sensor data.

At any step in the process, the SN might detect it is no longer being worn by a person and therefore return to the state of detecting whether somebody is wearing it. In the following sections, we describe, in detail, the purpose of each state in Figure 5.1.1.

Device authentication

To establish an authentic channel, the MN and SN need to have some existing information about each other. For example, each SN could be verified by some entity (e.g., the Food and Drug Administration (FDA)) and given a certificate along with a public and private key. An MN could store the SN's public key and use it to authenticate any communication from an SN using a digital signature scheme [125, 79, 46]. That is, each message sent over the channel would have an accompanying signature according to the chosen digital signature scheme. It is feasible for SN-class devices to conduct public-key cryptography [20]; it would only be needed once during a pairing operation.

Detecting wearers

Although each an SN could have a switch that would turn on and off power to the device, it is useful to know when the SN is attached to a person. Since SNs have

limited energy reserves, we can save energy by putting the an SN into a low-power state until it is worn by a person.

As seen in Figure 2.1.2, there are many locations on a person's body where they might wear an SN. More so, different locations might afford different types of ways an SN can be attached to a person's body. For example, a waist-worn SN might only slip into person's pocket, while a wrist-worn SN might contain a clasping mechanism. There are at least five such ways an SN could be worn: clipped on, strapped on, stuck on, or slipped into a pocket, or even implanted or ingested.

For devices that require skin contact, an SN could use bioimpedance since skin contact is easy to detect given a bioimpedance measurement. As shown in Figure 3.4.2, it is easy to detect when there is no contact with the wrist (or more generally, the skin) by measuring the impedance. When there is no contact, the magnitude of the impedance for such a measurement is effectively infinite across all frequencies. Thus, an SN could use a simple threshold to detect presence. This presence-detection scheme, however, requires a non-trivial amount of energy to execute. In the worst case, an application might require near-instant detection of wristband removal. Given that an impedance measurement takes about 98 mW of power to measure, this presence-detection method is not realistically feasible for continuous use. It might be feasible to look at an individual electrode configuration and single frequency, but we leave such methods for future work.

For devices that are clipped or strapped on to a person's body without contacting the person's skin, the easiest detection mechanism is a simple circuit. That is, when the device is strapped on to a person, the two ends could complete a circuit that would wake up the device. For a clip-on device, a magnet and Hall-effect sensor could be used to detect when the clipping mechanism is clipped. Note that neither of these mechanisms determine whether a person is wearing the SN. Rather, they are a proxy for whether the SN is being worn. A person could, for example, strap the

device to itself and set it on the bedside table.

Broadcasting presence

Given a method for detecting when it is worn, an SN would begin broadcasting its presence once it detects someone is wearing it. An MN would overhear this broadcast and bind with the SN. Binding establishes a communication channel that is confidential. This is accomplished by using a key exchange protocol. For example, an MN and SN would participate in a Diffie-Hellman key exchange [41] over the authenticated channel. Such a scheme is called Ephemeral Diffie-Hellman and gives us the *forward secrecy* property. Such an exchange allows the MN and SN to establish a key that can then be used to encrypt any further communications, ensuring the confidentiality and integrity of the channel. For better energy efficiency, we could employ Hide-n-Sense which is a "mHealth sensing protocol that provides strong security and privacy properties at the link layer, with low energy overhead, suitable for low-power sensors." [96].

Same-body verification

Once the MN and SN establish a confidential channel, the MN can then initiate verification. Verification is the process of determining whether the MN and SN are on the same body. Until an SN is verified, any sensor data from the SN is ignored by the MN. (As it may take some time for verification to succeed, in some implementations the MN may buffer the incoming data received between the moment of binding and the moment of verification, recording the data only once verification is assured. This "retroactive verification" of the early data is feasible because of our assumption that an SN can detect its own attachment and removal; if an SN is moved from one body to another before it was verified on the first body, the unbinding and rebinding events will clear the buffer on the first body's MN.) To achieve verification, our

protocol requires an algorithm that is able to decide whether two SNs are on the same body. We use the method described in Chapter 4.

Procedure 1 provides an overview of the process the MN uses to verify an SN. Because our method depends on recognizable acceleration events, our algorithm performs verification only when the user is walking. Thus, the MN must implement a recognition algorithm `UserIsWalking()` capable of recognizing walking using the accelerometer data.

To begin verification, the MN records acceleration data using its internal accelerometer for t seconds. Simultaneously, it asks the other SN to send it acceleration data for the same duration. The duration required depends on the level of confidence desired; a shorter duration may lead to more incorrect results (false positives and false negatives), but a longer duration makes the approach less responsive after the person first puts on the sensor. It then runs our algorithm from Chapter 4, `AreCorrelated` in Procedure 1, to determine whether its internal acceleration data correlates with the SN's acceleration data.

Sending data

Only when the accelerometer data correlates well does the MN begin to record that SN's other sensor data (e.g., electrocardiography data). Thus, the MN signals to the SN that it should start sending its sensor data according to preference of the application. For example, the MN could direct an SN that has a bioimpedance sensor to send bioimpedance samples. Using the method described in Chapter 3, the MN could use this data to establish the identity of the person wearing that SN. Additionally, because the MN has verified that it is on the same body as the SN, then it is must also be worn by that person wearing the SN. Thus, the MN can exploit this transitive relationship to automatically label any sensor data from an SN on the same body as it. It is not necessary for the MN nor all of the SNs to be able to identify

Procedure 1 Binding with and verifying SNs, from the MN's perspective

Notation:

B: set of bound SNs, initially empty

A_i: acceleration data from SN i, where $i = 0$ is the MN's acceleration data, and $i > 0$ re SNs.

`Record`(t): read MN's accelerometer for t seconds

`Recv`(b, t): read SN b's accelerometer for t seconds

`AreCorrelated`(x, y): determine whether acceleration data x and y

```
while { true } do
    if b := NewSensorNodeDetected() then
        B := B ∪ b
        { Mark SN b as unverified }
    end if
    for b ∈ B do
        if Disconnected(b) or Timeout(b) then
            B := B \ b
        else if d := RecvData(b) and IsVerified(b) then
            RecordData(b, d) { Save b's data d in our health record }
        end if
    end for
    if UserIsWalking() then
        for b | b ∈ B and not IsVerified(b) do
            { The next two lines are accomplished in parallel }
            A_0 := Record(t)
            A_b := Recv(b, t)
            if AreCorrelated(A_0, A_b) = true then
                { Mark SN b as verified }
                { Tell SN b to send sensor data }
            end if
        end for
    end if
end while
```

their wearer, rather, it is sufficient that they implement our verification method. It is the combinations of the methods described in Chapters 3 and 4 that gives our system the power to automatically label sensor data with the identity of the person that was wearing the SNs at the current time.

Unbinding

Unbinding occurs when a user removes an SN. In the ideal case, the following happens:

1. The user unstraps the SN from their body.
2. The SN detects that it was removed and notifies the bound MN of this fact.
3. The MN acknowledges this notification, thereby unbinding it with the SN.
4. Upon receipt of this acknowledgment (or upon timeout), the SN turns off.

An SN may lose power or go out of range of the MN, during or prior to the user unstrapping the SN. Thus, the MN periodically pings each SN (not shown in Procedure 1); if the SN does not reply (after some timeout period), the SN is likely not on the same body or not powered, and the MN treats it as unverified and unbound.

5.1.2 Hardware and software

We implemented these protocols on a Shimmer [132] and Google/Samsung Nexus S smart phone [153]. The Nexus S and Shimmer communicated wirelessly via Bluetooth 2.1 [23] using the Serial Port Profile. We relied on the underlying authentication and confidentiality mechanisms present in Bluetooth 2.1.

The Shimmer application is a 1570 line TinyC program capable of sensing the on-board accelerometer at 250 Hz and sending it via Bluetooth. It optionally includes

the capability to control an attached bioimpedance sensor board according to the method described in Chapter 3.

The Nexus S application is a 932 line Java Android 4.1 application. We used the OpenUAT toolkit [105] for some of the signal processing and the Weka toolkit [62] for the machine-learning aspects of the application. We implemented the algorithms described in Chapters 3 and 4. The models installed on the Nexus S were trained offline using Weka.

For our experiments, we instrumented the Android application to output time stamps related to when events were happening on the smart phone. This allowed us to correlate known events with our energy measurements. We could not, however, instrument the Shimmer because it require us either sending extra data via the radio or writing logging data to the SD card. Such extra instrumentation would adversely affect both the timing and energy profile of the application. To compensate, we coded the Shimmer application in such a way that would allow us to know where specific events took place based on known timing characteristics starting from the booting of the Shimmer.

5.2 Evaluation

Given the implementation of these protocols and underlying methods, we sought to evaluate system performance according to several factors. First, we evaluated whether such a system is feasible given that we require subjects to be walking for a period of time. Next, we study the energy requirements for such a system from the perspective of both the MN and SN. Finally, we end with a security analysis of our protocols under the security model described in Section 2.3.

Subject	Total (h)	Walking (min)	Count	Min (s)	Mean (s)	Median (s)	Max (s)	Std (s)
1	49.8	62.0	37	10.0	100.5	47.5	475.0	106.1
2	69.9	61.0	46	15.0	79.6	58.8	360.0	69.9
3	24.6	63.2	40	10.0	94.8	46.2	297.5	91.1
4	49.8	133.0	93	7.5	85.8	47.5	477.5	93.2
5	67.8	77.5	115	7.5	40.4	27.5	292.5	40.9
6	65.1	31.5	29	7.5	65.1	55.0	227.5	52.3
7	84.4	48.2	53	17.5	54.6	30.0	172.5	44.5
Avg	58.8	68.1	59	10.7	74.4	44.6	328.9	71.1
All	411.5	476.4	413	7.5	69.2	40.0	477.5	75.0

Table 5.2.1: Statistics for each subject about their particular walking habits. The first column, **Total**, is the total time in hours the subject carried the device. The next column, **Walking**, is the total time in minutes the subject spent walking according to our classifier. **Count** represents the number of the times the classifier detected walking, while **Min**, **Mean**, **Median**, and **Max** represent the length, in seconds, of the smallest, average, middle, and largest walking intervals. **Std** is the standard deviation of all walking intervals in seconds. The **Avg** row represents the average over the columns above, while the **All** row represents the concatenation (i.e., aggregation) of all subjects together as one.

5.2.1 Same-body verification intervals

Since we require about 40 seconds of acceleration data when a user is walking to verify whether sensors are on the same body according to Chapter 4, we would like to know how many times a day our system could perform same-body verification, on average. To find out, we continuously collected acceleration data from several subjects using an accelerometer placed in their pocket. Each subject was given a Shimmer [132], which continuously collected acceleration data, and was asked to carry the device at all times except when they were sleeping or performing an activity that could harm the device (e.g., showering, swimming, or working out).

We collected a total of 411 hours of acceleration data from seven students who each carried a single accelerometer that sampled at 100Hz. On average, a subject

carried the accelerometer for 58 hours. Table 5.2.1 shows statistics for each subject. In addition to this dataset, we also collected a one-hour training dataset consisting of mixed activities.

To classify walking intervals, we use the features described by He et al. [68]. They observe that there is a period of weightlessness a person experiences while performing activities like running or walking. They show how 6 features (mean peak height, mean weightless length, mean peak interval, mean weightless interval, ratio of peak number to weightless number, and ratio of weightless length to window length) can be used to achieve 98.54 % walking classification accuracy for an accelerometer placed in a subject's trouser pocket. They suggest extracting these features for every 5.12 seconds of data (with 50 % overlap between windows), so we extracted the described features for each 5 second window of acceleration data with 50 % overlap. Using the training dataset, we computed the features of that dataset, and trained a NB classifier. We then used the classifier to classify a window of features for each subject, drawing test cases from the test dataset. To improve results, we smoothed over classification results using a simple majority voting rule for every 3 classifications.

Figure 5.2.1 shows the acceleration intervals that were classified as walking for each subject. On average, a subject walked for 74 seconds per interval however the median walking interval length was 44 seconds. This means more than half of all classified walking intervals were at least 40 seconds long. It seems to me that what matters is the "time between walking intervals of at least 40 seconds." What we really want is mean-time-to-verification.

Figure 5.2.2 shows the distribution of walking interval lengths for all subjects. The distribution of walking interval lengths tends to cluster towards shorter lengths of 60 seconds or less. However, there do exist enough walking intervals greater than 40 seconds for each subject. Thus we would be able to perform same-body

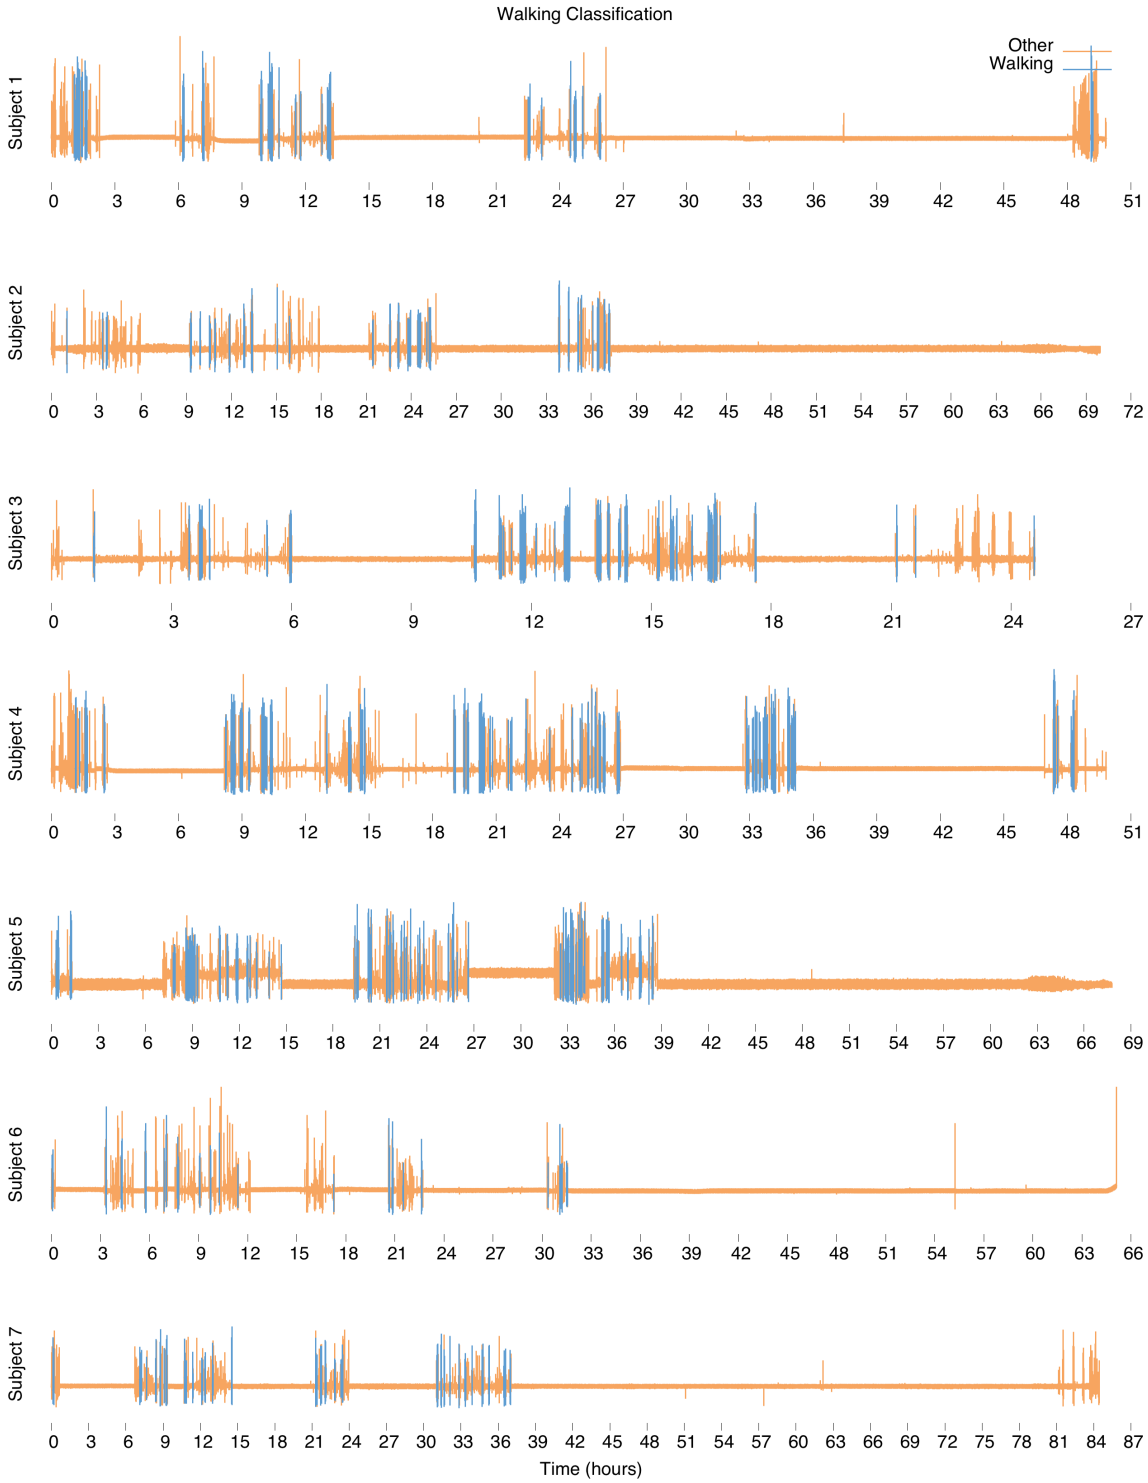

Figure 5.2.1: Walking classification for each subject for the whole time that they carried the sensor. According to the classification method, subjects walked for 74 seconds on average and in total walked an average of 2.3 % (1.9 % aggregated) of the collection period.

verification with smoothing for the given population.

5.2.2 Energy measurements

To capture energy measurements, we used the Monsoon Power Monitor [112] connected to a Windows laptop. The Power Monitor acts like a battery and samples the current drawn every 200 μs. In the graphs shown, we down-sampled the current drawn to 100 ms via averaging. We show energy measurements for the MN and SN for both same-body verification and bioimpedance recognition.

Same-body verification

Figure 5.2.3 shows how much current an SN draws when engaging in same-body verification with an MN. There are five distinct phases in the energy measurement. First is the idle sequence. The SN consumed 6.45 mA on average when idle. In the next phase, the SN turned on its Bluetooth radio at the 3.0 s dashed line. This radio consumed 10.0 mA while idle: the spikes present in this phase of the energy measurement correspond to times when the Bluetooth radio was attempting to discover the MN. Next, the SN and MN established a connection at the 11.5 s dashed line. On average, this phase consumed 27.0 mA. The MN instructed the SN to begin sampling the accelerometer and sending the data via Bluetooth at the 11.5 s dashed line. This phase consumed 25.8 mA on average. Next, the sampling and sending phase ended at the 53.2 s dashed line. Sampling and sending consumed 37.9 mA on average. Sampling the accelerometer accounted for less than 3 % of the current consumed. Finally, the Bluetooth radio was turned off at the 55.3 s dashed line. On average, this phase consumed 26.2 mA. The dominating factor in this energy measurement was the Bluetooth radio. Thus, it would make the most sense to limit the use of the radio or find lower-power ways of transmitting data. An SN

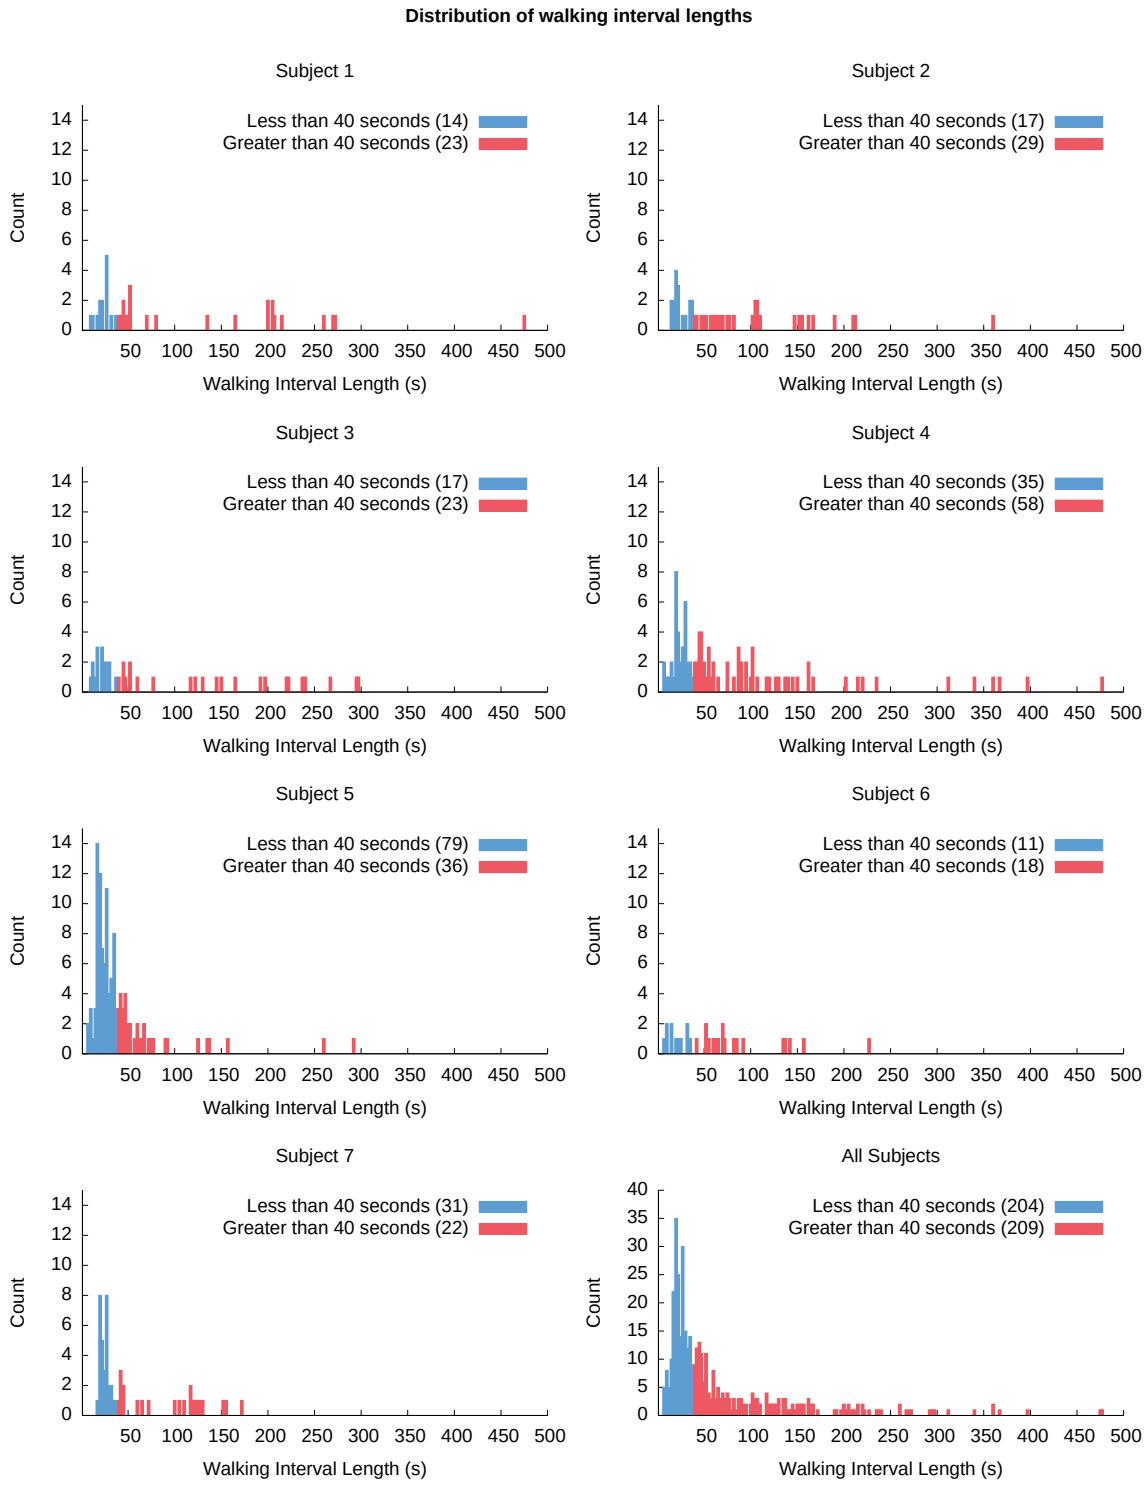

Figure 5.2.2: The distribution of walking interval lengths for all subjects and the aggregation of all subjects. Features were extracted every 5 seconds with 50 % overlap. The aggregated distribution tends to skew heavily towards shorter walking intervals for this particular set of subjects.

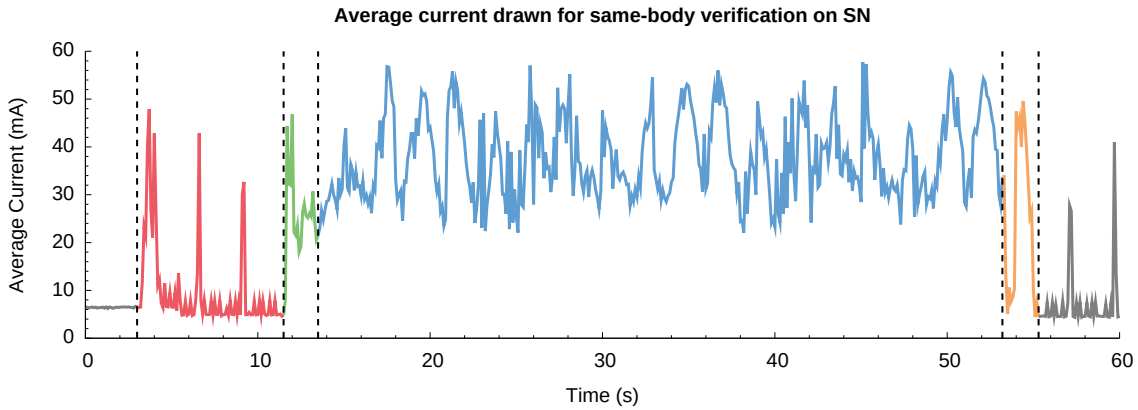

Figure 5.2.3: Energy measurement of an SN when collecting accelerometer data for verification. The dashed line at 3.0 s is when the SN turned on its Bluetooth radio. The dashed line at 11.5 s is when the MN connected to the SN. The dashed line at 13.5 s is when the MN told the SN to start sensing acceleration. The dashed line at 53.2 s is when the MN told the SN to stop sensing acceleration. The dashed line at 55.3 s is when the MN disconnected from the SN.

with a 450 mA h battery could last more than one day with the Bluetooth radio on continuously and with an hourly same-body verification.

Figure 5.2.4 shows how much current the MN consumed when engaging in the same-body verification protocol with an SN. This figure is not aligned in time with Figure 5.2.3 because we could only measure energy from one device at a time. There are four distinct phases in the energy measurement. The first phase shows the energy the MN consumed when idly waiting. On average, the MN consumed 57.1 mA. In the next phase, the MN turned on its Bluetooth radio at the 5.0 s dashed line and began connecting to the SN. This phase consumed 94.8 mA on average. The MN connected to the SN at the 7.3 s dashed line and immediately asked the SN to start sending data. In this phase, the application was both receiving acceleration data from the SN over Bluetooth and sensing the MN's internal accelerometer. Furthermore, it was also processing this data in real time according to the method described in Chapter 4. The dashed lines at 20.1 s, 26.4 s and 32.8 s in this phase of the energy measurement correspond to computation of feature coherences and

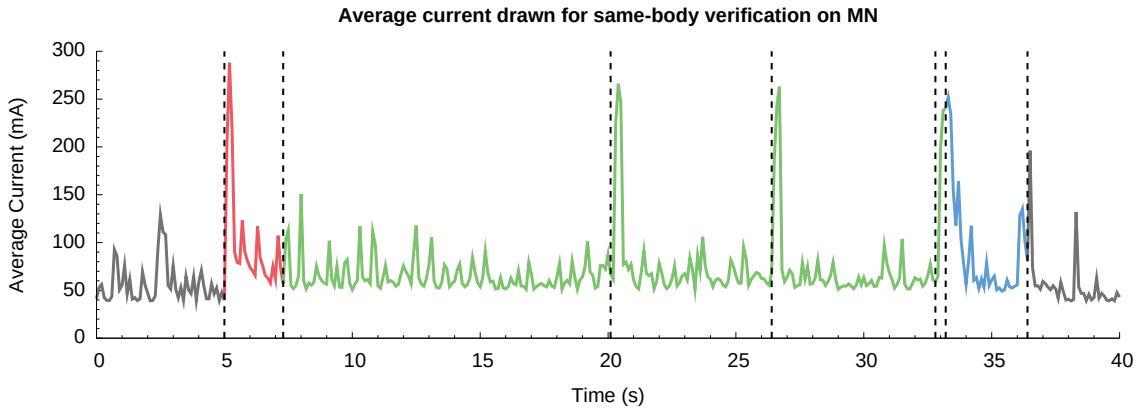

Figure 5.2.4: An energy measurement of an MN when running same-body verification. The dashed line at 5.0 s is when the MN turned on its Bluetooth radio. The dashed line at 7.3 s is when the MN connected to the SN. The dashed lines at 20.1 s, 26.4 s and 32.8 s are when the MN computed a feature coherence and classified them resulting in an current spike. The dashed line at 33.2 s is when the MN disconnected. The dashed line at 36.4 s is when the MN turned off the Bluetooth radio.

the classification of them. On average, this phase consumed 71.2 mA. At the 33.2 s dashed line, the MN disconnected from the SN and begun turning off its Bluetooth radio. This phase consumed 86.4 mA on average. Running our protocol does not incur significant overhead compared to how much overhead was required to run the Android operating system.

Bioimpedance recognition

Figure 5.2.5 shows the energy measurement of an SN sampling bioimpedance and sending these values to an MN. Here, we assume the MN has already verified that the SN is on the same body. Thus, this figure shows how much energy an SN consumed for just bioimpedance recognition. There are five distinct phases in the energy measurement. In the first phase, the SN was idle. This consumed 6.42 mA on average. Next, the SN turned on its Bluetooth radio at the 3.0 s dashed line and attempted to pair with the MN. On average, this phase consumed 9.33 mA. Again, the spikes in this phase correspond to the times when the Bluetooth radio

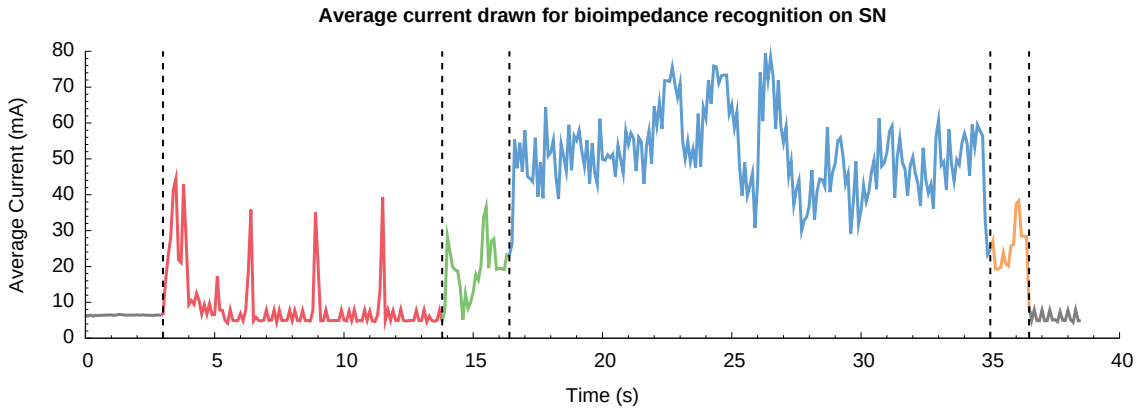

Figure 5.2.5: An energy measurement of an SN when collecting bioimpedance data for recognition. The dashed line at 3.0 s is when the SN turned on its Bluetooth radio. The dashed line at 13.8 s is when the MN connected to the SN. The dashed line at 16.3 s is when the MN told the SN to start sensing bioimpedance. The dashed line at 35.0 s is when the MN told the SN to stop sensing bioimpedance. The dashed line at 36.5 s is when the MN disconnected from the SN.

was searching for the MN. At the 13.8 s dashed line, the SN and MN established a connection. This phase consumed 27.0 mA on average. Next, the MN then instructed the SN 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 MN told the SN 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. An SN with a 450 mA h battery could last more than a 1 day with the Bluetooth radio on continuously and with an hourly bioimpedance recognition.

Figure 5.2.6 shows the energy measurement of an MN engaged in bioimpedance recognition. This figure is not aligned in time with Figure 5.2.5 because we could only

measure energy from one device at a time. Like the measurement in Figure 5.2.5, we assume the MN has already verified that the SN is on the same body. There are four phases in the energy measurement. This phase was a steady state energy measurement of the MN. This phase consumed 54.5 mA on average. In the next phase, the MN turned on its Bluetooth radio at the 5.0 s dashed line and began connecting to the SN. On average, this phase consumed 87.8 mA. At the 13.8 s dashed line, the MN was connected to the SN. In this phase, the MN collected bioimpedance samples from the SN. 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 SN started sensing a new electrode configuration. This phase consumed 67.9 mA on average. At the 29.6 s dashed line, the MN ran the recognition algorithm, told the SN to stop sensing, and began disconnecting from the SN. 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. In fact, bioimpedance recognition uses less energy than the same-body verification method because less time was spent receiving samples and the feature computations were simpler.

5.2.3 Security analysis

Recall that our security model outlined in Section 2.3 had four goals. In this section, we argue that our system meets each of these goals.

G1: Confidentiality

Our system should preserve the confidentiality of the sensed data and meta-data. That is, the sensed data and its meta-data should not be revealed to anyone but the user and service provider. This goal is the met for the following reasons. First,

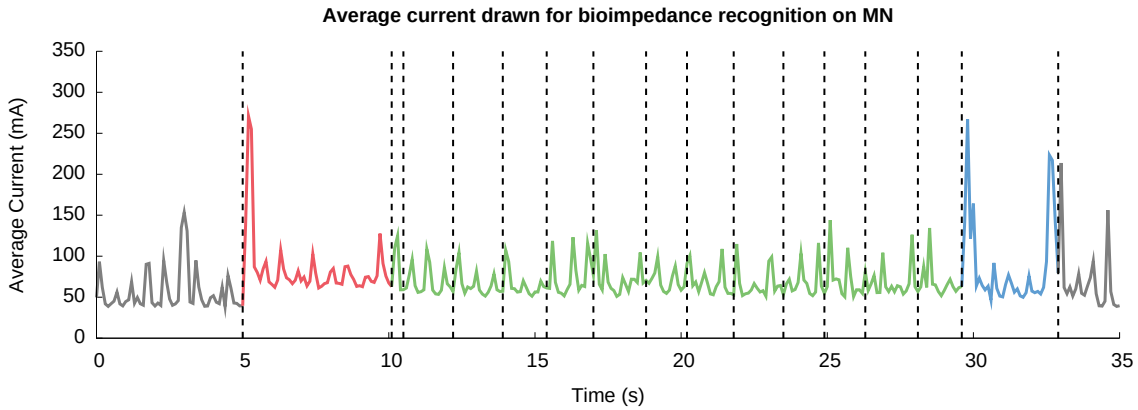

Figure 5.2.6: Energy measurement of an MN when running bioimpedance recognition. The dashed line at 5 s is when the MN turned on its Bluetooth radio. The dashed line at 10.1 s is when the MN and SN established a connection. The dashed lines at 10.5 s, 12.2 s, 13.9 s, 15.4 s, 17.0 s, 18.8 s, 20.2 s, 21.8 s, 23.5 s, 24.9 s, 26.3 s and 28.1 s are when the SN started sensing a new electrode configuration. The dashed line at 29.6 s is when the MN classified the bioimpedance sampled and disconnected. The dashed line at 32.9 s is when the MN turned off its Bluetooth radio.

the user trusts the SN to not reveal any of this sensed data or created meta-data to anyone but the MN. This is met under the assumption that an attacker cannot tamper with the hardware and software of the SN or MN. The SN confidentially reveals the sensed data and meta-data to the MN by establishing an encrypted channel between them to communicate any meta-data or sensed data. Before the establishment of this confidential channel, the only information communicated between the MN and SN is the presence of each device and our protocol does not rely on any identifying information being present in this broadcast. Thus, an attacker can only learn that there is an MN and SN present. While the MN and SN establish a confidential channel, the attacker cannot learn anything about sensed data or meta-data because neither device communicates any such data during this phase. Furthermore, the attacker cannot break the key exchange protocol used to establish this confidential channel because we assume the attacker is computationally bound. Once the encrypted channel is established, the attacker is still unable to learn the contents of that communication under the assumption that they are computationally bounded since

we rely on the computational hardness of the underlying cryptographic primitives. Once the data and meta-data has arrived at the MN from the SN, we rely upon our trust assumption about the MN to not reveal the data and meta-data to anyone else but the service provider.

G2: Integrity

Our system should preserve the integrity of the sensed data and meta-data, meaning that the service provider and user should be able to trust that the sensed data or meta-data has not been tampered with. This goal is the met for the following reasons. First, an SN will not tamper with sensed data or fabricate meta-data because we assume the integrity of the hardware and software on the SN. Because sensed data and meta-data is always sent over an authenticated and confidential channel, an attacker cannot tamper with either sensed data nor meta-data while it is in transit between the SN and MN. The attacker cannot tamper with the data because we assume it is computationally bounded and the underlying cryptographic primitives are computationally hard. An attacker cannot disrupt the flow of sensed data from the SN to the MN, because we assume an attacker will not engage in such denial-of-service type attacks. Because our trust assumptions state that an MN will not be tampered with, an MN will not modify or fabricate any sensed data or meta-data it received from an SN and thus it will maintain the integrity of the data it provides to the service provider.

G3: Authenticity

Our system should preserve the authenticity of the sensed data and meta-data. That is, the service provider should be able to trust that the sensed data and meta-data was acquired from the specified user and the specified SN. This goal is the central goal of this thesis and is the met for the following reasons. First, any means of

communicating the identity of the specified user to the service provider is met by goals **G1** and **G2** above. Our system establishes this identity through the methods described in Chapters 3 and 4. Because the MN verifies that it collects only from SNs on the same body, using the method described in Chapter 4, a service provider can be assured that all the sensed data and meta-data it receives from an MN has been collected from a single user. Under the assumption that at least one SN in our system implements bioimpedance recognition, as described in Chapter 3, the service provider can be assured of the identity of the user across all SNs. This must be the case due to the following transitive relationship: if all the SNs are on the same body as the MN and at least one SN has identified the current user using bioimpedance recognition and communicated this fact to the MN, then all SN must also be sensing that same person. It must be noted, however, that because the methods described in Chapters 3 and 4 are probabilistic, the MN is making a probabilistic statement about the authenticity of the sensed data and meta-data. In fact, our results only hold for small cohorts. We also do not take into account attacks against the subject (e.g., an attacker attempts to acquire a subject's bioimpedance). Such an attack would be more difficult to execute than with other biometrics like a fingerprint where a fingerprint can be lifted from nearly any object. A bioimpedance attacker would have to put electrodes on a person without the person knowing. Whether our system meets some guarantee of authenticity is application specific. Indeed, our methods can be tuned to provide different levels of confidence regarding of authenticity according to the desires of the application.

G4: Obscurity

Our system should preserve the obscurity of the sensors, meaning that the type of SN and MN the user is wearing should not be revealed to anyone but the service provider. This goal is the met for the following reasons. First, according to goal **G2**

our system preserves the confidentiality of any sensed data or meta-data between the user and service provider. Thus, an attacker cannot learn the type of SN and MN by inspecting the sensed data or meta-data. There is no other data communicated between the MN and SN (assuming a protocol like Hide-n-Sense [96]), thus the only avenue of attack is to do some traffic analysis on the data. For example, an attacker might be able to know that an MN and SN are engaging in same-body verification by inspecting the frequency of communication. However, an attacker already knows each MN and SN have an accelerometer since we assume each MN and SN must have an accelerometer. The same argument holds for bioimpedance recognition. For other types of sensors (e.g., heart rate sensor), however, the same argument does not hold since there is no *a priori* assumption about the presence of such a sensor. Such sensors are application dependent, so we leave this problem for the application to address.

5.3 Related work

The idea that WBANs ought to be simple to use is inspired from many other works [15, 47]. Venkatasubramanian et al. best define the usability of WBANs as needing to "activate on deployment, in a plug-n-play manner, with minimal (ideally none) initialization procedures" [147]. We have adopted this same definition of "usable" and have used it to guide our own usable security mechanisms in this prototype.

We described some simple mechanisms for detecting the presence of a human wearer. There are more complex technologies we could use. For example, we could integrate capacitive sensing technologies into our device to detect presence. In particular, the SemTech SX9300 [131] is an ultra low power, Specific Absorption Rate (SAR) controller that can discriminate between the human body and inanimate

objects. Intended for use in mobile phones to reduce transmitted power output when it is next to a human body (to comply with federal SAR standards), the chip can also be used as a general-purpose human-proximity detector. 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 areas on a printed circuit board). Open air and inanimate objects have a low permittivity (e.g., the permittivity of air is 1.0, of paper is 3.85, of concrete is 4.5, and of water is 80.1 [154]). Human tissues have high permittivity (e.g., generally above 1000 F/m at frequencies from 10 Hz to 100 kHz [55]). The integrated circuit comes in a 3 mm $\times$ 3 mm $\times$ 0.6 mm Quad-flat no-leads (QFN)-20 package and consumes only 459 μW in active mode. It can generate an interrupt to awaken a host micro-controller unit upon a "body close" or a "body far away" event, allowing the micro-controller unit to sleep until a human presence is detected or to react to wristband removal at a low power cost. 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. One could also integrate electric-field sensing technologies into an SN. Cohn et al. describe such a low-power wake-up method using electric-field sensing technology [34]. Their sensor requires contact with the skin, and using their low-power wake-up method only requires 9.3 μW total, a 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.

5.4 Limitations

The central limitation of our energy measurements is that we did not implement the presence detection method nor the walking recognition method on the SN. The

implementation of the presence detection method would require new hardware to integrate well with our prototype. Ideally such a method would boot the Shimmer when it detects a wearer. As for the walking classifier, there is a large body of literature on the efficient implementation of such activity recognition methods [68, 123, 103, 28, 16, 86, 93, 150, 94, 90, 31, 133]. In fact, the Android operating system now implements such activity recognition primitives [58]. Like the presence detection method, ideally such an activity recognition method would be intimately integrated with the hardware such that it wakes up the Shimmer when it detects the subject is walking.

Our energy measurements were also limited by the radios present in the MN and SN prototypes. Although the Shimmer platform has a 802.15.4 radio, the Samsung Nexus S does not, so we were forced to use the Bluetooth radio. Although Bluetooth and 802.15.4 operate at the same frequencies (2.5 GHz), implementations of 802.15.4 require less energy than Bluetooth. We could, however, use the new Bluetooth Low Energy standard to further decrease energy. Indeed, changing the wireless radio would be the most significant energy optimization step because the current radios consume so much energy.

The next best optimization would be to use better hardware components for, in particular, the bioimpedance sensor board. We did not design this sensor board with any sort of hard energy constraints. It might be possible, for example, to optimize the circuitry of the board for specific energy requirements. For example, the multiplexers could be replaced or removed if only a few electrode configurations are necessary to achieve the desired recognition rates. Additionally, the number of frequencies required for bioimpedance recognition could be reduced.

Finally, our system relies upon the existence of a pre-authenticated wireless medium. That is, we require the SNs to have a public/private key-pair and the MN to have these public keys already installed on the device. For some settings this might be

an unrealistic assumption. Ideally, we would not require a pre-authenticated wireless medium. Rather, there would be some kind of mechanism that would authenticate the wireless communication channel between the MN and SN. For example, by virtue of the two devices experiencing the same acceleration signal they might be able to authenticate each other by independently computing some shared key from this acceleration signal. In our experience, however, there is not enough entropy in the signal to establish such a shared key.

One promising way to pair is to use the data from the sensor to construct a shared key (e.g., using Photoplethysmograph (PPG) [144], ECG [146], or audio [61]). Mayrhofer et al. provide such a mechanism using accelerometers to exchange a cryptographic key between two devices by manually shaking the two devices together [106]. They use a method similar to the one described by Lester et al., which was used to determine whether two devices are being shaken together [91]. But, as they note, coherence "does not lend itself to directly creating cryptographic key material out of its results" [106]. To extract key material they extract quantized Discrete Fourier Transform (DCT) coefficients from the accelerometer data to use as entropy for generating a key. Bichler et al. provide a similar method using time-domain features quantized by a linear combination of pre-selected patterns that best represent the features [22, 21]. In our own preliminary experiments, this method did not work for accelerometers placed on the body (i.e., they were not deliberately shaken together).

Because these methods are inherently noisy (i.e., the accelerations do not match exactly), some protocol is necessary to resolve this problem. Mayrhofer et al. describe such a protocol they call the "Candidate Key Protocol" [104]. As the name implies, this protocol enables interactive generation of a candidate key by choosing among multiple candidate key parts. Attackers are unable to reconstruct this key since the communication of these candidate key parts are hashed with a

random nonce. Kirovski et al. describe a similar protocol called the “Martini Sync” that uses joint fuzzy hashing [85]. Their method quantizes the accelerometer stream and each parity communicates the parity of their measurement. The assumption is that if the parity matches, then their quantized values must also match. Because this assumption can be invalid, they use a progressive error-correcting code so that they can agree on a shared key. Venkatasubramanian et al. outline yet another protocol called “Physiological Signal based Key Agreement” [145]. Their protocol uses the fuzzy vault cryptographic primitive, which allows a user to lock a secret in a vault. This vault can only be unlocked it if a key is used that is similar enough to the key that was used to lock it [82]. The theoretical underpinnings of these cryptographic primitives (fuzzy extracts, fuzzy vaults, fuzzy commitments) have been described elsewhere [141]. There are many other protocols others have used [32, 17, 128].

A major alternative to all of these pairing methods is to use some characteristic of the wireless radio [83, 143, 107, 100, 101]. These methods harness the fact that in order for two devices to have knowledge of the radio environment they must be near each other. While these methods are promising, they require low-level access to the wireless radio, which may not be possible without assistance from the manufacturer.

5.5 Summary

In this chapter we described a prototype implementation that uses the methods described in Chapters 3 and 4 to realize a usably secure WBAN. Our MN prototype was a Samsung/Google Nexus S Android-based phone, while our SN prototype was a Shimmer. These two devices communicated via Bluetooth, a common short-range wireless network stack. We first described the protocols these two devices use to realize our security goals outlined in Section 2.3. We then empirically showed that our system can meet requirements of our methods via a study on the walking habits

of seven subjects. Next, we showed that our methods can be implemented in these phones in a way that does not incur significant overhead in terms of energy. Finally, we showed our protocols meet our security goals according to the trust assumptions and adversary model outlined in Section 2.3. When combined with an application, our system can provide statistical assurance that the WBAN is sensing whom it thinks it sensing.

6

Future Work

Throughout this thesis, we have outlined the limitations of our work. This chapter serves to summarize those limitations and provide direction for future work.

6.1 Bioimpedance recognition

The main limitation of our bioimpedance recognition method is its susceptibility to both movement and environmental factors. There are at least two ways of dealing with these problems. The first way is to augment our signal processing techniques to extract a more discriminating signal using better hardware (e.g., different electrodes, analog filters) or software (e.g., digital filters, more feature extraction). This is not

uncommon as most research in biometrics tends to be spent on feature engineering. The second approach is to rely on context and train a model for each context. For example, one model could be trained when a subject is moving and another when the subject is not moving. It might even possible to train a single model that incorporates context as just another feature. Clearly the limitation to such an approach is that it requires training for multiple contexts.

Although we do show bioimpedance recognition can work for long time scales, our experiments could benefit from longer term and larger population studies. It would be interesting to know whether bioimpedance recognition could be used in an enterprise setting. Having a passive biometric that could automatically authenticate and deauthenticate a user to a workstation has been the dream for many years. With sufficiently larger population and longer term studies, we believe bioimpedance in a wrist-worn form factor could be used to realize this dream. However, even if bioimpedance is not the solution to this dream, there might be room to integrate it with other passive biometrics. Vocal resonance, for example, uses a microphone to identify people based on how a person's chest cavity changes their voice [37]. It might be the case that the fusion of many biometrics is the solution to this goal.

Our wearable bioimpedance sensor can only capture bi-polar bioimpedance measurements. Although our results indicate that tetra-polar measurements provide no additional recognition performance, they might help with certain environmental factors. It might also be the case that the combinations of tetra-polar measurements we did not examine provide better recognition performance. There are, however, hardware challenges that must be overcome before tetra-polar measurements could be realized in a wearable device.

There might be other uses for bioimpedance now that we have a device that can continuously and unobtrusively measure it. We know, for example, that whole-body bioimpedance can be used to estimate fat content and other measurements of body

composition [51]. Others have shown how electro-dermal activity can be used to measure emotional and physical responses by use of a wearable device [121]. Such sensors passively measure the change in resistance of the skin, but there might be applications that benefit from skin impedance measurements.

Finally, we did not explore any kind of sophisticated spoofing attacks on our wearable bioimpedance device. A spoofing attack is an attack where the adversary seeks to fool our bioimpedance recognition method into believing it is sensing some target person. For example, an adversary could try to find a person with a similar bioimpedance to an intended target and attach the device to that similar person. Such an attack requires the adversary to put the device on many people until they find a matching bioimpedance as their target. Estimating the "expected number of subjects until success" remains future work and may be achieved with a larger uniqueness study. Since bioimpedance requires skin contact, it may be more feasible to attack the hardware and software of the device itself than put the device on a unknown number of subjects. Securing the hardware and software from such attacks remains future work. An adversary could also attempt to create a model of a target's wrist that mimics their bioimpedance. To do so, however, the adversary needs to acquire a subject's bioimpedance so that they can model it exactly. Again, it would be easier to target the hardware and software of the device than attempt to measure the target's bioimpedance at their wrist. Exploring the feasibility of such anatomical models for attack purposes remains future work.

6.2 Same-body verification

The central limitation of our same-body verification method is that we require the subject to be walking. Not only is this requirement somewhat specific, it's just not the case that all subjects will be able to walk. A subject, for example, might be confined

to a wheel chair. Ideally, the requirement would not specifically require walking but to require some kind of movement. Clearly, some kind of acceleration is necessary since a flat signal provides no information. However, it might be possible to use these times when a subject is not moving to accomplish same-body verification. That is, over long time scales, you might expect the accelerometers to experience periods of no acceleration when the subject is remaining sufficiently still perhaps even at specific frequencies. This a time-domain approach rather than the frequency-based approach we take.

Our method could also benefit from prior knowledge about the orientation of the sensor. Because of our desire to remain orientation-agnostic, we only examine the magnitude of the 3-axis accelerometer. However, many activity recognition methods examine all three axes independently since there might be information encoded in a particular axis. Exploring techniques for extracting information that would help to better discriminate same-body verification from each axis might improve the accuracy of our method. Alternatively, exploring techniques for determining the orientation of the sensor could also benefit our method.

Although our experiments indicate that the gyroscope is not as useful as an accelerometer, it may be useful for certain applications or for capturing activities that are not walking. It might also be the case that different features of the gyroscope data would yield better verification rates. For example, if two devices can be accurately tracked in space via an inertial measurement unit (i.e., a gyroscope and accelerometer combined), then so long as those two devices remaining within some distance of each other time we can say, with high probability, they are on the same body. However, algorithms for dead reckoning using an inertial measurement unit suffer from cumulative error. This problem might be overcome by assuming some model of movement according to how the human body moves.

Finally, our method leverages a single statistical measure called coherence.

There are, however, many such statistical measures that one can use [33]. One avenue of exploration is to use many statistical measures that can capture the similarity of two signals. Having many such measures allows each statistic to specialize on the relationship between two particular kinds of signal (e.g., linear versus non-linear). Combining these measures together might allow us to relax the walking assumption to a more general motion assumption.

6.3 The system as a whole

The primary limitation of our system is the lack of a formal usability study. While we claim our system is usable, our methods could benefit from a comparison with competing systems. It's important to validate claims about the effectiveness of such systems [74]. While our system does not make any claims about the effectiveness for a certain kind of application, it would be beneficial to try using these methods in a real-world system. This would allow us to not only understand the usability of such methods in comparison to other methods but also how they affect the usability of a system as a whole. The future of WBANs are in their use in medical-related applications. This work is an enabling technology.

The second limitation of our system is that it does not handle ephemeral SNs. Recall that ephemeral SNs are those SNs that a subject does not interact with very often and, when they do, only for a very short time. A good example of an ephemeral SN is a weight scale. Because the interaction time is very short, same-body verification will not work. Finding methods that allow usable interaction with ephemeral SNs remains an open research problem.

The largest outstanding problem is the pairing problem. It would be beneficial to integrate techniques that do not require the use of a public-key infrastructure. For WBANs, there are some of existing methods (e.g., Yuan et al. [160]) that could be

integrated into our system. Finding a solution that is generally applicable to WBANs remains an open research problem.

Finally, our system could benefit from integrated hardware. Our prototype is mainly based around the Shimmer platform. Because the Shimmer platform is meant to be used for a variety of applications, the energy usage is higher than a platform that is optimized for specific applications. More so, our system could benefit from a more wearable form factor that could enforce desirable properties like proper orientation of the sensors. The bioimpedance sensor, for example, could benefit from a bracelet that forces the subject to wear it a certain way (i.e., a subject's wrist is oval in shape and not perfectly circular). Tight integration with hardware might not only improve the comfort of the wearable device, but also the accuracy and energy profile of our methods. It might even be possible to integrate some form of kinetic energy harvesting to power our methods since our methods do require the subject to walk [59].

7
Conclusion

A network of pervasive wireless devices is the future. These devices will allow us to sense and collect an unprecedented amount of physiological data about ourselves. From these physiological signals we will learn to extract our psychological state as well. This era of "quantified self" will not be without its security and privacy problems if we do not look ahead. In this thesis, we have described two complementary security mechanisms for these networks of wireless sensors.

First, we showed how to identify who is wearing one of these devices using bioimpedance. Bioimpedance is a physiological property of tissue related to its ability to both store and a resist an electrical charge. We designed a wearable bracelet

capable of sensing the bioimpedance of a person's wrist using eight electrodes. We showed how this device could be used to either identify a subject from a known cohort or to verify an individual subject. In our experiments, our method achieved identification rates of 98 % BAC in a cross-validation study and 81 % BAC in a hold-out validation study. Our method also achieved verification rates of 12.9 % EER in a hold-out validation study, and for 3 subjects, our method achieved verification rates of 13.3 % EER even after more than four months had elapsed between the training and testing bioimpedance samples.

Second, we showed how to verify whether two of these devices are on the same body. To accomplish this verification, we make the reasonable assumption that each device has an accelerometer. Using these accelerometers, our method uses *coherence*, a frequency-based statistic, to determine whether these accelerometers, and by extension the wireless sensors, are on the same body. We showed that our method works when subjects are walking for at least 32 seconds and achieves a BAC of 86 %. However, performance degrades when subjects are walking in stride, where they experience increased false negatives and little increase in false positives.

Finally, we showed how to combine these two methods together to enable a WBAN to identify who is wearing it. If SN A knows person P is wearing it (by using bioimpedance recognition) and SN B knows it is on the same body as SN A (by using same-body verification), then SN B must also be worn by person P according to the transitive property. This relationship means only one SN needs to implement a biometric rather than all of them. We outlined the protocol that combines these two methods to achieve this relationship and show that this protocol achieves our desired security goals under the given trust assumptions and a realistic adversary model. We also showed that a prototype system implementing these methods and protocols does not incur significant overhead compared to the underlying operating system and wireless radio. Finally, we empirically showed that about half the time a

subject is walking they are doing so for at least 40 seconds.

While our proposed usable security mechanisms do not solve all the outstanding security problems in WBANs, we hope our contributions will be useful to the success of WBANs becoming adopted and ubiquitous, especially in the mobile health arena.

List of Abbreviations

AES	Advanced Encryption Standard
AL	Left Ankle
AR	Right Ankle
BAC	Balanced Accuracy
BAN	Body-Area Network
CL	Left Chest
DCT	Discrete Fourier Transform
ECG	Electrocardiography
EDR	Electro-dermal Response
EER	Equal Error Rate
EHR	Electronic Health Record
EM	Expectation-Maximization
EMG	Electromyography
FAR	False Accept Rate
FDA	Food and Drug Administration
FFT	Fast Fourier Transform

FN	False Negative
FOREST	Random Forest
FP	False Positive
FRR	False Reject Rate
GMM	Gaussian Mixture Model
HL	Left Hand
HR	Right Hand
IEEE	Institute of Electrical and Electronics Engineers
KNN	k-Nearest Neighbors
MN	Mobile Node
NB	Naive Bayes
PHR	Personal Health Record
PPG	Photoplethysmograph
QFN	Quad-flat no-leads
RAND	Random
SAR	Specific Absorption Rate
SD	Secure Digital
SI	International System
SN	Sensor Node
SVM	Support Vector Machine
TN	True Negative
TP	True Positive
WBAN	Wireless Body-Area Network
WR	Right Waist

List of Tables

List of Figures

List of Algorithms

Bibliography

[1] Topward 6303D. Online at http://www.topward.com/ps_6000.htm, visited Apr. 2013. Citation on page 42.

[2] 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. Citation on pages 44 and 51.

[3] Analog Devices AD7677. Online at http://www.analog.com/en/analog-to-digital-converters/ad-converters/ad7677/products/product.html, visited Apr. 2013. Citation on page 42.

[4] Analog Devices AD9754. Online at http://www.analog.com/en/digital-to-analog-converters/da-converters/ad9754/products/product.html, visited Apr. 2013. Citation on page 42.

[5] Analog Devices AD9852. Online at http://www.analog.com/en/rfif-components/direct-digital-synthesis-dds/ad9852/products/product.html, visited Apr. 2013. Citation on page 42.

[6] Analog Devices ADG1608 8-Channel Multiplexor. Online at http://www.analog.com/en/switchesmultiplexers/multiplexers-muxes/adg1608/products/product.html, visited Apr. 2013. Citation on page 46.

[7] Analog Devices ADSP-21065L. Online at http://www.analog.com/en/processors-dsp/sharc/adsp-21065l/products/product.html, visited July 2013. Citation on page 42.

[8] 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. Citation on page 83.

[9] N. Amini, M. Sarrafzadeh, A. Vahdatpour, and W. Xu. Accelerometer-based on-body sensor localization for health and medical monitoring applications. *Pervasive and Mobile Computing*, Sept. 2011. DOI 10.1016/j.pmcj.2011.09.002. Citation on page 111.

[10] 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. Citation on page 46.

[11] ANT Wireless. Online at http://www.thisisant.com/, visited Apr. 2012. Citation on page 3.

[12] Apex Fitness BodyBugg. Online at http://www.bodybugg.com/, visited Oct. 2010. Citation on pages 1 and 11.

[13] 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. Citation on page 30.

[14] H. Baldus, S. Corroy, A. Fazzi, K. Klabunde, and T. Schenk. Human-centric connectivity enabled by body-coupled communications. *IEEE Communications Magazine*, 47(6):172–178, June 2009. DOI 10.1109/mcom.2009.5116816. Citation on page 120.

[15] A. Banerjee, K. Venkatasubramanian, and S. K. S. Gupta. Challenges of implementing cyber-physical security solutions in body area networks. In *Proceedings of the Fourth International Conference on Body Area Networks (BodyNets)*, pages 1–8, Brussels, 2009. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering). DOI 10.4108/icst.bodynets2009.6031. Citation on page 147.

[16] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. In A. Ferscha and F. Mattern, editors, *Pervasive Computing*, volume 3001 of *Lecture Notes in Computer Science*, chapter 1, pages 1–17. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. DOI 10.1007/978-3-540-24646-6_1. Citation on pages 97 and 149.

[17] S. Bao, C. Poon Carmen, L. Shen, and Y. Zhang. Authenticated symmetric-key establishment for medical body sensor networks. *Journal of Electronics (China)*, 24(3):421–427–427, May 2007. DOI 10.1007/s11767-006-0152-z. Citation on page 151.

[18] A. T. Barth, M. A. Hanson, H. C. Powell, D. Unluer, S. G. Wilson, and J. Lach. Body-Coupled Communication for Body Sensor Networks. In *Proceedings of the 3rd International ICST Conference on Body Area Networks (BODYNETS)*. ICST, Mar. 2008. DOI 10.4108/ICST.BODYNETS2008.2964. Citation on page 120.

[19] H. Beigi, editor. *Fundamentals of Speaker Recognition*. Springer, 2011 edition, Dec. 2011. Online at http://www.worldcat.org/isbn/0387775919. Citation on page 83.

[20] D. Bernstein. Curve25519: New diffie-hellman speed records. In M. Yung, Y. Dodis, A. Kiayias, and T. Malkin, editors, *Public Key Cryptography - PKC 2006*, volume 3958 of *Lecture Notes in Computer Science*, pages 207–228. Springer Berlin Heidelberg, 2006. DOI 10.1007/11745853_14. Citation on page 128.

[21] D. Bichler, G. Stromberg, and M. Huemer. Synchronizing Shaking Sequences for Generating Symmetric Keys. In *Proceedings of the 2nd International Workshop on Nonlinear Dynamics and Synchronization (INDS)*, pages 75–80. IEEE, July 2009. Online at http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5227986. Citation on page 150.

[22] D. Bichler, G. Stromberg, M. Huemer, and M. Löw. Key generation based on acceleration data of shaking processes. In *Proceedings of the 9th International Conference on Ubiquitous Computing (UbiComp)*, volume 4717 of *LNCS*, pages 304–317. Springer, Sept. 2007. DOI 10.1007/978-3-540-74853-3_18. Citation on page 150.

[23] Bluetooth SIG. Online at http://www.bluetooth.com/, visited Apr. 2012. Citation on pages 3 and 133.

[24] R. M. Bolle, J. H. Connell, S. Pankanti, N. K. Ratha, and A. W. Senior. *Guide to biometrics*. Springer Professional Computing, 2003. Online at https://www.springer.com/computer/image+processing/book/978-0-387-40089-1. Citation on page 24.

[25] Bosch. BMI055: Small, versatile 6DoF sensor module. Online at http://ae-bst.resource.bosch.com/media/products/dokumente/bmi055/BST-BMI055-DS000-06.pdf, visited July 2013. Citation on page 96.

[26] 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 '92, pages 144–152, New York, NY, USA, 1992. ACM. DOI 10.1145/130385.130401. Citation on pages 19 and 39.

[27] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, Oct. 2001. DOI 10.1023/a:1010933404324. Citation on page 18.

[28] T. Brezmes, J.-L. Gorricho, and J. Cotrina. Activity recognition from accelerometer data on a mobile phone. In S. Omatu, M. Rocha, J. Bravo, F. Fernández, E. Corchado, A. Bustillo, and J. Corchado, editors, *Distributed Computing, Artificial Intelligence, Bioinformatics, Soft Computing, and Ambient Assisted Living*, volume 5518 of *Lecture Notes in Computer Science*, chapter 120, pages 796–799. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. DOI 10.1007/978-3-642-02481-8_120. Citation on pages 97 and 149.

[29] M. J. Burge and K. W. Bowyer, editors. *Handbook of Iris Recognition*. Springer London, London, 2013. DOI 10.1007/978-1-4471-4402-1. Citation on page 83.

[30] CC2420: Single-Chip 2.4 GHz IEEE 802.15.4 Compliant and ZigBee Ready RF Transceiver. Online at http://www.ti.com/product/cc2420, visited Aug. 2013. Citation on page 13.

[31] J. Cheng, O. Amft, and P. Lukowicz. Active capacitive sensing: Exploring a new wearable sensing modality for activity recognition. In P. Floréen, A. Krüger, and M. Spasojevic, editors, *Pervasive Computing*, volume 6030 of *Lecture Notes*

in Computer Science, pages 319–336. Springer Berlin / Heidelberg, 2010. DOI 10.1007/978-3-642-12654-3_19. Citation on page 149.

[32] S. Cherukuri, K. K. Venkatasubramanian, and S. K. S. Gupta. BioSec: A biometric based approach for securing communication in wireless networks of biosensors implanted in the human body. In *Proceedings of the International Conference on Parallel Processing Workshops*, pages 432–439. IEEE Computer Society, Oct. 2003. DOI 10.1109/ICPPW.2003.1240399. Citation on page 151.

[33] S.-S. Choi, S.-H. Cha, and C. C. Tappert. A survey of binary similarity and distance measures. *Journal on Systemics, Cybernetics and Informatics*, 8(1):43–48, 2010. Citation on page 157.

[34] 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 '12, pages 99–102, New York, NY, USA, 2012. ACM. DOI 10.1145/2370216.2370233. Citation on page 148.

[35] N. R. C. Committee on Identifying the Needs of the Forensic Sciences Community. *Strengthening Forensic Science in the United States: A Path Forward*. The National Academies Press, 2009. Online at http://www.nap.edu/openbook.php?record_id=12589. Citation on page 28.

[36] C. Cornelius and D. Kotz. On usable authentication for wireless body area networks. In *Proceedings of the USENIX Workshop on Health Security and Privacy (HealthSec)*, Aug. 2010. Online at http://www.cs.dartmouth.edu/~dfk/papers/abstracts/cornelius-healthsec10.html. Citation on pages 4 and 25.

[37] C. Cornelius, Z. Marois, J. Sorber, R. Peterson, S. Mare, and D. Kotz. Vocal resonance as a biometric for pervasive wearable devices. Technical Report TR2013-741, Dartmouth College Computer Science Department, Dec. 2013. Citation on page 154.

[38] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, Sept. 1995. DOI 10.1023/a:1022627411411. Citation on page 18.

[39] A. Criminisi, J. Shotton, and E. Konukoglu. Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning. *Foundations and Trends in Computer Graphics and Vision*, 7(2–3):81–227, Feb. 2012. DOI 10.1561/0600000035. Citation on page 18.

[40] 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. Citation on page 74.

[41] W. Diffie and M. Hellman. New directions in cryptography. *Information Theory, IEEE Transactions on*, 22(6):644–654, 1976. DOI 10.1109/TIT.1976.1055638. Citation on page 130.

[42] Digi-Key Corporation. BMI055 Bosch Sensortec | 828-1043-2-ND | DigiKey. Online at http://www.digikey.com/product-detail/en/BMI055/828-1043-2-ND/4196669, visited July 2013. Citation on page 96.

[43] Digi-Key Corporation. L3GD20TR STMicroelectronics | 497-12081-2-ND | DigiKey. Online at http://www.digikey.com/product-detail/en/L3GD20TR/497-12081-2-ND/2793099, visited July 2013. Citation on page 96.

[44] Digi-Key Corporation. MMA8653FCR1 Freescale Semiconductor | MMA8653FCR1TR-ND | DigiKey. Online at http://www.digikey.com/product-detail/en/MMA8653FCR1/MMA8653FCR1TR-ND/3831439, visited July 2013. Citation on page 94.

[45] DT9835. Online at http://www.datatranslation.com/products/dataacquisition/usb/DT9835/, visited Apr. 2013. Citation on page 42.

[46] T. El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Proceedings Advances in Cryptology (CRYPTO)*, pages 10–18, New York, NY, USA, 1985. Springer-Verlag New York, Inc. Online at http://dl.acm.org/citation.cfm?id=19478.19480. Citation on page 128.

[47] J. Espina, H. Baldus, T. Falck, O. Garcia, and K. Klabunde. *Towards Easy-to-Use, Safe, and Secure Wireless Medical Body Sensor Networks*, chapter 9, pages 159–179. IGI Global, Apr. 2009. DOI 10.4018/978-1-60566-332-6.ch009. Citation on page 147.

[48] T. Falck, H. Baldus, J. Espina, and K. Klabunde. Plug 'n Play Simplicity for Wireless Medical Body Sensors. *Mobile Networks and Applications*, 12(2):143–153, 2007. DOI 10.1007/s11036-007-0016-2. Citation on page 120.

[49] J. Fei and I. Pavlidis. Thermistor at a distance: Unobtrusive measurement of breathing. *IEEE Transactions on Biomedical Engineering*, 57(4):988–998, Apr. 2010. DOI 10.1109/tbme.2009.2032415. Citation on page 90.

[50] Fitbit.com. Fitbit. Online at http://www.fitbit.com/, visited Mar. 2010. Citation on page 1.

[51] K. R. Foster and H. C. Lukaski. Whole-body impedance–what does it measure? *The American Journal of Clinical Nutrition*, 64(3 Suppl), Sept. 1996. Online

at http://view.ncbi.nlm.nih.gov/pubmed/8780354. Citation on pages 83 and 155.

[52] Freescale Semiconductor. ±1.5g - 6g Three Axis Low-g Micromachined Accelerometer. Online at https://www.sparkfun.com/datasheets/Accelerometers/MMA7260Q-Rev1.pdf, visited July 2013. Citation on pages 92 and 109.

[53] Freescale Semiconductor. ±1.5g, ±6g Three Axis Low-g Micromachined Accelerometer. Online at http://www.freescale.com/files/sensors/doc/data_sheet/MMA7361L.pdf, visited July 2013. Citation on pages 92 and 103.

[54] Freescale Semiconductor. Xtrinsic MMA8653FC 3-Axis, 10-bit Digital Accelerometer. Online at http://cache.freescale.com/files/sensors/doc/data_sheet/MMA8653FC.pdf, visited July 2013. Citation on page 94.

[55] 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. Citation on pages 33 and 148.

[56] M. Garbey, N. Sun, A. Merla, and I. Pavlidis. Contact-free measurement of cardiac pulse based on the analysis of thermal imagery. *IEEE Transactions on Biomedical Engineering*, 54(8):1418–1426, Aug. 2007. DOI 10.1109/tbme.2007.891930. Citation on page 90.

[57] Google. Google Glass. Online at http://www.google.com/glass/start/, visited July 2012. Citation on page 10.

[58] Google. Location APIs | Android Developers. Online at https://developer.android.com/google/play-services/location.html, visited Aug. 2013. Citation on page 149.

[59] M. Gorlatova, J. Sarik, M. Cong, I. Kymissis, and G. Zussman. Movers and Shakers: Kinetic Energy Harvesting for the Internet of Things. Online at http://arxiv.org/abs/1307.0044, visited June 2013. Citation on page 158.

[60] 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. Citation on page 34.

[61] T. Halevi and N. Saxena. On Pairing Constrained Wireless Devices Based on Secrecy of Auxiliary Channels: The Case of Acoustic Eavesdropping. In *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS)*, pages 97–108. ACM, Oct. 2010. DOI 10.1145/1866307.1866319. Citation on page 150.

[62] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: an update. *SIGKDD Explor. Newsl.*, 11(1):10–18, Nov. 2009. DOI 10.1145/1656274.1656278. Citation on page 134.

[63] R. Halter, A. Hartov, and K. D. Paulsen. Design and implementation of a high frequency electrical impedance tomography system. *Physiological Measurement*, 25(1):379–390, Feb. 2004. Online at http://view.ncbi.nlm.nih.gov/pubmed/15005331. Citation on page 32.

[64] 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):650–659, Feb. 2008. DOI 10.1109/tbme.2007.903516. Citation on page 41.

[65] E. B. Hamida, M. Maman, B. Denis, and L. Ouvry. Localization performance in Wireless Body Sensor Networks with beacon enabled MAC and space-time dependent channel model. In *IEEE 21st International Symposium on Personal,*

Indoor and Mobile Radio Communications Workshops, pages 128–133. IEEE, 2010. DOI 10.1109/pimrcw.2010.5670416. Citation on pages 89 and 118.

[66] M. A. Hanson, H. C. Powell, A. T. Barth, K. Ringgenberg, B. H. Calhoun, J. H. Aylor, and J. Lach. Body Area Sensor Networks: Challenges and Opportunities. *IEEE Computer*, 42(1):58–65, Jan. 2009. DOI 10.1109/MC.2009.5. Citation on pages 2 and 10.

[67] 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. Citation on page 84.

[68] Z. He, Z. Liu, L. Jin, L.-X. Zhen, and J.-C. Huang. Weightlessness feature – a novel feature for single tri-axial accelerometer based activity recognition. In *Proceedings of the 19th International Conference on Pattern Recognition*, pages 1–4. IEEE, Dec. 2008. DOI 10.1109/icpr.2008.4761688. Citation on pages 136 and 149.

[69] Health-O-Meter. Health-O-Meter HDTM012DQ-69 Digital Tape Measure. Online at http://www.amazon.com/Health-Meter-HDTM012DQ-69-Digital-Measure/dp/B008CENXCS, visited Mar. 2013. Citation on page 63.

[70] Hirose Electric Co., Ltd. Hirose 3260-8S3(55) Connector. Online at http://www.hirose.co.jp/cataloge_hp/e23200014.pdf, visited Mar. 2013. Citation on page 46.

[71] Coppersmith Gordon Schermer and Brockelman. HITECH Act expands HIPAA privacy and security rules. Online at http://www.azhha.org/member_and_

media_resources/documents/HITECHAct.pdf, visited Nov. 2009. Citation on page 2.

[72] 3-Axis Digital Compass IC HMC5843. Online at http://www.honeywell.com/sites/servlet/com.merx.npoint.servlets.DocumentServlet?docid=DA9ACFE3C-F7C0-9998-6085-D9D84941499D, visited Aug. 2013. Citation on page 103.

[73] IEEE 802.15 Task Group 6 (TG6) Body Area Networks. Online at http://www.ieee802.org/15/pub/TG6.html, visited May 2012. Citation on page 10.

[74] S. S. Intille. Closing the evaluation gap in ubihealth research. *IEEE Pervasive Computing*, 12(2):76–79, Apr. 2013. DOI 10.1109/mprv.2013.28. Citation on page 157.

[75] InvenSense. IDG-500 Dual-Axis Gyro Product Specification. Online at http://invensense.com/mems/gyro/documents/PS-IDG-0500B-00-08.pdf, visited July 2013. Citation on pages 94 and 103.

[76] InvenSense. ISZ-500 Single-Axis Z-Gyro Product Specification. Online at http://invensense.com/mems/gyro/documents/PS-ISZ-0500B.pdf, visited July 2013. Citation on pages 94 and 103.

[77] H. Iwama, M. Okumura, Y. Makihara, and Y. Yagi. The OU-ISIR gait database comprising the large population dataset and performance evaluation of gait recognition. *IEEE Transactions on Information Forensics and Security*, 7(5):1511–1521, Oct. 2012. DOI 10.1109/tifs.2012.2204253. Citation on page 111.

[78] A. K. Jain, P. Flynn, and A. A. Ross, editors. *Handbook of Biometrics*. Springer-Verlag, 2007. Online at http://www.springer.com/computer/computer+imaging/book/978-0-387-71040-2. Citation on pages 26 and 83.

[79] D. Johnson and A. Menezes. The Elliptic Curve Digital Signature Algorithm (ECDSA), 1999. Citation on page 128.

[80] E. Jovanov and A. Milenkovic. Body Area Networks for Ubiquitous Healthcare Applications: Opportunities and Challenges. *Journal of Medical Systems*, 35(5):1245–1254, 2011. DOI 10.1007/s10916-011-9661-x. Citation on page 1.

[81] E. Jovanov, A. Milenkovic, C. Otto, and P. C. de Groen. A wireless body area network of intelligent motion sensors for computer assisted physical rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 2(6), 2005. DOI 10.1186/1743-0003-2-6. Citation on page 1.

[82] A. Juels and M. Sudan. A fuzzy vault scheme. In *Proceedings IEEE International Symposium on Information Theory*, pages 408+. IEEE, 2002. DOI 10.1109/isit.2002.1023680. Citation on page 151.

[83] A. Kalamandeen, A. Scannell, E. de Lara, A. Sheth, and A. LaMarca. Ensemble: Cooperative Proximity-based Authentication. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 331–344. ACM, June 2010. DOI 10.1145/1814433.1814466. Citation on page 151.

[84] 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. Citation on page 74.

[85] D. Kirovski, M. Sinclair, and D. Wilson. The Martini Synch: Joint Fuzzy Hashing Via Error Correction. In *Proceedings of the 4th European Workshop on Security and Privacy in Ad-hoc and Sensor Networks (ESAS)*, volume 4572 of *LNCS*, pages 16–30. Springer, July 2007. DOI 10.1007/978-3-540-73275-4_2. Citation on page 151.

[86] N. C. Krishnan and D. J. Cook. Activity recognition on streaming sensor data. *Pervasive and Mobile Computing*, July 2012. DOI 10.1016/j.pmcj.2012.07.003. Citation on page 149.

[87] A. Kumar, N. Saxena, G. Tsudik, and E. Uzun. A comparative study of secure device pairing methods. *Pervasive and Mobile Computing*, 5(6):734–749, 2009. DOI 10.1016/j.pmcj.2009.07.008. Citation on page 13.

[88] K. Kunze and P. Lukowicz. Dealing with sensor displacement in motion-based onbody activity recognition systems. In *Proceedings of the 10th International Conference on Ubiquitous Computing*, pages 20–29, New York, NY, USA, 2008. ACM. DOI 10.1145/1409635.1409639. Citation on pages 121 and 122.

[89] K. Kunze, P. Lukowicz, H. Junker, and G. Tröster. Where am I: Recognizing on-body positions of wearable sensors. In T. Strang and C. Linnhoff-Popien, editors, *Proceedings of the International Workshop on Location- and Context-Awareness (LoCa)*, volume 3479 of *Lecture Notes in Computer Science*, pages 264–275. Springer Berlin Heidelberg, 2005. DOI 10.1007/11426646_25. Citation on pages 111 and 121.

[90] O. D. Lara, A. J. Pérez, M. A. Labrador, and J. D. Posada. Centinela: A human activity recognition system based on acceleration and vital sign data. *Pervasive*

and Mobile Computing, 8(5):717–729, Oct. 2012. DOI 10.1016/j.pmcj.2011.06.004. Citation on page 149.

[91] J. Lester, B. Hannaford, and G. Borriello. "Are You with Me?" – Using Accelerometers to Determine If Two Devices Are Carried by the Same Person. In *Proceedings of the 2nd International Conference on Pervasive Computing (Pervasive)*, volume 3001 of *LNCS*, pages 33–50. Springer, Apr. 2004. DOI 10.1007/978-3-540-24646-6_3. Citation on pages 100, 120, 121, 125, and 150.

[92] S. Z. Li and A. K. Jain, editors. *Handbook of Face Recognition*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, second edition, Aug. 2011. Online at http://portal.acm.org/citation.cfm?id=1062383. Citation on page 83.

[93] T. Maekawa, Y. Kishino, Y. Sakurai, and T. Suyama. Activity recognition with hand-worn magnetic sensors. *Personal and Ubiquitous Computing*, pages 1–10, June 2012. DOI 10.1007/s00779-012-0556-8. Citation on page 149.

[94] T. Maekawa and S. Watanabe. Training data selection with user's physical characteristics data for acceleration-based activity modeling. *Personal and Ubiquitous Computing*, pages 1–13, Dec. 2011. DOI 10.1007/s00779-011-0491-0. Citation on page 149.

[95] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, editors. *Handbook of Fingerprint Recognition*. Springer Professional Computing. Springer-Verlag, New York, second edition, May 2009. DOI 10.1007/b97303. Citation on page 83.

[96] S. Mare, J. Sorber, M. Shin, C. Cornelius, and D. Kotz. Hide-n-Sense: Preserving Privacy Efficiently in Wireless mHealth. *Mobile Networks and Applications*,

pages 1–14, 2013. DOI 10.1007/s11036-013-0447-x. Citation on pages 130 and 147.

[97] M. Maróti, B. Kusy, G. Simon, and A. Lédeczi. The flooding time synchronization protocol. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, SenSys '04, pages 39–49, New York, NY, USA, 2004. ACM. DOI 10.1145/1031495.1031501. Citation on page 117.

[98] 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. Citation on page 27.

[99] 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. Citation on page 83.

[100] S. Mathur, R. Miller, A. Varshavsky, W. Trappe, and N. Mandayam. ProxiMate: Proximity-based Secure Pairing using Ambient Wireless Signals. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 211–224. ACM, June 2011. DOI 10.1145/1999995.2000016. Citation on page 151.

[101] S. Mathur, W. Trappe, N. Mandayam, C. Ye, and A. Reznik. Radio-telepathy: Extracting a Secret Key from an Unauthenticated Wireless Channel. In *Proceedings of the 14th ACM international conference on Mobile Computing and*

Networking (MobiCom), pages 128–139. ACM, 2008. DOI 10.1145/1409944.1409960. Citation on page 151.

[102] R. J. S. Matias, M. B. Cunha, A. M. Mota, and R. M. Martins. Modeling capacitive coupling systems for body coupled communications. In *Proceedings of the 7th International Conference on Body Area Networks*, BodyNets '12, pages 113–119, ICST, Brussels, Belgium, Belgium, 2012. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering). Online at http://portal.acm.org/citation.cfm?id=2442718. Citation on page 120.

[103] U. Maurer, A. Smailagic, D. P. Siewiorek, and M. Deisher. Activity recognition and monitoring using multiple sensors on different body positions. In *International Workshop on Wearable and Implantable Body Sensor Networks (BSN'06)*, pages 113–116. IEEE, 2006. DOI 10.1109/bsn.2006.6. Citation on pages 100 and 149.

[104] R. Mayrhofer. The Candidate Key Protocol for Generating Secret Shared Keys from Similar Sensor Data Streams. In *Proceedings of the 4th European Workshop on Security and Privacy in Ad-hoc and Sensor Networks (ESAS)*, volume 4572 of *LNCS*, pages 1–15. Springer, July 2007. DOI 10.1007/978-3-540-73275-4_1. Citation on page 150.

[105] R. Mayrhofer. OpenUAT: The Open Source Ubiquitous Authentication Toolkit. Online at http://www.openuat.org/, visited Aug. 2013. Citation on page 134.

[106] R. Mayrhofer and H. Gellersen. Shake Well Before Use: Authentication Based on Accelerometer Data. In *Proceedings of the 5th International Conference on Pervasive Computing (Pervasive)*, volume 4480 of *LNCS*, pages 144–161.

Springer, May 2007. DOI 10.1007/978-3-540-72037-9_9. Citation on pages 121 and 150.

[107] R. Mayrhofer, H. Gellersen, and M. Hazas. Security by Spatial Reference: Using Relative Positioning to Authenticate Devices for Spontaneous Interaction. In *Proceedings of the 9th International Conference on Ubiquitous Computing (UbiComp)*, volume 4717 of *LNCS*, pages 199–216. Springer, Sept. 2007. DOI 10.1007/978-3-540-74853-3_12. Citation on page 151.

[108] B. D. Mayton. Wristque: A personal sensor wristband for smart infrastructure and control. Master's thesis, Massachusetts Institute of Technology, Feb. 2013. Online at http://resenv.media.mit.edu/pubs/theses/bmayton_thesis_final.pdf. Citation on page 82.

[109] N. S. Mazloum. Body-coupled communications: Experimental characterization, channel modeling and physical layer design. Master's thesis, Chalmers University of Technology, Dec. 2008. Online at http://publications.lib.chalmers.se/records/fulltext/87937.pdf. Citation on page 120.

[110] Microchip Technology. Microchip MCP1252 Charge Pump. Online at http://ww1.microchip.com/downloads/en/DeviceDoc/21752B.pdf, visited Feb. 2013. Citation on page 46.

[111] Microsoft. Xbox 360 + Kinect. Online at http://www.xbox.com/en-US/Xbox360, visited Mar. 2013. Citation on page 1.

[112] Monsoon Power Monitor. Online at http://www.msoon.com/LabEquipment/PowerMonitor/, visited Aug. 2013. Citation on page 138.

[113] AccuFitness MyoTape Body Tape Measure. Online at http://www.accufitness.com/index.php/body-tape-measure-myotape, visited Apr. 2013. Citation on page 52.

[114] M. S. Nixon, T. Tan, and R. Chellappa. *Human Identification Based on Gait*. Springer US, Boston, MA, Oct. 2006. DOI 10.1007/978-0-387-29488-9. Citation on page 83.

[115] R. Paradiso, G. Loriga, and N. Taccini. A wearable health care system based on knitted integrated sensors. *IEEE Transactions on Information Technology in Biomedicine*, 9(3):337–344, Sept. 2005. DOI 10.1109/TITB.2005.854512. Citation on page 11.

[116] K. Partridge, B. Dahlquist, A. Veiseh, A. Cain, A. Foreman, J. Goldberg, and G. Borriello. Empirical Measurements of Intrabody Communication Performance under Varied Physical Configurations. In *Proceedings of the 14th annual ACM Symposium on User Interface Software and Technology (UIST)*, pages 183–190. ACM, Nov. 2001. DOI 10.1145/502348.502381. Citation on page 120.

[117] Q.-C. Pham, Y. Dhome, L. Gond, and P. Sayd. Video monitoring of vulnerable people in home environment. In S. Helal, S. Mitra, J. Wong, C. K. Chang, and M. Mokhtari, editors, *Smart Homes and Health Telematics*, volume 5120 of *Lecture Notes in Computer Science*, chapter 11, pages 90–98. Springer-Verlag, 2008. DOI 10.1007/978-3-540-69916-3_11. Citation on page 111.

[118] M. Z. Poh, D. McDuff, and R. Picard. A medical mirror for non-contact health monitoring. In *ACM SIGGRAPH 2011 Emerging Technologies*, New York, NY, USA, 2011. ACM. DOI 10.1145/2048259.2048261. Citation on page 90.

[119] M.-Z. Poh, D. J. McDuff, and R. W. Picard. Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Optics Express*, 18(10):10762–10774, May 2010. DOI 10.1364/oe.18.010762. Citation on page 90.

[120] M.-Z. Poh, D. J. McDuff, and R. W. Picard. Advancements in noncontact, multiparameter physiological measurements using a webcam. *IEEE Transactions on Biomedical Engineering*, 58(1):7–11, Jan. 2011. DOI 10.1109/tbme.2010.2086456. Citation on page 90.

[121] M.-Z. Poh, N. C. Swenson, and R. W. Picard. A wearable sensor for unobtrusive, long-term assessment of electrodermal activity. *IEEE Transactions on Biomedical Engineering*, 57(5):1243–1252, May 2010. DOI 10.1109/tbme.2009.2038487. Citation on page 155.

[122] M. Rahman, B. Carbunar, and M. Banik. Fit and vulnerable: Attacks and defenses for a health monitoring device, Apr. 2013. Online at http://arxiv.org/abs/1304.5672. Citation on page 3.

[123] N. Ravi, N. Dandekar, P. Mysore, and M. L. Littman. Activity recognition from accelerometer data. In *Proceedings of the National Conference on Artificial Intelligence*, pages 1541–1546. AAAI Press, 2005. Online at http://portal.acm.org/citation.cfm?id=1620107. Citation on pages 100 and 149.

[124] 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://portal.acm.org/citation.cfm?id=1005336. Citation on page 38.

[125] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, Feb. 1978. DOI 10.1145/359340.359342. Citation on page 128.

[126] J. Rosell, J. Colominas, P. Riu, R. Pallas-Areny, and J. G. Webster. Skin impedance from 1 Hz to 1 MHz. *IEEE Transactions on Biomedical Engineering*, 35(8):649–651, Aug. 1988. DOI 10.1109/10.4599. Citation on page 32.

[127] Ruby Programming Language. Online at http://www.ruby-lang.org/en/, visited Aug. 2013. Citation on page 43.

[128] R. V. Sampangi, S. Dey, S. R. Urs, and S. Sampalli. A Security Suite for Wireless Body Area Networks. *International Journal of Network Security & Its Applications*, 4(1):97–116, Jan. 2012. DOI 10.5121/ijnsa.2011.3103. Citation on page 151.

[129] L. A. Saxon. Ubiquitous Wireless ECG Recording: A Powerful Tool Physicians Should Embrace. *Journal of Cardiovascular Electrophysiology*, Feb. 2013. DOI 10.1111/jce.12097. Citation on page 10.

[130] T. C. W. Schenk, N. S. Mazloum, L. Tan, and P. Rutten. Experimental Characterization of the Body-Coupled Communications Channel. In *Proceedings of the International Symposium on Wireless Communication Systems (ISWCS)*. IEEE, Oct. 2008. DOI 10.1109/ISWCS.2008.4726053. Citation on page 120.

[131] 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. Citation on page 147.

[132] Shimmer Research. Online at http://www.shimmer-research.com/, visited Apr. 2013. Citation on pages 44, 48, 103, 133, and 135.

[133] S.-k. Song, J. Jang, and S. Park. An efficient method for activity recognition of the elderly using tilt signals of tri-axial acceleration sensor. In S. Helal, S. Mitra, J. Wong, C. K. Chang, and M. Mokhtari, editors, *Smart Homes and Health Telematics*, volume 5120 of *Lecture Notes in Computer Science*, chapter 12, pages 99–104. Springer-Verlag, 2008. DOI 10.1007/978-3-540-69916-3_12. Citation on page 149.

[134] 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 *Proceedings of the Workshop on Mobile Computing Systems and Applications (HotMobile)*, Feb. 2012. DOI 10.1145/2162081.2162092. Citation on pages 4, 8, and 126.

[135] SparkFun Electronics. WiTilt v2.5. Online at http://www.sparkfun.com/datasheets/Sensors/WiTilt_V2_5.pdf, visited October 2010. Citation on page 109.

[136] 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. Citation on pages 28 and 84.

[137] J. Sriram, M. Shin, T. Choudhury, and D. Kotz. Activity-aware ECG-based patient authentication for remote health monitoring. In *Proceedings of the International Conference on Multimodal Interfaces and Workshop on Machine Learning for Multi-modal Interaction (ICMI-MLMI)*, pages 297–304. ACM Press, Nov. 2009. DOI 10.1145/1647314.1647378. Citation on page 84.

[138] ST. L3GD20. Online at http://www.st.com/web/en/resource/technical/document/datasheet/DM00036465.pdf, visited July 2013. Citation on page 96.

[139] C. C. Tan, H. Wang, S. Zhong, and Q. Li. IBE-Lite: A Lightweight Identity-Based Cryptography for Body Sensor Networks. *IEEE Transactions on Information Technology in Biomedicine*, 13(6):926–932, Nov. 2009. DOI 10.1109/titb.2009.2033055. Citation on page 2.

[140] TinyOS. Online at http://www.tinyos.net/, visited Apr. 2013. Citation on page 46.

[141] P. Tuyls, B. Skoric, and T. Kevenaar, editors. *Security with Noisy Data*. Springer London, London, 2007. DOI 10.1007/978-1-84628-984-2. Citation on page 151.

[142] S. Ullah, H. Higgins, B. Braem, B. Latre, C. Blondia, I. Moerman, S. Saleem, Z. Rahman, and K. Sup. A comprehensive survey of Wireless Body Area Networks : on PHY, MAC, and Network layers solutions. *Journal of Medical Systems*, 36(3):1065–1094, June 2012. DOI 10.1007/s10916-010-9571-3. Citation on page 10.

[143] A. Varshavsky, A. Scannell, A. LaMarca, and E. De Lara. Amigo: Proximity-Based Authentication of Mobile Devices. In *Proceedings of the 9th International Conference on Ubiquitous Computing (UbiComp)*, volume 4717 of *LNCS*, pages 253–270. Springer, Sept. 2007. DOI 10.1007/978-3-540-74853-3_15. Citation on page 151.

[144] K. K. Venkatasubramanian, A. Banerjee, and E. K. S. Gupta. Plethysmogram-based Secure Inter-Sensor Communication in Body Area Networks. In *Proceedings of the Military Communications Conference (MILCOM)*. IEEE, Nov. 2008. DOI 10.1109/MILCOM.2008.4753199. Citation on pages 90 and 150.

[145] K. K. Venkatasubramanian, A. Banerjee, and S. K. Gupta. PSKA: Usable and Secure Key Agreement Scheme for Body Area Networks. *IEEE Transactions on Information Technology in Biomedicine*, 14(1):60–68, Jan. 2010. DOI 10.1109/TITB.2009.2037617. Citation on page 151.

[146] K. K. Venkatasubramanian, A. Banerjee, and S. K. S. Gupta. EKG-based Key Agreement in Body Sensor Networks. In *Proceedings of the 27th IEEE*

International Conference on Computer Communications (INFOCOM), pages 1–6. IEEE, Apr. 2008. DOI 10.1109/INFOCOM.2008.4544608. Citation on pages 90 and 150.

[147] K. K. Venkatasubramanian and S. K. S. Gupta. Physiological Value-Based Efficient Usable Security Solutions for Body Sensor Networks. *ACM Transactions on Sensor Networks (TOSN)*, 6(4), July 2010. DOI 10.1145/1777406.1777410. Citation on pages 12 and 147.

[148] W. Verkruysse, L. O. Svaasand, and J. S. Nelson. Remote plethysmographic imaging using ambient light. *Optics Express*, 16(26):21434–21445, Dec. 2008. DOI 10.1364/oe.16.021434. Citation on page 90.

[149] 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. Citation on page 84.

[150] Y. Wang, J. Lin, M. Annavaram, Q. A. Jacobson, J. Hong, B. Krishnamachari, and N. Sadeh. A framework of energy efficient mobile sensing for automatic user state recognition. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 179–192. ACM, June 2009. DOI 10.1145/1555816.1555835. Citation on page 149.

[151] Y. Wang, J. Yang, H. Liu, Y. Chen, M. Gruteser, and R. P. Martin. Sensing vehicle dynamics for determining driver phone use. In *Proceedings of the 11th Annual International Conference on Mobile Systems, Applications, and Services*, pages 41–54, New York, NY, USA, 2013. ACM. DOI 10.1145/2462456.2464447. Citation on page 122.

[152] Wikipedia. Brandon Mayfield — Wikipedia, The Free Encyclopedia, Aug. 2013. Online at http://en.wikipedia.org/wiki/Brandon_Mayfield. Citation on page 28.

[153] Wikipedia. Nexus S — Wikipedia, The Free Encyclopedia. Online at http://en.wikipedia.org/wiki/Nexus_S, visited Aug. 2013. Citation on page 133.

[154] Wikipedia. Relative permittivity — Wikipedia, the free encyclopedia. Online at https://en.wikipedia.org/wiki/Relative_permittivity, visited Apr. 2013. Citation on page 148.

[155] C. Wilson. *Vein Pattern Recognition: A Privacy-Enhancing Biometric*. CRC Press, first edition, Mar. 2010. Online at http://www.worldcat.org/isbn/1439821372. Citation on page 83.

[156] D. A. Winter. Kinematic and kinetic patterns in human gait: Variability and compensating effects. *Human Movement Science*, 3(1-2):51–76, Mar. 1984. DOI 10.1016/0167-9457(84)90005-8. Citation on page 105.

[157] H. Y. Wu, M. Rubinstein, E. Shih, J. Guttag, F. Durand, and W. Freeman. Eulerian video magnification for revealing subtle changes in the world. *ACM Transactions on Graphics*, 31(4), July 2012. DOI 10.1145/2185520.2185561. Citation on page 90.

[158] Xilinx Spartan XC2S30. Online at http://www.xilinx.com/support/documentation/data_sheets/ds001.pdf, visited Apr. 2013. Citation on page 42.

[159] R. V. Yampolskiy and V. Govindaraju. Behavioural biometrics: a survey and classification. *Int. J. Biometrics*, 1(1):81–113, June 2008. DOI 10.1504/ijbm.2008.018665. Citation on page 27.

[160] J. Yuan, L. Shi, S. Yu, and M. Li. Authenticated secret key extraction using channel characteristics for body area networks. In *Proceedings of the ACM conference on Computer and Communications Security (CCS)*, CCS '12, pages 1028–1030. ACM, 2012. DOI 10.1145/2382196.2382314. Citation on page 157.

[161] H. Zhang. The Optimality of Naive Bayes. In V. Barr and Z. Markov, editors, *FLAIRS Conference*. AAAI Press, 2004. Online at http://www.cs.unb.ca/profs/hzhang/publications/FLAIRS04ZhangH.pdf. Citation on page 17.

[162] L. Zhong, D. El-Daye, B. Kaufman, N. Tobaoda, T. Mohamed, and M. Liebschner. OsteoConduct: Wireless Body-Area Communication based on Bone Conduction. In *Proceedings of the 2nd International ICST Conference on Body Area Networks (BODYNETS)*. ICST, June 2007. DOI 10.4108/bodynets.2007.181. Citation on page 120.

[163] ZigBee Alliance. Online at http://www.zigbee.org/, visited Apr. 2012. Citation on page 3.