

Improved Forensic Medical Device Security through Eating Detection

A Thesis Presented for

The Master of Science

Degree

The University of Tennessee, Knoxville

Nathan Lee Henry

May 2014

© by Nathan Lee Henry, 2014
All Rights Reserved.

To one of the greatest men I have had the good fortune to know, my Grandfather, Zachary Adolphus Henry PhD, P.E. (1930-2013), a Christian, husband, father, grandfather, American, sailor, professor, engineer, deacon, commissioner, inventor, mentor, and farmer. He was disappointed that he would be unable to see me graduate.

Acknowledgements

I must thank Dr. Gregory Peterson for finding me this position, Dr. Nathanael Paul for letting me work with him, Dr. Nicole McFarlane for her invaluable sensor knowledge, Khandaker Mamun for allowing me learn from him, and Dr. Erik Ferragut for sharing his knowledge of machine learning.

Keep it simple stupid.

Abstract

Patients are increasingly reliant on implantable medical device systems today. For patients with diabetes, an implantable insulin pump system or artificial pancreas can greatly improve quality of life. As with any device, these devices can and do suffer from software and hardware issues, often reported as a safety event. For a forensic investigator, a safety event is indistinguishable from a potential security event. In this thesis, we show a new sensor system that can be transparently integrated into existing and future electronic diabetes therapy systems while providing additional forensic data to help distinguish between safety and security events. We demonstrate three bowel sound detection methods, the best of which has an 84.26% bowel sound classification accuracy. We provide additional contextual information by using detected bowel sounds to detect when a patient begins to eat. We achieved 100% eating detection accuracy in a laboratory environment. From the eating data, an algorithm or forensic investigator can identify potential malfeasance in a test subject.

Table of Contents

1	Introduction	1
1.1	Insulin Pumps and Security	2
1.2	Prior Work	3
1.3	Security Model	6
2	Understanding Insulin Pump System Activity	8
3	Experimental Approach	12
3.1	Experimental Data Acquisition	12
3.1.1	Pre-eating	12
3.1.2	Eating	13
3.1.3	Post-eating	13
3.1.4	Noise	13
3.2	First Steps	14
3.3	System Design	15
3.3.1	Bowel Sound Detection	15
3.3.2	Eating Detection	15
3.3.3	Event Classification and Policy Enforcement	15
4	Bowel Sound Detection	17
4.1	A Note on Bowel Sound Detection Rates	18
4.2	Sensor Data	21

4.2.1	Piezo Transducer Setup	21
4.2.2	Stethoscope Setup	21
4.3	Integration Approach	23
4.4	Filter and Threshold Method	24
4.5	Bayesian Discriminant Function Method	27
4.5.1	Windowing	28
4.5.2	Feature Extraction	29
4.5.3	Bayesian Bowel Sound Detection	30
4.6	Bowel Sound Training and Accuracy	32
4.7	Receiver Operating Characteristic	35
4.8	Noise	36
4.8.1	Talking	37
5	Eating Detection	40
5.1	Filter and Threshold Eating Detection	42
5.2	Bayesian Discriminant Function Eating Detection	43
5.3	Eating Detection and Errors in Bowel Sound Detection	44
5.4	Eating Detection Rate	45
5.5	Meal Indicating Factors	45
5.5.1	Cell Phone Factors	45
5.5.2	Insulin Therapy System Data	46
6	Event Classification and Policy Enforcement	47
6.1	Eating Model	48
6.2	Hyperglycemic Patient State	50
6.3	Hypoglycemic Patient State	51
6.4	Euglycemic Patient State	51
6.5	Insulin Therapy System Data	51
6.6	Cell Phone	51
6.7	Forensic Rules	52

6.7.1	Normal Food Bolus Event	52
6.7.2	No Bolus Forensic Event	52
6.7.3	No Food with Bolus Forensic Event	55
6.8	Classification of Safety and Security Events	56
7	Future Work	57
7.1	Bowel Sound Detection Improvements	57
7.1.1	Bowel Sound Detection in the Presence of Talking	58
7.2	Eating Detection Improvements	58
7.3	Forensic Rules and Event Classification	59
7.4	Sensor Integration	59
7.5	Implementation	61
7.6	Future Improvements Requiring Operational Changes	61
7.7	Comments on Experimental Method	62
8	Conclusions	64
8.1	Comparison to other works	64
	Bibliography	66
	Vita	70

List of Figures

2.1	An example report that is based on a leading manufacturers insulin pump system software. Dots represent finger sticks (blood glucose checks) taken over the period of a month. Specific periods of the day are assumed to correspond to a meal.	10
2.2	A Day in the Life of a Patient. This figure shows an example patient's blood glucose through a day. A lack of forensic information can harm a forensic investigator's ability to determine what contributed to a high or low blood glucose value.	11
3.1	The system is broken into three components: bowel sound detection, eating detection, and security event detection and event log.	14
4.1	System Design: This chapter addresses the bowel sound detection component.	17
4.2	Bowel Sound 1	19
4.3	Bowel Sound 2	19
4.4	Bowel Sound 3	20
4.5	Bowel Sound 4	20
4.6	The stethoscope head is attached to the body with medical tape. The stethoscope head is held in place with an a medical bandage wrapped around the body.	22

4.7	Bowel Sound Detection - Integration Approach: The integration method bowel sound detection component is composed of two subcomponents. This is a more detailed figure of the bowel sound detection component in Figure 3.1.	23
4.8	Bowel Sound Detection - Filter and Threshold Approach: The threshold and filter bowel sound detection component is composed of three subcomponents. This is a more detailed figure of the bowel sound detection component in Figure 3.1.	24
4.9	A detected bowel sound. This sound rose above the vertical threshold and was within the horizontal threshold.	25
4.10	A rejected sound. This sound rose above the vertical threshold but was longer than the horizontal threshold.	26
4.11	A potential false negative. This sound was within the horizontal threshold but fell below the vertical threshold. It was therefore not considered a bowel sound.	26
4.12	Bowel Sound Detection - Bayesian Approach: The Bayesian discriminant bowel sound detection component is composed of four subcomponents. This is a more detailed figure of the bowel sound detection component in Figure 3.1.	27
4.13	The circle represents the detected peak and the blue arrows represent the amount of time taken before and after the peak.	28
4.14	Feature extraction from wavelet decomposition filterbank.	29
4.15	Classification accuracy measurements for subject 1.	32
4.16	Classification accuracy measurements for subject 2.	33
4.17	Classification accuracy measurements for subject 3.	33
4.18	Classification accuracy measurements for subject 4.	34
4.19	Classification accuracy measurements for subject 5.	34
4.20	Receiver operating characteristic for each test subject.	35

4.21	Average power spectral density of talking for all five subjects. Subjects were asked to read continuously for a five minute period	37
4.22	Average power spectral density (over five minute period) of bowel sounds for all five subjects 20 minutes after start of eating. Subjects were asked to limit movement during this time to better record bowel activity.	38
5.1	System Design: This chapter addresses the eating detection component.	40
5.2	Eating Detection: Eating detection is composed of 3 subcomponents. This is a more detailed figure of the eating detection component in Figure 3.1.	41
5.3	Filter and threshold eating detection results. The vertical dashed line represents the start of eating.	42
5.4	Bayesian discriminant function method eating detection results. The square represents a value of 2503 bowel sounds detected in the 5 minute period. It is omitted so that the figure could be on the same scale as Figure 5.3. The vertical dashed line represents the start of eating. . .	43
6.1	System Design: This chapter address the security event detection and event log component.	48
6.2	Event Classification and Policy Enforcement: This component is composed of three subcomponents. This is a more detailed figure of the event classification and policy enforcement component in Figure 3.1.	49
6.3	Eating Model: The blood glucose check and bolus in the pre-eating period represents the recommended time period a compliant patient will administer an insulin bolus in relation to a meal. The large crooked line represents the start of eating. The small crooked line represents the end of eating. The dashed lines represent correction boluses and/or boluses given by a non-compliant patient.	50

6.4	Normal eating scenario: The patient has given checked her blood glucose value and administered an insulin bolus before the start of the of the meal.	53
6.5	No Bolus Eating Scenario: The meal begins and ends but no bolus is recorded.	53
6.6	No Meal Bolus Scenario: An insulin bolus is recorded but no meal is detected	55
7.1	Our goal is to integrate a sensor with components already used in insulin pump systems such as the infusion set, continuous glucose monitor, or the insulin pump device itself.	60

Chapter 1

Introduction

Diabetes* affects over 25.8 million people in the United States or approximately 8.3 % of the U.S. population (CDC, 2011). In 2007, the estimated cost of diabetes care in the United States was \$174 billion dollars including direct medical costs and indirect medical costs (e.g., disability, work loss, premature mortality).

Diabetes is characterized by the body's inability to produce insulin or to use insulin effectively. In a healthy patient insulin is produced to help regulate the amount of glucose in the body. Glucose is derived from the food (specifically carbohydrates) that a person ingests. The pancreas will secrete the insulin hormone to regulate a body's blood glucose level, and a higher level of insulin is needed to counteract an increased blood glucose level after consuming carbohydrates. The blood glucose level will rise as a patient consumes food.

When the blood glucose level rises above a safe level the patient is said to be in a state of hyperglycemia. This can have deleterious effects over long periods of time. Artificial insulin can be used to compensate for the body's lack of insulin. Insulin can lower the patient's blood glucose level to a safe level. However, if the amount of insulin administered to the patient is not correctly correlated to the amount of carbohydrates

*This work was first published at HealthTech 2013 under the title, *Using Bowel Sounds to Create a Forensically-aware Insulin Pump System* (Henry Jr et al., 2013). Portions of this thesis are drawn directly from the HealthTech paper.

consumed, then the blood glucose level may rise above safe levels (hyperglycemia) or fall below safe levels (hypoglycemia).

1.1 Insulin Pumps and Security

Various therapy methods exist to treat diabetes. For patients with type 2 diabetes, diet and exercise can typically help control blood glucose levels. For type 1 diabetes, patients require insulin therapy. This thesis focuses on devices that are commonly used by these patients. Insulin therapy can be administered via syringes, insulin pens, electronic insulin pumps, or closed loop artificial pancreas systems. Patients are increasingly using electronic insulin pump systems as a form of treatment.

Because blood glucose control is critical to addressing health complications related to diabetes, current insulin pump systems record blood glucose data and information related to blood glucose data. Patients and physicians use this data to improve patient blood glucose control by reprogramming the pump's insulin delivery. This data can be broken into two categories. First, the insulin pump, as the main system actuator device, records user information including user-adjusted settings, insulin basal rates, amount of delivered insulin, errors, and alarms. Second, many other system components generate data including: a continuous glucose monitor (CGM), a glucose meter, or the patient (e.g., by logging the type of food or amount of carbohydrates consumed).

These electronic insulin pump systems provide great benefits to patients. For example, rather than taking a number of large insulin injections at predetermined points in the day, the pump can programmatically and precisely release the needed amounts of insulin to a patient throughout the day. This continual release of insulin better mimics pancreatic insulin delivery. Patients can now use many devices that will benefit in blood glucose control. These pumps now use wireless communication to interact with PCs, remote controls, blood glucose meters, and continuous blood

glucose monitoring systems (in the future, they will interact with cars and other mobile systems (Ford, 2013)).

Recently, there have been several reports and demonstrations of unauthorized remote access of insulin pump systems (Benedict et al., 2004; Klonoff, 2008; Radcliffe, 2011; Barnaby, 2011; Li et al., 2011). Unauthorized remote access typically stems from insecure wireless access. Newer insulin pump system architectures are completely dependent on wireless communication. For example, some patch pump systems have no physical interface to the pump component. A user must use a remote control to issue any commands or to interact with the pump.

1.2 Prior Work

Several solutions have been suggested to prevent and detect implantable medical device security breaches. While our bowel sound system does not actively prevent security breaches, active prevention is an eventual goal. Some have used past glucose trends to detect anomalous insulin pump system behavior (Hei et al., 2013). Anomalous detection is promising in that the false positive rate (the rate at which a system identifies malicious abnormal usage when the pump usage is benign) is less than one percent. In practical use, a false positive would happen approximately once every 20-25 days (assuming a patient eats three meals a day and eats periodic snacks). This approach could potentially result in incorrect patient operational use.

As implantable medical devices make use of encrypted communication, key management is an emerging topic (Chang et al., 2012; Schechter, 2010). Dealing with key management in an emergency situation is especially difficult with respect to safety. Li et al. takes a simpler approach by suggesting rolling codes just as they are used in garage doors to protect against unauthorized entry (Li et al., 2011). This is a simple effective approach to protect against unauthorized access, but emergency access remains an issue.

In 2011, Gollakota et al. introduced an independent device that acts as a shield and transmitter (Gollakota et al., 2011). An independently worn secure communication device has been successfully prototyped for a cardiac device. Experimentation is needed to verify if the same approach would work for externally worn but subcutaneously implanted devices (e.g., insulin pump systems). For battery-powered systems, jamming would drain the battery resulting in more frequent battery changes. For insulin pump systems, requiring the patient to more frequently replace an expired battery may not be an acceptable burden.

Others have proposed using a wrist-worn watch (Sorber et al., 2012a) or a smart card (Sorber et al., 2012b) to protect against unauthorized communication. Some patients with diabetes have previously used a wrist watch for blood glucose control (the Glucowatch helped in non-invasive blood glucose control (Glu, 2013)). Similarly, a smart card could provide widely available access to stronger security primitives than what is currently available in current insulin pump systems. This is similar to the approach that some blood glucose monitoring devices use when insecurely pairing a blood glucose monitor to other devices (e.g., remote control or wireless glucose meter). One challenge to the adoption of these approaches is that portions of the patient population will not be able to use the device. Some patients will not own or use a watch, and others will not own a device capable of using a smart card (e.g., a phone).

All of these systems are promising. However, many systems that prevent security issues also require that patients make a significant operational change in how they use their medical device. When a significant operational change occurs, there is potential for a safety issue. Our goal is to provide security without weakening safety. A patient using this device would not need to operationally change how she used the pump. For our initial work, we provide improved forensic information about an implantable insulin pump system. This improved forensic information could serve as a deterrent to potential device misuse.

The most relevant forensic incident regarding intentional insulin pump misuse was in 2004 ([Benedict et al., 2004](#)). A nurse used an insulin pump to deliver fatal doses of the drug laudanone. The patient had type II diabetes but relied on an insulin pump for blood glucose control. In this case report ([Benedict et al., 2004](#)), the authors cite the potential for misuse of an insulin pump remote control. Given the increase of patients with diabetes and the growing use of insulin pump systems, the author theorizes that their misuse in homicides and/or suicides may become more prevalent.

While this earlier case provides specific evidence of malfeasance against a patient with an insulin pump system, we favor safety over rigid security rules in the design of a system that attempts to mitigate this malfeasance. The risk of negative safety events is well known ([Meier, 2010](#)). A potential safety issue that may arise from implementing a security mitigation is a design challenge.

Today, an insulin pump security incident could occur where an insulin bolus is used to cause hypoglycemia (i.e., a low blood glucose state that could lead to a diabetic coma). This security event is virtually indistinguishable from a safety event (e.g., where a patient accidentally overdoses). In the rest of this thesis, we introduce a new type of insulin pump sensor system that provides data about patient eating behavior. This thesis demonstrates the feasibility of enforcing forensic rules that are based on eating detection and describes our initial step in this process. We describe our implemented algorithm that uses eating data to demonstrate a new design for a forensically-aware insulin pump system that can detect potentially malicious events.

Our contributions include (1) a feasibility prototype for a forensically-aware insulin pump system that can help a forensic investigator focus on the events that could potentially indicate malfeasance; (2) as a side effect, we provide data from detected bowel sounds that can improve patient health; (3) we introduce a seamless method to integrate the bowel sound detection sensor into an insulin pump system; and (4) we introduce forensic rules that could be used to detect anomalous system use (i.e., when the patient’s own device is used against them).

1.3 Security Model

A security event from the misuse of an insulin pump system could cause euglycemia (a normal blood glucose value), hypoglycemia (a low blood glucose value), or hyperglycemia (a high blood glucose value). Euglycemia is normal, and the desired state of the patient. Our ultimate goal is to protect against hypoglycemia and hyperglycemia.

While hyperglycemia is harmful, the greater danger is hypoglycemia. Patients can become desensitized to hypoglycemia, and it can quickly negatively affect a patient. As we prioritize hypoglycemia, we make the following assumptions:

- We focus on those actors that can remotely interact with devices of the system (although our system can recognize anomalous behavior caused by someone with physical access to the system). Those that use a patient's own device to remotely interact with other parts of the system are of interest.
- Detection of conditions that can lead to hyperglycemia is less important than hypoglycemia, but any approach that can also detect hyperglycemia would be favored.

There are operational changes that could accomplish hypoglycemia and avoid detection policies. For example, a policy that identified a large insulin bolus given in the absence of food intake would not detect several smaller insulin boluses which could achieve the same amount of insulin delivery. To counter this issue, improved forensic policies could be implemented. For example, events could be marked where the intake of insulin over a period of time crosses a threshold where the patient has not eaten (i.e., keep the insulin in the body, insulin on board, to a safe level). This would stop the delivery of several smaller insulin boluses that sum to an amount that would be greater than a single safe insulin bolus.

This chapter introduced diabetes and the insulin therapy system as a method of diabetes management. The security problems associated with insulin therapy systems

as well as some security mitigation techniques have also been introduced. These are followed by our security model. The next chapter details how we approach insulin therapy system activity.

Chapter 2

Understanding Insulin Pump System Activity

In modern insulin pump systems, a patient will issue a bolus to correct a high blood glucose value, or she will give a bolus in response to, or in anticipation of, the consumption of food. In many insulin pumps these actions are logged and can be considered forensic information. However, if an insulin bolus causes hypoglycemia, it is not a straightforward process to determine whether the bolus was issued to prevent a case of hyperglycemia (high blood glucose) or in conjunction with food consumption. For instance, a patient could issue a bolus mistakenly (safety issue); someone other than the patient could also intentionally issue an unauthorized bolus (security issue). Improved logging of bolus events is needed. Thus, we explicitly define three forensic goals for a forensic investigator:

1. To determine what steps led to a negative patient event (e.g., miscounting of carbohydrates).
2. To determine what specifically caused a negative patient event (e.g., overdose of insulin caused hypoglycemia).

3. To take the steps that led to a negative patient event and the specific causes for that event and to determine the type of the negative patient event. The event will either be a safety event or a security event.

Fortunately, we can fulfill these three forensic goals using the information logged by an insulin pump in conjunction from the food consumption information gathered by the system demonstrated in this thesis. We show a system to forensically detect and identify the presence or absence of eating, and we use this data to determine if the patient is likely to experience an unsafe blood glucose level (hypoglycemia or hyperglycemia). Factors that could have an effect on glucose levels include food consumption, body temperature, and exercise. Food consumption is one of the primary methods by which a patient alters her blood glucose level. Our system uses bowel sounds to allow a forensic investigator to determine if a patient is eating. From the knowledge of food ingestion, an automated system could label an increased insulin delivery as anomalous and allow a forensics investigator to determine if the increased insulin is related to a benign, safety, or security event. For example, a patient could measure her blood glucose and note a low (hypoglycemic) value of 48 mg/dL. If the system had also detected that this patient had not eaten for 8 hours and that she had just issued 25 units of insulin, then the system could mark the 25-unit insulin bolus for further investigation.

Figure 2.1 shows an example graph of a patient's blood glucose values taken from finger sticks (the data has been populated with realistic but artificial values). This graph shows eating assumptions made by the software of a leading insulin pump manufacturer. The day is broken into periods corresponding to different meals. These periods are static, and a patient's behavior may not comply with assumed models of device use. Even if the system used associated bolus data to infer food consumption, the person giving the bolus is responsible for marking the bolus as one for food consumption. A forensic investigator does not have a reliable way to know why a




Figure 2.1: An example report that is based on a leading manufacturers insulin pump system software. Dots represent finger sticks (blood glucose checks) taken over the period of a month. Specific periods of the day are assumed to correspond to a meal.

large bolus is warranted. If the investigator could know when food consumption had taken place, this would help explain an instance of a potentially larger bolus.

In Figure 2.2, a patient receives a bolus at 1:00 am. The patient is asleep during this time. A forensically-aware pump that could detect and record patient eating instances would allow forensic investigators to note a possible case of hypoglycemia. The possibility of hypoglycemia coupled with the fact that the bolus was given early in the morning when a patient is assumed to be asleep could indicate an unauthorized bolus. At 6:30 am the patient's glucose level rises. This could be indicative of a meal or of the dawn phenomenon (ADA, 2011). With a forensically-aware pump, a forensic investigator may infer the absence of a meal and a possible case of hyperglycemia. At 6:30 pm a bolus is given without a corresponding meal. Abnormal bolus patterns that do not correspond with a meal could be indicative of a forgotten bolus, radio interference, or unauthorized third party interaction with the insulin pump.

As a secondary benefit, medical professionals could use this eating data to identify patient trends and identify vectors of treatment that could increase insulin pump system effectiveness. For example, a nurse may notice that a patient experiences




Figure 2.2: A Day in the Life of a Patient. This figure shows an example patient’s blood glucose through a day. A lack of forensic information can harm a forensic investigator’s ability to determine what contributed to a high or low blood glucose value.

hyperglycemia from continually failing to issue an insulin bolus for food that would require insulin. Because eating is a critical factor in a patient’s blood glucose level, eating information correlated with bolus and blood glucose data can yield much forensic information. Through experimentation we show that we are able to determine when a person is eating when they transition from a fasting (have not eaten food for at least two hours) to a fed state. The transition from a fasting to fed state should be the simplest to detect (e.g., breakfast). Later experiments should be run on test subjects that are not in a fasting state.

This chapter detailed our forensic goals as well as motivated the need for a method to detect eating to augment forensic information. The experimental approach to our method to detect eating is described in the next chapter.

Chapter 3

Experimental Approach

3.1 Experimental Data Acquisition

Our goal is to provide more forensic data about the use of an electronic diabetes therapy system (i.e., an insulin pump) to distinguish between benign, safety, and security events. This data will specifically address insulin boluses and eating. To detect eating, we experimentally detect someone eating through their acoustic bowel sounds. Using five different test subjects, we recorded patient abdominal acoustic activity before, during, and after food consumption. Data was recorded in segments with a length of five minutes. When a five minute segment ended another 5 minute segment would begin within approximately one minute.

3.1.1 Pre-eating

We required our test subjects to fast a minimum of two hours before the start of the experiment. This was to reduce the amount of bowel sound activity potentially resulting from a prior meal. Before beginning the eating phase, we measured a baseline fasting bowel sound signal while the participant sat quietly in a chair for a minimum of fifteen minutes.

3.1.2 Eating

During the eating stage, the participant was required to sit in a chair. Pizza was served to each subject. Chips and a soft drink were also offered to the subjects. The subjects could alternatively drink water or a diet soft drink. Most subjects finished the meal within 20 minutes.

3.1.3 Post-eating

After the initial eating period of up to 20 minutes, we continued to record bowel sounds for one hour and thirty minutes after the start of eating. During this time the subject was asked to remain still in a chair. The subject was allowed to read or use a personal computer. After one hour and thirty minutes, the subjects ate another piece of pizza. This was in order to determine if a meal could be detected shortly after a prior meal was eaten. We then recorded for approximately another fifteen minutes. In this work we do not address detecting an instance of eating from a non-fasting state. However, the extra 15 minutes of data may be used in future work to detect when a patient begins to eat from a non-fasting state.

3.1.4 Noise

Some noises could potentially cause false positives for our bowel sound detection algorithm. To test for potential false positive sounds, we had the participants go through a series of exercises that included, walking on a flat surface, walking up and down stairs, riding up and down an elevator, reading aloud, and having a cell phone ring in the front right pocket. For some subjects we had a conversation to simulate talking conditions. This was different from the reading aloud portion of the experiment because the subject would take natural pauses in the course of the conversation to listen.

3.2 First Steps

To annotate current insulin pump systems with forensic data about a patient's eating, a system needs to know when a patient eats. The overall approach to create a forensically-aware insulin pump system is depicted in Figure 3.1. The next three chapters will discuss each of each major figure components shown in the Figure 3.1.




Figure 3.1: The system is broken into three components: bowel sound detection, eating detection, and security event detection and event log.

3.3 System Design

We break up our forensic method into three separate components. We first record the signal from our sensor. We used an electronic stethoscope to gather the data for the results shown in this work. In the future we will employ a piezoelectric sensor.

3.3.1 Bowel Sound Detection

The bowel sound detection component is discussed in further detail in chapter 4. We implement three separate bowel sound detection techniques. In our first approach, we used an integration method influenced by Campbell ([Campbell et al., 1989](#)) (see [4.3](#)). In our second approach, we used filtering and thresholding method (see [4.4](#)). In our third and final approach, we used a Bayesian statistics method (see [4.5](#)).

3.3.2 Eating Detection

After bowel sound detection is completed, all sounds from the sensor have been time-stamped and classified as a bowel sound or not a bowel sound. The eating detector will take the time-stamped bowel sounds from the bowel sound detection component as input, and determine when the number of bowel sounds crosses a threshold. When this threshold is exceeded in some time period, eating has taken place. The eating detection algorithm will pass the time-stamped eating instances as output to the security event component. The eating detection component is discussed in further detail in chapter 5.

3.3.3 Event Classification and Policy Enforcement

This component detects and stores potential security event data. It would allow the forensic investigator to efficiently investigate negative patients events by tagging certain events as potential security events. This data could also be used as data to

train a normal use model that could be used by the security even detection component. This component is discussed in further detail in [chapter 6](#).

This chapter detailed how we acquired our experimental data. It also introduced the three components of our system model: bowel sound detection, eating detection, and event classification and policy enforcement. The first of these components, bowel sound detection, is described in the next chapter.

Chapter 4

Bowel Sound Detection




Figure 4.1: System Design: This chapter addresses the bowel sound detection component.

To produce forensic data, we record all sounds from sensor-provided data, interpret an eating event from this data by focusing on bowel sounds, and then distinguish system events that are related to eating. We interpret an instance of eating by the number of bowel sounds observed. Figure 3.1 shows how the sensor data is provided as input into a bowel sound detection algorithm that is part of the approach to create forensic data. Section 4.2 details two approaches that we evaluated to collect the sound data to test the bowel sound detection algorithm.

Bowel sounds vary in frequency, length, number of peaks, and amplitude. Figures 4.2, 4.3, 4.4, and 4.5 are several bowel sounds that show some of the variations seen in bowel sounds. Bowel sound detection algorithms must account for a large amount of variation. We attempted three bowel sound detection schemes to detect when someone was eating. These detection techniques are: 1) integration method (section 4.3), 2) filter and threshold method (section 4.4), and 3) the Bayesian discriminant method (section 4.5). We settled on the Bayesian discriminant method because it allows us to better distinguish between bowel sounds and non-bowel sounds. Most of the results in this paper are derived from the Bayesian discriminant method.

4.1 A Note on Bowel Sound Detection Rates

We learned what a bowel sound looked like by observing and listening to my own bowel sounds and the bowel sounds of the experimental subjects. We used this knowledge of bowel sounds to develop our bowel sound classification model. However, it is possible that my intuition of what a bowel sound looks like is incorrect. This uncertainty makes it difficult to calculate a number for the percentage of actual bowel sounds detected. However, we are able to calculate a classification accuracy using our classification methods. These results are presented in section 4.6.




Figure 4.2: Bowel Sound 1




Figure 4.3: Bowel Sound 2




Figure 4.4: Bowel Sound 3




Figure 4.5: Bowel Sound 4

4.2 Sensor Data

We used both a piezoelectric sensor and an electronic stethoscope to record abdominal sounds. Both the piezoelectric sensor and the stethoscope were in direct contact and attached via an adhesive to the skin on a test subject’s abdomen.

4.2.1 Piezo Transducer Setup

In our first experimental setup we used a piezoelectric film sensor to record bowel sounds. The film sensor was attached to the body using medical tape. The piezo sensor we used was unshielded and prone to electromagnetic interference from the environment. Instead of developing a sensor circuit to counter this, we opted to use a stethoscope and focus our attention on bowel sound and eating detection. In the future we hope to integrate a piezoelectric sensor into our system.

4.2.2 Stethoscope Setup

Today, physicians and nurses use stethoscopes to determine the presence of bowel sounds following surgery and during routine physicals. Figure 4.6 shows our experimental setup using a stethoscope. After the stethoscope detects and outputs the sound data, and the data is recorded, it is passed to a bowel sound detection algorithm.

Our intuition from prior work ([Campbell et al., 1989](#); [Craine et al., 1999](#)) is that we can use bowel sounds to detect when a patient is eating and use that information to infer system activity that could help in a forensic investigation. In an effort to detect eating instances, we recruited test subjects to eat a meal while we monitored their intestinal activity. We obtained IRB approval for our tests and all volunteers gave their signed consent. For the experimental equipment, we used a Thinklabs ds32a+ stethoscope to record bowel sounds ([Thinklabs, 2013](#)).




Figure 4.6: The stethoscope head is attached to the body with medical tape. The stethoscope head is held in place with an a medical bandage wrapped around the body.

The stethoscope head was set to diaphragm mode (a mode used to detect heart and bowel sounds) and placed in the right lower quadrant of the abdomen. The stethoscope head was secured to the abdomen by medical tape and pressure was applied to the stethoscope head via a medical bandage wrapped around the abdomen (see Figure 4.6). In order to encourage subjects to decrease the amount of activity that might cause extra noise, we had subjects sit at a table. The subjects could read or perform other activities that did not include standing or a large amount of movement. Five subjects participated in the experiment.

Using a PC, we recorded sounds from the stethoscope’s 2.5 mm output jack to a 3.5 mm input jack. With Audacity software ([Audacity, 2013](#)), we recorded the data at 22050 samples/second in a 16-bit PCM encoding. Through the use of National Instruments LabView ([National Instruments, 2013](#)) and MathWorks Matlab ([MathWorks, 2013](#)), we implemented prototype bowel sound detection algorithms with the 16-bit recorded audio signals. The stethoscope did not experience noticeable electromagnetic interference from the environment.

4.3 Integration Approach




Figure 4.7: Bowel Sound Detection - Integration Approach: The integration method bowel sound detection component is composed of two subcomponents. This is a more detailed figure of the bowel sound detection component in Figure 3.1.

Our first approach in eating detection was influenced by Campbell (Campbell et al., 1989). In this method, we obtained a signal from a piezo film sensor as detailed in section 4.2.1. We took the absolute value of the raw signal over a five minute period. A low-pass filter with a cutoff frequency of 20 Hz or 25 Hz was then applied to this signal. This signal was then integrated over the five minute period. Based on Campbell’s work we hypothesized that the single value resulting from the integration over a five minute period should increase after a person has consumed a meal (Campbell et al., 1989). We abandoned this approach because of the sensor

problems discussed in section 4.2.1 and because our integration method was unable to distinguish between different types of noise (see section 4.8 for a discussion of noise). The integration method was implemented in LabView.

4.4 Filter and Threshold Method




Figure 4.8: Bowel Sound Detection - Filter and Threshold Approach: The threshold and filter bowel sound detection component is composed of three subcomponents. This is a more detailed figure of the bowel sound detection component in Figure 3.1.

To account for the electromagnetic interference from the environment we began to record bowel sounds with an electronic stethoscope as detailed in section 4.2.2. In order to eliminate noises that are not bowel sounds we implemented a filtering and thresholding technique. In this technique we took the absolute value of the raw

data and filtered it using a 25 Hz low-pass filter. We then used thresholding in two dimensions. For a sound to be considered a bowel sound it had to rise above a vertical threshold, as well as have a width shorter than a horizontal threshold. This method was also implemented in LabView.

A bowel sound is shown in Figure 4.9. The algorithm uses a vertical threshold to eliminate noise, and it uses a horizontal threshold to eliminate long sounds (e.g., speaking or environmental noises with a large amplitude). Figure 4.10 shows a rejected sound that exceeded the horizontal threshold. This method was an improvement over the previous integration method, but it could still result in false negatives. Figure 4.11 shows a case where a potential bowel sound does not cross the vertical threshold.




Figure 4.9: A detected bowel sound. This sound rose above the vertical threshold and was within the horizontal threshold.

The algorithm calculates the horizontal threshold by determining the mean and standard deviation of the amplitude of the baseline fasting signal. The horizontal threshold was calculated by taking three standard deviations beyond the mean of the width of baseline fasting sounds at the horizontal threshold.




Figure 4.10: A rejected sound. This sound rose above the vertical threshold but was longer than the horizontal threshold.




Figure 4.11: A potential false negative. This sound was within the horizontal threshold but fell below the vertical threshold. It was therefore not considered a bowel sound.

The filter and threshold method was an improvement over the integration method in that it was able to account for some false positives. In the filter and threshold method, the false positives are mainly eliminated by simple thresholds. Thresholds alone cannot eliminate the number of false positives in a patient's environment. We thus implemented an improved method that can incorporate specific features of a bowel sound to distinguish a bowel sound from other non-bowel sounds.

4.5 Bayesian Discriminant Function Method




Figure 4.12: Bowel Sound Detection - Bayesian Approach: The Bayesian discriminant bowel sound detection component is composed of four subcomponents. This is a more detailed figure of the bowel sound detection component in Figure 3.1.

To address the false positives in bowel sound detection we implemented a Bayesian statistics method that is broken into several parts. The first part consists of windowing the raw data signal to focus on sounds that may be bowel sounds. The second part consists of extracting features from the windowed signal. Finally, the algorithm uses the extracted features to classify the windowed sound as a bowel sound or other type of sound using a Bayesian training model. This method was implemented in Matlab.

4.5.1 Windowing




Figure 4.13: The circle represents the detected peak and the blue arrows represent the amount of time taken before and after the peak.

Windowing is the process of finding the segment of the waveform over which to extract features. The signals that we analyze are recorded using a WAVE format and normalized to an amplitude of negative one to one. To find a sound in the signal, we find peaks above a certain amplitude threshold. All other peaks are discarded as

We first extract features from the raw signal by taking the mean, variance, energy, skewness, and kurtosis of the raw signal. This results in five features. We extract the same five features from the FFT of the raw signal. We then pass the raw signal through the filter bank shown in figure 4.14. The output of each high-pass filter is a vector. The vector output of the high-pass filters are labeled detail coefficients. From each vector of detail coefficients, we extract the five features: mean, variance, energy, skewness, and kurtosis. We extract these values from five levels of the filter bank for a total of 35 features.

Mean:

$$v = \frac{1}{n} \sum_{i=1}^n x_i \quad (4.1)$$

Variance:

$$v = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (4.2)$$

Energy:

$$e = \sum_{i=1}^n |x_i|^2 \quad (4.3)$$

Kurtosis:

$$k = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^4}{\left(\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2\right)^2} \quad (4.4)$$

Skewness

$$s = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^3}{\left(\sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2}\right)^3} \quad (4.5)$$

4.5.3 Bayesian Bowel Sound Detection

After a sound segment has been windowed and the features have been extracted the features are compared to the training set features using a Bayesian model. The probability that it belongs to the positive case is compared to the probability that it belongs to the negative case. The sound segment is then assigned to the case with the higher probability. The discriminant function for each case is shown below.

$$g_i(x) = -\frac{1}{2}(x - \bar{x}_i)^T \varepsilon_i^{-1}(x - \bar{x}_i) - \frac{1}{2} \ln |\varepsilon_i| + \ln P(\omega_i) \quad (4.6)$$

Where:

$$\bar{x}_i \quad (4.7)$$

is the mean vector for each case (positive or negative)

$$\varepsilon_i \quad (4.8)$$

is the covariance matrix of each case, and

$$P(\omega_i) \quad (4.9)$$

is the prior probability of each case.

We assume the same prior probability in the above equation to simplify our classifier method. In order to reduce the number of dimensions we project our data onto a one dimensional matrix using Fisher's linear discriminant ([Purdue, 2008](#)).

Fisher's linear discriminant:

$$w = ((n_{positive} - 1)\varepsilon_{positive} + (n_{negative} - 1)\varepsilon_{negative})^{-1} * (\bar{m}_{positive} - \bar{m}_{negative}) \quad (4.10)$$

Where:

$$n_{positive} \quad (4.11)$$

is the number of positive bowel sounds in the training set and

$$n_{negative} \quad (4.12)$$

is the number of negative bowel sounds in the training set.

This allows us to take our thirty-five dimension classification algorithm and reduce it to only one dimension.

4.6 Bowel Sound Training and Accuracy

In order to train the Bayesian classifier, sounds were manually identified as bowel sounds or not bowel sounds. Sounds were collected until 500 bowel sounds were detected and 500 non-bowel sounds were detected for a total of 1000 sounds. The sounds were gathered from the fasting and/or the start of eating from each test subject. All 1000 sounds were used for the Bayesian classifier discussed in Chapter 5.

To gain insight into the accuracy of our classification method and the sounds we use for training we plot the accuracy of our method over various sized training and testing sets. The sounds for the training sets are randomly selected from the 1000 sounds. The remaining sounds not used for the training set are used as the testing set. Test set sizes range from 50 sounds to 950 sounds in increments of 50 sounds. The accuracy of the training set for the testing set and training set are shown for each of the five subjects in Figures 4.15, 4.16, 4.17, 4.18, 4.19. To account for the randomness in our testing sound selection we performed this measurement 100 times for each subject and averaged the results.




Figure 4.15: Classification accuracy measurements for subject 1.




Figure 4.16: Classification accuracy measurements for subject 2.




Figure 4.17: Classification accuracy measurements for subject 3.

These figures give insight into the sounds used to create our model and the accuracy of our model. The faster the curve approaches the asymptote, the better the training data. A slower approach to the asymptote indicates a larger amount of




Figure 4.18: Classification accuracy measurements for subject 4.




Figure 4.19: Classification accuracy measurements for subject 5.

data is needed for the classification model. The height of the asymptote determines the accuracy of the classifier.

The training model results can be broken into 3 categories using their approach to the asymptote. First, the curve from subject one has accuracy data that quickly

approaches the asymptote. Second, the accuracies from subjects two and three approach the asymptote more slowly but reach the asymptote by approximately 500 training sound samples. Third, the data from subjects four and five approach the asymptote slowly indicating that more data samples are needed to create an accurate model.

The testing set accuracy for each subject with 950 sounds in the training set and 50 sounds in the testing set is: subject 1) 86.42%, subject 2) 85.26%, subject 3) 87.36%, subject 4) 76.38%, subject 5) 85.86%. Averaging this across all five subjects gives an 84.26% classification accuracy. It is important to note that this accuracy was obtained from 1000 samples taken from a silent fasting and silent eating state that may not represent every sound the sensor may detect (e.g., talking, vibrating cell phone).

4.7 Receiver Operating Characteristic




Figure 4.20: Receiver operating characteristic for each test subject.

Because our classification scheme returns a discrete value (i.e., bowel sound or not bowel sound) the receiver operating characteristic is composed of only one data

point. Figure 4.20 shows the receiver operating characteristic of the Bayesian method discussed above for each test subject. These curves were obtained from the 1000 sounds used to train the bowel sound classifier. The sounds were randomly divided into 900 sounds for a training set and 100 sounds for a testing set. This gave a single value for the false positive rate and the true positive rate. This was performed 100 times for each subject. The mean for each subject is shown in figure 4.20.

As demonstrated in section 4.6, Figure 4.20 shows the classification method performs better for certain subjects. The true positive rate varies from 88.35% to 95.57% and the false positive rate varies from 14.53% to 22.12%.

The differences in classification accuracies and the differences in the receiver operating curves motivates a personalized classification model that is tuned to each patient.

4.8 Noise

In our testing, we limited patient motion to decrease environmental noise by having the patients sit throughout most of the abdominal sound recording experiments. To be deployed as a sensor in an electronic insulin system, we need to address potential false positives that are caused by patient activity. To evaluate the potential issues of environmental activities that could cause false positives, we recorded certain activities (the sound passed through the test subject’s abdomen) that could cause false positives: walking, coughing, a vibrating cell phone, and talking.

We found that all four activities have the potential to cause false positives. Fortunately, there was a distinct difference in the active frequency ranges of the four activities in comparison to eating. Walking produced a periodic signal dependent on the test subjects gait. Coughing produced signals with dominant frequencies ranging from 0-50 Hz. Vibrating cell phones had strong frequencies of approximately 200 Hz. Bowel sounds had dominant frequencies in the approximate range of 75-125 Hz.

Talking did not have as much of a distinguishing characteristic to a bowel sound. We performed more analysis to address dealing with talking as a potential false positive.

4.8.1 Talking

Talking is aperiodic and in a similar frequency range as that of bowel sounds. Figure 4.21 shows the average power spectral density (PSD) of talking (i.e., reading aloud) for all five study participants over a five-minute period. We found that talking centered around two approximate frequencies of 100 Hz and 190 Hz. The power of talking frequencies drown out the dominant frequencies created by bowels sounds (See Figure 4.22).




Figure 4.21: Average power spectral density of talking for all five subjects. Subjects were asked to read continuously for a five minute period




Figure 4.22: Average power spectral density (over five minute period) of bowel sounds for all five subjects 20 minutes after start of eating. Subjects were asked to limit movement during this time to better record bowel activity.

If we can determine that a patient is talking, then a detected talking event could be eliminated as a false positive and labeled as a non-bowel sound. This could potentially introduce false negatives as talking and a bowel sound could occur at the same time. Figure 4.22 is the average five-minute PSD of all five subjects 20 minutes after the start of eating. During this time interval, the subject is finished eating and the detected number of bowel sounds should have reached its peak. This information should allow us to determine when a patient is talking.

This chapter detailed the method by which we recorded bowel sounds. It then explained the three bowel sound detection methods we used: the integration approach, the filter and threshold method, and the Bayesian discriminant function method. It then gives an analysis of the Bayesian discriminant function method. Finally, noise is

discussed in relation to bowel sound detection. The use of the detected bowel sounds is described in the next chapter.

Chapter 5

Eating Detection




Figure 5.1: System Design: This chapter addresses the eating detection component.




Figure 5.2: Eating Detection: Eating detection is composed of 3 subcomponents. This is a more detailed figure of the eating detection component in Figure 3.1.

In order to determine when a patient is consuming food we record the number of bowel sounds over a five minute period. If the number of bowel sounds over this five minute period is greater than a given threshold we assume the patient is eating. In the next two section we detail the results of our eating detection scheme using the filter and threshold bowel sound data (5.1) and the results of our eating detection scheme using the Bayesian discriminant function bowel sound data (5.2). Because we abandoned our integration bowel sound approach we do not have any applicable bowel sound data with which to apply our eating detection scheme. We do not address talking in relation to eating detection here. However, we realize that talking in relation to eating must be addressed and we recommend this as an opportunity for future work (see section 7.1.1). We recorded talking data before eating, and for some

subjects, while eating. This data could be used in future work and is omitted in the following results.

5.1 Filter and Threshold Eating Detection




Figure 5.3: Filter and threshold eating detection results. The vertical dashed line represents the start of eating.

In Figure 5.3 eating data has been time shifted so that eating begins at the same time for each subject. The amount of detected bowel sounds vary for each subject. However, for each subject there is an increase in the amount of detected bowel sounds. The smallest increase in the number of detected bowel sounds occurs for subject 3. The max baseline (at 0:19 minutes) fasting signal for subject 3 records 106 bowel sounds in the five minute period. At the start of eating (at 0:51 minutes) the number of bowel sounds increases to 156 bowel sounds. This is an increase by a factor of 1.47 from the max baseline fasting signal. The number of bowel sounds detected in the next five minutes (at 0:56 minutes) is 164 bowels sounds and is an increase by a factor of 1.55 from the max baseline fasting signal. From this we infer a minimum increase in the number of bowel sounds from the baseline fasting signal to the eating signal by a factor of 1.5 to record an instance of eating.

5.2 Bayesian Discriminant Function Eating Detection




Figure 5.4: Bayesian discriminant function method eating detection results. The square represents a value of 2503 bowel sounds detected in the 5 minute period. It is omitted so that the figure could be on the same scale as Figure 5.3. The vertical dashed line represents the start of eating.

In Figure 5.4, we see that the smallest increase in the number of bowel sounds again occurs for subject 3. The max baseline fasting number of bowel sounds detected in a five minute period (at 0:19 minutes) is 139 bowel sounds. This increases to 252 bowel sounds detected at the start of eating (at 0:51 minutes) an increase by a factor of 1.819. In the next five minutes (at 0:56 minutes) the number of bowel sounds rises to 320 for an increase by a factor of 2.30. From this we can infer that a minimum increase by a factor of 2 in the amount of detected bowel sounds from the baseline fasting signal to the eating signal is required to detect a meal.

Eating detection using bowel sound data from the Bayesian discriminant function gives a slightly higher increase in the amount of detected bowel sounds at the start of eating than the bowel sound data from the filter and threshold method. We recommend the Bayesian discriminant function method in future work.

5.3 Eating Detection and Errors in Bowel Sound Detection

We found that the position of the body had an effect on the recording of bowel sounds. This is best illustrated by subject 4. The number of bowel sounds peaks at 1 hour and 12 minutes. It was during this time that the subject placed his hands on his chest. This changed the position of the sensor. This caused an increase in the amount of bowel sound detected. In Figure 5.4, this spike is represented by a red square so that the graph maintains the same scale as Figure 5.3. Note that the spike is also present in Figure 5.3.

The reason the bowel sound detection scheme detected such a massive increase in the number of bowel sounds when the subject changes position is that the bowel sound detection scheme was trained on data when the subject is in one position and is unable to account for data recorded when the subject and/or the sensor changes position. This is an opportunity for future work (see section 7.1).

We have realized that despite having false positives and false negatives in our bowel sound detection schemes we are still able to detect eating. When eating, the number of bowel sounds detected must increase to a level that is at least 1.5 greater than the baseline fasting number of bowel sounds (for the filter and threshold bowel sound detection method) or by a factor of 2 (for the Bayesian discriminant function method). If the algorithm fails to detect some bowel sounds yet is still able to detect the relative increase in the number of bowel sounds, eating can still be detected, and the forensic inference of eating will still be correct.

As discussed above we do not address talking in relation to eating detection in this work but we do note a method by which we may be able to still detect eating in the presence of talking. We assume that there are times when a patient will not be talking (i.e., the patient must breathe). A person that is in a conversation will need to pause to listen, breathe, or she will eventually need to take a bite or swallow (from eating or drinking). We have the opportunity to detect bowel sounds during these

breaks from talking. As long as we can detect an increase in the number of bowel sounds, even if we cannot detect every bowel sound, we anticipate we can still detect an instance of eating. More work is needed on this subject (see section 7.1.1).

5.4 Eating Detection Rate

From the data presented in Figures 5.3 and 5.4 we claim 100% success in detecting the start of a meal. It should be noted that we can only detect the start of a meal from a fasting state. We do not have enough data to determine the end of an eating instance or an eating instance when the subject does not start in a fasting state. More data should be gathered to confirm our findings and further understand the relation of a meal to bowel sounds. This is again an opportunity for future work (see section 7.2).

5.5 Meal Indicating Factors

We believe our eating detection scheme could be improved by extra contextual information obtained from a patient's cell phone and/or her insulin therapy system.

5.5.1 Cell Phone Factors

Many cell phones have GPS capability. This could be used to add contextual data to the eating detection scheme. For instance, if the patient is at a fast food restaurant she is probably about to consume a meal or is in the process of consuming a meal. Other sensor information from the cell phone could also be used. For instance, if a cell phone sensor is able to detect when a patient is sitting, this could add useful information to eating detection.

5.5.2 Insulin Therapy System Data

In some insulin pump systems, the amount of carbohydrates or even the type of food may be entered by the user. This would be useful in eating detection. It may also be possible to record patient interaction with a diabetes therapy system before and during a meal to deduce pre-eating and eating patterns. This type of data could add to our eating detection method. Insulin therapy system data should be used with caution in an eating detection method as we are trying to protect against insulin therapy system misuse. This misuse could affect the insulin therapy system data used in an eating detection method.

This chapter detailed how the bowels sounds detected by the methods in the previous chapter are used to detect eating. It then described eating detection in relation to bowel sound detection errors. It then discussed eating detection rates. Finally, it detailed non-eating factors that could help indicate a meal. The use of the detected instances of eating is described in the next chapter.

Chapter 6

Event Classification and Policy Enforcement

There are two glycemic states that have a negative effect on a diabetes patient. The first is the hyperglycemic state (hyperglycemia) and the second is the hypoglycemic state (hypoglycemia). Although our system can detect both states, we focus on the hypoglycemic state. Patients may be unable to detect when they are experiencing a hypoglycemic state, and a hypoglycemic state can have an immediate negative effect on patients. If we can protect a patient from more severe hypoglycemia (e.g., a situation where the patient may be unable to treat the event without outside help), this is a benefit. Euglycemia is a third non-security state that our system can also detect.

The purpose of this thesis is to demonstrate the uses of eating detection in conjunction with current insulin therapy system data in forensic analysis. By combining the detection of eating with other insulin therapy system data, an algorithm can detect potential hyperglycemia and hypoglycemia states. In some situations it may be obvious that a security event is imminent, and a security mitigation action may be warranted. This chapter details the insulin pump system




Figure 6.1: System Design: This chapter address the security event detection and event log component.

data that the forensic analysis and rules will use and introduces three forensic rules for event classification.

6.1 Eating Model

We use the eating model shown in Figure 6.3 in our analysis of eating and insulin therapy system data. Our eating model is centered around the start and end of a meal. In the best case scenario a patient will check their blood glucose level with a finger stick and administer a pre-meal bolus. In this case the patient is being compliant to best insulin therapy system procedures. Patient compliance is not guaranteed however, and the patient may administer a bolus during the meal or after the meal

Event Classification and Policy Enforcement




Figure 6.2: Event Classification and Policy Enforcement: This component is composed of three subcomponents. This is a more detailed figure of the event classification and policy enforcement component in Figure 3.1.

has occurred. In our eating model there is a pre-determined amount of time before the start of a meal where, instead of countering the carbohydrates from a meal, a bolus can cause negative effects (pre-meal bolus period). In our model there is also a pre-determined amount of time after a meal when an insulin bolus is administered too late to offset negative patient effects caused by the consumption of carbohydrates (post-meal bolus period). The length of these pre-determined times are beyond the scope of this work.




Figure 6.3: Eating Model: The blood glucose check and bolus in the pre-eating period represents the recommended time period a compliant patient will administer an insulin bolus in relation to a meal. The large crooked line represents the start of eating. The small crooked line represents the end of eating. The dashed lines represent correction boluses and/or boluses given by a non-compliant patient.

6.2 Hyperglycemic Patient State

Hyperglycemia occurs when the blood glucose value of a patient rises above its recommended range. This can occur when a patient is consuming foods that contain carbohydrates. An insulin bolus can lower a patient's blood glucose value.

Hyperglycemia can also be caused by the suspension of insulin delivery (a common occurrence – this may happen when bathing), and it can also be caused by a malfunction in the physical components of the insulin therapy system (insertion site canula is clogged or bent, disrupting the flow of insulin). Tissue buildup at the insertion site can also cause hyperglycemia. Our system can only detect hyperglycemia in relation to a meal. The detection of hyperglycemia in other cases is an opportunity for future work (see section 7.3).

6.3 Hypoglycemic Patient State

Hypoglycemia occurs when the blood glucose value of a patient falls below its recommend range. This can occur when a patient administers an insulin bolus. In many cases an insulin bolus is administered in conjunction with a meal.

Hypoglycemia can also occur from a higher insulin basal rate and from exercise. Because our system can only detect hypoglycemia in relation to a meal, the detection of hypoglycemia in other cases is again an opportunity for future work (see [section 7.3](#)).

6.4 Euglycemic Patient State

Euglycemia is a safe blood glucose level and is the desired state of the patient. This is the third event that our system can detect and record.

6.5 Insulin Therapy System Data

Data from the insulin therapy system is crucial to detect patient events. The most important data from the insulin therapy system are the boluses given in relation to a meal. The relationship between a bolus and a meal are detailed in [sections 6.7.1](#), [6.7.2](#), and [6.7.3](#). Other insulin pump system data may also be useful (e.g., whether a command was issued remotely or by physically interacting with the pump itself, amount of insulin administered, and basal rates).

6.6 Cell Phone

Cell phone data could be integrated into the event detection system. As in [section 5.5.1](#), the GPS data of the cell phone could be used to interpret patient events. For example, a insulin bolus like that given after a meal high in carbohydrates would be

unlikely at a gym. Also, as in section 5.5.1, sensor data from the phone may also be able to yield useful information.

6.7 Forensic Rules

We use the patient eating data and data from the insulin therapy system to develop three forensic rules to record negative patient events. These rules are potential policies that could help a forensic investigator better understand how to evaluate a potential security (or safety) event. Each rule is based on the expected behavior of a patient. This behavior is shown in Figure 6.3. We provide these forensic rules as part of the policies of our system.

6.7.1 Normal Food Bolus Event

A normal food bolus is the benign use case scenario. The patient has consumed a meal in the presence of an insulin bolus. The insulin bolus does not necessarily have to occur before the meal as long as it occurs within the pre-meal and post-meal bolus periods.

Knowing that a patient administered a bolus and consumed a corresponding meal is important information because a forensic investigator could ignore this normal event when looking for security events. This information, coupled with a patient's historical patient data, could allow the forensic investigator to construct expected normal behaviors of the patient.

6.7.2 No Bolus Forensic Event

If food consumption is detected and no bolus is detected within the pre-meal or post-meal bolus period, then there is an increased probability of hyperglycemia.

There are several possibilities for a patient not to issue a bolus before eating (the best insulin therapy system use case). A patient could have forgotten to administer a

Blood Glucose Check




Figure 6.4: Normal eating scenario: The patient has given checked her blood glucose value and administered an insulin bolus before the start of the of the meal.




Figure 6.5: No Bolus Eating Scenario: The meal begins and ends but no bolus is recorded.

pre-meal (or pre-snack) insulin bolus, or she could have developed a poor habit where she may issue an insulin bolus after a meal. For a carbohydrate free food (e.g., nuts), a patient may eat without needing to bolus.

All of these situations can be dealt with similarly. To encourage patient compliance, the system should mark any event when food consumption is detected without a bolus administered in the pre-meal bolus period. By logging events where the patient forgets to issue an insulin bolus before a meal, the long-term effect will be a lower patient HbA1c value (a patient's average blood glucose value over the past three months), because blood glucose is better controlled. This policy could potentially radically improve patient health. Even a one percentage drop in HbA1c will reduce microvascular complications by 40% [CDC11]. Recording this event as a forgotten bolus will increase the number of incorrectly recorded negative patient events for those who consume a large amount of food that does not require an insulin bolus. We do not expect this case to be common across all patients. If eating foods that do not require a bolus is common in a patient, the pump can ask the patient when this situation is detected (e.g., a button press). We intend to address this issue in the future (see sections [7.3](#) and [7.6](#)).

Some patients may desire to give an insulin bolus while eating. For instance, a patient may wait to issue a bolus at a restaurant, because the food arrival time and portion size may not be known a priori. One way to address this issue is to use the patient's location (e.g., a GPS sensor) that could provide additional context about a patient's activity. If the GPS sensor indicates that a patient is at a location where they typically consume food, then this increases the probability that a patient is consuming food. Although we expect these instances to be uncommon, by using this contextual information, we can implement a system that could adapt to a patient's behavior to stop its alerts when a patient would want to suppress boluses before a meal. Wireless interference could also cause a deviation from the expected behavior. If one unintentionally or intentionally jammed a wireless bolus command, then it could block a patient from successfully issuing an insulin bolus. The result would be a recorded hyperglycemic event. Given that the hyperglycemia would occur at food consumption, a physician could better direct a patient to focus on issuing insulin correctly. If a patient were to issue insulin boluses and still observe the same results

from wireless interference (or jamming behavior), this could be noted. The patient could examine the data, take additional actions, or seek help.

6.7.3 No Food with Bolus Forensic Event

If an amount of insulin is given and there is no meal within an appropriate amount of time then there is increased risk for hypoglycemia and a potential security breach. The amount of insulin may be different for each patient.

Blood Glucose Check



Figure 6.6: No Meal Bolus Scenario: An insulin bolus is recorded but no meal is detected

If a pre-meal bolus is given without the detection of food consumption within the pre-meal bolus period, the probability of hypoglycemia is increased. This could potentially occur when the patient administers a correction bolus - a bolus given to decrease the current blood glucose level (typically after a blood glucose check). Food consumption is not always expected with a correction bolus. This would result in an incorrectly recorded negative patient event. To distinguish between a correction bolus and a pre-eating bolus, the patient could record the type of bolus (assuming that the patient is compliant). This would of course require a significant operational change on the part of the patient - something that our system avoids. This issue must be addressed in future work (see section [7.6](#)).

6.8 Classification of Safety and Security Events

Our system can identify the negative patient events discussed above. However, a method to classify a negative patient event as a safety or security event is more difficult. Some cases such as an insulin bolus given when the patient is typically asleep and in the absence of a meal may be obvious as a security event, but it is difficult to create a use model for every scenario. More work must be done to develop a classification model for safety and security events. Despite this, we believe our system offers enough information for a forensic investigator to determine the nature of a negative patient event.

This chapter detailed the eating model which we use to determine negative patient events. It then described three glycemic states and their relation to three forensic rules we developed to determine negative patient events. It also discussed the role the insulin therapy system or a cell phone could play in negative patient event detection. Finally it discussed safety and security classification. Possible improvements to our experiment and direction for future work are described in the next chapter.

Chapter 7

Future Work

There is still much work to accomplish before this system can be deployed in an insulin therapy system. First, the bowel sound detection methods must be able to account for environmental noise (e.g., talking, sensor noise). Second, more research is necessary to understand what may constitute an eating false positive or negative (i.e., whether a meal from a non-fasting state will be detected). Third, further work must be performed to develop forensic rules and models for negative event classification (i.e., whether a safety or security event). Fourth, the system must be integrated into a mobile sensor. Fifth, the system must be implemented in a real time architecture that is low-powered, computationally efficient, and portable. Finally, system implementations requiring an operation change must be examined.

7.1 Bowel Sound Detection Improvements

Our current Bayesian discriminant method is able to determine whether a given sound is a bowel sound or is not a bowel sound in a laboratory setting. In [section 4.8](#) we detail possible methods by which we may be able to account for noises encountered in a non-laboratory setting. However, these methods must be tested. This will require further experimentation and data gathering from test subjects.

The average bowel sound classification rate from our training sets using a Bayesian discriminant function is 84.26%. We believe this could be improved by using a support vector machine classification model. It is also possible that we may be able to improve our model by having a more diverse amount of training data that can account for different activities the patient may perform or different body and sensor positions.

7.1.1 Bowel Sound Detection in the Presence of Talking

More work needs to be performed to build models to detect bowel sounds in the presence of talking. We suggest that it may be possible to detect enough bowel sounds in the pauses when a person is talking to identify an eating instance (section 5.3) but this theory must be tested. It may even be possible to detect bowel sounds in the presence of talking.

7.2 Eating Detection Improvements

Using our current eating detection method we are able to detect eating with 100% accuracy using the Bayesian discriminant function method. This shows the feasibility of our eating detection scheme. However, more testing must be performed to confirm that our eating detection method works for a larger number of test subjects. Our results thus far only show the feasibility of detecting the start of a meal after a fasting state of at least two hours. Ideally we would like to be able to detect both the start and the end of a meal.

Further research must also be performed to determine whether a meal can be detected when the patient does not start in a fasting state. Also, the minimum time between meals that our eating detection method can distinguish between two meals must be determined.

7.3 Forensic Rules and Event Classification

We only present three forensic rules. Our rules are unable to account for cases where eating does not require an insulin bolus and when a given insulin bolus is a correction bolus. Our rules must be improved to account for these cases. The three rules we present are by no means an exhaustive list. Further forensic rules can also be developed.

We feel that the information our system provides does a good job of providing forensic information for a forensic investigator. A model could even be made for simple cases to distinguish between a safety and security events. However, it would be difficult to account for every possible safety and security case. More research should be performed to identify the feasibility of a general model to classify negative patient events.

7.4 Sensor Integration

Current insulin pump patients have an insertion site (where insulin is delivered subcutaneously into the body) and the option of a continuous blood glucose monitor. With a continuous glucose monitor, the patient has the additional burden of wearing another device. However, patients derive great benefit from real-time data of blood glucose levels. Because patients are already burdened with a high number of devices, we constrain our sensor integration designs to those that do not increase this device burden. Any additional sensor should be seamlessly integrated (as much as possible). Our plan is to integrate our sensors onto already existing electronic insulin pump components (e.g., the continuous blood glucose monitor or insulin insertion site). To increase patient acceptance and to decrease patient discomfort, the sensor is noninvasive. The envisioned system design is shown in Figure 7.1. A patient who uses our system would not be required to change the way she uses her current infusion set. Usability will be at least equal to the usability of current systems. In fact, a

patient could use a forensically-aware system without knowledge of the forensic data capability. Patient compliance in using a forensic system will be the same as those of normal insulin pump systems, because the patient would not be required to adopt new usage procedures.




Figure 7.1: Our goal is to integrate a sensor with components already used in insulin pump systems such as the infusion set, continuous glucose monitor, or the insulin pump device itself.

7.5 Implementation

The system presented in this thesis is implemented on a computer using recorded data. Eventually we would like to create a portable real-time system. We believe that our Bayesian discriminant bowel sound method does a good job of detecting bowel sounds without being computationally expensive. The tradeoff between computational efficiency and classification accuracy needs to be determined. Different bowel sound detection methods should be compared to this effect. A computationally efficient method would save on device power. Available hardware will also play a role in determining the classification method.

7.6 Future Improvements Requiring Operational Changes

In the methods presented in this work we simply detect potentially damaging glycemic levels and record this information. One of the goals of the system presented in this thesis was to refrain from requiring an operational change. However, we realize that an operational change on the part of the patient could improve patient security. For instance, if multiple negative events are recorded, then the system could raise an alert to the patient or physician. Insulin therapy system operational use could also be changed. For example, when a bolus command is given that could potentially harm the patient, the insulin therapy system could require the user to physically interact (lock out remote pump interaction) with the system to confirm the bolus. The insulin therapy system could even ask the user when it is unsure how to classify an event.

7.7 Comments on Experimental Method

While it was possible to detect the start of a meal using the techniques described in this thesis, there are a number of possible improvements that could be applied to the experimental process used in this thesis.

In order to avoid invading the privacy of the test subject, the person conducting the experiment instructed the subject on how to attach the sensor to the subject's body. This resulted in unwanted variation in the sensor attachment and therefore in the recorded data.

We set out to determine that a person was eating. Test subjects were requested to remain as still and quiet as possible in a sitting position. This was to minimize vibrations from the sensor rubbing against the body when the subject moved and vibrations created when the subject talked.

Test subjects did not remain still and would fidget in their chairs. Test subjects would also occasionally talk. Test subjects were allowed to either read or operate a computer. The amplitude of the recorded sound would also change when the subject shifted positions in the chair. This created unwanted variability in the subject data. This is especially demonstrated in the case of subject three. The spike in the number of bowel sounds recorded at 1 hour and 12 minutes shows where the participant placed his hands on his stomach.

In future experiments, the person conducting the experiments should attach the sensors to every test subject to decrease variability resulting from sensor attachment. The subjects should also be asked to lie down as in Campbell's work (Campbell et al., 1989). This could result in less fidgeting. Subjects should also not be allowed to read or perform any activity that does involve their health (i.e., insulin bolus, medication).

It may also be wise to increase the compensation of the test subjects. In the experiments detailed in this thesis, the compensation for each test subject was the meal they consumed. However, the experiment consumed approximately four hours

of the subject's time. In the future tests, the compensation for this experiment might be reconsidered.

This chapter detailed possible improvements to bowel sound detection, eating detection, and forensic rules and event classification. it also introduced ideas for sensor integration, implementation in hardware, and future improvements that might require an operational change on the part of the patient. Finally, it discussed possible general improvements to the experimental procedure. This work is concluded in the next chapter.

Chapter 8

Conclusions

In this work we demonstrate that we are able to successfully detect the start of a meal from a fasting state with 100% accuracy using bowel sounds in a laboratory environment. We also demonstrate three bowel sound detection methods. Of the three bowel sound detection methods detailed in this paper we found that the Bayesian discriminant method performed the best with a classification accuracy of 84.26% in a laboratory environment.

The success of our method to detect the start of a meal allows us to develop three forensic rules that can be used to identify negative patient events. We believe these rules greatly improve the ability of a forensic investigator to determine what led to a patient event, what caused the patient event, and whether it was a safety, security, or benign event.

8.1 Comparison to other works

To the best of our knowledge our work is the only work that uses bowel sounds to detect an instance of eating. The most accurate bowel sound classification rate we found in other work is 94.84 % found by Chauhan ([Chauhan et al., 2009](#)). This is better than the classification accuracy of 84.26% obtained in this work. To the best

of our knowledge our work is also the only work that uses bowel sounds as a means to forensically identify patient events.

Bibliography

- (2013). Glucowatch automatic glucose biographer. fda device approval and clearances. Available from: <http://www.fda.gov/MedicalDevices/ProductsandMedicalProcedures/DeviceApprovalsandClearances/Recently-ApprovedDevices/ucm089158.htm>. 4
- ADA (2011). Dawn phenomenon. Available from: <http://www.diabetes.org/living-with-diabetes/treatment-and-care/blood-glucose-control/dawn-phenomenon.html>. 10
- Audacity (2013). Audacity is free, open source, cross-platform software for recording and editing sounds. Available from: <http://audacity.sourceforge.net/>. 22
- Barnaby, J. (2011). Threatening vulnerabilities. Proceedings of the Fourth Hacker Halted USA, Miami, Florida. 3
- Benedict, B., Keyes, R., and Sauls, F. C. (2004). The insulin pump as murder weapon: a case report. *The American journal of forensic medicine and pathology*, 25(2):159–160. 3, 5
- Campbell, F., Storey, B., Cullen, P., and Cuschieri, A. (1989). Surface vibration analysis (sva): a new non-invasive monitor of gastrointestinal activity. *Gut*, 30(1):39–45. 15, 21, 23, 62
- CDC (2011). National diabetes fact sheet, 2011. Available from: http://www.cdc.gov/diabetes/pubs/pdf/ndfs_2011.pdf. 1
- Chang, S.-Y., Hu, Y.-C., Anderson, H., Fu, T., and Huang, E. (2012). Body area network security: robust key establishment using human body channel. In

- Proceedings of the USENIX conference on Health Security and Privacy*, pages 5–5. 3
- Chauhan, N., Ravi, V., and Karthik Chandra, D. (2009). Differential evolution trained wavelet neural networks: Application to bankruptcy prediction in banks. *Expert Systems with Applications*, 36(4):7659–7665. 64
- Craine, B. L., Silpa, M., and O’Toole, C. J. (1999). Computerized auscultation applied to irritable bowel syndrome. *Digestive diseases and sciences*, 44(9):1887–1892. 21, 29
- Ford (2013). Ford in-car health and wellness solutions. Available from: http://media.ford.com/images/10031/health_wellness.pdf. 3
- Gollakota, S., Hassanieh, H., Ransford, B., Katabi, D., and Fu, K. (2011). They can hear your heartbeats: non-invasive security for implantable medical devices. In *ACM SIGCOMM Computer Communication Review*, volume 41, pages 2–13. ACM. 4
- Hei, X., Du, X., Lin, S., and Lee, I. (2013). Pipac: Patient infusion pattern based access control scheme for wireless insulin pump system. pages 3030–3038. *Proceedings of the 32nd IEEE International Conference on Computer Communications*. 3
- Henry Jr, N. L., Paul, N. R., and McFarlane, N. (2013). Using bowel sounds to create a forensically-aware insulin pump system. In *Proceedings of the Fourth Workshop on Health Technology*. Usenix. 1
- Klonoff, D. C. (2008). Designing an artificial pancreas system to be compatible with other medical devices. *Journal of diabetes science and technology (Online)*, 2(5):741. 3
- Li, C., Raghunathan, A., and Jha, N. K. (2011). Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system. In *e-Health Networking*

- Applications and Services (Healthcom)*, 2011 13th IEEE International Conference on, pages 150–156. IEEE. 3
- MathWorks (2013). Matlab: The language of technical computing. Available from: <http://www.mathworks.com/products/matlab/>. 22
- Meier, B. (2010). Fda steps up oversight of infusion pumps. *the New York Times*. 5
- National Instruments (2013). Labview system design software. Available from: <http://www.ni.com/labview/>. 22
- Purdue (2008). Fisher linear discriminant. Available from: https://kiwi.ecn.purdue.edu/oldkiwi/index.php/Fisher's_Linear_Discriminant. 31
- Radcliffe, J. (2011). Hacking medical devices for fun and insulin: Breaking the human scada system. In *Black Hat Conference presentation slides*. 3
- Schechter, S. (2010). Security that is meant to be skin deep using ultraviolet micropigmentation to store emergency-access keys for implantable medical devices. 3
- Sorber, J., Shin, M., Peterson, R., Cornelius, C., Mare, S., Prasad, A., Marois, Z., Smithayer, E., and Kotz, D. (2012a). An amulet for trustworthy wearable mhealth. In *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*, page 7. ACM. 4
- Sorber, J. M., Shin, M., Peterson, R., and Kotz, D. (2012b). Plug-n-trust: practical trusted sensing for mhealth. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 309–322. ACM. 4
- Thinklabs (2013). The ds32a+ stethoscope. Available from: <http://www.thinklabsmedical.com/electronic-stethoscope.html>. 21

Vita

Nathan Lee Henry Jr. was born in Hendersonville, Tennessee, to Nathan and Melissa Henry. He is the second of six children: Leah, Josie, John, Rose, and Andrew. When he was nine he moved with his family to Vietnam. He was homeschooled throughout high school and returned to the United States to attend the University of Tennessee, Knoxville. Nathan obtained a Bachelors of Science in 2012 in Computer Engineering. He continued into a graduate research assistanceship at the University of Tennessee in Computer Engineering under Dr. Gregory Peterson. Nathan graduated with a Master of Science degree in Computer Engineering in May 2014.