ABSTRACT

BRUGAROLAS BRUFAU, RITA. Towards Automated Canine Training: Wearable Cyber-Physical Systems for Physiological and Behavioral Monitoring of Dogs. (Under the direction of Dr. Alper Bozkurt.)

Dogs play several important roles in our modern society while working in partnership with humans, such as assisting disabled people, search and rescue, and narcotic and explosive detection. In parallel, man-made devices and instruments continue to improve to carry out such functions, however, the efficiency of working dogs has not been rivaled. Dogs' more precise sensory perception sensitivity and their ability to exchange information with humans make them superior to any synthetic system. This capability of communicating with humans has enabled their training; a process by which animals learn to perform desired tasks when instructed. So far, the efficacy of the communication, therefore the training process, relies on the dog's and human's skills in mutually interpreting each other's behaviors and body language. In addition to demonstrating certain trained behaviors actively, dogs communicate their emotional state to humans through passive demonstration of certain postures underlined by a physiological response. The latter expressed through behavioral signs, such as trembling, restlessness or yawning, and also through certain physiological parameters, such as heart rate, heart rate variability, and respiratory events like panting. In addition, many dogs work beyond the visual range of their handlers. Thus, a mechanism to facilitate feedback and communication from a distance using a wireless system would greatly enhance the effectiveness and welfare of dogs.

Our goal is to improve the efficiency of the communication between dogs and humans by reducing the subjectivity of the interpretation of their body language and physiological condition. We do this by providing human trainers and handlers with sensed objective data related to their behavior and physiological state. In this dissertation, we introduce the concept of a canine body area network (cBAN), which combines multimodal wearable sensing of behavioral and physiological parameters, with a centralized control unit providing bidirectional communication capability and transmission of the sensor information to a remote data aggregator for further analysis. The insights developed from this analysis are shared with

the handler to assist with making decision during the training process and assess the welfare of the dog. Such an infrastructure also enables direct canine-machine interaction.

The proposed cBAN includes inertial measurement units (IMUs) at optimized locations for dog posture detection. Furthermore, electrocardiogram (ECG), photoplethysmogram (PPG) and e-stethoscopes enable the monitoring by the handler of canine heart rate, heart rate variability and respiratory patterns (sniffing and panting). To overcome the limitations imposed on the ECG recordings by the efficiently insulated skin and dense hair layers of dogs, we introduced novel electrode configurations with metal electrodes coated with a conductive polymer to allow for signal recordings without shaving the dog's hair. We studied the incorporation of light guides and optical fibers for efficient optical coupling of PPG sensors to the skin. We also demonstrated the feasibility of continuous auscultation for the detection of sniffing and panting by identifying discrete features in recordings from the chest and neck using a commercially available stethoscope.

As an application of the cBAN, this dissertation establishes the first results towards automated canine training. Posture information was classified in real time based on inertial measurements and a computer-controlled reward dispenser was used to reinforce every occurrence of the desired behavior. We evaluated the trade-off between low latency and accuracy on the reward delivery and empirically showed that dogs successfully learned from computer-delivered rewards that had latencies less than 0.3 s. It remains as the future work, to leverage the physiological feedback in the training process and further optimize the human-computer interface. The use of the cBAN for computer-assisted training can potentially speed up learning and reduce the cost of training.

Towards Automated Canine Training: Wearable Cyber-Physical Systems for Physiological and Behavioral Monitoring of Dogs.

by
Rita Brugarolas Brufau

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Electrical Engineering

Raleigh, North Carolina

2016

APPROVED BY:

_____	_____
Dr. Alper Bozkurt Chair of Advisory Committee	Dr. David L. Roberts
_____	_____
Dr. Barbara L. Sherman	Dr. Veena Misra
_____	_____
Dr. Troy H. Nagle	Dr. Omer Oralkan

DEDICATION

To my family

BIOGRAPHY

Rita Brugarolas Brufau received a B.Sc. degree in Telecommunications Engineering in 2011 from Polytechnic University of Valencia in Spain. Later that year she began her PhD at North Carolina State University, and joined the Integrated Bionic MicroSystems Laboratory (iBionicS Lab) under the direction of Dr. Alper Bozkurt. Her research at iBionicS included the development of wearable technologies for physiological and behavioral monitoring of canines to enhance human canine interactions. Also, her interest in medical instrumentation led to her developing an auto-adjusting mandibular repositioning device for the treatment of Obstructive Sleep Apnea. Rita is a recipient of a graduate fellowship from Fundación Caja Madrid and the 2015 American Academy of Dental Sleep Medicine graduate student research award.

ACKNOWLEDGMENTS

I would like to thank in the first place my advisor, Dr. Alper Bozkurt. Five years ago you gave me the opportunity to join you from overseas to contribute in the launching of the integrated Bionic MicroSystems laboratory (iBionics Lab). As the first PhD student to graduate from the lab, I can say that building from the ground up was difficult at times, but as most new projects, it also came with unique opportunities for innovation and brought out the greatest enthusiasm among the team members, making it a fun and rewarding experience. I feel very fortunate and honored that you trusted me in this important task. I am also thankful that you took personally the effort of training me not only in technical matters but also in the ideation process. I admire your passion and creativity when approaching new projects and challenges, and I hope that I will be able to apply what I have learned in my next research endeavors. I also appreciate your unconditional support for my ideas and research interests along my PhD journey.

I would also like to thank the rest of my advisory committee for their helpful guidance and support: Dr. Barbara Sherman, Dr. Veena Misra, Dr. David Roberts, Dr. Omer Oralkan, and Dr. Troy Nagle. I feel immensely fortunate of having been able to work in this interdisciplinary research project with Dr. Barbara Sherman and Dr. David Roberts, two exemplary academic roles that have merged their passion with their work. You are great mentors and I have learned a lot through our frequent research discussions and by working closely with you and your students.

This research has been possible thanks to the generous participation of volunteers that brought their dogs to participate in our experiments and the funding from National Science Foundation (NSF) under the award 1329738 and the NSF Engineering Research Center for Advanced Self-Powered Systems of Integrated Sensors and Technologies (ASSIST). In addition I would also like to thank the ASSIST center for providing the students with resources and translational opportunities to complement our training, and Fundación Caja Madrid for their graduate fellowship that allowed me to pursue my research interest.

I would also like to thank each and everyone of my colleagues and friends at the iBionics Lab, Ciigar Lab, the ASSIST center, and across the different departments at NC State, especially Electrical and Computer Engineering, Biomedical Engineering, Chemical Engineering, and Computer Science. Our collaborations have been among the most enriching experiences I have had in grad school and I hope we will continue to have more opportunities in the future.

Last but not least, I want to thank my dearest love and PhD ally Jose Manuel. We embarked in this adventure together and I could not have asked for a better companion. You have always been there for me and I am looking forward to see what the future is holding for us. I also want to thank my parents Concepción and Antonio for your example, support and encouragement since day one, and my siblings: Jimmy, Paul, Claudia, María Javier, Pedro, Teresa, Miguel, and Juan for your advice and encouragement through the journey.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1. Introduction

Numerous advances have been made in robotic sensing including the demonstration of systems mimicking canine locomotion [1] and olfaction [2]. However, the efficient sampling, agile mobility, and odor discrimination acuity of dogs have not been met yet [3]. Dogs are able to learn to perform a range of remarkable tasks thanks to their intrinsic cognitive capacity, high sensory perception sensitivity, and survival instinct. When working for humans, we benefit from all these capabilities without the development, powering and control challenges faced when designing complex synthetic systems.

Dogs play several important roles in our modern day society from assisting disabled people to search and rescue, from narcotic and improvised explosives device (IED) detection to companionship. This partnership has been made possible through dogs' and humans' abilities to communicate signals to each other. This communication enables dog training, a process by which animals learn to perform desired specific tasks when instructed.

Most efficacious training techniques are those based on positive reinforcement [4], which involves giving a primary reinforcer such as food, to reward a sought-after behavior. The timing in delivering the reinforcer [5] and accuracy of feedback are extremely important to achieve effective training [6]. Reinforcing multiple behaviors, failing to recognize correct behaviors or reinforcing the wrong behavior can slow down the learning process or inadvertently teach the dog to do an undesirable action.

One limiting step that leads to ineffective training, therefore limiting the extent to which we can benefit from a dog's performance, is the subjectivity of the interpretation by human handlers of the passive or active signals communicated by the dog. Canines communicate to humans with trained behaviors that were taught to them by human trainers; and through untrained intrinsic and instinctual responses originating from their emotional state (e.g. displaying excitement, anxiety through tailwag...).

Traditionally, the effectiveness of the communication has relied on the skill of the humans in interpreting dog's body language by monitoring postures, breathing patterns, vocalizations and behavioral signs of excitement or anxiety. Considering that handlers also need to be aware of other environmental stimuli at all times and consider the safety when working in potentially hazardous environments, this could be a challenging task. Furthermore, potential physical limitations in the communication channel such as being out of sight or earshot could hinder the effectiveness of the translation of information further.

In an attempt to address the subjectiveness of communication between handlers and dogs, we designed two novel wearable cyber-physical systems: a) a posture detection system based on inertial measurement units (IMUs) and b) a heart rate and respiratory pattern monitoring system based on electrocardiogram (ECG), photoplethysmogram (PPG) and stethoscopes. These systems provide quantified data from the dog during trained and untrained communications respectively that is not available during traditional human-canine interactions. IMUs, when placed at optimized locations and combined with classification algorithms enable the detection of canine postures and behaviors. The physiological parameters obtained from the second system enable more accurate interpretation of canine emotional response [7]–[10]. In addition, we have explored the use of actuators to provide haptic and auditory feedback to enable remote bidirectional communication between handlers and dogs when a direct physical channel between them is not available.

Making such additional information available for handlers improves their awareness and helps them make more informed decisions. To demonstrate that such sensor data and properly designed computer systems would help achieve a more effective and objective communication channel with the dog, we have designed a computer-assisted training system that automates the reinforcement of a desired posture. This system comprises wearable IMUs and a base station for processing the data to classify canine posture. When the algorithm detects the sought-after behavior, a treat is delivered from a computer-controlled dispenser. We determined the "sit" behavior as our model behavior for this study which is one of the first behaviors that working and companion animals are trained for traditionally. By making a

computer system to close the loop between the observation of the desired behavior and the reward delivery, we optimized the trade-off between the timing and accuracy of feedback. We characterized the improvement in the learning of the desired behavior by recording the growth in the rate of offering this behavior during the experiment. In our particular experiment, the number of offered sits was an indication of the learning process. The future work includes adding the physiological recordings to the feedback loop to gain more and quantified insight about the welfare and emotional state of the dog.

When developing wearable technologies for canines, it is of the utmost importance to consider the comfort and ergonomics of the dogs. These requirements translates into a limited number of sensor sites available, because most dogs will not tolerate sensors strapped to certain parts of their body especially within the reach of their teeth such as paws, tail or around their snout. It is worth noting that the efforts presented in this thesis tried to follow a participatory design process with veterinary behaviorists. Such design criteria should be followed to ensure that this new space of technologies for canines are providing value for not only humans but also animals.

In the following section, we address the specific objectives of this dissertation and how our contributions are distinguished from existing solutions both in the market and in research field.

1.1 Objectives

1) Development of a novel canine body area network (cBAN) infrastructure consisting of a set of sensors and actuators attached to a commercial-off-the-shelf (COTS) harness.

Achievement: A BeagleBone Black (BBB) running Ubuntu GNU/Linux serves as the centralized control unit on a “smart harness” and provides bidirectional communication capability based on IEEE 802.11in order to collect information from the sensor nodes and to transmit it to a remote computer running behavior classification algorithms. This infrastructure hosts the physiological and behavioral sensing systems introduced in the next objectives, 8 DC vibration motors and a speaker to enable handlers to instruct the dogs even when out of sight

or earshot, and gas sensors, camera and GPS to monitor the environment. A subset of peripherals can be selected depending on the application.

2) Optimization of the number and location of inertial measurement units for accurate behavior recognition while considering comfort and kinetics of the canine.

Achievement: The IMU number and locations were minimized after the identification of independently moving parts in the body during various postures. We designed a cascade of Decision tree classifiers and Hidden Markov Models for the detection of static postures (sitting, standing, lying down, standing on two legs, and eating of the ground) and dynamic activities (walking, climbing stairs and walking down a ramp) based on heuristic features of the accelerometer and gyroscope data achieving average accuracy above 95%. The system was evaluated on ten dogs of various age, sex, size and breed.

3) Overcoming the challenges of integrating a wearable physiological sensing system to the infrastructure for remote continuous monitoring of key parameters.

Achievement: We combined an electrocardiogram (ECG) and photoplethysmogram (PPG) to the smart harness and assessed the limitations of each modality. To overcome the limitations imposed by the efficiently insulated skin and dense hair layers of dogs, we introduced novel electrode configurations that allow for optimal contact and signal recording without the necessity of shaving dog's hair. We evaluated two types of electrodes to improve the tissue-electrode contact: thick-tapered electrode and comb-shaped arrays of thin spring-loaded pins. We demonstrated that surface modification of the metal electrodes by coating with poly (3,4-ethylenedioxythiophene) poly-(styrene-sulfonate) (PEDOT:PSS) conductive polymer can enhance their electrical property by increasing the charge injection capacity and decreasing the impedance of the tissue-electrode interface. Also, we studied the incorporation of light guides and optical fibers for efficient optical coupling of PPG sensors to the skin. Introduction of inertial measurement units (IMUs) enabled the tracking of the motion of the animal and deploy signal processing to eliminate these when needed. The system was evaluated on six dogs of various age, size, breed, and hair density.

4) Building of a system for continuous auscultation of respiratory behavior and identification of discrete features of sniffing and panting in the time domain.

Achievement: The key body sounds were recorded at the neck and chest by means of a commercially available electronic stethoscope. We identified discrete features of sniffing and panting in the time domain and used those to build a decision tree classifier.

5) Automated reinforcement of "sit" behavior in dogs.

Achievement: The wearable IMU-based system and posture detection machine learning algorithms were combined with a computer controlled treat dispenser to provide the primary reinforcer when the behavior is detected. We explored the trade-off between low latency and accuracy on the reward delivery. We empirically showed that dogs failed to learn the desired behavior when high classification accuracies (99% on average) were targeted as these came at the cost of latencies greater than 1.1 seconds. However, dogs successfully learned from computer-delivered rewards when slightly lower accuracies (91% on average) were achieved that caused shorter latencies (less than 0.3 seconds). The system was evaluated on four dogs of various age, size and breed.

1.2 Dissertation Outline

Chapter 2 provides a general description of our canine body area network (cBAN) as a platform that combines multimodal wearable sensing technologies and computational modeling to provide meaningful information on dog behavior, physiology, and the environment to handlers for objective decision support. The cBAN also includes actuators enabling the remote communication between handlers and dogs when a direct physical channel between them is not available. In this chapter, the cBAN is described for the specific application of augmenting the capabilities of search and rescue teams in the field by enabling a new type of Cyber-Enhanced Working dog. However, these capabilities apply to many uses of working dogs beyond search and rescue. This particular example was picked to set the design considerations and features required for the cBAN and in the subsequent chapters the design of each of the subsystems and sensors are described.

Chapter 3 describes a low-power posture detection system based on inertial measurement units (IMUs) that utilizes Bluetooth Low Energy to transfer the data to the computational node that runs the classification algorithms. We show how we optimized inertial sensor sites by considering the kinetics of the canine to identify independently moving locations on the body. The design and performance of the algorithms for the detection of static postures (sitting, standing, lying down, standing on two legs, and eating of the ground) and dynamic activities (walking, climbing stairs and walking down a ramp) is also reported. For this and each of the following chapters a literature review of the state of the art, the experimental setup used during validation and testing of the system, and the algorithms involved in the data processing are provided.

Chapter 4 describes a wearable heart rate sensor system combining electrocardiogram (ECG), photoplethysmogram (PPG), and IMUs. We present our efforts to overcome the limitations imposed by the dense hair layers of dogs which includes the introduction of novel electrode configurations and the incorporation of light guides for efficient optical coupling of PPG sensor to the skin.

Chapter 5 focuses on the subsystem for continuous auscultation of respiratory behavior by recording sounds at the neck and chest using a commercially available electronic stethoscope and the identification of discrete features of sniffing and panting in the time domain.

Chapter 6 describes our efforts towards automated training by using a computer which detects a desired posture from IMU data and trigger the delivery of a primary reinforcer (a food treat) from a computer-controlled disperser upon detection.

Finally, Chapter 7, summarizes the contributions of this dissertation, provides recommendations for future work, and some discussion of the potential broader impact that this work would offer.

Chapter 2. Canine Body Area Network (cBAN)

2.1 Introduction

After a catastrophic man-made or natural disaster impacts an urban environment, tremendous effort is spent locating and reaching people trapped under rubble. Time is of the essence when it comes to finding survivors. Well-trained search and rescue (SAR) dog and handler teams are becoming increasingly critical in the process of locating survivors rapidly; however, as a consequence of the incredible time, skill, and effort required to train SAR dog and handler teams, they are still relatively inaccessible. Working dogs, depending on their duty, have hundreds of hours of training and conditioning, and can cost between $15,000 and $100,000 [11]. Much of this cost is due to the amount of highly-skilled labor and time required for training.

SAR dogs, due to their intrinsic cognitive capacity, visual acuity, auditory range, olfactory capability, and physical abilities, are capable of learning to perform a range of remarkable tasks that any current robotic system would not be able to achieve. There are some robotics efforts trying to mimic dogs' locomotion [1] and their olfactory ability to detect components at very low concentrations [12] but none of these have approached the efficiency and performance of canines. Furthermore, by working with dogs, we gain superior capabilities without the challenges of powering, controlling, and developing a complex system. SAR dogs sometimes work on-leash in close proximity to handlers, and other times they are given autonomy off-leash to range out and explore wider areas or climb rubble piles at greater distances from their handlers. Such remote searches require complex and intelligent coordination between dogs, handlers and first responders. At all times, handlers need to

monitor the safety of the environment and the welfare of their dogs by continuously reading their body language for signs of anxiety, overheating, or any critical health condition. At present, monitoring is dependent on visualization of the dog.

To improve the efficiency of SAR training and to augment the current capabilities and safety of SAR teams in the field, we present techniques and technological platforms to enable a new type of *Cyber-Enhanced Working Dog* (CEWD). CEWDs are enabled by the use of sensors and actuators worn by dogs forming a "*canine body area network*" (cBAN) that provides detailed real-time monitoring of the dog and environment. A remote computational node that incorporates intelligent context-aware sensing algorithms and canine models retrieves and interprets the sensor data from the cBAN on the dog and presents meaningful information on dog behavior, physiology, and the working environment to the handler for a more intelligent decision making. In addition CEWDs benefit from a computer-mediated communication between dogs and their handlers, even when they are out of sight or earshot. This canine machine interface (CMI) would also enable computer-assisted intelligent dog training to speed up learning, reduce the cost of training, and increase the availability of these valuable assets.

Our cBAN consists of a set of sensor and actuator packages attached to a commercial-off-the-shelf (COTS) harness (Figure 2.1). A centralized control unit on this "smart harness" provides bidirectional communication capability, collecting information from the sensor nodes and communicating this to a remote computer running intelligent behavior classification and training algorithms. We have demonstrated the following capabilities that can be summarized and grouped into three categories:

Monitoring the dog: wearable inertial measurement units (IMUs) allow computers to detect and recognize posture and behaviors in real time. Further, electrocardiogram (ECG), photoplethysmogram (PPG), and thermocouples enable real time monitoring of vital signs.

Communicating with the dog: a combination of haptic and auditory commands can be sent to dogs using DC vibration motors and a speaker, which enable handlers to instruct their dogs

even when out of sight or earshot. Further, using remotely-operated treat dispensers enables us (or computers) to reward dogs for desirable behaviors from a distance.

Monitoring the environment: sensors on the CEWDs (*e.g.*, gas sensors, GPS, cameras, and microphones) provide handlers with real-time information about potential environmental dangers that confront their working dogs.

These capabilities do not only apply to handler-to-dog communications, but also computer-to-dog communications for many uses of working dogs beyond SAR (*e.g.*, guide dogs for the blind, therapy dogs, guard dogs, *etc*.). Throughout the rest of this article, we describe our cBAN's design, including results from evaluations where available. We also describe future capabilities that are currently in development.

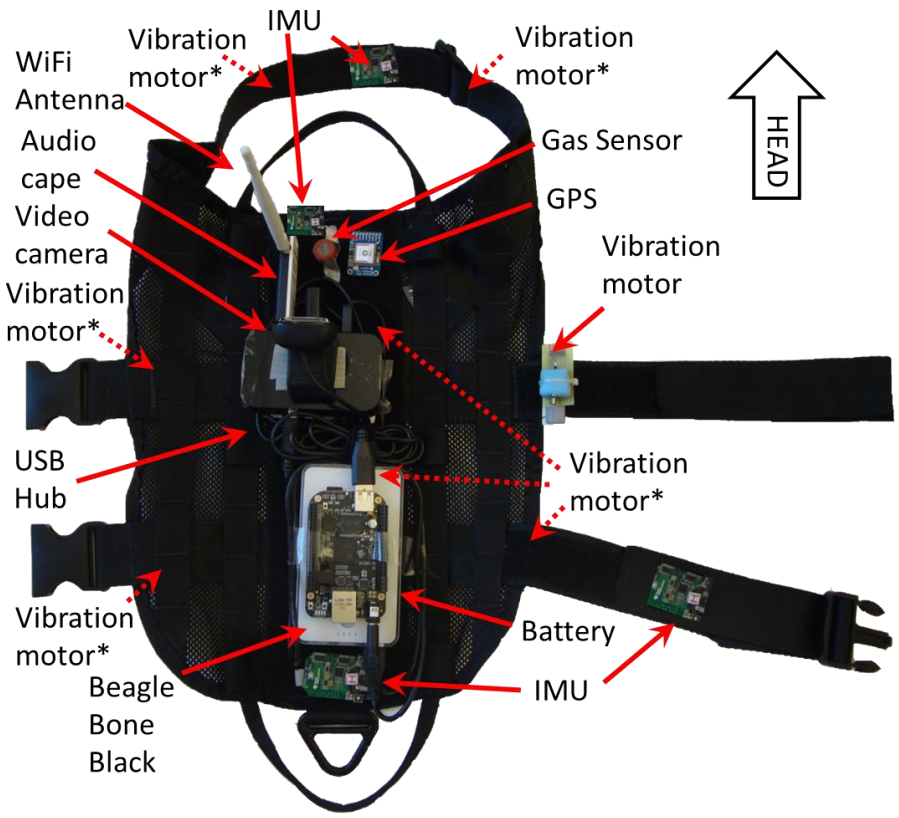

Figure 2.1 The Smart Harness Platform for remote CEWD handling and environmental monitoring. The "*" indicates the location of the vibration motors that aren't visible in the picture. The visible motor was exposed for illustration purposes, and is actually located underneath the strap during harness operation

2.2 Base system

All the electronic sensors and actuator technologies constituting the cBAN and connecting human and canine intelligences are currently mounted on a harness. To handle sensor information and communication from handlers, the harness is equipped with a small BeagleBone Black (BBB) computer. The communication link between the central unit and the remote computer is based on IEEE 802.11 transceivers and this can be provided by mobile base stations, such as unmanned ground or aerial vehicles (UGV or UAV), during real life disaster scenarios. Alternatively, the BBB can be outfitted with cellular or other RF communication hardware. The BBB includes a 1GHz processor from Texas Instruments (TI) (AM3358BZCZ100), 2GB of on-board flash storage and 512 MB of DDR3 RAM. It also includes up to 65 general purpose input output accessible pins, 8 PWM channels, and 8 channels of 12-bit analog to digital converters which are used to interface with the IMUs, vibration motors, and a variety of environmental sensors respectively. The BBB runs Ubuntu GNU/Linux, giving access to the utilities and development tools found on a standard GNU/Linux system. Our system uses COTS software for video streaming functionality; however, we developed all of the other interface software using Python. Most of our communication protocols are ASCII text to allow as many types of clients as possible to control the system, independent of programming languages or operating systems. In addition, most of the communication to and from the base station is done using User Datagram Protocol (UDP) in order to increase speed. We include a watchdog program that is responsible for monitoring the wireless connection and other services, and restarting them if they encounter an error or fail to respond. In addition, the watchdog program accepts commands over the wireless connection to allow the base station to activate, deactivate, or restart any service on demand.

To miniaturize the base system even further, we are in the process of integrating our sensors into a distributed wireless sensor node around the harness. These sensor nodes consist of COTS components that are leveraged on miniaturized flexible printed circuit boards containing a rechargeable battery and antenna unit for short-range signal transmission. Communications between the sensor nodes on the vest and the center unit will come via a

lower power Zigbee link provided by TI CC2530 chips which offer flexible power modes. In the long run, we hope this power conservation may enable recharging batteries on the smart harness through solar or possibly magneto-inductive or biomechanical energy harvesting systems worn by CEWDs for a self-powered (or at least highly energy efficient) operation.

2.3 Monitoring Dog Behavior and Physiology

SAR dogs are trained to perform certain behaviors, such as pointing, crouching, or lying in sternal recumbency ("cover") near the detection site to inform their handlers that they have detected a specific target odor. These trained behaviors are easily recognized by handlers in close proximity to their working dogs, but with the range CMIs will enable in CEWDs, there is a need to give handlers detailed posture and behavior information in real time, even if the dog is out of sight. In laboratory environments canine behavior can be monitored using high-resolution video recordings [13]. However, videos are fundamentally limiting for two main reasons: 1) vision-based systems require computationally expensive image processing to autonomously detect canine behavior; and 2) instrumented environments limit where this information can be obtained, making those techniques useless for practical working dog applications. IMUs offer an alternative to mitigate these problems. Accelerometers, gyroscopes, and magnetometers have been extensively used to monitor activity level in experimental animals in laboratory conditions [14]. Commercially available physiological monitoring systems for canines include accelerometers [15] but these monitors provide processed summary information about physical activity levels, not detailed characteristics of individual behaviors. Commercially available GPS technology has also been used to track a dog's position for sporting applications with a resolution around 0.5 m to 1 m with an accuracy of around 3 m. These systems are not suitable for indoor use, and only provide very coarse-grained location information without no posture detection capability. There is also some recent work using wearable tactile sensors that can be activated by working dogs with a nose push or gentle bite in order to send messages to their handlers [16] instead of presenting the detected canine posture via "passive" inertial measurements.

Remote estimation of canine posture and behavior requires collecting and fusing data from multiple IMUs at different locations on the body [17], [18]. IMU systems have been used extensively with humans for context awareness and activity profiling in health and exercise-related applications and to track the motion and predict the limb positions in real-time for neurorehabilitation, digital gaming, and movie production applications [19]. Canine posture estimation is a relatively new application area for IMUs. There has been some limited attempt to use multiple accelerometers to predict the posture of dogs remotely [20]. In that study, heuristic features and angle computation from accelerometer data were used, but the absence of detailed description of the algorithms and inefficient localization of sensors were limiting.

We studied the kinematics of dogs to identify independently moving locations on the body for locating the IMUs [21], then realized accurate posture estimation by cascading machine learning algorithms. Based on physical parameters, dog comfort, and algorithm performance, we determined the withers (shoulders), chest, abdomen, and back to be optimized locations. Using data collected from these IMUs, we were able to accurately classify five static postures (*sitting, standing, lying down, standing on two legs, and eating off the ground*) and three dynamic behaviors (*walking, climbing stairs, walking down ramps*) for eight different dogs: three privately owned dogs of different sizes and breeds (Kai Ken, Shiba Inu, and Labrador Retriever) and five dogs from a cohort of Labrador Retrievers being trained to detect improvised explosive devices. All procedures were consistent with NIH and USDA guidelines and were approved by the NCSU Institutional Animal Care and Use Committee. The accelerometer data include a static and a dynamic acceleration component. The static acceleration represents the projection of gravity over the axis and the dynamic acceleration is associated with the actual motion of the sensor. Gyroscope output corresponds to the angular acceleration, so during a static posture the output is zero and during dynamic activities it is meaningful.

We developed a cascade of three machine learning algorithms and a filtering stage. The first stage of the cascade consists of a set of Hidden Markov Models (HMMs), each associated with one of the dynamic behaviors of interest. HMMs account for the temporal structure of

dynamic behaviors and each model consists of a set from three to ten states and the probabilities of transitioning and starting at each of the states. Parameters of the HMMs were estimated using the iterative Baum-Welch algorithm. The input sensor data that is not classified as dynamic behavior is filtered with a moving average filter to remove high frequency noise and is fed to a two-level cascade classifier. Each stage consists of a C4.5 decision tree algorithm. The first level identifies the samples associated with transitions behaviors between postures, and the second identifies the postures. Classification accuracies using 10-fold cross validation for the two-level decision tree classifiers are presented in Figure 2.2. For all the reported results inertial information was logged in the computer while the session was video-recorded and algorithm efficiencies were estimated offline. In these, our approach to classification performed very well, achieving an average accuracy above 95 percent in all cases (98 percent excluding data from one trial where we suspect the IMUs were not properly mounted). As expected, the gyroscopes did not increase posture classification accuracy but were useful for identifying dynamic behaviors. The average accuracy of the HMMs in our model exceeded 96 percent.

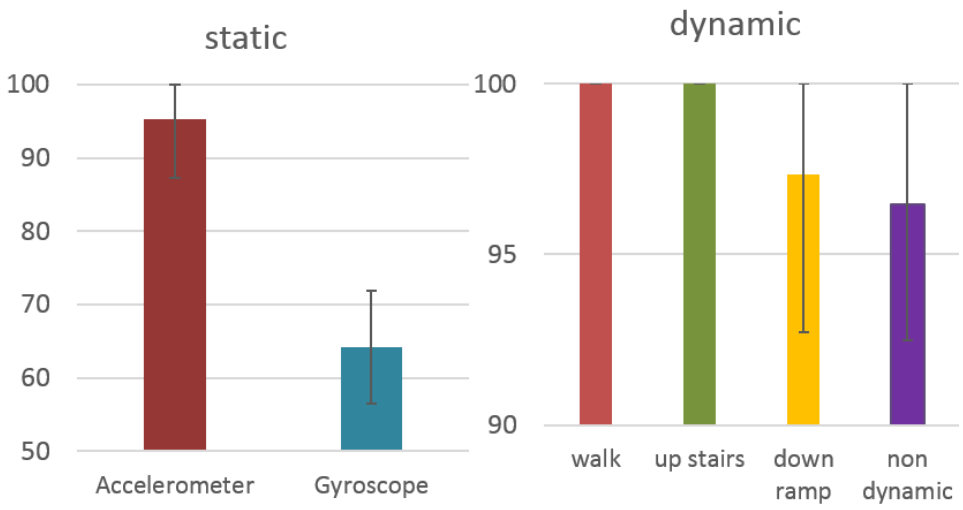

Figure 2.2 Average accuracies for static and dynamic behavior

Monitoring the emotional status and welfare of dogs in training would yield measurable benefits to their performance and attitude. For example, identifying stressors during the

training process may help trainers to adjust exposure to challenging environmental conditions in order to improve learning and reduce the occurrence of fear responses. Behavioral parameters may give clues about emotional states like stress or excitement. For example, the pace and direction and excursion of a dog's tail wag can indicate an invitation to approach or a caution to give space. However, these, as well as many other canine body language communications are subtle and require detailed postural and behavior monitoring. To improve the emotional state interpretation accuracy, behavioral monitoring needs to be correlated with physiological measurements.

The stress induced during the training or handling process has been reported to elicit responses in sympathetic adrenal medullary and hypothalamic pituitary adrenal systems that manifest in cardiovascular, endocrine, renal, gastrointestinal, and hematological parameter changes [22]. The most common methods to measure these responses have been monitoring cardiovascular performance via heart rate and heart rate variability. Respiratory rate may be correlated with dog anxiety as well as physical effort and temperature regulation. Skin temperature is also used to track core body temperature. To obtain consistent and reliable monitoring of dogs' responses to training, we also incorporated physiological sensors into the cBAN to noninvasively measure skin temperature, respiratory rate, heart rate, arterial oxygen saturation, and vocalizations. These physiological sensor nodes on the cBAN contain a surface thermocouple (RTD PT100, Omega), dry ECG contact electrodes; multiwavelength light emitting diodes (LEDs, Marubeni), photodiodes (TSL13T, AMS), and fiberoptic light guides to perform oximetry based PPG. The stainless steel ECG electrodes are similar to metal probes used in a common electronic training collar to avoid shaving the dog. ECG measurements require a differential measurement between three electrodes. The electrodes were incorporated into a chest strap to collect data from the axillary region, with a reference electrode on the abdomen. PPG detects changes in respiration and cardiac pulse through volumetric measurement. It is obtained optically by shining an infrared light into tissue and detecting the amount of light that is reflected to the photodiode. The PPG interbeat interval (IBI) at different frequencies provides cardiac and respiration waves as an indicator of autonomic activity where

the mean value and standard deviation of IBI is recorded as additional attributes of the presence of stress or excitement.

2.4 Computer-mediated remote communication with the Dog

In addition to developing technologies, we are developing training techniques to teach CEWDs to respond to computer-mediated signals from handlers. Despite the ongoing and significant ethical discussion about their use, electronic (shock) training collars are one of the most common remote canine training tools used in the market. They typically are used in two different ways. One type of collar is a feedback system without human intervention. The transmitter is controlled by an automatic electronic control system to keep the dog inside a fixed perimeter without physical barriers. When the dog approaches the perimeter, the shock stimulus is immediate, contingent, and predictable. A second type of collar involves a manual transmitter that delivers shocks at the press of a button. In this case, human error commonly results in delayed, inappropriate, and unpredictable stimulus. The latter type, in particular, may have negative impact on behavior and welfare, depending on the handler's skill and attentiveness [22].

Our CEWD training protocols are based on shaping by successive approximation [23]. Our approach to shaping is completely non-aversive, relying instead on food-based rewards to motivate desirable behaviors. Shaping is a technique used in animal training where desired behaviors are taught to learners by means of selectively requiring actions that are incrementally closer and closer approximations to the desired behavior. Eventually, the action necessary for a reward converges to the goal behavior. Shaping by successive approximation in animal training allows animals to learn much more rapidly than if rewards were given only when the exact target behavior was performed [23].

To deliver rewards, we use computer controlled treat dispensers that enable either humans or algorithms implemented on computers to train the dogs. Evaluation of our computer-controlled training algorithms is still on-going work; however, at the time of

submission of this article our algorithms have successfully trained two dogs to perform specific behaviors to receive treats from computer remote dispensers. In these, we also leverage CMIs by training CEWDs to respond to tactile and/or aural cues where we have incorporated simple aural and haptic communication actuators into the cBAN (Figure 2.1) to serve as a platform for both human- and computer-trained tasks to be performed without the direct presence of the handler [24]. Specifically for SAR applications, we are conditioning dogs to move in seven different directions based on patterns of haptic inputs. Our training pairs the haptic or audio cues with desired behaviors using either human- or machine-delivered rewards. We realize this using a COTS microphone (ADMP401) to record vocalizations, mini-speakers (CMT-1603), and mini vibration motors (306-109, Precision Micro) to provide training and command signals to dogs. The haptic vibration motors are distributed around the harness in pairs: two motors on each side, one pair on the back of the dog, and one pair in the front across the chest. The amplitude and vibration pattern of these motors is controlled from the cBAN central processing unit using pulse-width-modulated (PWM) signals. We are in the process of training dogs to perform commands upon feeling the motors activate in specific patterns.

2.5 Monitoring the environment

In addition to the physiological and inertial monitoring described earlier, our platform for CEWDs incorporates a range of environmental monitoring capabilities to help handlers understand the environment the dog is working in.

During disasters, the micro-environment in which CEWDs work may have toxic elements, which would negatively impact both dogs and human survivors. To monitor the environment, we included on the harness a connector to attach different gas sensors for carbon monoxide (MQ-7), hydrogen (MQ-8), methane (MQ-4), and liquidized petroleum gas (MQ-6, from Hanwei electronics). These gas sensors output a voltage that is proportional to the concentration of the detected gas. Software then reads the digitized value and sends it via UDP packets back to the base station. A GPS receiver (MTK3339 from Adafruit) and a forward-pointing video camera (Lifecam HD 3000) are also included to provide location information and improve context awareness. Upon request the camera streams video to the base station via

UDP. Combining the location information provided by the GPS and the terrain information obtained from the camera provides the operator with enough situational awareness to send commands to the canine in order to direct them to locations of interest. In the future, we are planning to incorporate streaming video from UAVs following the CEWDs. The system is also equipped with an audio adapter and speaker allowing the base station to send text to the BBB where eSpeak, an open source software speech synthesizer vocalizes the message through the speaker. This facilitates communications with nearby people or enables aural commands for the CEWD.

2.6 Discussion and Conclusion

Our cBAN prototype is fully implemented and has been used for evaluations of posture and behavior estimation, computer training, environmental awareness, and physiological monitoring but there are still many exciting avenues for future revision to its design, including power budgeting control algorithms, incorporating IMU measurements for noise cancelation in physiological measurements, additional haptic command training, and incorporation of the CEWD harness into disaster response teams. We are in the process of developing an intelligent model-based duty cycling algorithm that decides which sensors to turn on and communicate with the central unit under what conditions.

A large motivation for our work in developing computer-mediated communication technologies for dogs is to enable not just CEWDs, but also to enable richer interactions between dogs and a variety of other cyber-physical technologies that are gaining ground in the disaster recovery world. In particular, computer-mediated communication may enable CEWDs, UAVs, and UGVs to work together in semi-autonomous intelligent teams where unmanned vehicles monitor dogs' positions and physiology, and send action commands. Figure 3 highlights conceptual depictions of our CEWDs with a mobile robot and a drone to provide communications. Although progress is in its early stages, our vision for Cyber-Enhanced Working Dogs was partially demonstrated under Smart America Challenge to provide next generation emergency response capabilities (www.smartamerica.org).

The foundational work described in this article comprises the fundamental physical and algorithmic building blocks of a novel cyber-physical communication platform that enables intelligent computer-mediated two way communication between working dogs and their handlers. Our goal is to supplement and augment the capabilities of, and improve the welfare of, working dogs using an integrated cyber-enhanced dog-handler communication interface to enable handlers to safely work at distances from their dogs. Other applications in which dogs work remotely, such as detection of improvised explosive devices or cadavers, could similarly benefit from this Cyber-Enhanced Working Dog system. This system provides novel capabilities: 1) for human handlers to acquire a comprehensive physiological and behavioral picture of their dogs and their dogs' microenvironments in real time; 2) for dogs to acquire clear, non-aversive tactile and auditory inputs from humans to direct their search and insure their safety; and 3) for computers to perform positive reinforcement-based training. Using computer intelligence to connect human and canine intelligence would amplify the remarkable sensory capacities of SAR dogs that enable them to save lives.

The work presented here contributed to a multi-institution partnership between Boeing, Mathworks, National Instruments, University of Washington, MIT Media Lab, University of North Texas, and Worcester Polytechnic Institute to demonstrate a Smart Emergency Response System (SERS) for the Smart America Challenge organized by the Office of Science and Technology Policy under the White House.

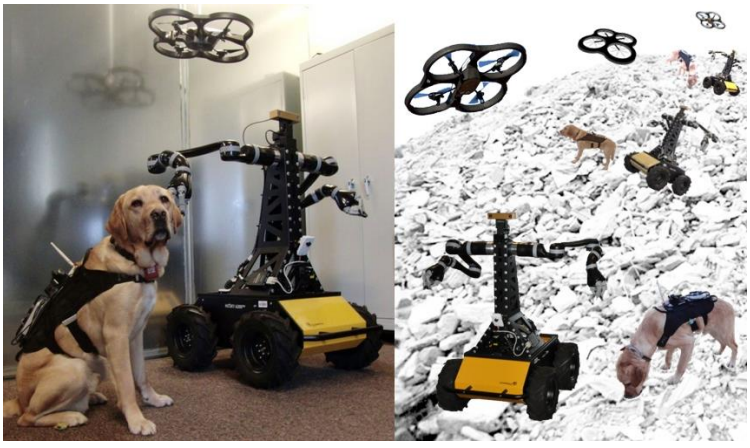

Figure 2.3 CEWD working dog and prototype base station. Conceptual depiction of CEWD and a sensor network formed at a disaster site

Chapter 3. Behavior recognition based on Machine Learning Algorithms and Inertial Measurement Units

3.1 Introduction

A low-power and low-cost bidirectional wireless infrastructure is required to incorporate context-aware and objective sensing and detection algorithms for more sensitive and specific awareness of canine's behavioral responses to training. Newly emerging system-on-chip (SoC) solutions are promising for this where analog, digital and mixed-signal circuits are combined with radio-frequency functions on a single substrate level. For this cBAN, we used a single chip solution from Texas Instruments (TI) [25]. CC2540 from TI combines an 8051 microcontroller with a high performance radio-frequency transceiver, while providing 8 KB of RAM and up to 256 KB of flash memory. It also provides tailored software to fit in with 2.4GHz Bluetooth Low Energy standards to establish connections with computers and smartphones. Optimizing the power budget is also possible with its flexible power modes. This single-chip system is an ideal solution for our canine machine interface with its 21 general-purpose input/output pins and 8 channel 12-bit analog to digital converter. CC2540 comes in a 6x6 mm2 package.

Two different inertial measurement sensors were connected to CC2540 through a serial-peripheral-interface (SPI) to set-up a sensor node. A MEMS-based three-axis ultra-low power accelerometer from VTI (CMA3000) [26] was used, where the unit was set-up for ±2 g measurement range and comes in a 2x2 mm^2 package. As the gyroscope, a MEMS-based three-axis device from ST Microelectronics (L3G4200D) was selected [27]. The package size of this system was 4x4 mm^2. For this study, we used commercial-off-the-shelf (COTS) evaluation boards to connect CC2540 with the inertial measurement units.

Data collection software was prepared using MATLAB to collect and store six sets of data (three-axis accelerometer and gyroscope) simultaneously received from each of the sensor nodes through the established Bluetooth link via a USB connected dongle (Figure 3.1).

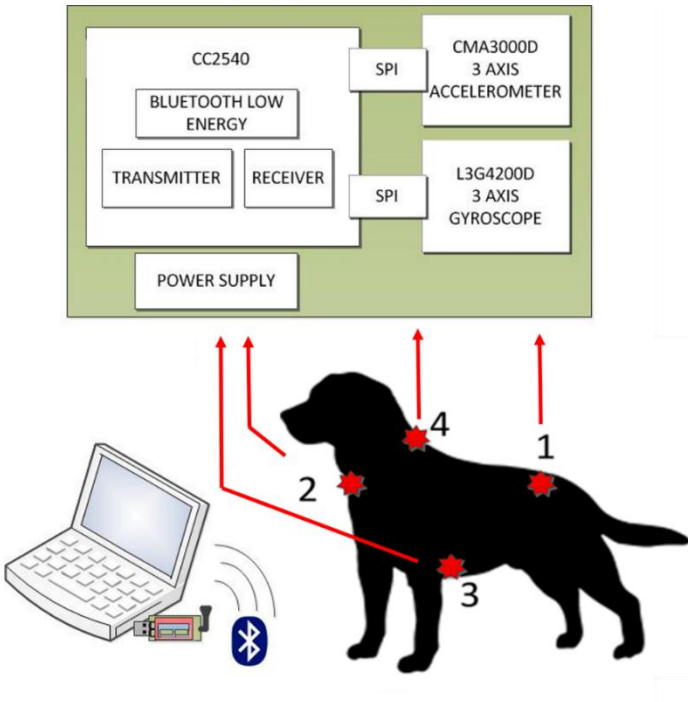

Figure 3.1 Block diagram of the Bluetooth Low Energy sensor system

The machine learning algorithms that we have developed use heuristic features from accelerometer and gyroscope for the static and dynamic posture recognition. Accelerometers

sensed static and dynamic acceleration. The static acceleration corresponds to the projection of gravity over the axes and the dynamic acceleration is associated with the vibration and actual motion of the sensor. During the performance of a static behavior, the accelerometer output shows the tilt angle of the sensor with respect to gravity, and because there is minimal motion, the dynamic component is very small. For dynamic activities, combinations of both components were present. The gyroscope provides the angular data rate so the output when performing a static pose will be zero.

3.2 Materials and Methods

Our platform has been tested with eight canines. Three of them were privately owned pets of different size and breed (Kai Ken, Shiba Inu and Labrador Retriever). The other five were Labrador Retrievers from a cohort being trained and tested as military working dogs to detect improvised explosive devices (IEDs) [28]. We had access to them in the College of Veterinary Medicine at North Carolina State University and all the animal procedures were approved by the NCSU Institutional Animal Care and Use Committee (IACUC).

For the privately owned pets the data collection was performed in their home environment. The owners of the animals asked the dogs to perform a battery of trained behaviors (sitting, lying down, standing, eating off ground, and standing on two legs) repeatedly. To have reliable correlation, the same posture was performed at least five times for approximately four seconds with the dog returned to standing between repetitions. The sampling rate was 10 Hz, resulting in approximately 200 instances of each behavior. Inertial information was logged in the computer while the session was video-recorded for offline analysis.

For the cohort of IED trained dogs, the same procedure was followed. The sessions were conducted in the College of Veterinary Medicine where we also had access to a flight of stairs and ramps. This enabled us to collect data on dynamic activities (climbing up stairs, walking, and walking down a ramp). The dogs were led by a handler through the following sequence three times each: walk up the stairs, walk across a platform to the ramp, walk down the ramp, and walk back to the starting position.

In the following sections the classification flow and accuracy for the detection of static postures and dynamic activities is presented. The static pose detection performance is evaluated on a group of privately owned pets, and the dynamic activities on the cohort of IED's dogs.

3.3 Experimental Results

3.3.1 Optimization of number and location for IMUs

When developing wearable technologies for canines there are limited number of sensor sites available to pick from. For example most dogs will not tolerate sensors strapped to their paws, tail or around their snout.

In [18], [21] we show how we optimized inertial sensor sites by considering the kinetics of the canine to identify independently moving locations on the body. Four locations were tested: the chest of the animal, the abdomen, and two locations at the back; one close to the head (around withers) and the other close of the tail (around rump) of the animal.

Below, we refer the sensor node on the rump with the x-axis of the accelerometer pointing towards the tail as 'Node1'; 'Node2' is the sensor node on the chest with the x-axis of the accelerometer pointing towards the face; 'Node3' is the sensor node located on the abdomen with the x-axis of the accelerometer pointing towards the tail and finally 'Node4' is on the back closer to the head with the x-axis of the accelerometer pointing towards the head.

Since accelerometers measure not only the dynamic acceleration due to actual motion, but also the static acceleration, which corresponds to the projection of gravity over the axes of the sensor, different postures may result in similar data depending on how the IMUs are aligned with respect to the dog's body and gravity, thereby requiring a multi-sensor measurement to assess the posture of the dog accurately.

In order to select the locations we looked at angle of change along three axes of accelerometer data for different postures relative to the baseline standing posture. Consistent

positioning of the accelerometers between trials helps maintain similar amount of angle change for similar task with different dogs. The gyroscope data can be used to correct the misalignments in the accelerometer positioning. However, even if the positioning is not very accurate and the multi-sensor correction is not applied, the patterns in the angle change still can be used to distinguish and estimate the separate postures to a certain extent. Figure 3.2 compares the angle change obtained in all three dogs for four postures (sitting, standing, walking, and standing on the two back legs) as collected by sensors on the back. We observed repeatable distinct patterns in the y-axis for all four posture types in all dogs. During standing, the horizontal position of the sensors was not altered where the average angle change was not significant. Sitting postures caused a change in the order of 5-10 degrees while two-leg standing changed the y-axis angle more than 25 degrees. Both Node1 and Node4 had very similar patterns (Figure 3.3).

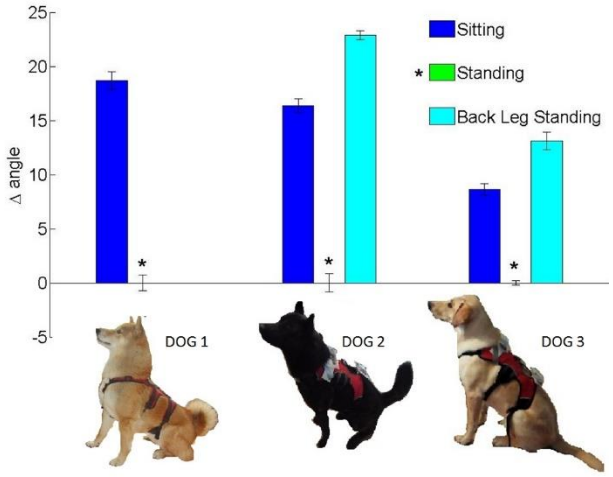

Figure 3.2 Three different canine breeds (Shiba Inu, Kai Ken and Labrador Retriever) were tested for two or three postures. Dog 1 is missing the “Back Leg Standing” posture. The data shows the y axis data obtained from the back sensors. Error bars represent the standard deviation of the mean

The two sensors on the back as it was suggested by [20] will lead into very similar sensor data due to both sites moving together during most postures. For example, during both standing and walking the back of the dog is horizontal, therefore a single system at the back is

not reliable by itself. To distinguish these two postures accurately, a second sensor at the chest was used providing a larger change during the walking behavior (Figure 3.4).

The optimal locations for IMUs leading to larger angle changes between postures were observed when sensors were placed at the chest of the animal (Node2) and at the end of the back around the rump (Node1).

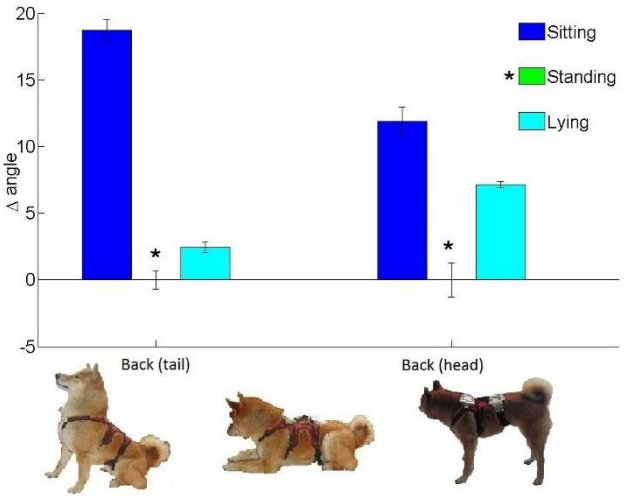

Figure 3.3 Comparison of two back sensors (Node1, Node4) for three different postures performed by Dog 1. Error bars represent the standard deviation of the mean

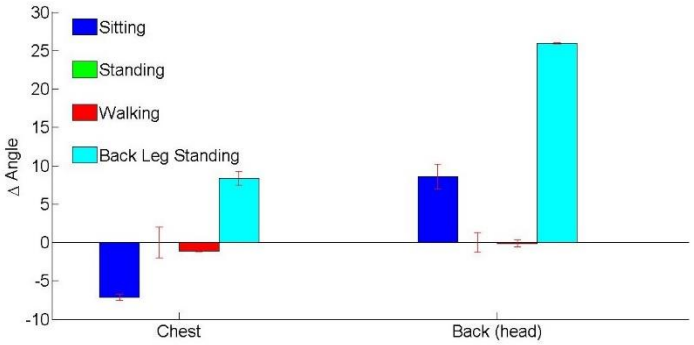

Figure 3.4 Comparison of chest and back sensors for four different postures of Dog 1. Error bars represent the standard deviation of the mean

3.3.2 Classification of static postures

Data was collected from two privately owned dogs performing a sequence of 5 postures (sit, stand, lie, stand on two legs, and eating off ground). To classify postures we used algorithms implemented in the WEKA Machine Learning toolkit developed at the University of Waikato [29]. We focused on three: Random Forest (RF), k-nearest neighbors (IBK), and Logistic Model Tree (LMT) which offer relatively better performance at a low computational cost for future microprocessor implementation. The data set used for classification was a discretized version of the time series data collected from the inertial sensors [17].

The data were preprocessed with a moving average filter. Using the videos, each sample was manually labeled as one of the five postures or a transition between postures. These labeled data were used to train and to evaluate the machine learning algorithms (Figure 3.5). A representative distribution of the samples in the x-axis acceleration space is shown in Figure 3.6.

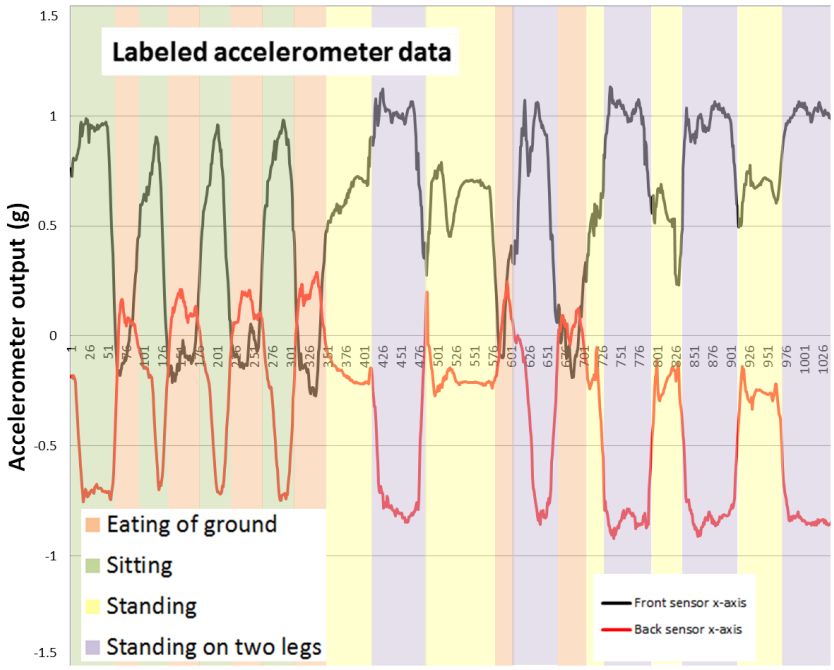

Figure 3.5 Representative accelerometer data with video-analysis based labeling of each posture

The spatial distribution of samples shows the consistency of the data and reliability of inertial measurements. Data corresponding to a static posture is closely tied to the gravity projection over the axes and has smaller standard deviation than inter-posture transitions. Note the widespread instances of stand in contrast to concentrated instances of lying posture. The wider range is associated with movement of the head.

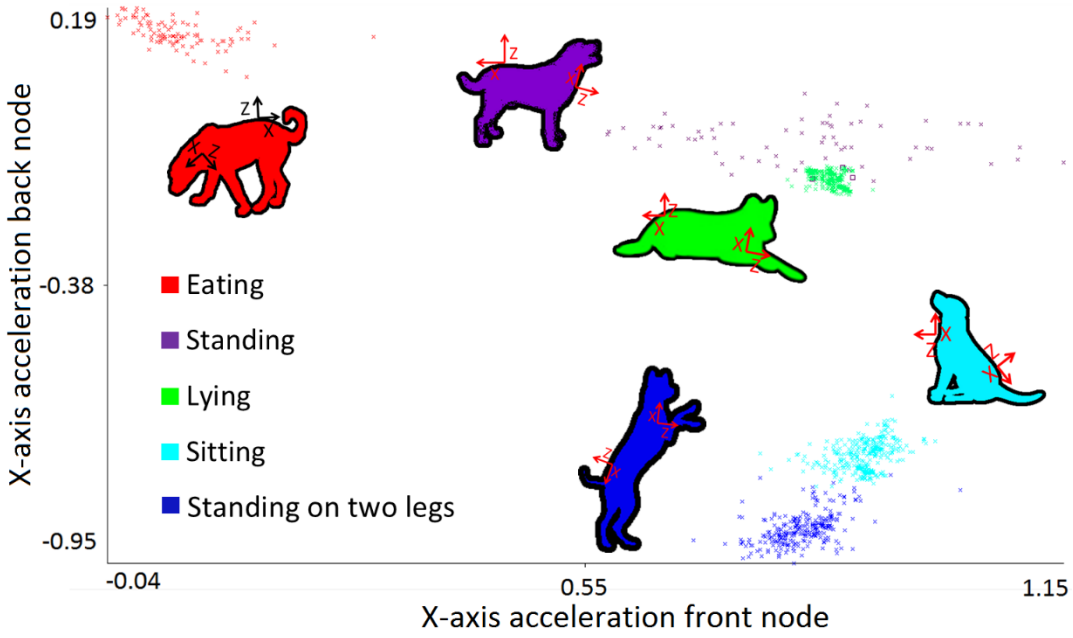

Figure 3.6 Representative distribution of samples in the x-axis acceleration space

Data used in this study is shown in Table 3.1 and were not independently and identically distributed (IID) due to differences in the durations of postures and transitions. A two-level cascade classifier [30] was used to handle this issue. The first (C1) separated postures and transitions. The instances labeled postures by C1 were input into the second (C2). C2 identified the five postures. Both accelerometer and gyroscope data were inputs for C1 and only accelerometer data for C2. The classification flow is shown in Figure 3.7

Table 3.1 Instances of each behavior used for the static posture classification analysis

	TRIAL 1	TRIAL 2	TRIAL 3	TRIAL 4
Subject	Labrador	Labrador	Kai Ken	Kai Ken
Total	772	1164	1336	1647
Sitting	72	476	375	400
Standing	141	280	70	221
Eating off the ground	152	148	100	123
Standing on two legs	299	132	286	286
Lying	0	0	294	292
Transitions	108	128	211	225

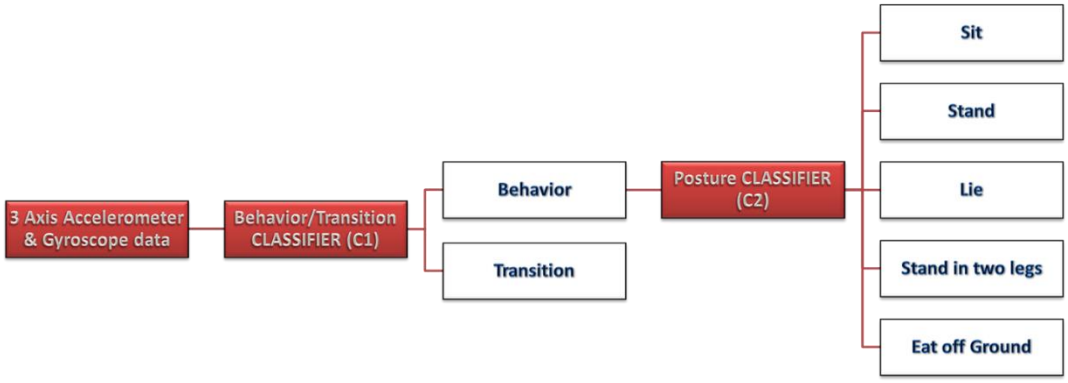

Figure 3.7 Classification flow for static posture detection

3.3.2.1 Algorithm performance

We evaluated the classification accuracy using the percentage of correctly classified instances. Most misclassified samples were at boundaries between postures and transitions and were likely errors in data labels. Discarding four samples at these boundaries improved classification accuracy up to 2.2%. Using gyroscope data improved C1 accuracy, especially for relatively still postures. Preprocessing gyroscope data and further optimizing the gyroscope location may also improve accuracy.

Table 3.2 contains accuracy values for C1 and C2. C1's best accuracy was achieved using the RF algorithm and C2's using IBK. LMT performed well in just one trial. We used 10-fold cross validation. For the remaining experiments, we present only results from RF for C1 and IBK for C2. We also examined performance using data from one trial for training and

another trial for testing. The sensors were removed and replaced to evaluate the algorithm's sensitivity to slight shifts in sensor location and orientation.

To evaluate inter-subject robustness, we trained classifiers using one dog’s data and tested it on the others. Data in

Table 3.4 suggest there was more variability between subjects during transitions than static postures. This makes sense given the dogs' different sizes, but similar kinematics. We also found that inter-dog classification errors were more likely between sitting and standing on two legs and between lying and standing postures. This is unsurprising as the angular positions of the dogs' joints for these pairs of behaviors are similar.

Table 3.2 Accuracy of C1, C2 and Combined algorithm. In parenthesis: best value providing algorithm

Dataset	**C1 (%)**	**C2 (%)**	**Combined (%)**
TRIAL 1	97.40 (IBK)	99.69 (LMT)	98.2
TRIAL 2	98.28 (RF)	99.61 (IBK)	98.28
TRIAL 3	98.80 (RF)	99.73 (IBK)	98.5
TRIAL 4	98.23 (RF)	98.79 (IBK)	98.39

Table 3.3 Accuracy with and without sensor repositioning (best and worst cases)

	Case	**Subject**	**Same position**	**Position changed**
C1	Worst	Labrador	98.19 %	81.08 %
	Best	Kai Ken	98.11 %	91.91 %
C2	Worst	Kai Ken	99.73 %	97.50%
	Best	Kai Ken	99.78 %	98.19 %

Table 3.4 Accuracy sensitivity to use of same or different subjects for training and testing (best/worst cases) Values in parenthesis correspond to cross-validation

	Case	**Training Subject**	**Testing Subject**	**Cross-canine**
C1	Worst	Kai Ken (98.11 %)	Labrador	81.01 %
	Best	Kai Ken (98.42 %)	Labrador	91.99 %
C2	Worst	Kai Ken (99.64 %)	Labrador	91.79 %
	Best	Labrador (99.8 %)	Kai Ken	100 %

3.3.3 Classification of dynamic activities

In this study in addition to the static poses we included dynamic activities such as walking, climbing up stairs and walking down a ramp. Figure 3.8 shows a representative sample of acceleration data from the x-axis of the accelerometer in each of the four locations. During the performance of a dynamic activity the accelerometer data presents higher variance that is the dynamic component of the accelerometer output.

Table 3.5 shows the distribution of instances available for the analysis. Dog1 to Dog5 are Labrador Retrievers from the cohort being trained as military working dogs. Dog6 and Dog7 correspond to recordings from additional animals for comparisons. These were two privately owned pets, a Labrador Retriever and a Kai Ken. For Dog1 and Dog2 data were available for the four sensor nodes whereas in other cases only 'Node1' and 'Node2' recordings were available. In Table 3.5 the set of instances per behavior and dog trial that were used for the analysis is presented.

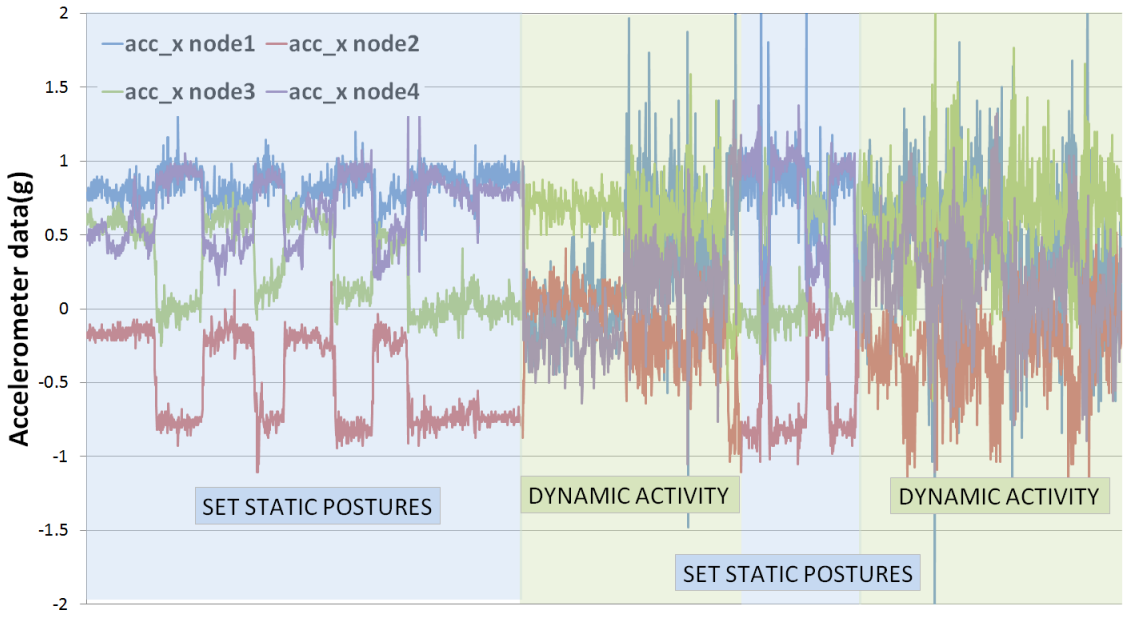

Figure 3.8 Representative accelerometer data of the x-axis in the four sensor locations showing static and dynamic activities

Table 3.5 Instances of each behavior available for the analysis

Behavior	Set of instances per trials								
	Dog1	Dog2	Dog3	Dog4	Dog5	Dog6_1	Dog6_2	Dog7_1	Dog7_2
Sitting	542	231	366	815	204	72	476	375	400
Standing	550	449	359	1155	368	141	280	70	221
Eating off ground	244	268	51	106	188	152	148	100	123
Standing two legs	215	245	204	180	210	299	132	286	386
Lying down	0	0	0	74	235	0	0	294	292
Walk	718	432	370	0	0	0	0	0	0
Climb stairs	102	72	94	0	0	0	0	0	0
Down ramp	75	54	92	0	0	0	0	0	0
Transition	82	185	181	272	188	108	128	211	225
Total	2528	1936	1717	2602	1393	772	1164	1336	1647

Figure 3.9 shows the classification flow we developed for this work. Maximum likelihood estimation with HMMs was used in the initial stage to identify each of the dynamic activities. HMM input was a combination of dynamic and postural recordings. Data not classified as a dynamic activity was processed with a moving average filter, and fed to a decision tree cascade classifier for the recognition of the transitions and the specific static poses. In this study behavioral classification was done offline. For the trials where ‘Node3’ and ‘Node4’ were available, those recordings were not used in the HMM implementation but they were used in the posture classification.

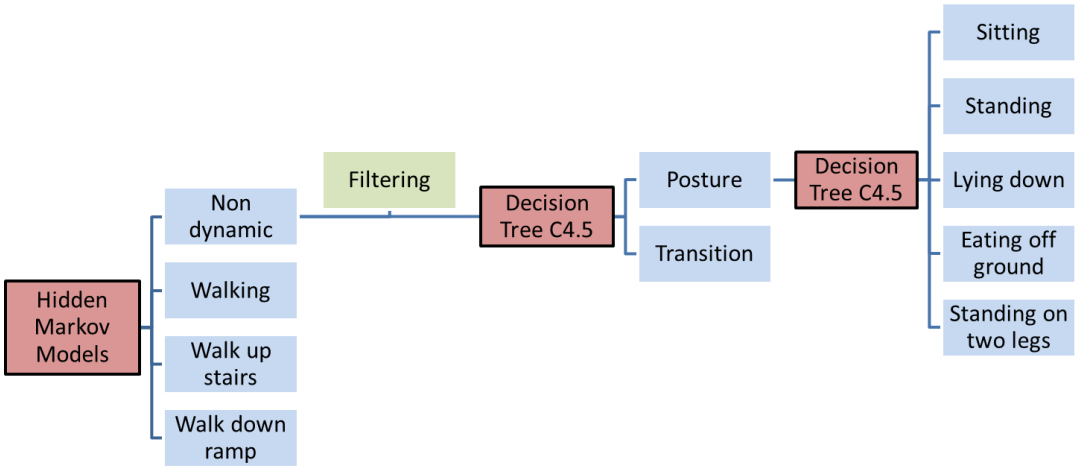

Figure 3.9 Classification flow for dynamic and static activity detection

3.3.3.1 Algorithm performance

3.3.3.1.1 Dynamic activity recognition

HMMs were used to represent and identify dynamic behaviors, as the classification methods used for static poses such as sitting do not explicitly account for the temporal structure of these behaviors.

Our HMMs consisted of a set of possible states that the dog could be in, along with the probabilities of transitioning from one state to the next at each time step and the probabilities of the model starting in each of the states. Each state generated an observation according to its observation distribution. In this case, the continuous sensor vectors were assumed to be drawn from a multivariate normal distribution with zero covariance, such that each feature was independent given the state.

To estimate the parameters of an HMM for a given set of sequences, the iterative Baum-Welch [31] algorithm was used to find parameters which maximized the probability of generating those sequences. As the Baum-Welch algorithm requires the number of states within the model to be known, multiple models were learned with between three and ten states,

and the best of those models was selected. The HMM algorithms were implemented in the Java programming language.

An HMM was learned for each behavior, as well as a single model for all sequences in which none of the three dynamic behaviors were being performed (Figure 3.9). The behavior occurring in an unlabeled sequence could then be found via maximum likelihood estimation, where the behavior was identified as the one for which the associated HMM has the highest probability of generating that sequence. This probability was computed using the Viterbi algorithm [32].

The signals from the sensors were split into ten step segments, using a ten step-sliding window such that the segments overlap. By also learning a model of behaviors that were not dynamic, it was possible to filter out those behaviors and pass them to the static classifiers for identification. The effectiveness of this learning approach was evaluated via five-fold cross validation, with the test sets consisting of thirty percent of the data.

The results in Table 3.6 show that, when learning to identify behaviors for a single dog, the HMM approach achieves high accuracy, and rarely misidentifies a sequence as being generated by another behavior. The results in Table 3.7, however, show that this approach was ineffective when trying to apply models learned for one group of dogs to another. This is expected given the small size of the dataset.

Table 3.8 shows the accuracy achieved by using only gyroscope data for the HMM and Table 3.9 using only accelerometer data. Accuracy results obtained with the gyroscope data were slightly better than with the accelerometer data, but neither of these achieved the accuracy of using both of the sensor data. These results support our intuition that the angular rate read by the gyroscope has an important role in recognition of cyclic activities [33].

Table 3.6 HMM classification accuracy

% correct	Walking	Up stairs	Down ramp	Non-dynamic
Dog1	100%	100%	100%	99.6%
Dog2	100%	100%	92.0%	97.8%
Dog3	100%	100%	100%	92.0%

Table 3.7 HMM accuracy sensitivity to algorithm transfer from a group of dogs to another

% correct	Walking	Up stairs	Down ramp
Testing against Dog1	97.9%	46.3%	29.6%
Testing against Dog2	66.8%	7.7%	100%
Testing againstDog3	97.7%	0%	0%

Table 3.8 HMM classification accuracy against gyroscope

% correct	Walking	Up stairs	Down ramp	Non-dynamic
Dog1	100%	100%	100%	97.9%
Dog2	97.2%	100%	92.0%	98.9%
Dog3	91.6%	100%	96%	93.3%

Table 3.9 HMM classification accuracy against accelerometer

% correct	Walking	Up stairs	Down ramp	Non-dynamic
Dog1	100%	100%	100%	100%
Dog2	100%	100%	100%	91.5%
Dog3	99.2%	100%	98%	83%

3.3.3.1.2 Static activity recognition

For static posture classification, a two-level cascade classifier [30] was used. The C4.5 algorithm implemented in the WEKA Machine learning toolkit developed at the University of Waikato [29] and 10-fold cross-validation were used for the study. This decision tree algorithm was used in the two levels, the first to distinguish between transitions and postures, and the second to classify the specific postures (sitting, lying, standing, eating off the ground and standing on two legs) from the instances classified as postures in the first level. C4.5 builds a classifier with a tree structure from the instances in the training set. The tree leaves represent class labels and branches represent conjunctions of features that lead to those class labels. The C4.5 algorithm uses “information gain” [34] as the splitting criterion for splitting the branch. At each splitting, the decision tree algorithm chooses the feature value providing the maximum reduction in uncertainty about the class labels. So, the feature at the root of the tree is the one

with the maximum information gain, and is therefore the best predictor. The feature used at the second level of the tree is the next best predictor given the value of the first [35].

Table 3.10 shows the average accuracy of the two cascade decision tree classifiers. For Dog1 and Dog2 the average of the accuracy in the two available trials was presented. The first column shows the accuracy achieved by using only the accelerometer data as feature and the second column using only the gyroscope data. From these results, we conclude that the gyroscope data were not required for accurate pose estimation.

Table 3.11 shows the best and second best predictors for the C4.5 algorithm. Results show that, in general, nodes 1 and 2 in the x direction offered the best prediction results. Though there were a few exceptions, these results were basically consistent with our initial studies showing those as the optimal sensor locations [21].

Table 3.10 2-Level decision tree classifier (C4.5) efficiency when using accelerometer or gyroscope data

	Accelerometer	Gyroscope
Dog1	99.5%	72.8%
Dog2	98.9%	63.3%
Dog3	94.4%	61.8%
Dog4	77.8%	63.4%
Dog5	97.5%	74.5%
Dog6	99.3%	62.2%
Dog7	99.4%	51.5%

Table 3.11 Location and orientation of sensor associated to best prediction results for 2-Level decision tree classifier (C4.5)

	The best predictor	The secondary best predictor
Dog1	node2_x	node1_x
Dog2	node3_x	node1_x
Dog3	node1_x	node2_x
Dog4	node1_x	node2_x
Dog5	node2_x	node1_x
Dog6	node1_x	node2_x
Dog7	node1_z	node1_x

3.3.4 Power consumption

In this chapter we have demonstrated the feasibility of a posture detection system using the CC2540 SoC with an embedded Bluetooth Low Energy transceiver. The current consumption of the system during transmission including peripherals (accelerometer and gyroscope) has been measured to be 29 mA at 3.3V. The system is compact in size (5x3 cm) and is powered from a 225mAh coin cell battery. The system proved to be suitable for posture detection, but the reduced number of peripheral interfaces available, and the short range of communication, limits its use as the centralized control unit of the cBAN. To overcome those limitations as we increase the number of sensing capabilities in the cBAN, we selected the Beaglebone Black (BBB) computer, as it includes 65 general-purpose input/output accessible pins, eight pulse-width modulation channels, and eight channels of 12-bit analog to digital converters, which are used to interface with the IMUs, vibration motors, and a variety of physiological and environmental sensors. The BBB transmits the sensor information to a remote data aggregator for further analysis using Wi-Fi. Table 3.12 shows the power consumption of both systems without optimization. The power consumption can be decreased through appropriate duty cycling strategies. The CC3200 is a recent single-chip Wi-Fi solution designed for the Internet-of-Things [36] and it is included for comparison purposes to illustrate that the power consumption of our system is comparable to other Wi-Fi solutions in the market. It is important to note that the CC3200 is based on an ARM Cortex-M4 Core at 80 MHz whereas the BBB is based on an ARM Cortex-A8 at 1GHz. The range of current consumption reported in the second raw for the CC2540 and CC3200 are extracted from datasheet and correspond to different output power conditions. The current consumption of the BBB reported on the datasheet includes peripherals (HDMI monitor connected, 4GB Thumbdrive) and Ethernet connected at 100M.

Table 3.12 Current consumption during transmission for both BLE and BBB platforms. CC3200 Wi-Fi SoC also provided for comparison

	CC2540 (BLE)	CC3200 (Wi-Fi SoC)	Beaglebone Black (Wi-Fi)
Supply Voltage (V)	3.3	3.6	5
Current consumption during transmission (reported datasheet) (mA)	21 - 31	166 - 229	430
Measured current consumption during transmission (includes power consumption of peripherals) (mA)	29	-	265
Average Transmission Power (W)	0.1	0.8	1.3

3.4 Discussion and Conclusion

In this study, we have used IMUs on 7 dogs to obtain 9 trials of data with an aim of estimating canine postures electronically. We have shown that a high level of accuracy in activity recognition can be achieved when building models for each individual dog, both for static postures and for dynamic activities like walking, climbing the stairs, and walking down a ramp. We observed that applying models learned for one group of dogs to another to be far less effective.

HMM algorithms used for dynamic activity recognition obtained better accuracy when using only gyroscope data compared to only using accelerometer data. That supports our intuitions that gyroscope data would play an important role on cyclic activity recognition. The combination of both sensor data provided the highest classification accuracy.

From the static pose classification, we conclude that the gyroscope data did not improve pose estimation accuracy, and we observe that nodes 1 and 2 in the x direction were, in general, the best predictors, consistent with our initial studies showing those as the optimal sensor locations [17], [21].

The implementation and computational cost of the HMM classifier were large in comparison to the decision tree classifiers for static postures, but once HMM models were learned, classifying new examples using these models were fairly computationally efficient when the Viterbi algorithm was used. Extension of HMMs to be able to recognize different levels of intensity in activities, such as walking versus running or galloping, is an ongoing work.

Chapter 4. Wearable Heart Rate Sensor Systems for Wireless Canine Health Monitoring

4.1 Introduction

Interpreting the physiological and emotional state underlying canine behavior is essential in human-canine interactions, to achieve effective training, and to improve canine welfare. Traditionally, dog trainers and handlers monitor postures, vocalizations, and behavioral signs of anxiety (e.g. trembling, restlessness, visual scanning, or yawning) to determine canines' emotional responses to certain stimuli or environmental conditions. There is an increasing interest in making this interpretive process less subjective by monitoring the physiological signals of working, service, and companion dogs. Vital signs, such as heart rate, may serve as indicators of emotions and welfare, and may also be used to aid in the identification of stress or excitement triggers outside of laboratory environments. This information could then be used by handlers to optimize methodologies to improve training effectiveness and efficiency. The overarching goal of our work is to achieve more accurate interpretation of canine emotional response by correlating behavioral responses with the physiological responses of heart rate (HR), heart rate variability (HRV), and respiratory rate [7], [8], [14]. In the long term, the physiological sensors presented here will be a part of a canine-body area network (cBAN) that combines multimodal wearable sensing technologies and computational modeling with behavioral assessment to effectively capture dogs' emotional responses [4], [9], [10], [17], [18], [21], [37], [38](Figure 4.1).

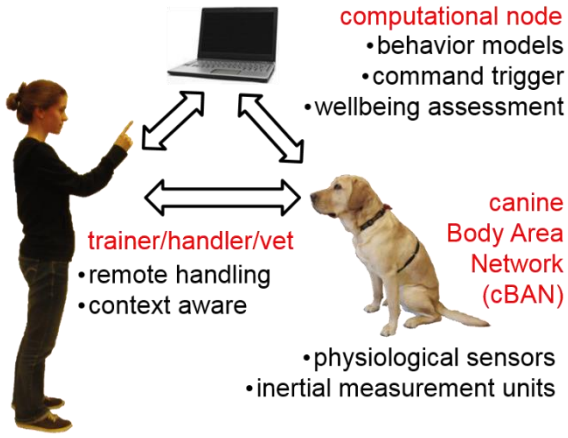

Figure 4.1 General description of the canine-machine interface to enhance human-canine interactions

When at work, dogs are very active. Consequently a heart rate monitoring system needs to be both resistive and robust. Reliably recording canine HR and HRV during static postures and dynamic activities in both interior and exterior environments in a context-aware manner is a challenging task. We present our initial efforts to improve the electrocardiogram (ECG) and photoplethysmogram (PPG) hardware for a better coupling with the tissue to provide more accurate and robust ECG and PPG monitoring outside of laboratory and veterinary environments. We additionally report the results of experiments characterizing the quality of the tissue coupling with these enhancements and comparing it with commercially available ECG electrodes used in clinics [39].

4.2 Background

In general, to maintain homeostasis, physical activity has a positive effect on HR and respiration rate. HR is the number of heart beats per minute and is based on the number of contraction of the ventricles. The variability in the interval between consecutive heart beats, referred to as HRV, is considered a valuable indicator of the fluctuations in autonomic nervous system inputs to the heart [9]. The ECG reveals the electrical activity of the heart. It is the gold standard to measure the HR as time between consecutive cardiac electrical impulses (heart beats), referred to as R-R intervals. Variability in R-R intervals represents interplay between

two branches of the autonomic nervous system, the parasympathetic branch, which has an extension effect on the R-R interval duration mediated by the release of acetylcholine by the vagus nerve, and the sympathetic branch, which has a shortening effect mediated by the release of epinephrine and norepinephrine [9]. In dogs, the R-R cycle length normally varies with respiratory activity. Called respiratory sinus arrhythmia, the R-R interval shortens during inspiration and lengthens during expiration [40]. R-R interval series are analyzed by time and frequency domain methods to calculate HRV measures. A number of studies have used HRV as an indicator of stress in animals [10][41].

The traditional ECG method in a veterinary setting involves pinching "alligator" clips applied to the skin to attach electrodes on the animal's body when it is anesthetized for surgery. These would be painful if used on awake or active animals and cause inflammation. Alternatively, commercially-available human ECG electrode patches are used for at-home cardiac monitoring. For these, the animal's hair must be shaved at the ECG surface electrode attachment sites to ensure optimal adherence and avoid detachment due to the movement of the animal [37]. Dense hair and subcutaneous fat for thermal insulation on the dog's body dampen electrical conductivity and limit the useful application of these patches. One possible solution for securing these electrodes is to use tape and bandages, as used in Holter monitoring, although it is time-consuming for daily use and could be uncomfortable for animals [42]. Moreover, most of these clinical systems only provide on-board recording capabilities, and the wireless telemetry systems are too expensive for home-use [43]. Implantable sensors provide another solution but are also expensive and require surgical placement [15].

In addition to traditional ECGs used in veterinary clinics, some new technologies are appearing in the consumer market for daily monitoring of HR and HRV in humans using dry electrodes. Some of these products from Polar® have been used by veterinary researchers and have shown reasonable correlation on stationary dogs, and also during dynamic activity when the hairs are clipped and additional conductive electrolyte gel is liberally applied [44]–[46]. PetPace® and Voyce® are two emerging commercial devices that use non-conventional acoustic and ultra-wideband radio-frequency technologies, respectively. At present, due to the

recent appearance of these products in the market, there is limited information available about their technologies and the clinical validation of their sensing capabilities and reliability.

To overcome some of the difficulties associated with the application of contact-based ECG electrodes to canines, we introduced novel electrode configurations that allow for signal recording without the necessity of shaving the dog's hair. We evaluated two types of electrodes to improve the tissue-electrode contact despite the existence of a dense hair layer: thick-tapered electrode and comb-shaped arrays of thin spring-loaded pins. We demonstrated that surface modification of the metal electrodes by coating with poly(3,4-ethylene-dioxythiophene) poly-(styrene-sulfonate) (PEDOT:PSS) conductive polymer can enhance their electrical property in terms of increasing the charge storage capacity and decreasing the impedance of the tissue-electrode interface [47]–[49].

We included PPG in our sensor system for self-calibrating the HR and HRV measurements. PPG is an optical volumetric measurement of blood flow obtained by shining infrared light into the tissue and detecting the amount of backscattered light with a photodiode. The traditional PPG recording has only been used on anesthetized animals during surgeries where traditional pulse oximetry clips were attached to locations, such as ear or tongue that is not practical for dogs' activities outside of veterinary environments [50]. We studied the incorporation of light guides and optical fibers to create an efficient optical coupling to a dog's skin through the hair when performing PPG measurement in wearable form factors at hairy locations on their chest. Placing the sensors on the chest enables us to use a single chest strap to contain all the sensors without restricting the range of motion. When the dog is at rest, the PPG based HR could be correlated with ECG based measurements to validate the accuracy of the measurements. Beyond this, concurrent measurement of PPG and ECG can track the arterial pulse transient time (PTT) which has been an indirect measure of relative changes of blood pressure or stress levels in humans and has a potential to be used in canines [51].

As the dogs' heart rate is correlated with its behaviors [52], it is essential to monitor the behaviors in parallel to HRV. In our earlier work, we demonstrated the capability of a cBAN to detect static dog postures (sitting, lying down, standing, eating off the ground, and

standing on two legs) and dynamic activities (walking down a ramp, climbing up stairs, and walking on flat ground) using a cascade of machine learning classifiers [30]. The classifiers used the linear and angular accelerations provided by IMUs worn by the dog in two locations. In particular, Hidden Markov Models were used for detection of dynamic behaviors, and decision tree classifiers for static postures, achieving an average accuracy above 95% [18], [33]. This provides the capability of identifying behaviors remotely and automatically without the need for camera recordings and video processing as traditionally used in veterinary behavior research and clinics. In addition to behavior monitoring, such an accelerometer can be placed at the site of the ECG electrodes to correlate activity levels to ECG signal to noise ratio to further understand the sources of noise. It is noteworthy that, this IMU can also be used to perform seismocardiography by registering the vibrations due to cardiac contraction and respiratory modulation and potentially calculate the resting HR and respiratory rate [53], [54].

4.3 Materials and Methods

To evaluate the performance of the ECG and PPG systems, we connected these to a wearable BeagleBone Black (BBB) platform integrated into an adjustable elastic belt placed around the thorax behind the two front legs. The BBB served as the processing unit of the cBAN to retrieve the sensor information and transmit it using a wireless communication link based on the IEEE 802.11 standard for Wi-Fi (Figure 4.2) [55]. The analog outputs of the ECG and PPG systems were directly connected to the 12 bit analog-to-digital converter (ADC) channels of the BBB and the digital output of the 3-axis accelerometer to the I2C serial bus. For experimental flexibility and to support scenarios where Wi-Fi is not available or battery power is limited, we also included a Texas Instruments system-on-chip (cc2541) in the cBAN to be able to connect to the computational node carried by the handler (e.g. smartphone) directly using Bluetooth low energy (BLE).

For this study we recruited five dogs of various size, breed, and hair density at the North Carolina State University (NCSU) College of Veterinary Medicine (CVM). Table 4.1 shows information on the five dogs participating. All animal procedures were approved by the Institutional Animal Care and Use Committee of NCSU.

Table 4.1 Dog information

Dog number	Breed/Breed Type	Age (years)	Sex
1	Labrador	3	M
2	Mixed	11	F-Spayed
3	Greyhound	9	M
4	Otterhound	6	M
5	Mixed	2	M

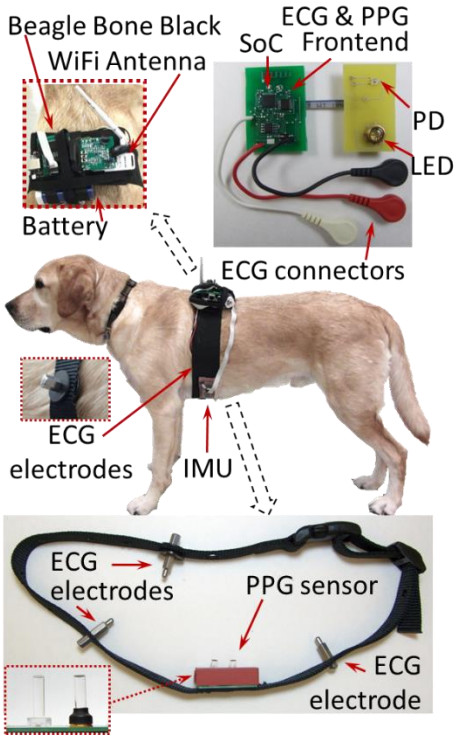

Figure 4.2 (Top) Top view of the BBB and the PCB for ECG and PPG front-end circuit and BLE SoC. (Middle) A labrador retriever with the cBAN harness with sensor locations indicated. (Bottom) Chest strap including ECG and PPG sensors with inset showing the light guides for PPG recording

4.3.1 Improving the ECG Electrodes

To overcome the limitations imposed by dog fur, we used two types of electrodes for the recording of ECG signals: a four millimeter thick stainless steel pointed style electrode, which is traditionally used in training collars, and electrodes that are comprised of arrays of one millimeter thick spring-loaded gold-coated pins similar to dry electrodes used in EEG caps [56] (Figure 4.3). We hypothesized that using an array of multiple thinner electrodes instead

of a thicker single one would reduce the tissue-electrode impedance and improve the hair penetration. It would also improve the comfort by distributing the applied pressure to a larger area. We used a commercial-off-the-shelf (COTS) pointed-style stainless steel electrode (RFA-529) with added silicone cushioning to expose only the narrower tip for favoring the attachment and comfort (Figure 4.2). For the second type of electrode, COTS copper alloy spring-loaded pins with 1.07 mm diameter, 8 mm height and 0.51 μm gold plating (Mill-Max Mfg.) were soldered on a mini printed-circuit-board (PCB) to create a comb-shaped array (Figure 4.3). The small dimensions of the pins would allow easy penetration through the hair and the spring mechanism would help in maintaining the skin-to-electrode contact under motion or rapid excursions of the chest during panting.

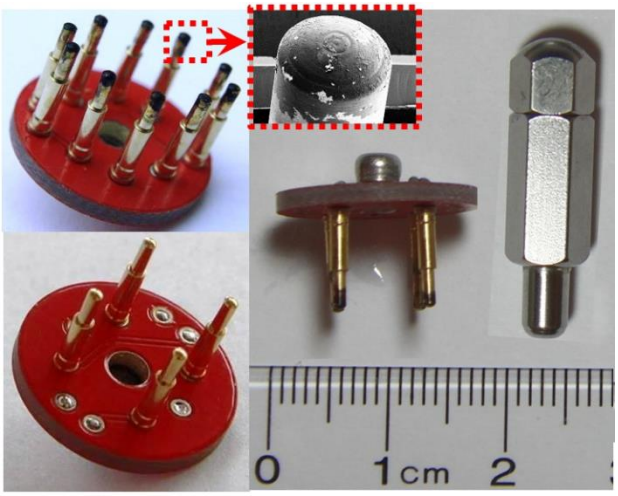

Figure 4.3 (Left) Comb-shape array electrode with four and ten gold spring-loaded pins. (Inset) SEM of the PEDOT-PSS coated tip of the ten pin array. (Right) Relative sizes of stainless steel electrode and four pin array.

To further reduce the impedance of the electrode array, we coated the tip of the gold-coated pins with PEDOT:PSS. We prepared the PEDOT:PSS monomer solution in the laboratory by using a magnetic stirrer to mix 35 mg of 3,4-ethylenedioxythiophene with 250 mg poly (styrenesulfonic acid sodium salt) in 25 ml of deionized water for two hours [47]. All the chemicals for both solutions were purchased from Sigma Aldrich, Inc. A prerequisite of the coating process involved the ultrasonic cleaning of the gold plated pins using acetone

followed by alcohol. For the electrochemical polymerization process, we dipped the tips of the spring-loaded pin arrays in the PEDOT:PSS solution along with a 500 μm platinum wire as counter electrode while applying galvanostatic charges of 100 $\mu A/mm^2$ for 10 minutes. Figure 4.3 shows one of the polymerized pins with the characteristic black coloration of PEDOT:PSS coating. It should be noted that direct electro-polymerization of PEDOT:PSS over stainless steel is not possible therefore, only the gold-coated pin arrays were coated with PEDOT:PSS in this study.

After assembly, we connected the three electrodes required for ECG recording to a COTS heart rate monitor analog front-end chip (AD8232, Analog Devices). This integrated circuit included an instrumentation amplifier, an operational amplifier for filtering and extra gain, and a right leg drive amplifier to compensate for the common-mode voltage variations to improve the common-mode rejection ratio of the system. We designed the narrowband filter bandwidth to be between 7-24 Hz. We selected this bandwidth to remove motion artifacts at the cost of some distortion of the ECG signal. This type of distortion can be afforded in our application as we were interested in the detection of HR and HRV through the R-R interval rather than observing the full PQRST waveform. In addition, AD8232 provided additional circuitry to reduce the settling time and fast signal restoration when a saturation condition was detected at the output of the instrumentation amplifier. The ADC in the BBB processed the filtered ECG signal at a rate of 350 Hz before transmitting it to the computational node.

4.3.2 Improving the skin to Device Coupling for PPG

In PPG measurements, canine fur becomes a barrier by reflecting most of the incident light applied for measurement. To penetrate through the hair, we incorporated light guides that utilize internal reflection to homogenize non-uniform light sources and optical fibers for efficient optical coupling to the skin (Figure 4.2). In our system, we used a light emitting diode (LED) working at 850 nm to shine light into the tissue and a COTS monolithic photodetector and transimpedance amplifier (TSL12T) to detect the modulated light exiting the tissue (Figure

4.4). We included a silicone structure to hold the light guides at a fixed distance and orientation with respect to each other. Silicone also acts as an optical barrier to prevent direct light from the LED to reach the photodetector directly.

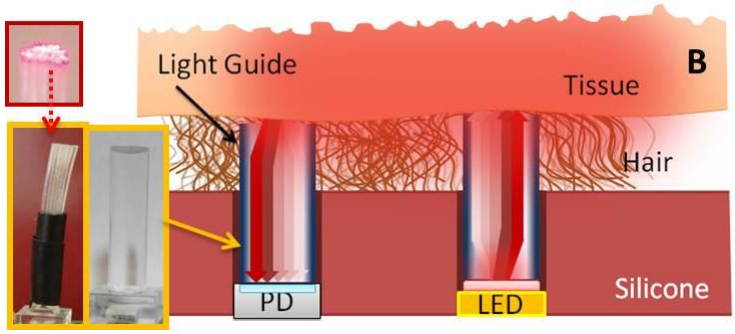

Figure 4.4 Operation of PPG sensor. Light guides in the photodetector and LED for tissue coupling. On the lower left, a close look of the light pipe and also optic fiber bundle for further enhancement. Silicone is used for structural support and to block direct light transmission from LED to photodetector

In order to reduce the 1/f noise, we set the LED pulse rate at 500Hz; the processing unit generated these pulses while the operational amplifier sourced the current. A sample and hold circuit demodulated the photon populations received by the photodiode to extract the arterial pulsation signal. Subsequently, we removed the DC component before amplifying and high-frequency filtering the signal. We added back the DC component before sending it to the microcontroller's ADC (Figure 4.5). Finally, a laptop node analyzed and stored the data received from the processing unit.

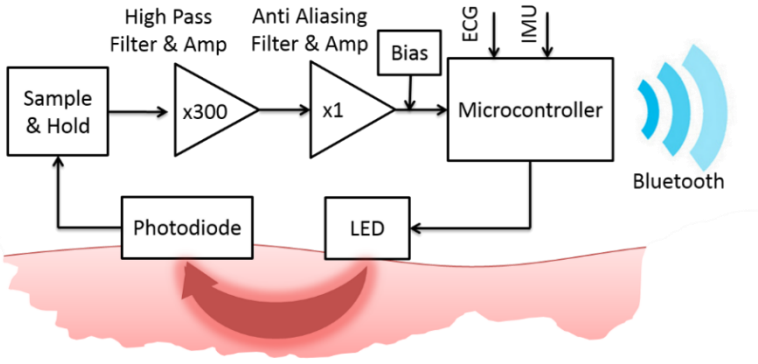

Figure 4.5 System block diagram of the various parts of the wireless PPG sensor.

4.4 Experimental Results

In order to validate the performance of the ECG and PPG systems we performed three separate characterizations. The first characterization (see section 4.4.1) involved the in vitro and in vivo impedance analysis of the two types of electrodes through electrochemical impedance spectroscopy (EIS) to demonstrate that a) the surface modification by coating the metal electrodes with PEDOT:PSS can reduce the impedance of the tissue-electrode interface, and b) to compare the quality of the tissue coupling using these electrodes and commercially available ECG electrodes used in clinics. The second characterization (see section 4.4.2), involved the in vivo evaluation of the ECG system performance by a) comparing the accuracy to a Holter monitor routinely used in veterinary clinics and b) performing recordings while dogs were performing activities with various intensities to show the reliability of the ECG under various conditions. Lastly, we present the effectiveness of the light guides to enable PPG measurements in a hairy chest region (see 4.4.3).

4.4.1 In vitro and in vivo electrode impedance characterization

We performed the electrode impedance characterization in two separate experiments. The first experiment involved in vitro EIS to reveal the effect of the number of gold spring-loaded pins and the electropolymerization with PEDOT:PSS on the overall impedance of the comb-shaped array electrode. For the second experiment we used *in vivo* EIS to compare the electrical properties of the tissue-electrode interface formed with patch electrodes, pointed style electrodes, and spring-loaded pin arrays under various interface conditions (i.e., presence of hair and application of an electrolyte gel). For each condition we iteratively fitted the obtained EIS graphs to the equivalent circuit model of the tissue-electrode interface.

4.4.1.1 In vitro electrode impedance characterization

We characterized the electrodes in a phosphate buffered saline solution (0.8% NaCl) with a 4.8×4.8 mm2 silver-silver chloride (Ag|AgCl) reference electrode using a GAMRY Reference

600 Potentiostat to determine and compare the impedance properties and evaluate the success of the electro-polymerization process. EIS recorded the complex impedance to characterize the charge transport mechanism. We used a 10 mV AC voltage (typical EIS small-signal value) and zero bias voltage as the input signal and recorded the impedance between 10 Hz and 10 kHz at 10-20 discrete frequencies per decade. We analyzed the impedance of a single pointed style electrode and multiple pin comb-shaped electrode arrays to investigate the effect of the number of pins in the array.

4.4.1.1.1 Results

Figure 4.6 shows the in vitro EIS characterization of the stainless steel pointed style electrode and spring-loaded gold-coated pin arrays. The stainless steel electrode and the spring-loaded six pin array presented a very similar impedance response. The impedance improved as the number of pins increased. The increment in the number of pins decreased the overall electrode impedance since the same type and size of pins are connected in parallel. However, it became relatively more difficult to penetrate through the hair as the number of pins increased.

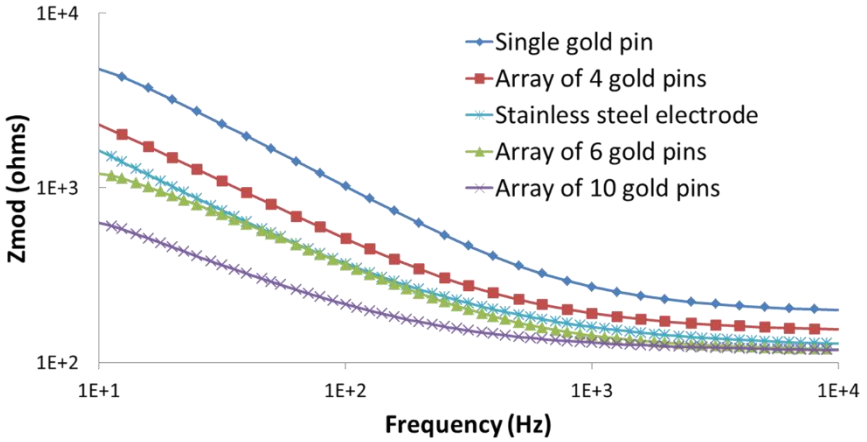

Figure 4.6 In vitro EIS measurements in four different gold spring-loaded pin arrays with one, four, six and ten pins.

EIS results in Figure 4.7 show the effect of PEDOT:PSS coating over gold coated spring-loaded pin arrays with various numbers of pins. The improvement in the impedance

with PEDOT:PSS coating was more pronounced at lower frequencies typical for HRV measurements. At 10 Hz, the single pin impedance decreased from 4.8kΩ to 670Ω and in the four pin array from 2.3kΩ to 300Ω, both showing a similar improvement factor of 7.2 and 7.6 respectively. These results showed a greater reduction effect in the impedance by PEDOT:PSS coating of the electrodes than by increasing the number of pins.

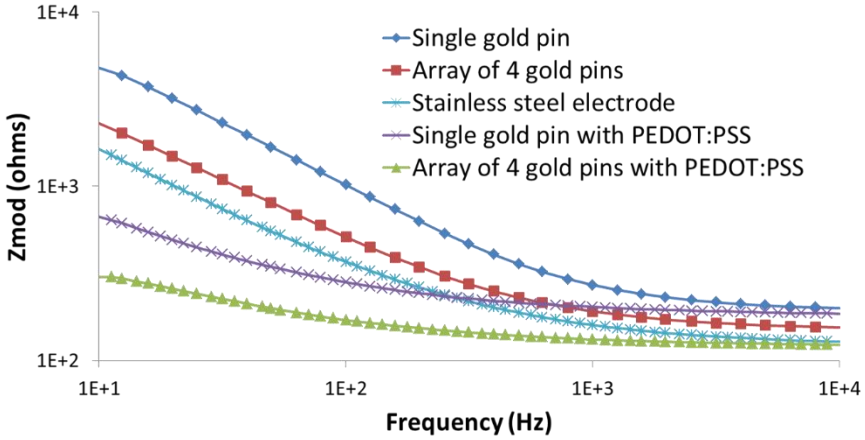

Figure 4.7 In vitro EIS measurements showing the effect of electropolymerization with PEDOT:PSS in a single gold plated spring-loaded pin, and a four pin array.

4.4.1.2 In vivo electrode impedance characterization

To characterize the skin-to-electrode interface impedance, we performed EIS measurements with two electrodes on the chest of Dog 1: the electrode under study (working electrode) and a reference/counter electrode separated by 15 cm. The reference/counter electrode for all these characterization measurements was an Ag|AgCl electrode with gel placed in a clipped region and the working electrode was switched between the traditional patch electrodes used in clinical ECG settings, pointed style stainless steel electrodes and the spring-loaded comb-shaped array. These tests included the effect of placing patch electrodes directly on the skin in a shaved area and over the fur and application of gel electrolyte to the regions where the stainless steel electrodes and PEDOT:PSS coated spring-loaded pin arrays were in contact with the skin.

4.4.1.2.1 Results

Figure 4.8 shows the impedance response of the pointed style stainless steel electrode and the array of four spring-loaded pins coated with PEDOT:PSS when placed on a hairy region of the dog without and after applying electrolyte gel. The same figure also contains the impedance of the traditional ECG patch electrode applied to a shaved spot on the dog. As in the in vitro experiments, we obtained lower impedance when using the four spring-loaded pins coated with PEDOT:PSS than the pointed style stainless steel electrodes. These electrodes both provided better impedance with respect to the traditional ECG patches on hairy regions. When we applied electrolyte gel to stainless steel and gold electrodes in hairy regions, this impedance further improved to a level similar to the patch electrodes applied on a shaved region during clinical monitoring.

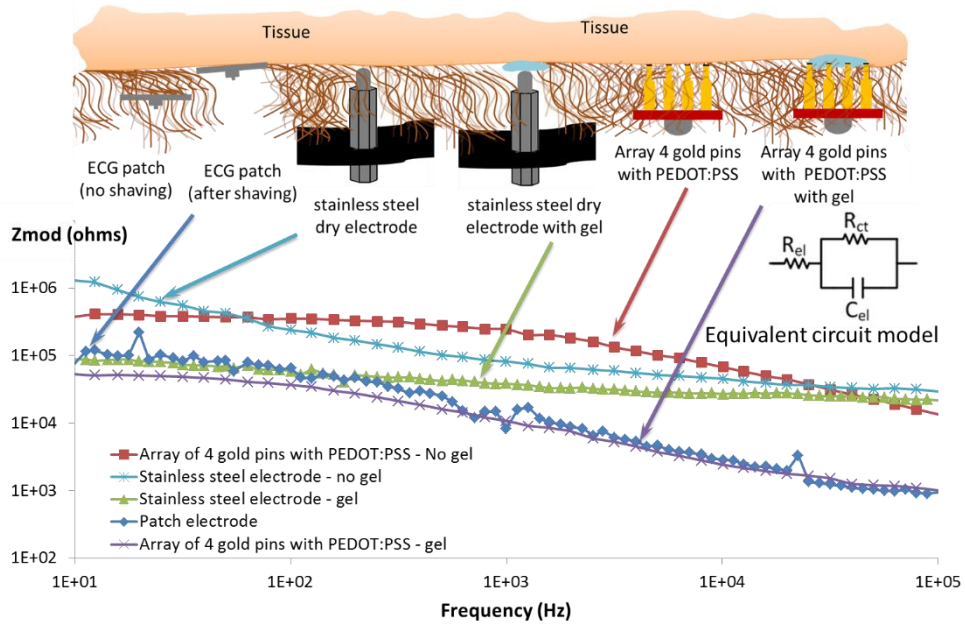

Figure 4.8 In vivo EIS measurements at the patch, stainless steel, and gold spring-loaded array interface with tissue.

The equivalent model of the impedance consisted of a resistance (R_{el}) in series with a parallel combination of a resistance (R_{ct}) and a capacitance (C_{el}) (Figure 4.8). R_{el} represents

the cumulative effect of electrodes, tissue, and additional resistances caused by the cabling and connections. The lower it is, the more voltage is transferred from the electrodes to the amplifiers and less thermal noise is generated. C_{el} is the electrode capacitance formed at the electrode–tissue interface and lower values of C_{el} results in less distortion of the signal. The charge transfer resistance R_{ct} can be interpreted as the ability of the interface to transfer a charge from the tissue to the electrodes. A low R_{ct} value is required to improve the recording performance. We iteratively fitted the obtained EIS graphs to the equivalent circuit model parameters of the tissue-electrode interface for further analysis by using the EIS300 software (Gamry Instruments, PA). We refined the parameter estimates until we obtained a 10% maximum difference between the measured and calculated impedance. Table 4.2 shows the fitted element values. We observed lower R_{el} and C_{el} in the four gold pins with PEDOT:PSS than in the stainless steel electrode under dry conditions. When we added the electrolyte gel R_{ct} and R_{el} decreased and the C_{el} increased due to an increase in the permittivity (ε).

Table 4.2 Equivalent circuit model parameters

Type electrode	Gel	C_{el}(nF)	R_{el}(kΩ)	R_{ct}(kΩ)
Stainless Steel	-	7	41	891
Stainless Steel	+	15	27	40
4 gold pin PEDOT:PSS	-	0.5	10	298
4 gold pin PEDOT:PSS	+	15	2	39
Patch with solid gel	+	17	0.6	23

4.4.2 In vivo ECG system validation with Holter monitor during different intensity activities

We conducted the in vivo ECG system performance characterization in two independent experiments. The first experiment involved the simultaneous recording of our system and a Holter device (Trillium 5000, Forest Medical Inc) that is routinely used in the NCSU CVM Cardiology Service, to assess accuracy. The second experiment, involved performing ECG on the recruited five dogs while performing activities with different intensities to assess functionality and utility of the system. In both experiments we included an IMU to characterize

activity intensity information. For one of the dogs in the later experiment we performed additional analysis on the ECG and accelerometer data to evaluate the use of the accelerometer to detect respiration and heart contraction in addition to activity intensity.

4.4.2.1 System validation with Holter monitor

In order to validate the system accuracy, we performed two independent ECG measurements simultaneously on Dog 2 over a 12 minute period involving two minutes of walking and ten minutes of resting in standing and lying postures. A Cardiac Holter Diagnostic Service Technician installed the Holter recorder on the dog following standard protocol. The protocol involved attaching the five leads of the Holter device and the ECG surface electrodes with incorporated electrolyte gel over shaved spots, and placing our ECG system with three comb-shaped array electrodes with six pins mounted on a chest-strap behind the front legs along with an IMU in close proximity (Figure 4.2). To facilitate the contact of the pins with the skin, we gently wiggled the electrodes and applied some electrolyte gel. We then wrapped both sets of electrodes (patch and comb-shape array) with elastic adhesive tape. The wireless wearable monitoring system transmitted the data coming from our ECG system and IMU to the base unit for further processing and storage.

The Trillium 5000 had only on-board recording capability requiring data to be transferred to computer at the end of the test for offline processing. The proprietary software of this Holter recorder provided the detected QRS complex peak and the R-R interval series. The QRS complex is the most prominent deflection in the ECG waveform and is associated to the electrical activity of the ventricular contraction. In order to compare both systems, we implemented a QRS detector based on an existing algorithm [57]. We filtered the ECG signal with a differentiator (notch filter) $y_o\ (n)=x(n)-x(n-4)$ and a digital low pass filter $y_1\ (n)=y_o\ (n)+4y_o\ (n-1)+6y_o\ (n-2)+4y_o\ (n-3)+y_o\ (n-4)$. We defined a threshold based on the signal amplitude and detected the crossings of the absolute value of the derivative of the threshold. We determined a QRS complex when four or fewer crossings occurred in a 200 ms window, and otherwise we considered it as noise.

4.4.2.1.1 Results

Figure 4.9 shows representative recordings of the R-R interval series obtained simultaneously from the Holter device and our system with the six pin comb-shape array electrodes during static and dynamic activities. During static activities, the R-R interval series obtained with our system tracks closely the commercially-available system validating the hardware accuracy. Even during dynamic activities, the difference between the two measurements is small. We used the difference signal between the two R-R series to observe the discrepancy between the two measurements. Table 4.3 shows for each activity, the recording duration, the mean R-R interval and the associated number of heart beats recorded with the Holter device. To quantify the difference between the two measurements, the number of beats missed by our system with respect to the Holter, and the mean and standard deviation (SD) of the difference signal are also shown. The mean difference during both activities was below 35 ms, which was less than 10% of the mean R-R interval value. During a two minute walking period, our system missed only 5% of beats compared to the beats recorded by the Holter. In those cases, the measured R-R interval showed an unlikely duration of approximately double the average R-R interval. These events can be easily compensated in the detection algorithm but are included on the table with the purpose of showing the system performance.

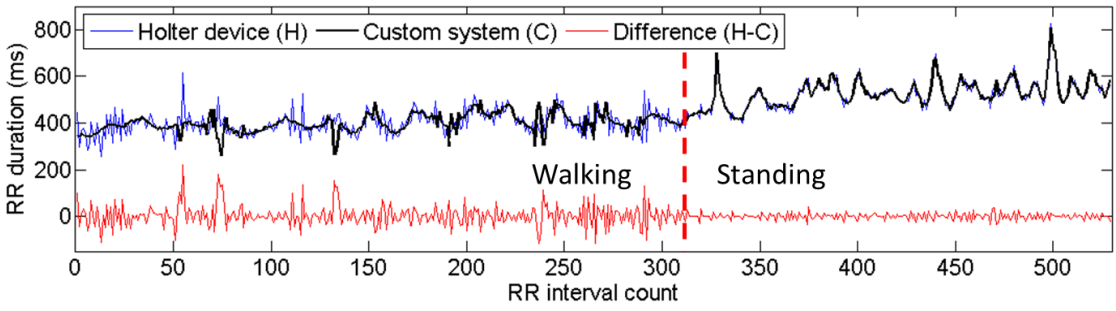

Figure 4.9 Sequence of R-R interval of two minutes of walking followed by standing. Data recorded with the holter device, our system, and the difference between these are shown

Table 4.3 Validation with holter monitor

Dog 2	Duration (s)	Beats recorded Holter	Mean (SD) R-R interval (ms)	Missed beats respect to Holter	Mean (SD) difference (ms)
Static	647	1155	558.10 (52.29)	4 (.35%)	4.75 (5.53)
Walk	122	313	397.16 (39.68)	17 (5.4%)	33.93 (31.94)

Note that, in order to keep the R-R interval sequence from the two systems in sync for comparison we interpolated an additional QRS complex at half the duration of the long R-R interval.

4.4.2.2 System validation during different intensity activities

We evaluated the performance of the system on the five dogs while performing static and walking activities during a 3-6 minute period. Because of the preliminary nature of the analyses we conducted, we did not base sample size on a formal power analysis. Instead, we estimated the smallest number of dogs that would allow us to obtain a representative sample and achieve a preliminary indication of functionality and utility. Table 4.4 summarizes the trials indicating the breed, approximate hair length at the electrode sites, the type of electrodes used, and the duration for the static and dynamic periods.

The system setup was the same as in the previous section. To quantify activity intensity, we used the integral modulus of accelerometer (IMA) obtained from the readings of the 3-axis of the accelerometer on the chest and defined as:

$$IMA = \frac{1}{fsT} \sum_{no}^{no+fsT} |a_x[n] + |a_y[n] + |a_z[n]|$$

IMA has been previously used for assessment of daily physical activity [58], [59]. We filtered the 3-axis acceleration using a band-pass filter (3-20 Hz) before integration. In order to emphasize short events of very high activity, such as naturally occurring shaking behavior, we used a small T equal to two seconds. In this experiment the QRS detector based on [31] was also used to detect QRS complexes and compute the R-R interval.

Table 4.4 Trial details

Dog number	Hair length (cm)	Duration: static/walk (s)	Type electrode
1	2	38 /116	Comb-shape array 6 pins
2	4	130/108	Comb-shape array 6 pins
3	0.5	169/159	Comb-shape array 6 pins
4	5	147/212	Pointed style stainless steel
5	2	84/148	Comb-shape array 6 pins

4.4.2.2.1 Results

Figure 4.10 presents a sample signal recorded from Dog 1 transitioning from a standing position to walking while wearing our system with the six pin comb-shape array electrodes. As seen here, the signal to noise ratio (SNR) at both IMA conditions was high enough to allow for QRS detection.

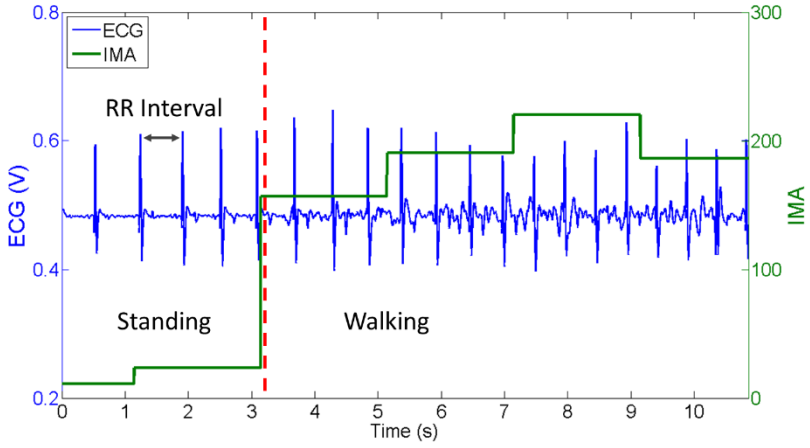

Figure 4.10 Sample ECG signal recorded with an array of six gold spring-loaded pins and IMA computed from 3-axis accelerometer showing different activity intensity for static and walking.

For each trial, we computed the statistical indices sensitivity (SE) and specificity (SP). SE is the ratio of the correctly detected true positive beats (TP), over the number of TP plus undetected false negatives beats (FN). This represents the percentage of the time the algorithm detected a beat relative to how often it should have. SP is the ratio of TP over the number of TP plus falsely detected false positives beats (FP) and represents the percentage of the time the algorithm correctly detected a beat. We used the IMA value to determine a high, or low,

intensity activity for each subject, and for each activity level, we obtained an average R-R interval value. We defined a normal range for R-R interval values when it fell within two standard deviations of the average R-R interval. Short R-R intervals below this range indicated a falsely detected beat (FP) and large R-R interval values exceeding this range indicated an undetected beat (FN). Table 4.5 shows the mean R-R interval duration, SE and SP for each dog and activity. Note that we used the six pin comb-shaped array electrode due to the impedance similarity to the stainless steel. Using this version of the electrodes enabled a more meaningful comparison of tissue adherence due to comparable impedance characteristics. SP was above 97.7 for all dogs during both static and dynamic activities whether SE presented lower values during dynamic respect to static but under all conditions SE was above 95.7.

Table 4.5 Statistical indices for various activity intensity

Dog	Static (Standing, laying down)			Walking		
	Mean R-R interval (SD) ms	SE (TP,FN)	SP (TP,FP)	Mean R-R interval (SD) ms	SE (TP,FN)	SP (TP,FP)
1	675 (110)	1 (43,0)	0.977 (43,1)	580 (114)	0.966 (171,6)	0.983 (171,3)
2	532 (57)	1 (204,0)	1 (204,0)	419 (83)	0.957 (310,14)	1 (310,0)
3	607 (63)	0.996 (285,1)	1 (285,0)	456 (41)	0.976 (320,8)	1 (320,0)
4	587(177)	0.988 (252,3)	0.996 (252,1)	490 (69)	0.988 (429,5)	1 (429,0)
5	612 (80)	0.979 (138,3)	1 (138,0)	522 (80)	0.965 (278,10)	1 (278,0)

To further validate the system, we performed an additional analysis on data corresponding to Dog 3 by observing physiologically well-known cardiovascular events. During this we also benefitted from the seismocardiographic capability of the accelerometer to detect respiration and heart contraction, beyond recognizing various activity intensities through IMA. The aim of this analysis was to observe: a) an elevated heart rate (shorter R-R interval) during dynamic activities with respect to static postures, b) a cyclic response in the R-R interval corresponding to the respiratory sinus arrhythmia, and c) positive correlation between the respiratory modulation and heart rate obtained from the ECG and

seismocardiography measurements. Top part of Figure 4.11 contains the R-R intervals detected with the comb-shaped array electrode and the accelerometer output. The transition between a static to a dynamic activity can be observed on the amplitude and variance of accelerometer. The transition to a higher activity intensity shows a reduction of the R-R interval indicating elevated heart rate. In the right bottom of Figure 4.11, the zoomed-in view of the accelerometer shows an almost sinusoidal modulation at the respiration frequency that illustrates how the IMU registered the chest movement associated with respiration. It can be seen that the cyclic oscillation in the accelerometer signal matches with the respiratory sinus arrhythmia reflected in the ECG based R-R interval. In the left bottom of Figure 4.11, zooming further, shows cardiac contraction on the accelerometer matching the beats recorded with our ECG electrodes when the dog was not moving. The two examples indicate and validate the successful deployment of our electrode system to monitor modulations in HR and HRV on moving dogs and show potential for the use of accelerometer and ECG for multimodal information fusion to achieve robust heart rate monitoring.

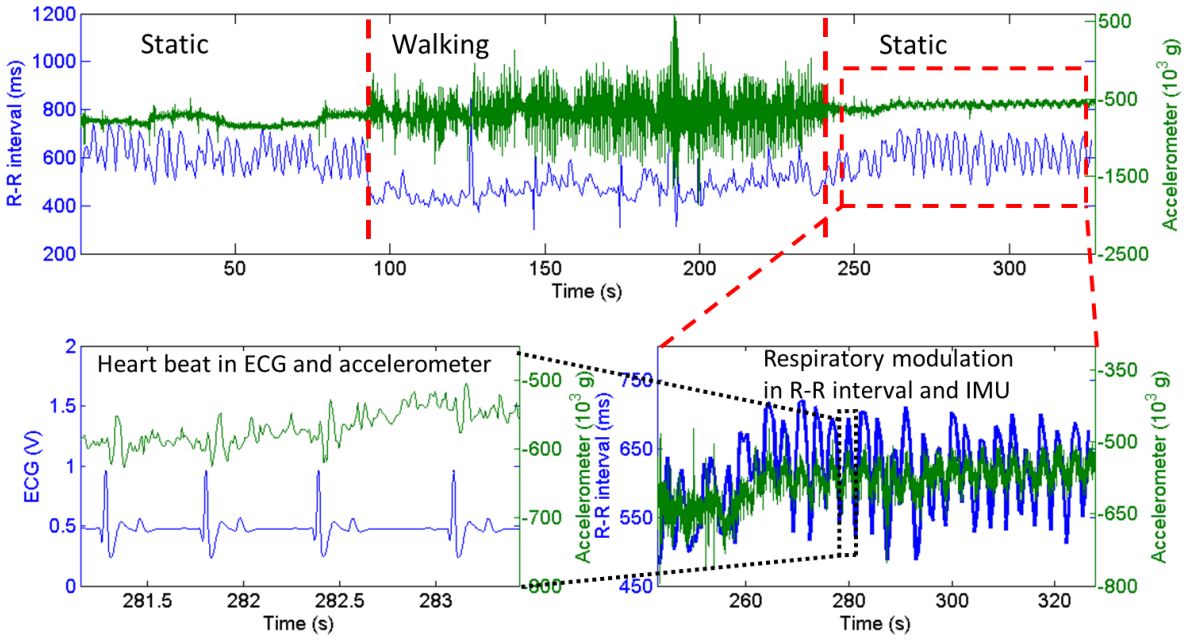

Figure 4.11 (Top) R-R interval series and x-axis of accelerometer during static and walking conditions. (Right bottom) Zooming into the data shows the overlapping respiratory sinus arrhythmia on the ECG based R-R interval sequence and the respiratory modulation from the chest accelerometer. (Left bottom) overlapping cardiac beats and contractions registered with ECG electrodes and chest accelerometer

4.4.3 In vivo PPG system performance

The performance characterization of the custom PPG system and the evaluation of the efficacy of the incorporation of lights guides in improving the tissue-device interface involved performing PPG recordings on a chest in the presence of hair.

We evaluated the performance of the system on Dog 1 during rest. We placed the PPG sensor equipped with light guides in the LED and photodetector against the skin in an unshaved region and held in place by the chest strap (Figure 4.2). We also used the PPG signal to extract the respiration rate from the respiration related modulation.

4.4.3.1 Results

Figure 4.12 shows a sample PPG signal that demonstrates the heart beat and respiration related modulations on the dog during rest. The signal significantly degraded during higher intensity activities. Although the light guide method was effective in static postures and can be used for self-calibration and PTT tracking, more studies are required for deployment during dynamic activities.

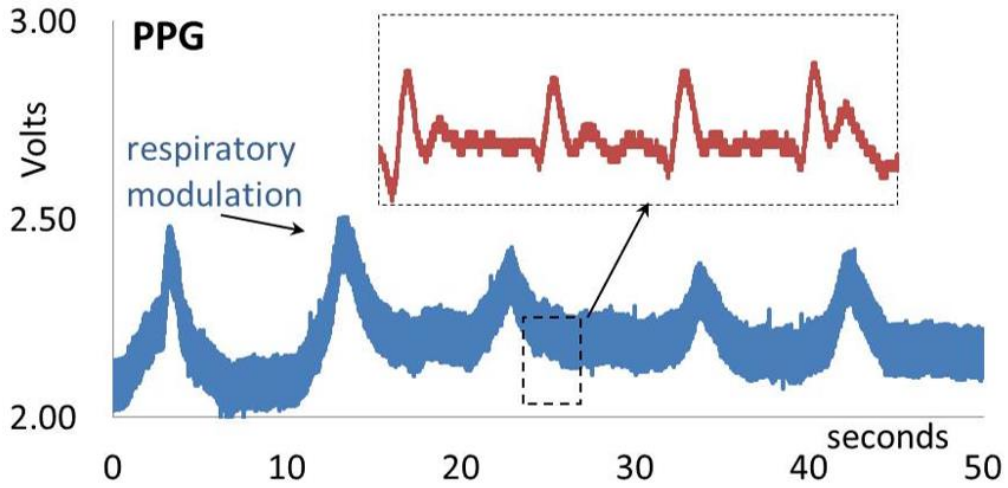

Figure 4.12 Sample PPG signal recorded on a resting Labrador Retriever demonstrates both heart beat and respiration related modulations.

4.5 Discussion and Conclusion

The traditional ECG electrode patches used in veterinary clinics require removal of dense layers of hair and are prone to detachment when the animal is on the move, thereby limiting their use during daily life activities. In this article, we studied the use of pointed style electrodes traditionally used in training collars and comb-shaped spring-loaded pin arrays to overcome these limitations. We performed *in vitro* EIS characterization of the electrodes showing that the impedance of the stainless steel pointed style electrodes is comparable to the comb-shaped electrode containing six spring-loaded pins. We showed that the impedances below the stainless steel electrode can be achieved by increasing the number of pins or by coating the pins with PEDOT:PSS.

We performed *in vivo* EIS measurements to compare the electrical properties of the tissue-electrode interface formed with patch electrodes, pointed style electrodes, and spring-loaded pin arrays under various interface conditions (i.e., presence of hair and application of an electrolyte gel). These *in vivo* experiments validated the improvement in the impedance with the use of spring-loaded PEDOT:PSS coated pin arrays in comparison to the pointed style stainless steel electrodes. These electrodes provided ECG signals without requiring the removal of hair and provided similar impedance to the traditional clinical ECG electrode patches attached to the shaved skin. We tested the electrodes by observing key HRV signatures on five dogs of various size, breed, and amount of hair.

We also presented our preliminary results in improving the tissue-device interface of a custom built PPG system intended for self-calibration, PTT measurements, and in the long term for continuous recording of HRV when ECG recordings are noisy. We validated the efficacy of incorporation of light guides and optical fibers to light sources and detectors for an improved optical coupling with the skin at new body locations other than the ear or tongue.

The next stage of this research involves the collection of physiological data and behavioral information during extended periods of typical daily life activities and perform multimodal information fusion to achieve robust heart rate monitoring in wearable form factors. The combination of physiological signals and canine body language recognition can

be used to provide more objective interpretation of canine emotional response to external stimuli and environments.

In conclusion, we have demonstrated a robust and successful heart rate monitoring system for dogs during low and moderate activity intensity enabled by the hardware improvements performed at the electrodes. This physiological data would enable handlers to monitor dog welfare and enhance the understanding of its emotional response to certain conditions.

Chapter 5. Towards a wearable system for continuous monitoring of sniffing and panting in dogs

5.1 Introduction

Dogs are widely and successfully used as olfactory detectors for explosives, narcotic drugs, and cadavers. While instrumental detection devices continue to improve to carry out these functions, these "electronic sniffing devices" suffer from lack of efficient sampling, poor selectivity in the presence of interfering odors, and limited mobility and tracking ability with respect to detector dogs [3].

For a dog, the odor detection process starts by air scenting until an odor plume is identified. Once a dog "enters" an odor plume, it begins to sniff in order to direct the odor onto the olfactory epithelium [60]. The reliability of dogs as detectors does not depend on the dog performance alone, but also on the dog handler's alertness and skill in interpreting the breathing patterns, and behavioral signals during the nose work. Sniffing and panting are two of the primary perceptually relevant behaviors monitored by handlers. Sniffing, on the one hand, is a behavior characterized by the quick alternating sequence of inspirations and expirations through the nose to increase the fraction of inspired airflow that reaches the olfactory region to actively sample odors [61]. This sequence of multiple sniffs, referred to as bouts, lasts from approximately 0.5s up to several seconds. The sniffing rate was reported for 7 dogs of different breeds (with a weight range between 6-50 kg) ranging from 4 to 7 Hz, and

was independent of body size [61]–[63]. Panting, on the other hand, is the characteristic rapid respiration through the open mouth to regulate body temperature by increasing evaporative cooling by the respiratory tract. The total number of pants per minute has been measured to be generally larger than the sniffing rate but the period of each sniff is smaller than the pants as a result of shorter sniffing bouts [64]–[66]. The panting rate depends on the body size [67]. It was also shown that the efficiency of dogs' olfactory work diminishes after strenuous physical activity due to an increase in panting behavior for body cooling [66]. Panting will cause a decrease in sniffing consequently reducing detection efficiency as these two behaviors cannot be performed simultaneously.

The sniffing behavior is a way that a dog communicates information, such as the presence of an olfactory gradient or the presence of complex odors or olfactory success (i.e. detection). The amount of time a dog spends panting during olfactory detection can be used as an indicator of dog detection efficacy since that would be an indicative of less time spent sniffing. The quantification of respiratory behavior patterns occurring during olfactory detection activity suggest that an electronic sensor system may prove to be a valuable interface between the dog and handler. Such a system would augment handlers' abilities to monitor and interpret their dogs' respiratory behavior more precisely, thereby improving the accuracy of the olfactory detection work.

Monitoring of respiratory behaviors has been attempted outside the laboratory environments during unrestricted activities. In [63] a thermistor fastened inside one of the nostrils measured airflow during various breathing patterns. In the study, only one of six subjects tolerated the instrumentation. Less invasive instrumentation was used by fastening microphones above the dog nose [62] and in the nose region of a custom designed muzzle [68]. In [62], sniffs could not be determined on grass because of noise caused by the dog's nose touching the grass. An omnidirectional microphone was used in [68] that is not immune to ambient noise. In addition, an acclimation period was required to allow the dog to habituate to wearing the custom muzzle.

Accordingly, in this paper we present our preliminary efforts towards a wireless wearable system for continuous auscultation of respiratory behavior by recording internal sounds at the neck or chest of the dog by means of a commercially available capacitive contact microphone based electronic stethoscope with amplified sound output. We attached two of these e-stethoscopes to a dog collar and chest strap to achieve higher immunity to ambient noise and to record pharynx and lung sounds separately Figure 5.1. For proof-of-concept, we present an empirical analysis of the data collected with this system on two dogs of different breeds. These contributions represent early steps towards a novel real-time olfactory detection monitoring system that will provide decision support for handlers in the field. Providing handlers with real time quantitative information about the occurrence of relevant behaviors could accelerate the training process, increase objectivity and enhance accurate decisions during critical missions.

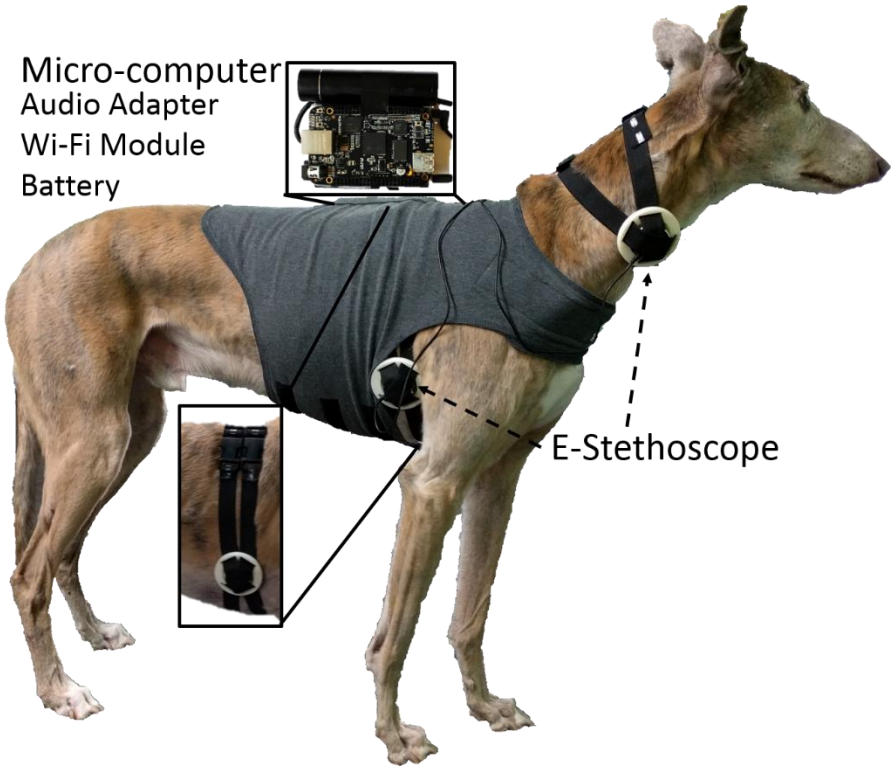

Figure 5.1 The two e-stethoscopes based system worn by a Greyhound

5.2 System description and experimental setup

Our auditory sensing e-stethescope (*One® by Thinklabs Medical LLC*) is used for recording the sound of airflow over the airways. This device has recently became commercially available to improve the auscultatory sensitivity of medical doctors. The principle of operation is similar to conventional stethoscopes but instead of transmitting sound via pressure changes produced by a vibrating diaphragm, the electromagnetic diaphragm of *One®* constitutes one of the plates of a capacitor. A high voltage is applied across the diaphragm and a metal plate behind it to create an electric field between these capacitive plates. The vibration of the diaphragm changes the distance between the capacitor plates, thereby changing the electric field intensity that is sensed as a change in voltage to generate the analog audio signal. The device includes an output stage with configurable amplification and filtering with a maximum cut-off frequency of 2 kHz. Through a USB sound adapter, we connect this audio output to a BeagleBone Black (BBB) microcomputer mounted on a commercial dog harness. Both One® sensors are held against the unshaved skin on the neck and chest by a custom 3D printed mounting structure and four silicone backed elastic bands that we shaped as a collar and chest-strap Figure 5.1. The BBB includes a 1 GHz processor, 2GB of on-board flash storage, 512 MB of DDR3 RAM, and runs Ubuntu GNU/Linux. It serves as the control unit receiving commands to start and stop recording the sensor output, and store and transmit the data using a wireless communication link based on the IEEE 802.11 standard for Wi-Fi. We used the Advanced Linux Sound Architecture framework to sample the audio at 16 kHz with 16 bit resolution.

For this preliminary study, we recruited two trained adult male dogs, a Labrador Retriever (age 4, weight 30.7kg) and a Greyhound (age 10, weight 34.1 kg). All animal procedures were approved by the Institutional Animal Care and Use Committee of North Carolina State University. We evaluated the performance of the system on these two dogs while they conducted an olfactory search for food treats, off-leash, in an indoor room. The dogs stayed outside of the room while the experimenter hid between 3 and 6 treats out-of-line-sight around the furniture of the room. Then, the handler entered the room with the dog and gave the cue to start the search. Both dogs were familiar with the task and sniffed the room

intensively until they found all the treats. Sessions were video-recorded for offline labeling of the data. The resulting recordings included different types of behaviors: normal breathing, panting, sniffing and air scenting. The shifting of the stethoscope against the skin, especially during high activity episodes of air scenting, resulted in occasional artifacts in the audio recording. The ultimate goal of this project is to identify features and build classifiers for all of these events, but in the present analysis we have focused on sniffing and panting patterns only. As we expected, our preliminary observation indicates that no additional habituation is required as most dogs trained for olfactory detection are already acclimated to wearing collars and harnesses.

5.3 Feature extraction and classification of respiratory behavior

Fig 2 shows a representative sample recording of a sequence including panting and sniffing events. In this study, we compute the features in the time domain: event duration, event rate (inverse of period), event mean energy, and the number of consecutive events in a row, and evaluate their efficacy for sniffing and panting detection.

In order to compute the features, the audio recordings must be segmented by events of sniffing and panting. For this, we compute short-time energy contour of the signal using an adaptive window size and apply an adaptive threshold over the energy to identify the boundaries of each event to determine the duration of the event. The mean energy is computed between those boundaries where the inverse of the time between one event and the prior event defines the event mean energy and event rate, respectively.

To compute the short-time energy contour, we use a window which slides over the signal while we calculate the root mean square of the signal amplitude inside the window. In each iteration, the window slides 1/5 of the window size resulting in 80% overlap respect to the signal in the previous computation. If the window size is constant, it performs poorly when sweeping over events of different duration. A window size resulting in a good envelope for a

panting event would lead to a distorted envelope for a sniffing event which is not tolerable since the features are computed from the envelope/energy estimation.

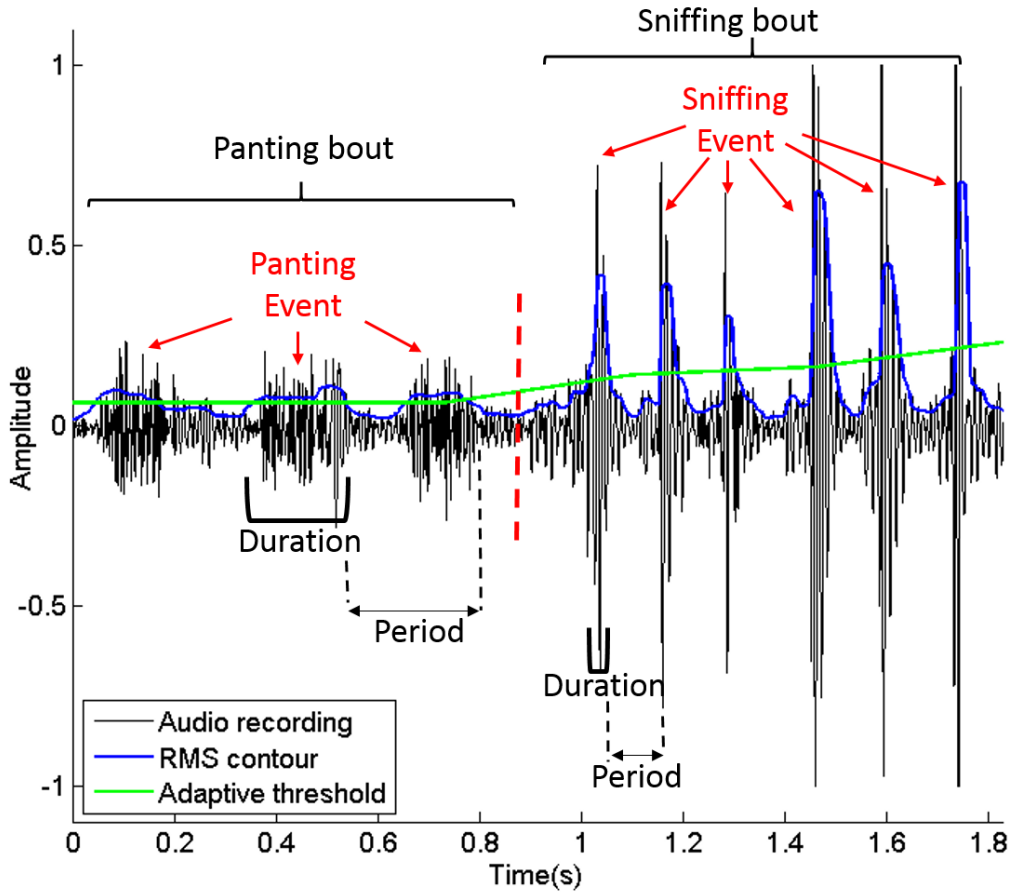

Figure 5.2 A Fragment of a panting bout including three events, and a fragment of a sniffing bout including six events recorded from the chest of a Labrador Retriever. Features are labeled.

To compute the short-time energy contour, we use a window which slides over the signal while we calculate the root mean square of the signal amplitude inside the window. In each iteration, the window slides 1/5 of the window size resulting in 80% overlap respect to the signal in the previous computation. If the window size is constant, it performs poorly when sweeping over events of different duration. A window size resulting in a good envelope for a

panting event would lead to a distorted envelope for a sniffing event which is not tolerable since the features are computed from the envelope/energy estimation.

Our adaptive window algorithm calculates the window size based on a linear function were the independent variable is the energy of the prior window and the dependent variable is the window size. The independent variable bounds range from the typical energy of a pant to a typical energy of a sniff and the bounds in the window size were fixed to the optimized window size for the typical sniff and panting events. Using this approach, the energy estimate was found to be in agreement with the energy curves obtained for both sniffing and panting using two optimized different window sizes. As the adaptive window algorithm sweeps through the signal, it keeps track of the adaptive window sizes used in each run. This is used to interpolate each of the energy calculation points to linearly increasing time scale after the sweep. Due to the large difference in the energy of a sniff and a panting event, the use of an adaptive threshold is required. This threshold is calculated to be the average of the signal energy within a previous window that has double the duration of a typical panting event.

To classify between panting and sniffing using the aforementioned features, we use the C4.5 algorithm implemented in the WEKA Machine learning toolkit [29]. This algorithm builds a decision tree classifier where the leaves are the class labels (sniffing or panting) and the branches are the conjunctions of feature values that lead to those class labels. The C4.5 algorithm uses information gain as splitting criterion choosing the feature values providing the maximum reduction in the uncertainty about the class labels. To evaluate the individual impact of the features, we built a set of decision tree classifiers leaving out a different set of features in each case.

Table 5.1 shows the distributions of events included in the study, and the mean and standard deviation for the different features. Figure 5.3 is the scatter plot of those features where color indicates the type of event (sniffing or panting) and the marker indicates the dog. As can be seen in the figure, the features are very separable between the behaviors across the two dogs. The individual differences among the dogs, in particular, the higher panting rate of the Labrador Retriever and the shorter duration of his sniffs are also noteworthy. The observed

sniffing and panting rates fall between the expected values found in literature and are distinct between the two behaviors. Since the panting frequency is dependent on body size [67] it could result in similar value to the sniffing rate of a different dog. This concern could be easily addressed by calibrating the system for the animal specific frequencies. Also, [61] describes sniffing generally occurring in episodes with a crescendo and decrescendo, which we have also observed in our data. It is shown in Figure 5.3 as a larger dispersion of the mean energy feature. The mean number of consecutive events is included in Table 5.1 and matches what is described in [61]: the panting can occur for prolonged periods of time where sniffing happens in shorter bouts. It is worth noting that the discrepancy in behavior duration is significantly larger for panting since the events occur at approximately half the speed of the sniffs. Although we have data for only two dogs, the agreement with the literature suggests that this measuring technique and the selected features are promising for the detection of these behaviors.

Table 5.1 Distribution of sniffing and panting samples used for study.

	Greyhound		**Labrador Retriever**	
	Panting	***Sniffing***	***Panting***	***Sniffing***
# of Events	125	33	21	150[a]
# of Events Training	39	10	10	52
# of Events Testing	86	23	11	98
Duration (sec) (m,std)	0.153,0.03	0.061,0.025	0.133, 0.03	0.036, 0.012
Rate (Hz) (m,std)	2.48, 0.6	6.52, 0.68	3.39, 0.3	7.10, 1.34
Mean Energy (m,std)	0.16, 0.03	0.46, 0.13	0.09,0.01	0.47, 0.15
Mean # of events per bout	30.2, 15.5	8.2, 2.9	9.5, 0.5	8.7, 3.8

a. Includes sniffing events recorded at the chest and at the neck

The recordings were performed with one e-stethoscope on the neck and the other on the chest. The sniffing features were very similar at both locations so they were combined for the analysis. On the other hand, when panting was recorded at the neck, inspiration and expiration sounds appeared as separate sounds as shown in Figure 5.4. We hypothesize that the sniffing at both sites is similar because the inspiration phase is more intense with respect to expiration during sniffing and some of the air volume that is inspired during the sequence of sniffs would be expired when the episode of sniffs is completed. Since the features of the

panting events recorded at the neck are different, we left those recordings out of the classification analysis in the scope of this paper. This observation was true for both dogs (Figure 5.4).

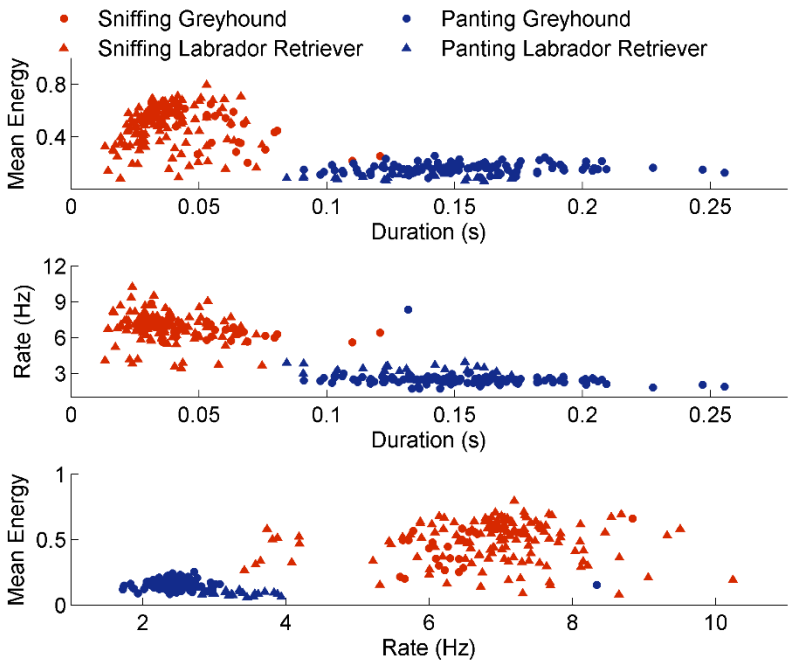

Figure 5.3 Scatter plots showing the distribution of the features for the different behaviors (different colors) and the different dogs (different marker).

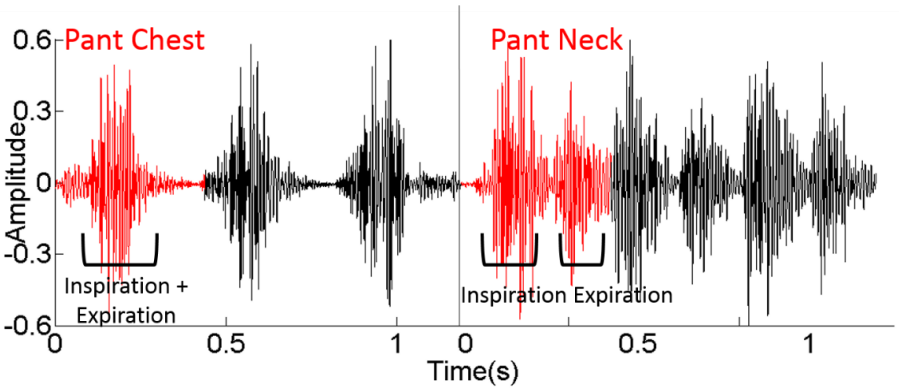

Figure 5.4 Panting recordings at the neck (right) show distinctive inspiration and expiration phases not separable at the chest location (left).

5.4 Classification performance

To evaluate the classification performance, we first removed the sections of the audio recordings corresponding to normal breathing and air scenting episodes and unidentified fragments. Unidentified fragments represented noise, likely due to device shifting against the dog's body, or fragments not correlated with behavior in the video recordings and sections. Then, we manually fragmented the audio recording into the sniffing and panting events. For each of the events, we computed the set of features.

Since the computation of the rate requires the time difference between an event and the prior event, we cannot fully randomize the samples to apply cross-validation. So instead, we split the audio recordings in 1/3 for training and 2/3 for testing. This division was done randomly maintaining the 1:2 ratio in the number of events included for each dog and the number of events of each class. The distribution shown in Table 5.1.

In the training phase, we used the features derived from the 1/3 of the audio recordings reserved for training to build a decision tree classifier using the C4.5 algorithm implementation in Weka [29]. To observe the impact of the features in the classification accuracy, we built different models using different subsets of the features. When duration was included in the set of input features, the C4.5 algorithm built a tree that included the duration features and none of the other features. Therefore, accuracy for any subset of feature types that includes duration is identical to duration only (second row of Table 5.2).

Table 5.2 Auto-segmentation and classifier accuracy

Subset of classifier features	*TP*	*FP*	*FN*
Auto-Segmentation	211	17	7
Model I (Duration only)	193	8	10
Model II (Mean_E & Rate)	204	7	0
Model III (Mean_E only)	205	6	0
Model IV (Rate only)	203	4	4

In the testing phase, we evaluated the accuracy in two steps. First, the accuracy of the segmentation function and second, the accuracy of the classifier when the output of

segmentation function was fed to the model. In the first step, we performed the adaptive segmentation on the 2/3 of the recordings reserved for testing and we compared the results to the manual segmentation to obtain the number of false negatives (FN, the missed events), and false positives (FP, the extra events). Table 5.2 first row shows that out of the 218 events in the testing set, 211 were successfully segmented. For these 211 segments that coincided with the manually labeled segments, we computed the features and ran them through the set of classifiers. We compared the classification results to the manual labels to obtain the accuracies in Table 5.2. The lower accuracy for Model I is due to the segmentation function leading to more variability in event boundaries, therefore in the duration feature, with respect to the manually labeled data used in training showing that a model including a larger subset of features is more robust to the segmentation

5.5 Discussion and Conclusion

In this paper, we present our preliminary efforts towards a wireless wearable system for continuous auscultation of respiratory behavior by recording internal sounds by means of a commercially available electronic stethoscope. We have focused in the identification of features and building classifiers for sniffing and panting, leaving the other behaviors as future work. We presented a computational method to segment the audio recordings into individual sniffing and panting events. Our segmentation method computes short-time energy of the signal using an adaptive window size and applies an adaptive threshold over the energy to identify the boundaries of each event. We compute duration, rate, mean energy and the number of consecutive events in a row and feed these into a decision tree classifier we built using the C4.5 algorithm implementation in Weka. Although we present data for only two dogs as a pilot study, the features obtained from the data reveal differences between sniffing and panting and match findings reported in the literature for these behaviors. This system holds promise for real world scenarios where handlers could detect the olfactory behavior of their working dogs more precisely.

Chapter 6. Efforts towards Automated Training: Computer-delivered reinforcement of sit behavior in dogs

Work presented in this chapter is a preliminary analysis to the publication: J. Majikes, R. Brugarolas, M. Winters, S. Yuschak, S. Mealin, K. Walker, P. Yang, B. Sherman, A. Bozkurt, and D. L. Roberts, "Balancing Noise Sensitivity, Response Latency, and Posture Accuracy for a Computer-Assisted Canine Posture Training System," *Int. J. Hum. Comput. Stud.*, 2016.

6.1 Introduction

For as long as dogs have been domesticated, we have spent time and effort to train them to assist us with a multitude of tasks. In this chapter, we give an overview of the hardware and software used in our autonomous canine training system, which (to the best of our knowledge) is the first training system that does not require human intervention to reinforce postures. We discuss experiments that we conducted which involved teaching dogs to sit, the successful results of which illustrate the efficacy of our system.

Our system consists of a harness outfitted with inertial measurement units (IMUs), which is worn by the dog. Labeled accelerometer and gyroscope data from the IMUs are used to classify postures as quickly and accurately as possible. To close the loop and provide rewards to the dogs for desirable postures, two remotely-operated treat dispensers are triggered based on the output of classification algorithms.

When performing dog training, one of the most effective techniques is positive reinforcement [4], which involves giving the dog something desirable such as food to reward sought-after behavior. Two of the most important qualities in a good dog trainer are consistency and accuracy. The trainer must be consistent when "marking" the behavior, which is the process of indicating the desired behavior via a short sound, such as a click. Marking

multiple behaviors can slow down the learning process, while marking the wrong behavior can teach the dog to do an undesirable action. The trainer must also be accurate, since that sound must occur as close to the exhibition of the behavior as possible; a delay of greater than 0.5 seconds can retard the dog's ability to create the mental connection between the behavior and reward [5].

These two factors directly translate to response latency and noise sensitivity when designing algorithms for capturing and reinforcing behaviors. If the algorithms aggressively filter the data to reduce noise and increase accuracy, they run the risk of missing the critical 0.5 seconds window to reinforce a desirable behavior. On the other hand, algorithms tuned to respond to posture changes more quickly run a risk of being overly sensitive to noise in the sensor data and incorrectly reinforcing the wrong behavior. For example, a dog with "snappy execution" may go from a standing position to a laying down position in 0.7 seconds, briefly passing through a sitting position on the way. If the algorithm to detect the sit posture emphasizes accuracy and applies aggressive filtering of the data, the dog may be laying down before the algorithm detects the sit and triggers reinforcement.

Our experiments consisted of trials with four different dogs, each participating in one, two, or three trials for a total of seven. Over the course of the trials, we evaluated the performance of different data smoothing techniques and classification algorithms. We evaluated a Simple Moving Average filter in comparison to Univariate Spline Smoothing. Additionally, we compared the performance of Decision Trees, Random Forests, and a Threshold-based Decision Rule classifier. After successfully teaching one dog to sit for rewards using the best-performing combination of techniques on the sixth trial, we were able to replicate those results with a different dog on the seventh trial. Below we will provide more details on the training platform, algorithms used, as well as experimental design and results for all seven trials.

6.2 Related Work

The use of wearable inertial measurement units on humans has become increasingly popular recently. These studies use a wide variety of machine learning algorithms to classify behavior

[69]. There has been work in applying similar classification techniques using sensors to various animals [70]–[72], however these studies have primarily focused on monitoring activity levels, not identifying specific activities. Position estimation involves keeping track of canines' exact positions via a GPS receiver. Orientation and velocity are used for dead reckoning when a GPS signal isn't available [73]. There has also been progress made in remote communication with canines including sending commands via audible and tactile signals [74], and receiving feedback via sensors the dog can operate with its nose or mouth [16].

Behavior recognition uses inertial sensors to classify the activity the dog is currently performing. In one study, recognition was done in real-time, however the positioning of the sensors were placed inefficiently for behavior recognition [75]. In prior work, we identified positions for sensors on a canine in order to optimize for posture classification [21]. In other work, accelerometer data was used to identify when and for how long canines exhibited a total of seven static postures and dynamic behaviors, however the recognition was not done in real-time due to the reliance on data that needed to be manually extracted from video, and the need to connect to the sensors to get the data [76]. Our previous work on the evaluation of a machine learning algorithm based on a two-stage cascade classifier used raw sensor data [33] to accurately recognize five static postures and three dynamic behaviors in near real-time [18]. To our knowledge there is no existing work on using computers to automatically capture and reinforce postures for dog training.

6.3 Platform

During experimentation we refined the design of our platform to promote flexibility and efficiency. Both versions involved three-axis accelerometers and three-axis gyroscopes mounted on a commercial off-the-shelf (COTS) harness. Additionally, both versions involved using a laptop for classification and controlling two remotely-operated treat dispensers. Figure 6.1 illustrates the platform and one of our study subjects wearing the custom harness with the sensors.

Version 1: We mounted four IMUs on a COTS "webmaster" harness from RuffWear Inc—on the rump, chest, abdomen, and shoulders. Earlier work helped us identify these

locations as providing good differentiation of the gravity projection over accelerometer axes during different postures. Each IMU connected via a Bluetooth Low Energy link to a Dell Latitude E6410 with an Intel core i5 processor and 4GB of RAM. The laptop was running a Simplified Moving Average filter and Decision Tree classifier (C4.5) using Matlab. This platform was described in more detail in Chapter 3

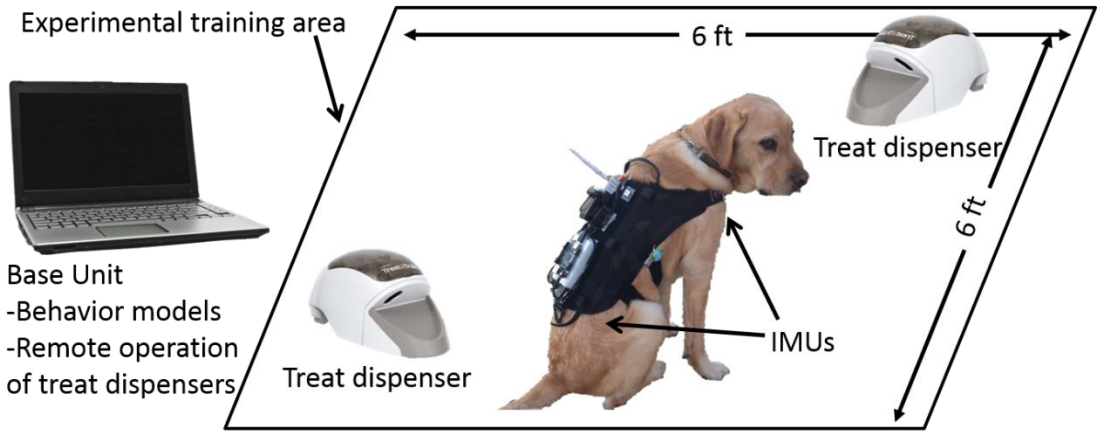

Figure 6.1 Automated canine training system platform and one of the subjects of our study wearing the custom harness outfitter with the sensors

Version 2: In the second design we consolidated data aggregation and communication responsibilities to a BeagleBone Black (BBB) development platform, and transitioned to using 802.11 for communications. Offline analysis indicated two IMUs (on the back and the chest) yielded nearly identical classification accuracy compared to four. The BBB and IMUs were mounted on a COTS lift-load-carry harness from Ray Allen Manufacturing. The BBB sent IMU data to a Macbook Pro with an i5 processor and 8GB of RAM running Univariate Spline Smoothing and either Random Forests or our Threshold-based Decision Rule classifier.

6.4 Experiment Design

Our experiments included trials with four privately-owned dogs: a Labrador Retriever, a Kai Ken, a Retriever Mix, and a Weimaraner Mix. The experiments were conducted over two days

separated by five months. On the first day the first dog ran two trials ($D1^{DT}1$ and $D1^{DT}2$) and the second and third ran one trial ($D2^{DT}$ and $D3^{DT}$ respectively). On the second day the first, second and fourth dog ran one trial ($D1^{TC}$, $D2^{TC}$ and $D4^{RF}$). Note that the abbreviations DT, RF, and TC in the trial names refer to the classification method used (Decision Trees, Random Forests, and Threshold-based Decision Rule classifier respectively, details presented below). All procedures involving dogs followed guidelines set out in the Animal Welfare Act of 1966, and were approved by the Institutional Animal Use and Care Committee.

The procedure included four phases:

1) Collecting labeled posture data and building a classifier
2) Validating the classification model
3) Introducing the dogs to the reward delivery mechanism
4) Capturing and reinforcing the sit posture

After completing all four phases, we conducted an offline analysis to characterize performance.

Phase 1: The first steps in our protocol involved collecting labeled data to train our supervised learning algorithms. We collected labels in real-time by having an experimenter lead the dog through a sequence of three postures: standing, sitting, and eating off the floor or laying down (depending on the trial). While demonstrating these behaviors, someone sat at the keyboard to apply the appropriate labels. Each posture was demonstrated for several seconds, before moving to the next posture. The entire sequence was repeated between five and 10 times depending on the dog. Table 6.1 contains the number of samples of "sit" and "other" behaviors for each of the trials. Note that for $D1^{DT}1$ and $D1^{DT}2$, a single model was built from a labeled data set that was reused for a second trial. To mitigate the effects of noise introduced by human perception delays in applying labels when postures changed, in the first four trials we dropped five samples (0.5 seconds worth of data) for every label change. We didn't continue this practice during the last three trials. Once all of the labeled data had been collected, we built one of the three classification models described below.

Table 6.1 Number of training instances of each posture

Trial	Sit	Other Behavior
$D1^{DT}1$ and $D1^{DT}2$	246	370
$D1^{TC}$	227	238
$D2^{DT}$	121	191
$D2^{TC}$	336	566
$D3^{DT}$	391	705
$D4^{RF}$	235	409

Phase 2: The second phase of the procedure was a "sanity check" to ensure we weren't setting the dogs up for failure. Because we can't erase a dog's memory like we can a machine learning algorithm's, it was imperative we had confidence that our classification model was going to be reasonably accurate. This would avoid completely confusing dogs by having treats dispensed at seemingly random times. This validation of the model consisted of repeating the same sequence of behaviors two to four times, and visually comparing the labels output by the classification algorithm with the postures the dogs were displaying. In all seven cases, we were satisfied with the model's performance on the first try and proceeded to the next phase.

Phase 3: Clicker training is an extremely popular paradigm for dog training [77] whereby a repeatable "click" sound is paired with a "primary reinforcer" [23] (a food reward in our case). The repeatable sound is used to "mark" behaviors, which indicates to the dog that their most recent behavior was desirable and there is a reward on the way. This process of pairing the click with the reward is known as "charging the clicker" and results in the sound becoming a "conditioned secondary reinforcer." Our process involved repeatedly playing a click sound and dispensing a treat 5-10 times from one treat dispenser, then 5-10 times from the other dispenser, and finally 5 times alternating between dispensers. To avoid biases, for dogs that conducted multiple trials, we used a different click sound on the different trials.

Phase 4: The final phase of the protocol involved stepping aside and allowing the algorithm to capture and reinforce the sit behavior. For every instance of data where the algorithm classified a sit posture and the previous instance was not classified as a sit, the algorithm would trigger the click sound to mark the behavior followed by dispensing a treat. The delay from the time the dog offered the sit until the algorithm triggered the click is the latency we were interested in keeping below 0.5 seconds. Alternating treat dispensers were

used to reinforce subsequent sits. Qualitatively, the goal was to observe the dogs focusing on the treat dispensers, sitting, getting up to eat the dispensed treat, and quickly offering another sit. To avoid confusion and stress on the dogs, an experimenter stood in the room with them during this phase. If the dogs appeared confused or stressed at the beginning of the phase, the experimenter would use body language to encourage the dogs to reposition and sit.

Offline Data Analysis: After all of the trials were complete, we examined both log and video data. To extract ground truth from the videos, we had three raters code the videos for the timing of sit behaviors and clicks. The exact time of each sit behavior was defined by the first frame that shows the canine's upper thigh making contact with the ground. The time of the click was determined by the audio track. We averaged the times of all three raters. They also annotated false positives and negatives.

6.5 Algorithms

In order to meet our strict timing requirements, we experimented with two filtering algorithms and three different classification algorithms. We found the more heavily filtered the data was, the more accurate the classification of postures; however, as a consequence of the increased accuracy, increased latencies must be tolerated. Our efforts in exploring these options were targeted toward identifying a fast and accurate approach.

Filtering/Smoothing Methods For our first four trials, we used a Simplified Moving Average (SMA) filter to smooth our data. SMA filters are tunable based on the number of historical samples used to compute the average. The higher the number the more aggressively the data was filtered, resulting in slower response to legitimate changes in raw data. SMA was calculated by $\frac{\sum_{i=0}^{n-1} v_{t-i}}{n}$ where v_t indicates the value at time *t* and *n* was the window size. We experimented with a range of values and found that $n = 11$ yielded the best performance in terms of classification accuracy [18].

For trials five through seven, we used Univariate Spline Smoothing (USS). A USS is a numerical function defined by piecewise polynomials with a single variable, which can smooth data with different intensities. The joint points between the polynomials, known as knots, give

the curve freedom to bend; the greater the number of knots, the more the filtered data will resemble the unfiltered data. USS provides an efficient and flexible smoothing method due to its simplicity to calculate, making it beneficial for real-time data processing. It is also flexible since its smoothing aggressiveness could be controlled by a single parameter. We aimed to reduce outliers to filter out some noise, but did not want to smooth the data too much in order to allow the classifier to maintain a low latency. We did this by applying a very high number of knots, but did not interpolate through every data point. This resulted in the filtered data containing a very low (non-zero) mean squared error; functionally, the outliers were almost eliminated.

Classification Methods: In addition to using different filtering methods, we also examined how three different classification algorithms worked in terms of accuracy as well as response latency. Even though the data from the harness was time-series data, because we were classifying postures and not behaviors we were able to focus on instance-based classification algorithms. We opted to compare Decision Trees (C4.5), Random Forests, and a Threshold-based Decision Rule classifier. We chose these three algorithms because of their speed and simplicity; however, we found great differences in their performance that we will discuss below.

For the first four trials, in combination with the SMA filter, we used the Decision Tree classification algorithm that is part of the Matlab Statistics Toolbox. The Decision Tree classifier splits the data by testing the value of one of the 12 or 24 variables (depending on whether two or four IMUs were used) against some threshold that was learned by splitting the continuous values of the data. In effect, the learning process produces a sequence of contingent thresholds which, when compared against that sequence, lead to a classification with reasonably high accuracy. Figure 3.6 shows the distribution of samples in the x-axis acceleration space providing some insight about these thresholds [33].

For the fifth through seventh trials, we tested the two remaining methods: Random Forests and Threshold-based Decision Rule classification. Random Forest classification is an ensemble method which generates many Decision Trees with a randomly sampled subset of

training data and features. Each decision tree makes a prediction and then all of the decision trees vote for the final prediction. Random Forests run very fast and provide a competitive prediction accuracy when compared to other complex classifiers. There is a trade-off between prediction latency and prediction accuracy. In general, the larger number of trees the more prediction latency and less prediction error. After investigating a range of number of trees, we found that Random Forests had extremely high variance in latency depending on the data and parameters. In general, we found that 11 trees resulted in a good trade-off between latency and accuracy.

For the Threshold-based Decision Rule classifier, we first split the data into groups based on each of the posture labels. For each group we computed the mean and standard deviation of each of the 12 attributes (three linear and three rotational accelerations for two locations). For each of the attributes, *a*, and labels, *l*, we computed a range as $[\mu_{al} \pm c \cdot \sigma_{al}]$ where *c* is a constant. We then identified ranges of attributes such that the range of the sit label was unique from the other labels. We found that setting $c = 1.5$ provided two unique ranges for the first trial and three unique ranges for the second. For the two trials, the axis parallel to the dog's spine on the IMU mounted to its back was common as a unique range classifier for both trials.

6.6 Results

Here we discuss three objective measures of performance: classification accuracy on a per-instance basis, classification accuracy on a per-posture basis, and response latency. We will also discuss the overall performance qualitatively.

Data for classification accuracy of instances came from the human-applied labels on the training data, using 10-fold cross-validation. For accuracy on posture recognition and response latency, we report coded data from the subsequent video analysis that represents the average of three independent coders. Finally, the qualitative analysis combines our impressions during the trials, performance numbers, as well as video analysis.

6.6.1 Instance-based accuracy

To evaluate the performance of the classification techniques we performed 10-fold cross-validation over the human-labeled training data sets. Each instance corresponds to a sensor reading where at 10 Hz a three-second sit would produce 30 instances. By dividing the data set based on instances, the data remained IID and therefore cross-validation is a true test of classification accuracy. As we expected, with the aggressive moving average filter and decision tree classification scheme, we found the instance-based accuracy to be extremely high. Table 6.2 contains the confusion matrices for all seven trials. Please note that $D1^{DT}1$ row also applies for $D1^{DT}2$ since the same model was used for both dog training trials. When using USS with Random Forest or Threshold-based Decision Rules, we found the accuracy dropped slightly. The cross validation results in Table 6.4 rows $D1^{TC}$ and $D2^{TC}$ show that the accuracy for the Threshold Classifier was 89.69% and 92.61%. The classifier was more sensitive to noise which lowered the accuracy. Overall, all of these accuracies were encouraging.

Table 6.2 Instance-based confusion matrices for six trials

Dog Number and Classification/Filtering type	Sitting Posture		Other Posture	
	True Sit	False Sit	False Other	True Other
$D1^{DT}1$ SMA	246	0	0	370
$D1^{TC}$ USS	244	1	92	565
$D2^{DT}$ SMA	122	1	1	190
$D2^{TC}$ USS	199	13	28	315
$D3^{DT}$ SMA	391	0	0	705
$D4^{RF}$ USS	233	0	2	409

6.6.2 Posture Accuracy

Looking at instance-based accuracy of the classification schemes only tells part of the story. When capturing behaviors offered by dogs, it was very important that the false negative and false positive rate be extremely low, especially early on in the training process.

Table 6.3 contains the confusion matrices for six of the seven trials ($D2^{DT}$ had too few data points to be meaningful) and for the two successful trials ($D1^{TC}$ and $D2^{TC}$) zero false

negatives occurred. In these two trials we repeatedly observed a second true positive within 0.5 seconds of the first and in those cases we considered that as one. This phenomenon can be explained by the fact that these low latency algorithms captured the sit very early in the performance of the behavior, while the dog was still transitioning between a different posture to a sit. Note that we predict every 0.1 second and the dog may take longer than that to transition between two postures. This explains why the more noise sensitive classifiers had more variability associated with this transition, and oscillated between two predictions until the dog completed the sit.

Table 6.3 Confusion matrices describing the accuracy of detecting offered sit postures. Note we excluded $D2^{TC}$ from this table due to having too few data points

Dog and Algorithm	True Positive	False Positive	False Negative
$D1^{DT}1$	6	0	2
$D1^{DT}2$	5	0	3
$D1^{TC}$	46	0	0
$D2^{TC}$	23	6	0
$D3^{DT}$	12	0	1
$D4^{RF}$	16	1	8

The elapsed time for each trial depended on many factors including the dog's willingness to participate. $D1^{DT}1$ lasted 18:15, $D1^{DT}2$ lasted 10:39, $D1^{TC}$ lasted 15:18, $D2^{TC}$ was 17:20, $D3^{DT}$ lasted 12.14, and $D4^{RF}$ lasted 18:31.

6.6.3 Response Latency

The data discussed above was useful for understanding the performance of the filtering and classification approaches from a machine learning perspective; however, accuracies and confusion matrices were only half of the story when it comes to measuring performance of an autonomous canine training system. Equally, if not more, important than accuracy was response latency. Timings assigned by our coders were measured by frame numbers in the video—a resolution of $\frac{1}{30}$ th of a second.

Table 6.4 shows a summary of the computed classification accuracies obtained from 10-fold cross-validation on the training set (Phase 1) and the reinforcement delay obtained from video analysis (Phase 4). The last column corresponds to the reward delay between when the sit behavior was offered and when the sound marked the posture. There were a few interesting things to note. First, as we expected, the trials with the highest classification accuracy that resulted from aggressive SMA filtering and the use of Decision Trees, also had the highest response latencies. Second, note that the reinforcement delay in row $D2^{TC}$ was negative. Because our coders marked the time when the dog's upper-thigh touched the ground, it was possible for the algorithm to detect the posture change prior to it being completed, i.e., when the dog was half-way between standing and sitting. If we use the posture accuracies from Table 6.3 (phase 4), TC would result in the best classifier in terms of classification accuracy and latency. As was mentioned, the instance-based accuracy is a true test of classification accuracy. The longer latencies we experienced when running those models in Phase 4 were associated with lower instance-based accuracies, and that is the tradeoff we're referring to.

Table 6.4 Classification accuracy (on filtered data), response latency, and response standard deviation for six of seven trials. $D2^{DT}$'s data was excluded for having too few samples

	Instance based Accuracy (%)	Latency Mean (StdDev)
$D1^{DT}1$	100	1.492±(0.253)
$D1^{DT}2$	100	1.126±(0.251)
$D1^{TC}$	89.69	0.290±(0.251)
$D2^{TC}$	92.61	-0.141±(0.108)
$D3^{DT}$	100	3.863±(1.690)
$D4^{RF}$	99.69	1.233±(2.398)

Next, note the large value of the standard deviation measured for $D4^{RF}$. This trial was the only one for which we attempted to use the Random Forest classifier. In our testing outside of the training environment, we found that the Random Forest classifier was very sensitive to the parameter settings and the data set. For example, reducing the number of trees in the forest might actually increase the response latency by an order of magnitude. Further, we found that

the Random Forest classifier often was able to detect one class label with far lower latency than the other.

Lastly, note that the best response latency occurred when using the Threshold-based Decision rules with USS. This can be observed by examining rows $D1^{TC}$ and $D2^{TC}$ in Table 6.4. In those two cases, we successfully achieved response latency less than the 0.5 second criteria identified in the literature on dog training. As we will see in the next section, it was these two trials where we successfully captured the sit behavior from the dog.

6.6.4 Qualitative Results

The final analysis of our system's performance was in the form of characterizing the dogs' behaviors. As stated above, ideally the dogs would have rapidly learned to sit, receive a treat (which required breaking the sit posture), and repeat immediately. We found this to be the case in two of the seven trials, ($D1^{TC}$ and $D2^{TC}$) which we consider a big success.

Figure 6.2 illustrates the rate at which each of the dogs offered sits. Ideally we would see super-linear growth in these data; however, faster linear growth was better than slower linear growth as it indicates the dog was more enthusiastically offering the desired behavior. At the beginning of phase 4 in our procedure, when the dogs appeared confused or stressed the experimenter used body language to encourage the dog to reposition the sit. Figure 6.2 also indicates, for the trials where the experimenter was involved, when he stopped giving body language feedback to the dog. Note that the data in Figure 6.2 seems to indicate that $D4^{RF}$ performed similarly to $D2^{TC}$; however, what wasn't represented in the figure was that the dog was bored, laying down by the time the high-latency reinforcement occurred, and required encouragement from the experimenter to get up to receive the treat (despite being highly-food motivated early in the trial). For $D1^{TC}$ it was very noticeable that the dog had learned the concept since he performed more than double the number of sits respect to the other dogs in the same period of time.

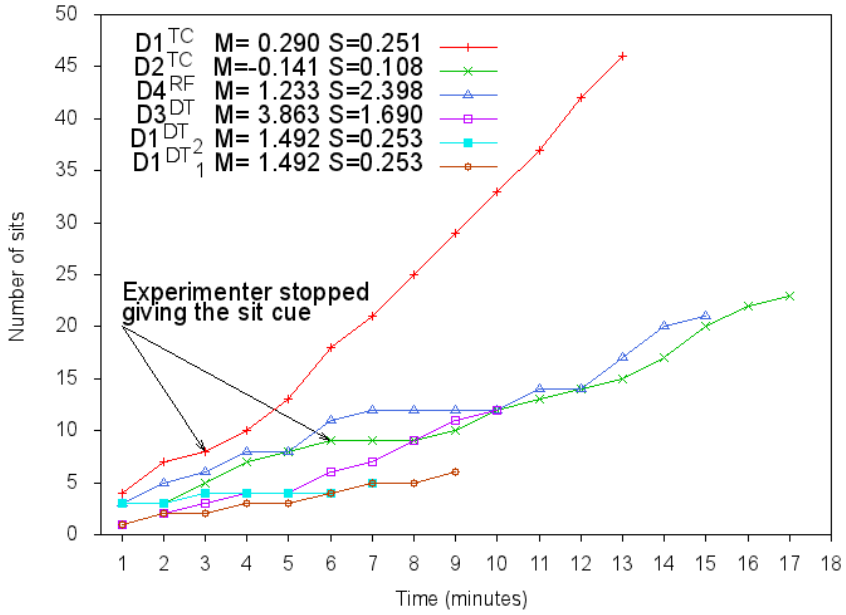

Figure 6.2 The cumulative number of sit behaviors offered by each dog over time. Series labels include the mean and standard deviation of the response latency.

One of the interesting things we observed for both $D1^{TC}$ and $D2^{TC}$ was that initially they tended to sit while orienting towards the experimenter who was located in the corner of the room. Over the course of the trial, both dogs started to ignore the experimenter in favor of focusing on the treat dispensers. Figure 6.3 shows how $D1^{TC}$ started to focus less on the experimenter, and instead focus on the treat dispensers over time. We believe that this was indicative of the dog realizing that the experimenter was not in control of the treats, and subsequently lost interest in them. This is strong evidence that the dog learned from the training system free from influence by the experimenter.

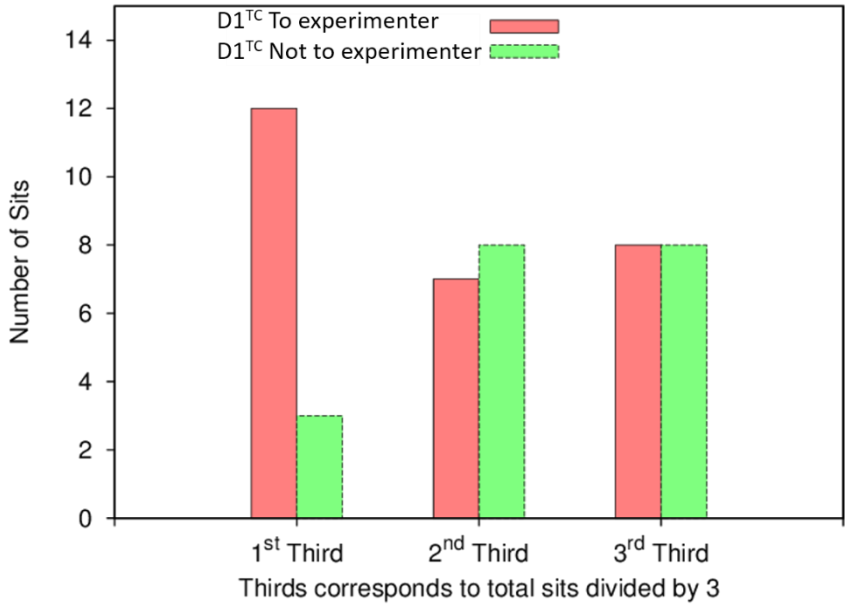

Figure 6.3 Orientation of D1TC to experimenter vs. treat dispenser grouped by thirds of the experimental trial

6.7 Discussion and Conclusion

Building on our success to this point, we have identified several improvements that can be made to both the algorithms and the training protocol while expanding the scope of the system. While the algorithms proved to be effective for teaching the sit behavior, we will need to continue developing them in order to increase the system's accuracy for teaching multiple postures, while keeping the response latency within the critical 0.5 second window.

We are also interested in having the ability to request specific behaviors from the dogs. This would entail the pairing of a unique stimulus with each behavior, which is known as a "cue." The most significant barrier to implementing this ability is timing the cue so the dog learns to associate the cue with the behavior or posture. Teaching the association between the cue and behavior will require the development of an algorithm that can predict the dogs' next behavior before the dog can complete it, while having a very high accuracy rate. We believe that development of this algorithm should be possible through the analysis of data collected

during our current and future experiments. Further, research in veterinary behavior indicates that the timing of cues is much more flexible than the timing of rewards during training [5], simplifying the predication.

In this chapter we have described, what is to our knowledge, the first autonomous canine training system capable of capturing and reinforcing postures. We briefly discussed the platform for our research, as well as the design considerations which include high-accuracy and low-latency recognition of postural changes from multiple inertial measurement units. We reported on a set of experiments with varying success using two different data filtering techniques and three different classification algorithms that illustrate the importance of the speed-accuracy trade-off.

Results of our experiments indicate that low response latency consistent with research in veterinary behavior [5], even at the cost of more than a 10% decrease in instance-based accuracy, resulted in the best performance. Further, despite the drop in instance based accuracy, achieving no false negatives when sit postures were offered also appears to correlate with successful training. To achieve these conditions, we used Univariate Spline Smoothing to process the data and remove a small amount of the noise. We then used a Threshold-based Decision Rule classifier constructed by looking for hard boundaries that existed between the accelerometer data values.

In conclusion, we have demonstrated that computer delivered rewards can be used effectively to capture and reinforce canines' postures in controlled environments. The combination of a custom-designed harness with wearable sensors and a small computer for wireless communications, instrumented treat dispensers that can be actuated remotely by a computer, and tuned algorithms for data filtering and posture classification can achieve performance on par with that required to effectively teach dogs. Our early-stage system for automated training of canines shows great promise as we look to augment its capabilities to go beyond capturing a single posture in a controlled environment.

We have extended the study described in this chapter to include a larger number of dogs and use more consistent filtering and classification technique among trials. The results

have been recently published [78]. The published manuscript provides more detail on the design process and the performance of the algorithms.

Chapter 7. Conclusions

Dogs and humans have worked in partnership for thousands of years, thanks to their ability to communicate with each other. This communication is inherently subjective because it relies on the mutual interpretation of behaviors and body language. Humans try to interpret the signals communicated by the dogs in two ways: looking for trained behaviors to indicate a certain condition or situation, and observing the untrained instinctual signs of behaviors mostly indicating the dog's emotional state. In order to reduce the ambiguity in human-canine communication, we have developed two novel wearable and wireless cyber-physical systems that can provide handlers with quantitative data related to their behavior and physiology. These are (a) a posture detection system based on IMUs and (b) a heart rate and respiratory pattern monitoring system based on ECG/PPG and electronic stethoscopes. The physiological parameters obtained from the second system (heart rate, heart rate variability and respiratory events such as sniffing or panting) have been identified as useful for the interpretation of canine emotional responses [7]–[10], [14]. Hereinafter, the contributions highlighted in each chapter are summarized along with some recommendations for future work. This chapter concludes explaining how the results of this interdisciplinary research have the potential to impact several domains of knowledge and practice.

In Chapter 3, we described an approach for canine behavior recognition based on Machine Learning Algorithms and IMUs. IMUs have been extensively used for activity profiling in human and animals, however, their use for canine posture detection is a relatively new application area for IMUs with just a few limited and inefficient attempts [75][76]. We optimized the number and location of the inertial measurement units for accurate behavior recognition while considering comfort and movement of the canine. We also provided details of the design of the IMU system. We designed a cascade of Decision tree classifiers and Hidden Markov Models for the detection of static postures (sitting, standing, lying down, standing on two legs, and eating of the ground) and dynamic activities (walking, climbing stairs and

walking down a ramp) based on heuristic features of the accelerometer and gyroscope data reducing the computational complexity with respect to [76]. The system was evaluated on ten dogs of various age, size and breed achieving average accuracy above 95%. This IMU based study contributes to our overarching goal of reducing the subjectivity in human-canine communications by enabling the detection of behaviors that are commonly taught to animals to alert their handlers. The computer-mediated communication makes this information available to handlers even when they are out of sight or earshot.

In Chapter 4, we described a wearable heart rate sensor system for wireless canine health monitoring. Electrocardiogram reveals the electrical activity of the heart and is the gold standard to measure heart rate. The variability in the interval between consecutive heart beats, referred as heart rate variability, is considered a valuable indicator of the fluctuations in autonomic nervous system [9] and a number of studies have used this parameter as an indicator of stress in animals [10][41]. The traditional ECG method in a veterinary setting involves pinching "alligator" clips to the skin on the animal while they are anesthetized. For cardiac monitoring outside the clinics, commercially-available human electrode patches are placed against the skin after shaving the hair and secured with tape and bandages. To overcome the limitations imposed by the efficiently insulated skin and dense hair layers of dogs, we introduced novel electrode configurations that allow for signal recording without the necessity of shaving dog's hair. We evaluated two types of electrodes to improve the tissue-electrode contact: thick-tapered electrode and comb-shaped arrays of thin spring-loaded pins. We demonstrated that surface modification of the metal electrodes by coating with poly (3,4-ethylenedioxythiophene) poly-(styrene-sulfonate) (PEDOT:PSS) conductive polymer can enhance their electrical property by increasing the charge injection capacity and decreasing the impedance of the tissue-electrode interface. Also, we combined ECG with PPG and studied the incorporation of light guides and optical fibers for efficient optical coupling of PPG sensors to the skin. The data coming from the IMUs, as explained in the earlier chapter, enabled the tracking of the animal's motion and deploy signal processing to eliminate the motion artifacts

on ECG data when needed. The system was evaluated on six dogs of various age, size, breed, and hair density and validated against a holter monitor used in veterinary practice.

In addition to the traditional ECGs used in veterinary clinics, some new technologies have been emerging in the consumer market for monitoring of cardiac activity. Most recently, PetPace® and Voyce® have been the most popular products which use unconventional acoustic and ultra-wideband radio-frequency technologies respectively. Due to their recent appearance, there is limited information about their technologies and clinical validation. Also, these systems provide only averaged metrics over long periods of time therefore instantaneous heart rate information is not available for researchers. In our case, these information is needed to measure the recovery time to normal heart rate after a stimuli. Chapter 4 contributes to the dissertation's overarching goal of reducing the subjectivity in human-canine communications by providing measurements of physiological parameters that have been shown to be useful indicators of the emotional state of the dog. It is important to emphasize that dogs do communicate their emotional responses with behavioral signs (e.g. trembling, restlessness, visual scanning, or yawning), however those can be subtle and difficult to perceive by the average dog handler or the behavior estimation system presented in Chapter 3. The heart rate detection system presented here can provide objective feedback regarding parameters correlated with those emotional states.

In Chapter 5, we described the last component of the physiological feedback, namely the system for continuous auscultation of respiratory behavior for the identification of respiratory patterns such as sniffing and panting. Past research has attempted to monitor sniffing and panting outside laboratory environments in a non-invasive manner using microphones placed in the nose region [62], [68]. These microphones are not immune to ambient noise and the form factor required an acclimation period to habituate to wearing the custom muzzle. To overcome this difficulty, we recorded the key body sounds at the neck and chest by means of a commercially available electronic stethoscope. We identified discrete features of sniffing and panting in the time domain and utilize event duration, event rate, event mean energy, and the number of consecutive events in a row to build a classifier. As the

classifier, we selected a decision tree classifier capable of differentiating between sniffing and panting. The algorithm design and the evaluation of the system performance on two dogs were presented in this chapter. Chapter 5 contributes to the dissertation's overall goal by providing information about when a dog is breathing, sniffing or panting to reduce the subjectivity in human-canine communications. Each of those respiratory events communicates a different message yet it is often difficult to be realized by handlers. For example, the sniffing behavior can inform about the presence of an olfactory gradient, the presence of complex odors or an olfactory success when they detect the trained target scent. Panting, on the other hand, can indicate a thermoregulatory process for cooling their body temperature or a sign of anxiety depending on the context. Therefore providing information about the respiratory state is a complimentary part of our system to complete the big picture of augmented human-canine communication.

In Chapter 6, the wearable IMU-based system was put into action to achieve an automated training system based on Skinner's theory of behavior training to augment human-canine communication by introducing computer-mediation. For this, the posture detection machine learning algorithms were combined with a computer-controlled treat dispenser to provide a primary reinforcer (food treat) when a "sit" behavior is detected as a model behavior. We explored the trade-off between the low latency and accuracy of the reward delivery with a group of four dogs of various age, size and breed. We empirically showed that dogs failed to learn the desired behavior when high classification accuracies (99% on average) were targeted coming at the cost of latencies greater than 1.1 seconds due to the computational complexity. However, slightly lower accuracies (91% on average) caused shorter latencies (less than 0.3 seconds) where dogs learned from computer-delivered rewards successfully. To our knowledge there is no existing work examining the trade-off between noise, response latency, and posture accuracy of real-time classification algorithms for canine training. The contribution of this chapter to the overall goal of the dissertation is that the computer mediation can effectively bring the correct message to the dog during the training process by optimizing

the timing and accuracy of reinforcement. This provides an important potential to demonstrate a more objective communication channel between humans and dogs.

This section has reviewed the main contributions of this dissertation, which comprise the development of a wearable system with IMU-based real-time posture detection, as well as, capability to monitor canines' heart rate, heart rate variability and respiratory parameters. The former has been leveraged into a semi-autonomous training system capable of reinforcing a particular behavior.

7.1 Recommendations for future work

As an enabling technology, the infrastructure presented in this dissertation lends itself to envisioning several further studies to augment the canine-machine interfaces and applications of these interfaces. What follows are three example suggestions on (a) how the semi-automated training system could be made fully automated and expanded to other behaviors, (b) how the physiological feedback can be integrated with the automated training platform, and (c) how the sensor and algorithms used for respiratory monitoring could be improved to make it suitable in a wider range of real-world scenarios.

In Chapter 6, we described a semi-automated training system containing four elements: a smart harness outfitted with IMUs worn by a dog, a laptop that collects data from the IMUs, algorithms for classifying postures, and a computer-controlled treat dispenser to provide reinforcement. Through a set of experiments, we demonstrated that these four components are enough to close the loop and enable researchers to observe learning of the desirable behavior. In our experiments, we automated the reinforcement of the sit behavior, but a future study could expand the platform so that specific behaviors can be requested from dogs using a "discriminative stimulus" [79]. The process of associating a behavior with a stimulus control requires a unique stimulus to be paired to each behavior and serve as a "cue". In traditional training, cues are generally verbal (such as short words) or physical (such as an up-sweeping motion with the hand). Technology and computational systems provide a new dimension, where computers can provide the appropriate timing and consistency in giving cues but also

innovate with the nature of the cues, like electronically generated sounds, images on a display, etc.

As we discussed throughout the Chapter 6, the proposed system is semi-autonomous because at the initial phase it requires human assistance to provide labels for supervised classification of postures. To make it fully autonomous, the use of unsupervised classifiers would be required on which we have demonstrated some preliminary efforts [80].

Also, it remains as future work to include the physiological feedback in the training loop. Canine training can benefit from the physiological data in two ways. First, by having a more accurate interpretation of the canine emotional response to a certain stimulus or conditions. This understanding could be used to define more comfortable methods or environments where training and interaction with dogs could be more efficiently. Second, the multi sensor data from ECG, PPG and stethoscope could enable the identification of patterns in those signals that correlate to the presence of the stimulus. For example, an automated scent discrimination training system could be realized by utilizing a computer-controlled scent delivery apparatus. Similarly to the experiment in Chapter 6, algorithms could capture and reinforce a desirable behavioral response upon detection of scent. It is important to note that this system should not only rely on the physiological feedback to provide the reinforcement but look for a postural indication following the detection of the scent. This is necessary because the goal is not only to train the dog to identify the scent, but also to communicate that to the handler. Dogs are generally trained to alert by standing still performing a specific posture. Moving or pawing around the target scent could also put them in harmful situations if the search was for explosive detection task.

In Chapter 5 we described how we used a commercially e-stethoscope (*One*® by Thinklabs) to evaluate the feasibility of continuously recording respiratory sounds from the chest and neck and detecting sniffing and panting respiratory patterns. We identified discrete features of sniffing and panting in the time domain that enabled classification. The stethoscopes were held in place by a custom 3D printed mounting structure and four silicon backed elastic bands that we shaped as collar and chest strap. The height of the e-stethoscope

that was originally designed for medical doctors examining human patients, imposed some limitations in the range of activities that the dog could perform while getting good quality data. To get around this issue for our analysis, we designed our experiment so the dog would be sniffing bins instead of a whole area. Generally when dogs have to search a large space they move to cover the area and find the scent plume. In addition to miniaturization, a further study could assess if features (such as Mel frequency cepstral coefficients) and techniques that are commonly used in speech recognition could be used as effectively to improve the respiratory pattern classification [81], [82].

7.2 Broader impact of this dissertation

The results of this interdisciplinary research have the potential to impact several domains of knowledge and practice. We will focus on the benefits our technologies provide to the training and care of working and service dogs, to the veterinary community, and to the market segment developing pet products.

Throughout the various chapters of the dissertation, we have discussed different ways in which human-canine interactions can be enhanced by using technology. Specifically, in Chapter 2, we discussed three particular benefits that these technologies provide to search and rescue teams. First, the sensors on the dog, both for environmental and health monitoring, augment the capabilities of handlers by supplying them with objective decision support so they can improve their efficiency and provide better care to their dogs. Second, our cBAN provides computer-mediated communication between handlers and dogs even when there are physical boundaries between them. Finally, we discussed how the combination of wearable sensors and computer system can automate certain training tasks. Training military or service dogs costs several tens of thousands of dollars [11], requiring extensive efforts of specialized personnel during long periods of time, and does not always achieve the desired results. By using technologies, we can reduce the cost and the time required by specialized personnel to achieve training, thereby allowing more people to benefit from these specially trained dogs. These capabilities extend beyond the search and rescue dogs to many other uses of working dogs, such as guide dogs for the blind, therapy dogs, and so on.

The veterinary community is the second group that could benefit from the technology developments presented here. Our cyber-physical systems generate a continuous stream of data that is sent through a wireless network to the processing unit. The insights gained from the algorithms, and even the raw data, can be shared with the end user in real time or stored for later analysis. These features allow veterinarians to have access to data in real time, but also outside the laboratory environment. This is important because a visit to the clinic does not always reflect the general condition of the dog. Studies have shown that dogs present signs of anxiety during hospital visits [83], [84]. In addition, a clinical visit reflects a single point in time, while our technologies can provide continuous data during daily situations.

To the best of our knowledge, there is no commercial tool available for the monitoring of sniffing and panting outside the laboratory environment. The tools available for laboratory analysis (barometric whole-body plethysmograph) impose restrictions on animals' natural behavior [85]. As mentioned earlier, these behaviors can provide information about the emotional status of the dog and could be very useful for specialists in behavior medicine to look at the frequency and circumstances of the occurrence. Also, these technologies can enable further analysis of these behaviors to gain better understanding of them from a physiological perspective.

Lastly, our design process and the lessons we learnt by developing this system for dogs may serve as an example in the process of developing other canine technologies, especially in the Internet of Things (IoT) space. For the past two decades, the Internet and the advancements in wireless technologies have changed the way we access information and communicate with each other, enabling more frequent, effective, and ubiquitous communication means. We are now seeing this flow of information shifting towards "things", allowing objects to be sensed and controlled remotely across existing network infrastructure. We are experiencing a rapid proliferation of IoT-enabled devices in the market, which are affecting most sectors of the global economy, including retail, healthcare, transportation, and energy. These devices, which generally include low-cost sensors, embedded microcontrollers, and wireless connectivity, take advantage of cloud services through Internet connections to distribute computing tasks

and quickly analyze massive amounts of data and take action on the insights developed from the analysis. These are also known as cyber-physical systems, and they are finding their way into consumer pet products to enhance human-canine interactions. These devices are taking the form of smart feeders (*RightBite*®), adaptive interactive puzzles (*Clever Pet*®), GPS trackers (*Tagg*®), connected cameras to watch pets while away from home (PetCube®), activity and health monitor collars (*Whistle*®, *Voyce*® and *PetPace*®), and even smart litter boxes that monitor the animal's weight and the amount of waste produced to provide insights on wellness (*Tailio* ®). In each of these examples, the aim is to increase handlers' awareness about their pets. The systems described in this dissertation can be considered part of this IoT ecosystem. Through the entire course of this dissertation, we followed a participatory design process that involved dogs, the end users of our technologies, to ensure the feasibility of the design. However, designing wearable technologies that are comfortable for dogs imposes many challenges. Most dogs will not tolerate sensors strapped to their paws, tail, or around their snout. Thus, future design and engineering efforts to develop wearable technologies for dogs in the IoT space can benefit from our findings and design flow. The findings of our study provide insight on how one might set their system requirements and ensure they provide benefits for both humans and animals.

REFERENCES

[1] M. Raibert, K. Blankespoor, G. Nelson, R. Playter, and others, “BigDog, the rough-terrain quadruped robot,” in *Proceedings of the 17th World Congress*, 2008, pp. 10823–10825.

[2] S. E. Stitzel, D. R. Stein, and D. R. Walt, “Enhancing vapor sensor discrimination by mimicking a canine nasal cavity flow environment,” *J. Am. Chem. Soc.*, vol. 125, no. 13, pp. 3684–3685, 2003.

[3] K. G. Furton and L. J. Myers, “The scientific foundation and efficacy of the use of canines as chemical detectors for explosives,” *Talanta*, vol. 54, no. 3, pp. 487–500, 2001.

[4] E. F. Hiby, N. J. Rooney, and J. W. S. Bradshaw, “Dog training methods: their use, effectiveness and interaction with behaviour and welfare,” *Anim. Welf.*, vol. 13, no. 1, pp. 63–70, 2004.

[5] M. Yamamoto, T. Kikusui, and M. Ohta, “Influence of delayed timing of owners’ actions on the behaviors of their dogs, Canis familiaris,” *J. Vet. Behav. Clin. Appl. Res.*, vol. 4, no. 1, pp. 11–18, 2009.

[6] B. Bailey and P. Farhoody, “Discrimination Workshop,” 2015. [Online]. Available: http://behaviormatters.com/workshops/. [Accessed: 01-Jan-2016].

[7] S. Venkatraman, X. Jin, R. M. Costa, and J. M. Carmena, “Investigating neural correlates of behavior in freely behaving rodents using inertial sensors.,” *J. Neurophysiol.*, vol. 104, no. 1, pp. 569–75, Jul. 2010.

[8] R. a Galosy, L. K. Clarke, and J. H. Mitchell, “Cardiac changes during behavioral stress in dogs.,” *Am. J. Physiol.*, vol. 236, no. 5, pp. H750–8, May 1979.

[9] M. Malik, “Heart Rate Variability,” *Ann. Noninvasive Electrocardiol.*, vol. 1, no. 2, pp. 151–181, 1996.

[10] E. Mohr, J. Langbein, and G. Nürnberg, “Heart rate variability: a noninvasive approach to measure stress in calves and cows,” *Physiol. Behav.*, vol. 75, no. 1, pp. 251–259, 2002.

[11] W. Konrad, “An Aide for the Disabled, a Companion, and Nice and Furry,” *The New York Times*, 21-Aug-2009.

[12] K. Arshak, E. Moore, G. M. Lyons, J. Harris, and S. Clifford, “A review of gas sensors employed in electronic nose applications,” *Sens. Rev.*, vol. 24, no. 2, pp. 181–198, 2004.

[13] P. Martin, P. P. G. Bateson, and P. Bateson, *Measuring behaviour: an introductory guide*. Cambridge University Press, 1993.

[14] B. D. Hansen, B. D. X. Lascelles, B. W. Keene, A. K. Adams, and A. E. Thomson, “Evaluation of an accelerometer for at-home monitoring of spontaneous activity in dogs,” *Am. J. Vet. Res.*, vol. 68, no. 5, 2007.

[15] “Integrated Telemetry Systems.” [Online]. Available: www.itstelemetry.com. [Accessed: 21-Jul-2014].

[16] M. M. Jackson, C. Zeagler, G. Valentin, A. Martin, V. Martin, A. Delawalla, W. Blount, S. Eiring, R. Hollis, Y. Kshirsagar, and others, “FIDO-facilitating interactions

for dogs with occupations: wearable dog-activated interfaces," in *Proceedings of the 17th annual international symposium on International symposium on wearable computers*, 2013, pp. 81–88.

[17] R. Brugarolas, D. Roberts, B. Sherman, and A. Bozkurt, "Machine Learning Based Posture Estimation for a Wireless Canine Machine Interface," in *IEEE Topical Conference on Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireless)*, 2012, pp. 3–5.

[18] R. Brugarolas, R. T. Loftin, P. Yang, D. L. Roberts, B. Sherman, and A. Bozkurt, "Behavior Recognition Based on Machine Learning Algorithms for a Wireless Canine Machine Interface," in *10th annual Body Sensor Networks Conference*, 2013, pp. 1–5.

[19] "xSense." [Online]. Available: www.xsense.com/en/mtw.

[20] C. Ribeiro, A. Ferworn, M. Denko, and J. Tran, "Canine Pose Estimation: A Computing for Public Safety Solution," *2009 Can. Conf. Comput. Robot Vis.*, pp. 37–44, May 2009.

[21] R. Brugarolas, D. L. Roberts, B. Sherman, and A. Bozkurt, "Posture Estimation for a Canine Machine Interface Based Training System," in *Proceedings of The Engineering in Medicine and Biology Conference*, 2012, pp. 4489–4492.

[22] B. Beerda, M. B. H. Schilder, J. A. van Hooff, and H. W. de Vries, "Manifestations of chronic and acute stress in dogs," *Appl. Anim. Behav. Sci.*, vol. 52, no. 3, pp. 307–319, 1997.

[23] B. F. Skinner, *Science and Human Behavior*. Macmillan, 1953.

[24] J. Miller, G. Flowers, and D. Bevly, “A System for Tracking an Autonomously Controlled Canine,” *J. Navig.*, vol. 65, no. 03, pp. 427–444, 2012.

[25] T. Instruments, “2.4 GHz Bluetooth low energy System-on-chip,” 2011. [Online]. Available: http://www.ti.com/product/cc2540.

[26] Murata Electronics Oy, “CMA3000 Ultra Small Size, Low Power 3-axis Accelerometer,” 2012. [Online]. Available: http://www.muratamems.fi/products/accelerometers/cma3000-accelerometers.

[27] STMicroelectronics, “L3G4200D low power 3-axis Gyroscope,” 2012. [Online]. Available: http://www.st.com/internet/analog/product/250373.jsp.

[28] S. Weinberger, “IDE Detector Dog Research Aims to Save Lives of U.S. Combat Troops -- Human and K-9,” *August 31*, 2012.

[29] I. H. W. Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, *The WEKA Data Mining Software: An Update*. SIGKDD Explorations, Volume 11, Issue 1, 2009.

[30] P. Viola and M. Jones, “Robust Real-time Object Detection,” in *Second International Workshop on Statistical and Computational Theories of Vision*, 2001, pp. 1–25.

[31] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, “A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains,” *Ann. Math. Stat.*, pp. 164–171, 1970.

[32] A. Viterbi, “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm,” *Inf. Theory, IEEE Trans.*, vol. 13, no. 2, pp. 260–269, 1967.

[33] R. Brugarolas, D. Roberts, B. Sherman, and A. Bozkurt, "Machine learning based posture estimation for a wireless canine machine interface," in *2013 IEEE Topical Conference on Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireleSS)*, 2013, pp. 10–12.

[34] T. M. Mitchell, *Machine Learning*. The Mc-Graw-Hill Companies, Inc., 1997.

[35] J. R. Quinlan, "Induction of Decision Trees," *Mach. Learn.*, vol. 1, no. 1, pp. 81–106, Mar. 1986.

[36] T. Instruments, "CC3200 SimpleLink ™ Wi-Fi ® and Internet-of-Things Solution , a Single-Chip Wireless MCU," 2014. [Online]. Available: http://www.ti.com/product/CC3200. [Accessed: 01-Jan-2016].

[37] "Alba Medical Systems, Inc." [Online]. Available: www.albamedical.com/transmitholter.html. [Accessed: 21-Jul-2014].

[38] A. Bozkurt, D. L. Roberts, B. L. Sherman, R. Brugarolas, S. Mealin, J. Majikes, P. Yang, and R. Loftin, "Toward Cyber-Enhanced Working Dogs for Search and Rescue," *Intell. Syst. IEEE*, vol. 29, no. 6, pp. 32–39, Nov. 2014.

[39] R. Brugarolas, T. Latif, J. Dieffenderfer, K. Walker, S. Yuschak, B. L. Sherman, D. L. Roberts, and A. Bozkurt, "Wearable Heart Rate Sensor Systems for Wireless Canine Health Monitoring," *IEEE Sens. J.*, vol. 16, no. 10, pp. 3454–3464, May 2016.

[40] L. P. Tilley and F. W. K. Smith Jr, *Blackwell's five-minute Veterinary consult: canine and feline*. John Wiley & Sons, 2011.

[41] B. Beerda, M. B. H. Schilder, J. A. R. A. . van Hooff, H. W. de Vries, and J. A. Mol,

"Behavioural, saliva cortisol and heart rate responses to different types of stimuli in dogs," *Appl. Anim. Behav. Sci.*, vol. 58, no. 3, pp. 365–381, Feb. 2015.

[42] "Corsience GmbH & Co. KG." [Online]. Available: www.corscience.de/en/medicalengineering/products/telemedicine/te-sys-ecg.html. [Accessed: 21-Jul-2014].

[43] "EMKA Technologies." [Online]. Available: www.telemetry.emka.fr/prod.php?prod=4. [Accessed: 21-Jul-2014].

[44] V. S. M. Jonckheer-Sheehy, C. M. Vinke, and A. Ortolani, "Validation of a Polar® human heart rate monitor for measuring heart rate and heart rate variability in adult dogs under stationary conditions," *J. Vet. Behav. Clin. Appl. Res.*, vol. 7, no. 4, pp. 205–212, Jul. 2012.

[45] A. Essner, R. Sjöström, E. Ahlgren, and B. Lindmark, "Validity and reliability of Polar® RS800CX heart rate monitor, measuring heart rate in dogs during standing position and at trot on a treadmill," *Physiol. Behav.*, vol. 114–115, no. 0, pp. 1–5, Apr. 2013.

[46] D. S. Quintana, J. a J. Heathers, and A. H. Kemp, "On the validity of using the Polar RS800 heart rate monitor for heart rate variability research.," *Eur. J. Appl. Physiol.*, vol. 112, no. 12, pp. 4179–80, Dec. 2012.

[47] A. Bozkurt and A. Lal, "Low-cost flexible printed circuit technology based microelectrode array for extracellular stimulation of the invertebrate locomotory system," *Sensors Actuators A Phys.*, vol. 169, no. 1, pp. 89–97, Sep. 2011.

[48] K. A. Ludwig, N. B. Langhals, M. D. Joseph, S. M. Richardson-Burns, J. L. Hendricks, and D. R. Kipke, "Poly(3,4-ethylenedioxythiophene) (PEDOT) polymer coatings facilitate smaller neural recording electrodes," *J. Neural Eng.*, vol. 8, no. 1, p. 14001, 2011.

[49] X. Cui and D. C. Martin, "Electrochemical deposition and characterization of poly(3,4-ethylenedioxythiophene) on neural microelectrode arrays," *Sensors and Actuators B-chemical*, vol. 89, no. 1, pp. 92–102, 2003.

[50] N. S. Matthews, S. Hartke, and J. C. Allen, "An evaluation of pulse oximeters in dogs, cats and horses.," *Vet. Anaesth. Analg.*, vol. 30, no. 1, pp. 3–14, Jan. 2003.

[51] H. Gesche, D. Grosskurth, G. Küchler, and A. Patzak, "Continuous blood pressure measurement by using the pulse transit time: comparison to a cuff-based method," *Eur. J. Appl. Physiol.*, vol. 112, no. 1, pp. 309–315, 2012.

[52] K. Maros, A. Dóka, and Á. Miklósi, "Behavioural correlation of heart rate changes in family dogs," *Appl. Anim. Behav. Sci.*, vol. 109, no. 2, pp. 329–341, 2008.

[53] P. Castiglioni, A. Faini, G. Parati, and M. Di Rienzo, "Wearable seismocardiography," in *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, 2007, pp. 3954–3957.

[54] J. Zanetti and others, "Seismocardiography: a new technique for recording cardiac vibrations. Concept, method, and initial observations," *J. Cardiovasc. Technol.*, vol. 9, no. 2, pp. 111–118, 1990.

[55] Beaglebord.org, "BeagleBone Black." [Online]. Available: Beaglebord.org.

[Accessed: 02-Jan-2015].

[56] G. Technologies, “g.tec Sahara dry electrode.” [Online]. Available: http://www.gtec.at/. [Accessed: 02-Jan-2015].

[57] W. A. H. Engelse and C. Zeelenberg, “A single scan algorithm for QRS-detection and feature extraction,” *Comput. Cardiol.*, vol. 6, no. 1979, pp. 37–42, 1979.

[58] S. T. Boerema, L. van Velsen, L. Schaake, T. M. Tönis, and H. J. Hermens, “Optimal sensor placement for measuring physical activity with a 3D accelerometer,” *Sensors*, vol. 14, no. 2, pp. 3188–3206, 2014.

[59] C. V. C. Bouten, K. Koekkoek, M. Verduin, R. Kodde, and J. D. Janssen, “A triaxial accelerometer and portable data processing unit for the assessment of daily physical activity,” *Biomed. Eng. IEEE Trans.*, vol. 44, no. 3, pp. 136–147, 1997.

[60] W. S. Helton, *Canine ergonomics: the science of working dogs*. CRC Press, 2009.

[61] B. a Craven, E. G. Paterson, and G. S. Settles, “The fluid dynamics of canine olfaction: unique nasal airflow patterns as an explanation of macrosmia.,” *J. R. Soc. Interface*, vol. 7, no. 47, pp. 933–943, 2010.

[62] A. Thesen, J. B. Steen, and K. B. Doving, “Behaviour of dogs during olfactory tracking,” *J. Exp. Biol.* , vol. 180 , no. 1 , pp. 247–251, Jul. 1993.

[63] J. B. Steen, I. Mohus, T. Kvesetberg, and L. Walløe, “Olfaction in bird dogs during hunting.,” *Acta Physiol. Scand.*, vol. 157, no. 1, pp. 115–9, 1996.

[64] E. C. Crawford, “Mechanical Aspects of panting in dogs.,” *J. Appl. Physiol.*, vol. 17, no. 2, pp. 249–251, 1962.

[65] K. Schmidt-nielsen, W. L. Bretz, and C. R. Taylor, “Panting in dogs: unidirectional air flow over evaporative surfaces.,” *Science*, vol. 169, no. 950, pp. 1102–1104, 1970.

[66] I. Gazit and J. Terkel, “Explosives detection by sniffer dogs following strenuous physical activity,” *Appl. Anim. Behav. Sci.*, vol. 81, no. 2, pp. 149–161, 2003.

[67] W. E. Hull and E. C. Long, “Respiratory impedance and volume flow at high frequency in dogs,” *J. Appl. Physiol.*, vol. 16, no. 3, pp. 439–443, 1961.

[68] I. Gazit, Y. Lavner, G. Bloch, O. Azulai, A. Goldblatt, and J. Terkel, “A simple system for the remote detection and analysis of sniffing in explosives detection dogs.,” *Behav. Res. Methods. Instrum. Comput.*, vol. 35, no. 1, pp. 82–89, 2003.

[69] S. J. Preece, J. Y. Goulermas, L. P. J. Kenney, D. Howard, K. Meijer, and R. Crompton, “Activity identification using body-mounted sensors--a review of classification techniques.,” *Physiol. Meas.*, vol. 30, no. 4, pp. 1–33, Apr. 2009.

[70] C. Cornou and S. Lundbye-Christensen, “Classifying sows activity types from acceleration patterns: an application of the multi-process Kalman filter,” *Appl. Anim. Behav. Sci.*, vol. 111, no. 3, pp. 262–273, 2008.

[71] D. J. Wrigglesworth, E. S. Mort, S. L. Upton, and A. T. Miller, “Accuracy of the use of triaxial accelerometry for measuring daily activity as a predictor of daily maintenance energy requirement in healthy adult Labrador Retrievers,” *Am. J. Vet. Res.*, vol. 72, no. 9, pp. 1151–1155, 2011.

[72] M. Moreau, S. Siebert, A. Buerkert, and E. Schlecht, “Use of a tri-axial accelerometer for automated recording and classification of goats’ grazing behaviour,” *Appl. Anim.*

Behav. Sci., vol. 119, no. 3–4, pp. 158–170, Jul. 2009.

[73] J. Miller and D. M. Bevly, "Position and orientation determination for a guided K-9," in *Proceedings of the 20th international technical meeting of the satellite division of the institute of navigation ION GNSS*, 2007, pp. 1768–1776.

[74] W. R. Britt, J. Miller, P. Waggoner, D. M. Bevly, and J. A. Hamilton Jr, "An embedded system for real-time navigation and remote command of a trained canine," *Pers. Ubiquitous Comput.*, vol. 15, no. 1, pp. 61–74, 2011.

[75] C. Ribeiro, A. Ferworn, M. Denko, and J. Tran, "Canine Pose Estimation: A Computing for Public Safety Solution," in *Canadian Conference on Computer and Robot Vision, 2009. CRV'09*, 2009, pp. 37–44.

[76] L. Gerencsér, G. Vásárhelyi, M. Nagy, T. Vicsek, and A. Miklósi, "Identification of Behaviour in Freely Moving Dogs (Canis familiaris) Using Inertial Sensors," *PLoS One*, vol. 8, no. 10, 2013.

[77] J. Barry, M. Emmen, and S. Smith, *Positive Gun Dogs: Clicker Training for Sporting Breeds*. Sunshine Books, 2007.

[78] J. Majikes, R. Brugarolas, M. Winters, S. Yuschak, S. Mealin, K. Walker, P. Yang, B. Sherman, A. Bozkurt, and D. L. Roberts, "Balancing Noise Sensitivity, Response Latency, and Posture Accuracy for a Computer-Assisted Canine Posture Training System," *Int. J. Hum. Comput. Stud.*, 2016.

[79] J. O. Cooper, T. E. Heron, and W. L. Heward, *Applied Behavior Analysis (2nd ed.)*. Pearson, 2007.

[80] M. Winters, R. Brugarolas, J. Majikes, S. Mealin, S. Yuschak, B. Sherman, A. Bozkurt, and D. Roberts, “Knowledge Engineering for Unsupervised Canine Posture Detection from IMU Data,” in *Advances in Computer Entertainment*, 2015.

[81] S. Chauhan, P. Wang, C. S. Lim, and V. Anantharaman, “A computer-aided MFCC-based {HMM} system for automatic auscultation,” *Comput. Biol. Med.*, vol. 38, no. 2, pp. 221–233, 2008.

[82] R. Palaniappan and K. Sundaraj, “Respiratory sound classification using cepstral features and support vector machine,” in *Intelligent Computational Systems (RAICS), 2013 IEEE Recent Advances in*, 2013, pp. 132–136.

[83] S. G. Moesgaard, A. V Holte, T. Mogensen, J. Mølbak, A. T. Kristensen, A. L. Jensen, T. Teerlink, A. J. Reynolds, and L. H. Olsen, “Effects of breed, gender, exercise and white-coat effect on markers of endothelial function in dogs,” *Res. Vet. Sci.*, vol. 82, no. 3, pp. 409–415, 2007.

[84] A. M. Belew, T. Barlett, and S. A. Brown, “Evaluation of the White-Coat Effect in Cats,” *J. Vet. Intern. Med.*, vol. 13, no. 2, pp. 134–142, 1999.

[85] J. Talavera, N. Kirschvink, S. Schuller, A. Le Garrérès, P. Gustin, J. Detilleux, and C. Clercx, “Evaluation of respiratory function by barometric whole-body plethysmography in healthy dogs,” *Vet. J.*, vol. 172, no. 1, pp. 67–77, 2006.

APPENDIX

APPENDIX A - PUBLICATIONS

Publications produced from the content of this dissertation.

Journal Papers:

1. R. Brugarolas, T. Latif, J. Dieffenderfer, K. Walker, S. Yuschak, B. L. Sherman, D. L. Roberts, and A. Bozkurt, "Wearable Heart Rate Sensor Systems for Wireless Canine Health Monitoring," *IEEE Sensors Journal*, vol. 16, no. 10, pp. 3454–3464, May 2016.
2. J. Majikes, R. Brugarolas (co-first author), M. Winters, S. Yuschak, S. Mealin, K. Walker, P. Yang, B. Sherman, A. Bozkurt, and D. L. Roberts, "Balancing Noise Sensitivity, Response Latency, and Posture Accuracy for a Computer-Assisted Canine Posture Training System," *International Journal Human-Computer Studies*, 2016.
3. Bozkurt, D. L. Roberts, B. L. Sherman, R. Brugarolas, S. Mealin, J. Majikes, P. Yang, and R. Loftin, "Toward Cyber-Enhanced Working Dogs for Search and Rescue," Intelligent Systems IEEE, vol. 29, no. 6, pp. 32–39, Nov. 2014.
4. R. Brugarolas, S. Yuschak, B.Sherman, D. Roberts, and A. Bozkurt, "Wearable system for continuous monitoring of heart rate and respiratory behavior in dogs during olfactory detection", in preparation to be submitted to IEEE Sensors Journal.

Conference Papers:

1. R. Brugarolas, D. Roberts, B. Sherman, and A. Bozkurt, "Posture estimation for a canine machine interface based training system," in 2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2012, pp. 4489–4492.
2. R. Brugarolas, D. Roberts, B. Sherman, and A. Bozkurt, "Machine learning based posture estimation for a wireless canine machine interface," in 2013 IEEE Topical Conference on Biomedical Wireless Technologies, Networks, and Sensing Systems, 2013, pp. 10–12.

3. R. Brugarolas, R. T. Loftin, P. Yang, D. L. Roberts, B. Sherman, and A. Bozkurt, “Behavior recognition based on machine learning algorithms for a wireless canine machine interface,” in 2013 IEEE International Conference on Body Sensor Networks, 2013, pp. 1–5.
4. R. Brugarolas, J. Dieffenderfer, K. Walker, A. Wagner, B. Sherman, D. Roberts, and A. Bozkurt, “Wearable wireless biophotonic and biopotential sensors for canine health monitoring,” in IEEE SENSORS 2014 Proceedings, 2014, pp. 2203–2206.
5. R. Brugarolas, T. Agcayazi, S. Yuschak, D. L. Roberts, B. Sherman, and A. Bozkurt, “Towards a wearable system for continuous monitoring of sniffing and panting in dogs,” Accepted for presentation at IEEE Body Sensor Networks, 2016.