

ABSTRACT

MOHAMMADZADEH, FARROKH FABIAN. Power Efficient Physiological Response Prediction for Wearable Health Monitoring Platforms. (Under the direction of Dr. Edgar Lobaton and Dr. Chang S. Nam.)

The main objective of this dissertation is the development of a power-efficient framework for prediction of physiological responses such as breathing rate in a wearable health monitoring platform when the user is performing various activities. In developing the prediction model, we established that using all available features for prediction may result in fast battery depletion. We specifically aimed our attention at feature selection to optimize power consumption for prediction of breathing rate. We performed dynamic feature selection by identifying the activity state of the subject using features extracted from a single accelerometer. In each activity state, a group lasso regression model was used to predict breathing rate in a power-efficient way by selecting most important features. Moreover, we evaluated different algorithms such as a nonlinear support vector machine (SVM) regression model for prediction of future physiological response and discovered that physiological responses are highly correlated to past measurements. We also concluded that environmental factors can help prediction of physiological responses under real-world conditions. The physiological responses we aimed to predict depend on the activity/motion-behavior of the subject as well as environmental factors. Accordingly, we first focused on the characterization of different activities, then on the exploratory analysis of physiological response prediction, and finally an energy efficient and activity-aware physiological response framework is introduced.

© Copyright 2018 by Farrokh Fabian Mohammadzadeh

All Rights Reserved

Power Efficient Physiological Response Prediction for Wearable Health Monitoring Platforms

by
Farrokh Fabian Mohammadzadeh

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

2018

APPROVED BY:

Dr. Alper Bozkurt

Dr. Gregory Buckner

Dr. Edgar Lobaton
Co-chair of Advisory Committee

Dr. Chang S. Nam
Co-chair of Advisory Committee

BIOGRAPHY

Farrokh Mohammadzadeh received his Bachelor of Science degree in Microbiology & Immunology and Physics from the University of Miami, Coral Gables, FL in 2010 and his Master of Science degree in Electrical Engineering in 2014 from North Carolina State University.

ACKNOWLEDGEMENTS

I would like to express my sincere appreciation and thanks to my advisor, Dr. Edgar Lobaton, without whom this would not have been possible and Dr. C.S. Nam for his mentorship and support. I would also like to thank the rest of my advisory committee: Dr. Alper Bozkurt and Dr. Gregory Buckner for their support.

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	vii
Chapter 1 Introduction	1
Chapter 2 Feasibility of a Wearable, Sensor-Based Motion Tracking System	11
2.1 Introduction	11
2.2 Development of the wearable motion tracking system	14
2.2.1 Overview	14
2.2.2 Tracking algorithm	15
2.3 Method	17
2.3.1 Subject	17
2.3.2 Experiment Procedure	17
2.3.3 Data collection and analysis	19
2.4 Results and Discussion	19
2.5 Conclusion and Extensions	20
Chapter 3 Prediction of Physiological Response over Varying Forecast Lengths with a Wearable Health Monitoring Platform	23
3.1 Introduction	24
3.2 Dataset	25
3.2.1 Sensors	25
3.2.2 Data Acquisition Protocol	27
3.3 Prediction of Physiological Responses	27
3.3.1 Problem Description	27
3.3.2 Prediction Model	28
3.3.3 SVM Regression	28
3.3.4 Linear SVM Regression	28
3.3.5 Nonlinear SVM Regression	29
3.4 Experimental Results	31
3.5 Conclusion and Extensions	33
Chapter 4 Sensor Selection for Power-Efficient Prediction of Physiological Responses in a Wearable Health-Monitoring Platform	38
4.1 Introduction	39
4.2 Dataset	40
4.2.1 Data Acquisition Protocol	40
4.3 Proposed Framework	41
4.3.1 Activity State Identification	42
4.3.2 Prediction Model: Group Lasso for Linear Regression	43
4.4 Experimental Results	47

4.4.1	Clustering for Activity Identification	47
4.4.2	Feature Selection Versus Prediction Accuracy	48
4.4.3	Dynamic Trade-Off Between Feature Set and Power	52
4.5	Comparison of the Methods and Power Consumptions	54
4.6	Conclusion and Extensions	60
BIBLIOGRAPHY		70

LIST OF TABLES

Table 2.1	Results for all subjects.	22
Table 4.1	RMSE and Number of Sensors at λ_{min} for 10s forecast length.	49
Table 4.2	RMSE of GL and SVM, number of sensors, and power consumption for λ_{min} with and without clustering for forecast length of 10s.	58
Table 4.3	List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 10s. . .	59
Table 4.4	RMSE of GL and SVM and number of sensors for λ corresponding to 3% more error in the training set with and without clustering for forecast length of 10s. .	60
Table 4.5	List of sensors selected by group lasso with λ corresponding to 3% more error in the training set for each day used as test set and for each cluster as well as without clustering for forecast length of 10s.	61
Table 4.6	RMSE of GL and SVM and number of sensors for λ corresponding to 5% more error in the training set with and without clustering for forecast length of 10s. .	62
Table 4.7	List of sensors selected by group lasso with λ corresponding to 5% more error in the training set for each day used as test set and for each cluster as well as without clustering for forecast length of 10s.	63
Table 4.8	RMSE of GL and SVM, number of sensors, and power consumption for λ_{min} with and without clustering for forecast length of 30s.	64
Table 4.9	List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 30s. . .	65
Table 4.10	RMSE of GL and SVM, number of sensors, and power consumption for λ corresponding to 3% more error in training set with and without clustering for forecast length of 30s.	66
Table 4.11	List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 30s. . .	67
Table 4.12	RMSE of GL and SVM, number of sensors, and power consumption for λ corresponding to 5% more error in training set with and without clustering for forecast length of 30s.	68
Table 4.13	List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 30s. . . .	69

LIST OF FIGURES

Figure 1.1	Relation between breathing rate (bpm) and trait anxiety scores during anticipatory anxiety. A positive correlation between the increase in breathing rate and trait anxiety scores is observed [MH01].	4
Figure 2.1	System Overview Flowchart.	15
Figure 2.2	Experiment set-up. Left: Subject wearing OptiTrack Velcro suit with markers and IMU sensors; Right: Experiment layout with 12 OptiTrack cameras.	16
Figure 2.3	Elbow flexion-extension motion.	18
Figure 2.4	Elbow flexion-extension motion.	20
Figure 2.5	Joint angle and visualization of an elbow flexion motion.	21
Figure 3.1	Breathing rate prediction results using all sensing modalities with a 10 second forecast. First row: results from trained and constant predictors. Second row: residuals of the trained and constant predictors as a function of breathing rate. The curves represent the 5, 50 and 95 percentiles of the residuals. Third row: histograms of the residuals for the trained and constant predictors. The mean is shown as a dotted line. Fourth row: mean absolute error of the trained and constant predictors as a function of breathing rate.	34
Figure 3.2	Breathing rate prediction results using all sensing modalities with a 20 second forecast.	35
Figure 3.3	Breathing rate prediction error for various forecast lengths. The constant predictor is compared to trained predictors with six different sets of features.	36
Figure 3.4	Residual values of breathing rate predictions for a 20 second forecast on testing data as a function of the actual breathing rate. Results are shown for four days.	37
Figure 4.1	Block diagram of the signal processing flow in the proposed framework	42
Figure 4.2	Explicit solution of ridge (left) and lasso (right) for the orthonormal feature matrix. The red dotted line shows the unrestricted least square estimation [4].	45
Figure 4.3	Elliptical contours of the least squares cost function and contours of the regularization region for ridge(left) and lasso(right).	46
Figure 4.4	Visualization of different clusters of the acceleration features in reduced dimension in space with the number of data points in each cluster.	48
Figure 4.5	Results of clustering over time in the test set.	49
Figure 4.6	GCV and SLLCV versus λ in the training set obtained via cross validation with forecast length of 10s.	50
Figure 4.7	Top: RMSE (root mean square error) and number of sensors versus λ for all four clusters in the training set obtained via cross validation. Bottom: λ at which the sensors data are excluded from the model in each cluster as well as no clustering. λ_{min} is showed with a dashed line and the green and orange bars show sensors included in and excluded from the prediction at λ_{min} , respectively.	51

Figure 4.8	Predicted and actual breathing rate vs time using the respective cluster's λ_{min} obtained from LCV.	52
Figure 4.9	GCV and SLICV versus λ in the training set obtained via cross validation with forecast length of 20s.	53
Figure 4.10	Top: RMSE (root mean square error) and number of sensors versus λ for all four clusters in the training set obtained via cross validation. Bottom: λ at which the sensors data are excluded from the model in each cluster as well as no clustering. λ_{min} is showed with a dashed line and the green and orange bars show sensors included in and excluded form the prediction at λ_{min} , respectively with forecast length of 20s.	53
Figure 4.11	Predicted and actual breathing rate vs time using the respective cluster's λ_{min} obtained from LCV with forecast length of 10s.	54
Figure 4.12	RMSE (root mean square error) and number of sensors versus λ for all four clusters in one day obtained via cross validation on the remaining three days with forecast window of length 10s. Each row corresponds to one day, and each column corresponds to one cluster.	55
Figure 4.13	RMSE (root mean square error) and number of sensors versus λ for all four clusters in one day obtained via cross validation on the remaining three days with forecast window of length 20s. Each row corresponds to one day, and each column corresponds to one cluster.	56

CHAPTER

1

INTRODUCTION

Recent progress in technology has resulted in the development of various sensing, computation, and transmission devices which have evolved into fundamental elements of our everyday life, establishing mobile sensory platforms. These pervasive sensing devices have shown to be practical in various areas such as health monitoring which is the main application we focus on in this dissertation. The main objective is the development of a power-efficient framework for prediction of physiological responses such as breathing rate in a wearable health monitoring platform when the user is performing various activities.

Wearable health monitoring devices can capture real-time data about user's physiological responses and activities. Respiratory and cardiovascular data are the most crucial physiological information for health monitoring. Additionally, breathing rate and heart rate analyses are the most common noninvasive tools in cardiopathy and exercise physiology research [Ach06; Nic17]. Therefore, continuous monitoring of physiological signals such as respiratory rate and heart rate as well as

prediction of their future values are invaluable in detecting and diagnosing various cardiovascular and pulmonary diseases in the early stages [Maj17]. It is established in medical science that patients experience changes in vital signs before cardiac arrest. Therefore, early detection of deterioration in the patient's vital signs including heart rate and breathing rate and subsequent early therapeutic intervention is reported to reduce the probability of occurrence of cardiac arrest [Moc17]. Moreover, exercise physiology research has shown that tracking and controlling respiratory rate and heart rate during exercise can be critical for preventing overtraining [Hal14]. However, users are generally reluctant to wear many sensors or sensors that are uncomfortable to carry or interfere with their daily activities. It is therefore valuable to be able to predict some of the physiological responses such as breathing rate from other physiological or inertial measurements. Deviations of actual physiological responses from predicted values obtained using other modalities can be an indication of a potential problem. For example, if inertial measurements indicate that a user is in a resting state and environmental factors are nominal and therefore the predicted breathing rate is low but the actual breathing rate is significantly higher, this can be a sign of an anomaly in the respiratory system.

Breathing, heartbeat, blood pressure, skin tension and eye reaction are physiological responses describing the state of human bodies. While it is established that continuous monitoring of physiological responses is valuable, predicting these responses in the future is also powerful. For instance, predicting heart rate and breathing rate in endurance exercises can help athletes keep a desired intensity level during their training. Moreover, in medical applications, asthmatic patients can control their behavior such as performing physical activities, resting or staying indoors if they are aware of their future air volume intake via breathing rate prediction. In this particular application, real time information about environmental factors such as particulate matter can be helpful in prediction of the patient's physiological response. Additionally, prediction of physiological responses can help determine abnormalities in a person's health condition by tracking their inconsistencies with the expected values. Breathing rate can also be used to infer our feelings, such as happiness, fear, anxiety or calm. Physiological responses depend on various aspects of human's daily life such as

physical activities, environmental factors (e.g., temperature, air pollution, humidity) and cognitive and affective states (e.g., stress levels). As a result, predicting these responses into the future is a challenging task. Subject dependency is another main challenge in prediction of physiological responses. Each individual's physiological responses depend on their physical characteristics such as fitness level, age, gender, etc. Additionally, different individuals can have different reactions to environmental factors or can be affected differently by various activities depending on their physical characteristics. Accordingly, personalized prediction models are going to be inevitable for wearable health monitoring devices to determine the normal responses for each user in contrast with the norm for the whole population.

Breathing includes two main phases, inspiration and expiration. During inspiration, the diaphragm contracts and moves downward, and the external muscles pull the ribs upwards to enlarge the rib cage. These movements increase the volume of the chest. As a result, the air pressure in the lungs becomes lower than atmospheric air. Since air flows from a higher pressure to a lower pressure area, it circulates into the body, going through the airway into the lungs. During the expiration, the diaphragm returns to its original state and the chest restores to its lower volume. These motions push the air out of the lungs. Motor commands for the contraction of these muscles are produced in the brain neuronal networks. Different afferent inputs are integrated to generate respiration that are mainly results of metabolic demands. In addition to metabolic breathing, respiratory motor function is also affected by internal changes and environmental factors, which is called behavioral breathing. Breathing is not only a response to metabolic demands but also a regular result of changes in emotions, such as sadness, happiness, anxiety and fear [HM08]. Emotions typically result physiological changes within the whole body. Studies have shown the behavioural responses during fear and the anxiety state [Dav92]. Masaoka et al showed in [MH01] that the emotion of anxiety affects respiratory patterns without changes in metabolism. Figure 1.1 shows the relationship between breathing rate and trait anxiety scores. The figure illustrate that breathing rate and anxiety have a positive correlation ($r = 0.756, P < 0.05$). Relationships between breathing rate and emotions have shown higher breathing rate during an arousal state [Boi98]. Studies have also shown that




Figure 1.1 Relation between breathing rate (bpm) and trait anxiety scores during anticipatory anxiety. A positive correlation between the increase in breathing rate and trait anxiety scores is observed [MH01].

respiration varies when subjects look at photos, that intrigue emotional responses [Boi94]. It has also been established that there are relationships between breathing rate and emotional reactions to natural noises or unpleasant sounds have been explored in [GD04].

During physical activities and exercise, the body's need for oxygen increases. As a result of cellular respiration, the amount of carbon dioxide in the body also increases. As the concentration level of CO_2 raises, the body responds by elevating the breathing rate in order to dissipate the CO_2 . Since exercising muscles also need more oxygen, the respiratory rate also increases to expedite the transmission of oxygen to the blood stream, where it is then delivered to the muscles. Breathing rate is also shown to increase in naturally occurring asthma [Kes90]. Asthma is an inflammation in the lungs that narrows the airways, making it difficult for air to flow. The reason that asthma causes an elevated breathing rate is that it reduces the amount of air flowing in the lungs with each breath. Therefore, the patient needs to take more frequent breaths in order to provide the body with its required oxygen. Since breathing rate is affected by emotional states, physical activities, and preexisting conditions such as asthma, we are interested in measuring and using these factors in

predicting breathing rate.

There is a body of literature on estimating breathing rate from other physiological responses such as Photoplethysmography (PPG) and Electrocardiography (ECG) [Bir17; Her17]. However, they do not focus on the prediction of future values of breathing rate from the current and past physiological measurements nor do they incorporate measurements of environmental factors. Although machine learning methods are shown to be highly effective in health informatics [Dua14], the studies for breathing rate prediction have been limited. In [BG98], a simple recurrent neural network is trained to predict dynamic respiratory state in the apneic sleeping patient using heart rate, respiration force, blood pressure, blood oxygen saturation, and cardiac stroke volume data. Similarly for heart rate prediction, activity data captured by three axis accelerometer is used in [Xia10] to forecast heart rate using a feedforward neural network. The study in [YJ08] also proposes a heart rate prediction model based on the relationship between heart rate and physical activity using a feedforward neural network. ECG and arterial pressure, fasting blood sugar, heart rate, cholesterol, and age/gender are used in [Rav14] to train a neural network for classification of different heart diseases and to predict abnormalities in the heart. However, the effect of heart rate on breathing rate prediction are not evaluated in any of these studies. In [Lok17], breathing rate is predicted based on heart rate and activity information in a controlled lab environment.

Smart, wearable platforms need to satisfy certain requirements to be practical, including performance accuracy, cost, size, usability, privacy, and memory and power consumption[Lym03]. Since mobile sensory devices are battery-powered and often need to operate for extended periods of time in critical scenarios, power efficiency is one of the most significant aspects of their performance. Sensing various physiological and inertial signals, filtering and processing of the acquired signals, and transmitting extracted features all consume considerable amount of power in wireless sensory platforms. If power optimization is not considered in the algorithm design, batteries can be depleted within a shorter period of time and continuous monitoring would not be feasible. In wearable platforms, sensing and processing units require significant power [Gha15]. This calls for performing power optimization in the processing unit and optimizing the amount

of sensing in order to extend the functioning time of these devices before requiring a recharge. In machine learning regression models, different features have different importances and relevance values. Accordingly, we hypothesized that excluding less important features from the model does not significantly worsen the prediction performance but helps reduce the power consumption by disabling the corresponding sensors.

The physiological responses we aim to predict depend on the activity/motion-behavior of the subject as well as environmental factors. We therefore hypothesized that including information about subjects' activity as well as environmental measurements such as temperature, air pollution, humidity in our prediction process will improve its performance. Accordingly, we first focused on the characterization of different activities, then on the exploratory analysis of physiological response prediction, and finally an energy efficient and activity-aware physiological response framework is introduced. To also take into account both the motion information and environmental factors in prediction of physiological responses, we use multi-modal data streams including accelerometer from the wrist and ankle, body and environmental temperature, humidity, as well as heart rate, breathing rate, and electrodermal activity (EDA). Activity classification is the main component of many data analysis applications in a wide variety of fields such as fitness and health monitoring, rehabilitation, navigation and localization. In general, there are two different approaches to activity recognition. The first one is using external sensors where devices are mounted in fixed locations in the environment. The second approach, which is the focus of this study, makes use of wearable sensors that are attached to the user. Activity recognition can be performed by utilizing different sensory modalities such as vision-based ones, inertial sensors, or a combination of the them [Che12]. In this study, we make use of inertial sensors to dynamically control the prediction model in order to conserve power.

As part of our study of activity/motion behavior, we developed a motion-capture platform, and analyzed the motion behavior of cockroaches equipped with inertial sensors. We first evaluated the feasibility of an accurate, low power, wearable, sensor-based motion tracking system to monitor upper arm motion [Moh15]. The proposed system is also low cost, marker-free and easy to install

and use. The applications of wireless wearable sensors enable healthcare users to have flexibility and mobility with regards to convenience and reduce healthcare costs [Ren10]. National health care spending in the United States is estimated to reach \$4.8 trillion in 2021, which will consume nearly 20% of GDP [MS14]. Our proposed monitoring system is able to provide an economical and quantitative means of recording upper limb motion for physical rehabilitation. This wearable system uses only data from two gyroscopes for precise motion tracking of elbow joint angles while reducing the power consumption by minimizing the types of sensors used and calibrating without the need of additional sensors. We developed a motion trajectorization algorithm with the results showing that a gyroscope-based motion tracking system can be practically used to precisely track a patient's motion without requiring a large number of sensors or an excessively complex set-up.

Next, we focused on activity recognition in classifying motion modes for Madagascar Hissing cockroaches in a search and rescue scenario [Col17]. In a debris-filled disaster environment, it is vital to quickly locate and save trapped people. Traditional search operations often require other individuals to intervene, putting additional lives in danger. In response, we developed a biobotic insect platform comprised of a Madagascar hissing cockroach (*Gromphadorhina portentosa*) equipped with a battery-powered circuit board [Boz16]. Previously, these biobots have been successfully controlled via external commands to move along a desired path [LB12]. The Madagascar hissing cockroach was the prime choice for several reasons: It is small enough to fit into small crevices while it is also large enough to carry a circuit board with communication and sensing capabilities as well as a small battery. Cockroaches are exploratory insects by nature, making them ideal for such an operation. They can be easily reared and are relatively docile while being manipulated. Robotic platforms have difficulties negotiating non-uniform terrain whereas cockroaches are notoriously good at navigating cluttered areas. To use an electronically-augmented cockroach effectively as a search and rescue agent, we also need a means of localizing it. To do so, we equipped a cockroach with a custom-printed board containing a three-axis accelerometer, a two-axis gyroscope, and a wireless, low-power system-on-chip that helps reduce the power consumption. For localization purposes, it is critical to be able to determine the roach's movement state. Using only the inertial

measurement units (IMUs) mounted on the roach, we can recognize the roach's current mode of motion. Using the existing model of roach behavior and standard gait analysis techniques, we constructed a classification model that accurately predicts the gait or motion mode of the roach.

The main contributions of this dissertation are summarized below:

- Due to the dependence physiological responses to activity and motion behavior of the user, we performed an exploratory analysis of motion data and activity recognition. This exploratory analysis includes the development of a wearable sensor-based motion tracking system that minimizes the types of sensors used and is capable of calibrating without the need of additional equipment, performing additional anthropometric measurements, or measuring the location of the sensors. In order to determine how human activity recognition methods can generalize, we then developed an algorithm for classification of motion modes of Madagascar hissing cockroaches in the search and rescue application.
- Designed a protocol for evaluating motion behavior and physiological response under unconstrained condition similar to activities and exercise in daily life. Performed the collection of data using a variety of multimodal sensor devices and used this data later on for evaluation of performance and effectiveness of prediction and optimization technique developed in this study.
- Evaluated different algorithms including nonlinear regression models to accurately predict the future physiological response using multi-modal data streams that include motion information, environmental factors as well as physiological data, namely body and environmental temperature, humidity, accelerometer from the wrist and ankle, electrodermal activity (EDA), breathing rate and heart rate. The results of this analysis revealed that physiological responses are highly correlated to past measurements. We also concluded that environmental factors can help prediction of physiological responses under real-world conditions.
- Proposed an activity-aware and power-efficient framework for physiological response prediction that combines state-of-the-art techniques on sensor selection, and compared it against

standard techniques and showed that our approach provides a significant decrease on the number of sensors required on different activity modes. We investigated the tradeoff between the accuracy in prediction of physiological responses and power consumption in the wearable health monitoring platform in order to maintain prediction accuracy, while reducing the power consumption. Specifically, we developed a dynamic feature selection method by intelligently deactivating different sensors depending on the activity state of the user which is obtained using unsupervised activity classification. Using a reduced set of features in our signal processing and machine learning algorithm for physiological response prediction, we reduced the amount of power consumption in the wearable health monitoring platform.

One limitation of the study in this dissertation is its use of data from only one subject. As a result, we do not explore the variability of physiological prediction across subjects. Extensions of this work may include analysis of subject dependency in prediction of physiological responses since each user's physiological responses are affected by their health and fitness, age, gender, etc. This requires data collection on a large set of subjects of various BMI, age and gender. Moreover, our data collection protocol does not include different emotional states. Since user's feelings, such as happiness, fear, anxiety or calm affect their physiological responses such as breathing rate and heart rate, another direction for extending this work is focusing on a data collection protocol where the subject experiences excitement or stress in order to better evaluate the effect of user's cognitive and affective states in prediction of physiological responses. Furthermore, the temporal information on the breathing rate time series can be better exploited for example by using an auto-regressive model for prediction and fusing information from other modalities.

The rest of the dissertation is organized as follows: Chapter 2 introduces a wearable sensor-based motion tracking system that is able to calibrate without the need of additional equipment using a reduced number of sensors. In Chapter 3, we use multi-modal data streams including breathing rate, body and environmental temperature, heart rate, humidity, accelerometer from the wrist and ankle and electrodermal activity (EDA) to predict breathing rate over varying forecast lengths. In Chapter 4, we focus on power optimization in predicting breathing rate using a dynamic feature

selection method by deactivating different set of features depending on the activity state of the user. This implementation results in significant reduction in power consumption while retaining similar prediction accuracy.

CHAPTER

2

FEASIBILITY OF A WEARABLE, SENSOR-BASED MOTION TRACKING SYSTEM

The content of this chapter is adapted from the conference paper published in the 6th International Conference on Applied Human Factors and Ergonomics (AHFE 2015) [Moh15].

2.1 Introduction

The advent of smartwatches marked the beginning of minimally obtrusive health monitoring platforms. Most of the healthcare systems were designed for different types of patients based on age, disease type, biological signals measured, and other factors. Target users of previous healthcare

systems included patients, doctors, therapists and others. Lee et al. developed a healthcare monitoring system for elderly clinical and trauma patients [Lee06] which records physiological parameters such as ECG and body temperature. Their platform connects wirelessly to a base station to transmit data if an abnormal response is detected. Shnyder et al. created a medical monitoring platform consisting of custom developed pulse oximeter, two-lead ECG, and an IMU-based sensorboard [Shn05]. Navarro et al. designed a monitoring system for elderly and infirm patients [Nav09], and Suryadevara et al. integrated a wearable sensor-based healthcare system to monitor health perception and daily activity behavior for the elderly [Sur12]. Monitoring systems can track a person's individual limbs by strategically placing a set of inertial measurement units (IMUs) on them. Motion tracking techniques are applied in many fields, ranging from animation [Zha13] to clinical applications [Whi07; Bou12]. Current sensor-based motion tracking systems use a variety of sensors aimed to monitor the motion patterns of a patient. These IMUs are generally comprised of accelerometers and gyroscopes, but are often also augmented by magnetometers. In addition to tracking 3D motion, wearables-based monitoring systems can also create inferences about the patient's motions. For example, simple inferences can be made by merely determining whether movements are above a certain threshold. This may provide a crude estimate of a person's state of activity/inactivity. Slightly more sophisticated algorithms that detect periodicity in a patient's movement can suggest whether they are still, walking, or running. For purposes of remote health monitoring, a broader range of states are needed to make diagnostic inferences. This can be achieved by combining a wide set of sensor modalities and more advanced machine learning algorithms. Such techniques require extensive research and training data in order to understand a wide range of activities and to accurately classify the patient's current activity.

Previous work has investigated the feasibility of using low-cost, off-the-shelf sensors to track upper limb motion of a user. Often, this is achieved by mounting accelerometers, gyroscopes, and magnetometers and multiple key locations on the body. Current sensor-based motion tracking systems use a variety of sensors aimed to monitor the motion patterns of a patient. Generally, gyroscope and accelerometer data are combined to determine the pose (orientation and position)

of a tracked body segment, modeled as a rigid body [Shn05]. Aside from the extra power consumed from continuously sampling multiple sensors, using an accelerometer can produce extra sources of errors, primarily because the accuracy of the measured acceleration is highly sensitive to the accuracy of the measured rotation. It also becomes necessary to correctly measure the location of the sensor on the body in these studies. In addition to the gyroscope and accelerometer, motion tracking systems may use other sensors such as magnetometers to calibrate the orientation of the sensor which can further increase the cost of these systems. To avoid these issues, the wearable sensor-based motion tracking system introduced in this chapter minimizes the types of sensors used to track a patient's upper limb motion and is capable of calibrating without the need of performing additional anthropometric measurements or measuring the location of the sensors. Although there are other motion tracking systems on the market (e.g., Xsens [Zha13]), our motion tracking system is low cost, marker-free, and easy to install and use.

While the electronics driving these wireless systems have become sufficiently small for unobtrusive monitoring, battery and energy harvesting technologies are still unable to power these devices for more than a few hours. For example, in the case of the previously described platform, Shnayder et al. acknowledge that their monitoring system only has limited battery life with some components lasting only 5 hours. While recharging and/or replacing batteries frequently may be acceptable in medically deployed devices, a remote health monitoring system should remain relatively unobtrusive and require minimal maintenance from the user. For this reason we will focus on techniques that can be employed to reduce power consumption in wireless, wearable sensor networks. The wearable sensor-based motion tracking system introduced in this chapter minimizes the types of sensors used and is capable of calibrating without the need of additional equipment, performing additional anthropometric measurements, or measuring the location of the sensors.

This chapter is structured as follows: Section 2.2 outlines our IMU system, experimental set-up, and the algorithm used to extrapolate the upper limb motion from IMU data. In Section 2.3, the experimental procedure and the method of analyzing the data is described in detail. Section 2.4 presents the results and discusses the findings. Lastly, Section 2.5 explains the significance of the

research and describes its potential for further investigation.

2.2 Development of the wearable motion tracking system

2.2.1 Overview

We developed an upper limb motion tracking system using three low-cost SensorTag wireless sensor units from Texas Instruments. The on-board IMUs were used to track the upper arm and forearm, transmitting the sampled data via the device's Bluetooth Low Energy (BLE) module. The bluetooth module used on the SensorTag is the CC2541 by Texas Instruments. Bluetooth Low Energy is a new Bluetooth specification that addresses the need for robust wireless communications under low power settings with a focus on consumer and healthcare applications [Gom12]. With this set-up, it is possible to use a Bluetooth enabled laptop to connect to and gather data from multiple SensorTags simultaneously. The SensorTag uses the IMU-3000 gyroscope by InvenSense. Each gyroscope was set to a range of ± 250 degrees per second with a sampling frequency of 50 Hz, which is sufficient to track a subject's motion with high accuracy. The SensorTag uses the KXTJ9 accelerometer by Kionix. Each accelerometer has its range set to $\pm 4G$ and its sampling rate set to 5 Hz. To validate the accuracy of our system and model, we used the OptiTrack System, a commercially available, marker-based, optical motion capture system by NaturalPoint. It uses a Velcro suit to which 36 retroreflective markers were attached. These markers were tracked by 12 specialized cameras equipped with infrared LEDs. Figure 2.1 outlines the underlying principle guiding the experiment through which our system's accuracy will be assessed. The OptiTrack system produces joint trajectories with a temporal resolution of 100 Hz and a spatial resolution of 1 millimeter [Ker12].

The Velcro suit also offers a means for attaching the IMU units to the subject as shown in Figure 2.2. A chest IMU was also used during the motion capture recording, but data collected from it was not used in the data analysis. A video camera was set up to record the subjects' motions. This was implemented in order to facilitate time synchronization between the IMU and the OptiTrack motion capture systems when processing the data. The set-up is shown in Figure 2.2.

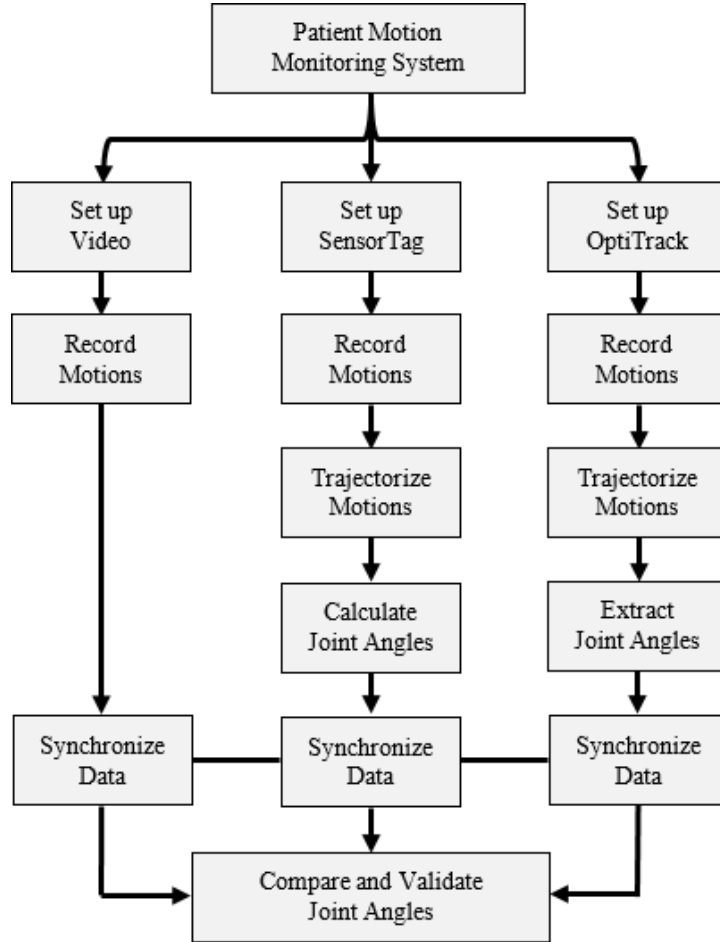


Figure 2.1 System Overview Flowchart.

2.2.2 Tracking algorithm

The human body’s posture can be fully described by knowing all of its joint angles. The data obtained from an IMU’s gyroscope represent the angular velocity (ω_t) given in x-, y-, and z-rotation components with respect to the IMU’s own reference frame. In order to calculate the sensor’s angular displacement, these velocity samples can be chained together akin to mathematical integration. Due to the high sensitivity of the accuracy of the orientation (angular position) to sampling error, we use quaternions to represent rotations instead of Euler angles since the latter contain points of singularities [Die06]. These singularities are equivalent to the phenomenon of gimbal lock and any

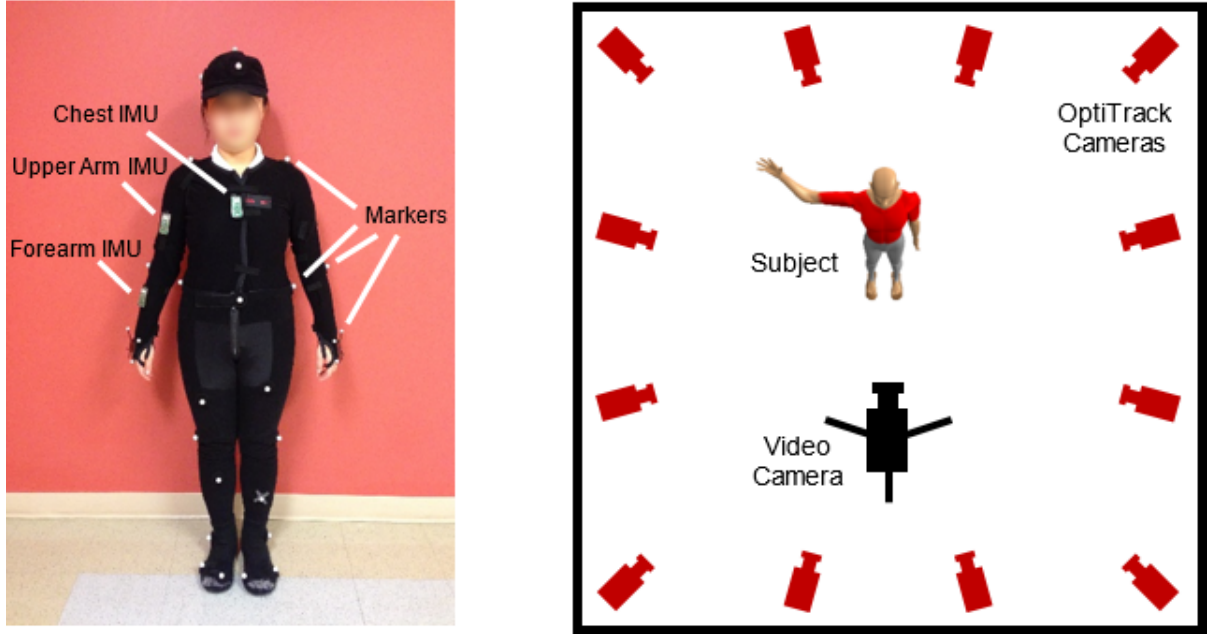


Figure 2.2 Experiment set-up. Left: Subject wearing OptiTrack Velcro suit with markers and IMU sensors; Right: Experiment layout with 12 OptiTrack cameras.

small error in measurement or rounding near these points will be magnified immensely, yielding poor estimations of orientation. By knowing the sampling time interval (Δt), we can calculate the quaternion representing angular displacement between samples as shown in Equation 2.1.

$$r_{S_{t-1}}^{S_t} = \cos\left(\frac{\theta_t}{2}\right) + \sin\left(\frac{\theta_t}{2}\right) \cdot (\omega_{tx}i + \omega_{ty}j + \omega_{tz}k) \quad (2.1)$$

Where $\theta_t = |\omega_t| \Delta t$ denotes the angle step, script S denotes the sensor frame, subscript t is the current time step, and i, j, and k are the quaternion basis elements. In this convention, r_a^b is the rotation quaternion describing the orientation of frame b with respect to frame a. The process of trajectorizing the sensor's orientation can be achieved as follows:

$$r_{S_0}^{S_t} = r_{S_0}^{S_1} \cdot r_{S_1}^{S_2} \cdot \dots \cdot r_{S_{t-1}}^{S_t} \quad (2.2)$$

Here, $r_{S_0}^{S_t}$ represents the orientation of the current sensor frame (S_t) relative to the initial sensor

frame (S_0). Since the orientation of the IMU relative to the subject's limb is arbitrary or unknown, a calibration step needs to be performed to relate the two. While a multitude of calibration methods are described throughout the literature which rely on additional sensors [Pér10; Zho08], we developed a unique calibration process that only requires the subject to perform a known, specified motion recorded using only the gyroscope and the accelerometer. The subjects were instructed to hold a neutral stand posture with the arm pointing straight down. In this posture, the SensorTag's accelerometer was used to detect the direction of gravity. Next, the subjects were asked to flex their shoulder while keeping the elbow fully extended. These motions are sufficient to determine the orientation of the arm with respect to the world and the orientation of the IMU with respect to the arm. The elbow joint angle can be calculated by knowing the orientation of the forearm and upper arm in space.

2.3 Method

2.3.1 Subject

Eight healthy adults (4 female, 4 male, age 22 - 28 years) participated in the experiments. The average height was 169.69 cm (SD = 10.14 cm). All of the subjects were right-handed. None of the subjects had any previous upper limb injuries.

2.3.2 Experiment Procedure

After setting up the video camera and the twelve OptiTrack cameras surrounding the recording arena, subjects were asked to don the OptiTrack motion capture suit and cap. Next, the 36 retroreflective markers were attached according to the OptiTrack's marker guide. Three SensorTags were then fastened to the body as depicted in Figure 2.2 - one to the forearm, one to the upper arm, and one to the chest - using the suit's Velcro fabric. The suit is elastic, designed to conform to the user's body while preventing attached markers from slipping. This feature was also exploited to rigidly secure the IMUs to their respective body segments. The IMUs were powered on, and a PC with




Figure 2.3 Elbow flexion-extension motion.

a BLE host dongle (BLED112 USB Bluetooth dongle from Bluegiga) was used to connect to each sensor. All motions that each subject had to perform were first described to them before data collection began. To maintain consistency between subjects, the motions that each subject had to perform were presented in front of them by a demonstrator so that they could simply mirror the motions. At the beginning of each recording session, the first motion that the subjects performed was a rapid abduction-adduction of the shoulder. This motion allows for easy identification of the exact frame of highest arm elevation in the IMU and OptiTrack data. This helps to synchronize the measurements between the two systems. Subsequently, the IMU calibration procedure was executed as described in the tracking algorithm. Next, repeated elbow flexion-extension motions were performed as depicted in Figure 2.3. Each recording session consisted of ten such cycles.

2.3.3 Data collection and analysis

During the experiment, eight subjects were asked to perform ten elbow flexion-extension cycles as described above. Data from our motion tracking system and OptiTrack motion capture system were recorded. The trajectorization algorithm was used to obtain the elbow joint angles. To validate our motion tracking system, the trajectorized data of the IMU system were compared to those of the OptiTrack system. To evaluate the similarity between measurements of our system and OptiTrack system, correlation coefficient (CC) analysis was used. Root-mean-square error (RMSE) was calculated to indicate the difference of joint angles between the two systems [Zho08]. As a scale for how much the difference in angle is distributed, the median absolute deviation (MAD) was calculated.

2.4 Results and Discussion

The results showed that our motion tracking system measured the elbow flexion-extension (joint angles) with high consistency to those of the OptiTrack. Figure 2.4 shows the results of all subjects obtained from our IMU system and from the OptiTrack system.

Figure 2.5 shows a single flexion motion for one subject with a visually recreated representations of three key frames of the motion. For the joint angles measured from the two systems, results of eight subjects are shown in Table 1. The correlation coefficient (CC) for joint angles of two systems was high (> 0.99 for all subjects). The high CC presents a strong correlation between two data sets. RMSEs for all subjects were low ($2.06^\circ - 5.53^\circ$). Also, MADs between measures from two systems were low for all subjects ($2.21^\circ - 3.47^\circ$).

Compared to previous studies, our system is able to provide quantitative measurements of elbow motion with relatively high accuracy using wearable sensors. Moreover, due to the simplicity of the presented system with regards to the number of types of sensors used and the ease of donning the sensors, our IMU motion tracking system provides notable advantages over previously developed systems. The readily available hardware and software permit this system design to be a practical




Figure 2.4 Elbow flexion-extension motion.

solution for any limb motion monitoring system.

2.5 Conclusion and Extensions

In this study, we developed a wearable, sensor-based motion tracking system and elbow joint angles of eight subjects were tracked using two IMUs attached to the upper arm and forearm. A motion trajectory algorithm was developed that is able to calculate elbow joint angles. This algorithm can be used with arbitrary placement of the IMUs as long as each IMU is placed on the appropriate limb. Thus, no manual anthropomorphic measurements need to be performed. Results of the elbow flexion-extension experiment yielded results consistent with that of the optical motion tracking

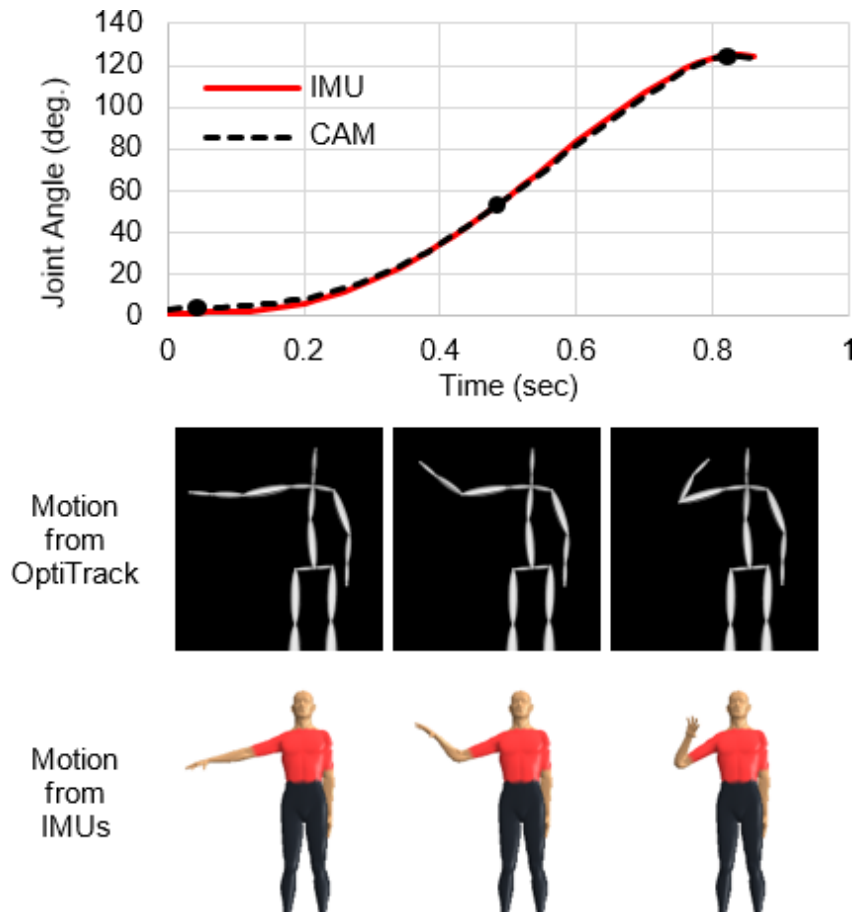


Figure 2.5 Joint angle and visualization of an elbow flexion motion.

system. Specifically, the IMU-based tracking system and the camera-based tracking system yielded RMSE values between 2.06° and 5.53° . These results suggest that a gyroscope-based motion tracking system can be realistically used to accurately track a patient's motion without the need for numerous sensors or an overly complicated set-up. In our experiment, we only used data from two IMUs to track the elbow joint angle for some simple motions for one arm. A continuation of this work should include more IMUs to track more body parts and joint angles. Additionally, integration of this motion tracking system with advanced kinematic modeling software, such as OpenSim can further aid in processing, visualizing, and analyzing patient motions to aid in physical rehabilitation.

Table 2.1 Results for all subjects.

Subject	1	2	3	4	5	6	7	8
CC	0.99	0.99	0.99	1.00	1.00	0.99	1.00	0.99
RMSE (°)	2.46	3.14	2.06	3.07	3.87	4.01	4.53	5.53
MAD (°)	2.21	2.62	2.23	2.38	3.22	2.69	3.47	3.12

CHAPTER

3

PREDICTION OF PHYSIOLOGICAL
RESPONSE OVER VARYING FORECAST
LENGTHS WITH A WEARABLE HEALTH
MONITORING PLATFORM

The content of this chapter is adapted from the conference paper published in the 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC'18) [Moh18].

3.1 Introduction

Respiratory and cardiovascular data are the most crucial physiological information for health monitoring. Additionally, breathing rate and heart rate analyses are the most common noninvasive tools in cardiopathy and exercise physiology research. Therefore, continuous monitoring of these physiological responses and predicting their future values are highly valuable. There is a body of literature on estimating breathing rate from other physiological responses such as Photoplethysmography (PPG) and Electrocardiography (ECG) [Bir17; Her17]. However, they do not focus on the prediction of future values of breathing rate from the current and past physiological measurements. Although machine learning methods are shown to be highly effective in health informatics [Dua14], the studies for breathing rate prediction have been limited. In [BG98], a simple recurrent neural network is trained to predict dynamic respiratory state in the apneic sleeping patient using heart rate, respiration force, blood pressure, blood oxygen saturation, and cardiac stroke volume data. Similarly for heart rate prediction, activity data captured by three axis accelerometer is used in [Xia10] to forecast heart rate using a feedforward neural network. The study in [YJ08] also proposes a heart rate prediction model based on the relationship between heart rate and physical activity using a feedforward neural network. However, the effect of heart rate on breathing rate prediction are not evaluated in any of these studies. In [Lok17], breathing rate is predicted based on heart rate and activity information in a controlled lab environment.

Selective sampling methods are investigated in [Fre07] to conserve power in wearable sensor platforms. Additionally, the study presented in [Zap08] uses sensor selection for energy efficient activity recognition by minimizing the number of sensors required to achieve a desired classification accuracy. For activity classification using wearable sensors, the authors in [GS13] implemented a power-aware feature selection framework. In this study, we use multi-modal data streams including breathing rate, body and environmental temperature, heart rate, humidity, accelerometer from the wrist and ankle and electrodermal activity (EDA) to predict breathing rate. Experimental results are carried out using data we collected from an individual going through his regular daily routine for four

days. We utilize a nonlinear Support Vector Machine (SVM) regression model to predict the future breathing rate and heart rate using the measurements described above. Since the multi-modal data streams used in this study are all measured simultaneously in real time, our framework has the potential to be utilized in a variety of applications such as cardiopathy research and diagnosis, heart attack danger notice, sports capability measure, and mental activity assessment.

The remainder of this chapter is organized as follows: In Section 4.2, we describe our dataset and the devices used for data collection. We present our problem statement and the prediction model in Section 4.3. Experimental results are included in Section 4.4, and Section 4.6 concludes the chapter.

3.2 Dataset

3.2.1 Sensors

The platform is composed of a number of sensors.

AirBeam: The AirBeam¹ is a portable, environmental sensor capable of measuring particulate matter, humidity, temperature, environmental sound levels and GPS coordinates. The particulate sensor uses a light-scattering technique to detect PM2.5 particulate matter in the environment, an air pollutant whose high levels have negative effects on a population's health [Xin16]. As air flows through the device, an LED shines light that scatters off particles in the air stream. The amount of scattered light can be detected and used to measure the quantity of particles. PM2.5 is one of the main air pollutants regulated by the US Environmental Protection Agency [Uni17] that threatens the public's health. In this study, we only use environmental temperature and humidity measured by this sensor with a sampling rate of 1Hz.

BioHarness: The BioHarness is a device developed by Zephyr². It contains a wearable, health-monitoring module designed for measuring, storing, and transmitting physiological data, including a single-lead ECG, heart rate, respiratory rate, and activity levels. The main advantage of using this

¹<http://aircasting.org>

²<https://www.zephyranywhere.com>

device over other physiological sensors is that this device is capable of acquiring measurements reliably even during intense activities and heavy exercises. The BioHarness consists of a chest strap integrated in a sports shirt and a detachable module secured to the strap. The conductive pads integrated in the module measure ECG in mV sampled at 250Hz from which we extracted heart rate in BPM at sampling rate of 1Hz. The left side of the strap contains a pressure sensor that measures the chest expansion due to breathing to provide respiratory rate waveform. An internal three axis accelerometer measures the subject's activity level. The output values of the BioHarness have been shown empirically to have a strong relationship to actual values of heart rate, breathing rate, and acceleration ($r = 0.89$ to 0.99 , $p < 0.01$) [Joh12]. In this study, we use heart rate and breathing rate sampled at 1Hz.

Empatica E4: The Empatica E4³ wristband is a wearable wireless device suitable for continuous physiological data acquisition. The E4 device includes photoplethysmography (PPG) sensor that provide blood volume pulse at 64 Hz using light in the red and green wavelength spectrum. The PPG signal is computed internally in the device by subtracting the green signal from the red signal, with the justification of reducing baseline drift. Heart rate, heart rate variability, and other cardiovascular features are obtained from the provided PPG. The E4 also includes electrodermal activity (EDA) sensor, that measures sympathetic nervous system arousal that is used to obtain information about subject's stress, engagement, and excitement level. Additionally, three axis accelerometer and infrared thermopile, reading body temperature are integrated in the device. These sensors generate blood volume pulse signal sampled at 64Hz, inter beat interval time, electrodermal activity and skin temperature sampled at 4 Hz, and three axis raw acceleration at 32Hz. We use body temperature, acceleration of the wrist, and EDA as measured by this device.

SensorTag: The TI SensorTag⁴ is a development board by Texas Instruments containing ambient light, ambient temperature, humidity, barometer, accelerometer, gyroscope, and magnetometer sensors. It is most notable for being a low power device capable of operating for a year under its recommended, nominal operation on a single C2032 coin cell battery. The sensors used for this

³<https://www.empatica.com/en-eu/research/e4/>

⁴<http://www.ti.com/sensortag>

study include the on-board accelerometers and gyroscopes. Although the default sampling rate of the SensorTag is 1Hz, we modified its firmware to sample and transmit at 100Hz in order to capture higher frequency movement features involved in fast-paced activities.

3.2.2 Data Acquisition Protocol

We recorded physiological, activity and environmental data on a healthy, 30-35 year old male subject over a period of four days. The subject was wearing the four aforementioned sensors in the following manner: The E4 was worn on the left wrist, following the manufacturer's instructions. The SensorTag was worn on the right ankle to capture lower body movement. The BioHarness was worn using the accompanying sports shirt. The Airbeam was carried on the subject's back except during workout routines in which the device could be obstructive. Since each device keeps track of time locally, the subject was asked to jump at the beginning and at the end of each recording session. This event is captured by all devices and is used to synchronize all data channels. The subject performed periods of different exercises as part of his daily routine, such as walking and running, which lead to highly varying heart rate and breathing rate data in an uncontrolled environment.

3.3 Prediction of Physiological Responses

We now present the problem, summarize the features used and physiological responses measured from the subject, and review the prediction model.

3.3.1 Problem Description

We make use of the following set of features acquired from the subject: The breathing rate in breaths per minute, body temperature in Fahrenheit, heart rate in beats per minute, environmental temperature in Fahrenheit, relative humidity in percentage, three axis accelerometer data collected on the wrist and ankle, and EDA measurements. For each axis of the accelerometer, we compute the mean, variance, covariances, skewness, kurtosis, peak spectral magnitude, and the corresponding peak frequency over a window of 5 seconds. For EDA measurements, we extract the linear trend,

variance, skewness, Kurtosis, and range between maximum and minimum values over a window of 30 seconds. Heart rate and respiratory rate are available at 1 Hz. The objective is to predict breathing rate using the features described above, so the features are computed every second for the corresponding window size. Prediction of physiological responses at each instance of time in the future is done using the regression model described below.

3.3.2 Prediction Model

In this study, we use Support Vector Machines (SVM), which are tools from statistical learning theory and were originally developed by Vladimir Vapnik in 1995 . SVM can be used for both classification and regression, and can be linear or nonlinear which is a kernelized modification of linear SVM.

3.3.3 SVM Regression

In this section, we briefly describe linear and nonlinear Support Vector Machines (SVM) regression methods. We first explain linear SVM regression and then provide an overview of the more capable nonlinear SVM which is a kernelized modification of linear SVM.

3.3.4 Linear SVM Regression

Suppose the training set is represented as $\{x_i, y_i\}, i = 1, 2, \dots, l$, where $x_i \in \mathbb{R}^m$ is the features or predictors and $y_i \in \mathbb{R}$ is the target or response obtained given an input x_i . Assuming that x and y are linearly related, there exists a vector $w \in \mathbb{R}^m$ such that $f(x) = w \cdot x + b$. The SVM regression problem is formulated as

$$\min \frac{1}{2} \|w\|^2 \quad (3.1)$$

subject to $y_i - w x_i - b \geq \epsilon$ and $w x_i + b - y_i \geq \epsilon$. Since existence of such $f(x)$ is not guaranteed, some errors are tolerated to have a possible solution. This is incorporated by adding slack variables ζ_i and $\zeta_i^*, i = 1, 2, \dots, l$ to the optimization problem as $\min \frac{1}{2} \|w\|^2 + C \sum_{i=1}^l (\zeta_i + \zeta_i^*)$ subject to $y_i - w x_i - b \geq \epsilon + \zeta_i, w x_i + b - y_i \geq \epsilon + \zeta_i^*$ and $\zeta_i, \zeta_i^* \geq 0$, where C is the penalty factor.

Using Lagrangian multiplier λ_i and λ_i^* , $i = 1, 2, \dots, l$ for the constraints results in the dual problem below:

$$L_D = -\frac{1}{2} \sum_{i=1}^l \sum_{j=1}^l (\lambda_i - \lambda_i^*) (\lambda_j - \lambda_j^*) (x_i x_j) - \epsilon \sum_{i=1}^l (\lambda_i + \lambda_i^*) + \sum_{i=1}^l y_i (\lambda_i - \lambda_i^*) \quad (3.2)$$

subject to

$$\sum_{i=1}^l (\lambda_i - \lambda_i^*) = 0 \quad (3.3)$$

$$\lambda_i, \lambda_i^* \in [0, C].$$

After solving λ_i and λ_i^* , the function for predicting new values depends only on the support vectors:

$$f(x) = \sum_{i=1}^l (\lambda_i - \lambda_i^*) (x_i x) + b. \quad (3.4)$$

3.3.5 Nonlinear SVM Regression

In numerous applications, the observed targets are not linearly related to the predictors. Linear SVM is generalized to adapt to numerous applications where the observed targets are not linearly related to the predictors. Intuitively, the data points are mapped to a higher dimensional space using so called kernel functions and linear regression is applied in that space.

The objective function L_D for nonlinear SVM regression is modified by replacing $(x_i x_j)$ with $\phi(x_i)\phi(x_j)$, where $\phi(\cdot)$ is the aforementioned transformation function. In theory, the regression problem would be resolved by obtaining the transformation function ϕ . Direct calculation of $\phi(x_i)$ is however computationally expensive for high dimensional problems where the number of features is large. Since the nonlinear objective function L_D requires the product of $\phi(x_i)\phi(x_j)$ and not $\phi(x_i)$ and $\phi(x_j)$ separately, computation of the transformation function ϕ can be avoided by selecting $\phi(\cdot)$ such that the inner product of any pair of points in the higher dimensional space can be represented

as a function of the pair of points in the initial space. In other words, if a kernel function K exists such that $\phi(x_i)\phi(x_j) = K(x_i, x_j)$, it would be sufficient to calculate $K(x_i, x_j)$ instead of directly finding $\phi(x_i)$ and $\phi(x_i)\phi(x_j)$ in order to find the optimal function $f(x)$ in the transformed higher dimensional space [Vap95].

The dual problem for nonlinear SVM regression can be obtained by replacing $(x_i x_j)$ with $K(x_i, x_j)$. Thus, nonlinear SVM regression finds the coefficients that minimize

$$L_D = -\frac{1}{2} \sum_{i=1}^l \sum_{j=1}^l (\lambda_i - \lambda_i^*)(\lambda_j - \lambda_j^*) K(x_i, x_j) - \epsilon \sum_{i=1}^l (\lambda_i + \lambda_i^*) + \sum_{i=1}^l y_i (\lambda_i - \lambda_i^*) \quad (3.5)$$

subject to

$$\sum_{i=1}^l (\lambda_i - \lambda_i^*) = 0 \quad (3.6)$$

$$\lambda_i, \lambda_i^* \in [0, C].$$

The function for predicting new values in nonlinear SVM regression is represented as:

$$f(x) = \sum_{i=1}^l (\lambda_i - \lambda_i^*) K(x_i, x) + b. \quad (3.7)$$

The most common nonlinear kernels are Gaussian and Polynomial. Gaussian kernel SVM is a nonparametric technique while polynomial kernel is parametric. Gaussian kernel is formulated as $K(x_i, x_j) = \exp(-\|x_i - x_j\|^2)$ while the polynomial kernel function is represented as $K(x_i, x_j) = (1 + x_i' x_j)^q$, where $q \in \{2, 3, \dots\}$. Since Gaussian kernel SVM is a nonparametric technique that assigns a score proportional to the nearness of the query point to the support vector points, it performs better than other types of kernels for prediction of highly varying data [KL03]. Due to these reasons and experimental evaluation of different kernels, we choose Gaussian kernel in nonlinear SVM regression framework for this study.

3.4 Experimental Results

We now present the findings of our experiments and report the breathing rate prediction performance. The complete list of features used in this study is described in Section 4.3. We compare the results for the prediction using different feature combinations. We also compare the results to a constant predictor which simply uses the current value of breathing rate or heart rate as a naive estimate of the future value.

One of the challenges in applying regression techniques for prediction of the physiological responses is that we cannot select training and test sets randomly which is a common method in machine learning for separating training and test set. The reason is that the data points in our problem include measurements of breathing rate and heart rate, which are not independent between samples. We therefore choose training and testing set segments to be two segments of data that are separate from each other in time to minimize their correlation and influence. In this experiment, we select the first 80% of the data from a single day as the training set and the last 20% of the data from the same day for testing. We train and test with data from a single day in order to avoid the variation across days due to changing environmental factors. We would be able to use the models across days if more repetitions were available (e.g., perhaps by observing the individual for a month). On the training set, we use 5-fold cross validation to select the parameters of the SVM regression model. We normalize the values for all the features and target to have zero mean and standard deviation of one before applying the regression in order to avoid the bias. We then apply to the test set the model that we validated on the training set.

Figure 3.1 shows the breathing rate prediction performance on the test data using heart rate, breathing rate, acceleration, EDA, body temperature, environmental temperature and humidity with the forecast length of 10 seconds. The first row of the Figure shows the predicted and actual future breathing rate values versus time, as well as the results from the constant predictor. The predicted breathing rate values are closer to the actual future values than the constant predictor, which validates our prediction performance. The second row compares the residuals of the trained

prediction to the residuals of the constant predictor as breathing rate varies. The three curves on the plots show the 5, 50 and 95 percentiles of the residuals, which are narrower in the trained predictor than the constant predictor. Note that there is slight positive trend (i.e., lower breathing rates are overestimated, and higher rates are underestimated). The third row of the figure shows the histograms of the residuals with the mean residual shown with a dotted vertical line, illustrating that the error of our predictor is has a smaller variance around zero than the constant predictor. The fourth row represents the absolute error as a function of breathing rate, showing that the predicted values have less absolute error compared to the constant predictor. Figure 3.2 shows the breathing rate prediction performance using the same set of feature as Figure 3.1 with the forecast length of 20 seconds. Although the error increases as the length of the forecast window increases from 10s to 20s, our prediction results are more accurate than the constant predictor despite the fluctuations of the breathing rate values with time. However, the relative improvement between the methods has decreases substantially due to the fact that predicting for longer forecast is a more challenging task that is also less deterministic.

Figure 3.3 compares the breathing rate prediction error for different sets of features as the length of forecast window increases. The Normalized Root Mean Squared Error (NRMSE) defined as
$$\text{NRMSE} = \frac{\sqrt{\frac{1}{n} \sum_{t=1}^n (\hat{y}_t - y_t)^2}}{y_{\max} - y_{\min}}$$
 is employed for performance evaluation as it is widely used in regression studies. These findings indicate that we can predict future breathing rate more accurately using all the measured features. Figure 3.3 also shows that not including previous values of breathing rate in the regression model increases the prediction error. Actually, the performance of methods that do not include breathing rate seem to be the same for all forecast lengths suggesting that they are not help for prediction on their own. Moreover, the constant predictor performs worse than all the predictors that used current values of breathing rate. Therefore, having access to heart rate and other physiological and inertial measurements in addition to current breathing rate can improve the accuracy of future prediction. Similar trends were observed in [Lok17] in a controlled lab environment.

We repeated the experiment for prediction of breathing rate on each of the four days of data

collected. Figure 3.4 shows the residual values of prediction versus actual breathing rate values for each day of the experiment using the forecast length of 20 seconds. According to this Figure, a similar pattern is observed for each day.

3.5 Conclusion and Extensions

In this study, we predicted breathing rate using multi-modal data streams collected via a wearable health monitoring platform. These features include body and environmental temperature, humidity, accelerometer from the wrist and ankle, EDA, breathing rate and heart rate at current and previous times. Our experimental results revealed that using other physiological responses and activity information captured by inertial measurements in the regression model improve the breathing rate prediction accuracy. An extension of this work could investigate how to better exploit temporal information on breathing rate time series (e.g., by using an auto-regressive model for prediction and fusing information from other modalities). Another extension of this work can focus on a protocol where the subject experiences excitement or stress in order to better evaluate the effect of EDA features in prediction of physiological responses. By acquiring and analyzing larger datasets including more modalities, we can clarify the dependencies of physiological responses on other factors. Another continuation of this research could be evaluating how motion artifacts in the sampled data affect prediction accuracy and how these artifacts could be detected filtered out. In the next chapter, we present the results of physiological response prediction on a larger dataset.

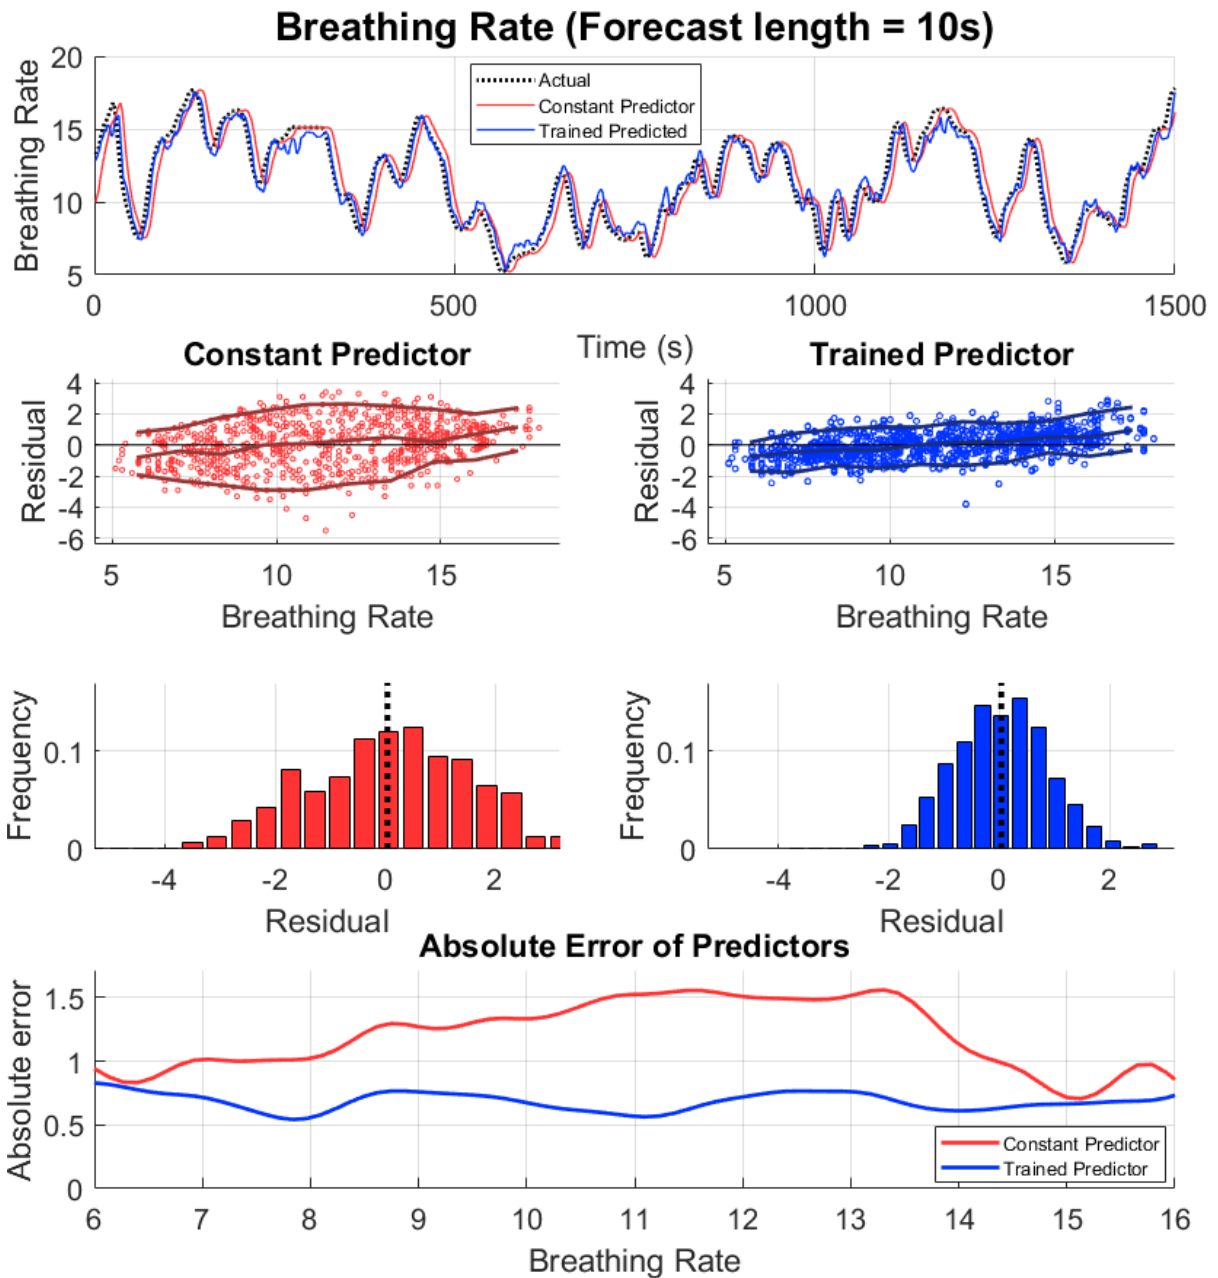


Figure 3.1 Breathing rate prediction results using all sensing modalities with a 10 second forecast. First row: results from trained and constant predictors. Second row: residuals of the trained and constant predictors as a function of breathing rate. The curves represent the 5, 50 and 95 percentiles of the residuals. Third row: histograms of the residuals for the trained and constant predictors. The mean is shown as a dotted line. Fourth row: mean absolute error of the trained and constant predictors as a function of breathing rate.

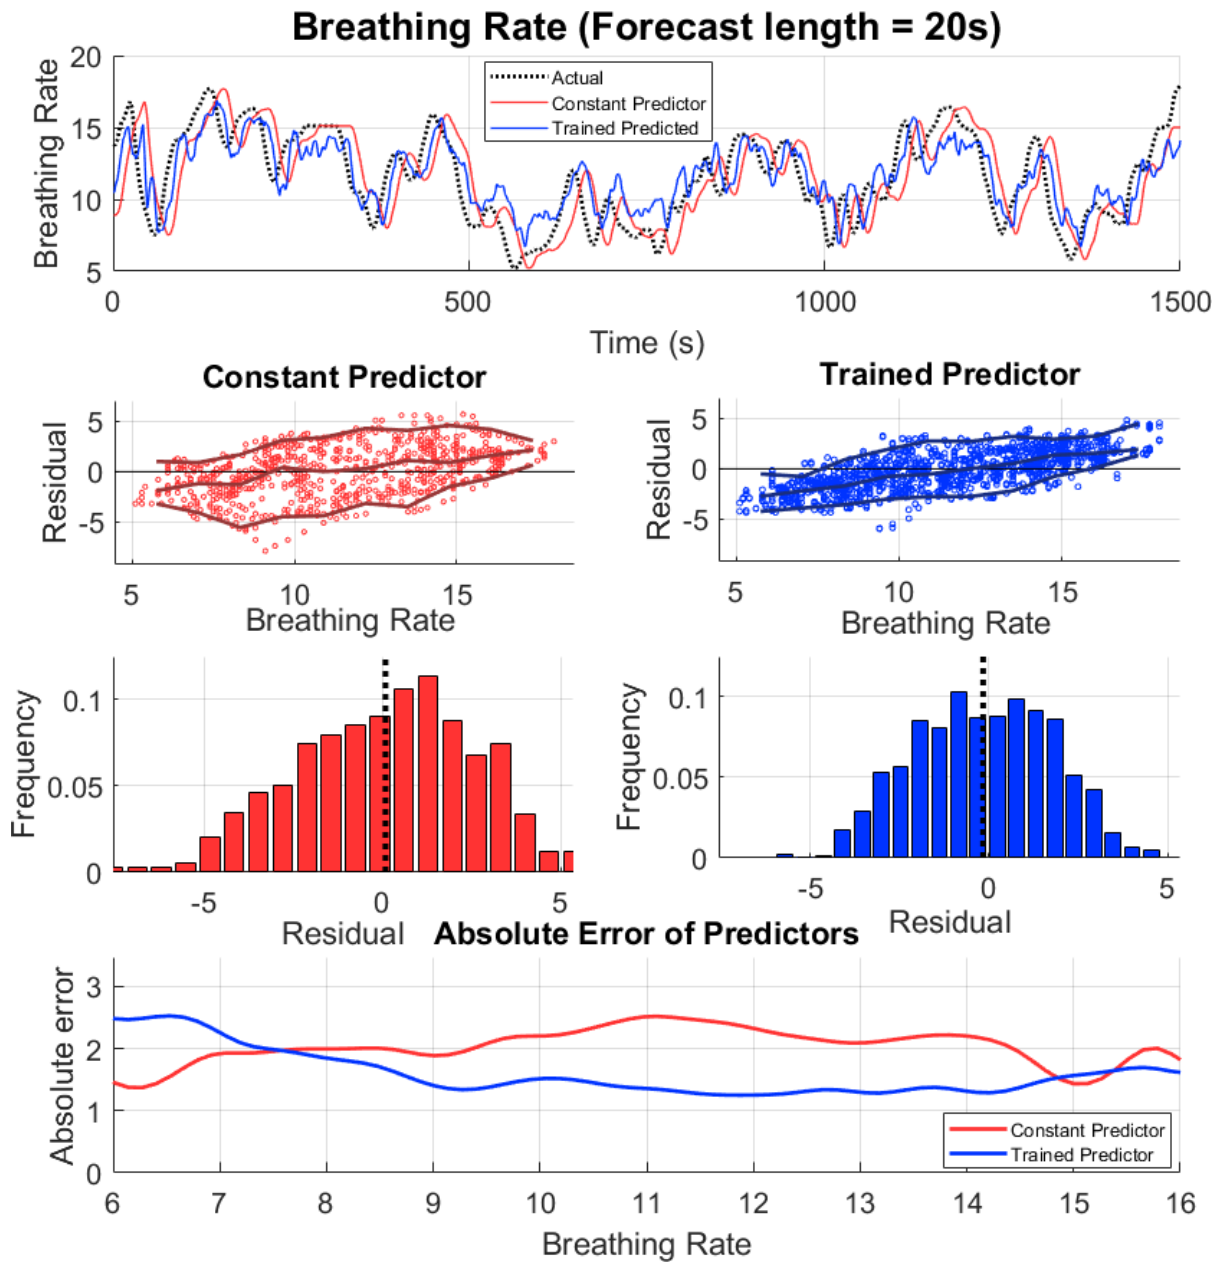


Figure 3.2 Breathing rate prediction results using all sensing modalities with a 20 second forecast.

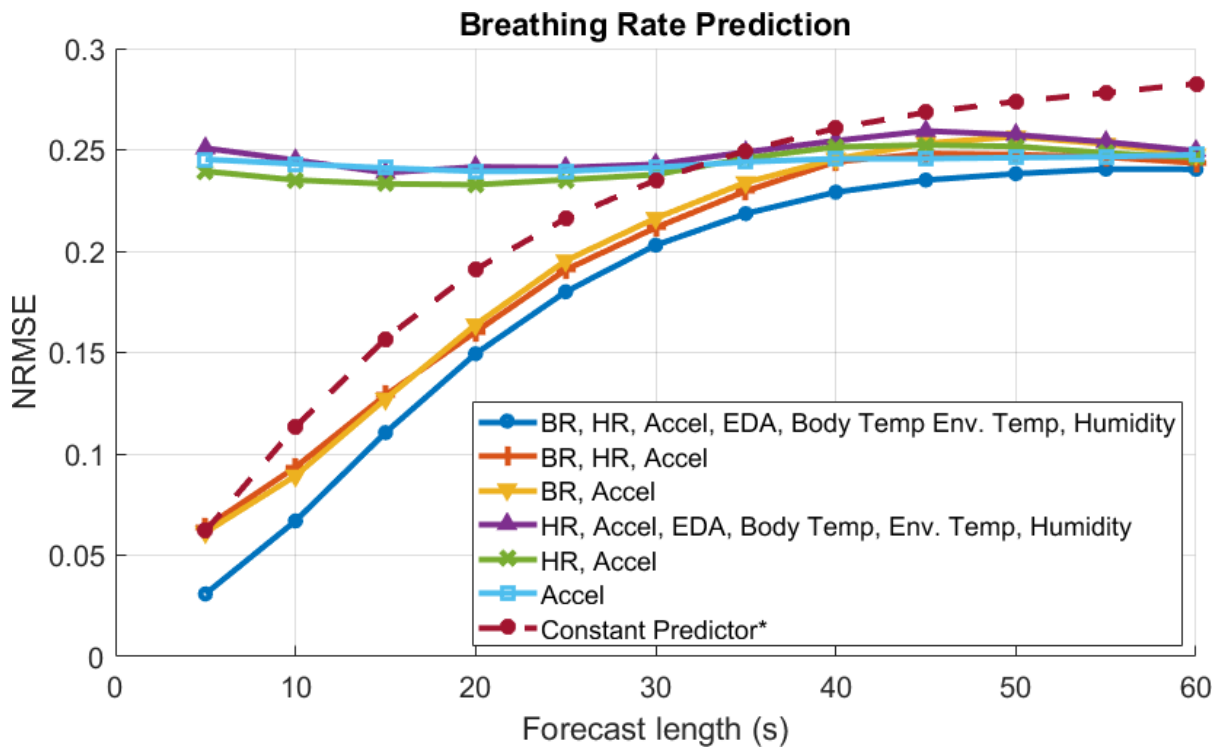


Figure 3.3 Breathing rate prediction error for various forecast lengths. The constant predictor is compared to trained predictors with six different sets of features.

Residuals of Breathing Rate Predictions (Forecast time = 20s)

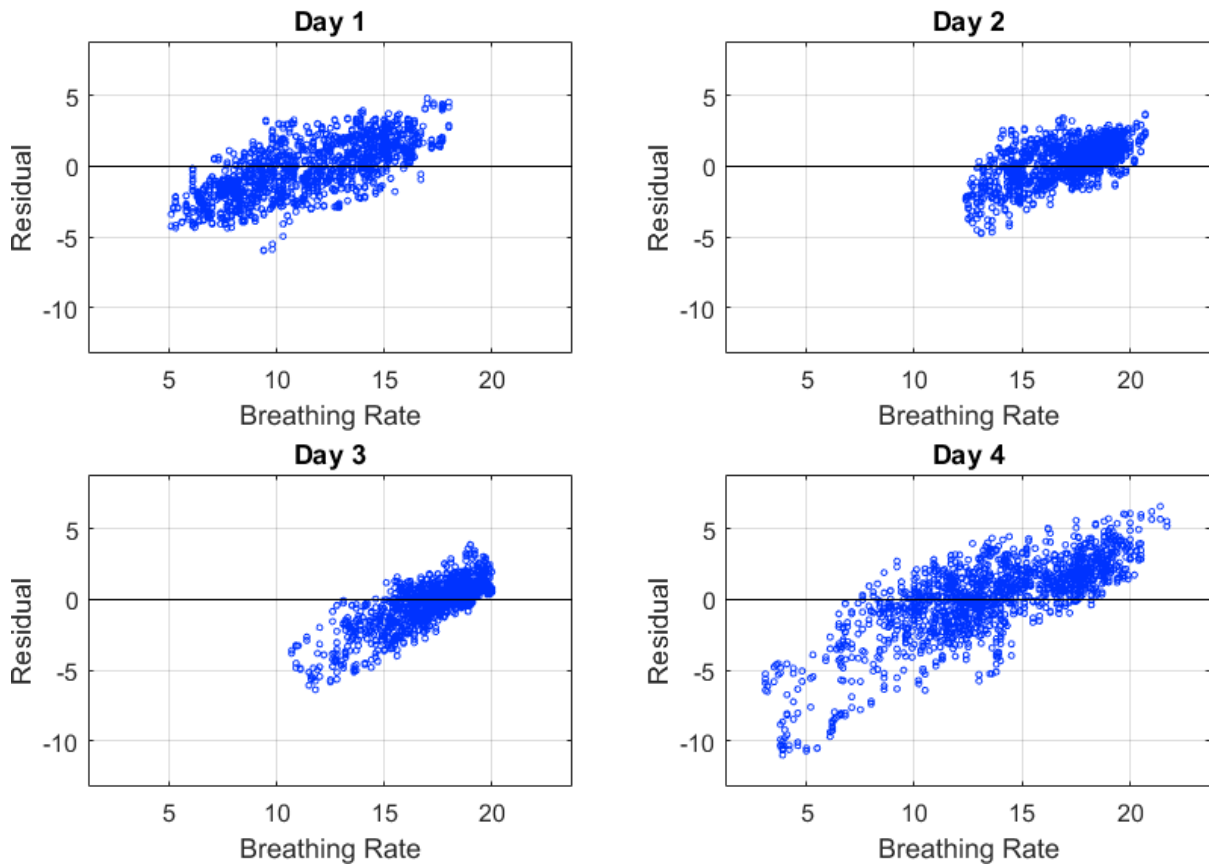


Figure 3.4 Residual values of breathing rate predictions for a 20 second forecast on testing data as a function of the actual breathing rate. Results are shown for four days.

CHAPTER

4

SENSOR SELECTION FOR
POWER-EFFICIENT PREDICTION OF
PHYSIOLOGICAL RESPONSES IN A
WEARABLE HEALTH-MONITORING
PLATFORM

4.1 Introduction

All wearable sensors have the requirement of remaining tether-free to not be obtrusive to the wearer. In order to accomplish this, these sensors require a local energy source capable of providing sufficient power to the sensors, processing units, and wireless transmitters. This is almost exclusively accomplished using a rechargeable lithium-ion (Li-ion) or lithium-polymer (Lipo) type battery as these contain sufficiently large energy and power densities for wearable sensors. Remote health monitoring systems, however, require long periods of continuous sensing and transmitting in order to track a patient's activities. Most sensor nodes, however, consume too much energy to remain powered for more than a few hours. Much effort has been made to minimize energy consumption of every hardware component in a sensor node. To accompany these efforts, it is similarly important to make judicious use of the available energy.

As medical wearable devices are utilized for regular offsite supervision of patients, power efficiency is one of the most significant aspects of their performance [Blo07]. Continuous power usage by any of these modules will quickly deplete its battery. In general, efforts to reduce power cost for each module may be as simple as turning on the display when a user is looking at it, reducing the sampling rate of sensors, or disabling wireless communication. Individual components in consumer electronics can for the most part be independently analyzed to determine their respective power consumption. Displays are by far the largest power consuming modules on most devices that contain a lit display. This includes smartphones, smartwatches, GPS navigators, tablets, laptops. Other components include processors, speakers, and wireless transmitters/receivers. Although wireless wearable sensors often do not contain a display and are generally power efficient, they consume a lot of power due to high sampling rates and large amounts of data transmission. Moreover, these devices are typically continuously active resulting in higher power usage. Recently, energy harvesting methods have been given much attention in attempts to solve the energy problem. Nonetheless, today's state-of-the-art body energy harvesters are still not able to reach the 1mW threshold [Sua16] needed to continuously drive even the most energy efficient body-worn sensors that are of interest

in remote patient monitoring. It is clear that novel and creative methods of optimizing power usage need to be developed.

In this Chapter, we focus on power optimization in predicting breathing rate (for a 10s forecasting window) using multi-modal data streams including body and environmental temperature, heart rate, breathing rate, humidity, accelerometer from the wrist and ankle, and electrodermal activity (EDA). We focus on a dynamic feature selection method by deactivating different set of features depending on the activity state of the user. We identify the activity state using the features extracted from the subject's ankle-mounted accelerometer. Various features are extracted from this one device, dimensionality reduction is applied, and clustering is performed on the data points in a reduced dimension feature space. A group lasso regression model is then used within each cluster to select the reduced set of relevant features for prediction in a power-efficient way while features from each sensor are included in or excluded from the model together. Our framework has the potential to be utilized for prediction of a wide variety of physiological responses in applications such as cardiopathy research and diagnosis, heart attack danger notice, sports capability measure, and mental activity assessment.

The remainder of the Chapter is organized as follows: In Section 4.2, we explain our dataset and the devices used to collect the measurements. Section 4.3 describes the problem and presents our proposed framework to address it. Experimental results are included in Sections 4.4 and 4.5, and Section 4.6 concludes the chapter.

4.2 Dataset

We used the wearable health monitoring platform described in the previous chapter with a different data acquisition protocol described below to obtain the dataset used in this chapter.

4.2.1 Data Acquisition Protocol

We recorded physiological activity and environmental data on a healthy, 25-30 year old male subject over a period of four days. The subject wore the aforementioned devices in the following manner:

The E4 was worn on the left wrist, following the manufacturer's instructions. The SensorTag was strapped to the right ankle to capture lower body movement. The BioHarness was worn using the provided sports shirt that contains the built-in chest strap. The AirBeam was carried on the subject's back except during workout routines in which the device would be obstructive. During exercise, the AirBeam was kept close to the subject to measure his environment's temperature and humidity.

Since each device keeps track of time locally, the subject was asked to jump at the beginning and at the end of each recording session. This event is captured by all recording devices and is used to synchronize all data channels. The subject started the data collection by walking outdoors for 10min. Then he alternated between walking one lap and running one lap on a 250m long indoor track for 20 minutes. He then rested for 10min and performed two 15min long sessions of various exercises with 5min of resting in between. The exercises include chest, leg, and abdominal workout routines established by a third party. Next, the subject walked indoors for 10 mins and then left the building and walked outdoors for another 10 mins. The main intention of the protocol was to obtain highly varying breathing rate and heart rate data during various activities in an uncontrolled environment. The same data collection routine was repeated in the second day.

4.3 Proposed Framework

We use the following set of features acquired from the subject in our study: the breathing rate in breaths per minute, body temperature in Fahrenheit, heart rate in beats per minute, environmental temperature in Fahrenheit, relative humidity in percentage, various accelerometer features from the wrist and ankle and EDA features. For each axis of the accelerometer, we compute mean, variance, covariances, skewness, kurtosis, peak spectral magnitude, and the corresponding peak frequency. For EDA measurements, we extract the linear trend, variance, skewness, Kurtosis, and range between maximum and minimum value.

The first objective is to predict physiological responses, namely breathing rate using the features described above. We focus on the case of forecasting breathing rate 10s, 20s, and 30s into the future by including the aforementioned features as well as the current breathing rate value. The



Figure 4.1 Block diagram of the signal processing flow in the proposed framework

second and main objective is to dynamically identify a subset of features that can best predict the physiological responses and use the reduced feature space in order to conserve the required processing power. We aim to perform a dynamic feature selection based on the user's activity. The acceleration data from one location on body can be used to identify the activity. Within each activity state, prediction of physiological responses at each instance of time in the future can be performed using a regression model. By using a lasso regression technique for each activity state, we find a reduced set of features most predictive of the physiological responses. The power consumption will be consequently lowered by having to process only those relevant features.

Figure 4.1 shows the block diagram of the signal processing flow in the proposed framework. We first preprocess the measurements obtained from the sensors on the wearable platform. We then extract features and pass them to the activity identification module. Based on the identified activity state of the user, we select a subset of features to use for prediction of breathing rate. Note that several features can be obtained from one sensor such as 21 features derived from one accelerometer. Therefore, in the feature selection process, the set of features obtained from each sensor need to be included in or excluded from the model together. We accordingly use group lasso, where each group represents features obtained from one sensor.

4.3.1 Activity State Identification

We perform activity state identification on one sensor node specifically the SensorTag on the ankle since it is the device with the lowest power consumption. However, this choice is arbitrary and other sensors could have been used. We use mean, variance, covariances skewness, kurtosis, peak spectral magnitude, and the corresponding peak frequency for each axis of the accelerometer data

collected on ankle to identify the activity state of the user. We first reduce the dimension of the features space from 21 to 3 using Locality Preserving Projections (LPP) method [HN04]. In the obtained three-dimensional space, we perform k-means clustering to partition the data into four activity states present in our collected dataset, namely walking, running, resting and exercising. Note that the optimal reduced set of features relevant for prediction of physiological responses are not necessarily the same for each activity state. Therefore, prediction is performed differently based on the activity state to dynamically optimize the set of features.

4.3.2 Prediction Model: Group Lasso for Linear Regression

Although the idea of L_1 penalty is not new, its application to regularization formed a particularly powerful tool for sparse regression after the least absolute shrinkage and selection operator (or lasso) technique was proposed in [Tib96]. In this section, we briefly describe our utilized prediction model based on L_1 regularization for linear regression which is the most commonly used form of Lasso. Consider a dataset with N observations and p predictors, where the response vector is $\mathbf{Y} = (y_1, \dots, y_N)^\top$ and the feature matrix is $\mathbf{X} = [\mathbf{x}_1 | \dots | \mathbf{x}_p]$, with predictors $\mathbf{x}_j = (x_{1j}, \dots, x_{Nj})^\top$. The common linear regression model is represented by

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}. \quad (4.1)$$

The errors ϵ_i are assumed to be independent and identically distributed Gaussian random variables with mean and variance of $(0, \sigma^2)$. The Lasso framework includes the minimization of the cost function below subject to an L_1 constraint:

$$\hat{\boldsymbol{\beta}} = \min_{\boldsymbol{\beta}} \|\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \quad (4.2)$$

subject to

$$\|\boldsymbol{\beta}\|_2^2 \leq s, \quad (4.3)$$

where $\beta \in R^p$, $s \geq 0$ and $\|\cdot\|_q$ denotes the q -norm. The Lagrangian representation of the lasso is given as follows:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_1. \quad (4.4)$$

where $\|\cdot\|_q$ denotes the q -norm. Hence, the lasso technique replaces the L_2 penalty term in the well known ridge regression [HK70] with L_1 penalty. Although the lasso optimization formulation is convex, it is non-differentiable at zero as a result of L_1 penalty term. Lasso is specifically advantageous in reducing the number of features. Using the non-differentiable characteristic of L_1 term at zero, lasso chooses a subset of features by forcing some of the regression coefficients to be equal to zero. λ determines the degree of freedom in the model, i.e. using a small λ close to zero provides a solution similar to the least squares estimation and by increasing λ , more and more coefficients in β become equal to zero. Accordingly, the value λ can be utilized to manage the number of features included in the regression process. When the number of features is larger than the number of observations, the problem is not necessarily convex and therefore lasso regression will have several solutions. These solutions however have same sparsity structure.

Another difference of lasso and ridge regression is that in general, the lasso problem does not have an explicit solution except for an orthonormal input matrix \mathbf{X} , where the explicit solution can be obtained as

$$\hat{\beta}_j = \operatorname{sign}(\hat{\beta}_j^{ls}) (\hat{\beta}_j^{ls} - \lambda)_+, j \in \{1, \dots, p\}, \quad (4.5)$$

where $\hat{\beta}_j^{ls}$ denotes the least square estimation of the j -th feature and $(\cdot)_+$ denotes the positive part. The above method is referred to as soft thresholding and is similar to the ridge regression closed form for orthonormal feature matrix shown below:

$$\hat{\beta}_j = \frac{\hat{\beta}_j^{ls}}{1 + \lambda}, j \in \{1, \dots, p\} \quad (4.6)$$

Figure 4.2 illustrates a comparison of lasso and ridge regressions for a space with one feature. The X and Y axis represent the least square coefficients $\hat{\beta}_1^{ls}$ and the regularized coefficients $\hat{\beta}_1$, respectively. The red lines depict the unrestricted least square estimations.




Figure 4.2 Explicit solution of ridge (left) and lasso (right) for the orthonormal feature matrix. The red dotted line shows the unrestricted least square estimation [4].

Although it may seem that reducing the number of utilized features may reduce the degree of freedom of the model, lasso's ability to choose particular features and eliminate the rest provides a balance. Figure 4.3 shows the optimization problem for the space with two features using L_1 and L_2 penalties. The solution of the regularized model is the intersection point of the elliptical contours representing the least square cost function and the constraint which is a circle for ridge and a rhombus for lasso both centered at $(0, 0)$. The feature selection only happens in lasso and not in ridge. The reason can be explained as follows. When the regularization term is non-differentiable at $\beta_j = 0$, it shapes a corner at j -axis. This forces the cost and penalty functions to intersect exactly on the axis. Therefore, the corresponding β_j would be equal to zero. On the other hand, the ridge regression does not have any corners on the axes. As a result, it is probabilistically impossible for β_j to be precisely equal to zero. Feature selection only happens for L_q where $q \leq 1$ due to the regularization term being non-differentiable at zero, while the corresponding optimization problem is convex for $q \geq 1$. Since lasso is the only intersection of the two restrictions it can perform feature selection without sacrificing convexity.

Group lasso was proposed in [YL06] to include in or exclude predetermined groups of features.




Figure 4.3 Elliptical contours of the least squares cost function and contours of the regularization region for ridge(left) and lasso(right).

Formalization of the group lasso is a generalized form of standard lasso:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \left\| \mathbf{Y} - \sum_{j=1}^J \mathbf{X}_j \beta_j \right\|_2^2 + \lambda \sum_{j=1}^J \|\beta_j\|_{K_j}, \quad (4.7)$$

where J is the number of groups, K_1, \dots, K_J are positive definite matrices and for a vector η , $\|\eta\|_K = (\eta' K \eta)^{1/2}$. In group lasso representation, the matrices \mathbf{X}_j and β_j represent the features and coefficients for each group. If each group contains only one feature, i.e. $K_j = I$, then equation (4.7) reduces to the standard lasso. On the other hand, if all the features are in one group, it reduces to ridge regression. λ determines the degree of freedom in the model, i.e. using a small λ close to zero provides a solution similar to the least squares estimation and by increasing λ , more and more groups of coefficients become equal to zero. Accordingly, the value λ can be utilized to manage the number of groups included in the regression process.

Note that the data needs to be normalized before applying group lasso in order for the penalty to be independent of the scale of the features and targets values. If the values are not centered at zero, the vector β needs to include an intercept term as well.

4.4 Experimental Results

For the initial analysis presented in this section, we rely on data obtained from the first two recording days. Here, the first day was used for training the prediction model and the second day was used for evaluating the model's accuracy. In order to determine valid values for λ in Equation 4.7, we employ various cross validation approaches that are described in Subsection 4.4.2. Section includes further analysis and results that use all four days of data.

4.4.1 Clustering for Activity Identification

We first apply LPP to the set of 21 features extracted from 3-axis accelerometer data collected on the subject's ankle, including mean, variance, covariance, skewness, kurtosis, peak spectral magnitude, and the corresponding peak frequency computed for each axis. The dimension of the feature space is reduced from twenty-one to three. We then apply k-means clustering to partition the points in the three dimensional feature space into clusters, where each data point is assigned to the cluster with the nearest mean. We select the number of clusters to be four since in our data collection protocol the subject was doing four different activities, namely walking, running, exercising and resting. In order to enforce independence between our training and test sets, we use the data from the first day for training and the data collected on the second day for testing.

Figure 4.4 illustrates the clustering results in three dimensional space with the number of data points in each cluster after performing clustering on the acceleration features in reduced dimension. Different colors are used to show different clusters. This Figure depicts a clear separation between the four clusters corresponding to four different activity states. Figure 4.5 shows the clustering over time in the test set. This figure illustrates that the points in each cluster are mainly close to each other in time. This is expected since clusters represent activity states and each activity happens in specific chunks of time. For example the activity corresponding to the first cluster has mainly occurred between minutes 10 and 30 as well as between minutes 64 and 88.

Clustering Visualization




Figure 4.4 Visualization of different clusters of the acceleration features in reduced dimension in space with the number of data points in each cluster.

4.4.2 Feature Selection Versus Prediction Accuracy

The complete list of features used in this study is described in Section 4.3. On the training set, we perform three types of K -fold (with $K = 5$) cross validations. The first method is called global cross validation (GCV) where the data points are split into folds before clustering. The problem with this method is that the data in the training folds may not be representative of the data in the testing fold since each activity state may not be present in all folds. To address this issue, we also perform a local cross validation (LCV) after clustering, where the parameter λ in Equation 4.7 is selected by minimizing the root mean square error (RMSE), $\text{RMSE} = \sqrt{\frac{1}{n} \sum_{t=1}^n (\hat{y}_t - y_t)^2}$, for each cluster. We

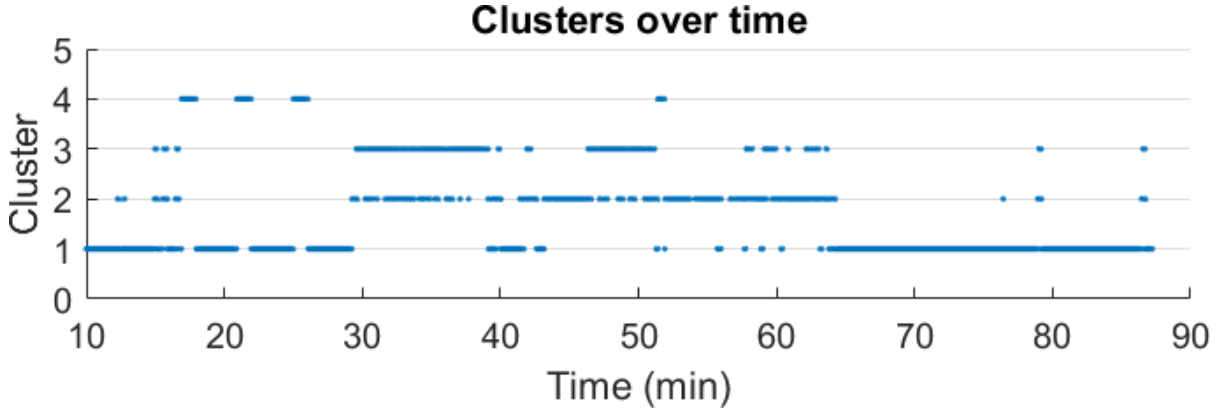


Figure 4.5 Results of clustering over time in the test set.

Table 4.1 RMSE and Number of Sensors at λ_{min} for 10s forecast length.

	$RMSE_{min}$	N of sensors
Cluster 1	1.5303	5
Cluster 2	2.2590	3
Cluster 3	1.2232	3
Cluster 4	0.9059	3
LCV	1.6585	3-5
SLLCV	1.6611	3-5
GCV	1.6659	3-5
No clustering	1.6598	6

also perform a single lambda local cross validation (SLLCV) method that computes the weighed average of MSE values obtained in LCV across clusters to compute the total RMSE. In the cross validation process, we also select training and testing folds to be two separate segments of data in time instead of random selection since our target values are not independent at different times.

Figure 4.6 shows the RMSE of prediction accuracy using GCV and SLLCV cross validation methods as a function of λ in Equation (4.7) for a forecast length of 10s. This figure illustrates that although GCV and SLLCV methods result in similar accuracies for many values of λ , GCV results in higher RMSE for a significant portion of λ values. The reason is that in GCV, each activity state may not be present in all folds and as a result the data in the training folds may not be representative of the data




Figure 4.6 GCV and SLLCV versus λ in the training set obtained via cross validation with forecast length of 10s.

in the testing fold. In order to address this issue, we have performed experiments with LCV cross validation method as well. Figure 4.7 shows in the top row the prediction accuracy obtained using LCV for the training data (using cross-validation) and the test data, as well as the corresponding number of relevant features. It is clear in the Figure that the number of relevant features decreases as λ increases. The smallest value of λ results in all the features being used in the model. When using all the features from all sensors, the prediction RMSE is relatively higher. As λ increases and the model uses features from less sensors, the error decreases as a result of irrelevant features being removed from the regression model. However, the Figure shows that error begins to increase after reaching its minimum since λ is passing its optimal value denoted by λ_{min} and some relevant features get removed from the prediction model. We therefore consider the optimal λ as the point where all the irrelevant features are removed and perform the rest of feature selection for power optimization on the λ values greater than λ_{min} . Accordingly, since the power consumed by the platform increases by using more sensors, there is a trade-off between performance and consumed power for $\lambda > \lambda_{min}$.

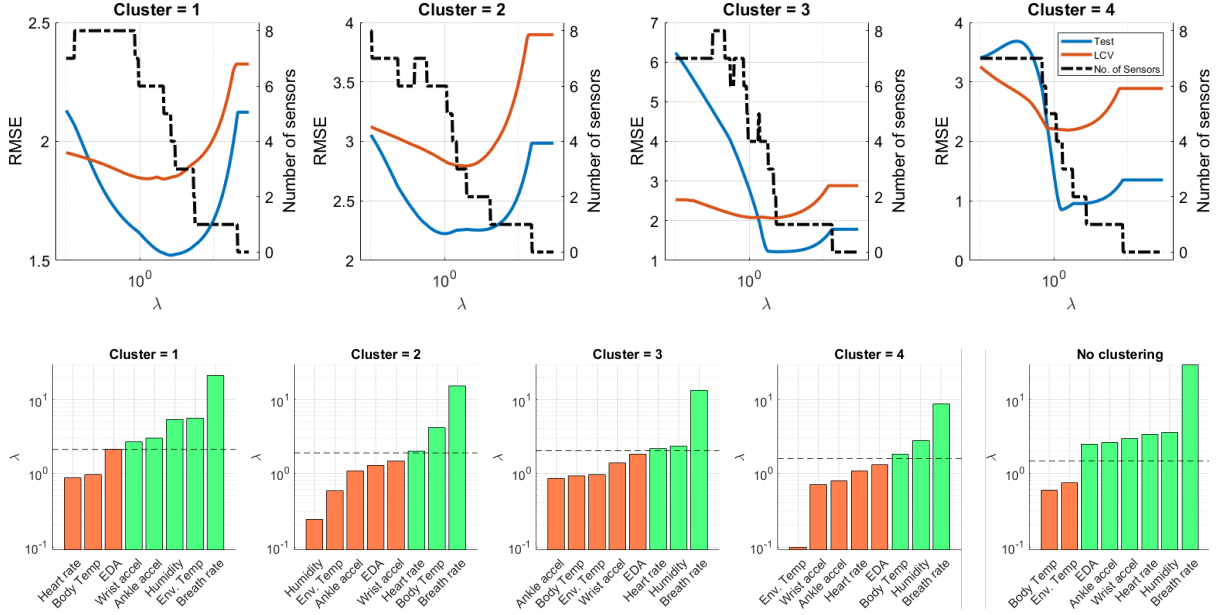


Figure 4.7 Top: RMSE (root mean square error) and number of sensors versus λ for all four clusters in the training set obtained via cross validation. Bottom: λ at which the sensors data are excluded from the model in each cluster as well as no clustering. λ_{min} is showed with a dashed line and the green and orange bars show sensors included in and excluded form the prediction at λ_{min} , respectively.

Table 4.1 shows the RMSE of the test set the using λ_{min} obtained from crossvalidation and the number of sensors used. It shows the total RMSE obtained using LCV (total and per cluster), GCV and SLLCV cross validation methods. Note that LCV, GCV and SLLCV are the same when there is no clustering. According to Table 4.1, the LCV has slightly lower RMSE than GCV and SLLCV due to the fact that it ensures that training folds have data representative of the testing fold. Also, SLLCV performs slightly worse than LCV since it tries to find a single λ for all four clusters while LCV optimizes λ per cluster. Table 4.1 also shows that dynamic feature selection performed using the clustering method keeps the error low while significantly reduces the number of sensors used to achieve it.

Second row of Figure 4.7 shows λ at which the sensors features are excluded from the model as well as the value of λ_{min} with a dashed line. The green and orange bars show sensors included in and excluded form the prediction at λ_{min} , respectively. These plots illustrate that only 5, 3, 3 and 3 sensors are used in clusters 1, 2, 3, and 4, respectively to achieve the minimum error. On the other




Figure 4.8 Predicted and actual breathing rate vs time using the respective cluster's λ_{min} obtained from LCV.

hand, the data from 6 sensors is required without clustering to achieve the minimum error as shown in Figure 4.7. Note that the overall minimum error using clustering is as low as without clustering as shown in Table 4.1 while the proposed clustering based dynamic method uses significantly fewer sensors and requires much lower power.

Figure 4.8 shows the predicted and actual breathing rate vs time in the test set using the respective cluster's λ_{min} obtained from LCV. As it is clear in the Figure, predicted and actual breathing rate values are very close showing that using λ_{min} in each cluster results in a promising over-all prediction performance in time.

We repeated the experiment described above for a forecast length of 20 seconds. Figures 4.9-4.11 show the results.

4.4.3 Dynamic Trade-Off Between Feature Set and Power

As an illustration, suppose that we are in cluster 1. We have already computed how RMSE and number of sensors vary with λ using the training set as shown in top row of Figure 4.7. We can derive λ_{min} that results in the smallest RMSE in the training set. Then, we determine how much performance deterioration we are willing to accept. If we are willing to accept 5% more RMSE, we find the corresponding $\lambda_{5\%}$ and coefficients. For cluster 1, this corresponds to an increase of RMSE

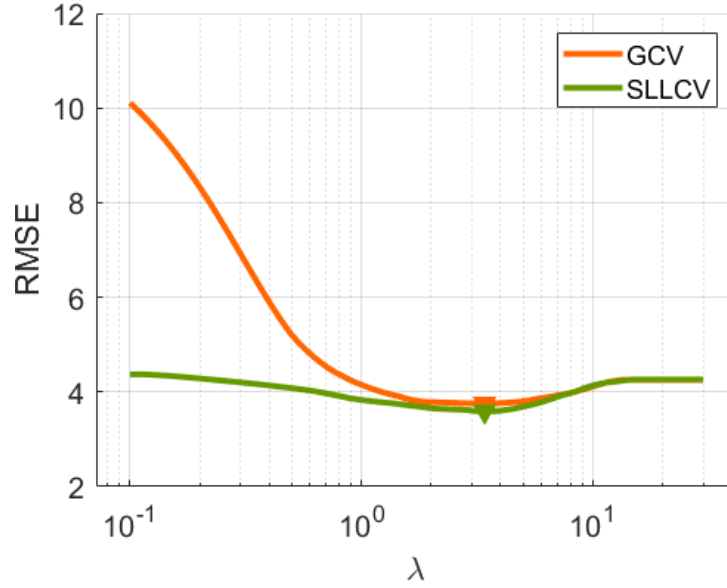


Figure 4.9 GCV and SLLCV versus λ in the training set obtained via cross validation with forecast length of 20s.

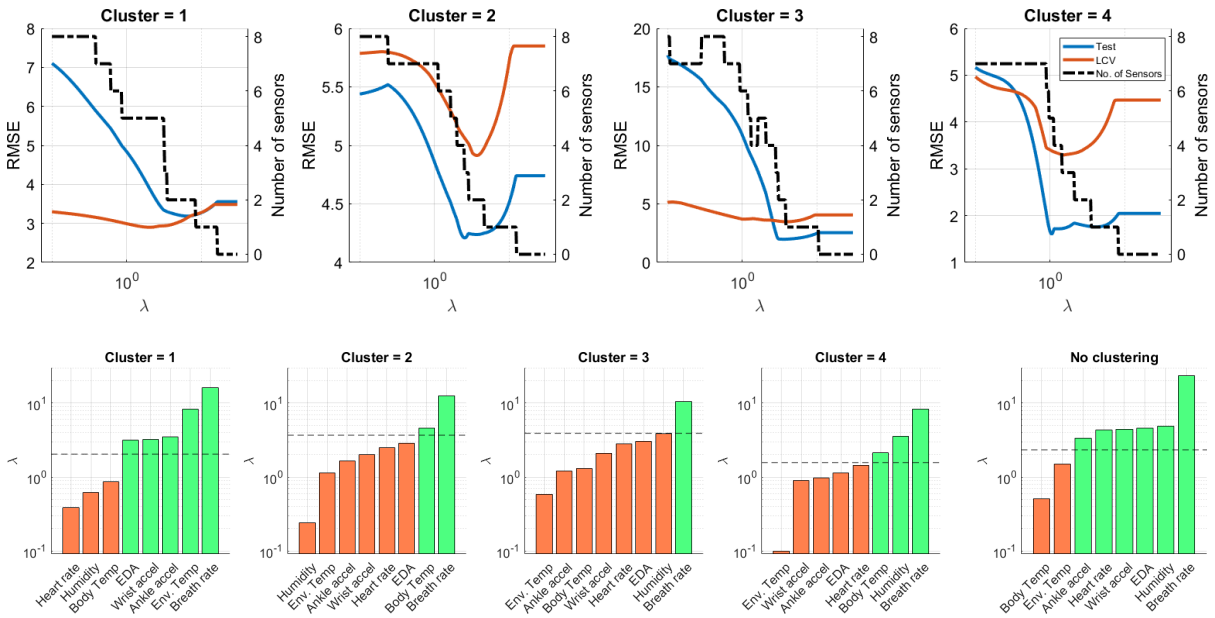


Figure 4.10 Top: RMSE (root mean square error) and number of sensors versus λ for all four clusters in the training set obtained via cross validation. Bottom: λ at which the sensors data are excluded from the model in each cluster as well as no clustering. λ_{min} is showed with a dashed line and the green and orange bars show sensors included in and excluded form the prediction at λ_{min} , respectively with forecast length of 20s.

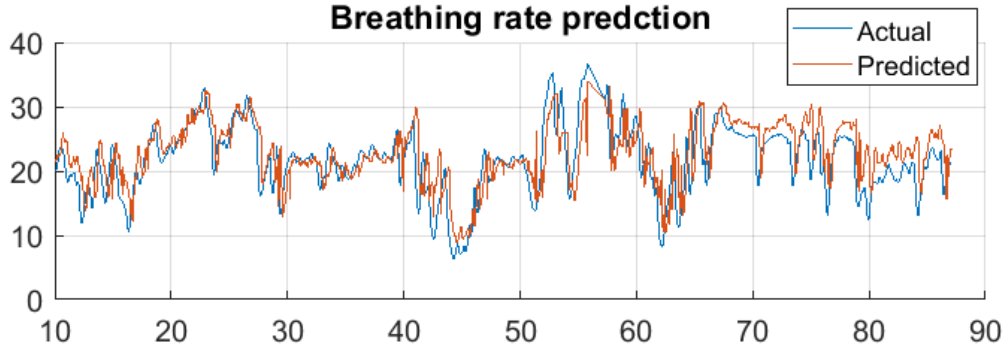


Figure 4.11 Predicted and actual breathing rate vs time using the respective cluster's λ_{min} obtained from LCV with forecast length of 10s.

in test set from 1.53 to 1.58 bpm (breaths per minute). It also decreases the power consumption from five sensors to one sensor. If 5% more RMSE is not tolerable in the application, we can aim for 3% more error and achieve more accurate prediction with an RMSE of 1.54 bpm. This lower error comes with the cost of using three sensor, which is still lower than the five sensors used to achieve minimum error.

4.5 Comparison of the Methods and Power Consumptions

In order to better investigate the trade-off between the number of features used and the consumed power, we perform the rest of our experimental analysis on the full dataset collected over 4 days in a similar manner as before. This provides a richer dataset for better evaluation of our framework. We use the data collected on one day for testing and the data from the remaining 3 days for training. We repeat this procedure for each of the 4 days and compute the error averaged over 4 days. In the training set that contains data from 3 days, we perform a 3-fold cross validation where the data from each day forms a separate fold. Based on the results in Section 4.4, we employ LCV (local cross validation) as the cross validation approach here. That is, we first perform clustering and then implement cross validation on each cluster in order to obtain an optimal value of λ per cluster.

Figure 4.12 shows the prediction accuracy obtained using LCV for the 3 days of training data and one day of test data, and the corresponding number of relevant features. Similar to the 2 days

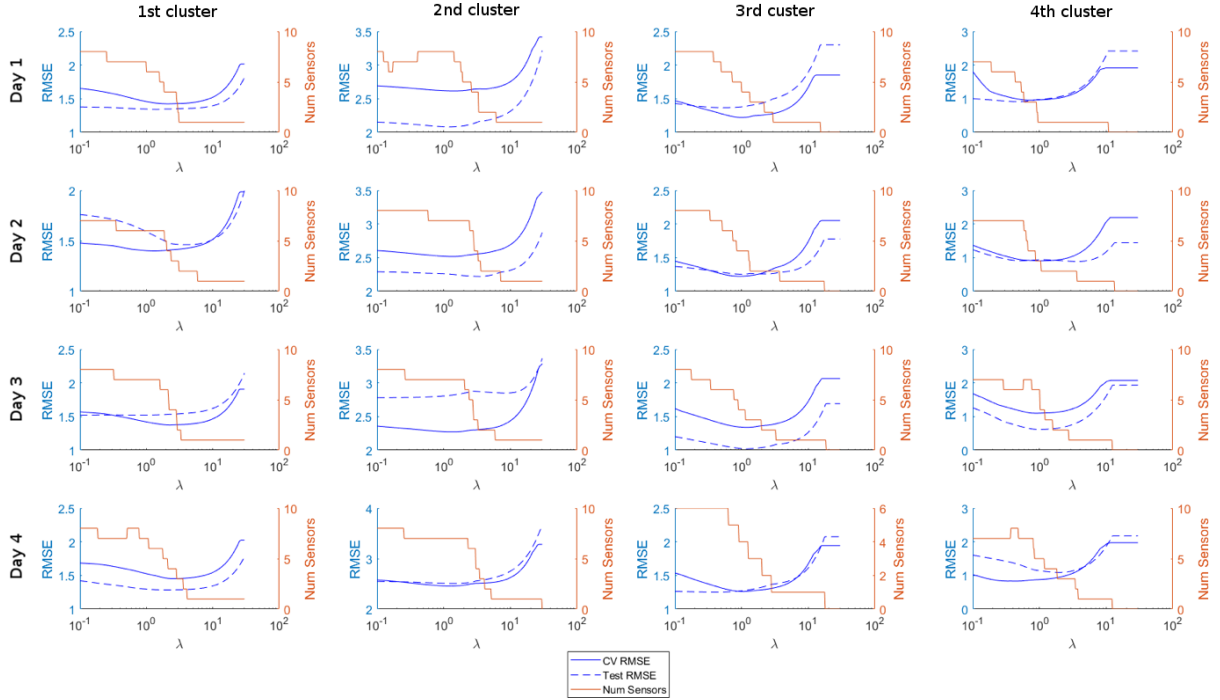


Figure 4.12 RMSE (root mean square error) and number of sensors versus λ for all four clusters in one day obtained via cross validation on the remaining three days with forecast window of length 10s. Each row corresponds to one day, and each column corresponds to one cluster.

dataset we used before, the number of relevant features decreases as λ increases. The smallest value of λ results in all the features being used in the model. When using all the features from all sensors, the prediction RMSE is relatively higher. As λ increases and the model uses features from less sensors, the error decreases as a result of irrelevant features being removed from the regression model. This is similar to the behavior we observed in the 2 days dataset. One difference between using data from 4 days (Figure 4.12) and using data from 2 days (Figure 4.7) is that using data from 4 days results in lower RMSE for the first value of λ in all clusters due to having a richer dataset. Also similar to our observation before, the error begins to increase after reaching its minimum since λ is passing its optimal value denoted by λ_{min} and some relevant features are removed from the regression model. As it is clear in Figure 4.12, for those values by removing more features the error increases. Hence, there is a trade-off between performance and consumed power for $\lambda > \lambda_{min}$ as the power consumed by the wearable platform decreases by using less sensors.

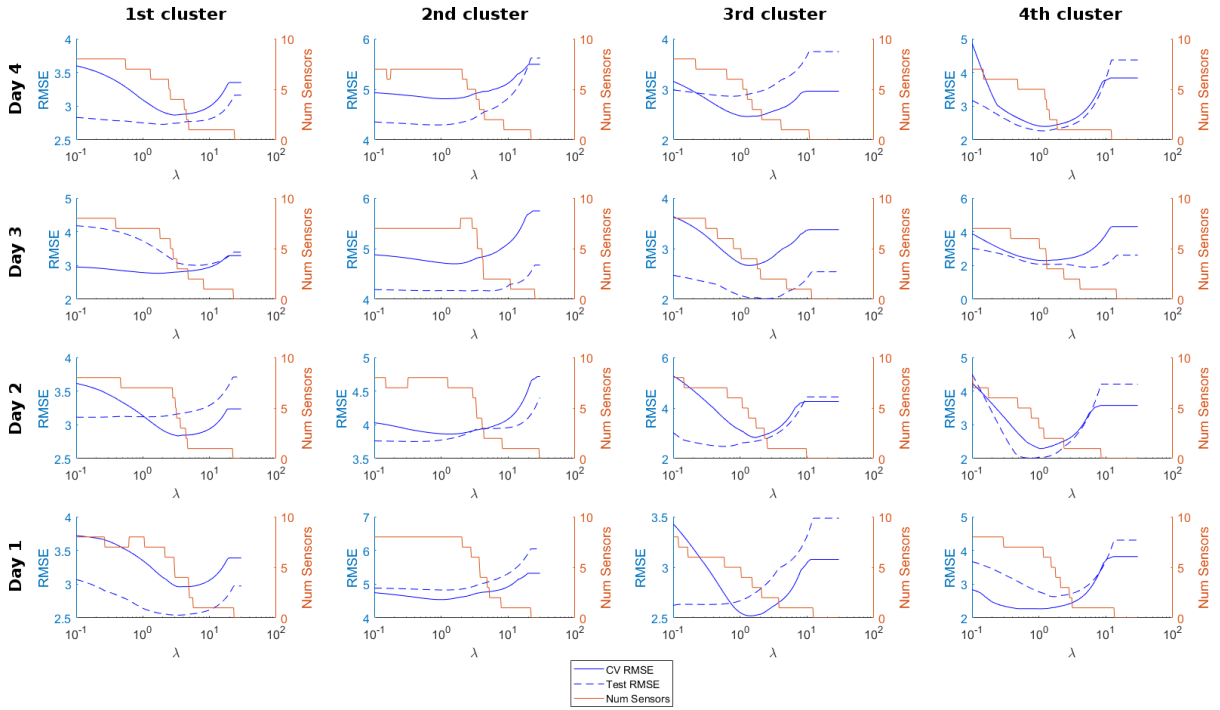


Figure 4.13 RMSE (root mean square error) and number of sensors versus λ for all four clusters in one day obtained via cross validation on the remaining three days with forecast window of length 20s. Each row corresponds to one day, and each column corresponds to one cluster.

Note that the trajectory of cross validation RMSE on three days training set is closer to the test set in the four days dataset compared to training only on one day’s data since in 3 days training set we have data from all clusters in each of the folds of cross validation while we did not necessarily have that when training in only one day’s data. This is the main advantage of using 4 days data that provides a more accurate training process and better optimized parameters.

We repeated the experiment above with forecast window of length 20s and shown the results in Figure 4.13. As expected, the error in each cluster increases as forecast length increases. However, the points explained above regarding the tradeoff between accuracy and power consumption are all still valid when using a longer forecast window.

Table 4.2 shows Group Lasso (GL) RMSE obtained using the optimal λ for each day and each cluster with the forecast window length of 10s. The optimal λ obtained from LCV on 3 days’ data is used on the test day data. This table also includes the weighted average of RMSE over clusters,

where the weighting is computed based on the size (the number of data points) in each cluster. For comparative purposes, it also shows the GL RMSE when no clustering is performed and therefore there is no dynamic sensor selection. The power consumption for the selected sensors is also shown, along with the fractional improvement of power consumption (in parentheses) compared to no clustering. This value can be interpreted as the improvement of power consumption due to activity-based clustering. Table 4.2 also shows the power consumption and number of sensors used to achieve the presented RMSE values for each cluster, the weighted average number of sensors over all clusters, and the number of sensors required when not clustering. The power values in parentheses show the fractional improvement in power consumption achieved due to clustering. The size of each cluster, which is the number of data points in the corresponding cluster is also shown in the table for each cluster, the average, and for the whole day's dataset.

As an additional technique, we combined the previous two frameworks, namely nonlinear gaussian SVM regression and group lasso. We used group lasso the similar way as before for selecting an optimal set of sensors to use and then instead of using the coefficients of group lasso for prediction, we used gaussian SVM regression. The results of this method that are included in Table 4.2 show that in 3 out of 4 days data SVM results in higher prediction error and in one day SVM performs slightly better than group lasso. The respective power consumption values for each column were obtained by measuring the average power usage of a Shimmer device, a small, wearable, wireless multi-modal sensor. We assume that all sensing modalities utilize the same power as the Shimmer device at their respective sampling and transmission rates. While this approximation may be unsuitable for vigorous comparison purposes, we include these values for illustrative purposes as a proof-of-concept for the presented framework. On average, using group lasso model performs better for prediction and results in smaller RMSE as show in the last two lines of Table 4.2. Note that we cannot average RMSE of group lasso or SVM for each cluster over 4 days since a specific cluster in different days may not correspond to the same activity state. This is the reason for empty cells in the last two rows of the table. Table 4.2 also includes the actual value of the power in mW consumed in each scenario. As it is clear in the tables, our proposed dynamic sensor selection method significantly

Table 4.2 RMSE of GL and SVM, number of sensors, and power consumption for λ_{min} with and without clustering for forecast length of 10s.

Test Day		1st cluster	2nd cluster	3rd cluster	4th cluster	Avg. over clusters	No clustering
1	GL RMSE	1.347	2.083	1.379	0.941	1.533	1.542
	SVM RMSE	1.403	2.259	1.448	1.148	1.632	2.252
	Num sensors	4	8	4	3	4.84	6
	Power (mW)	426.0	842.8	426.0	321.2	514.7 (21.6%)	656.4
	Size	2822	1145	817	285	1267	5069
2	GL RMSE	1.549	2.251	1.253	0.938	1.669	1.711
	SVM RMSE	1.720	2.343	1.633	0.937	1.837	1.704
	Num sensors	6	7	4	3	5.74	7
	Power (mW)	609.6	749.6	434.8	286.0	597.1 (9.0%)	656.4
	Size	2689	1072	805	223	1197	4789
3	GL RMSE	1.533	2.822	1.019	0.617	1.759	1.773
	SVM RMSE	1.472	2.749	1.151	1.047	1.729	1.685
	Num sensors	4	7	3	6	4.5	7
	Power (mW)	379.2	749.6	332.8	647.6	433.0 (42.0%)	749.6
	Size	4071	662	121	100	1238	3925
4	GL RMSE	1.286	2.508	1.276	1.353	1.644	1.633
	SVM RMSE	1.294	2.675	1.504	1.515	1.734	1.730
	Num sensors	4	7	4	8	4.85	5
	Power (mW)	426.0	749.6	426.0	842.8	517.0 (8.2%)	563.2
	Size	3038	1117	584	215	1238	4954
Avg.	GL RMSE	-	-	-	-	1.65	1.665
	SVM RMSE	-	-	-	-	1.73	1.843
	Num sensors	-	-	-	-	4.98	6.25
	Power (mW)	-	-	-	-	520.0 (23.5%)	679.7

reduces the consumed power while maintaining the prediction error low.

Table 4.3 illustrates the list of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 10s. This Table shows that the regression model always used previous values of breathing rate for prediction of its future values. Interestingly, EDA is also always included in the model as a useful measurement for prediction of breathing rate. The Table also shows that humidity is the least important feature and it is often excluded from the regression model.

Table 4.3 List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 10s.

Day	cluster	Body Temp	Heart Rate	Env. Temp	Humidity	Wrist accel	Ankle accel	EDA	Breath rate
1	1st	•	•					•	•
	2nd	•	•	•	•	•	•	•	•
	3rd		•		•			•	•
	4th						•	•	•
	No clustering	•	•			•	•	•	•
2	1st	•		•		•	•	•	•
	2nd	•	•	•		•	•	•	•
	3rd		•			•		•	•
	4th			•				•	•
	No clustering	•	•	•		•	•	•	•
3	1st	•		•				•	•
	2nd	•	•	•		•	•	•	•
	3rd		•					•	•
	4th	•	•	•			•	•	•
	No clustering	•	•		•	•	•	•	•
4	1st		•	•				•	•
	2nd	•	•	•		•	•	•	•
	3rd		•		•			•	•
	4th	•	•	•	•	•	•	•	•
	No clustering		•			•	•	•	•

We have repeated the experiment for the λ corresponding to 3%, and 5% more error in the training set, the results of which on the test sets are shown in Tables 4.4 - 4.7. These tables are useful to determine which sensors need to be turned on based on the desired accuracy that is set by the application as illustrated in Section 4.4.3. The Tables also provide information about how much power in mW will be used when the corresponding sensors are functioning. Tables 4.8 - 4.13 show similar results with a forecast length of 30s. In some instances, the no clustering approach may result in lower power consumption. This only occurs in the extreme case in which only a single sensor (the breathing rate sensor) is needed. Additionally, these negative values of average power reduction only occur for short forecast lengths. Such negative values do not occur for 30s forecast

length predictions in which more sensing modalities are generally needed.

Table 4.4 RMSE of GL and SVM and number of sensors for λ corresponding to 3% more error in the training set with and without clustering for forecast length of 10s.

Day		1st cluster	2nd cluster	3rd cluster	4th cluster	Avg. over clusters	No clustering
1	GL RMSE	1.368	2.245	1.440	1.011	1.605	1.619
	Num sensors	1	1	3	1	1.33	1
	Power (mW)	99.6	99.6	332.8	99.6	137.6 (-38.1%)	99.6
	Size	2810	1150	825	284	1267	5069
2	GL RMSE	1.466	2.245	1.258	0.9170	1.624	1.607
	Num sensors	2	2	2	2	2	3
	Power (mW)	192.8	192.8	239.6	192.8	200.5 (45.5%)	368.0
	Size	2722	1062	784	221	1197	4789
3	GL RMSE	1.579	2.855	1.065	0.6540	1.798	1.809
	Num sensors	1	1	2	2	1.21	1
	Power (mW)	99.6	99.6	239.6	192.8	125.8 (-26.3%)	99.6
	Size	2379	736	583	227	981	3925
4	GL RMSE	1.314	2.619	1.351	1.220	1.698	1.691
	Num sensors	1	1	2	7	1.38	1
	Power (mW)	99.6	239.6	426.0	842.8	141.3 (-41.9%)	99.6
	Size	3038	1116	585	215	1238	4954
Avg.	GL RMSE	-	-	-	-	1.681	1.681
	Num sensors	-	-	-	-	1.48	1.5
	Power (mW)	-	-	-	-	151.3 (9.2%)	166.7

4.6 Conclusion and Extensions

We investigated the tradeoff between the accuracy in prediction of physiological responses and power consumption in a wearable health monitoring platform in order to maintain prediction accuracy, while reducing the power consumption. The features measured by the platform include body and environmental temperature, humidity, accelerometer from the wrist and ankle, EDA, breathing rate and heart rate. We performed a dynamic feature selection based on the activity

Table 4.5 List of sensors selected by group lasso with λ corresponding to 3% more error in the training set for each day used as test set and for each cluster as well as without clustering for forecast length of 10s.

Day	cluster	Body Temp	Heart Rate	Env. Temp	Humidity	Wrist accel	Ankle accel	EDA	Breath rate
1	1st								•
	2nd								•
	3rd		•					•	•
	4th								•
	No clustering								•
2	1st			•					•
	2nd			•					•
	3rd		•						•
	4th			•					•
	No clustering		•				•		•
3	1st								•
	2nd								•
	3rd		•						•
	4th			•					•
	No clustering								•
4	1st								•
	2nd								•
	3rd		•						•
	4th	•	•	•	•		•	•	•
	No clustering								•

state of the subject identified by clustering the features from accelerometer on ankle in a reduced dimension space. We then used one group lasso regression model for each activity state to efficiently choose the relevant features. Experimental results on inertial and physiological data collected from actual subjects demonstrated significant power savings. The framework presented in this chapter can be adopted for prediction or classification applications in battery-powered, wearable devices in various fields such as cardiopathy research, heart attack prediction, athletic performance evaluation, and mental health monitoring.

Table 4.6 RMSE of GL and SVM and number of sensors for λ corresponding to 5% more error in the training set with and without clustering for forecast length of 10s.

Day		1st cluster	2nd cluster	3rd cluster	4th cluster	Avg. over clusters	No clustering
1	GL RMSE	1.381	2.307	1.503	1.034	1.641	1.648
	Num sensors	1	1	1	1	1	1
	Power (mW)	99.6	99.6	99.6	99.6	99.6 (0%)	99.6
	Size	2810	1150	825	284	1267	5069
2	GL RMSE	1.482	2.291	1.257	0.9090	1.646	1.623
	Num sensors	1	1	2	2	1.21	1
	Power (mW)	99.6	99.6	239.6	192.8	126.8 (-27.3%)	99.6
	Size	2722	1062	784	221	1197	4789
3	GL RMSE	1.603	2.848	1.090	0.6810	1.812	1.819
	Num sensors	1	1	1	2	1.06	1
	Power (mW)	99.6	99.6	99.6	192.8	105.0 (-5.0%)	99.6
	Size	2379	736	583	227	981	3925
4	GL RMSE	1.333	2.694	1.377	1.151	1.734	1.724
	Num sensors	1	1	1	5	1.17	1
	Power (mW)	99.6	99.6	192.8	842.8	116.9 (-17.4%)	99.6
	Size	3038	1116	585	215	1238	4954
Avg.	GL RMSE	-	-	-	-	1.708	1.703
	Num sensors	-	-	-	-	1.11	1
	Power (mW)	-	-	-	-	112.1 (-12.5%)	99.6

Table 4.7 List of sensors selected by group lasso with λ corresponding to 5% more error in the training set for each day used as test set and for each cluster as well as without clustering for forecast length of 10s.

Day	cluster	Body Temp	Heart Rate	Env. Temp	Humidity	Wrist accel	Ankle accel	EDA	Breath rate
1	1st								•
	2nd								•
	3rd								•
	4th								•
	No clustering								•
2	1st								•
	2nd								•
	3rd		•						•
	4th			•					•
	No clustering								•
3	1st								•
	2nd								•
	3rd								•
	4th			•					•
	No clustering								•
4	1st								•
	2nd								•
	3rd								•
	4th	•	•	•				•	•
	No clustering								•

Table 4.8 RMSE of GL and SVM, number of sensors, and power consumption for λ_{min} with and without clustering for forecast length of 30s.

Test Day		1st cluster	2nd cluster	3rd cluster	4th cluster	Avg. over clusters	No clustering
1	GL RMSE	3.673	5.608	3.709	3.807	4.181	4.319
	Num sensors	4	7	4	4	4.65	7
	Power (mW)	426.0	749.6	434.8	423.2	497.5 (33.6%)	749.6
	Size	2794	1081	830	284	1247	4989
2	GL RMSE	4.496	5.514	2.503	3.456	4.444	4.596
	Num sensors	7	8	4	3	6.54	8
	Power (mW)	749.6	842.8	434.8	286.0	696.4 (17.4%)	842.8
	Size	2694	1028	776	221	1180	4719
3	GL RMSE	3.951	5.580	3.238	3.559	4.146	4.161
	Num sensors	4	8	2	3	4.29	6
	Power (mW)	379.2	842.8	239.6	332.8	430.9 (34.3%)	656.4
	Size	2348	614	571	227	940	3760
4	GL RMSE	3.354	5.993	3.280	4.100	4.140	4.226
	Num sensors	2	8	6	7	4.07	6
	Power (mW)	239.6	842.8	656.4	749.6	449.8 (31.5%)	656.4
	Size	3004	1141	574	215	1233	4934
Avg.	GL RMSE	-	-	-	-	4.228	4.325
	Num sensors	-	-	-	-	4.89	6.75
	Power (mW)	-	-	-	-	518.6 (28.6%)	726.3

Table 4.9 List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 30s.

Day	cluster	Body Temp	Heart Rate	Env. Temp	Humidity	Wrist accel	Ankle accel	EDA	Breath rate
1	1st	•	•					•	•
	2nd	•	•	•		•	•	•	•
	3rd		•			•		•	•
	4th					•	•	•	•
	No clustering	•	•		•	•	•	•	•
2	1st	•	•	•		•	•	•	•
	2nd	•	•	•	•	•	•	•	•
	3rd		•			•		•	•
	4th			•				•	•
	No clustering	•	•	•	•	•	•	•	•
3	1st	•		•				•	•
	2nd	•	•	•	•	•	•	•	•
	3rd		•						•
	4th		•					•	•
	No clustering	•	•			•	•	•	•
4	1st		•						•
	2nd	•	•	•	•	•	•	•	•
	3rd	•	•			•	•	•	•
	4th	•	•	•		•	•	•	•
	No clustering		•	•		•	•	•	•

Table 4.10 RMSE of GL and SVM, number of sensors, and power consumption for λ corresponding to 3% more error in training set with and without clustering for forecast length of 30s.

Test Day		1st cluster	2nd cluster	3rd cluster	4th cluster	Avg. over clusters	No clustering
1	GL RMSE	3.714	6.008	3.845	3.728	4.335	4.523
	Num sensors	1	2	2	3	1.50	1
	Power (mW)	99.6	192.8	239.6	321.2	155.7 (-56.3%)	99.6
	Size	2794	1081	830	284	1247	4989
2	GL RMSE	3.765	5.330	2.469	3.328	3.975	4.058
	Num sensors	2	5	2	3	2.70	6
	Power (mW)	192.8	507.6	239.6	286.0	273.4 (58.3%)	656.4
	Size	2694	1028	776	221	1180	4719
3	GL RMSE	4.083	5.734	3.273	4.004	4.286	4.347
	Num sensors	1	2	2	1	1.31	1
	Power (mW)	99.6	192.8	239.6	99.6	136.1 (-36.6%)	99.6
	Size	2348	614	571	227	940	3760
4	GL RMSE	3.447	5.972	3.550	3.624	4.186	4.325
	Num sensors	1	8	2	6	2.95	2
	Power (mW)	99.6	842.8	239.6	656.4	312.0 (-30.2%)	239.6
	Size	3004	1141	574	215	1233	4934
Avg.	GL RMSE	-	-	-	-	4.195	4.313
	Num sensors	-	-	-	-	2.11	2.50
	Power (mW)	-	-	-	-	219.3 (19.9%)	273.8

Table 4.11 List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 30s.

Day	cluster	Body Temp	Heart Rate	Env. Temp	Humidity	Wrist accel	Ankle accel	EDA	Breath rate
1	1st			•					•
	2nd			•					•
	3rd		•						•
	4th						•	•	•
	No clustering								•
2	1st			•					•
	2nd			•	•		•	•	•
	3rd		•						•
	4th			•				•	•
	No clustering		•	•		•	•	•	•
3	1st			•					•
	2nd			•					•
	3rd		•						•
	4th								•
	No clustering								•
4	1st	•	•	•	•	•	•	•	•
	2nd		•						•
	3rd		•						•
	4th		•	•		•	•	•	•
	No clustering		•						•

Table 4.12 RMSE of GL and SVM, number of sensors, and power consumption for λ corresponding to 5% more error in training set with and without clustering for forecast length of 30s.

Test Day		1st cluster	2nd cluster	3rd cluster	4th cluster	Avg. over clusters	No clustering
1	GL RMSE	3.746	6.217	3.915	3.697	4.423	4.578
	Num sensors	1	2	2	3	1.50	1
	Power (mW)	99.6	192.8	239.6	321.2	155.7 (-14.6%)	99.6
	Size	2794	1081	830	284	1247	4989
2	GL RMSE	3.773	5.217	2.485	3.299	3.947	3.934
	Num sensors	1	3	2	3	1.69	5
	Power (mW)	99.6	286.0	239.6	286	172.0 (69.5%)	563.2
	Size	2694	1028	776	221	1180	4719
3	GL RMSE	4.145	5.749	3.365	4.101	4.342	4.389
	Num sensors	1	2	2	1	1.31	1
	Power (mW)	99.6	192.8	239.6	99.6	136.1 (-36.6%)	99.6
	Size	2348	614	571	227	940	3760
4	GL RMSE	3.496	6.030	3.685	3.537	4.240	4.404
	Num sensors	1	7	2	6	2.72	1
	Power (mW)	99.6	702.8	239.6	656.4	279.6 (-181%)	99.6
	Size	3004	1141	574	215	1233	4934
Avg.	GL RMSE	-	-	-	-	4.238	4.326
	Num sensors	-	-	-	-	1.80	2
	Power (mW)	-	-	-	-	185.8 (13.8%)	215.5

Table 4.13 List of sensors selected by group lasso with λ_{min} for each day used as test set and for each cluster as well as without clustering for forecast length of 30s.

Day	cluster	Body Temp	Heart Rate	Env. Temp	Humidity	Wrist accel	Ankle accel	EDA	Breath rate
1	1st			•					•
	2nd								•
	3rd		•						•
	4th						•	•	•
	No clustering								•
2	1st			•					•
	2nd				•				•
	3rd		•						•
	4th			•				•	•
	No clustering		•	•		•	•		•
3	1st			•					•
	2nd								•
	3rd		•						•
	4th								•
	No clustering								•
4	1st	•		•	•	•	•	•	•
	2nd		•						•
	3rd		•	•					•
	4th		•	•		•	•	•	•
	No clustering								•

BIBLIOGRAPHY

- [Ach06] Acharya, U. R. et al. "Heart rate variability: a review". *Medical and biological engineering and computing* **44.12** (2006), pp. 1031–1051.
- [Bir17] Birrenkott, D. et al. "A robust fusion model for estimating respiratory rate from photoplethysmography and electrocardiography". *IEEE Transactions on Biomedical Engineering* (2017).
- [Blo07] Blount, M. et al. "Remote health-care monitoring using Personal Care Connect". *IBM systems journal* **46.1** (2007), pp. 95–113.
- [BG98] Bock, J. & Gough, D. A. "Toward prediction of physiological state signals in sleep apnea". *IEEE Transactions on Biomedical Engineering* **45.11** (1998), pp. 1332–1341.
- [Boi98] Boiten, F. A. "The effects of emotional behaviour on components of the respiratory cycle". *Biological psychology* **49.1-2** (1998), pp. 29–51.
- [Boi94] Boiten, F. A. et al. "Emotions and respiratory patterns: review and critical analysis". *International journal of psychophysiology* **17.2** (1994), pp. 103–128.
- [Bou12] Boulmalf, M. et al. "A lightweight middleware for an e-health WSN based system using Android technology". *Multimedia Computing and Systems (ICMCS), 2012 International Conference on*. IEEE. 2012, pp. 551–556.
- [Boz16] Bozkurt, A. et al. "A Biobotic Distributed Sensor Network for Under-Rubble Search and Rescue". *Computer* **49.5** (2016), pp. 38–46.
- [Che12] Chen, L. et al. "Sensor-based activity recognition". *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* **42.6** (2012), pp. 790–808.
- [Col17] Cole, J. et al. "A study on motion mode identification for cyborg roaches". *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE. 2017, pp. 2652–2656.
- [Dav92] Davis, M. "The role of the amygdala in fear and anxiety". *Annual review of neuroscience* **15.1** (1992), pp. 353–375.
- [Die06] Diebel, J. "Representing attitude: Euler angles, unit quaternions, and rotation vectors". *Matrix* **58.15-16** (2006), pp. 1–35.
- [Dua14] Dua, S. et al. *Machine learning in healthcare informatics*. Springer, 2014.

- [Fre07] French, B. et al. “Selective sampling strategies to conserve power in context aware devices”. *Wearable Computers, 2007 11th IEEE International Symposium on*. IEEE. 2007, pp. 77–80.
- [GS13] Ghasemzadeh, H. & Shirazi, B. “Context-aware signal processing in medical embedded systems: A dynamic feature selection approach”. *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*. IEEE. 2013, pp. 642–645.
- [Gha15] Ghasemzadeh, H. et al. “Power-aware computing in wearable sensor networks: An optimal feature selection”. *IEEE Transactions on Mobile Computing* **14.4** (2015), pp. 800–812.
- [Gom12] Gomez, C. et al. “Overview and evaluation of bluetooth low energy: An emerging low-power wireless technology”. *Sensors* **12.9** (2012), pp. 11734–11753.
- [GD04] Gomez, P. & Danuser, B. “Affective and physiological responses to environmental noises and music”. *International Journal of psychophysiology* **53.2** (2004), pp. 91–103.
- [Hal14] Halson, S. L. “Monitoring training load to understand fatigue in athletes”. *Sports Medicine* **44.2** (2014), pp. 139–147.
- [HN04] He, X. & Niyogi, P. “Locality preserving projections”. *Advances in neural information processing systems*. 2004, pp. 153–160.
- [Her17] Hernando, A. et al. “Finger and forehead PPG signal comparison for respiratory rate estimation based on pulse amplitude variability”. *Signal Processing Conference (EUSIPCO), 2017 25th European*. IEEE. 2017, pp. 2076–2080.
- [HK70] Hoerl, A. E. & Kennard, R. W. “Ridge regression: Biased estimation for nonorthogonal problems”. *Technometrics* **12.1** (1970), pp. 55–67.
- [HM08] Homma, I. & Masaoka, Y. “Breathing rhythms and emotions”. *Experimental physiology* **93.9** (2008), pp. 1011–1021.
- [Joh12] Johnstone, J. A. et al. “BioHarness multivariable monitoring device: part. I: validity”. *Journal of sports science & medicine* **11.3** (2012), p. 400.
- [KL03] Keerthi, S. S. & Lin, C.-J. “Asymptotic behaviors of support vector machines with Gaussian kernel”. *Neural computation* **15.7** (2003), pp. 1667–1689.
- [Ker12] Kertis, J. D. “Biomechanical evaluation of an optical system for quantitative human motion analysis” (2012).
- [Kes90] Kesten, S. et al. “Respiratory rate during acute asthma”. *Chest* **97.1** (1990), pp. 58–62.

- [LB12] Latif, T. & Bozkurt, A. “Line following terrestrial insect biobots”. *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE. 2012, pp. 972–975.
- [Lee06] Lee, D.-S. et al. “Vital sign monitoring system with life emergency event detection using wireless sensor network”. *Sensors, 2006. 5th IEEE Conference on*. IEEE. 2006, pp. 518–521.
- [Lok17] Lokare, N. et al. “Activity-Aware Physiological Response Prediction Using Wearable Sensors”. *Inventions* **2.4** (2017), p. 32.
- [Lym03] Lymberis, A. “Smart wearable systems for personalised health management: current R&D and future challenges”. *Engineering in Medicine and Biology Society, 2003. Proceedings of the 25th Annual International Conference of the IEEE*. Vol. 4. IEEE. 2003, pp. 3716–3719.
- [Maj17] Majumder, S. et al. “Wearable sensors for remote health monitoring”. *Sensors* **17.1** (2017), p. 130.
- [MH01] Masaoka, Y. & Homma, I. “The effect of anticipatory anxiety on breathing and metabolism in humans”. *Respiration physiology* **128.2** (2001), pp. 171–177.
- [MS14] Medicare, C. for, Services, M., et al. *National Health Expenditure Projections 2011-2021, Forecast Summary*. 2014.
- [Moc17] Mochizuki, K. et al. “Importance of respiratory rate for the prediction of clinical deterioration after emergency department discharge: a single-center, case–control study”. *Acute Medicine & Surgery* **4.2** (2017), pp. 172–178.
- [Moh18] Mohammadzadeh, F. et al. “Prediction of physiological response over varying forecast lengths with a wearable health monitoring platform”. *2018 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE. 2018.
- [Moh15] Mohammadzadeh, F. F. et al. “Feasibility of a wearable, sensor-based motion tracking system”. *Procedia Manufacturing* **3** (2015), pp. 192–199.
- [Nav09] Navarro, K. F. et al. “Medical MoteCare: A distributed personal healthcare monitoring system”. *eHealth, Telemedicine, and Social Medicine, 2009. eTELEMED’09. International Conference on*. IEEE. 2009, pp. 25–30.
- [Nic17] Nicolò, A. et al. “Respiratory Frequency during Exercise: The Neglected Physiological Measure”. *Frontiers in physiology* **8** (2017).
- [Pér10] Pérez, R. et al. “Upper limb portable motion analysis system based on inertial technology for neurorehabilitation purposes”. *Sensors* **10.12** (2010), pp. 10733–10751.

- [Rav14] Ravish, D. et al. “Heart function monitoring, prediction and prevention of heart attacks: Using artificial neural networks”. *Contemporary Computing and Informatics (IC3I), 2014 International Conference on*. IEEE. 2014, pp. 1–6.
- [Ren10] Ren, Y. et al. “Monitoring patients via a secure and mobile healthcare system”. *IEEE Wireless Communications* **17.1** (2010).
- [Shn05] Shnayder, V. et al. “Sensor networks for medical care” (2005).
- [Sua16] Suarez, F. et al. “Designing thermoelectric generators for self-powered wearable electronics”. *Energy & Environmental Science* **9.6** (2016), pp. 2099–2113.
- [Sur12] Suryadevara, N. et al. “Intelligent sensing systems for measuring wellness indices of the daily activities for the elderly”. *Intelligent Environments (IE), 2012 8th International Conference on*. IEEE. 2012, pp. 347–350.
- [Tib96] Tibshirani, R. “Regression shrinkage and selection via the lasso”. *Journal of the Royal Statistical Society. Series B (Methodological)* (1996), pp. 267–288.
- [Uni17] United States Environmental Protection Agency. *2012 National Ambient Air Quality Standards (NAAQS) for Particulate Matter (PM)*.
<https://www.epa.gov/pm-pollution/2012-national-ambient-air-quality-standards-naaqs-particulate-matter-pm>. 2017.
- [Vap95] Vapnik, V. “The nature of statistical learning theory Springer New York Google Scholar” (1995).
- [Whi07] Whitchurch, A. K. et al. “Design and development of a wireless remote point-of-care patient monitoring system”. *Region 5 Technical Conference, 2007 IEEE*. IEEE. 2007, pp. 163–166.
- [Xia10] Xiao, F. et al. “Heart rate prediction model based on physical activities using evolutionary neural network”. *Genetic and Evolutionary Computing (ICGEC), 2010 Fourth International Conference on*. IEEE. 2010, pp. 198–201.
- [Xin16] Xing, Y.-F. et al. “The impact of PM_{2.5} on the human respiratory system”. *Journal of thoracic disease* **8.1** (2016), E69.
- [YL06] Yuan, M. & Lin, Y. “Model selection and estimation in regression with grouped variables”. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68.1** (2006), pp. 49–67.
- [YJ08] Yuchi, M. & Jo, J. “Heart rate prediction based on physical activity using feedforward neural network”. *Convergence and Hybrid Information Technology, 2008. ICHIT'08. International Conference on*. IEEE. 2008, pp. 344–350.

- [Zap08] Zappi, P. et al. “Activity recognition from on-body sensors: accuracy-power trade-off by dynamic sensor selection”. *Wireless sensor networks*. Springer, 2008, pp. 17–33.
- [Zha13] Zhang, J.-T. et al. “Concurrent validation of Xsens MVN measurement of lower limb joint angular kinematics”. *Physiological measurement* **34.8** (2013), N63.
- [Zho08] Zhou, H. et al. “Use of multiple wearable inertial sensors in upper limb motion tracking”. *Medical engineering & physics* **30.1** (2008), pp. 123–133.