Title
Non-Intrusive Analytical Approaches to Obstructive Sleep Apnea Detection

Author
Samy, Lauren

Publication Date
2016

Peer reviewed|Thesis/dissertation

University of California
Los Angeles

Non-Intrusive Analytical Approaches to Obstructive Sleep Apnea Detection

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Computer Science

by

Lauren Samy

2016

Abstract of the Dissertation

Non-Intrusive Analytical Approaches to Obstructive Sleep Apnea Detection

by

Lauren Samy

Doctor of Philosophy in Computer Science

University of California, Los Angeles, 2016

Professor Majid Sarrafzadeh, Chair

Obstructive sleep apnea (OSA) is the most common type of sleep apnea and the most prevalent sleep disorder in general. Approximately 22 million people in the United States and more than 100 million people worldwide are affected by this serious disorder. The disorder is characterized by repeated involuntary breathing cessations (apnea) or reductions in breathing (hypopnea) during sleep, accompanied by oxygen desaturation (hypoxia). The apneas occur due to the obstruction of the upper airway as a result of the relaxing of throat tissue. The hypoxia triggers a brain activity arousal to restore the normal breathing pattern until the next apneic episode occurs. OSA patients can experience hundreds of apneas per night. The frequent sleep disruptions and the resulting fragmented sleep pattern can have serious physical and psychological consequences and can lead to premature death.

Because the apnea episodes rarely trigger a full awakening, patients are often unaware of having difficulty breathing at night. As a result of this unawareness, OSA is severely underdiagnosed. This problem is exacerbated by the intrusiveness of current diagnosis methods. Polysomnography—the traditional gold-standard diagnosis method—requires close overnight monitoring of patients' body functions, making it an uncomfortable and intrusive diagnostic test. Intrusive diagnosis methods discourage patients from getting tested for OSA until they experience its serious health effects. As a result, many patients remain untreated and OSA continues to be a major public health problem that places a significant burden on health care systems worldwide.

In an effort to increase the level of diagnosis and treatment of the disorder, this work investigates the design and implementation of effective non-intrusive OSA diagnostic methods. We explore two kinds of detection methods: nocturnal methods that rely on diagnosing OSA by detecting apneic episodes from overnight recordings, as well as daytime methods which use features and signals that can be obtained while the subject is awake. This work involves identifying the data to be collected, building the appropriate systems to obtain them, and using machine learning techniques to analyze them.

The effectiveness of our methods was proven using real data from clinical trials. Daytime methods quickly revealed their advantages over the common nocturnal methods and the few existing daytime methods. Experimental results show that our daytime methods perform the initial diagnosis of OSA non-intrusively without significantly affecting the diagnosis accuracy compared to the current state-of-the-art nocturnal methods.

The dissertation of Lauren Samy is approved.

Paul Michael Macey

Miryung Kim

Mario Gerla

Majid Sarrafzadeh, Committee Chair

University of California, Los Angeles

2016

To my parents, for their continual love and support
and for all the opportunities they've given me over the years
and to my sister, for her encouragement and motivation
and the hours of "peaceful piano" that have sustained me throughout this process.

Table of Contents

List of Figures

List of Tables

Acknowledgments

I would like to express my gratitiude to my advisor, Professor Majid Sarrafzadeh, for giving me the opportunity to work in his group and for his supervision throughout my graduate career. I would also like to express my thanks and appreciation to my committee members—Professor Gerla, Professor Kim, and Professor Macey—for taking time to serve on my committee and for providing me with insightful feedback. I would also like to extend my sincere gratitude to Professor Paul Macey for his guidance throughout the process. His discussions, ideas, and feedback have been absolutely invaluable.

I would also like to give a very special thanks to my undergraduate research advisor, Professor Frank Vahid. It is under his supervision that I first got introduced to Computer Science and to academic research. I will always be thankful for his mentorship, encouragement and support, which have helped me grow in knowledge and in confidence. He was and remains a great role model for a scientist, mentor, and teacher.

Throughout my PhD studies, I have had the opportunity to collaborate with numerous researchers. I would like to thank Amir Vajid, Mars Lan, Nabil Alshurafa, Jason Liu, and Suneil Nyamathi for contributing to different aspects of this work and my lab-mates, Bobak Mortazavi, Diane Woodbridge and Ivan Lee, for their help when I first joined the group.

I would also like to thank Turker Garip for all the times he interrupted his work to help me with mine, for his tremendous patience and moral support whenever I needed it, and for continuing to be an amazing friend. I would like to thank Amir Vajid for being a great friend and study partner and for the numerous after-work trips he took from San Diego to Los Angeles to help me with a project we started before he graduated. I would also like to thank my roommate and friend, Marian Antonious, for always being there for me and for sharing the graduate school experience with me.

Finally, I would like to express my gratitude to my family. I thank my parents and my sister for their never-failing love and support, which helped me get to where I am today, and my cousins for always being ready to have fun, comfort me, and deliver home-made food to my apartment.

Vita

2011	B.S. (Computer Science), University of California, Riverside.
2014	M.S. (Computer Science), University of California, Los Angeles
2011-present	Teaching Assistant/Associate/Fellow, Computer Science Department, UCLA. Taught sections of Computer Science 35L (Introductory software construction course) under direction of Professor Paul Eggert. During summer 2012, taught the Data Structures and Algorithms course. During summer 2015, taught an introductory Computer Science course.

Publications

Lauren Samy, Paul M. Macey, Majid Sarrafzadeh. A Daytime Obstructive Sleep Apnea Severity Assessment Framework. Engineering in Medicine and Biology Society (EMBC). 2016

Lauren Samy, Paul M. Macey, Majid Sarrafzadeh. A Gender-Aware Framework For The Daytime Detection Of Obstructive Sleep Apnea. Engineering in Medicine and Biology Society (EMBC). 2015

Constantinos Sideris, Nabil Alshurafa, Mohammad Pourhomayoun, Farhad Shahmohammadi, Lauren Samy, Majid Sarrafzadeh, A Data-Driven Feature Extraction Framework For Predicting The Severity Of Condition Of Congestive Heart Failure Patients. Engineering in Medicine and Biology Society (EMBC). 2015

Lauren Samy, Ming-Chun Huang, Jason Liu, Wenyao Xu, Majid Sarrafzadeh. Unobtrusive Sleep Stage Identification Using A Pressure-Sensitive Bed Sheet. IEEE Sensors Journal Special Issue on Material-Integrated Sensing (Sens. J.). 2013

Lauren Samy, Paul M. Macey, Nabil Alshurafa, Majid Sarrafzadeh. An Automated Framework For Predicting Obstructive Sleep Apnea. Conference On Pervasive Technologies Related To Assistive Environments (PETRA). 2015 Conference on Healthcare Informatics, Imaging and Systems Biology (HISB). 2012

Ming-Chun Huang, Wenyao Xu, Jason Liu, Lauren Samy, Amir Vajid, Majid Sarrafzadeh. Inconspicuous On-Bed Respiratory Rate Monitoring. Conference On Pervasive Technologies Related To Assistive Environments (PETRA). 2013

Myung-Kyung Suh, Jonathan Woodbridge, Tannaz Moin, Mars Lan, Nabil Alshurafa, Lauren Samy, Bobak Mortazavi, Hassan Ghasemzadeh, Alex Bui, Sheila Ahmadi, Majid Sarrafzadeh. Dynamic Task Optimization in Remote Diabetes Monitoring Systems.

Mars Lan, Lauren Samy, Nabil Alshurafa, Myung-Kyung Suh, Hassan Ghasemzadeh, Majid Sarrafzadeh. WANDA: An End-to-End Remote Health Monitoring and Analytics System for Heart Failure Patients. ACM Conference on Wireless Health (WH). 2012

CHAPTER 1

Introduction

1.1 Sleep Apnea (SA)

Sleep apnea is a sleep disorder that is characterized by episodes of breathing cessation (apnea) or abnormally low breathing (hypopnea) during sleep. Each apnea or hypopnea event can last from a few seconds to a few minutes, and may occur 5 to 30 times or more an hour. The lack of breathing causes a reduction in blood oxygen saturation and an increase in the carbon dioxide levels in the blood. Oxygen depletion, especially to the brain, can have detrimental health effects since brain cells need constant oxygen to live. Furthermore, the imbalanced levels of respiratory gases signal the brain to wake the sleeper up to start breathing again. Although the sleeper does not fully wake up, these frequent arousals cause a transition into a lighter stage of sleep and result in a fragmented sleep pattern. Fragmented sleep interrupts the physical and mental repair and regeneration work carried out during sleep and can have serious physiological and psychological effects. Even though the manifestation of SA is always the same, namely repetitive episodes of decreased or total cessation of respiratory airflow during sleep, the cause of these apnea/hypopnea episodes is often different. Based on the cause of the breathing problem, SA can be categorized into three different types, which are briefly explained in the next subsections.

1.1.1 Obstructive Sleep Apnea (OSA)

Obstructive sleep apnea is by far the most common type of SA and the most common sleep disorder in general [Org]. When defined as >5 episodes of apnea or hypopnea per hour of sleep, OSA affects 24% and 9% of middle-aged men and women, respectively [YPD93]. OSA arises from what is essentially a mechanical problem. During sleep, the

muscle tone of the body relaxes. For OSA patients, when the throat muscles relax, the tongue and soft tissues of the throat fall back into the airway and block airflow. Despite efforts to breathe, breaths are hampered by a physical blockage to airflow. This work is focused on this type of SA but, for the sake of completeness, the other two types are explained in the next subsections.

1.1.2 Central Sleep Apnea (CSA)

Unlike OSA, central sleep apnea is characterized by a diminished or absent effort to breathe. While OSA is caused by a physical obstruction in the airway, CSA is largely caused by an imbalance in the brain and central nervous system during sleep. The brain is responsible for monitoring and adjusting the oxygen levels in the blood. In healthy people, this feedback mechanism results in a consistent intake of oxygen and outtake of carbon dioxide. In CSA, the brain is slow to react to changes in the oxygen and carbon dioxide levels in the body during sleep. As a result, the sleeper experiences long missed breaths, followed by short rapid breaths (hyperpnea) as the brain attempts to increase the oxygen levels by overcompensating for the missed breaths. Unlike obstructive apneas, there is no effort made to breathe during a central apnea and there are no chest movements. CSA is a much rarer type than OSA.

1.1.3 Mixed Sleep Apnea (MSA)

Mixed sleep apnea is the least common type of SA and, as the name implies, is a combination of OSA and CSA. MSA episodes usually begin with a reduced central respiratory effort, followed by obstruction. When the symptoms of OSA are treated, the brain is still unable to control the patient's breathing properly. In other words, while treated for OSA, the patient still suffers from CSA.

1.2 Obstructive Sleep Apnea

This section focuses on the OSA syndrome and explains the risk factors, symptoms, and morbidities associated with it. It then goes on to describe the conventional method used

in diagnosing the disorder, as well as its many shortcomings and drawbacks. Finally, the treatment options for OSA are discussed and briefly explained. Understanding the information presented in this section was essential for developing the solutions proposed in this work.

1.2.1 Risk Factors

OSA is a very common disorder. In the US alone, 22 million people are affected by it, making it as common as asthma (20 million) and type II diabetes (23.6 million) [ASAa], not counting people who have the disorder but have not been diagnosed. Even though OSA can strike anyone at any age (even children), there are some risk factors that increase the probability of a person developing OSA. Some of the risk factors include the male sex, obesity (BMI >30), excessive use of alcohol and sedatives, smoking, family history of OSA, being age 40 or older, ethnicity (African-Americans, Pacific-Islanders and Hispanics), and upper airway and facial abnormalities, like having a large tongue or tonsils, a recessed chin, a large overbite, a small jaw, or a larger neck size (>17" for men, >16" for women).

1.2.2 Symptoms

OSA is very difficult to detect because the disorder only occurs during sleep and because the same signs and symptoms of OSA can also occur in healthy people. Most of the time, OSA is discovered by the patient's bed partner and not by the patient himself/herself. A strong indicator of OSA is chronic snoring. The snoring sound results from air passing through a partially blocked airway, which a sleep apneic would experience during a hypopnea episode. Also, at the end of an apnea episode, a sleep apneic can be heard gasping for air, choking or coughing as breathing is resumed. Due to their fragmented sleep pattern, which prevents them from attaining deep, regenerative sleep, OSA patients tend to be sleep deprived and often suffer from excessive daytime sleepiness (EDS). EDS often results in falling asleep at work, on the phone or while driving, as well as a wide range of other symptoms, such as difficulty concentrating, lack of energy, morning headaches, irritability, sexual dysfunction, learning and memory difficulties, and depression.

1.2.3 Comorbidity & Mortality

Disorders or diseases that commonly co-occur with SA are called comorbidities of the disorder. Even though there aren't necessarily any causal relationships between the co-morbid conditions and SA, the associations may reflect a role of SA in their origination. Drug-resistant hypertension is one of the most prevalent comorbidities of SA and is characterized by high blood pressure despite adherence to medication. 83% of people with drug-resistant hypertension are also OSA patients [LPM01]. Other very common comorbidities include congestive heart failure (CHF), coronary artery disease (CAD), and type II diabetes mellitus. 76% (40% CSA and 36% OSA) of all CHF patients, 30.5% (all OSA) of all CAD patients, and 48% of all type II diabetes patients were found to have SA [OLF06], [SKE99], [ESE07]. Other physiological comorbidities include stroke, cardiac arrhythmia, heart attack, atrial fibrillation, as well as a three-fold increase in mortality risk [YFP08]. Aside from the physiological comorbidities, OSA and SA in general, is associated with a variety of psychological comorbidities, including depression, anxiety and stress.

One explanation for OSA's accompanying diseases is the fragmented, non-restorative sleep pattern that OSA causes. During a normal sleep cycle, human beings go through different sleep stages. The two most important ones are the Non-Rapid-Eye-Movement-3 (NREM3) and Rapid-Eye-Movement (REM) sleep stages. The NREM3 stage is critical for the renewal and repair of the body. In NREM3, the body's immune system is repaired and the pituitary gland releases growth hormone, which stimulates tissue growth and muscle repair. In the REM sleep stage, on the other hand, the body is in complete paralysis and the brain is the center of attention. During REM, the brain dreams, reorganizes information and consolidates memory, which facilitates learning and neural growth. In short, NREM3 helps with physical recovery while REM sleep helps with mental recovery. Because of their frequent arousals from these two essential sleep stages, OSA patients don't get the necessary physical and mental recovery and therefore develop an increased risk of the aforementioned physiological and psychological illnesses, as well as mortality. The scientific processes by which OSA leads to the comorbidities mentioned above are outside the scope of this work.

1.2.4 Conventional Diagnosis: Polysomnography (PSG)

Polysomnography (PSG) is the conventional method for diagnosing OSA. PSG is an overnight sleep test that is conducted in a specialized sleep laboratory. During the test, the sleeper's biophysical changes that occur during the night are monitored and later analyzed and interpreted by a sleep technician to determine the presence of OSA.
The next two paragraphs explain the PSG mechanism in more detail and point out its drawbacks that make it ineffective as an early-stage diagnostic tool, further contributing to the underdiagnosis of OSA.

PSG Mechanism Polysomnography uses a comprehensive set of sensors to monitor bodily functions during sleep and accurately identify the sleep stage that the sleeper is in. These sensors are attached to the sleeper's body and the wires for each sensor converge into a central box, which in turn is connected to a computer system for recording, storing and displaying the data. Figure 1.1 shows all PSG sensors and their respective locations on the sleeper's body. The recorded signals include:

- Electroencephalogram (EEG): Brain waves are recorded in order to detect sleep stages, arousals and other aspects of the sleeper's sleep quality. As shown in Figure 1.1 a), 8 EEG electrodes are attached to the sleeper's head to measure brain activity. The electrodes measure the frontal (red), central (blue), medial (yellow) and occipital (green) EEG signals which have different characteristics and frequencies during the various stages of sleep.

- Electrooculogram (EOG): To measure eye movement, two EOG electrodes are attached near the eyes as shown in Figure 1.1 b); one is attached above the right eye and the other is attached below the left eye. Eye movement can help determine if the sleeper is in the REM (Rapid Eye Movement) stage.

- Electromyogram (EMG): EMG electrodes measure muscle activity and are used in two places. Three EMG electrodes are attached to the sleeper's chin to detect chin muscle movement which can be indicative of teeth grinding (bruxism) – a sign of sleep apnea [Fou]. As shown in Figure 1.1 b), the two EMG electrodes to the left and right of the chin midline are the negative and positive leads, respectively,

while the one in the middle of the chin is the reference EMG lead. Figure 1.1 d) shows another pair of EMG electrodes which is attached to the legs, centered between the knee and the ankle, to record leg movement. An excessive amount of leg movements during sleep could be indicative of periodic limb movement disorder (PLMD). EMG also helps determine when sleep occurs and when the REM stage occurs. During sleep, the body is relaxed and there's a significant decrease in muscle tension. During REM sleep, the sleeper becomes partially paralyzed and there's an even bigger decrease in skeletal muscle tension. Since REM is when dreaming takes place, the partial paralysis is a defense mechanism that prevents the sleeper from acting out dreams.

- Electrocardiogram (ECG): As shown in Figure 1.1 c), two ECG electrodes are attached to the sleeper; one is attached under the right clavicle (collarbone) and the other one under the rib cage on the left side of the body. ECG measures the electrical activity of the heart as it contracts and expands and can be an indicator for the sleep stage the sleeper is in since the rate and regularity of heartbeats vary for the different sleep stages.

- Nasal and oral airflow: As shown in Figure 1.1 b), a thermistor is placed directly below the nostrils. The thermistor measures the nasal and oral airflow as the sleeper inhales and exhales. In addition, a nasal cannula is placed on top of the thermistor and directly inside the nostrils. This sensor detects the fluctuations in pressure caused by inhalation and exhalation. These two sensors allow the sleep technician to monitor the sleeper's respiration rate and identify interruptions in breathing.

- Respiratory effort: In addition to nasal and oral airflow, respiratory effort is also measured in order to monitor respiration. Respiratory effort is measured by means of effort belts, as can be seen in Figure 1.1 c). The belts are worn around the thorax and abdomen to measure the expansion of the chest and abdomen during breathing. Respiratory effort is a better way of monitoring respiration rate because nasal and oral airflow can produce false positives. Some OSA patients will open and close their mouths during an apnea episode as they struggle to breath. This forces air in and out of the mouth, but due to the obstruction in the airway, no

air will enter the airway or lungs. In that case, the cannula and thermistor will detect this diminished airflow and the respiratory event may be falsely identified as a hypopnea, instead of an obstructive apnea. The belts, on the other hand, will only expand if there's air going in and out of the lungs and the sleeper is breathing.

- Pulse oximetry: As shown in Figure 1.1 d), a pulse oximeter is attached the the ring finger of the sleeper. The pulse oximeter determines changes in blood oxygen saturation that often occur with sleep apnea.

- Snoring: Figure 1.1 b) shows a microphone placed on the neck, lateral to the larynx, to detect snoring. Snoring indicates airflow and can be used during hypopneas to determine whether the hypopnea is due to OSA.

- Video: In addition, most sleep labs have a video camera in the room so the technician can observe the patient visually from an adjacent room.

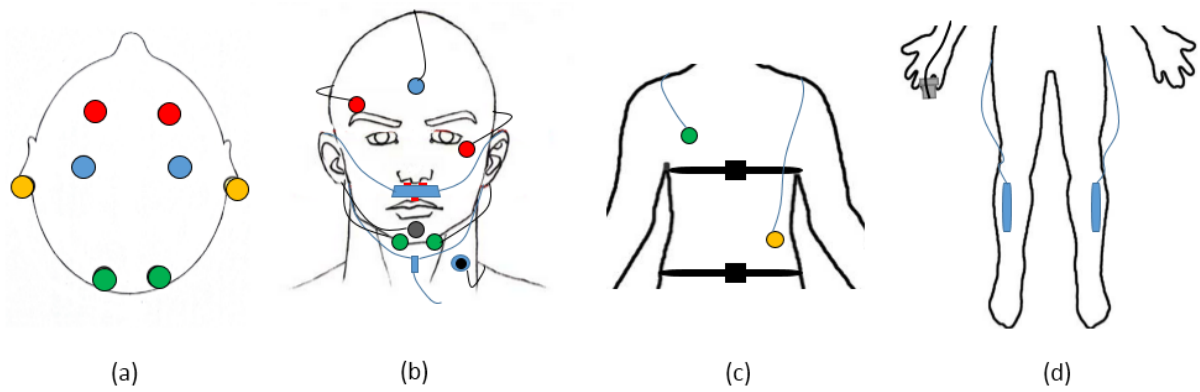

Figure 1.1: a) Birds eye view of a human head with 4 pairs of EEG electrodes attached to it to measure frontal (red), central (blue), occipital (green), and medial (yellow) EEG signals. b) 1 electrode is attached to the middle of the forehead for grounding, 2 EOG sensors are attached diagonally above the right eye and below the left eye to measure eye movement. 3 EMG electrodes are attached to the chin to measure chin movement. A microphone is attached to the neck for snore detection. A thermistor is placed below the nostrils to measure nasal and oral airflow and a nasal cannula is placed directly in the nostrils to measure pressure. c) 4 sensors are placed on the chest: 2 ECG electrodes (under clavicle and below rib cage) to measure heart rhythm and 2 belts to measure thorax as well as abdominal effort. d) 3 sensors are place on the lower body: 2 EMG electrodes between the knee and the ankle to measure leg movement and a pulse oximeter on the ring finger of the right hand to measure SpO2, pulse and plethysmography.

PSG Drawbacks Even though PSG is a multi-parametric test that measures and monitors a comprehensive set of signals and is considered the gold standard for sleep disorder diagnosis in sleep medicine, there are a number of serious drawbacks associated with that sleep test.

Because OSA does not trigger a full awakening and because it's a disorder that only manifests during sleep, an individual with OSA is rarely aware of having difficulty breathing. As a result, many people in the U.S. and worldwide remain undiagnosed and untreated. It is estimated that only 10-20% of people with OSA have been diagnosed [Sle08]. Usually, the primary physician will evaluate a patient's symptoms first and, based on it, he/she will recommend PSG – the only definitive diagnostic tool for determining the presence of OSA.

Cost At a few thousand dollars per sleep study, PSG is prohibitively expensive for many patients. Many patients are discouraged to spend that kind of money to get tested for a disorder that they might not even have and doctors are sometimes hesitant to recommend the test, let alone repeat it as is sometimes required.

Discomfort: First-Night Effect Another drawback is a phenomenon referred to as the first-night effect [SMY07]. PSG is an overnight test conducted at a sleep laboratory. The unfamiliar environment of the sleep laboratory as well as the inconvenience that results from equipment and sensors attached to the participant's body can greatly decrease the participant's quality of sleep and lead to distorted results. This is one of the reasons a test might need to be repeated. The very nature of PSG can interfere with the results and accuracy of the measured signals, especially on a patient's first night at the sleep laboratory.

Scarcity & Long Wait Times If a patient is willing to pay for PSG and endure the discomfort of sleeping at a sleep laboratory with wires attached, he/she is faced with another problem. The available resources for performing the sleep test, which include the sleep technician, the laboratory, and the equipment, are relatively scarce. This leads to very long wait times for patients, especially those with subtle symptoms who understandably end up being de-prioritized.

All these drawbacks contribute to a large population of OSA patients remaining undiagnosed. Undiagnosed, and thus untreated, OSA is estimated to cause $3.4 billion in additional medical costs in the U.S. [KBS99].

1.2.5 Treatment Options

Surgery The goal of surgery is to increase the size of the airway and thus reduce the apnea episodes. To do that, the surgeon may either reconstruct the jaw to enlarge the upper airway or he/she can identify the site of obstruction in the upper airway and remove excess tissue in that place. As this is a challenging task, if the surgeon does not treat the correct area or only one of many areas of obstruction, it is unlikely that the OSA will

be treated or even diminished to a degree that improves the patient's sleep quality. In some rare cases, OSA symptoms can even become worse after surgery. The success rate of surgery this about 50% and in many cases, the patient would have to undergo multiple operations before OSA is fully treated, which increases the risks of complications and infections [ASAb]. After surgery, the patient would have to do a sleep study (PSG) to get re-assessed for OSA to determine the success of the surgery.

Positional Therapy For some patients with mild OSA, the airway only collapses when they are sleeping in the supine position. Those patients can eliminate or reduce airway obstruction by learning to sleep on their side. To achieve that, the patient can sleep wearing one of the many different kinds of positioners available. A positioner is a belt that is worn around the upper torso and helps the sleeper stay off of his/her back. Alternatively, the sleeper can place three or four tennis balls in a pocket sewn on the back of a pajama top.

Continuous Positive Airway Pressure (CPAP) The most widely used treatment for moderate and severe sleep apnea is CPAP. The CPAP device is a face mask that is worn over the mouth and nose during sleep and whose purpose is to supply pressurized air continuously into the sleeper's throat. The increased air pressure keeps the sleeper's lungs full and prevents the airway from collapsing and causing obstruction. In order for CPAP to be effective, it needs to be worn for at least 6 hours every night. If regularly worn, CPAP is a very effective, but very intrusive, form of treatment. Due to the discomfort of sleeping with a face mask, a lot of patients stop using the device during the first few days of treatment. When adherence is defined as greater than 4 hours of nightly use, 46 to 83% of patients with OSA have been reported to be non-adherent to treatment [WG08a].

Mandibular Advancement Devices (MAD) An MAD is an oral appliance that is worn in the mouth during sleep. This appliance is meant to position the lower jaw slightly forward of its usual rest position. For some people, this small change is enough to keep the airway open during sleep. These oral devices are a bit more comfortable than CPAP but they're not always as effective.

Weight Loss The link between excess weight and OSA is well established. Overweight people have extra tissue in the back of their throat, which can fall down over the airway and block the airflow into the lungs while they sleep. About 70% of people with OSA are overweight or obese. Despite the dearth of studies that show the effectiveness of weight loss in diminishing the apnea episodes, practitioners report significant improvements in OSA among patients who lose weight. A study from 1985 ([SGM85]) and another from 1991 ([SGS91]) show a significant reduction in apnea frequency (apneas/hour) as patients decreased their weight. In some cases, physicians prescribe weight loss medications to obese patients with OSA.

1.3 In Closing

This chapter served as a brief overview of the sleep apnea disorder, in general, and obstructive sleep apnea in particular. The chapter discussed the conventional method of diagnosis and its disadvantages, the diseases and disorders associated with OSA, as well as the available methods of OSA treatment. The research in the wireless health field that is focused on OSA attempts to either 1) come up with new, less intrusive methods of diagnosing OSA, or 2) develop new approaches to improve OSA symptoms and treat comorbidities to improve patients' quality of life, or 3) prevent the progression of OSA or reverse the disorder by designing new treatment methods that are more comfortable and easier to adhere to.

This work focuses on the diagnosis problem. In the next two chapters, a variety of novel diagnosis methods are presented.

CHAPTER 2

Non-Intrusive Nocturnal Diagnosis Methods

2.1 Introduction

Since OSA is a disorder that manifests only during sleep, most diagnosis techniques have focused on monitoring subjects while they are sleeping. These nocturnal techniques attempt to "witness" apnea episodes through the monitoring of biophysical changes and symptoms that are good diagnostic predictors of an apnea episode. Polysomnography– discussed in Section 1.2.4–is the conventional diagnosis method which monitors a comprehensive set of biophysical signals that change during the night. The amount of wires and connections required to conduct such a multi-parametric assessment, as well as other reasons discussed in Section 1.2.4, makes PSG an undesirable and intrusive test. This chapter presents novel nocturnal methods that attempt to "witness" apnea episodes using completely non-intrusive monitoring tools and techniques.

2.2 Inconspicuous On-Bed Respiratory Rate Monitoring

As discussed in section 1.2.4, respiratory effort is one of the signals monitored during a sleep study. The respiratory effort signal can capture episodes of breath cessation. For OSA, the respiratory effort signal is more reliable than the airflow signal since some OSA patients will open and close their mouths during an apena episode. While this motion will create an airflow in and out of the mouth, it does not actually reflect any breathing. Respiratory effort, on the other hand, will only be detected if the air enters the airway and reaches the lungs and is therefore a better indicator of respiration. During a sleep study, the respiratory signal is acquired by means of two effort belts worn around the chest and abdomen. The pressure exerted on the belts corresponds to respiratory effort. Work in this chapter provides an inconspicuous method for respiratory effort

monitoring that can be used for unobtrusive OSA detection. An e-textile bed sheet embedded with a dense pressure sensitive sensor array system is introduced to measure human respiratory rate under a clinical or home environment. This system continuously detects a patient's pressure distribution on the bed. Respiratory rate is extracted via analyzing time-stamped pressure map sequences. Using this system, a subject can sleep freely without any wires or sensors attached to his/her body. Subjects under test may not even notice that they are being monitored. This system is therefore a type of ambient intelligence (AmI) that is integrated in a daily-life object and fits seamlessly into the users' environment without interfering with their everyday habits.

2.2.1 Related Works

In a typical hospital setting, respiratory rate is usually measured by manual observation or, if continuous monitoring is required, by attaching sensors to the patient. In the former case, the measurement is often inaccurate due to the different ways of measuring respiratory rate. For instance, measuring respiratory rate using a stethoscope yields different results than manually tracking the movement of a patient chest. Discrepancies were also reported between counting the number of breaths within a 30-second period and multiplying the result by 2 and counting breaths for 60 seconds [Fle10]. Also, if patients are aware that their respiratory rate is being monitored, they may involuntarily control their breathing, resulting in inaccurate measurements.

Because measuring respiration rate accurately is a challenging task, there are not many sensor products commercially available to solve this problem. Most existing sensors are for use in hospitals and are part of a larger medical system that aggregates data from other sensors. One example is the Acoustic Respiration Rate monitor (RRa) designed by Masimo [Mas], which uses an adhesive sensor to be placed on the patient's neck while he/she is sleeping in order to measure respiration rate. The sensor is an acoustic transducer that generates an audio signal that is then interpreted to identify inhalation and exhalation. Bates et al. placed 3-axis MEMS sensors composed of accelerometers and gyroscopes on the patient's torso to continuously infer respiratory rate from the acceleration and angle acceleration information [BLM10]. Both authors considered their system setup to be non-intrusive even though it clearly interferes with people's sleep

habits.

Because of the lack of commercial non-intrusive solutions, researchers have been attempting to offer solutions where the sensors are hidden from the subjects as to not disrupt their daily routine. One such solution is the 40-kHz ultrasound transmitter-receiver pair proposed by Yamana et al [YTM11]. The transmitter-receiver is placed under the bed mattress for monitoring. The transmitted ultrasound signal is reflected on the bottom surface of the mattress and the amplitude of the received ultrasonic wave is modulated by the particular shape of the mattress. This reflected signal is used to identify respiratory activity. Though non-intrusive, this method of respiratory rate monitoring was severely affected by human body movement, which was an unavoidable source of noise that negatively affected the accuracy of the system. Another non-intrusive respiration rate monitoring approach was proposed by Zhu et al [ZCN06]. Water-filled vinyl tubes were placed under the user's pillow to retrieve pressure signals. This pressure signal has a static component that corresponds to the weight of the head and a dynamic component that reflects the weight fluctuation of the head due to breathing movements. The subjects had to sleep in the supine and recumbent positions for data collection. Though non-intrusive, the system is limited to the mentioned sleep positions and the sleeping subjects have to be well-controlled so as to have their whole head on the pillow at all times. To slightly relax the restriction on human movement, Townsend et al. utilized an array of pressure sensors which was placed above the bed mattress to identify respiratory activity [THG09]. The device they used is a bed occupancy sensor equipped with 24 pressure sensors. The sensor is to be placed just below the pillow region in order to localize on the areas of respiration, namely the chest. The authors had users sleep in prone, supine and side positions and asked them to sleep apnea by holding their breath. The results showed that there was correlation between pressure variance and the occurrences of apnea. Using a pressure sensor array slightly relaxed the restriction of the human body movement, but the sleeper is still limited to stay within the area covered by the sensor array in order to be monitored.

The work proposed in this section describes a system and a series of methods that utilize a full-size, e-textile-based, high density pressure-sensitive bed sheet that solves all the issues mentioned above. With 8192 pressure sensors embedded in it, the proposed

high-density sensor can capture the full pressure distribution of the sleeper's body. High-resolution, full-body pressure information enables accurate human sleep posture detection and chest area localization. Hence, the system enables non-intrusive respiration rate monitoring without any restrictions on sleeping posture or body movement. In addition, the proposed e-textile-based bed sheet is a type of ambient intelligence that can be seamlessly embedded in the environment and is an example of how technology should adapt to the user's behavior and not the other way around.

2.2.2 System Architecture

The e-textile-based hardware, as well as the overall operation of the accompanying software are described below.

2.2.2.1 Hardware Architecture

E-textiles are fabrics that have electronics and interconnections woven into them. E-textiles are perfectly suitable for monitoring applications because of their flexibility and cost-effectiveness. E-textiles feel just as comfortable as regular fabrics but have one important difference that makes them suitable for serving as pressure sensors: the resistance of e-textiles changes when a force is applied to them. E-textiles are built by sandwiching a sheet of very thin piezoresistive fabric between two sheets of traditional textile fabric [XLH11]. The middle e-textile layer has a couple of useful characteristics: 1) the electrical resistance of the piezoelectric fabric decreases with increasing pressure force, 2) the fabric's thickness is on the order of microns, comparable to the diameter of a human hair. These two characteristics enable us to build a highly flexible, comfortable and low-cost pressure-sensing system that perfectly fits our respiration rate monitoring application. The proposed system consists of an e-textile bed sheet that records the pressure distribution of the body. The prototype bed sheet is 1.25 m wide x 2.5 m long and can fit easily on any standard-size bed. The thickness of the bed sheet is 1.5 mm, which makes it flexible and suitable for non-intrusive applications. The bed sheet is a matrix of 8192 pressure sensors generated by the intersections of 64 columns and 128 rows of conductive buses. To build the sensors, we use a three-stacked-layer structure as shown in Figure

2.1. The top layer is a conductive bus that is connected to a voltage supply via an analog multiplexer (M_1). The middle layer consists of the pressure-sensitive e-textile piezoresistive fabric. Lastly, the bottom layer is a conductive bus that is orthogonal to the top bus. This bottom layer connects to an analog-to-digital converter (ADC) via a second analog multiplexer (M_2) such that each input to the ADC is grounded via an offset resistor (R_0). The intersection areas between the orthogonal buses form the individual sensors. A microcontroller is used to set the addresses of both M_1 and M_2 to uniquely select a pressure sensor. For example, when M_1 connects bus i on the top layer to a voltage supply, and M_2 connects bus j to the ADC, the sensor located at row i, column j, which is denoted as V_{ij} in Figure 2.1 a), will be read. A sampling unit is connected to all conductive lines and performs matrix scanning to measure pressure map sequences. Pressure map signals retrieved from each of the 8192 sensors are converted by the ADC from a voltage to an 8-bit integer with a value of 0 representing no pressure (highest e-textile resistance) and a value of 255 representing maximum pressure (lowest e-textile resistance). The sensor values are then sent to an Android tablet over USB for storage and further analysis. The sampling rate of the system is adjustable up to 10Hz, but 1.5Hz was used for the respiratory rate measurement application. This allows the system to achieve a maximum detectable breathing rate of up to 45 breaths per minute (bpm) according to the Nyquist rule. As shown in Table 2.1, newborns have the fastest respiratory rate of 44 bpm due to their small lung capacity, which results in frequent breathing. Since the highest frequency component of our signal is 44 bpm, a low sampling rate of 1.5 Hz is more than sufficient for our application and will reduce the power consumption of the system.

Figure 2.1 b) shows a cross-section view of the bed sheet. The e-textile piezoresistive fabric is sandwiched between two orthogonal conductive bus layers, as described above. The advantage of this design is that the top and bottom layers can be made out of traditional fabric coated with parallel conductive buses with the e-textile piezoresistive fabric embedded in between. Another advantage of this design is that the MxN sensor structure only requires 2(M+N) I/O pins. Also, the matrix structure of the bed sheet enables random access to an arbitrary sensor in the system. Figure 2.2 a) shows an example of a user lying on the bed sheet and Figure 2.2 b) shows the corresponding pressure map. Compared to other systems that use bed mattresses, sheets or pads for

Table 2.1: The normal respiratory rates of the different age groups. As evident by the table, the human respiratory rate becomes fixed after the 7th year of life. This is because the lung volume have grown fully by that age. Newborns have the smallest lung volume and therefore also the highest respiration rate in order to inhale enough oxygen for sustaining their metabolism.

Normal Respiratory Rates By Age	
Newborns	44 respirations per minute
Infants	20-40 respirations per minute
Children (1-7 years)	18-30 respirations per minute
Adults	12-20 respirations per minute

health applications ([MMK10], [KMB10]), the described system has the flexibility of e-textiles as well as a dense high-resolution pressure map that enables high-quality medical diagnosis.

2.2.2.2 Software Architecture

An Android application was created to acquire data from the bed sheet to a tablet, as well as visualize them. The data can be sent to a server for analysis or the analysis can be done locally on the tablet if no network is available. To visualize the data, the array of pressure sensors is displayed in a 64 x 128 grid that matches the layout of the bed sheet. Each cell in the visualized pressure map corresponds to a sensor on the bed sheet. The cells are color coded to represent the magnitude of the data byte retrieved from the corresponding sensor on the bed sheet. The different colors represent pressure values of significantly different magnitudes, as shown in Figure 2.2 b). The high-level software flow of the respiration rate monitoring system is depicted by the state machine in Figure 2.3.

NO USER: Initial System State The NO USER state is the initial state of the system which indicates that no user is occupying the bed sheet. The system remains in the NO USER state as long as the sum of pressure values in the current frame is below an empirically determined threshold. From $P = \frac{F}{A}$, where F is the weight of the subject and A is the surface area covered by his or her body, we can see that pressure and area

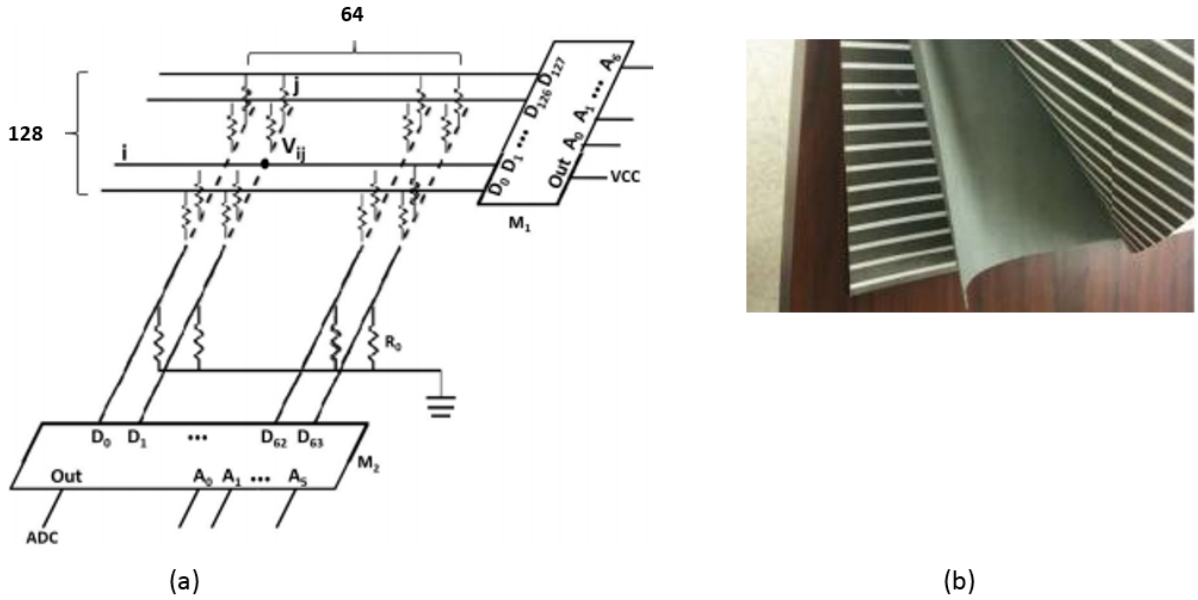

Figure 2.1: E-textile bed sheet system hardware: a) The circuit used to scan the pressure distribution of the bed sheet sensor array. b) Cross-section of the three-stacked-layer structure of the e-textile bed sheet. The e-textile piezoresistive fabric is sandwiched between two orthogonal conductive bus layers.

are directly proportional. Therefore, the threshold for the NO USER state was derived from the pressure values of the subject with the lightest weight (45 kg/99 lb). To account for an even bigger range of subjects the threshold was determined to be one third of the threshold of the lightest subject in our dataset.

INITIALIZATION: Identifying Breaths Once the sum of the pressure values goes over the preset pressure threshold described above, the system state immediately changes to the INITIALIZATION state. As the label of the state indicates, this state is where the system initializes itself to get ready to detect subjects' breaths. During initialization, the system waits 10 seconds for the incoming signal to settle because the movements of getting on/off the bed are not relevant to the breath identification process. After the signal settles, the system starts the breath detection algorithm. After successfully detecting 3 breaths within a time interval of 20 seconds, corresponding to a minimum breathing rate of 9 bpm, the system enters the RESPIRATION state. This requirement ensures that only human beings, rather than heavy objects, could trigger the transition to the RESPIRATION state.

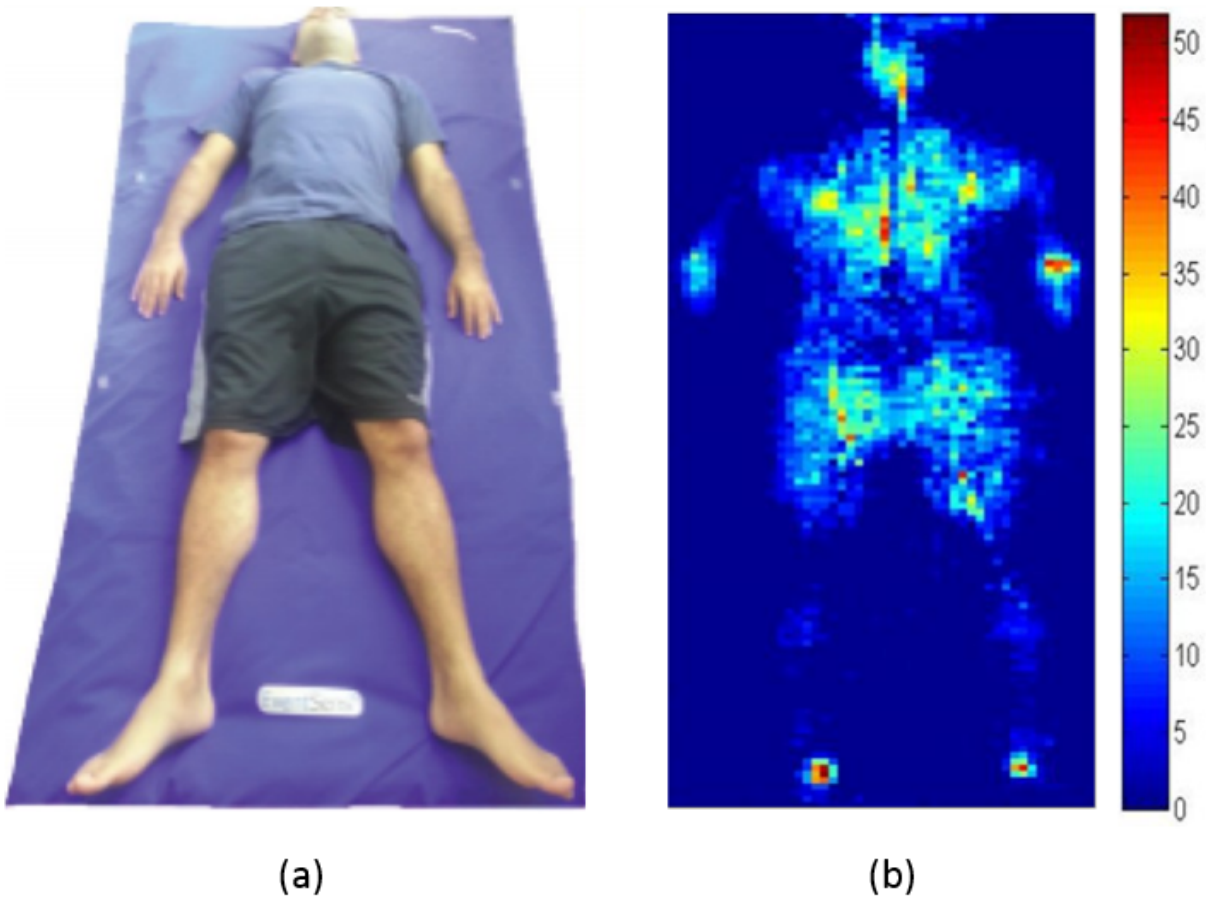

Figure 2.2: The bed sheet system as it is being used. a) A subject sleeping on the pressure-sensitive bed sheet in the supine position. b) The corresponding pressure map as displayed by the Android application on the tablet.

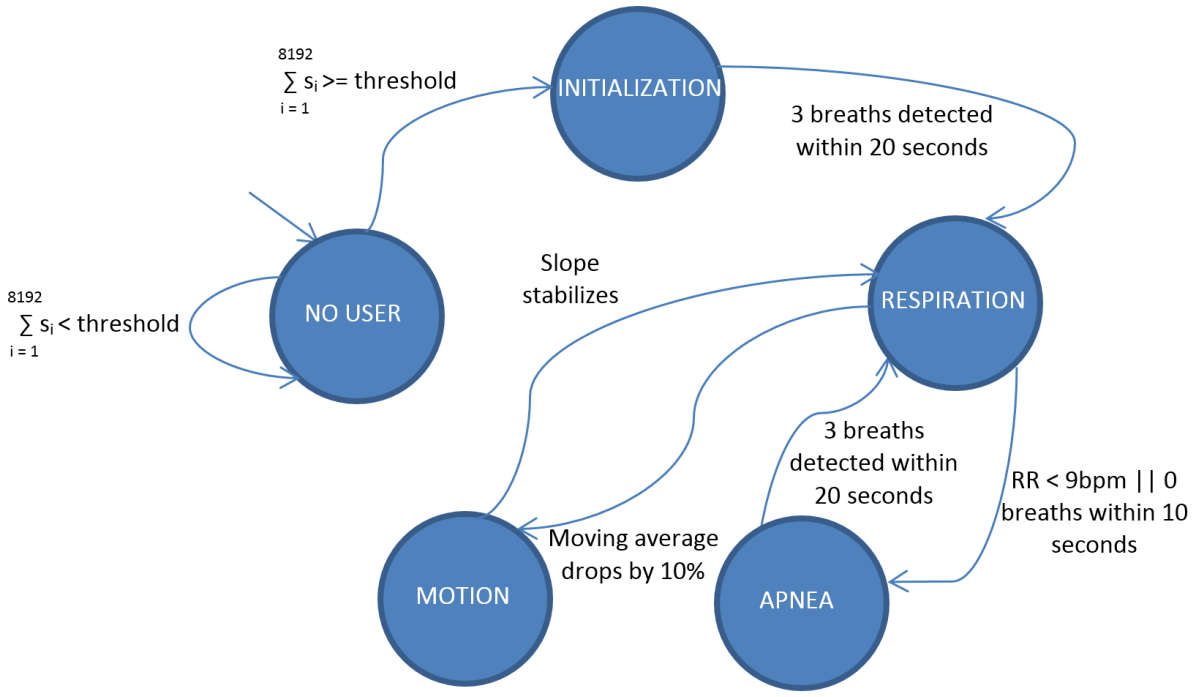

Figure 2.3: High-Level Software State Machine

RESPIRATION: Detecting Breaths The RESPIRATION state is entered once a stable representation of the subject's breaths is identified. This state could be entered either from the INITIALIZATION state or from the MOTION state. This is the state where the breath detection algorithm described in section 2.2.3 is activated. The algorithm calculates the instantaneous respiration rate of the user and determines if he/she is experiencing an apneic episode.

APNEA: Respiratory Cessation There are two ways to enter the APNEA state: either the user's current respiration rate falls below 9 bpm or no breaths are detected within 10 seconds. This period was chosen because sleep apnea is clinically defined as the complete cessation of air flow for at least 10 seconds [Fou11]. To leave the APNEA state and transition back to the RESPIRATION state, the system has to detect at least 3 breaths over a period of 20 seconds. This is again an approximation of a breathing rate of 9 bpm.

MOTION: Observing Users' Movements If a large amplitude of human movement occurs (e.g. due to a change in sleep posture during sleep), the state changes to MOTION. As in the INITIALIZATION state, the system waits until the pressure signal

settles, except the waiting time is not preset in the MOTION state. Instead, a sleep posture detection algorithm, which is part of the breath detection algorithm, is executed continuously to identify when a user gets in a stable sleep posture. Once a stable sleep posture is recognized, the state changes back to RESPIRATION.

2.2.3 Breath Detection Algorithm

The breath detection algorithm is partitioned into five stages, as shown in Figure 2.4. In stage I, pressure matrices from the bed sheet platform are collected and formatted. Both bed sheet pressure data and the ground truth information are collected and aligned with the time-indexed frames. In stage II, geometric feature-based sleep posture recognition is applied to the raw pressure matrix sequences in order to identify the sleeper's sleep posture. In this work, the sleep posture recognition algorithm is also used to identify movement. Sleep posture recognition is a machine learning process described in the paper proposed by Liu et al [LXA13]. Once the sleep posture is identified, the next stage (stage III) locates the chest area of the sleeper. The chest localization algorithm continuously tracks and calculates the location of the chest for each incoming frame. Hence, the pressure values around the chest area can be extracted to increase the accuracy of breath detection and respiratory rate identification.

Before extracting respiration-related features from the chest area, the time-drifting phenomenon of e-textile materials needs to be compensated for in stage IV of the breath detection algorithm. Sensor drift is an inherent property of piezoelectric pressure sensors that often causes a signal to grow logarithmically over time until the system reaches operating temperature [Mey08]. Such drift may cause inconsistent results, which is unacceptable in devices used for health monitoring where high accuracy is required at all times. After drift compensation, four respiration-related features are analyzed in time series: the maximum pressure sensor value, the sum of all pressure sensor values, the standard deviation of all pressure sensor values and the maximum singular value of the chest area. The respiratory signal can be observed directly by searching for peak values of these timeseries feature signals. To automate this process, a peak detection algorithm written in Matlab is applied in stage V to filter out noise and detect peaks.

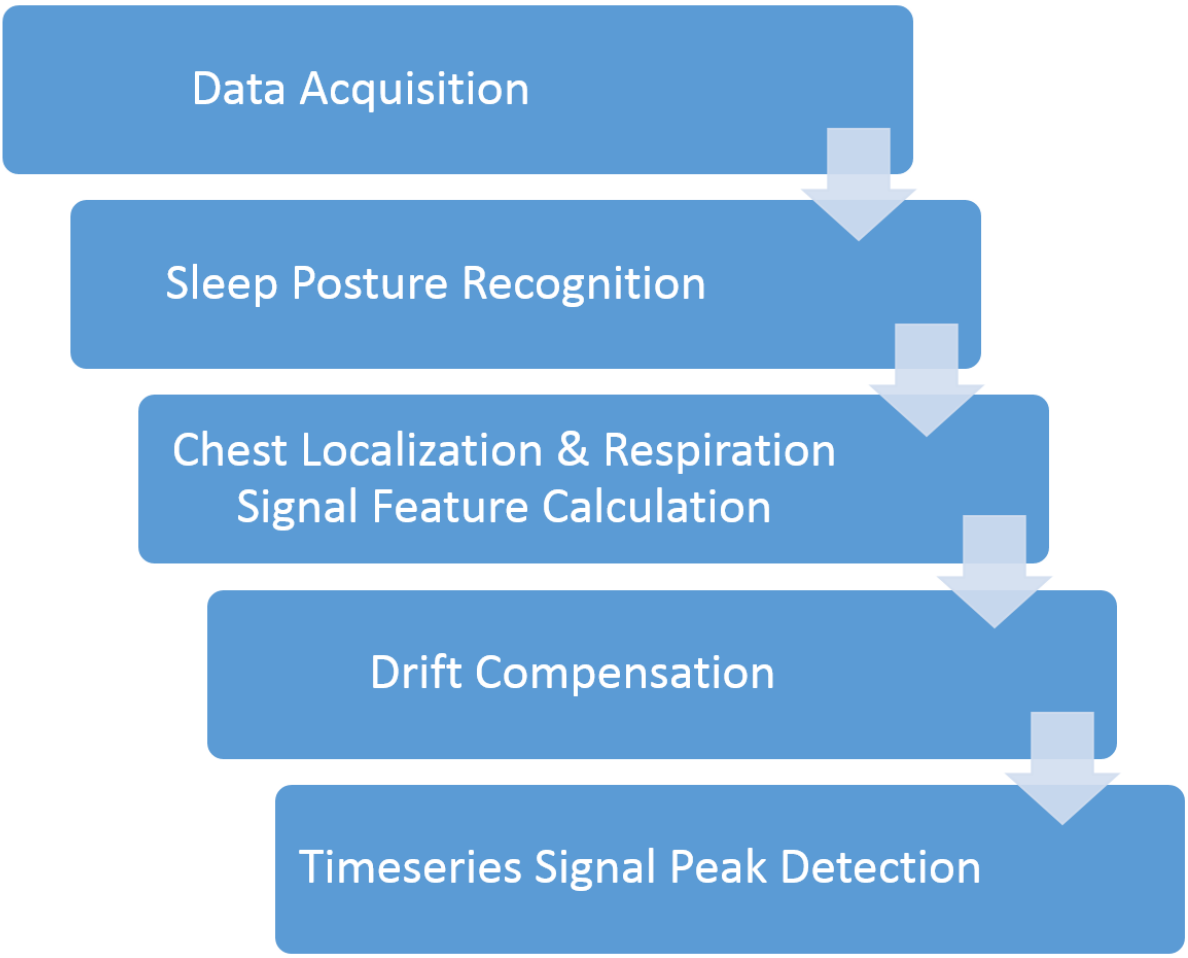

Figure 2.4: Algorithm flowchart for detecting respiration from the e-textile bed sheet.

2.2.3.1 Stage I: Data Acquisistion

The Android application records pressure data acquired from the pressure sensitive bed sheet over time into a log file. Any new frame (64 X 128 integer pressure matrix) read from the bed sheet is stored to that log file. Each new row of a frame starts at the beginning of a new line, and the individual values of a single row are separated by commas. An additional blank line is printed between each data frame. The result is a CSV file that could be easily imported into MATLAB or R as a two-dimensional matrix for post-analysis. Because the collected data is utilized to identify features that represent breaths of a typical subject, some type of ground truth was required to be synchronized with the current data acquisition framework. The ground truth mechanism needed to be able to identify every breath the user makes so that each one can be extracted from the acquired data during the post-analysis review. This was achieved through manual annotation by an observer as well was through a video recording of the sleeping participant. An annotation was made whenever the participant breathed, since this can be clearly seen through observation. To make the annotation process simple and synchronized with the recorded data, a button was added to the data collection application so that the observer can press it every time a breath is observed. The application then logs these respiration events by saving the frame numbers and time of the annotated breaths to an additional log file. This made it easy to correlate the observed breaths to the pressure data frames being logged. To synchronize the video recordings with the logged data, the current system frame number was included at the bottom of each recorded video frame. This way, the video could be reviewed in a post-analysis session to observe any abnormal data and synchronize it to the actual frame number it occurred at. This video recording could be used to identify any unexpected changes in pressure (e.g. during movements) as well as any possible breaths that the annotator may have missed. It is worth emphasizing that the manual annotation and video recordings are for testing and evaluation purposes only and are not needed by the proposed bed sheet system to measure respiratory rate.

2.2.3.2 Stage II: Sleep Posture Recognition

For sleep posture analysis, Liu et al. extract a set of geometric features from the pressure images obtained from the e-textile bed sheet and develop three heuristics based on sparse representation to classify sleep postures. First, the pressure images are preprocessed to remove noise; the images are aligned to a common center of mass and relocated to the center of the image, then a smoothing filter is applied and the images are normalized so that the sum of pixel weights is one. This step attempts to counteract the differences in patient body mass. Next, the algorithm calculates geometric features from the pressure images that reflect the physical characteristics of the body shape. The geometric features include spatial features and bodypart features. Spatial features are features that describe global aspects of the image such as the proportion of the image that is covered by the subject, how symmetric the image is, and direction of any curvature in the pressure image. Bodypart features are features that describe location and size of expected body parts such as the hip and shoulder. [LXA13] has a list of all the geometric features. The feature vectors for each of the samples form a sparse matrix. More formally, given a data set of n samples, with each sample having m dimensions (features), the data matrix $A \in R^{nxm}$ is comprised of these m-element column vectors arranged side by side. Given a new sample $y \in R^m$, a solution $x \in R^n$ needs to found such that x is described in terms of the data set, i.e. x needs to satisfy:

$$y = Ax \tag{2.1}$$

So y is a linear combination of the columns in the data set, and $x = [x1, x2, ..., xn]^T$ is an unknown vector of coefficients. Because there are more unknowns than equations, this linear system in underdetermind and there are infinitely many solutions to x. The paper imposes some constraints so that a unique solution to x can be found. In the last step of the posture recognition algorithm, a heuristic is applied to select the class label for a new sample y given a sparse solution for x. The classification has 83.5% precision, 82.9% recall, and 83.2% accuracy.

2.2.3.3 Stage III: Chest Localization & Feature Calculation

Once the sleep posture is identified, the location of the subject's chest area can be derived to increase the accuracy of breath detection. This is done by first calculating the subject's center of mass. The center of mass is the point at which all the mass can be considered to be "concentrated." The center of mass of an object can be calculated by summing up the partial masses at each position of the object and dividing by the total mass of the object. Since the only information that can be obtained from the bed sheet is the pressure at each pixel, we use the pressure values instead of mass. This is safe to do because pressure is proportional to mass:

$$P = \frac{F}{A} = \frac{mg}{A} \tag{2.2}$$

As seen in equation (2.2), pressure is defined as force per unit area. The force applied to the bed sheet is the gravitational force, with a magnitude equal to the mass m times the gravitational acceleration g. This force is applied to each sensor/pixel of the bed sheet. Since the area of each of the 8192 pressure sensors and the gravitational acceleration are constant, it can be seen that the pressure at each sensor is just a scaled mass.

To calculate the location (x, y coordinates) of the center of mass for the sleeping subject, we use the following formulas:

$$com_x = \frac{\sum_{i=0}^{127}\sum_{j=0}^{63} x_{ij}p_{ij}}{p_{total}} \tag{2.3}$$

$$com_y = \frac{\sum_{i=0}^{127}\sum_{j=0}^{63} y_{ij}p_{ij}}{p_{total}} \tag{2.4}$$

, where x_{ij}, y_{ij} and p_{ij} are the x and y coordinates and the pressure value of the sensor at row i and column j, respectively.

The location of the center of mass is com = (com_x, com_y). This is of course an approximation of a person's true center of mass because a subject's clothing will contribute to the mass and will cause a shift in the center of mass, but it is nevertheless a very close approximation.

In the supine and prone positions, the center of mass lies around the umbilicus, both in the x and y directions (Figure 2.5 a)). In the left and right fetal positions, the center of mass is slightly higher (above the umbilicus) and more shifted towards the left or right, respectively (Figure 2.5 b)).

After locating the center of mass, we make the assumption that the shoulder is located in the top-most quarter of the pressure image. Since the chest area we are trying to locate is below the shoulders and above the center of mass, we use the connected component labeling algorithm to extract the part of the body between the shoulders and the center of mass. The connected component analysis is meant to extract the pixels of the image that have non-zero pressure values and are connected to each other (none of the pixels is separated from another by a zero-valued pixel). The restriction for the non-zero pixels to be connected attempts to prevent parts of the arms to be included in the localized chest area. If the position reported by the posture recognition algorithm is the prone or supine position, then the connected component analysis is restricted to the portion of the image bounded by the y position of the center of mass from below and the y position of the beginning of the second quarter of the pressure image from above. If the subject is in one of the fetal positions, the connected component labeling algorithm is additionally bounded by the center of the pressure image along the y-axis. This additional restriction in the fetal positions reduces the chances of capturing the back of the subject as part of the chest. Algorithm 1 shows the pseudocode for the chest localization. In the pseudocode, the "bounding box" refers to the restricted area that can be labeled for the prone, supine and side positions. Figure 2.6 shows an example of a pressure image where the chest has been localized using the above algorithm.

After the chest area is localized, respiration-related features can be derived from it. We extract four different respiration-related features and analyze them in timeseries. The four feature are 1) the maximum pressure value of the chest area (max), 2) the sum of all pressure values in the chest area (sum), 3) the standard deviation of all pressure sensor values in the chest area (std), and 4) the maximum singular value of the chest area (svd). Each of the 4 features are computed for each pressure image and plotted over time. An example sum-of-pressures signal is shown in Figure 2.7.

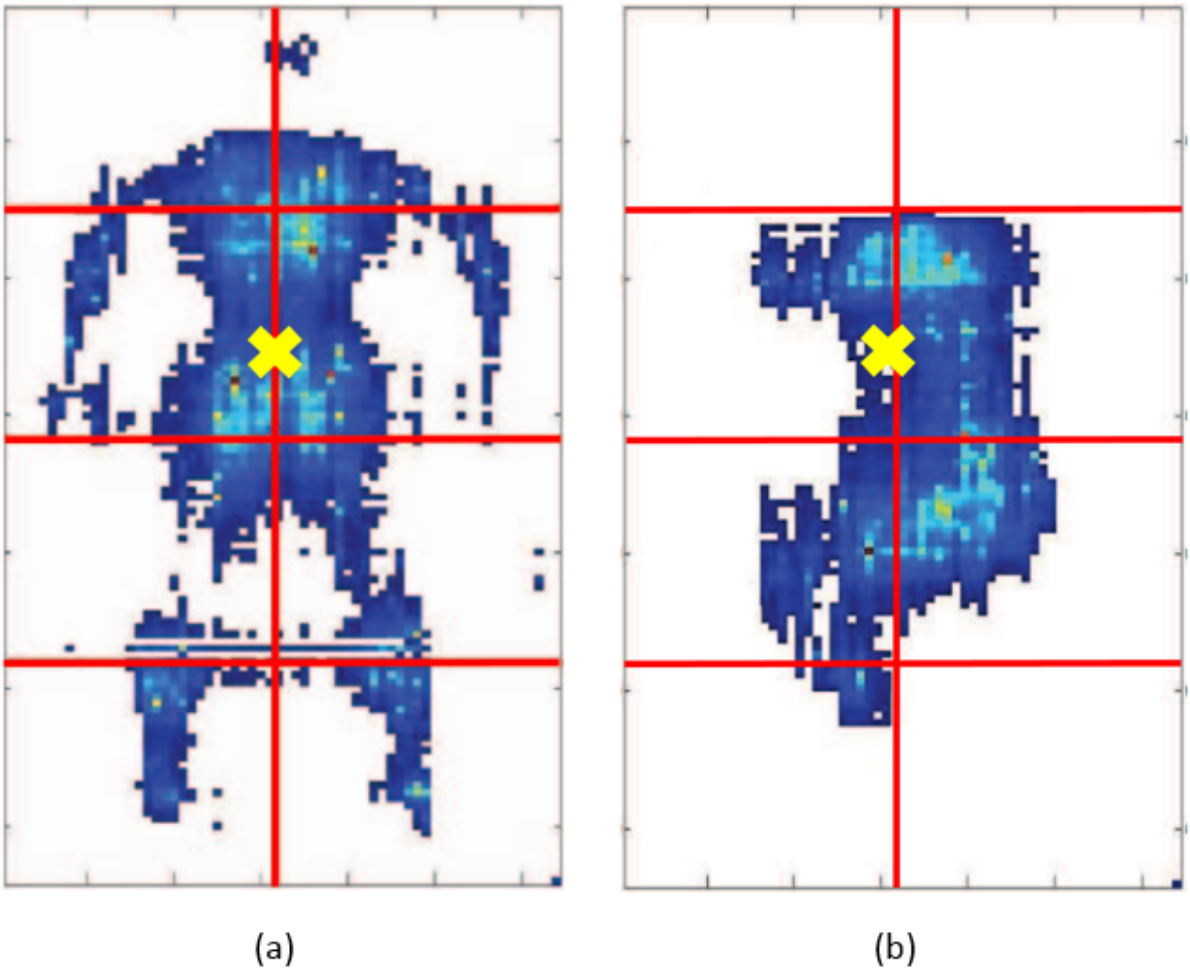

Figure 2.5: The location of the center of mass in different sleeping postures. a) A pressure map of a subject sleeping in the supine position. The center of mass is located at the umbilicus as indicated by the yellow cross. b) A pressure map of a subject sleeping in the right fetal position. The center of mass is located slightly above and to the right of the umbilicus as indicated by the yellow cross.

Algorithm 1 Chest localization algorithm

```
stack[]
chestPixels[]
currPixel = pressureFrame[com_x][com_y]
chestPixels.add(currPixel)
stack.push(currPixel)

while stack is not empty do
    currPixel ← stack.pop()
    if currPixel is within bounding box then
        N = currPixel's neighbors
        for neighbor in N do
            if neighbor ≠ 0 && neighbor ∉ chestPixels then
                chestPixels.add(neighbor)
                stack.push(neighbor)
            end if
        end for
    end if
end while
```

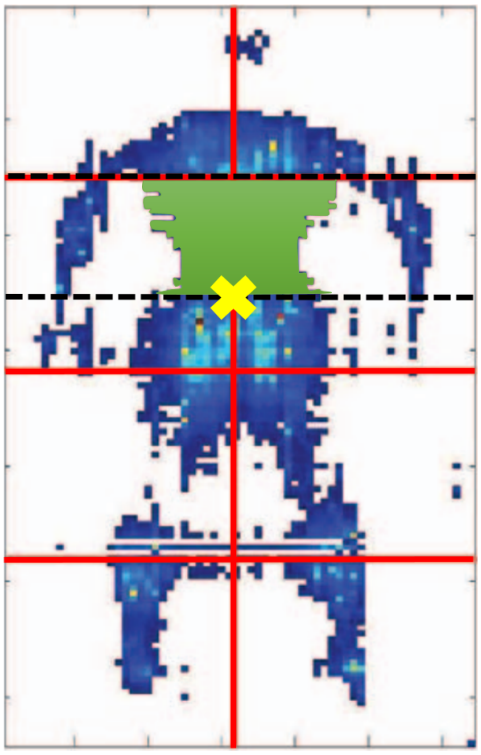

Figure 2.6: A pressure image of a subject sleeping in the supine position. The localized chest area is shown in green. The dashed black lines are the boundaries of the bounding box used in the chest localization algorithm.

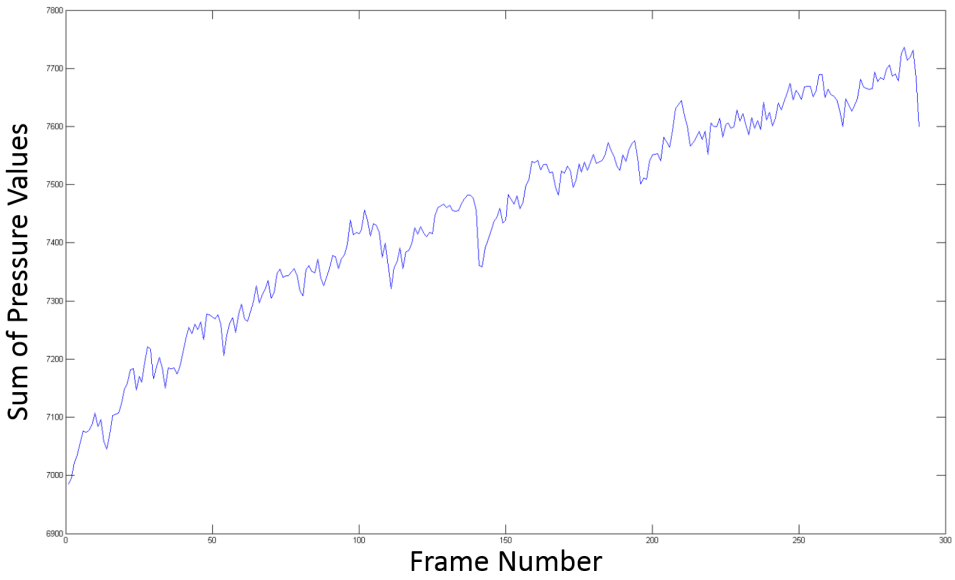

Figure 2.7: An example sum-of-pressures signal: the sum of all pressure values in a frame are plotted over time.

2.2.3.4 Stage IV: Drift Compensation

Piezoelectric pressure sensors inherently suffer from the time drifting phenomenon. The effect of drift is that when repetitive cycles of compressing and relaxing of the piezoelectric fabric are performed, the pressure is different between the cycles at the same point of compression. It has been shown that the pressure will continue to grow logarithmically over time until the system reaches operating temperature. Figure 2.8 shows an example sum-of-pressure signal and its best-fitting curve which corresponds to a logarithmic function. The drift phenomenon specifically affects our motion detection which is a simple threshold-based algorithm that uses a moving average of the sum of pressures over the last 5 frames. This is used to represent the average pressure the user exerts on the bed sheet for his/her current position. When the user moves (e.g. to change sleeping posture), he/she will lift some parts of the body off the bed sheet, causing a significant drop in pressure and a noticeable drop in the moving average. The logarithmic drift, however, will cause the moving average to always increase, making it harder to clearly identify deviations from the average. To compensate for the e-textile drift, the first-order derivative of the signal is calculated. The first-order derivative curve has the same number of peaks and troughs and succeeds in projecting the signal into a horizontal plane. Figure 2.9 shows the sum-of-pressure signal before and after drift compensation.

Even though the first-order derivative curve has different y values than the original curve, it preserves the shape of the curve (peaks and troughs) and eliminates the logarithmic growth of the signal. For the thresholding algorithm used to detect motion, it is not important what the absolute amplitudes of the signal peaks are, but rather that there are no big deviations among the different peak amplitudes. This way, if any big deviation occurs, it can be classified as motion.

2.2.3.5 Stage V: Timeseries Signal Peak Detection

A single oscillation in any of the extracted signals represents a single breath made by the user. This is the case because a breath involves an inhalation period followed by an exhalation period, which corresponds to increases and decreases in the pressure exerted on the bed sheet. The respiratory signal can therefore be observed directly by searching

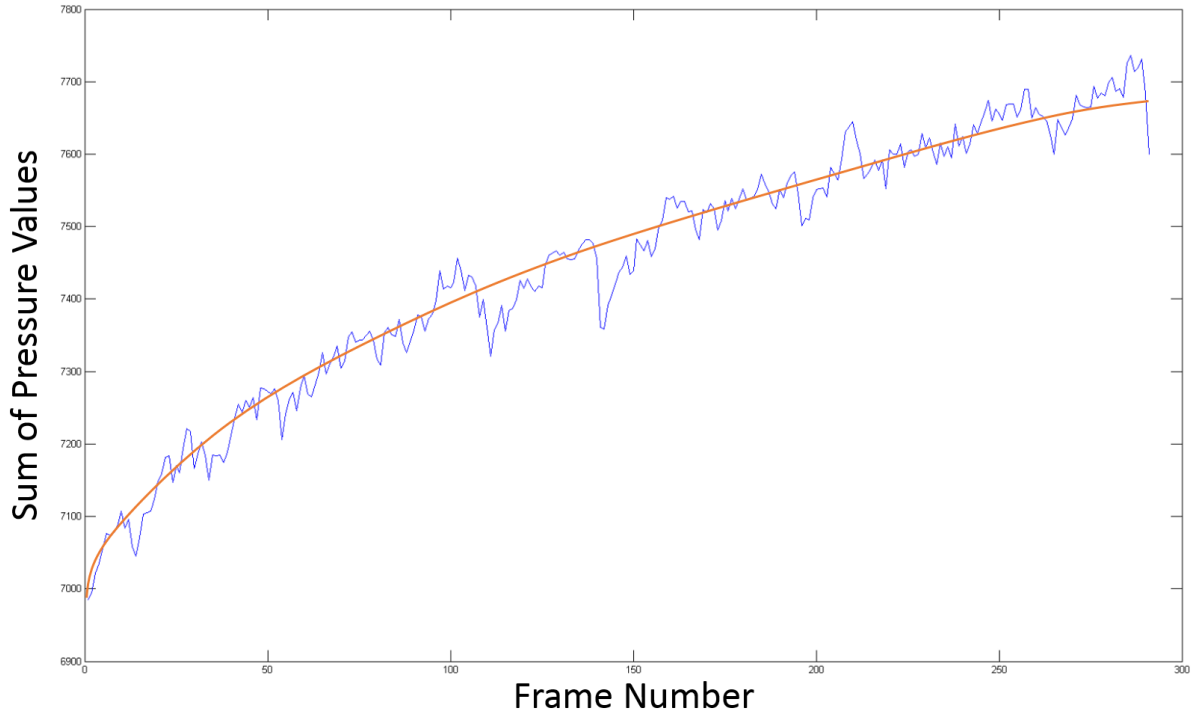

Figure 2.8: The sum of pressures signal with a best-fit line that models the logarithmic drift of the pressure sensors.

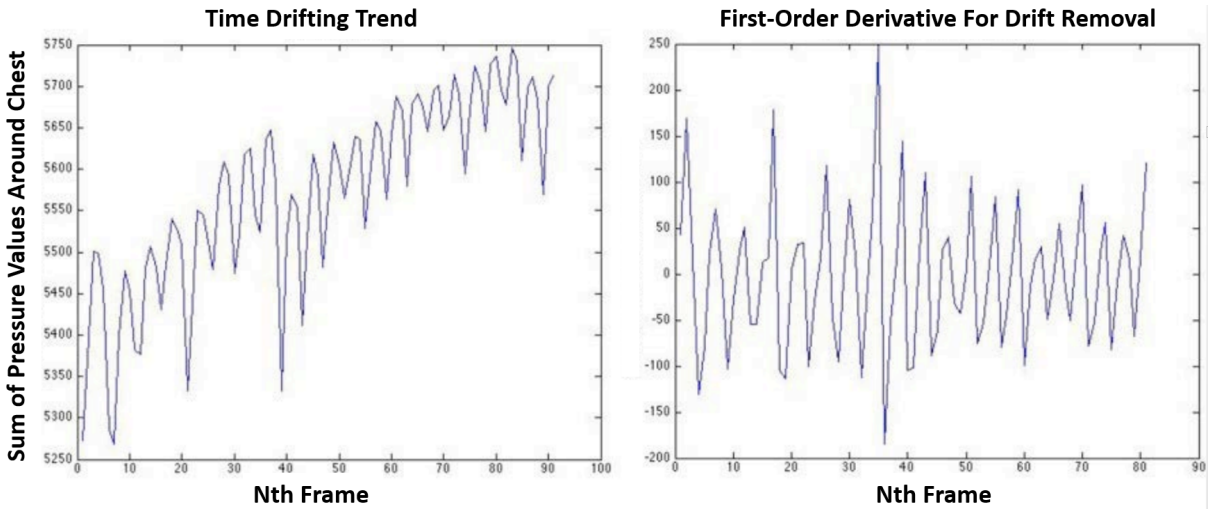

Figure 2.9: The sum of pressures signal before and after drift compensation via first-order derivation.

for the peaks in the timeseries feature curves (max, sum, std, and svd). Since every oscillation corresponds to a breath, identifying these breaths will allow us to calculate an instantaneous respiration rate and determine if the sleeper is experiencing an apneic episode. A peak detection algorithm was implemented to capture the occurrence of each breath. A data structure was created to hold the occurrences of breaths over a 30-second wide moving window. With a sampling rate of 1.5Hz, 30 seconds corresponds to 45 frames. The instantaneous respiration rate is calculated by multiplying the number of detected breaths within the last 30-second period by 2 to obtain the number of breaths per minute. An apneic episode is detected when the instantaneous respiration rate falls below 9 bpm or when no breaths are detected within 10 seconds.

2.2.4 Results & Discussion

Data was collected from 14 subjects (males and females), ranging from 20 to 30 years of age. The participants were asked to sleep on the bed sheet in three different positions: prone(face down), supine (face up), and on their sides. A total of at least 6 minutes of data was collected from each user (2 minutes per position). While the participants were simulating sleeping, the system data acquisition framework was utilized to record their pressure data over time and to create an annotated log of their breathing events.

The first experiment attempted to test if chest localization helped in respiratory identification. 14 subjects were asked to statically lie on the bed in supine position without any posture changes. In order to assess whether chest localization helped in getting clean respiratory signals, two pressure value matrices were derived from the raw pressure value matrix sequence. One was a sub matrix which only related to each subject's chest area and the other one was the original pressure value matrix. The four respiration-related features described earlier (sum, max, std, and svd) were utilized to visualize waveforms from these two matrices. Figure 2.10 shows the sampled curves of all four features with and without chest localization. As can be seen in Figure 2.10 b), the respiration signal can be inferred from the chest area even via just visual identification. This is because the movements around the chest area are directly related to the respiration signal, while movements from other body parts simply disturb the respiration-related features, as can be seen in Figure 2.10 a).

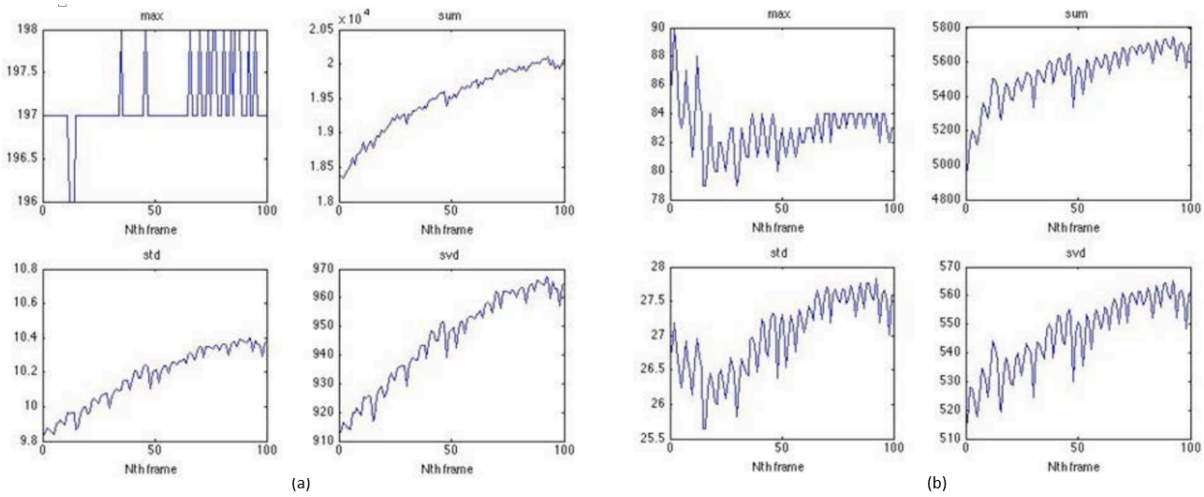

Figure 2.10: Sampled curves of all four respiration-related features with and without chest localization: a) The signals corresponding to the four respiration-related features when extracted from the full bed sheet. b) The signals corresponding to the four respiration-related features when extracted from the localized chest area.

In the second experiment, we evaluate the effectiveness of first-order derivation to remove the logarithmic drift of the respiration-related features and to improve the accuracy of the breath detection. After drift compensation, the number of breaths missed by the peak detection algorithm dropped significantly. Figure 2.11 shows the four respiratory signals of one of the subjects before and after drift compensation. We observe that, after drift compensation, the oscillations become more pronounced which greatly facilitates breath detection.

In the third experiment, and only using the localized pressure signals, we attempt to demonstrate that the proposed bed sheet system system can identify the respiratory signal in the supine, prone, and side positions. This is essential since people constantly change their sleeping posture during the night and the respiratory monitoring needs to be continuous. In addition, the test subject could sleep on any corner of the bed sheet without any restriction. In the posture detection step , the pressure distribution maps are aligned to a common center of mass and shifted to the center of the pressure image. Although the respiratory signals were always captured, different postures indeed generated different pressure distributions on the bed sheet when the subject breathed. Respiration-related features could still be visually compared with the ground truth in the prone and supine postures. However, since, in the side postures, the contact area

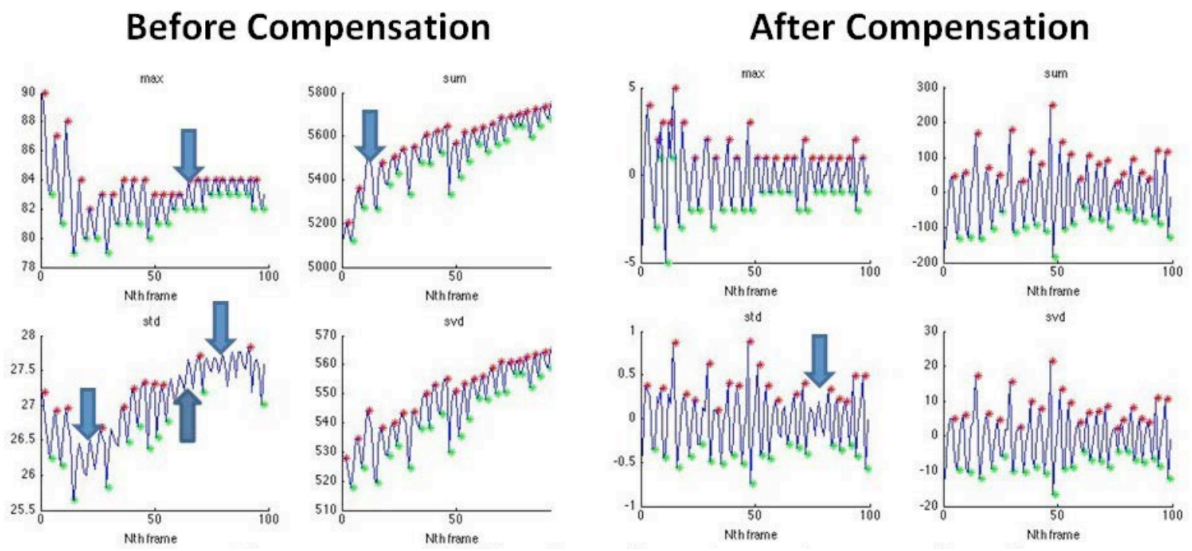

Figure 2.11: The chest-localized signals corresponding to the four respiration-related features before and after drift compensation. The blue arrows mark the missed breaths.

between the human body and the bed sheet is the arm area or the side of the body rather than the chest, side postures tended to obscure the breath signal. Therefore, even if breathing made the chest area expand and shrink, the pressure on the bed sheet only changed slightly. Figure 2.12 provides accuracy analysis among these sleep postures for each of the four features and Figure 2.13 shows the overall average accuracy for the fours features compared to ground truth.

2.2.5 Conclusion

In this work, we presented a non-intrusive nocturnal framework for respiration rate monitoring and OSA detection. An e-textile based bed sheet was used to collect high density pressure images. The bed sheet can be seamlessly integrated in a home or clinical environment, thereby eliminating any interferences with the subject's sleep. In addition, the bed sheet provides high-density pressure images and can capture the full body pressure distribution. This is particularly important for localizing specific body parts, like the chest area where the respiratory signal becomes more prominent. Body part localization is only possible using a high-density array-based bed sheet, rather than sparse pressure sensors on the bed as was done in some of the discussed related works. The system continuously tracks the sleeper's chest area based on his/her posture and extracts

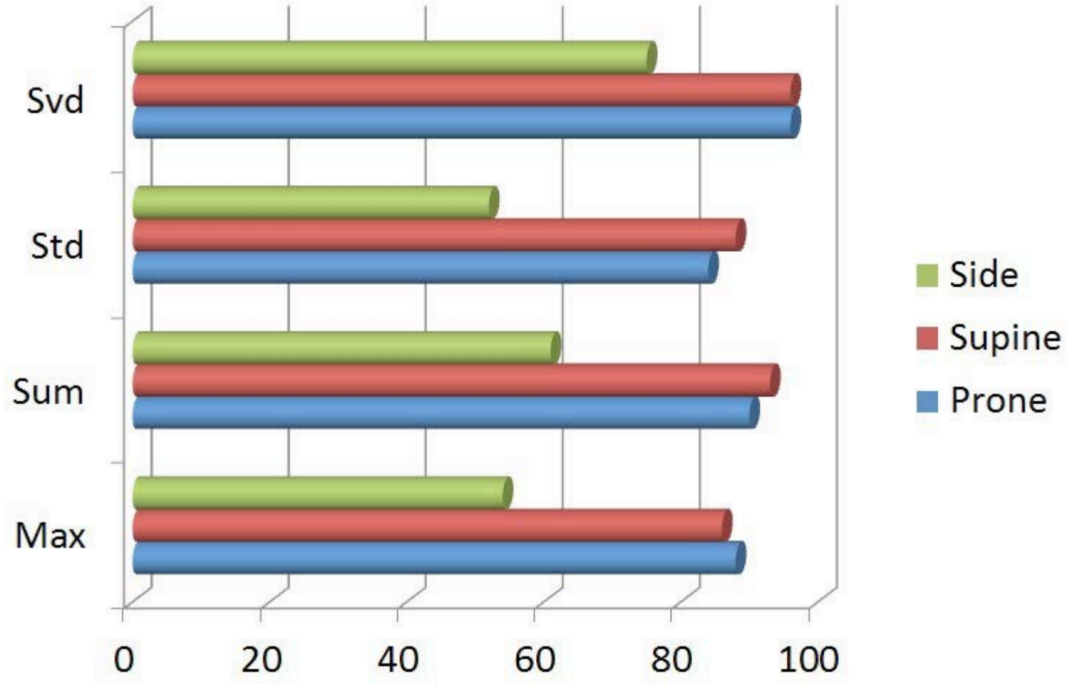

Figure 2.12: Accuracy of breath detection for the four respiratory related features in the 3 different sleeping postures.

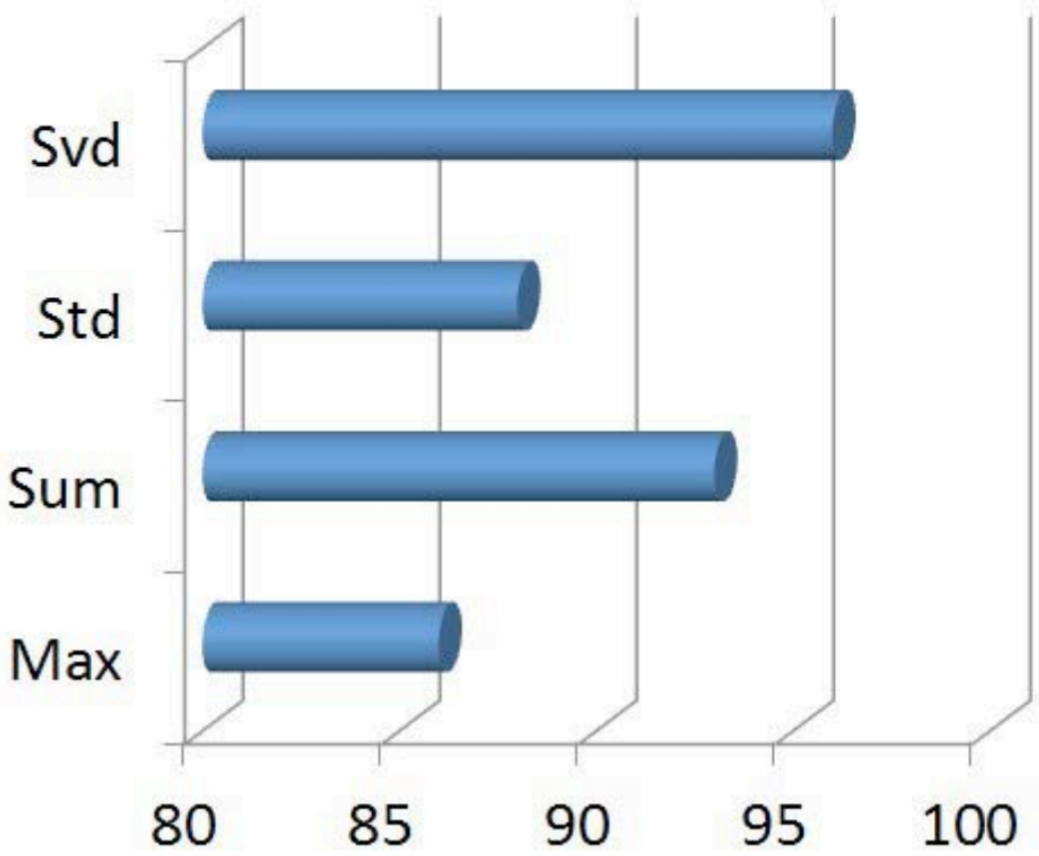

Figure 2.13: Average accuracy of breath detection compared to ground truth.

a clean, localized pressure signal. After drift compensation, a peak detection algorithm is applied to extract individual breaths within a 30-second period and calculate the instantaneous respiration rate of the sleeping subject. Using SVD as the feature gives an overall accuracy of 96%.

2.3 Unobtrusive Sleep Stage Identification Using a Pressure-Sensitive Bed Sheet

Since the e-textile bed sheet system proved to be effective at capturing the respiratory signal, our goal was to try to extract other signals from the bed sheet. The more biophysical features we can extract from the bed sheet, the better the decisions we can make about the presence of OSA. In this work, we focus on sleep stage analysis. As described in Chapter 1, when an OSA patient experiences an apneic episode and the blood oxygen levels start to drop, the brain is signaled to wake the sleeper up to start breathing again. This does not result in a full awakening, but rather a transition from the deeper, more restorative stages of sleep to lighter, shallower stages. These arousals and the accompanying sleep stage transitions occur frequently throughout the night and result in a fragmented, non-restorative sleep pattern. In a PSG, the many signals being monitored allow the sleep technician to interpret and analyze the subject's sleep and draw conclusions about his/her sleep cycle patterns and how they affect the subject's sleep quality. In this work, we extract other signals from the e-textile bed sheet, in addition to the respiratory signal, and attempt to draw conclusions about the sleep stage pattern of the subject being monitored. Identifying the sleeper's sleep stages can help with diagnosing OSA as well as with understanding more about how OSA affects a patient's sleep and, thereby, his/her quality of life.

2.3.1 Sleep Cycle Introduction

Sleep can be classified into two broad categories: Non-rapid eye movement (NREM) and rapid eye movement (REM) sleep. The American Academy of Sleep Medicine (AASM) further divides NREM sleep into three distinct stages: N1, N2, and N3 [SAB07]. As shown in Figure 2.14, a healthy adult's sleep cycle, which lasts between 90 and 100

minutes, begins with 3 stages of NREM sleep followed by REM sleep. The first stage (N1) is the lightest and shortest stage of sleep (1-7 minutes) and marks the transition from wakefulness (W) to sleep. N1 is followed by N2, which lasts anywhere from 10 to 25 minutes. This stage is where the body reaches a state of complete relaxation in preparation for the deeper sleep to come. After N2, a healthy adult enters N3, the last stage of NREM which is also referred to as deep sleep. N3 lasts 20 to 40 minutes and is the stage where the body does most of its repair and regeneration work. Following the N3 stage of sleep, a healthy adult ascends to lighter NREM sleep stages, typically N2, for 5 to 10 minutes before entering the REM sleep episode. REM sleep is characterized by high brain activity and is where memory consolidation predominantly occurs. REM sleep comprises about 20 to 25% of total sleep in typical healthy adults [Har], [Ame07]. In healthy subjects, this pattern continues to repeat in a cyclical fashion throughout the night. Since each stage fulfills a vital biological function, sleep stage analysis is crucial to the evaluation of the quality of sleep and is a proven biometric in diagnosing sleep apnea, cardiovascular disease, diabetes, and depression [Tha06].

In sleep medicine, in order to evaluate sleep quality, a person gets a PSG. The standard practice is then to divide the sleep time into 30-second epochs, and based on the recorded signals, each epoch can be scored by sleep technicians as W, N1, N2, N3 or REM. Since the pattern of the Wake-NREM-REM stages is highly modified for subjects with sleep disorders and other diseases, sleep stage analysis can provide valuable information for diagnosing these diseases. In this work, we again use the bed sheet system described in Section 2.2.2.1 to do sleep stage analysis. Our contributions in this work can be summarized as follows:

- To the best of our knowledge, our work is the first to perform sleep stage analysis using a completely contact-free non-intrusive system.
- Our sleep stage analysis results are validated against over 50 hours of gold standard polysomnography data–the current state-of-the-art in sleep analysis–and our results are in the range of inter-rater agreements reported in the literature (70% and 72% [DAZ09]).

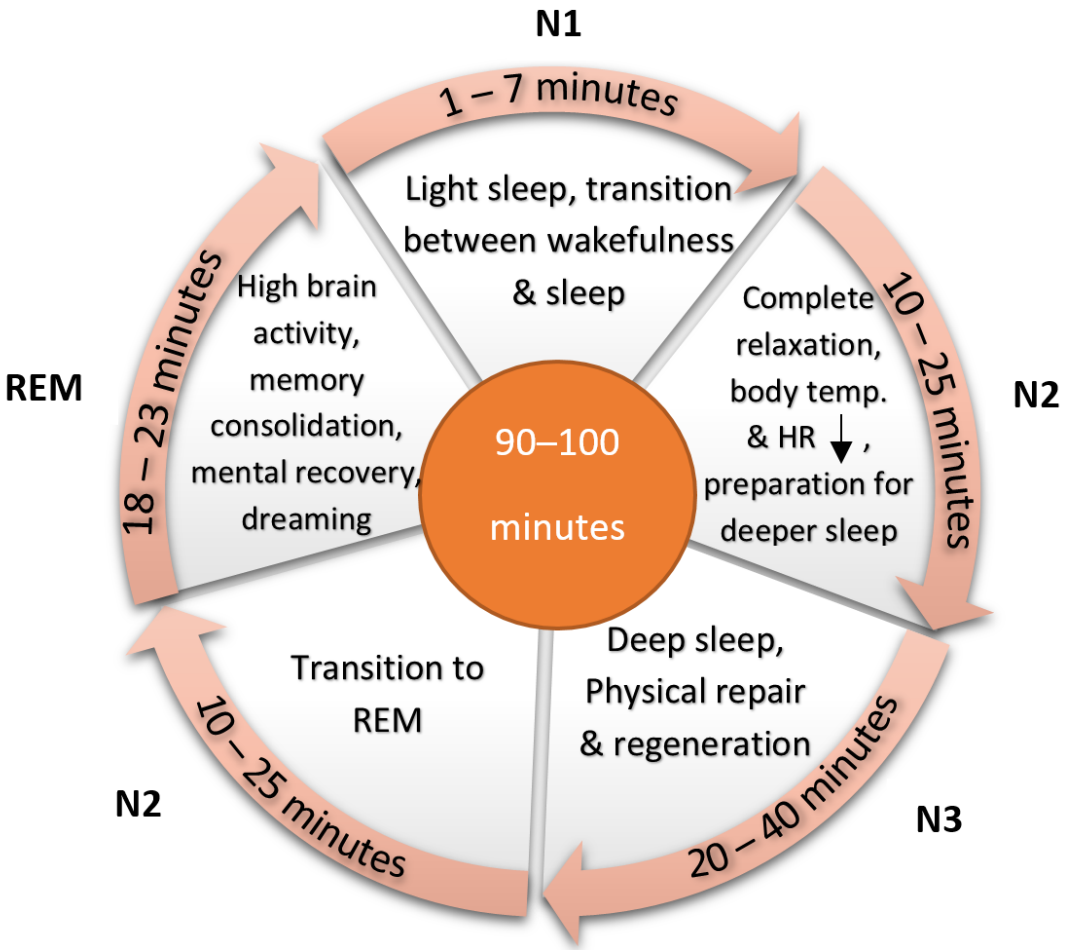

Figure 2.14: A healthy adults sleep cycle: The cycle starts at the NREM substage N1 and ends with REM.

2.3.2 Related Works

While there are many academic and commercial sleep monitoring tools available, there is a dearth of low-cost, non-intrusive solutions. This section surveys the major categories of solutions and describes their characteristics and limitations.

Existing sleep monitoring tools can be classified into three categories. The first category of tools extracts sleep stages from directly-measured physiological signals. Because many body functions like breathing, heart rate and movement change during sleep, tracking these changes throughout the night can provide a good indication of the sleep stage that a patient is in. Many of these tools distinguish themselves from full PSG by using only a small subset of the PSG sensors. [RH06] shows that ECG and respiratory effort alone can be sufficient to distinguish between the WAKE, NREM and REM stages with moderate accuracy (67%). [EN10] uses features extracted from the EEG signal in addition to heart rate variability to detect the different sleep stages. [FLK12], on the other hand, uses only the EEG signal to study brain activity and determine sleep stages based on it. An accuracy of 83% is reported. Another example is Zeo, a commercial sleep monitoring product that detects sleep stages based on brain activity. A headband, to be worn during sleep, analyzes brain signals and identifies sleep stages based on the signals' frequencies. The company went out of business as of 2013. Although these systems use fewer sensors than a full PSG, they still require equipment to be attached to the patient's body during sleep and are often expensive.

The second category of tools tries to infer sleep stages from body movement. Actigraphy is a commonly used technique for sleep monitoring; it uses an accelerometer embedded in a watch-like device to monitor activity and identify sleep stages [SBP03], [SA02], [HPP04]. Sleeptracker [Tra], Actiwatch [Phi], and UP [Jaw] are a few examples of the many commercial products available in this category. Though less invasive than the products in the first category, these products still require a device to be attached to the patient's wrist. To overcome this limitation, [HDS13] presents a sleep monitoring system based on RFID technology. WISP tags are added along the edge of the bed mattress and accelerometer data is collected from these tags by an RFID reader. This system does not require any attachments to the patient's body but might represent a problem if it were

to be deployed in a clinical setting since the RFID frequency can interfere with other medical devices [YL10].

The third category of products uses audio and video signals to identify sleep stages. [PLS06] uses audio and video sensors to infer sleep-wake stages. These systems are expensive and raise obvious privacy concerns. In [BGK12], the authors attempt to solve the privacy problem by using an infrared camera. They analyze the temperature maps acquired from the camera to detect the patient's body and to extract body movement information. Infrared cameras are a controversial technology, however, as the high-resolution thermal images make it very simple to observe exactly what the patient is doing in bed and can be considered a breach of privacy. Our proposed system opens the way for a fourth category of sleep monitoring that is completely contact-free, non-intrusive, comfortable, cheap, and avoids privacy violations.

2.3.3 System Architecture

The bed sheet used in this work is the same as the one described in Section 2.2.2.1, except that the sampling rate was adjusted to be 1 Hz instead of 1.5 Hz in order to reduce the amount of collected data. In addition to sleeping on the e-textile bed sheet, the test subjects also had to get a polysomnography. The signals acquired by the PSG were used as the reference gold standard for assessing the accuracy of our sleep identification algorithm. The PSG system used in this work is the "SOMNOscreen plus" system manufactured by SOMNOmedics [SOM]. In this work, all the sensors used and signals collected are as described in Section 1.2.4, except for video. The data collection on the bed sheet system and the PSG system started simultaneously and synchronized.

2.3.4 Sleep Stage Identification Algorithm

In this section, we describe the features used for classification, how they are extracted from the bed sheet pressure images and why they are suitable for sleep stage identification. The pressure images from the entire night are divided into 30-second epochs, as is the convention in sleep medicine, and the features are extracted for each epoch. Following feature extraction, we describe the two-phase procedure used for identifying the Wake,

NREM and REM stages.

2.3.4.1 Feature Extraction

As discussed earlier, several physiological signals change during normal sleep and vary with each sleep stage. Three of those signals are respiratory effort, leg movement and body movement. The following subsections describe each one in more detail and explain how they were extracted from the pressure images.

Respiration Rate Respiration rate, as measured by the number of breaths observed per minute, is another biophysical signal that changes during sleep. Respiration rate is considerably faster during the REM and Wake sleep stages than in the NREM stage. So, one respiratory feature used for sleep stage identification is the respiration rate observed during a 30-second epoch. The breathing rate is extracted from the respiratory signal by counting the number of breaths and multiplying by two to get the number of breaths per minute. This process is described in Section 2.2.

Respiration Rate Variability Respiration is a physiological signal that undergoes significant changes during sleep. During the NREM sleep stage, the breathing pattern is regular, both in amplitude and frequency [HSK87], while in REM the signal becomes irregular, much more rapid, and sudden changes in both amplitude and frequency can be observed [KPR03]. To capture this behavior of the respiratory signal, we first extract the respiration signal from the bed sheet.

In this work, we evaluate the accuracy of the extracted respiration signal by using the thorax effort signal obtained from the PSG system as ground truth. The thorax signal is generated by measuring the amount of pressure applied to the effort belt that the user wears around the chest. The thorax effort signal is shown in Figure 2.15. It is noteworthy to mention that, in contrast to the pressure signal obtained from the bed sheet, this signal indicates an inhalation event at each local maximum when the subject is in the supine position. The signals are inverted because, in the supine position, when a subject inhales and his/her chest rises, the pressure exerted on the bed sheet decreases while the pressure exerted on the effort belt increases. Another difference between the

two signals under comparison is the sampling rate. While the bed sheet pressure image is sampled at 1 pressure image per second, the thorax effort signal has a sampling rate of 32 Hz. While the bed sheet signal in Figure 2.16 shows 30 samples on the x-axis, Figure 9 shows 960 samples for the same period of time. As shown in the figure, the thorax effort signal indicates a total of 8 inhalations and 8 exhalations (local maxima and minima, respectively), for a total of 8 breaths. During the same time period, the pressure signal obtained from the bed sheet indicates the same number of inhalations and exhalations. Both systems would therefore result in a respiration rate of 8 breaths per 30 seconds, or 16 breaths per minute.

After extracting the respiration signal, we compute the mean amplitude and the mean frequency of that signal in each 30-second epoch. Then, for each portion of the signal in the 30-second epoch, we compute its variability from the mean in terms of both amplitude and frequency. This is done as follows. Starting from the first sample in the signal, we compute its standard deviation from the mean amplitude and the mean frequency of that epoch. We incrementally grow the size of the signal by adding one sample at a time and computing the amplitude standard deviation and the frequency standard deviation until all 30 samples of an epoch are covered. We then move to the next epoch and repeat the process. That way, we get a measure of the dispersion of the respiration signal from the mean amplitude and the mean frequency of an epoch at each pressure frame. The variability of the respiration signal is a very good feature to distinguish between the NREM stage where the signal is regular and shows very little variability, and the Wake-REM stages where the amplitude and frequency of the respiration signal vary significantly and rapidly.

Leg Movement Leg movement is also an important indicator of sleep stage. During the first stage of NREM, sleepers can experience sudden jerks of their legs. These jerks are common while falling asleep but, if excessive, can also be a symptom of sleep disorders like Periodic Limb Movements (PLM) and Restless Legs Syndrome (RLS) [Nat10]. In either case, these jerks are associated with the NREM sleep stage and can be used as a feature to distinguish the NREM sleep stage from the other stages.

To extract leg movement from the bed sheet, we assume that a subject's legs while

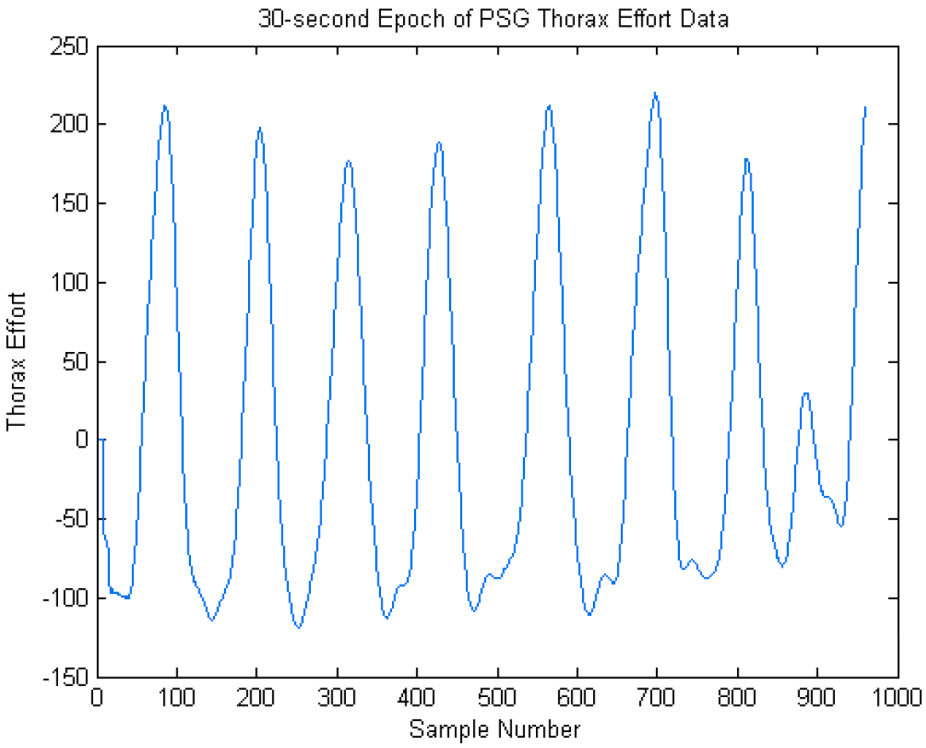

Figure 2.15: The thorax effort signal obtained from the PSG system. Each local maximum (peak) indicates an expansion of the chest which corresponds to an inhalation event.

lying on the bed sheet will occupy the lower half of the sheet. To extract leg movement, we sum all the pixels in the lower half of the pressure image and mark a leg movement when a significant drop or increase in pressure is detected in that portion of the bed sheet. This is a crude approximation of leg movement and can be improved in the future by localizing the patient's legs in the pressure images.

The ground truth for leg movement events is obtained from the PLM channel of the PSG recording, as shown in Figure 2.17. This signal measures leg muscle movement over time. The signal is obtained by sampling the leg EMG electrodes at 256 Hz. When compared to the PLM signal obtained from the PSG system, our approximation accurately detects 80.7% of all leg movements.

Body Movement In addition to leg movement, whole-body movement can also occur during sleep. The movements are associated with the Wake stage and light stages of NREM as a result of changes in sleep posture that occur every 5-10 minutes. REM sleep, on the other hand, is characterized by muscle immobility and body paralysis to prevent

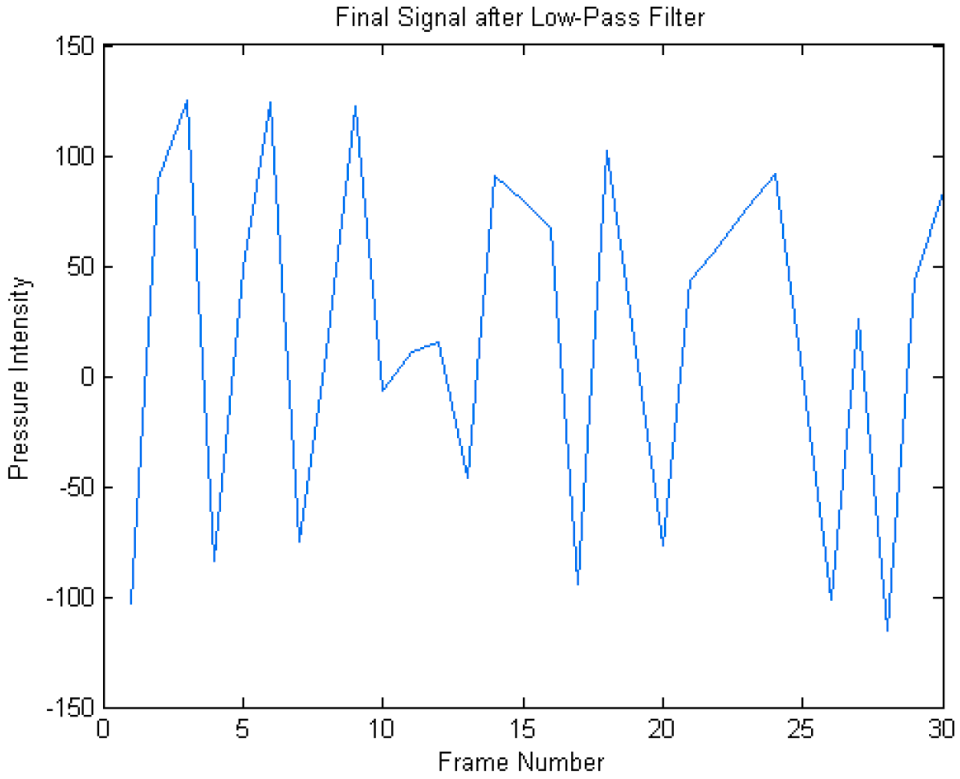

Figure 2.16: The pressure signal obtained from the bed sheet while the subject is in the prone position. Each local maximum (peak) indicates an expansion of the chest which corresponds to an inhalation event.

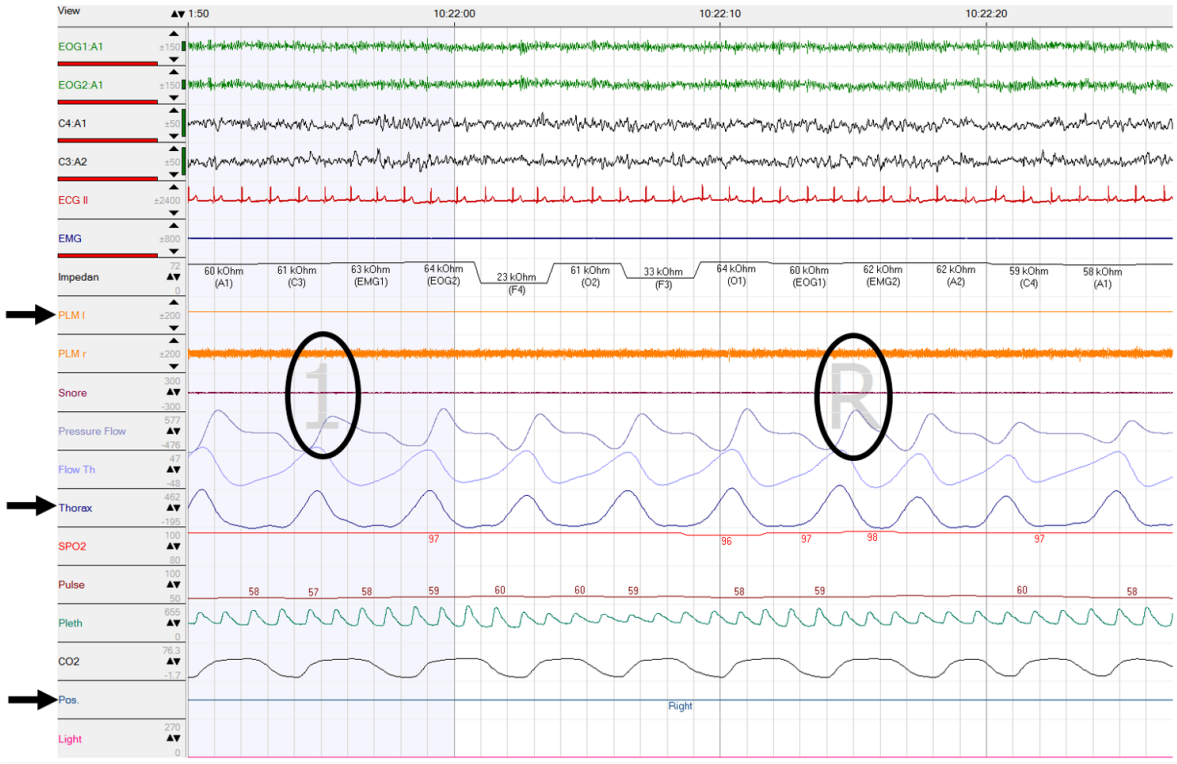

Figure 2.17: A sample polysomnogram that shows a subject transitioning from sleep stage N1 to REM as indicated by the circles. The polysomnogram shows 19 of the PSG signals and the three arrows are pointing to the three channels that are used as ground truth for leg movement (PLM), respiration (Thorax) and whole-body movement (Posture).

sleepers from acting out their dreams and hurting themselves [PAF01]. Therefore, the lack of movement can be a good indication of the REM sleep stage. During movements, certain body parts like arms, hands, elbows or knees are lifted off the bed sheet. This abrupt release in pressure results in a significant drop in pressure intensity. This sudden and significant change in pressure can be detected by a simple thresholding technique that keeps track of the difference between local maxima and minima over a sliding window and reporting a movement if the difference drops significantly. Our algorithm reports a movement if there is a change of more than an order of magnitude in this peak-to-peak amplitude. The detected movements are validated against the posture information provided by the PSG system. Our thresholding algorithm correctly detects 96.5% of the whole-body movement events.

Posture and Body Orientation Features In addition to the biophysical features described above, we also extracted some geometric features from the pressure images. The geometric features are motivated by the fact that the orientation of the body during sleep as well as sleep posture can affect sleep stages. This is mainly because we are likely to go into NREM and the deeper stages of sleep if our bodies are situated in a comfortable position, especially when the body is in a "mid-line" position, where both the head and neck are kept straight. Furthermore, according to the British Snoring and Sleep Apnoea Association [Ass], patients who sleep in the side positions often demonstrate a decrease in the amount of NREM and an increase in the amount of REM sleep. We used 32 geometric features including body symmetry, balance, hip location and shoulder location. These features summarize posture and body orientation and are explained in more detail in [LXA13].

2.3.4.2 Two-Phase Classification

To perform sleep stage identification, the pressure images from a full-night sleep are divided into groups of 30 frames. Since the sampling rate for the bed sheet is 1 Hz, 30 frames corresponds to 30 seconds of sleep. It is common in polysomnography analysis to split the night into 30-second epochs, so the same practice was followed in this work. For each 30-second epoch, the respiration rate, the amplitude and frequency variability of the

respiratory signal, the leg movement events as well as the geometric features described in Section 4.1 are extracted. Since those features mostly distinguish the NREM stage from the rest of the stages, these features are used for phase I classification into NREM and Wake + REM. Phase I, therefore, acts as a filter for the NREM sleep stage. In phase 2, the algorithm tries to distinguish between the Wake and REM stages. Since no strong body movement can occur in the REM stage, as mentioned before, it can be used as a feature to distinguish between Wake and REM. For all the epochs that were classified as Wake + REM in phase I of the algorithm, body movement events are extracted and used to determine if a given epoch should finally be classified as Wake or REM. The decision tree for the described process is shown in Figure 2.18.

Three statistical classifiers were used to evaluate this two-phase procedure of classifying a sleep epoch to the NREM, Wake or REM sleep stage. The parameters for each classifier were obtained by learning from our training dataset. In each phase of classification, each target/non-target sample was labeled with the binary values 1 or 0, respectively; each classifier was used to assign a testing data point - in our case a sleep epoch - to one of these two labels which map to two different sleep stages. The overall accuracy is determined by the percentage of the testing data that was assigned to the correct sleep stage. The three classifiers used in this work are based on different statistical principles and are described below.

The binary Support Vector Machine (SVM) classifier is a non-probabilistic linear classifier. It constructs an optimized hyperplane in the feature space such that the separation between two different types of samples is maximized. The hyper-plane is obtained by learning from the training samples. Each training sample contains a binary label to indicate the group it belongs to. Once the hyper-plane is calculated, the testing data can be projected onto the feature space and classified by the hyper-plane into one of the two categories.

The K-Nearest-Neighbor classifier uses the notion of distance between data points in the feature space as the basis for classification. It assigns a testing data point to the class which the majority of the k nearest neighbors are from. Similar to SVM training, the labels of the nearest neighbors are obtained from the training data. In this work, Euclidean distance is used as the distance metric and the majority rule is applied to the

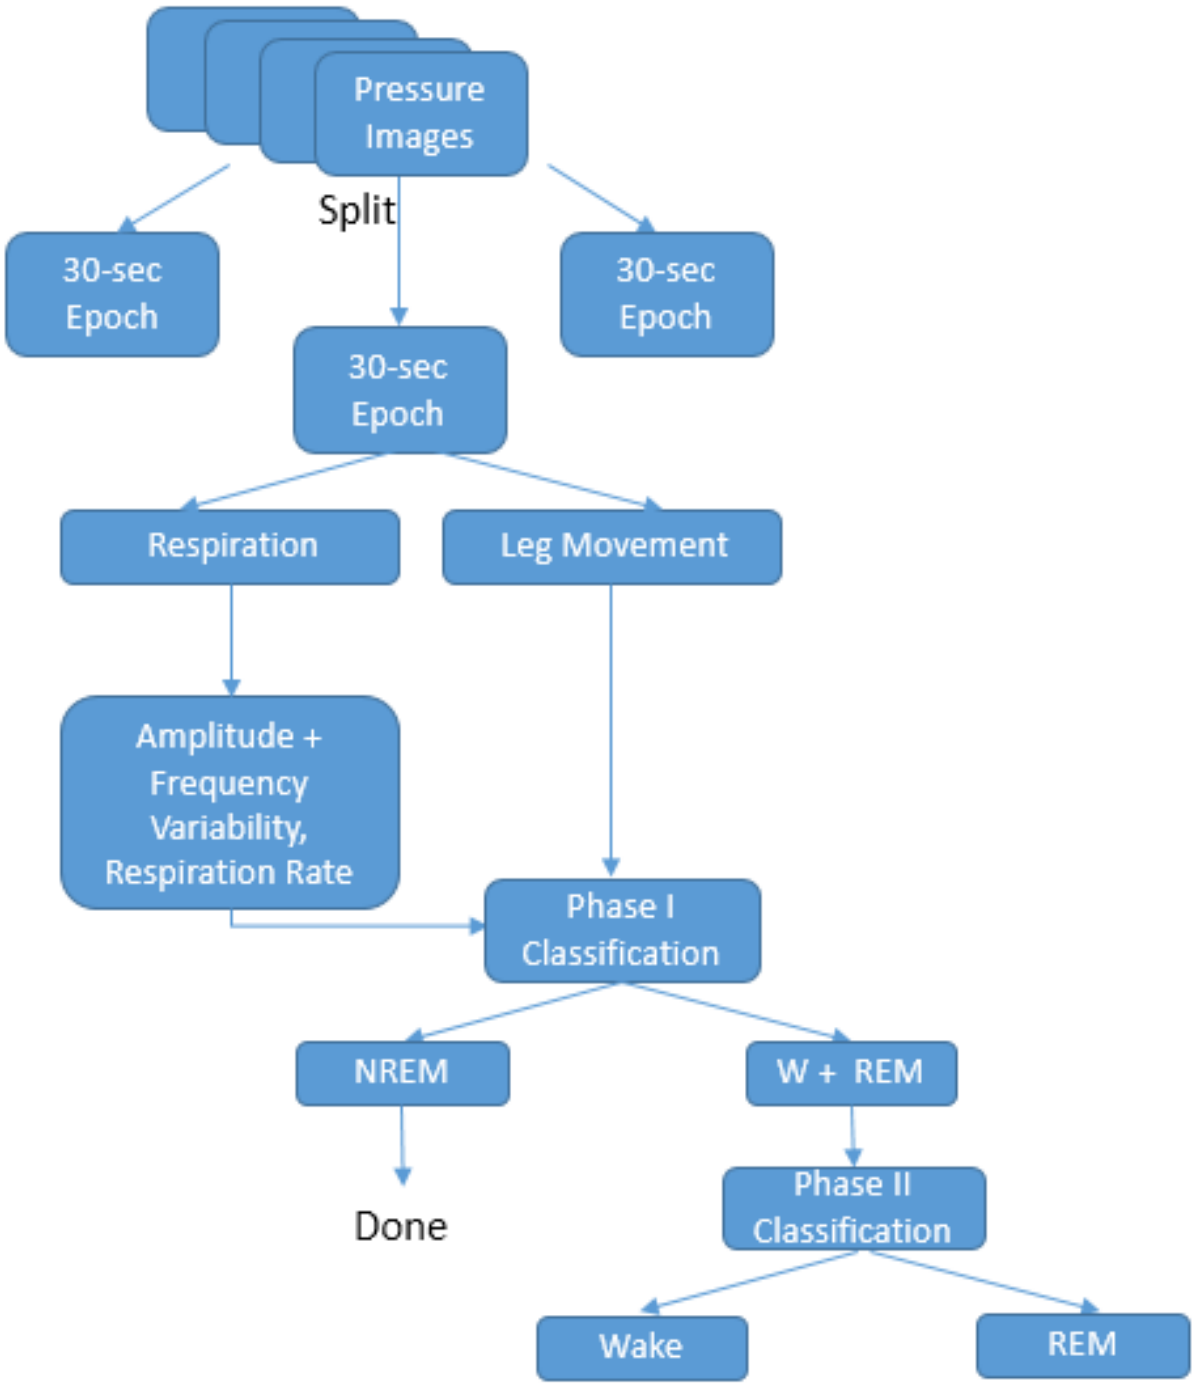

Figure 2.18: The steps of the sleep stage identification process. The pressure images of a whole night are split into 30-second epochs from which biophysical signals are extracted. Phase I classification groups epochs into NREM and Wake+REM. In Phase II, the Wake+REM epochs are then classified separately based on the extracted movement information.

k candidates in order to determine the group that the sample belongs to.

Naive Bayes is a probabilistic classifier based on the Bayesian theorem. It is particularly suitable when the dimensionality of the inputs is high. It assumes that features contribute independently to the probability of a given sample belonging to a certain class. Bayes' Theorem finds the probability of an event given the probability of another event that has already occurred.

2.3.5 Experimental Setup

Seven subjects participated in the sleep study. Three of the subjects were female, four were male. Their ages ranged from 21 to 60 years and their weights from 93 to 190 pounds. Each subject underwent a full-night PSG study and had all the sensors described in Section 1.2.4 attached to him/her. At the same time, each subject slept on the e-textile bed sheet and his/her pressure images were continuously recorded overnight. The 50+ hours of PSG recordings from the seven subjects were scored by SOMNOmedics Domino software. The thorax effort signal, leg EMG signal, sleep posture and sleep stage information provided by the software served as ground truth for all extracted features. The experiments for all 7 subjects were conducted under the same conditions. To ensure the setting was comfortable for sleeping, the light was turned off, the air in the room was in circulation and nearby noise sources were eliminated. The room temperature was set to 68 degrees Fahrenheit, which is a comfortable temperature for sleeping. Figure 2.19 shows a subject wearing the PSG sensors and laying down on the bed sheet shortly before the lights were turned off. Once the patient got in bed and before starting the overnight recording, the so-called bio-calibrations were performed. Bio-calibrations are short tests whose purpose is to discover any poor PSG signals or incorrect hookups prior to the start of the study. The subjects were asked to blink 5 times in order to test the EOG signal. They were then asked to point their toes towards their nose. This bio-cal tests the EMG electrodes placed on the legs. Clenching the teeth verifies the accuracy of the chin EMG signal, and breathing through the nose and mouth verifies the flow, pressure and effort signals. To test the microphone, the patient was asked to make a snoring sound, and to verify the posture detection provided by the PSG system, the subject was asked to change the sleeping position. In the morning and before disconnecting the sensors,

the same bio-calibrations are repeated to ensure that all the sensors remained in place throughout the night. One out of the 7 all-night PSG recordings had a problem with the morning bio-calibrations; the thorax effort belt loosened during the night, resulting in a poor respiration signal. That recording had to be repeated the following night. The pressure images obtained from the bed sheet were labeled using the sleep stage information provided by the PSG software. Testing was done using Leave One Out Cross Validation (LOOCV). One subject's pressure images are left out for testing and the other subjects' images are used for training.

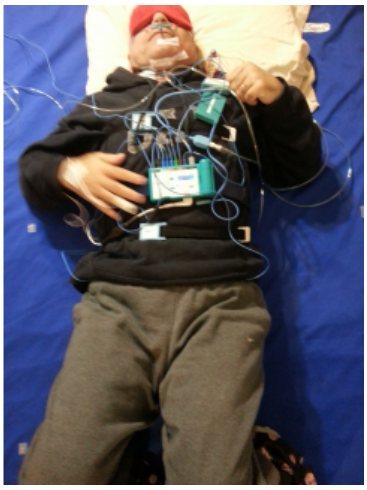

Figure 2.19: A male subject lying on the bed sheet with all PSG sensors attached to his body shortly before the lights were turned off.

2.3.6 Results & Discussion

Table 2.3.6 shows the performance results of our sleep stage identification process based on signals derived from the bed sheet. The results of three different classifiers after phase I and phase II, as well as the overall results, are shown. It is noteworthy to mention that the Naive Bayes classifier outperforms both SVM and KNN in terms of precision, recall and accuracy despite the oversimplified assumptions it makes about the underlying probability model of the data. This is likely the case because in each of the two phases of classification, the correct class is more probable than the other class. In phase I, for example, where we try to separate the NREM stage from the other stages, classifying more samples as NREM than REM+Wake is likely to produce good results since the NREM stage constitutes 75% - 80% of the night and is therefore more probable. The

same argument can be made for phase II.

Classifier	Precision	Recall	Phase I Accuracy	Phase II Accuracy	Overall Accuracy
K-Nearest Neighbor	55.9%	56.9%	68.75%	65.67%	67.12%
Support Vector Machines	67.7%	68.6%	73.1%	68.5%	70.33%
Naive Bayes	70.3%	71.1%	75.23%	69.2%	72.2%

Table 2.2: Comparison of different classifiers in terms of precision, recall, phase I performance, phase II performance, and overall accuracy

Another point to observe is that the performance of phase II is significantly worse than the performance of phase I. This performance degradation is expected since the physical changes that occur during REM sleep are very similar to the ones that occur during the Wake stage, making the task of separating REM and Wake epochs in phase II a challenging task. Because of the similarity between the REM and Wake stages, the REM sleep stage is sometimes referred to as paradoxical sleep. Even though it is one of the stages of sleep, it is characterized by a brain wave pattern and physical signals that are similar to that of wakefulness. Figure 2.20 visually shows the similarity between REM and Wake, as well as the dissimilarity between those two stages and the NREM stage. The figure shows the respiratory signal acquired by the PSG system during three epochs of the N3 stage of NREM (top), REM (middle) and Wake (bottom) of a single patient. As can be seen, the respiratory signal in the NREM stage is regular with very little variation in amplitude and frequency. Both the REM and Wake stages, on the other hand, show the same irregularity with clear variations in the amplitude and frequency of the respiratory signal.

Figure 2.21 shows the hypnogram of one of the seven subjects. A hypnogram is a graph commonly used in polysomnography that represents the stages of sleep as a function of time, specifically as a function of epochs. In a clinical setting, a polysomnography record is usually scored by more than one sleep technician and the hypnogram provides a visual way to show the agreement between the scoring of the different technicians. Here, we use a hypnogram to visually show the agreement between the sleep stages obtained from the PSG system (top) and the ones obtained from the bed sheet after feature extraction and classification (bottom). The classifier used in the creation of this hypnogram is SVM. The hypnogram from the bed sheet shows 77.48% agreement with the

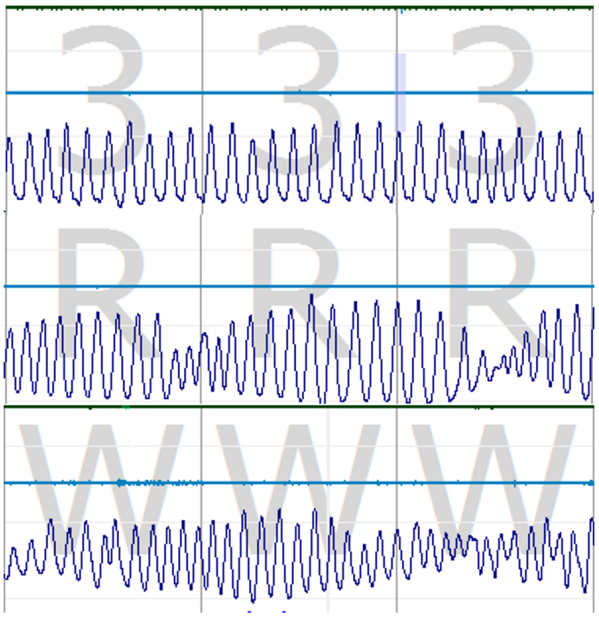

Figure 2.20: The respiration signal during three different epochs: one epoch during the N3 stage of NREM (top), another epoch during REM (middle), and another during wakefulness (bottom).

gold-standard hypnogram. It can be observed from the figure that a significant majority of the NREM epochs are classified correctly as indicated by the matching NREM portions of the hypnograms. It can also be observed that the REM and Wake stages are often misclassified. The red bolded lines in the bottom hypnogram show the epochs that were incorrectly classified as Wake as a result of our classification algorithm. The red bolded lines in the top hypnogram show that these misclassified epochs were scored as REM epochs by the PSG system.

Finally, Tables 2.3 and 2.4 show two confusion matrices for the same patient as in the hypnograms. In Table 2.3, the confusion matrix shows the precision and recall values for the two-phase classification algorithm described in Section 2.3.4.2. Again, we can observe that phase II of the classification performs poorly compared to phase I due to the similarity of the REM and Wake stages. If no significant body movements occur during the Wake stages then separating REM and Wake becomes even more difficult. In Table 2.4, the confusion matrix corresponds to a one-phase classification procedure where we attempt to classify all three stages without filtering out the NREM stage first. This one-phase procedure leads to significantly worse precision and recall values because it leaves

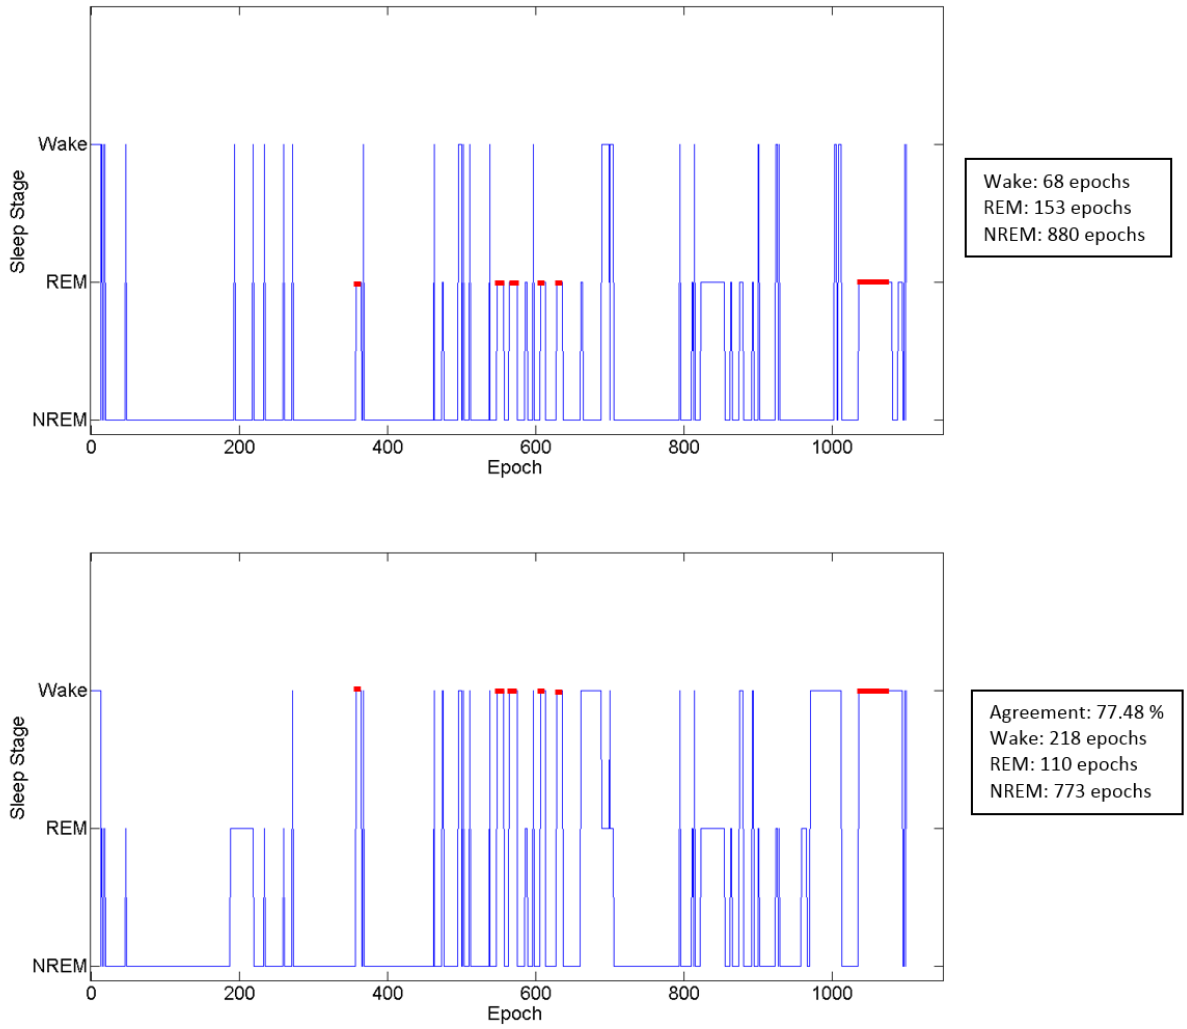

Figure 2.21: The top hypnogram shows sleep stages over time (epochs) as provided by the PSG system. (Wake: 68 epochs, REM: 153 epochs, and NREM:880 epochs). The bottom hypnogram shows sleep stages over time as obtained by the two-phase classification procedure. (Wake: 218 epochs, REM: 110 epochs, and NREM: 773 epochs, Agreement: 77.48%)

	NREM	REM	Wake	Recall
NREM	**773**	38	69	87.8%
REM	0	**42**	111	27.5%
Wake	0	30	**38**	55.9%
Precision	100%	38.2%	17.4%	77.5%

Table 2.3: Confusion matrix for one of the subjects when the SVM classifier is used in two-phase classification.

	NREM	REM	Wake	Recall
NREM	**729**	76	113	79.4%
REM	2	**12**	139	7.8%
Wake	27	25	**16**	23.5%
Precision	96.2%	10.6%	6.0%	68.8%

Table 2.4: Confusion matrix for one of the subjects when the SVM classifier is used in one-phase classification.

more room for error between all three classes, whereas the two-phase process eliminates the epochs that are well distinguishable from the other ones first before it proceeds to the more difficult task of separating the two similar classes - REM and Wake.

2.3.7 Conclusion

In this work, we proposed a non-intrusive, completely wireless and contact-free sleep stage identification system. We extracted a set of sleep-related biophysical features as well as geometric features from pressure images obtained from an e-textile bed sheet. These features were used as part of a two-phase classification procedure that first separates NREM from REM and Wake and then separates the two latter classes. The superiority of the two-phase procedure over the one-phase procedure was shown. The system achieved up to 70.3% precision and 71.1% recall on average. The proposed unobtrusive system opens the way to a cheap, contact-free sleep diagnosis solution, eliminating some of the drawbacks that the traditional PSG method presents. The combination of a flexible and non-intrusive bed sheet system with simple signal processing and classification makes the

described system a portable sleep screening solution that can be used in a clinical as well as a home environment. The affordability of the system can make sleep screening accessible to a bigger population which could lead to early diagnoses of sleep disorders and chronic diseases.

2.4 In Closing

This chapter presented two nocturnal methods for diagnosing OSA. The methods utilized a pressure-sensitive bed sheet that fit seamlessly into a person's environment without disrupting sleep. Although non-intrusive, these overnight methods continuously monitor a patient over the duration of their sleep and produce a large amount of data that need to be stored and analyzed. The next chapter discusses daytime diagnosis methods that, in addition to being non-intrusive, also produce considerably less data for analysis.

CHAPTER 3

Non-Intrusive Daytime Diagnosis Methods

3.1 Introduction

Chapter 2 discussed non-intrusive, nocturnal methods for the diagnosis of OSA. Nocturnal diagnosis methods attempt to diagnose OSA by detecting signs of the disorder that occur while the patient is sleeping. The great majority of OSA diagnosis methods available in the literature fall under the nocturnal category. More recently, however, the research community started shifting its focus to daytime diagnosis methods, which can detect the OSA disorder while the patient is awake. These methods are very rare, but if proven effective, they would have many advantages compared to nocturnal methods. The first advantage is that the patient can stay awake while being tested. This means that the patient's sleep quality will not be affected in any way. Since the nocturnal methods presented in Chapter 2 are contact-free and do not interfere with the subject's sleep, this particular benefit is not really applicable to the nocturnal methods presented in this work. Another advantage of a daytime diagnostic test is the fact that it has to rely on an inherent characteristic of OSA in order to be reliable. Since a daytime test cannot detect apneic episodes, it would have to rely on detecting more permanent physiological or psychological changes caused by the OSA disorder. This has two benefits: first, it can help us understand more about the OSA disorder and its consequences and secondly, it can detect the disorder in situations where a nocturnal test wouldn't be able to detect it. Some patients, for example, have positional OSA, where the apneic episodes do not occur when the patient sleeps in the side position. If a patient with positional OSA sleeps in the side position during the diagnostic test, nocturnal methods, that rely on breathing cessations to make a diagnosis, will not be effective. Daytime methods, however, rely on manifestations of the disorder that transcend sleep and can therefore always make a diagnosis. Another advantage of daytime diagnosis methods is that they usually produce

a lot less data. Unlike nocturnal methods that monitor patients and record their signals over an 8-hour period, daytime methods are much shorter tests that produce small data that can be stored and analyzed more efficiently.

This chapter presents novel daytime methods that attempt to detect permanent changes caused by the OSA disorder using completely non-intrusive monitoring tools and techniques.

3.2 An Automated Framework for Predicting Obstructive Sleep Apnea Using a Brief, Daytime, Non-Intrusive Test Procedure

Many different questionnaires have been designed for OSA screening. These questionnaires try to identify patients at risk of OSA by asking them questions about symptoms (snoring, EDS, etc.), comorbidities (e.g. hypertension), risk factors (age, BMI, neck circumference, etc.) and family history. One of the most common validated questionnaires for OSA screening is the Berlin questionnaire. It consists of 11 questions and stratifies patients into high or low risk of having OSA based on their endorsement of symptom severity [NSN99]. Though questionnaires like the Berlin questionnaire are validated and commonly used by clinicians, they still primarily rely on self-report, which is prone to error. Patients may not answer questions very accurately because of genuine forgetfulness and lack of judgment. However, some patients may also deliberately provide the wrong answers to the questions to exaggerate or minimize the severity of their condition. Even though these questionnaires have the advantage of being short and providing a quick way to exclude low-risk patients, they still cannot replace objective forms of evaluation. In this work, we propose a daytime screening tool that combines the questionnaires with a more objective protocol to assess the presence of OSA. Though it may seem counterintuitive to diagnose a sleep disorder during wakefulness, the proposed protocol relies on detecting changes that transcend the period of sleep. Specifically, it leverages the brain structural changes that are caused by the OSA disorder. Studies have shown that OSA is accompanied by brain injury as a result of oxygen depletion to the brain. The mechanisms by which OSA leads to brain injury have been demonstrated in both animal and

human studies [MMQ03], [MKW08]. It has been shown that the injury specifically affects central autonomic regulatory regions of the brain, which are responsible for regulating cardiovascular functions, like blood pressure and heart rate. In this work, we present the subjects with a set of autonomic challenges that are designed to trigger cardiovascular responses. Analyzing the cardiovascular responses signaled by a healthy brain versus the responses signaled by an injured brain provides insights into the existence of OSA.

3.2.1 Related Works

The serious negative impact of obstructive sleep apnea on human health makes monitoring and diagnosing it a necessity. Due to its debilitating nature and in an effort to reduce cost, waiting time, and the discomfort associated with the traditional OSA diagnostic method, a plethora of novel OSA detection methods have been proposed in the literature. The great majority of these methods are nocturnal ones. Some of these methods seek to detect OSA acoustically by analyzing breathing and snoring sounds ([YM09], [AGK12]). Others use respiratory effort to detect if a subject has stopped breathing ([ANK00], [HXL13]). Other techniques rely on brain waves (EEG) and sleep stage information ([MSK09], [SHL14]), heart rate ([CHS03], [SZE00], [MPI00]) , or oximetry ([RRK14], [MHA09]), or a combination of these signals to detect OSA. A more detailed explanation of these methods was provided in Section 2.3.2. Only very few works have attempted to detect sleep apnea during wakefulness and they all have considerable drawbacks. In [LRH09], the authors use ultrasound imaging to measure the width of the subject's tongue. They claim that the severity of the OSA disorder is correlated with the width of the tongue base; however, no results to support that claim were shown. In [SMK93], ultrafast MRI was used in an attempt to detect and localize the sites of obstruction in the pharyngeal airway. While the technique achieved good results (87%) when performed on a sleeping subject, it performed worse than a random guess (47%) when conducted during wakefulness. Another study tried to predict OSA during wakefulness by using EEG and pupil size as predictors ([LPL08]). Sleepiness affects the beta and theta brain wave patterns and the size of the pupils. The test relies on the OSA patients exhibiting signs of excessive daytime sleepiness when they come in for the test. Even though this test can be performed during wakefulness, it is still restrictive as it needs

to be performed 12 hours after the subject's mid-sleep period to increase the chance of sleepiness occurring, and if sleepiness does not occur, the test is inconclusive. In [MM11], a method based on tracheal breath sound analysis was proposed. The breath sounds resulting from breathing through the nose and mouth were recorded using a microphone placed over the suprasternal notch of the trachea in both the supine and upright positions. Based on the features extracted from the digitized acoustic signals, the authors pictorially showed that the extracted features achieve good separability between the OSA and non-OSA classes, but again, no quantitative classification results were provided.

Given all the aforementioned methods, the main contributions of this work are:

- An end-to-end OSA screening framework that can identify OSA patients while they are awake using a novel technique that triggers brain responses that are indicative of the OSA disorder. The technique does not rely on a specific time or patient condition (e.g. sleepiness) and does not disrupt the patient's sleep like the nocturnal methods.
- The screening method is simple, non-intrusive, and requires no specialized personnel. It can be performed in a home environment or at the physician's office.

3.2.2 Obstructive Sleep Apnea Screening Tool

In this section, we describe our OSA screening tool that consists of a) an Android phone application for data collection and b) a data analytics component for OSA prediction. This section also describes the clinical trial conducted to validate the viability of the proposed screening tool.

3.2.2.1 Data Collection Framework

Our data collection system consists of a smartphone loaded with our data collection software, a Bluetooth-enabled blood pressure monitor, and a Valsalva device. The data collection software is an Android application that guides users through a 15-minute protocol. The structure of the protocol is shown in Figure 3.1. Phase I of the application includes some setup instructions for the patient. It ensures that 1) patients are ready to

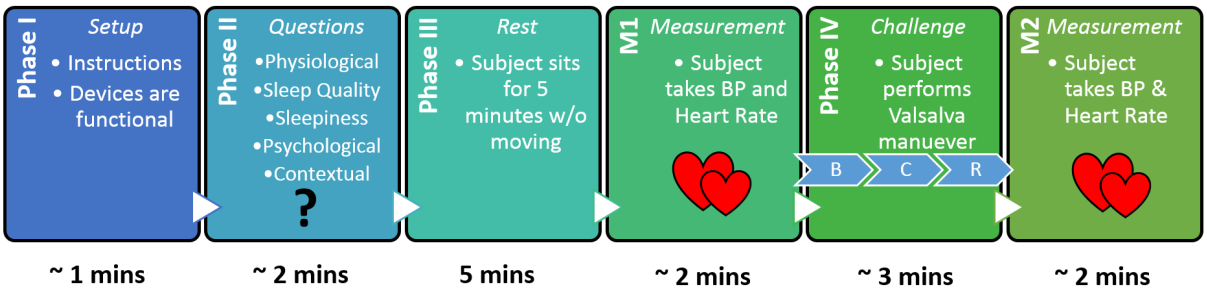

Figure 3.1: 15-minute protocol implemented by the data collection system.

spend an uninterrupted 15 minutes taking the test, 2) the blood pressure (BP) cuff is on and the BP monitor is ready to use and 3) the Valsalva device has been tested.

In phase II of the protocol, the application presents the user with a questionnaire. The questionnaire consists of 15 questions that aim to assess the user's sleep quality, sleepiness, psychological status, as well as some aspects of the user's context. The questionnaire consists of a short form of the Pittsburgh Sleep Quality Index (PSQI) [BIM89], the sleepiness screening (SS) tool [ZRF08], the Perceived Stress Scale (PSS-4) [CKM83], as well as the Patient Health Questionnaire (PHQ-4) [KSW09], which are all short forms of validated questionnaires used by health care professionals as symptom measures for sleep quality, sleepiness, perceived stress, and psychological health, respectively. The questionnaire also includes two contextual questions that aim at assessing the user's level of physical activity and if any unusual events have occurred in the user's life on a given day. Table 3.2.2.1 lists each of the 15 questions and what measure they each try to assess. Following the questionnaire, the application visually guides the patient through a 5-minute rest period and indicates the end of the rest period with a beeping sound.

The phase III rest period comes in preparation for the BP measurement to follow. In M1, the patient is instructed and walked through the steps of taking his/her resting BP (bp_baseline) and resting heart rate (HR_baseline) using the Bluetooth-enabled blood pressure monitor.

After the BP measurement, phase IV – the challenge phase – starts. The challenge phase consists of 3 parts: baseline (B), challenge (C), and recovery (R). In the baseline period, users are instructed to stay calm for 60 seconds in order to reach their normal/resting BP and to prepare for the following challenge. In the challenge period, the user is visually guided to perform the Valsalva maneuver wherein patients blow against

Measure	QID	Question	Validated Questionnaire
Sleep Quality	*q_1*	What time did you go to bed?	PSQI
	q_2	Approximately how long did it take for you to get to sleep?	
	q_3	Approximately how many hours did you sleep?	
	q_4	What time did you get out of bed in the morning?	
Sleepiness	*q_5*	Please rate your sleepiness today (0 = none, 10 = extremely high).	SS
Perceived Stress	*q_7*	In the last day, how often have you felt that you were unable to control the important things in your life?	PSS-4
	q_8	In the last day, how often have you felt confident about your ability to handle your personal problems?	
	q_9	In the last day, how often have you felt that things were going your way?	
	q_10	In the last day, how often have you felt difficulties were piling up so high that you could not overcome them?	
Depression	*q_11*	In the last day, how often have you experienced little interest or little pleasure in doing things?	PHQ-4
	q_12	In the last day, how often have you felt down, depressed or hopeless?	
Anxiety	*q_13*	In the last day, how often have you felt nervous, anxious or on edge?	
	q_14	In the last day, how often have you not been able to stop or control worrying?	
Context	*q_6*	Did you experience anything out of your normal routine in the past 24 hours?	N/A
	q_15	Rate your level of physical activity over the last 24 hour period.	

Table 3.1: Questionnaire presented at phase II of the data collection system. Column 1: measure to be assessed, Column 2: question ID, Column 3: question text, Column 4: the name of the validated questionnaire containing the questions

a closed airway connected to a pressure measuring device. Using the Valsalva device, the patient is instructed to exhale for 18 seconds at a pressure of 40 mmHg, which is indicated by a blue light on the Valsalva device used in this study. The Valsalva maneuver leads to changes in blood pressure and heart rate and is usually used as a test of cardiovascular function [PBT84]. Since OSA is associated with a dysfunction in the cardiovascular regulation, comparing the cardiovascular responses caused by the Valsalva maneuver as signaled by an injured brain to the changes signaled by a healthy brain could provide evidence of the presence of OSA. The challenge period is followed by a 90-second period in which the user recovers from the challenge.

After phase IV, a second set of BP and heart rate measurements (bp_end and HR_end) is taken at the end of the test. The BP and heart rate measurements (M1 and M2), as well as the answers to the questionnaire, are collected by the application and uploaded to a HIIPA-compliant server.

3.2.2.2 OSA Prediction Framework

The ultimate goal of the developed system is to classify a patient into one of the two following classes: non-OSA (healthy) and OSA. This section describes the prediction system which takes a set of features as input and, after analyzing the features, makes a

decision about a specific patient.

A. Feature Extraction The first step of the prediction process is to extract a useful set of features from the collected signals. This step is critical to the success of any prediction system as an effective set of features decreases the computational complexity of the system and increases classifier performance. Five categories of features are used in the analysis: physiological, sleep quality-related, sleepiness-related, psychological, and contextual features. The physiological features are extracted from the BP and heart rate measurements, while the other features are extracted from the questionnaire responses.

Category I - Physiological The physiological features used in the analysis are bp_sys_diff, bp_dias_diff, and HR_diff, which correspond to the difference between the systolic BP measurements M1 and M2, the difference between the diastolic BP measurements M1 and M2, and the difference between the heart rate measurements M1 and M2, respectively. The Valsalva maneuver performed in phase IV of the protocol is known to evoke a sequence of blood pressure and heart rate changes. These changes involve a rapid increase in blood pressure at the time of exhalation (BP rises above baseline), followed by a rapid decrease in blood pressure at the time of air pressure release when the user starts breathing normally (BP falls below baseline) and finally, recovery (BP goes back to baseline). The heart rate shows a similar response but in the opposite direction of blood pressure. Studies have shown that these same changes occur in healthy as well as OSA subjects, but that a sleep apneic's response is delayed and less pronounced. For example, [MKW13] shows that the OSA subjects failed to raise heart rate to the same extent as the healthy subjects and that the start of the increase is delayed compared to healthy subjects. The three aforementioned features aim to capture the extent of change in blood pressure and heart rate between the baseline measurement M1 and the post-Valsalva measurement M2 to help distinguish between OSA and healthy subjects.

Category II - Sleep Quality The next set of features aims to distinguish between the two populations based on the quality of their sleep.

- sleep latency: This metric is a measure of sleep deprivation and is derived from q

2. Sleep latency is the amount of time it takes a person to fall asleep. Due to their disorder, OSA patients are usually sleep deprived and fall asleep much faster than a healthy person. We expect OSA patients to have a shorter sleep latency. Based on PSQI, we differentiate among 4 different latency levels: $\leq$ 15 minutes (sleep deprived), 16-30 minutes (normal), 31-60 minutes (long normal), and $>$ 60 minutes (insomniac).

- sleep efficiency: This metric is an important measure of sleep quality. It is the ratio of time spent asleep to the amount of time spent in bed ($\frac{total\_time\_spent\_sleeping}{total\_time\_spent\_in\_bed}$). It is derived from questions q_1, q_3 and q_4 and is calculated as $\frac{q\_3}{q\_4 - q\_1}$. OSA patients can have a very high sleep efficiency since, due to sleep deprivation, they fall asleep quickly and stay asleep for most of the time while they are in bed, even though not in NREM3 and REM sleep, which are the sleep stages where the physical and mental regeneration occur, respectively. Some OSA patients, however, can have a low sleep efficiency if they suffer from depression – a common comorbidity of OSA. Insomnia – a symptom of depression – can cause OSA patients to have a lower sleep efficiency than healthy subjects. In both cases, OSA subjects' sleepiness (explained under Category III) will be negatively affected. We follow the same scoring scheme used in PSQI and differentiate among 4 different levels of sleep efficiency: $>$85% (very good), 75% - 85% (good), 65% - 74% (bad), and $<$65% (very bad).

- sleep duration: This metric is derived from q_3. Again based on PSQI, we differentiate between 4 levels of durations: $>$7 hours, 6-7 hours, 5-6 hours, and $<$5 hours. We expect sleep apnea patients to report longer sleep durations because their poor sleep quality does not give them the rest they need, causing them to stay in bed longer.

Category III - Sleepiness Another feature we use in our prediction system is sleepiness. This metric is a measure of daytime sleepiness and is derived from q_5. We expect OSA patients to have higher sleepiness values than healthy subjects since OSA results in excessive daytime sleepiness.

Category IV - Psychological The next set of features aims to distinguish between OSA and non-OSA subjects based on their psychological health which is negatively affected by the lack of efficient, regenerative sleep, as well as by the brain injury caused by OSA. We expect OSA patients to have poorer psychological health than healthy subjects.

- perceived stress: This metric is a measure of the degree to which a person is experiencing stress in his/her life. This feature is derived from questions q_7, q_8, q_9 and q_10 and scored according to PSS-4. Each of the perceived stress questions has 5 response alternatives corresponding to 0=never, 1=almost never, 2=sometimes, 3=fairly often, and 4=very often. This measure is calculated as q_7 ¬ q_8 + ¬ q_9 + q_10, where the ¬ q_x is the negation of the question's answer, calculated as 4 minus the provided answer. The higher the calculated perceived stress metric, the more stressed a subject is.

- depression: This metric is used by health care personnel to screen for depression. The metric is derived from q_11 and q_12 and scored according to PHQ-4. Each of the depression questions has 4 response alternatives: 0=not at all, 1=several times, 2=very often, and 3=the whole day. This metric is calculated by summing up the answers to both questions: q_11 + q_12. Based on PHQ-4, scores are rated as depressive (3-4) and normal (0-2).

- anxiety: This metric is used by health care personnel to screen for anxiety. The metric is derived from q_13 and q_14 and scored according to PHQ-4. Each of the anxiety questions has 4 response alternatives: 0=not at all, 1=several times, 2=very often, and 3=the whole day. This metric is calculated by summing up the answers to both questions: q_13 + q_14. Based on PHQ-4, scores are rated as anxious (3-4) and normal (0-2).

- psychological distress: This metric combines the depression and anxiety scores into one global measure of psychological distress. It is calculated by adding the depression and anxiety scores (depression + anxiety). Based on PHQ-4, scores are rated as normal (0-2), mild (3-5), moderate (6-8), and severe (9-12) and we differentiate between these 4 levels of distress in our data.

Category V - Contextual The next couple of features help us understand the user's context.

- atypical events: This metric indicates whether or not patients have experienced anything out of their normal routine during the last 24 hours. This question helps us understand the user's daily context and is very useful for eliminating outliers. Unusual events can elicit OSA symptoms in non-OSA subjects. A non-OSA patient who just got fired, for example, may exhibit the same symptoms of depression, anxiety and poor sleep quality as an OSA subject. This metric is derived from q_6.

- physical activity: This metric is derived from q_15 and reflects how energetic a subject is. The higher the level of physical activity of a certain subject, the more likely it is that he/she is happier and less stressed and anxious.

B. Classification In the classification step, the feature vector including the aforementioned features for a specific subject on a specific day is provided to the classifier as input. The classifier then maps the input to one of the two possible output classes: OSA or non-OSA. Two statistical classifiers were used to classify a feature vector belonging to a subject into the OSA or non-OSA category. A training dataset is used to train the classifier and to obtain the parameters that will provide the best classification performance. The training dataset contains multiple instances of feature vectors and their corresponding class labels. A separate testing dataset contains unlabeled feature vectors that the classifier will assign to a class label. The overall accuracy of the classifier is determined by the percentage of the testing data that was assigned to the correct output class. The classifiers used in this work are based on different statistical principles and are described below.

- The Support Vector Machine (SVM) classifier is a binary non-probabilistic linear classifier. Training a support vector machine consists of constructing an optimal hyperplane in the feature space. The optimal hyperplane is the one that maximizes the separation between the nearest training samples of the two different classes. Once the hyperplane is constructed, the testing data can be projected onto the feature space and classified by the hyperplane into one of the two classes. In this

study, we used the sigmoid kernel function with a cost of 0.8 and an intercept constant of 0.1.

- The k-Nearest-Neighbor (k-NN) classifier uses the notion of distance between data points in the feature space as the basis for classification. In the training phase, the algorithm simply stores the labeled feature vectors and the class labels of the training samples. In the testing phase, it assigns a testing data point to the class that is the most common among the testing point's k nearest neighbors based on a majority vote of its neighbors. In this work, k = 3 and the Manhattan distance is used as the distance metric.

3.2.2.3 UCLA Obstructive Sleep Apnea Study

Our system has been deployed in a UCLA IRB approved pilot study of 16 subjects. 3 of the subjects had been diagnosed with sleep apnea through the gold standard PSG method and the remaining 13 are healthy control subjects. All subjects were given kits that consist of an Android phone, preloaded with our application, a Bluetooth-enabled BP monitor, and a Valsalva device. Every subject completed the test protocol described in Section 3.2.2.1 daily for a total period of 42 days. The data recorded by the phone application is uploaded to our HIIPA-compliant server when each daily test is completed. The uploaded data is then split into a training set and a testing set. The training set is used by our prediction system to train the classifiers, while the testing set is used as input to our trained classifiers to make predictions and to evaluate the performance of our classifiers.

3.2.3 Results & Discussion

Experiment I To evaluate the performance of our classifiers, testing was done using Leave-One-Out-Cross-Validation (LOOCV). For this study, this is the only kind of cross-validation that can evaluate the accuracy of the classifiers without bias. Due to the fact that the data includes 42 observations per subject, an n-fold cross-validation approach would yield overly-optimistic results since a subset of those observations will be in the training set and the rest of the observations will be in the testing set. Having observations

	Predicted Outcome			
	SVM		k=3 NN	
OSA?	Yes	No	Yes	No
Yes	**68**	34	**80**	22
No	72	**352**	109	**315**
Sensitivity/Specificity	66.7%	83.0%	78.4%	74.3%
Overall Accuracy	79.8%		75.1%	

Table 3.2: Confusion matrix for the SVM and k-NN Classifiers

for the same subject in both the training and testing sets introduces bias into the results since we expect observations for the same patient to be similar even if the test was performed on different days. In this study, we employ a special case of LOOCV, namely a leave-one-patient-out-cross-validation (LOPOCV) in which we omit all observations for the same patient and train the classifier on the remaining data. The omitted subject's observations are then used for testing. After filtering out outliers, we are left with 424 observations for the non-OSA class and 102 observations for the OSA class. Due to the imbalance in our dataset (3 OSA subjects, 102 observations vs. 13 control subjects, 424 observations), we also randomly undersampled the training dataset to balance the OSA and non-OSA observations. Using the features described in Section 3.2.2.2 and LOPOCV, we were able to achieve an overall classification accuracy of 75.1% and 79.8% using the k=3 NN and SVM classifiers, respectively. Table 3.2.3 shows the confusion matrices for the two classifiers.

We can see that even though the overall accuracy of the SVM classifier is $\approx$ 80% as compared to the $\approx$ 75% accuracy of the k=3 NN classifier, the sensitivity of SVM is much lower than the sensitivity of k=3 NN. For a medical classification task, it is very important for the results to be both sensitive (all OSA patients are correctly classified as OSA) and specific (all healthy subjects are correctly classified as healthy) since we would not want people with the disorder to go unnoticed or people without the disorder to think they're sick when they really aren't. With 78.4% sensitivity and 74.3% specificity and an overall screening accuracy of 75.1%, the k=3 NN proves to be the better classifier for this classification task.

The SVM classifier builds a global model of the data, meaning that it uses the complete training dataset to compute a global function that maps a feature vector to a class label. k-NN, on the other hand, takes a local approach to classification where prediction is done by local functions using only a subset of the training dataset – the neighboring points in this case. The fact that the local performs better than the global approach for the prediction of OSA suggests that not all members of the same class, whether OSA or non-OSA, are similar to each other, but rather that within the same class there are different clusters of subjects. For example, some healthy subjects can have very poor sleep quality and some OSA subjects can be psychologically healthier than non-OSA subjects, but in both cases, these subjects will share other properties with subjects from their class. SVM constructs a hyperplane that partitions the feature space into two regions belonging to each class. A test point has to fit into one of these two precomputed regions. In that sense, SVM tends to oversimplify the two groups. The local approach works well because decisions are made based on small local neighborhoods of similar clusters. This approach assigns more significance to differences between subjects in the same class.

Experiment II Another important metric for evaluating a classifier, especially one that is applied in the medical field, is the variance. In our study, each subject repeated the screening test daily for a period of 42 days. Though the answers to the questionnaire and the BP measurements may vary from day to day for the same subject, we don't expect the intra-subject variability to be very high and we certainly don't expect the decision as to whether or not a subject is suffering from OSA to change from one day to the next. Therefore, the ideal classifier would assign all the observation for the same patient to the same class label. The repeatability of a medical test is crucial to its validity and reliability. To measure the repeatability of our test, we use statistical variance to measure the variation in the predictions of a classifier; the lower the variance, the more consistent the predictions are and the more repeatable the test is. Figure 3.2 shows the variance for each subject's predictions and the overall average variance of the k=3 nearest neighbor classifier. Notice that a test can have a perfect (zero) variance if it always gives the wrong prediction. In our results, we show that our classification results are both accurate (75.1% accuracy) and repeatable (0.17 variance).

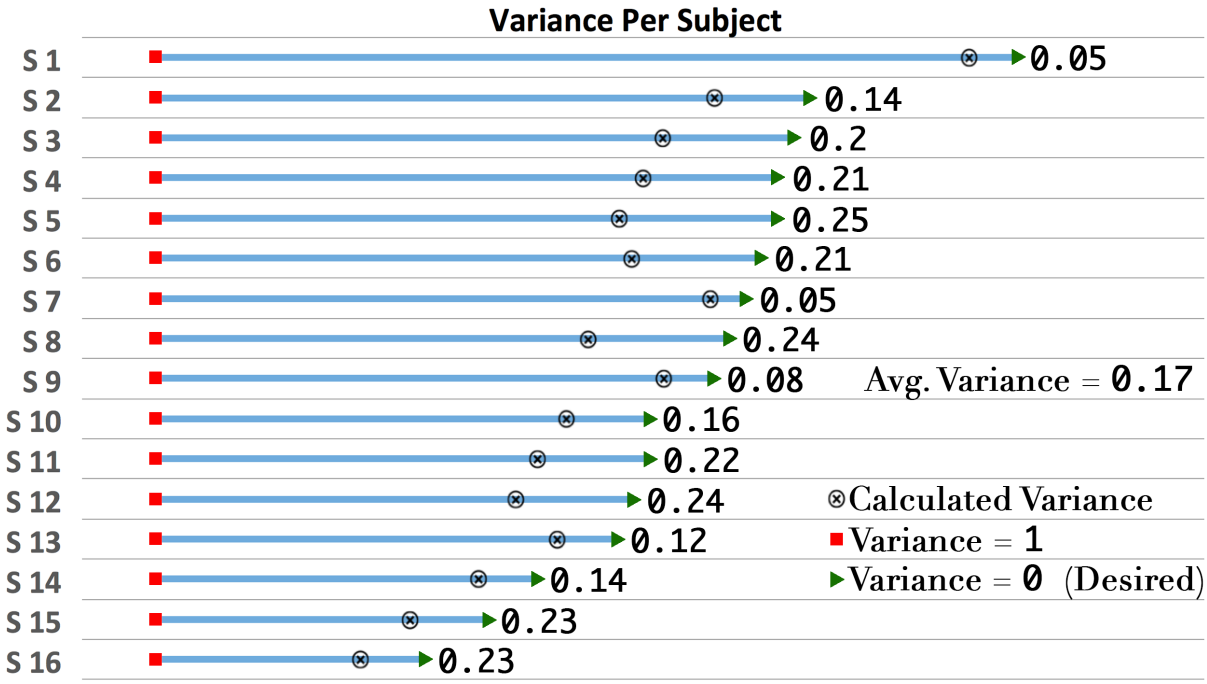

Figure 3.2: Shows the variance for each of the 16 subjects (S1-S16). The length of each line is proportional to the number of observations for each subject after the removal of outliers.

Experiment III In an attempt to better understand the importance of each feature category (see Section 3.2.2.2) for the prediction of OSA, we successively eliminated one feature category from our dataset. We then repeated the training and testing steps for the reduced dataset using the k=3 NN classifier. Figure 3.3 shows the accuracy, sensitivity, and specificity of the classification after eliminating each of the 5 feature categories. The experiment results, as shown in the figure, indicate a decrease in classification accuracy when any of the 5 feature categories is removed, compared to when all categories are preserved. The decrease in accuracy is mainly due to the dramatic decrease in sensitivity. As the figure shows, when any one of the 5 feature categories is missing, the test's ability to correctly identify OSA patients diminishes from $\approx 78\%$, when all features are present, to $\approx 24\%$ on average when one of the categories is missing. This shows that all the feature categories play a vital role in the diagnosis of OSA. The results reveal the intricacy of the OSA disorder which affects first, our sleep quality, and as a result, our body, mind and daily activities.

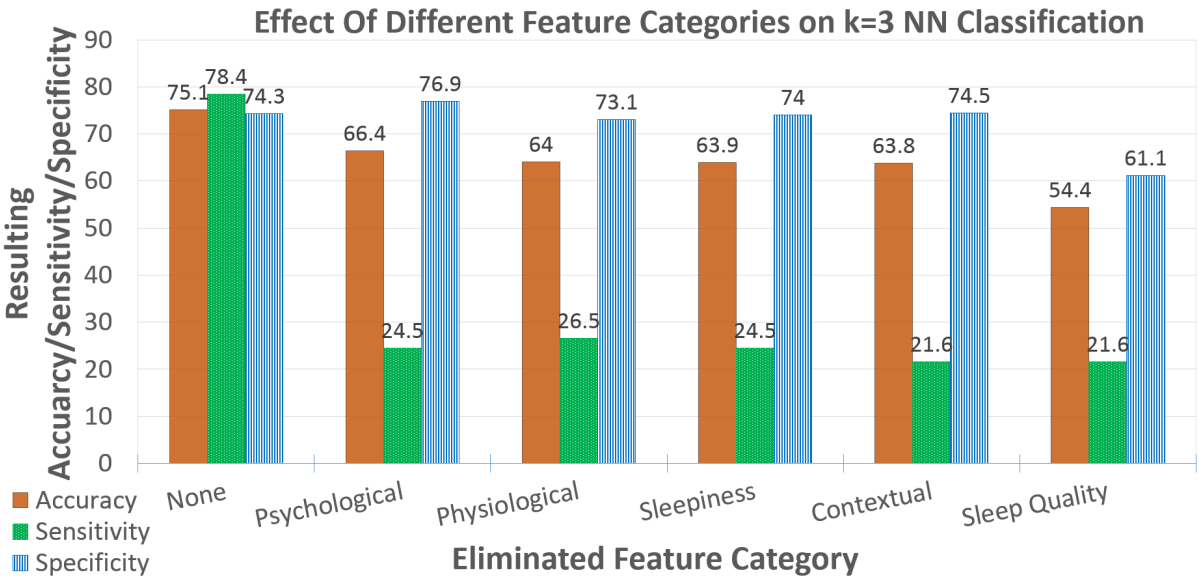

Figure 3.3: Shows the accuracy, sensitivity, and specificity of classification before eliminating any feature categories and after eliminating each category at a time.

3.2.4 Conclusion

In this work, we proposed a daytime OSA screening tool, which is non-intrusive, affordable, and time-efficient. The presented method relies on triggering autonomous responses that are indicative of the disorder. The accuracy of the proposed OSA classification system has been validated in a UCLA IRB approved study and the results demonstrate considerable potential in applying the k-NN classifier to data obtained from our Android application to help a sleep specialist in the initial assessment of patients with suspected OSA. The affordability and non-nocturnal nature of the tool will make OSA screening accessible to a larger population, leading to an improved detection of this notoriously under-diagnosed sleep disorder and the reduction of health care costs.

3.3 A Sex-Aware Framework for the Daytime Detection of Obstructive Sleep Apnea

To improve the accuracy of the daytime screening method described in Section 3.2.2, we extend it by leveraging the insights we gained from the experiments we conducted. From the experiments in that section, we conclude the following:

- The physiological category of features plays a significant role in the classification of OSA.

- From the fact that a simple k-NN classifier outperforms SVM alludes to the existence of different clusters within the same class.

Based on these two conclusions, we included more physiological features in the analysis by adding a pulse oximeter as an additional sensor, as well as by adding a couple of new autonomous challenges to the protocol described in Section 3.2.2.1. In addition to the Valsalva challenge, which was described in Section 3.2.2.1, this work extends the protocol by adding a breath hold challenge and a stroop challenge. These three challenges have been shown to elicit clear cardiovascular responses in the human body and are therefore ideal for assessing the state of the brain that signaled those responses. Also, the added pulse oximeter enables us to extract more features from all three challenges that were not possible before with the BPM as the only data source. The pulse oximeter provides a blood oxygen saturation (SpO_2) signal, as well as a continuous heart rate signal, which can be derived from the plethysmographic waveform generated by the pulse oximeter.

Furthermore, to address the clustered nature of the data, this work utilizes unsupervised learning and clustering techniques to find intrinsic hidden structures in the data. Our analysis shows that, within each class of subjects, there exist clusters that happen to correspond to the two different sexes. Based on this finding, we designed a different classification model for each of those clusters, rather than one global model like we did in our previous work.

3.3.1 Related Works

In an effort to reduce cost, waiting time, and the discomfort associated with the traditional OSA diagnostic method, a plethora of OSA detection methods have been proposed in the literature. Some techniques seek to detect sleep apnea acoustically by analyzing breathing and snoring sounds ([YM09], [AGK12]). Others use respiratory effort to detect if a subject has stopped breathing [ANK00], [HXL13]). Other techniques rely on brain waves (EEG) and sleep stage information ([MSK09], [SHL14]), heart rate ([CHS03], [SZE00], [MPI00]) , or oximetry ([RRK14], [MHA09]), or a combination of these signals to detect OSA.

Despite this wide variation in detection methods, the aforementioned techniques are all nocturnal detection methods predicated on identifying the presence of OSA by detecting apneic episodes. Since apneas only occur during sleep, these methods must be performed overnight while the subject being tested for the disorder is sleeping. Existing methods that attempt to detect sleep apnea during wakefulness have considerable drawbacks. In [LRH09], the authors use ultrasound imaging to measure the width of the subject's tongue, based on a relationship between OSA severity and the width of the tongue base; however, no results to support that claim were shown. In [SMK93], ultrafast MRI was used to localize sites of obstruction in the pharyngeal airway. While the technique achieved good results (87%) when performed on a sleeping subject, it performed worse than a random guess (47%) when conducted during wakefulness. Another study tried to predict OSA during wakefulness by using EEG and pupil size as predictors ([LPL08]), based on sleepiness affecting the beta and theta brain wave patterns and the pupil size. However, the test relies on OSA patients exhibiting signs of excessive daytime sleepiness at the time of the test, which is therefore recommended to be performed 12 hours after the subject's mid-sleep period to increase the chance of sleepiness occurring. If sleepiness does not occur, the test is inconclusive. In [MM11], breath sounds resulting from airflow through the nose and mouth were recorded using a microphone placed over the suprasternal notch of the trachea in supine and upright positions. Based on the features extracted from the digitized acoustic signals, the authors pictorially showed that the extracted features achieve good separability between the OSA and non-OSA classes, but again, no quantitative classification results were provided. The method in this work addresses shortcomings in these existing approaches.

The main contributions of this work are:

- An OSA screening framework that can identify OSA patients while they are awake using a novel technique that triggers brain responses that are indicative of the OSA disorder. The technique does not rely on a specific time or patient condition (e.g. sleepiness) and does not disrupt the patient's sleep like the nocturnal methods.

- The screening method is simple, non-intrusive, and requires no specialized personnel. It can be performed in a home environment or at the physician's office.

- The diagnosis accuracy is improved by taking into account sex-related differences, which none of the existing diagnosis techniques consider.

3.3.2 Extended OSA Screening Protocol and Data Collection Framework

The OSA Screening protocol used in this work is an extension of the protocol explained in Section 3.2.2. More specifically, phase IV of the protocol shown in Figure 3.1 was extended to include the two new challenges. Phase IV is the core of the screening tool, which visually and acoustically guides users through a series of tasks, which we call challenges. These challenges have been shown to elicit clear cardiovascular responses in the human body and are therefore ideal for assessing the state of the brain that signaled those responses. Each challenge is preceded by a baseline period that is used as a reference point when analyzing the changes that happen during the challenge. Each challenge is also followed by a recovery period, which is designed to relax the subject and prepare him/her for the next challenge. The three challenges and the cardiovascular responses they trigger are described below. As shown in Figure 3.4, in this extended protocol, a pulse oximeter continuously measures heart rate (3Hz), blood oxygen saturation (SpO_2) (3Hz) and photoplethysmograph waveform (75 Hz) starting at phase III until the end of the 20-minute screening protocol. At the end of the protocol, the smartphone transmits all the data it received from the devices over Bluetooth to a secure HIIPA-compliant server for data storage and processing.

Challenge I: Valsalva Maneuver The Valsalva maneuver is usually used as a test of cardiac function. The maneuver is performed by forcefully blowing against a closed airway. In our study, the subject is provided with a Valsalva device, which is a tube that is closed on one end and connected to a pressure measuring device. Using the Valsalva device, the subject is instructed to exhale against the tube for 18 seconds at a pressure of 40 mmHg, which is indicated by a blue LED on the Valsalva device. The Valsalva maneuver has been shown to result in a distinct pattern of heart rate changes that is mainly characterized by rapid tachycardia (increase in heart rate) during exhalation, an undershoot after release, and a slow but steady increase towards the baseline heart rate during the recovery period, as can be seen in Figure 3.5. Studies have shown that the

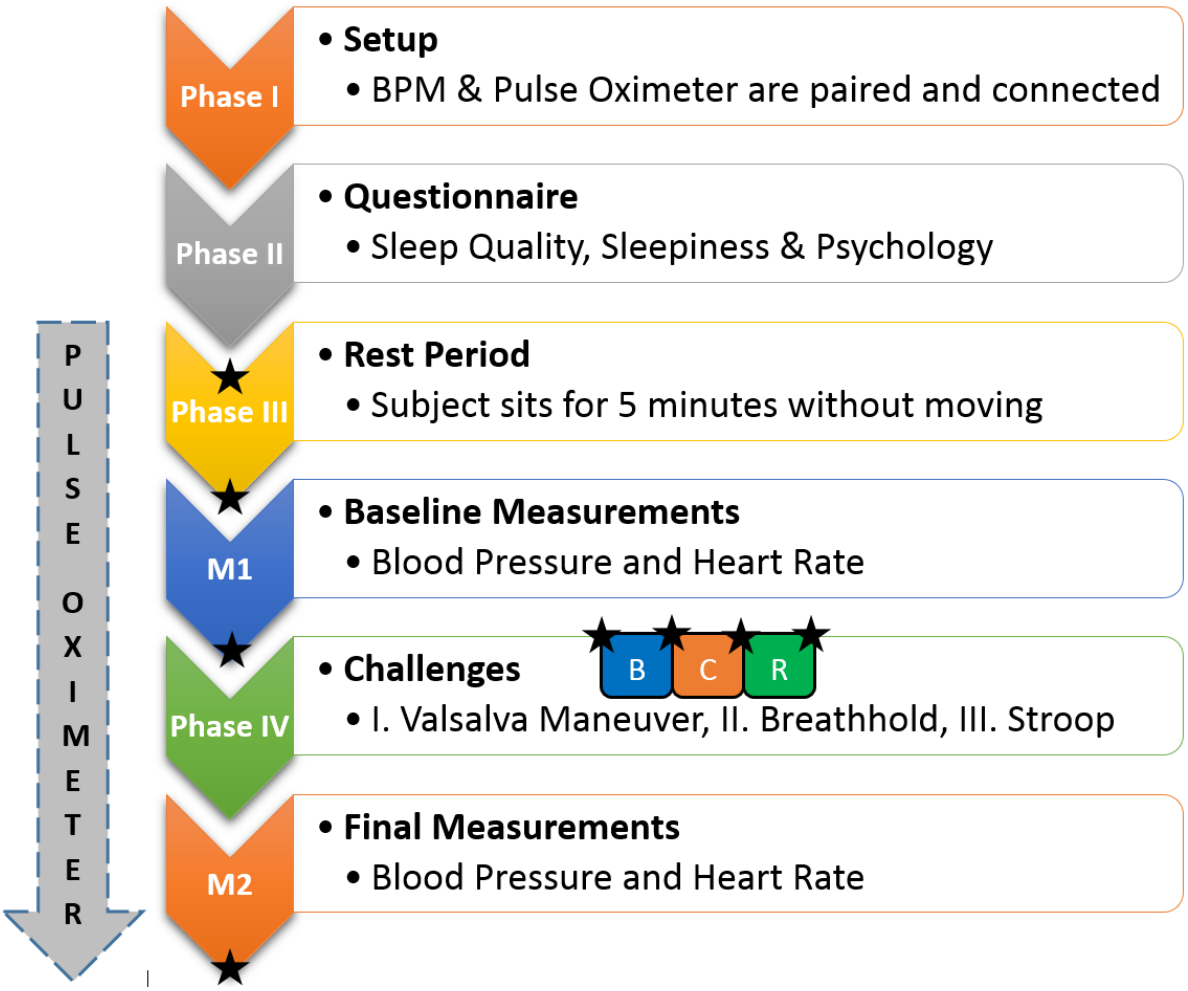

Figure 3.4: The extended protocol implemented by the data collection system. The extension includes two additional challenges in phase IV and continuous pulse oximeter measurements starting at phase III.

effect of the Valsalva maneuver is less pronounced and delayed in OSA patients than in healthy controls [MKW13], reflecting the inability of the injured OSA brain to raise the heart rate to the same extent as a healthy subject's during the exhalation phase or lower it to the same extent in the release phase, as shown in Figure 3.5.

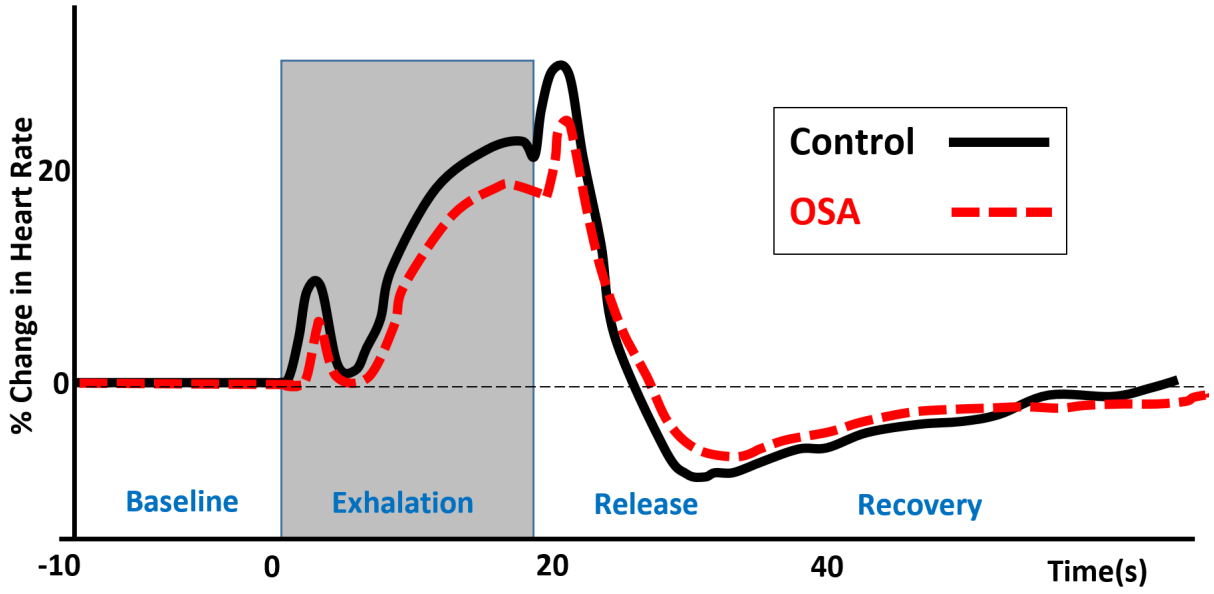

Figure 3.5: The heart rate response to the Valsalva maneuver in Control and OSA subjects. The x-axis shows the time relative to the start of the Valsalva maneuver and the y-axis shows the % change in heart rate.

Challenge II: Breath Hold During the breath hold challenge, the subject is instructed to hold his/her breath for 30 seconds. The heart rate response to the breath hold challenge is opposite to that of the Valsalva maneuver. The response is characterized by bradychardia (slowing of the heart rate) during the breath hold, an overshoot after release, and a steady decrease towards the baseline heart rate. Again here, we expect to see a less pronounced cardiovascular response in the OSA patients.

Challenge III: Stroop The stroop challenge is a test in which the subject is presented with the name of a color that is printed in a different color than the name indicates (e.g. the word blue is printed in green). In this study, we have a stroop baseline phase, in which the name and the color match (congruent stroop), followed by the challenge phase with mismatched names and colors (incongruent stroop). Subjects have 2 seconds to

name each color that appears on the phone screen. Studies have shown that the stroop test is characterized by heightened heart rate levels as a result of the stress it creates [RB97]. Again, because stress is strongly linked to autonomic nervous system functions, including heart rate and BP regulation, we expect to see an altered heart rate response to stress in OSA patients.

3.3.3 OSA Analytics and Prediction Framework

21 subjects (15 controls and 6 OSA) were enrolled in a UCLA Institutional Review Boards-approved pilot study. All controls are healthy without any major illnesses or diagnosed cardiovascular disease. The OSA patients had been diagnosed with sleep apnea using gold standard PSG, with AHI values ranging from 8 to 55. All subjects were given kits consisting of a Bluetooth-enabled blood pressure monitor (BPM), a Blueooth-enabled pulse oximeter, a Valsalva device, and a smartphone preloaded with our OSA screening protocol and data collection application. Each of the 21 subjects performed the screening protocol described in the previous section every day for a period of 42 days, for a total of 882 observations. In this section, we describe how these data are used to make a diagnosis.

3.3.3.1 Feature Extraction

After preprocessing the data and excluding flawed recordings and outliers, we perform feature extraction. From the questionnaire responses, we extract a total of 15 sleep quality, sleepiness and psychological features, which we described in our previous work. In this work, we focus on the new physiological features that reflect brain functionality. We developed a peak detection algorithm in order to extract the instantaneous heart rate signal from the plethysmographic waveform provided by the pulse oximeter. Instantaneous heart rate is used in all heart rate-related feature computation. Some of the features are standard features and some were empirically derived. The most prominent features derived from each challenge are presented below:

Valsalva The purpose of the features extracted from the Valsalva challenge is to capture the differences in both the magnitude and the timing of the autonomic response

between OSA and control. To capture the magnitude or extent of heart rate increase and undershoot during the Valsalva maneuver, we compute the tachycardia ratio and the Valsalva ratio, respectively. The tachycardia ratio measures the heart rate increase between baseline and the peak of the challenge, while the Valsalva ratio measures the extent of the heart rate undershoot in the release period, compared to the peak during the challenge. The area-under-the-curve feature captures the sustained increase in heart rate during the challenge; the increase curve is usually blunted in OSA patients and would result in a smaller area under the curve. To express the timing of the heart rate increase, which is usually delayed in OSA patients, we calculate the slope of the increase curve. These four feature are shown in Figure 3.6.

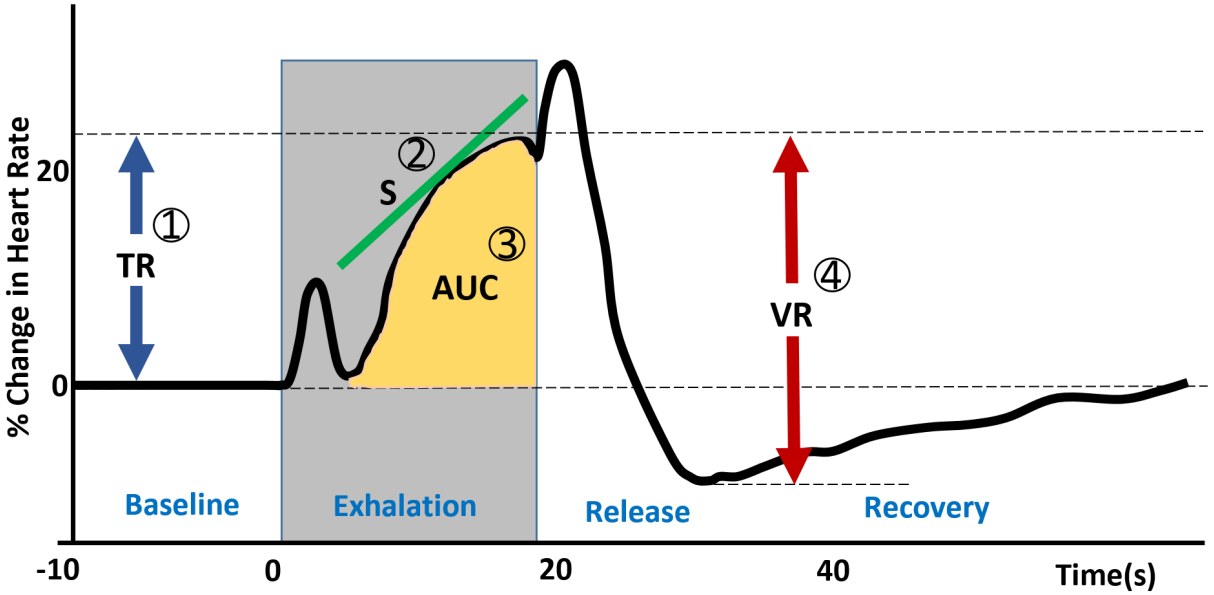

Figure 3.6: Features of the Valsalva challenge. 1. TR is the tachycardia ratio: (peak HR during challenge/baseline HR), 2. S is the slope of the HR increase curve during the challenge, 3. AUC is the area under the HR increase curve, 4. VR is the Valsalva ratio: (peak HR during challenge/minimum HR during undershoot)

Breath Hold From the breath hold challenge, we extract the breath hold ratio feature, which is similar to the valsalva ratio described above, except the heart rate decreases in this case. The breath hold ratio captures the magnitude of heart rate decrease during the challenge, compared to the baseline heart rate. Another feature we extracted is the difference in blood oxygen saturation (SpO_2) between the baseline period and the

challenge period. Because of their disorder, OSA patients are more susceptible to oxygen desaturation in response to pauses in breathing. This feature, therefore, captures the rapid blood oxygen desaturation characteristic of OSA.

Stroop The features extracted from this challenge are the difference in heart rate between baseline and congruent stroop, and between baseline and incongruent stroop.

Other features used include average resting heart rate (usually higer in OSA), average resting SpO_2 (usually lower in OSA), and heart rate variability (HRV). HRV is a global indication of how capable the brain is of responding to cardiovascular challenges by decreasing and increasing the heart rate. The higher the HRV, the more responsive the brain is to external stimuli. We use two measure of HRV: the standard deviation and the interquartile range (IQR).

3.3.3.2 Sex-Aware Feature Selection And Classification

Though OSA is more prevalent in males, studies have shown that it has a more detrimental effect on females. Females showed a more pronounced heart rate response impairment (lower amplitude, delayed onset, and slower rate changes) than males [MKW13]. Given the different effects that OSA has on members of the different sexes and the different resulting heart rate response patterns, designing a single global model for diagnosing OSA patients of both sexes would be insufficient. In this work, we build two different prediction models for male and female subjects and perform features selection for each group separately. To make a diagnosis, we first cluster a subject into his/her appropriate sex group and then perform classification using the specific model for the subject's sex. This sex-aware classification achieves more superior results than the sex-agnostic model. We found that the Bayesian Network classifier achieved the best results for the male cluster, while the k Nearest Neighbor (kNN) classifier worked the best for the female group.

A Bayesian Network is a probabilistic graphical model in which the feature space is represented as a directed acyclic graph (DAG). The nodes of the DAG correspond to the features, while the edges of the DAG represent the influence of one feature on another. Each node is annotated with a conditional probability distribution that represents $p(X_i|Pa(X_i))$, where X_i is a feature and $Pa(X_i)$ is the parent of that feature in the

DAG. Probabilistic inference can then be used to make predictions about the class that a feature vector belongs to.

The kNN classifier is a type of lazy learning that defers all computation till classification. In the training phase, a kNN classifier merely stores the feature vectors and class labels. In the testing phase, it assigns a testing data point to the class label of the majority of its k nearest neighbors based on a distance metric. In this work, k=3 and the distance metric is the Manhattan distance.

To avoid pseudo-replication, we perform Leave-One-Subject-Out Cross-Validation (LOOCV), where all observations for one subject are used for testing and the remaining observations are used to train the gender-specific models.

3.3.4 Results And Discussion

To evaluate the merit of a sex-aware model for the diagnosis of OSA, we show the Receiver Operating Characteristics (ROC) curves for the sex-agnostic and the sex-aware OSA classifications in Figure 3.7. The classifier that achieved the best results for the entire dataset (male + female subjects) was the Random Forest classifier. The Random Forest classifier is an ensemble learning method that constructs multiple decision tree predictors at training time. Each tree casts a vote about the class label for a specific feature vector and the most popular class is assigned to the feature vector at testing time. As can be seen in the figure, both the sex-agnostic and sex-aware models achieve excellent prediction accuracy (AUC >0.97) as a result of the novel screening protocol and the extracted physiological features. However, The importance of a sex-aware approach becomes clear when comparing the sensitivity of the classification. Table 3.3.4 shows a more detailed evaluation of the sex-aware and sex-agnostic approaches. While the sex-aware models were successful at detecting OSA patients over 90% of the time, the sex-agnostic model was only successful ≈74% of the time, making the sex-aware approach more suitable for a medical diagnostic test. On closer investigation, we found that more than 90% of the OSA instances that the sex-agnostic approach misclassified as Control are male subjects. A possible explanation for this is that, when one model is used for males and females, the male OSA patients appear healthy compared to their female counterparts who, as stated earlier, get affected more seriously by the disorder. Again, this shows the need for

a sex-aware approach.

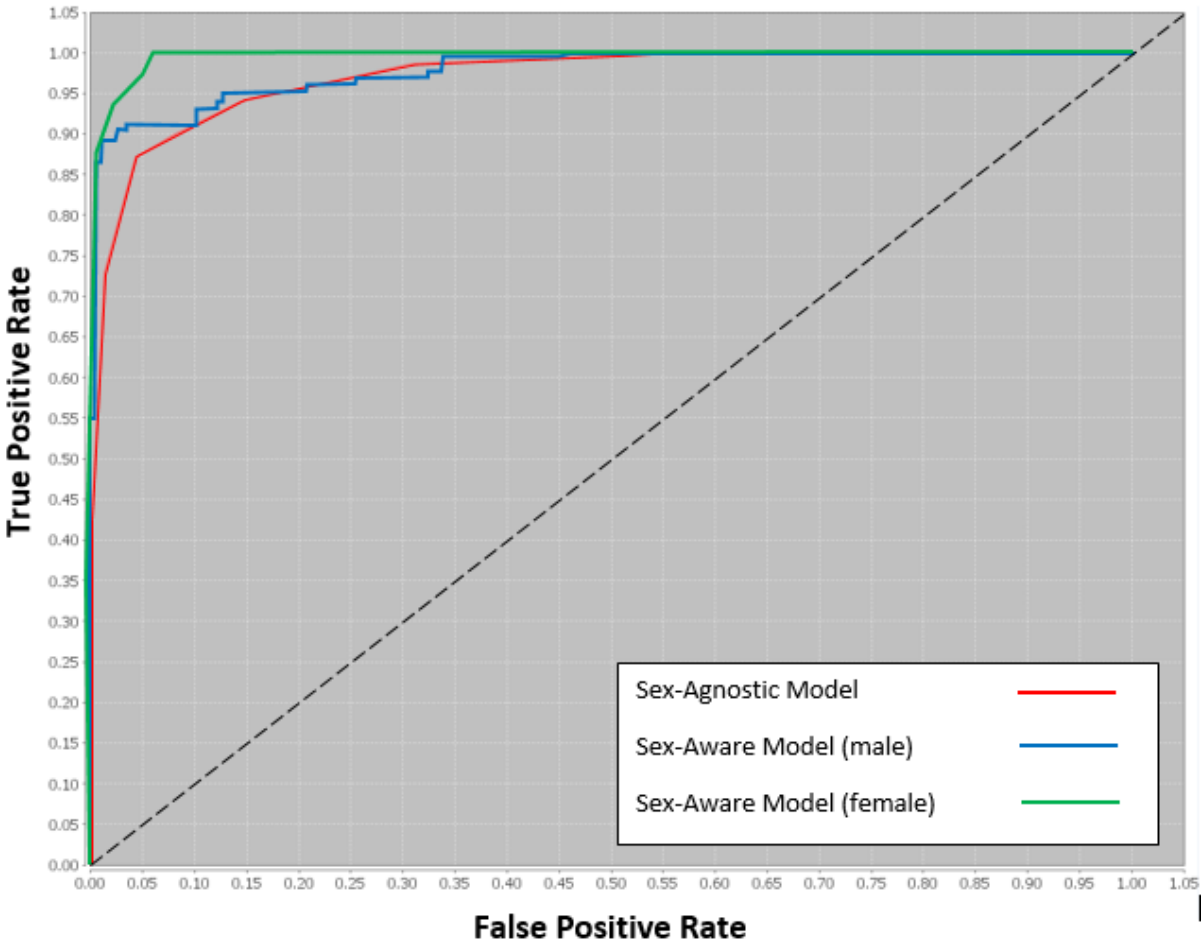

Figure 3.7: ROC curves comparing the prediction accuracy of the gender-aware models vs. the single gender-agnostic model.

3.3.5 Conclusion

In this work, we proposed a novel OSA screening method that can identify OSA patients while they are awake without needing to monitor them during the night and disrupting their sleep quality. Our approach relies on presenting the subjects with a sequence of autonomic challenges that are designed to trigger specific cardiovascular responses. When analyzed, these responses are indicative of brain alterations caused by the OSA disorder. We also showed that a sex-aware classification approach performs much better than a sex-agnostic model at detecting OSA patients. The proposed diagnosis framework offers an affordable, and noninvasive screening tool for the prediction of OSA during wakefulness.

Model	Sensitivity	Specificity	Area Under ROC	Overall Accuracy
Sex-Agnostic	**0.736**	0.989	0.972	91.8 %
Sex-Aware (M)	**0.915**	0.968	0.977	94.8 %
Sex-Aware (F)	**0.933**	0.984	0.995	97.7 %

Table 3.3: Classification results of the sex-agnostic and sex-aware approaches

3.4 A Daytime Obstructive Sleep Apnea Severity Assessment Framework

Section 1.2.5 demonstrated the different treatment methods for the OSA disorder. Since the therapeutic approach suitable for treating a specific patient largely relies on the severity of the disorder, it is important for a diagnostic or screening tool to assess the severity of the condition, as opposed to just its presence or absence. Assessing the severity of the OSA condition can also enable a more optimal utilization of scarce PSG resources by making it possible for patients with severe OSA to be more efficiently scheduled for a PSG and to be given priority over other patients with a milder degree of the disorder. A screening method that can provide severity analysis would also provide more detailed results for the patient being screened and the severity of the condition might encourage patients to go for a full PSG and get diagnosed. Furthermore, based on the guidelines of the American Society of Anesthesiologist (ASA), patients undergoing surgery should be screened for OSA to prevent possible perioperative (preoperative, intraoperative, and postoperative) complications, which increase with the severity of the disorder [GBC06]. Since the screening needs to be done in a timely manner prior to surgery, the daytime screening method proposed in Sections 3.2 and 3.3 would be ideal since it enables fast screening without the need for a lengthy overnight study. In this work, we extend our daytime screening tool to identify the different OSA severity levels. This is a challenging task for two main reasons:

- Unlike our previous work, the problem of severity assessment is a multi-class classification problem that is significantly more challenging than the binary classification problems investigated so far.

- In sleep medicine, the severity of sleep apnea is determined by the Apnea Hypopnea

Index (AHI), which is the number of apneas or hypopneas recorded during PSG per hour of sleep. Based on the AHI, the severity of OSA is classified as follows:

- None/Minimal: AHI <5/hour
- Mild: AHI $\geq$ 5/hour and <15/hour
- Moderate: AHI $\geq$ 15/hour and <30/hour
- Severe: AHI $\geq$ 30/hour

Because the AHI is, by definition, a nocturnal measure of the severity of OSA, all the methods for the prediction of OSA severity available in the literature are nocturnal methods. In this work, the goal is to use our daytime diagnosis method to predict an inherently nocturnal measure.

3.4.1 Related Works

In an effort to overcome some of the limitations associated with PSG, research focused on alternative diagnostic and severity identification methods has notably increased. New techniques for OSA severity detection usually rely on the analysis of a reduced set of data compared to PSG.

[MDV03] and [MHA12], for example, use only overnight oximetry to predict the AHI. Because apneas and hypopneas are usually accompanied by oxygen desaturation due to the lack of airflow, OSA patients tend to present unstable blood oxygen saturation (SpO_2) signals. In [MDV03], time-domain features like the number of O_2 desaturations below a certain threshold and the cumulative time spent below the normal saturation level were used as severity predictors, while in [MHA12], frequency-domain features like the power spectral density (PSD) of the oximetry signal were used in addition to time-domain features to capture signal fluctuations due to repetitive apneas.

In [BTZ12] and [DTZ13], the signal of choice was an indirect measure of snoring obtained overnight from a directional condenser microphone. The biological instability of the upper airway structure across the night, especially during obstructive events, is thought to change the snoring characteristics of apneic snorers compared to simple snorers. Acoustic features representing intra- and inter-snore qualities were used to capture

these characteristics and assess OSA severity.

In [SN14], disposable adhesive sensor patches embedded with ECG electrodes, a tri-axial accelerometer and a Bluetooth Low Energy (BLE) module were used to collect ECG and actigraphy data from sleeping patients. Heart rate, respiratory, posture and movement features were extracted from 150-second sleep epochs and a Support Vector Machine (SVM) classifier was trained to detect the presence/absence of apneas/hypopneas for each epoch.

By reducing the number of sensors and the set of data that need to be analyzed compared to a full PSG, and by eliminating the need for costly, complex equipment, these methods achieve cheaper, more comfortable, and more portable early-stage OSA severity screening that can significantly reduce screening wait times and increase access to OSA diagnosis.

Despite the variation in the type and number of analyzed signals and extracted features, two overarching themes characterize all the aforementioned screening methods. First, they all rely on signals obtained overnight. The nocturnal nature of the methods is expected since sleep is the time during which the obstructive events manifest and since AHI is, by definition, a nocturnal measure. However, the dependence on sleep makes nocturnal methods dependent on the subject's sleep quality, which might be affected by a variety of factors other than the OSA disorder. Also, the time required for a nocturnal screening method to provide severity results is not practical for perioperative risk stratification. Second, all these methods first use regression analysis to estimate AHI values, and then use the predicted values to categorize patients into the corresponding severity group. This way, the assignment of patients into the correct severity category is as accurate as the AHI prediction can be.

In the next section, we describe our *daytime* severity screening framework, which relies on measurements obtained *exclusively while subjects are awake* without the need for a lengthy overnight test. In addition, our method first solves the simpler problem of assigning a subject into a coarse-grained OSA severity category using multi-class classification, and then uses regression analysis to predict that subject's AHI within the assigned category.

3.5 OSA Severity Screening Framework

It might seem counterintuitive that a disorder that only manifests during sleep can be screened for while a subject is awake. While it is true that an OSA patient won't experience any disordered breathing while awake, studies have shown that OSA is associated with chronic physiologic changes that persist during wakefulness. One possible mechanism for sustained physiologic changes is hypoxic brain injury. The lack of airflow during apnea episodes and the accompanying hypoxia prevents oxygen from reaching the brain. As a result, brain tissue is damaged [MKW08], and the damage has been shown to target central autonomic regulatory regions of the brain, which are responsible for regulating cardiovascular functions, like blood pressure and heart rate.

Based on these findings, we implemented an application that is designed to trigger specific cardiovascular responses in the subject under test. The magnitude and delay of the subject's autonomic response reveals the presence/absence of brain dysfunction, which is related to the underlying brain injury. [MKW13]. The more impeded the response, the more severe the brain damage and the OSA condition. In our earlier work, we used this application to collect data from which we can extract features that capture this modified cardiovascular response and designed a binary classifier that can distinguish between healthy subjects and OSA patients using those features. In this work, we use the same data collection software but extend our earlier work to a multi-class classification problem to distinguish different levels of OSA severity. Sections 3.2.2.2 and 3.3.3.1 have a detailed description of the software architecture, the monitored signals, as well as the computed features used for our earlier work.

3.5.1 Screening Protocol and Data Collection Framework

In addition to the 28 features described in 3.2.2.2 and 3.3.3.1, we compute four more features for the severity classification. The four additional features are BP level, lowest resting HR, lowest resting SpO_2 and Body Mass Index (BMI), which are all correlated with OSA severity, as shown in [WLR14], [CLG12], [ZMJ14] and [VBD07], respectively. The BP level, lowest resting HR and SpO_2 are all obtained during the rest period portion of the application. We use the five BP ranges recognized by the American Heart

Association as values for the BP level feature [Ame].

3.5.2 Two-Layer Analytics and Severity Prediction Framework

The second component of our severity screening framework is the prediction unit, which we further divide into two subtasks. The first subtask is to categorize subjects into one of the 4 OSA severity groups. As mentioned in Section 3.4, the severity groups are defined by the AHI and are categorized as follows: healthy (AHI $<$ 5 events/hr), mild (5 $\leq$ AHI $<$ 15), moderate (15 $\leq$ AHI $<$ 30), and severe (AHI $\geq$ 30). Unlike these broad severity categories, AHI can provide a measure of progress during the treatment stage. Therefore, the second subtask of severity prediction is to also estimate the AHI value for a subject. In the next sections, we describe how we approach those two subtasks.

3.5.2.1 Layer 1: Coarse-Grained Severity Prediction

Unlike past research surveyed in Section 3.4.1, we solve the coarse-grained problem of predicting the categorical severity classes first, and use that knowledge to then solve the more fined-grained problem of predicting AHI values within each severity class. To solve the coarse-grained prediction problem, we designed a multi-class classifier that consists of multiple binary classifiers. For each pair of severity classes, we trained a different binary classifier to accurately distinguish between them. Due to the insufficient number of mild OSA subjects, we had to eliminate the mild class from our analysis, but the methodology could be applied to a larger dataset with more low severity patients. Leaving out mild subjects is still consistent with common medical practice since healthy and mild cases are often symptomatically indistinguishable. In fact, it is a controversial topic whether mild OSA should be treated or not [Bro07][Lit07].

With the healthy, moderate and severe classes remaining, we built a healthy-vs.-moderate (HvM), a healthy-vs.-severe (HvS) and a moderate-vs.-severe (MvS) model. The HvM and HvS models are SVM classifiers. SVM is a binary classifier that constructs an optimal hyperplane in the feature space to separate the samples belonging to two different classes. The optimal hyperplane is the one that maximizes the distance between the nearest training samples of the two classes. Once the hyperplane is constructed, the

testing data can be projected onto the feature space and classified by the hyperplane into one of the two classes. In this study, we used a third degree polynomial kernel function with a cost of 1.1 and an intercept constant of 0.1 for both models.

The MvS model, on the other hand, is a Decision Tree classifier (DTC). A DTC is a hierarchical structure that consists of directed edges, as well as internal and leaf nodes. The internal nodes contain conditions that test a certain feature and separate the input samples based on those test conditions until a final decision about class membership is made in the leaf nodes. Once the tree is constructed, new samples can be classified by starting from the root node of the tree and applying the test condition to the sample and following the appropriate branch based on the outcome of the test until a leaf node is reached. We used the Gini impurity as the measure split quality.

After the 3 pair-wise binary classifiers are trained, they are combined to form a multi-class severity classifier. Fig. 3.8 (Layer 1) shows how a subject is classified into one of the 3 severity classes. First, the feature vector is computed from the collected data and then fed into each of the 3 classifiers. Each binary classifier assigns the test sample to one of the 3 severity classes. If two out of three classifiers agree on a severity class, the combined classifier predicts the class that got the most votes. In the case where there's no majority and each binary classifier predicts a different severity class, the final decision is made based on the quality, rather than the quantity of the votes. The quality of a classifier's vote is determined by the classifier's confidence about its prediction. For the SVM classifiers, we use the distance of the test sample to the separating hyperplane as the confidence score. The further the distance of the sample from the hyperplane, the more confident the classifier is about the class membership. For the DT classifier, the confidence score is the probability of the test sample belonging to the predicted class. In case of a tie, the decision of the classifier with the highest confidence score is chosen.

3.5.2.2 Layer 2: Fine-Grained AHI Prediction

After assigning a test sample to one of the broad severity classes, we further predict the AHI of the sample within its severity class. We solve this fine-grained prediction problem using regression analysis. A regression model fits a curve to data points such that the differences in the distances of data points from the curve are minimized. In this case, that

best-fit curve is expressed by a function that, given the features as independent variables, outputs the AHI as the dependent variable. We used an ElasticNet regression model to predict the AHI for an input sample. ElasticNet is a hybrid regression approach that penalizes both the L1 norm and L2 norm in its objective function. It is especially useful when the input contains multiple correlated features.

For the AHI prediction step, we built two different ElasticNet regression models for each of the OSA classes (moderate and severe). If, in the coarse-grained prediction step, one of those classes is predicted, the corresponding sample is fed into the appropriate regression model to predict the AHI, as shown in Fig. 3.8 (Layer 2). To further improve the AHI estimation, we apply a heuristic that takes advantage of the coarse-grained classification results. The heuristic adjusts the AHI predicted by the regression model (AHI_{pred}) by a factor in the range $[0.8, 0.9]$ or $[1.1, 1.2]$, depending on the direction of the desired adjustment, i.e. whether AHI_{pred} should be decreased or increased. The direction of the adjustment is determined by the class that lost the majority vote in Layer 1 (C_{lost}). If C_{lost} is a more severe class than the final predicted class (C_{final}), the adjustment factor will increase AHI_{pred}. If C_{lost} is less severe than C_{final}, the adjustment factor decreases AHI_{pred}. The magnitude of the adjustment is determined by the probability of C_{lost}. The final adjustment factor is calculated as $1 \pm (P_{np} \times \alpha)$, where the sign is the direction of the adjustment, P_{np} is 1 minus the probability of C_{lost}, and α is an empirically determined value that depends on the severity and probability of C_{lost}. The more severe and more probable C_{lost} is, the bigger the adjustment. Also, since both the direction and magnitude of the adjustment depend on C_{lost}, an adjustment is only made if C_{lost} was predicted with a probability ≥ 0.9. Fig. 3.8 clarifies how this heuristic works by showing an example adjustment.

3.6 Results and Discussion

To evaluate the accuracy of our prediction framework with its coarse-grained and fine-grained components, data were collected from 24 subjects (16 healthy, 8 OSA) every day for 42 days using our Android application. For every day's data, we calculate our daytime features described in Section 3.5.1. Since the data includes 42 observations

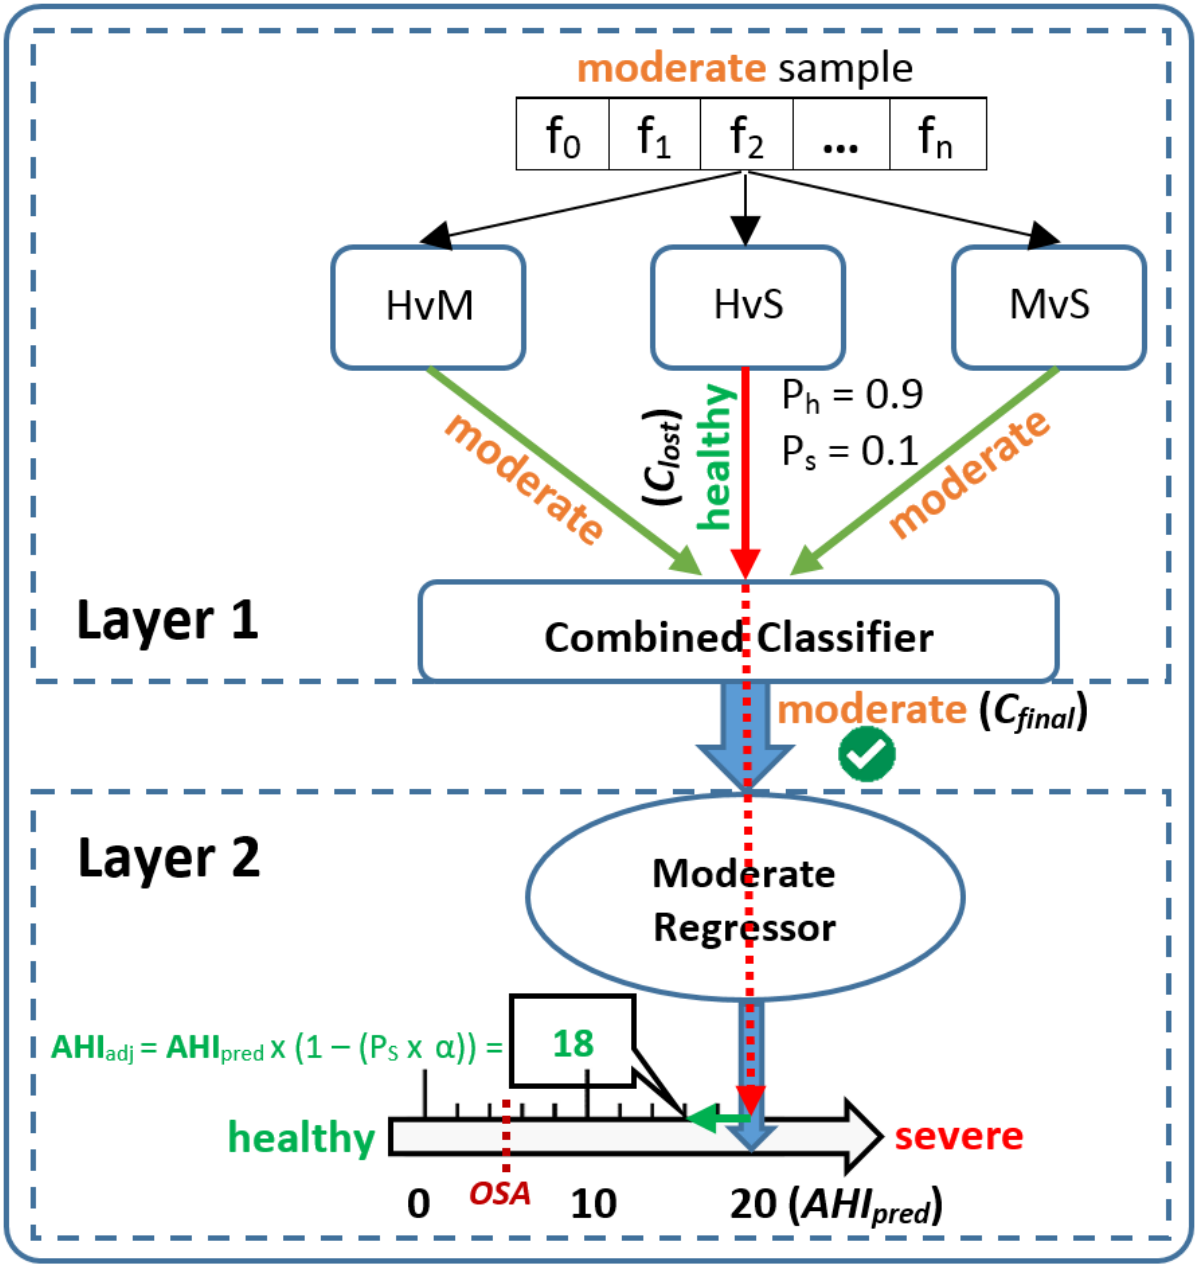

Figure 3.8: AHI Adjustment Heuristic: In this example $\alpha = 1$

per subject, we perform Leave-One-Subject-Out-Cross-Validation (LOOCV), where all observations for one subject are used for testing and the remaining observations are used for training, to avoid pseudo-replication. The binary classifiers in the coarse-grained prediction step are trained using observations from only the pair of classes the classifier is meant to distinguish between. Each of the two regression models in the fine-grained prediction step were trained using observations from one severity class. Feature selection was performed separately for each classification and regression model and features were z-score standardized.

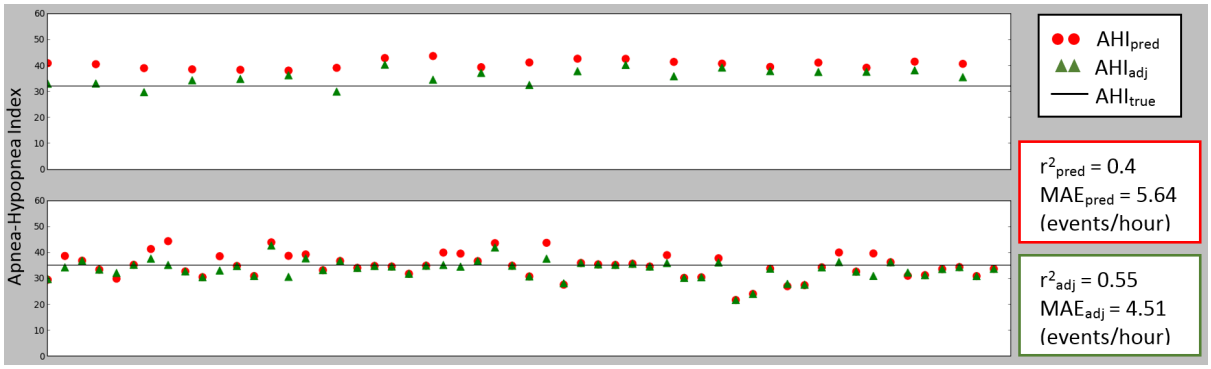

Figure 3.9: Predicted AHI and Adjusted AHI For 2 Severe OSA Patients

Table 3.6 shows the classification results for the 3 pair-wise binary classifiers and the final combined severity classifier. The MvS classifier has the lowest sensitivity out of the 3 binary classifiers. This is expected since the moderate and severe classes are both OSA classes and are more similar to each other than the healthy class is to either of them. The combined classifier has high accuracy and only misclassifies a few moderate instances as severe.

Classifier	**Sensitivity**	**Specificity**	**Accuracy**
HvM	100%	100%	100%
HvS	99.4%	100%	99.5%
MvS	71.1%	100%	99.4%
Combined	undef	undef	99.6%

Table 3.4: Classification results of pair-wise binary classifiers and final multi-class severity classifier

Fig. 3.9 shows the AHI values predicted by the ElasticNet regression model, as well as the AHI values adjusted by the heuristic described in Section 3.5.2.2. As can be seen in the figure, the heuristic brings the predicted AHI values closer to the true AHI, which reduces the mean absolute error (MAE) from ≈ 5.6 to ≈ 4.5 events/hr on average. The small mean absolute error is due to the accuracy of the coarse-grained classifier. The coarse-grained classifier predicts the severity class for a specific sample, and then that

sample is fed into the regression model belonging to the predicted severity class. This order guarantees that the predicted AHI will stay within the range of values for the predicted severity class. If the fine-grained AHI prediction were to be performed first, the predicted AHI values would show a larger variation, and relying on those values for categorizing patients into the severity groups would decrease the classification accuracy.

Table 3.5 shows the most important features used in the severe regression model. It is interesting to see that four of our daytime physiological features appear in the list of top five important features, while only one questionnaire feature, namely sleepiness, made the list. A similar observation was made for the moderate regression model, where 3 out of 5 features were physiological. This shows the value of our daytime physiological features for AHI prediction and their advantage over traditional questionnaire screening.

Feature	**Description**
avg_resting_SpO_2	Average SpO_2 during rest period
avg_val_SpO_2	Average SpO_2 during Valsalva challenge
sleepiness	Daytime sleepiness on a 0-10 scale
valsalva_ratio	The extent of the heart rate undershoot in release period, compared to the peak during challenge
avg_resting_hr	Average heart rate during rest period

Table 3.5: Top-Ranked Features For Severe Regression Model (most important to least important)

3.7 Conclusion

In this work, we proposed an OSA screening method that can accurately identify OSA patients, the severity of their condition, and their AHI without any need for overnight monitoring. Our approach presents the subjects with a sequence of autonomic challenges that trigger specific cardiovascular responses. These responses are indicative of brain alterations caused by the OSA disorder. The magnitude and delay of the responses are correlated with the severity of the brain damage and the OSA condition. Features extracted from the responses are used in a two-layer prediction framework to first predict the coarse-grained severity class (overall accuracy of 99.6%), followed by a fine-grained prediction of the AHI (mean absolute error of 4.51 events/hour). The accuracy of our AHI prediction is comparable to the best existing *nocturnal* method of AHI estimation,

which achieved a mean absolute error of 3.5 events/hour [BTZ12]. Once being deployed and validated in a large-scale study, the proposed severity screening framework will offer affordable, noninvasive, and fast severity screening, resulting in more efficient PSG scheduling, wider access to diagnosis and better perioperative risk stratification.

3.8 In Closing

This chapter presented non-intrusive daytime methods for diagnosing and assessing the severity of OSA. The accuracy of our classification models has been validated in a UCLA IRB approved pilot study and the results demonstrate considerable potential in using our daytime features to design classification and regression models that provide an accurate early-stage diagnosis of OSA and an assessment of its severity. Deploying this daytime screening tool in a bigger clinical trial will prove its reliability as a cheap, fast and non-intrusive OSA screening method.

CHAPTER 4

Stress Reduction for OSA Patients With Resistant Hypertension

4.1 Introduction

Chapters 2 and 3 discussed the main focus of this research, namely to design screening methods that can serve as cheaper, more convenient, non-intrusive alternatives to polysomnography, as well as be more reliable than the existing screening questionnaires that rely mostly on self report. The goal of designing such systems is to make OSA diagnosis more accessible and affordable, so that more people can be aware of their condition and seek treatment.

In this chapter, we briefly touch on another aspect of the OSA disorder, namely its comorbidities, i.e. its accompanying illnesses, which make the disorder a life-threatening one. We specifically focus our attention on resistant hypertension, which is the most prevalent comorbidity of sleep apnea, as shown in Figure 4.1.

4.2 Stress Reduction For Resistant Hypertension Treatment

Resistant hypertension is defined as blood pressure that remains above goal despite lifestyle modification and administration of three anti-hypertensive drugs of different classes. Section 1.2.3 discussed physiological and psychological comorbidities of sleep apnea. Figure 4.1 shows the prevalence of sleep apnea in 8 different physiological comorbidities. As shown in the figure, 83% of all patients with resistant hypertension also have sleep apnea–making resistant hypertension the most prevalent physiological comorbidity of sleep apnea.

Unsurprisingly, studies have shown that treatment of the OSA disorder by the use of

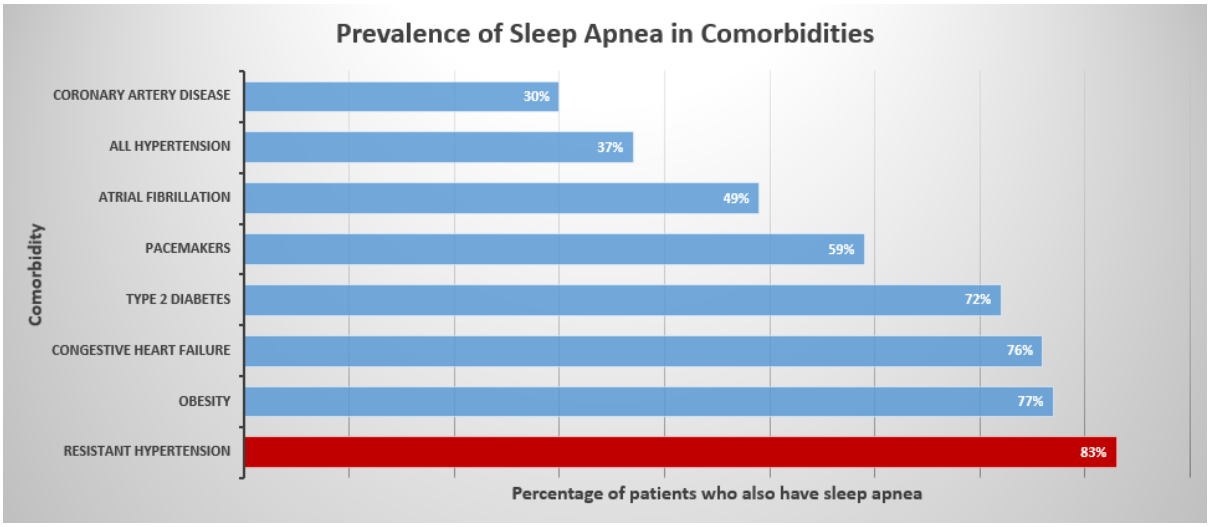

Figure 4.1: Health Conditions (Comorbidities) Associated With OSA

CPAP (see Section 1.2.5) results in a reduction of blood pressure, although the degree of that reduction remains disputed [KPO13], [PDP13]. Because of the intrusiveness of the CPAP mask, however, studies have shown that adherence to CPAP treatment is meager. When adherence is defined as greater than 4 hours of nightly use, 46 to 83% of patients with OSA have been reported to be non-adherent to treatment [WG08b]. As a result, the research community has been investigating alternative strategies for the treatment of resistant hypertension.

The fact that mental and emotional stress can raise blood pressure is long established. Many studies addressed the effects of anxiety, stress, and depression on cardiac disease in general and hypertension specifically [WB86], [KOE98], [GS98]. In this work, we investigate the other side of the equation. What effect can the treatment of stress and other emotional disorders have on resistant hypertension?

This chapter investigates the use of a smartphone protocol to guide hypertensive OSA patients through stress reduction exercises in order to achieve a reduction in hypertension. In the next sections, we discuss previous research on drug-resistant hypertension reduction, as well as our proposed method and the software system we designed to achieve a reduction in hypertension.

4.2.1 Related Works

In the medical field, there is a substantial number of studies on stress reduction and its effect on blood pressure. The studies have shown that stress reduction exercises can in fact result in a decline in blood pressure. In one study, a group of hypertensive African American men and women practiced a form of meditation called transcendental meditation. After 3 months, women showed a decline of $\approx$ 10 mm Hg in systolic blood pressure and $\approx$ 6 mm Hg in diastolic blood pressure, while men showed a decline of $\approx$ 13 mm Hg and $\approx$ 8 mm Hg, respectively [ASS96]. A similar result was shown in a study where a group of hypertensive employees received stress reduction intervention at the workplace. Three months post-intervention, the employees exhibited a reduction of $\approx$ 11 mm Hg in systolic blood pressure and $\approx$ 6 mm Hg in diastolic blood pressure [MAT04]. In [RSN07], 17 interventions employing different stress reduction techniques were surveyed and the results were consistent with the aforementioned studies.

This study differentiates itself from previous research as follows:

- To the extent of our knowledge, there haven't been any clinical trials that investigated the effect of stress reduction on **resistant** hypertension. A few medical papers state that stress reduction can be a way to reduce resistant hypertension but this claim has not been backed up by data from clinical trials.
- The aforementioned studies considered stress reduction to be successful if an average reduction in hypertension was observed among the participating patients. What these studies failed to do is investigate why stress reduction is successful for some people and not for others.

4.2.2 Stress Reduction Software System

Our stress reduction routine (SRP) is implemented as an extension of our Android application described in Sections 3.2.2.1 and 3.3.2. After the challenge phase (phase IV) of the protocol, the Android application guides hypertensive subjects through breathing awareness meditation exercises, records their blood pressure before and after the exercises, and transmits the blood pressure data to a HIPAA-compliant server.

The stress reduction routine begins by prompting subjects to take their blood pressure. The application textually and pictorially guides subjects to take a blood pressure measurement. After the measurement is complete, the application acquires the data from the blood pressure monitor over Bluetooth and transmits it over WiFi to a server for storage.

The application then presents the subjects with information about stress, the physical changes that stress sets off in humans, and the harmful health effects–physiological and psychological–that sustained stress can cause or exacerbate. In this part of the routine, subjects are also educated on what breathing awareness means and how it can help reduce stress. They are also provided with instructions on how to practice breathing awareness, including how to sit and how to focus on breathing. Subjects are also provided with answers to a number of frequently asked questions about practicing breathing awareness. The information is presented textually and pictorially to the subjects and is meant to raise awareness about the effects of prolonged stress on our health, to encourage subjects to adhere to the stress reduction routine, and to provide them with the necessary knowledge that can enable them to perform the breathing awareness exercises correctly.

After the informative section is completed, the user can start the breathing awareness exercises for stress reduction. The application presents the user with a choice between a guided and an unguided breathing awareness routine. The guided routine consists of 10 minutes of Ananda music with a voice overlay guiding subjects through the breathing awareness exercises. Once subjects are familiar with the exercises, they can choose the unguided routine, which only consists of the Ananda music without the vocal instructions.

Breathing awareness is a form of meditation that uses the process of breathing as its primary focus. During the exercises, subjects are instructed to pay attention to their breath – how shallow or deep it is, its rhythm, if, how and when it stops, etc. Since our emotional state has a direct effect on our breathing patterns–when we are upset or anxious our breathing speeds up, when we are relaxed or sleeping, it slows down–becoming aware of the breath or learning to breathe consciously can have a calming effect on the body and mind. When subjects learn to observe the fact that breathing is a rhythmic process and that it is constantly changing, it teaches them to be comfortable with change, which is an essential aspect of stress management. The Ananda music accompanying the exercises is

designed not to stir up one's emotions, which makes it ideal for meditation and being in a state of relaxation and complete focus on breathing.

After subjects are done with the breathing awareness exercises, the application prompts them for another blood pressure measurement, which is again transmitted to the server. After a subject completes the breathing awareness routine, the application displays a message reminding subjects to take 30 seconds or more to focus on their breath at times throughout the day. The next time the subject uses the application to practice breathing awareness, they will be presented with a question about how often they took some time to focus on their breath outside of the breathing awareness routine. This question is meant as both a reminder for the subjects to focus on their breathing more often, and as a way to gage how much they are benefiting from the routine and how it correlates with the goal of reducing stress and the ultimate goal of reducing their blood pressure. Figure 4.2 shows the software protocol described above.

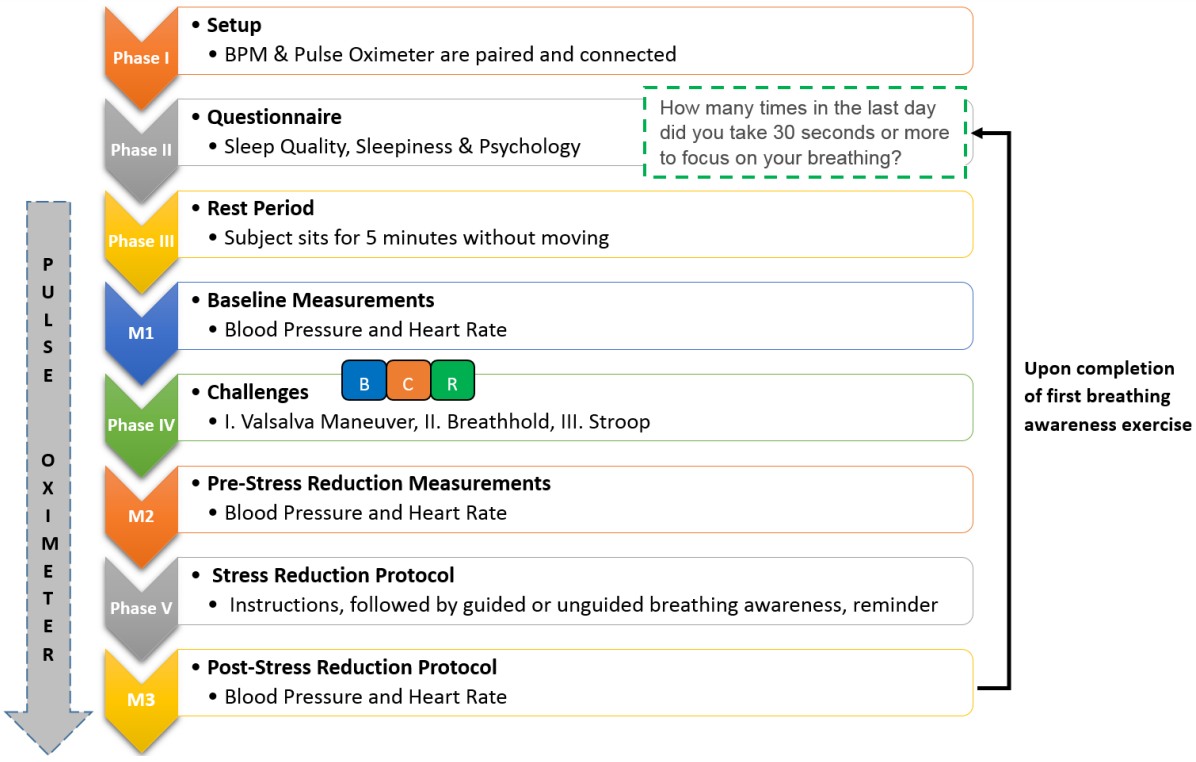

Figure 4.2: The Android protocol extended with the stress reduction routine.

4.2.3 Stress Reduction Pilot Study and Preliminary Results

We are currently conducting a pilot study in collaboration with the School of Nursing to test the effectiveness of stress reduction on lowering blood pressure. To date, one hypertensive OSA patient and three control subjects used our software application to practice breathing awareness. Each subject completed at least 28 days of the stress reduction routine following a 14-day baseline.

Figures 4.3 and 4.4 show some preliminary results of the stress reduction routine and its effect on blood pressure for 3 control subjects and one hypertensive OSA patient. Although we cannot draw any conclusions from the data about the success of stress reduction due to the small sample size, it is interesting to see that in Figure 4.4, the trend of increase in systolic blood pressure is attenuated after the patient begins the stress reduction routine.

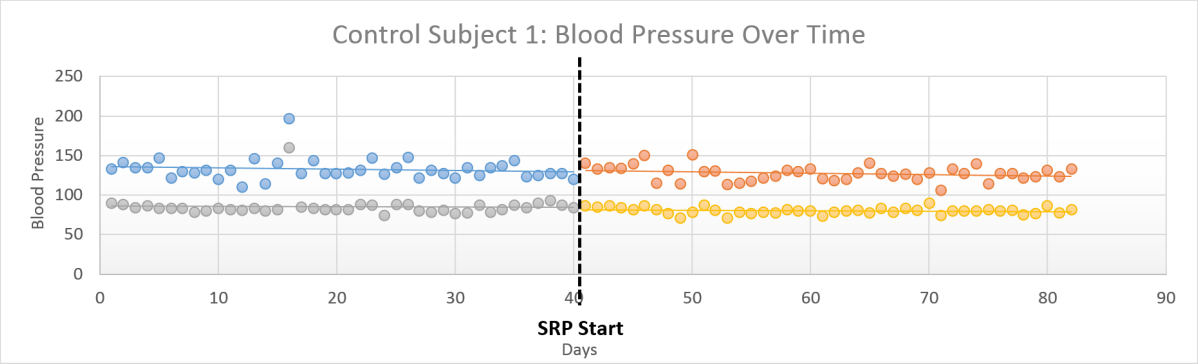

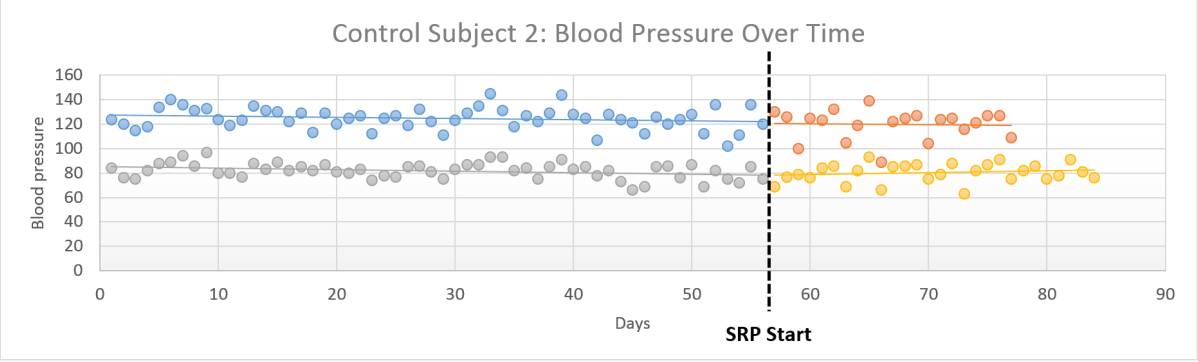

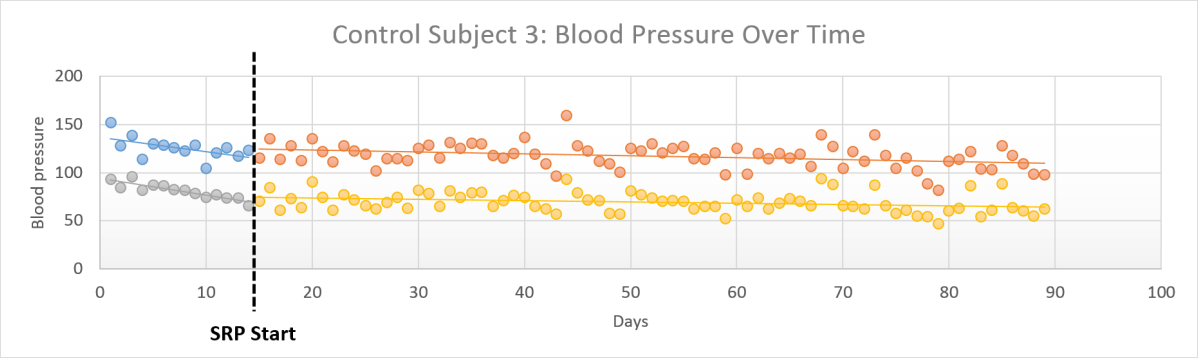

Figure 4.3: The systolic blood pressure (top) and diastolic blood pressure (bottom) values over time of 3 control subjects before and after the stress reduction routine (SRP)

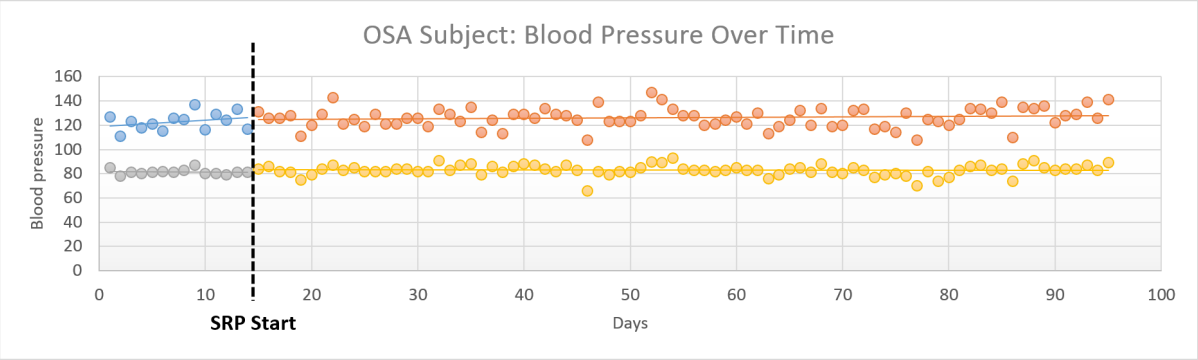

Figure 4.4: The systolic blood pressure (top) and diastolic blood pressure (bottom) values over time of a hypertensive OSA patient before and after the stress reduction routine (SRP)

Given a larger sample size, if stress reduction proves to be successful at reducing drug-resistant hypertension for some of the patients, the goal of this work is to build machine learning models to predict the success of stress reduction for any given hypertensive patient. Based on those models, we can explore which features play a significant role in the successful outcome of stress reduction with the goal of possibly securing positive outcomes for other hypertensive patients.

To date, there are not enough data (only one hypertensive OSA patient) to conduct this analysis, so we discuss possible methods of analysis and give general research guidelines on how to build machine learning models to predict the successful outcome of stress reduction in the next chapter.

4.3 In Closing

This chapter introduced the main comorbidity of sleep apnea, namely resistant hypertension, and the problems with the traditional method of managing it. We introduced a stress reduction routine as an alternative method of reducing blood pressure. The stress reduction routine is a breathing awareness exercise that is implemented as an Android application and is meant to reduce stress and lower blood pressure. The goal of this work is to evaluate the effectiveness of the stress reduction routine on resistant hypertension and, more importantly, to understand when and under which circumstances stress reduction is successful, which was not possible due to the small sample size. In the next chapter, we propose a few research guidelines for the analysis of the stress reduction data and discuss other research directions for future work.

CHAPTER 5

Conclusion and Future Work

This dissertation addressed the most serious problem in the wide spectrum of sleep-disordered breathing, namely the undertreatment of OSA, driven by the underdiagnosis and poor treatment options. Despite being the most prevalent sleep disorder and the most common type of apnea, OSA remains to be notoriously under-recognized and under-diagnosed [Org]. When left undiagnosed and untreated, OSA significantly increases the sufferer's risk of developing life-threatening physical and psychological illnesses. According to the American Academy of Sleep Medicine, approximately 80% to 90% of adults with OSA remain undiagnosed [Sle08]. The main reason for this alarming statistic is the prohibitive cost and high level of intrusiveness of polysomnography–the sleep test used to diagnose sleep apnea.

To remedy the phenomenon of underdiagnosis and make diagnosis accessible to a wider population, this dissertation presented a number of cheap, non-intrusive methods for the diagnosis of obstructive sleep apnea. The contributions of this work can be summarized as follows:

- The design and implementation of a non-intrusive *nocturnal* diagnosis method that utilizes an e-textile bed sheet. The bed sheet is seamlessly embedded with a 128 x 64 matrix of pressure sensors. Using signal and image processing techniques on the pressure images acquired from the bed sheet, a number of PSG signals, including respiration rate, leg movement, body movement, and sleep stages were extracted. A number of biophysical features were extracted from these features and used, along with other geometric features like body symmetry and sleep posture, to predict the presence or absence of OSA. The system does not require any sensors to be attached to patients while they're sleeping and doesn't need to be performed at a specialized lab or using any specialized equipment, achieving both goals of affordability and

non-intrusiveness.

- The design and implementation of a non-intrusive *daytime* diagnosis method that does not require overnight monitoring of patients. As opposed to nocturnal methods, which make their diagnoses based on detected biophysical changes that occur during sleep as a result of the obstructive events, this method leverages chronic physiologic changes that persist during wakefulness. As a result of the oxygen depletion to the brain during apneic episodes, brain tissue damage. This damage was found to target central autonomic regions of the brain, which are responsible for regulating cardiovascular functions like blood pressure and heart rate. Our daytime method presents subjects with autonomic challenges that are designed to trigger a response from the damaged regions of the brain. Physiological features are extracted from a blood pressure monitor, and a pulse oximeter while subjects perform the autonomic challenges. Using those physiological features, in addition to psychological features obtained from questionnaire responses, a diagnosis is made. The test takes 20 minutes to complete and can be conducted at any time of day without disrupting patients' sleep.

- The implementation of a sex-aware diagnosis method as an extension to the aforementioned daytime method. The sex-aware extension takes into account the different effects that OSA has on autonomic responses between males and females. Instead of applying one global model to classify subjects into the OSA or healthy categories, two different sex-aware models were designed to achieve better classification results for both sexes.

- The design an implementation of an OSA severity assessment framework. Because the severity of the OSA disorder plays an important role in its treatment, we designed an extension to the daytime diagnosis framework, which assesses the severity of the disorder, instead of just its presence or absence. Based on the severity level a patient gets classified into, the framework can also predict the Apnea-Hypopnea Index for a specific patient. In addition to its use in determining the appropriate therapeutic approach, such a severity assessment framework is useful for perioperative risk stratification and for enabling a more optimal utilization of scarce PSG

resources.

- The deployment of the aforementioned diagnosis systems in clinical trials to validate their success. In the field of wireless health research, it is crucial for the technological solutions to undergo clinical trials and be tested on real patients to confirm their effectiveness.

- The implementation of a breathing awareness stress reduction routine for the treatment of the most prevalent comorbidity of sleep apnea, namely resistant hypertension. The goal is to investigate whether or not stress reduction can help reduce blood pressure in OSA patients with resistant hypertension. If stress reduction proves successful for some patients, the gaol is to investigate the criteria that correlate with a successful outcome.

In addition to the contributions listed above, which mainly focused on the problem of non-intrusive diagnosis, this work has identified a number of new problems related to obstructive sleep apnea that can guide future research in this field. In the next sections, we describe those research problems and offer some guidelines for future work.

5.1 Stress Reduction For Resistant Hypertension Treatment

Before any analysis concerning the success of stress reduction in reducing resistant hypertension can take place, the following questions need to be answered: What constitutes a successful outcome? Should any reduction in blood pressure be considered a success? To answer these questions, future work would have to account for normal daily fluctuations in blood pressure and investigate the use of appropriate statistical measures to define outcome success. In addition to comparing the final blood pressure, after stress reduction is complete, to the initial blood pressure, future work could also analyze changes in blood pressure over time. Such an analysis may reveal some information about when stress reduction ceases to be effective or when it starts having an effect.

After outcome success is clearly defined, appropriate features (physiological, psychological, demographical) need to be extracted and used to build a classification model that can successfully predict outcome success. The next step would be to explore which

of those features are the most significant for classifying a hypertensive patient into the success group. Different feature ranking algorithms, including both filter and wrapper methods, can be experimented with. The feature ranking results may reveal that certain modifiable features contribute the most to the success of the resistant hypertension treatment. Unlike age and sex, for example, modifiable features are those that can be controlled and changed. If such modifiable features, e.g. physical activity or weight, prove to play a significant role in the effectiveness of treatment, patients with hypertension can be advised to change their behavior in order to benefit from stress reduction to lower their blood pressure. Alternatively, if the most effective features turn out not to be modifiable, this knowledge can still help physicians to target stress reduction therapies at patients who are more likely to respond to treatment. This could potentially save health care costs, as well as clinician and participant time and resources.

5.2 Weight Loss Intervention For Curing Obstructive Sleep Apnea

As explained in Section 1.2.1, obesity is a major risk factor for OSA. It is also the second most prevalent comorbidity of sleep apnea, as shown in Figure 4.1. The extra tissue and fat in the neck area compress the airway and lead to obstruction. Studies have shown that weight loss is accompanied by improvement in or a complete reversal of the OSA condition, suggesting that weight loss might be a cornerstone of the treatment of OSA [RCL10].

However, the main obstacle standing in the way of weight loss for many OSA patients is the sleep disorder itself. As a result of the breath stoppages and following arousals during the night, many OSA patients are sleep deprived and suffer from excessive daytime sleepiness (EDS). For patients who barely have enough energy to stay awake during simple conversation, the concept of exercise can seem ludicrous.

The other commonly used treatment is CPAP, which is used to manage the OSA disorder and prevent its adverse health effects. The CPAP device is a face mask that is worn over the mouth and nose during sleep to supply pressurized air continuously into the sleeper's throat. CPAP keeps the airway open and unobstructed and improves sleep

quality. Because of the intrusiveness of the CPAP mask, however, studies have shown that adherence to CPAP treatment is meager. When adherence is defined as greater than 4 hours of nightly use, 46 to 83% of patients with OSA have been reported to be non-adherent to treatment [WG08b]. Also, since CPAP does not cure OSA, the prospect of life-long CPAP is very discouraging for patients.

A possible solution lies in the combination of CPAP and exercise. Using CPAP on a daily basis will enable patients to manage daytime sleepiness, which will in turn make exercise and weight loss possible. Using Bluetooth-enabled CPAP machines, a mobile software application can be designed to track a patient's nightly use of CPAP. Information such as usage duration and usage trend over time can be displayed to the patients and reported to their physicians. Physicians will have a chance to track their patients' progress and follow up with them when necessary. The application can also send automated reminders and notifications to encourage patients to use CPAP. Based on the statistics collected by the application on daily CPAP usage and adherence over time, the application can predict the type, duration and intensity of the workout a specific patient would be able to perform and can recommend an appropriate workout routine. One day the application might recommend 10 minutes of running, another day it might suggest the patient to do 20 minutes of simple tongue and throat exercises to strengthen throat muscles and decrease OSA symptoms. The application could also monitor the activity performed by the patients and adjust the workout prediction model based on the patients' performance. Finally, the application could communicate with a Bluetooth-enabled weight scale to keep track of patients' weight and encourage them to keep losing, so that they can reverse their condition.

While the approaches that facilitate behavior change remain a challenge, the greatest benefit-to-cost ratio to society is possible through increasing health behaviors that prevent or reverse OSA and its chronic comorbidities.

5.3 Alertness Prediction For Preventing Motor Vehicle Accidents Caused by Excessive Daytime Sleepiness

One of the most dangerous symptoms of obstructive sleep apnea is excessive daytime sleepiness. Due to their frequent arousals and fragmented sleep pattern, OSA patients tend to be sleep deprived and often suffer from excessive daytime sleepiness (EDS). EDS often causes OSA patients to fall asleep at work, on the phone, during a conversation, and–more seriously–while driving. A 2015 study found that OSA patients were 2.45 times as likely to be involved in motor vehicle accidents than people without the disorder [KHH15].

One future research direction is to use a sleep tracker to track an OSA subject's sleep patterns. Sleep tracking usually relies on actigraphy–the analysis of muscular movement during sleep–or brain wave analysis. Sleep trackers come in the form of wristbands embedded with accelerometers to continuously measure activity and movement, or in the form of headbands embedded with EEG electrodes that continuously monitor brain activity. Based on the activity or brain wave data, those trackers can determine the sleep stages a sleeper is in because each sleep stage has special markers that are reflected in muscle movement and brain wave data. Based on this information, sleep trackers can provide useful information on patients' sleep patterns, like for instance, how long it took them to fall asleep, how much time they spent awake during the night, the number of arousals, the amount of time spent in each of the sleep stages, etc.

Based on the information provided by the sleep trackers, future work could extract features that can capture the sleeper's sleep quality. These features may include sleep efficiency, sleep latency, sleep duration, number of arousals per hour of sleep, average circadian cycle, time of night the sleeper goes to sleep, etc. Based on these features, machine learning models could be developed to predict the alertness and psychomotor abilities of a given user on a specific day. Based on the predicted level of alertness, a patient may be advised not to drive after a night with a poor sleep quality. Validated psychomotor skill tests, such as the ones used to select pilots [Fle56], can be used as ground truth. These psychomotor and alertness tests usually evaluate reaction time, attention, visuo-spatial functions, multi-tasking, which are all important skills for operating a motor vehicle.

Such a system would alert sleep deprived OSA patients to their decreased alertness and motor skills level for a given day and would advise them not to drive if the risk is too high. This can help protect patients from sleep-deprived driving and decrease the number of OSA-related motor vehicle accidents.

References

[AGK12] S. Alqassim, M. Ganesh, S. Khoja, M. Zaidi, F. Aloul, and A. Sagahyroon. "Sleep Apnea Monitoring using mobile phones." In *IEEE International Conference on e-Health Networking, Applications and Services*, 2012.

[Ame] American Heart Association. "Understanding and Managing High Blood Pressure, 2014." https://www.heart.org/idc/groups/heart-public/@wcm/@hcm/documents/downloadable/ucm_461840.pdf.

[Ame07] American Academy of Sleep Medicine. *The AASM manual for the scoring of sleep and associated events: rules, terminology and technical specifications.*, 2007.

[ANK00] I. Ayappa, R.G. Norman, A.C. Krieger, A. Rosen, R.L. O'malley, and D.M Rapoport. "Non-Invasive detection of respiratory effort-related arousals (REras) by a nasal cannula/pressure transducer system." *Sleep*, **23**(6):763–771, Sep 2000.

[ASAa] ASAA. "American Sleep Apnea Association." http://www.sleepapnea.org/learn/sleep-apnea.html.

[ASAb] ASAA. "OSA Treatment Options." American Sleep Apnea Association.

[Ass] British Snoring & Sleep Apnoea Association. "Sleep Position." https://www.britishsnoring.co.uk/why_do_i_snore/sleeping_position.php.

[ASS96] C.N. Alexander, R.H. Schneider, F. Staggers, W. Sheppard, B.M. Clayborne, M. Rainforth, J. Salerno, K. Kondwani, S. Smith, K.G. Walton, and B. Egan. "Trial of Stress Reduction for Hypertension in Older African Americans." *Hypertension*, **28**(2):228–237, Aug 1996.

[BGK12] J. Bak, N. Giakoumidis, G. Kim, H. Dong, and N. Mavridis. "An intelligent sensing system for sleep motion and stage analysis." *International Symposium on Robotics and Intelligent Sensors (IRIS)*, **41**:1128–1134, Aug 2012.

[BIM89] D.J. Buysse, C.F. Reynolds III, T.H. Monk, S.R. Berman, and D.J. Kupfer. "The Pittsburgh Sleep Quality Index: A new instrument for psychiatric practice and research." *Journal of Psychiatric Research*, **28**(2):193–213, 1989.

[BLM10] A. Bates, M.J. Ling, J. Mann, and D.K. Arvind. "Respiratory Rate and Flow Waveform Estimation from Tri-axial Accelerometer Data." In *Body Sensor Networks (BSN)*, 2010.

[Bro07] L.K. Brown. "Mild Obstructive Sleep Apnea Syndrome Should be Treated." *Journal of Clinical Sleep Medicine*, **3**(3):259–262, Apr. 2007.

[BTZ12] N. Ben-Israel, A. Tarasiuk, and Y. Zigel. "Obstructive Apnea Hypopnea Index Estimation by Analysis of Nocturnal Snoring Signals in Adults." *Sleep*, **35**(9):1299–1305, Sep 2012.

[CHS03] P. de Chazal, C. Heneghan, E. Sheridan, R. Reilly, P. Nolan, and M. O'Malley. "Automated processing of the single-lead electrocardiogram for the detection of obstructive sleep apnoea." *Biomedical Engineering*, **50**(6):686–696, Jun 2003.

[CKM83] S. Cohen, T. Kamarck, and R. Mermelstein. "A global measure of perceived stress." *Journal of Health and Social Behavior*, **24**(4):385–396, Dec 1983.

[CLG12] D. Cicek, H. Lakadamyali, S. Gkay, I. Sapmaz, and H. Muderrisoglu. "Effect of Obstructive Sleep Apnea on Heart Rate, Heart Rate Recovery and QTc and P-wave Dispersion in Newly Diagnosed Untreated Patients." *American Journal of Med Sci*, **344**(3):180–185, Sep. 2012.

[DAZ09] H. Danker-Hopfe, P. Anderer, J. Zeitlhofer, M. Boeck, H. Dorn, G. Gruber, E. Heller, E. Loretz andD. Moser, S. Parapatics, B. Saletu andA. Schmidt, and G. Dorffner. "Interrater reliability for sleep scoring according to the Rechtschaffen & Kales and the new AASM standard." *Journal of Sleep Research*, **18**(1):74–84, Mar 2009.

[DTZ13] E. Dafna, A. Tarasiuk, and Y. Zigel. "OSA Severity Assessment based on Sleep Breathing Analysis using Ambient Microphone." In *Engineering in Medicine and Biology Society (EMBC)*, 2013.

[EN10] E. Estrada and H. Nazeran. "EEG and HRV signal features for automatic sleep staging and apnea detection." In *Conference on Electronics, Communications and Computer (CONIELECOMP)*, 2010.

[ESE07] D. Einhorn, D.A. Stewart, M.K. Erman, N. Gordon, A. Philis-Tsimikas, and E. Casal. "Prevalence of Sleep Apnea in a Population of Adults With Type 2 Diabetes Mellitus." *Endocrine Practice*, **13**(4):355–362, 2007.

[Fle56] E.A. Fleischman. "Psychomotor Selection Tests: Research and Application in the United States Air Force." *Personal Psychology*, **9**(4):449–467, Dec. 1956.

[Fle10] S. Fleming. *"Measurement and fusion of non-invasive vital signs for routine triage of acute paediatric illness."*. Master's thesis, Oxford University, 2010.

[FLK12] L. Fraiwan, K. Lweesy, N. Khazawneh, H. Wenz, and H. Dickhaus. "Automated sleep stage identification system based on time-frequency analysis of a single EEG channel and random forest classifier." *Computer Methods and Programs in Biomedicine*, **108**(1):10–19, Oct 2012.

[Fou] National Sleep Foundation. "Teeth Grinding." National Sleep Foundation.

[Fou11] National Sleep Foundation. "Sleep Apnea and Sleep." https://sleepfoundation.org/sleep-disorders-problems/sleep-apnea, 2011.

[GBC06] J.B. Gross, J.L. Benumof, R.A. Caplan, R.T. Connis, C.J. Cote, D.G. Nickinovich, V. Prachand, D.S. Ward, E.M. Weaver, L. Ydens, and S. Yu. "Practice Guidelines for the Perioperative Management of Patients with Obstructive Sleep Apnea." *Anesthesiology*, **104**:1081–1093, 2006.

[GS98] A.H. Glassman and P.A. Shapiro. "Depression and the course of coronary artery disease." *American Journal of Psychiatry*, **155**(1):4–11, Jan 1998.

[Har] Division of Sleep Medicine at Harvard Medical School. "Natural Patterns of Sleep." http://healthysleep.med.harvard.edu/healthy/science/what/sleep-patterns-rem-nrem.

[HDS13] E. Hoque, R.F. Dickerson, and J.A. Stankovic. "Monitoring Sleep with WISP Tags." In P. Peris-Lopez, J.C. Hernandez-Castro, and T. Li, editors, *Security and Trends in Wireless Identification and Sensing Platform Tags: Advancements in RFID.* Information Science Reference, 2013.

[HPP04] J. Hedner, G. Pillar, S.D. Pittman, D. Zou, L. Grote, and D.P. White. "A novel adaptive wrist actigraphy algorithm for sleep-wake assessment in sleep apnea patients." *SLEEP*, **27**(8):1560–1566, Dec 2004.

[HSK87] R.M Harper, V.L. Schechtman, and K.A. Kluge. "Machine classification of infant sleep state using cardiorespiratory measures." *Electroencephalography and Clinical Neurophysiology*, **67**(4):379–387, Oct 1987.

[HXL13] M.C. Huang, W. Xu, J. Liu, L. Samy, A. Vajid, N. Alshurafa, and M. Sarrafzadeh. "Inconspicuous on-bed respiratory rate monitoring." In *PErvasive Technologies Related to Assistive Environments*, 2013.

[Jaw] Jawbone. "UP." https://jawbone.com/up/trackers.

[KBS99] V. Kapur, D.K. Blough, R.E. Sandblom, R. Hert, J.B. de Maine, S.D. Sullivan, and B.M. Psaty. "The Medical Cost Of Undiagnosed Sleep Apnea." *SLEEP*, **22**(6):749–755, Sep 1999.

[KHH15] M. Karimi, J. Hedner, H. Hbel, O. Nerman, and L. Grote. "Sleep Apnea Related Risk of Motor Vehicle Accidents is Reduced by Continuous Positive Airway Pressure: Swedish Traffic Accident Registry Data." *SLEEP*, **38**(3):341–349, Mar. 2015.

[KMB10] J.M. Kortelainen, M.O Mendez, A.M. Bianchi, M. Matteucci, and S. Cerutti. "Sleep Staging Based on Signals Acquired Through Bed Sensor." *IEEE Transactions on Information Technology in Biomedicine*, **14**(3):776–785, Apr 2010.

[KOE98] S. Kulkarni, I. O'Farrell, M. Erasi, and M.S. Kochar. "Stress and hypertension." *Wisconsin Medical Journal*, **97**(11):34–38, Dec 1998.

[KPO13] A. Khan, N.K. Patel, D.J. O'Hearn, and S. Khan. "Resistant Hypertension and Obstructive Sleep Apnea." *International Journal of Hypertension*, **2013**, 2013.

[KPR03] J.W. Kantelhardt, T. Penzel, S. Rostig, H.F. Becker, S. Havlin, and A. Bunde. "Breathing during REM and non-REM sleep: correlated versus uncorrelated behaviour." *Physica A: Statistical Mechanics and its Applications*, **319**:447–457, Mar 2003.

[KSW09] K. Kroenke, R.L. Spitzer, J.B. Williams, and B. Lowe. "An ultra-brief screening scale for anxiety and depression: the PHQ-4." *Psychosomatics*, **50**(6):613–621, 2009.

[Lit07] M.R. Littner. "Mild Obstructive Sleep Apnea Syndrome Should Not Be Treated." *Journal of Clinical Sleep Medicine*, **3**(3):263–264, Apr. 2007.

[LPL08] D. Liu, Z. Pang, and S. R. Lloyd. "A Neural Network Method for Detection of Obstructive Sleep Apnea and Narcolepsy Based on Pupil Size and EEG." *IEEE Transactions on Neural Networks*, **19**(2):308–318, Feb 2008.

[LPM01] A.G. Logan, S.M. Perlikowski, A. Mente, A. Tisler, R. Tkacova, M. Niroumand, R.S. Leung, and T.D Bradley. "High prevalence of unrecognized sleep apnoea in drug-resistant hypertension." *Journal of Hypertension*, **19**(12):2271–2277, Dec 2001.

[LRH09] Y. Lahav, E. Rosenzweig, Z. Heyman, J. Dolijansky, A. Green, and Y. Dagan. "Tongue base ultrasound: a diagnostic tool for predicting obstructive sleep apnea." *Annals of Otology, Rhinology and Laryngology*, **118**(3):179–184, Mar 2009.

[LXA13] J. Liu, W. Xu, M.C. Huang andN. Alshurafa, and M. Sarrafzadeh. "A Dense Pressure Sensitive Bedsheet Design for Unobtrusive Sleep Posture Monitoring." In *Pervasive Computing and Communications (PerCom)*, 2013.

[Mas] Masimo. "Masimo." www.masimo.com.

[MAT04] R. McCraty, M. Atkinson, and D. Tomasino. "Impact of a Workplace Stress Reduction Program on Blood Pressure and Emotional Health in Hypertensive Employees." *Journal of Alternative and Complementary Medicine*, **9**(3):355–369, Jul 2004.

[MDV03] U.J. Magalang, J. Dmochowski, S. Veeramachaneni, A. Draw, M.J. Mador, A. El-Solh, and B.J.B. Grant. "Prediction of the Apnea-Hypopnea Index From Overnight Pulse Oximetry." *Chest*, **124**(5):1694–1701, Nov 2003.

[Mey08] Jan Meyer. *"Textile Pressure Sensor: Design, Error Modeling and Evaluation."*. Master's thesis, ETH Zurich, 2008.

[MHA09] J. Victor Marcos, Roberto Hornero, Daniel Alvarez, Felix del Campo, and Carlos Zamarron. "Assessment of four statistical pattern recognition techniques to assist in obstructive sleep apnoea diagnosis from nocturnal oximetry." *Medical Engineering and Physics*, **31**(8):971–978, Oct 2009.

[MHA12] J.V. Marcos, R. Hornero, D. Alvarez, M. Aboy, and F. del Campo. "Automated Prediction of the Apnea-Hypopnea Index from Nocturnal Oximetry Recordings." *IEEE Transactions on Biomedical Engineering*, **59**(1):141–149, Jan 2012.

[MKW08] PM. Macey, R. Kumar, MA. Woo, EM. Valladares, FL. Yan-Go, and RM. Harper. "Brain Structural Changes in Obstructive Sleep Apnea." *Sleep*, **31**(7):967–977, Jul 2008.

[MKW13] PM. Macey, R. Kumar, MA. Woo, FL. Yan-Go, and RM. Harper. "Heart Rate Responses to Autonomic Challenges in Obstructive Sleep Apnea." *PLoS One*, Oct 2013.

[MM11] A. Montazeri and Z. Moussavi. "Obstructive Sleep Apnea Prediction During Wakefulness." In *Engineering in Medicine and Biology Society*, 2011.

[MMK10] M.O Mendez, M. Migliorini, J.M. Kortelainen, D. Nistico, E. Arce-Santana, S. Cerutti, and A.M. Bianchi. "Evaluation of the sleep quality based on bed sensor signals: Time-variant analysis." In *Engineering in Medicine and Biology Society (EMBC)*, 2010.

[MMQ03] M.J. Morrell, D.W. McRobbie, R.A Quest, A.R. Cummin, R. Ghiassi, and D.R. Corfield. "Changes in brain morphology associated with obstructive sleep apnea." *Sleep Medicine*, **4**(5):451–454, Sep 2003.

[MPI00] J.E. Mietus, C.K. Peng, P.Ch. Ivanov, and A.L. Goldberger. "Detection of obstructive sleep apnea from cardiac interbeat interval time series." In *Computers in Cardiology*, 2000.

[MSK09] G. Matsuoka, T. Sugi, F. Kawana, and M. Nakamura. "Automatic detection of apnea and EEG arousals for sleep apnea syndrome." In *ICCAS-SICE*, 2009.

[Nat10] R. Natarajan. "Review of periodic limb movement and restless leg syndrome." *Journal of Postgraduate Medicine*, **56**(2):157–162, 2010.

[NSN99] N.C. Netzer, R.A. Stoohs, C.M. Netzer, K. Clark, and K.P Strohl. "Using the Berlin Questionnaire to identify patients at risk for the sleep apnea syndrome." *Annals of Internal Medicine*, **131**(7):485–491, Oct 1999.

[OLF06] O. Oldenburg, B. Lamp, L. Faber, H. Teschler, D. Horstkotte, and V. Topfer. "Sleep-disordered breathing in patients with symptomatic heart failure: a contemporary study of prevalence in and characteristics of 700 patients." *European Journal of Heart Failure*, **9**(3):251–257, Mar 2006.

[Org] World Health Organization. "Chronic Respiratory Diseases." http://www.who.int/gard/publications/chronic_respiratory_diseases.pdf.

[PAF01] D. Purves, G.J. Augustine, D. Fitzpatrick, L.C. Katz, A-S. LaMantia, J. O McNamara, and S.M. Williams, editors. *Neuroscience*. Sinauer Associates, 2 edition, 2001.

[PBT84] CJ. Porth, VS. Bamrah, FE. Tristani, and JJ. Smith. "The Valsalva maneuver: mechanisms and clinical implications." *Heart and Lung*, **13**(5):507–518, Sep 1984.

[PDP13] R.P. Pedrosa, L.F. Drager, L.K. de Paula, A.C. Amaro, L.A. Bortolotto, and G. Lorenzi-Filho. "Effects of OSA treatment on BP in patients with resistant hypertension: a randomized trial." *Chest*, **144**(5):1487–1494, Nov 2013.

[Phi] Philips. "Actiwatch." http://www.usa.philips.com/healthcare/product/HC1044809/actiwat
2-activity-monitor.

[PLS06] Y. Peng, C. Lin, and M. Sun. "Multimodality Sensors for Sleep Quality Monitoring and Logging." In *International Conference on Data Engineering Workshops*, 2006.

[RB97] P. Renaud and J.P. Blondin. "The stress of Stroop performance: physiological and emotional responses to color-word interference, task pacing, and pacing speed." *International Journal of Psychophysiology*, **27**(2):87–97, Sep 1997.

[RCL10] A. Romero-Corral, S.M. Caples, F. Lopez-Jimenez, and V.K. Somers. "Interactions Between Obesity and Obstructive Sleep Apnea." *Chest*, **137**(3):711–719, Mar. 2010.

[RH06] S. Redmond and C. Heneghan. "Cardiorespiratory-based sleep staging in subjects with obstructive sleep apnea." *IEEE Transactions on Biomedical Engineering*, **53**(3):485–496, Mar 2006.

[RRK14] A. Romem, A. Romem, D. Koldobskiy, and S. M. Scharf. "Diagnosis of Obstructive Sleep Apnea Using Pulse Oximeter Derived Photoplethysmographic Signals." *Journal of Clinical Sleep Medicine*, **10**(3):285–290, 2014.

[RSN07] M.V. Rainforth, R.H. Schneider, S.I. Nidich, C. Gaylord-King, J.W. Salerno, and J.W. Anderson. "Stress reduction programs in patients with elevated blood pressure: A systematic review and meta-analysis." *Current Hypertension Reports*, **9**:520, Dec 2007.

[SA02] A. Sadeh and C. Acebo. "The role of actigraphy in sleep medicine." *Sleep Medicine Reviews*, **6**(2):113–124, Apr 2002.

[SAB07] M. Silber, S. Ancoli-Israel, M. Bonnet, S. Chokroverty, M. Grigg-Damberger, M. Hirshkowitz, S. Kapen, S. Keenan, M. Kryger, T. Penzel, M. Pressman, and C. Iber. "The visual scoring of sleep in adults." *Journal of Clinical Sleep Medicine*, **3**(2):121–131, 2007.

[SBP03] L. de Souza, A. Benedito-Silva, M. Pires, D. Poyares, S. Tufik, and H. Calil. "Further validation of actigraphy for sleep studies." *SLEEP*, **26**(1):81–85, 2003.

[SGM85] P.L. Smith, A.R. Gold, D.A. Meyers, E.F. Haponik, and E.R. Bleecker. "Weight Loss in Mildly to Moderately Obese Patients with Obstructive Sleep Apnea." *Annals of Internal Medicine*, **103**(6):850–855, Dec 1985.

[SGS91] A.R. Schwartz, A.R. Gold, N. Schubert, A. Stryzak, R.A. Wise, S. Permutt, and P.L. Smith. "Effect of Weight Loss on Upper Airway Collapsibility in Obstructive Sleep Apnea." *American Thoracic Society*, **144**(3):494–498, Sep 1991.

[SHL14] L. Samy, M.C. Huang, J.J. Liu, W. Xu, and M. Sarrafzadeh. "Unobtrusive Sleep Stage Identification Using a Pressure-Sensitive Bed Sheet." *IEEE Sensors Journal*, **14**(7):2092–2101, Jul 2014.

[SKE99] H. Schafer, U. Koehler, S. Ewig, E. Hasper, S. Tasci, and B. Luderitz. "Obstructive sleep apnea as a risk marker in coronary artery disease." *Cardiology*, **92**(2):79–84, 1999.

[Sle08] American Academy of Sleep Medicine (AASM). "Obstructive Sleep Apnea." http://www.aasmnet.org/resources/factsheets/sleepapnea.pdf, 2008.

[SMK93] Y. Suto, T. Matsuo, T. Kato, I. Hori, Y. Inoue, S. Ogawa, T. Suzuki, M. Yamada, and Y. Ohta. "Evaluation of the pharyngeal airway in patients with sleep apnea: value of ultrafast MR imaging." *AJR: American Journal of Roentgenology*, **160**(2):311–314, Feb 1993.

[SMY07] M. Suetsugi, Y. Mizuki, K. Yamamoto, S. Uchida, and Y. Watanabe. "The effect of placebo administration on the first-night effect in healthy young volunteers." *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, **31**(4):839–847, 2007.

[SN14] N. Selvaraj and R. Narasimhan. "Automated prediction of the apnea-hypopnea index using a wireless patch sensor." In *Engineering in Medicine and Biology Society (EMBC)*, 2014.

[SOM] SOMNOmedics. "SOMNOmedics." http://somnomedics.eu/.

[SZE00] M. Schrader, C. Zywietz, V. Einem, B. Widiger, and G. Joseph. "Detection of sleep apnea in single channel ECGs from the PhysioNet data base." In *Computers in Cardiology*, 2000.

[Tha06] M.E. Thase. "Depression and sleep: pathophysiology and treatment." *Dialogues in Clinical Neurosciences*, **8**(2):217–226, Jun 2006.

[THG09] D.I Townsend, M. Holtzman, R. Goubran, M. Frize, and F. Knoefel. "Simulated central apnea detection using the pressure variance." *Engineering in Medicine and Biology Society (EMBC)*, pp. 3917–3920, 2009.

[Tra] Sleep Tracker. "Sleeptracker." https://www.sleeptracker.com/.

[VBD07] W. Vosa, J. De Backer, A. Devolder, O. Vanderveken, S. Verhulst, R. Salgado, P. Germonpre, B. Partoens, F. Wuyts, P. Parizel, and W. De Backer. "Correlation between severity of sleep apnea and upper airway morphology based on advanced anatomical and functional imaging." *Journal of Biomechanics*, **40**(10):2207–2213, 2007.

[WB86] W.B. White and L.H. Baker. "Episodic hypertension secondary to panic disorder." *Archives of Internal Medicine*, **146**(6):1129–1130, Jun 1986.

[WG08a] T.E. Weaver and R.R. Grunstein. "Adherence to Continuous Positive Airway Pressure Therapy." *Proceedings of the American Thoracic Society*, **5**(2):173–178, Feb 2008.

[WG08b] T.E. Weaver and R.R. Grunstein. "Adherence to Continuous Positive Airway Pressure Therapy." *American Thoracic Society*, **5**(2):173–178, Feb 2008.

[WLR14] H.K. Walia, H. Li, M. Rueschman, D.L. Bhatt, S.R. Patel, S.F. Quan, D.J. Gottlieb, N.M. Punjabi, S. Redline, and R. Mehra. "Association of Severe Obstructive Sleep Apnea and Elevated Blood Pressure Despite Antihypertensive Medication Use." *Journal of Clinical Sleep Medicine*, **10**(8):835–843, Aug. 2014.

[XLH11] W. Xu, Z. Li, M.C. Huang, N. Amini, and M. Sarrafzadeh. "ecushion: An etextile device for sitting posture monitoring." In *Body Sensor Networks (BSN)*, 2011.

[YFP08] T. Young, L. Finn, P.E Peppard, M. Szklo-Coxe, D. Austin, F.J. Nieto, R. Stubbs, and K.M. Hla. "Sleep Disordered Breathing and Mortality: Eighteen-Year Follow-up of the Wisconsin Sleep Cohort." *SLEEP*, **31**(8):1071–1078, Aug 2008.

[YL10] W. Yao and C.H. Chu andZ. Li. "The use of RFID in healthcare: Benefits and barriers." In *IEEE International Conference on RFID-Technology and Applications (RFID-TA)*, 2010.

[YM09] A. Yadollahi and Z. Moussavi. "Acoustic obstructive sleep apnea detection." In *Engineering in Medicine and Biology Society*, 2009.

[YPD93] T. Young, M. Palta, J. Dempsey, J. Skatrud, S. Weber, and S. Badr. "The occurrence of sleep-disordered breathing among middle-aged adults." *The New England Journal of Medicine*, **328**:1230–1235, Apr 1993.

[YTM11] Y. Yamana, S. Tsukamoto, K. Mukai, H. Maki, H. Ogawa, and Y. Yonezawa. "A sensor for monitoring pulse rate, respiration rhythm, and body movement in bed." In *Engineering in Medicine and Biology Society (EMBC)*, 2011.

[ZCN06] X. Zhu, W. Chen, and T. Nemoto. "Real-Time Monitoring of Respiration Rhythm and Pulse Rate During Sleep." *IEEE Transactions on Biomedical Engineering*, **53**(12):2553–2563, Dec 2006.

[ZMJ14] M.A.B. Zapata, E.D. Meza, C. Jaramillo, D.M. Gomez, and C. Torres-Duque. "Sleep apnea and oxygen saturation in adults at 2640 m above sea level." *Sleep Science*, **7**(2):103–106, Sep. 2014.

[ZRF08] S.N. Zallek, R. Redenius, H. Fisk, C. Murphy, and E. O'Neill. "A single question as a sleepiness screening tool." *Journal of Clinical Sleep Medicine*, **4**(2):143–148, Apr 2008.