12-2-2014

Automated Detection of Cigarette Smoking Puffs from Mobile Sensors - A Multimodal Approach

Amin Ahsan Ali

AUTOMATED DETECTION OF CIGARETTE SMOKING PUFFS FROM MOBILE SENSORS - A MULTIMODAL APPROACH

by

Amin Ahsan Ali

A Dissertation

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

Major: Computer Science

The University of Memphis

December 2014

ACKNOWLEDGMENTS

I would like to thank my supervisor Santosh Kumar who has always been there to motivate me and provide guidance for my research. He has been an extremely supportive supervisor. He ensured my continuous research funding which allowed me to focus on my studies.

Special thanks to my long time friend and lab mate Syed Monowar Hossain without whose help I could not have reached this point. I would also like to thank my lab mates Md. Mahbubur Rahman, Hillol Sarker, and former post-docs of the lab Karen Hovsepien and Kurt Plarre for their help in the different projects I was involved in the lab.

I would like to thank my dissertation committee members for their valuable comments and suggestions that has helped me immensely.

I would like to thank Emre Ertin (at The Ohio State University), who designed the AutoSense sensors, Mustafa al'Absi (at University of Minnesota) for providing access to the lab stress data sets used in this work, and Saul Shiffman (at University of Pittsburgh) for many invaluable conversations on smoking research. This work was supported in part by NSF grant CNS-0910878 funded under the American Recovery and Reinvestment Act of 2009 (Public Law 111-5) and NIH Grant U01DA023812 from NIDA.

Finally, I would like thank my family for their encouragements and unconditional support.

ABSTRACT

Ali, Amin Ahsan. Ph.D. The University of Memphis. December 2014. Automated Detection of Cigarette Smoking Puffs from Mobile Sensors - A Multimodal Approach. Major Professor: Santosh Kumar.

Smoking has been conclusively proved to be the leading cause of preventable deaths in the United States. Extensive research is conducted on developing effective smoking cessation programs. Most smoking cessation programs achieve low success rate because they are unable to intervene at the right moment. Identification of high-risk situations that may lead an abstinent smoker to relapse involves discovering the associations among various contexts that precede a smoking session or a smoking lapse. In the absence of an automated method, detection of smoking events still relies on subject self-report that is prone to failure to report and involves subject burden. Automated detection of smoking events in the natural environment can revolutionize smoking research and lead to effective intervention.

We investigate the feasibility of automated detection smoking puff from measurement obtained from respiratory inductive plethysmography (RIP) sensor. We introduce several new features from respiration that can help classify individual respiration cycles into smoking puffs or non-puffs. We then propose supervised and semi-supervised support vector models to detect smoking puffs. We train our models on data collected from 10 daily smokers and show that our model can still identify smoking puffs with an accuracy of 86.7%.

We further show accuracy of smoking puff detection can be improved by fusing measurements from RIP and inertial sensors. We use measurements obtained from wrist worn accelerometer and gyroscope to find segments when the hand is at mouth. The segments are used to identify respiration cycles that can be potentially puff cycles. A SVM classifier is trained using 40 hours of data collected from 6 participants. The 10-fold cross validation results show that at 90.3% true

positive rate, respiration feature based classifier produces on average 43.8 false positives puff per hours which is reduced to 3.7 false positives per hour when both wrist and respiration features are used. We also perform leave one subject out cross validation and show that the method generalized well.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

Introduction

1.1 Motivation

Since the first U.S. Surgeon General's report in 1964 there has been overwhelming and conclusive evidence that use of tobacco, especially in the form of cigarette smoking, causes cancer in different organs throughout the body, leads to cardiovascular and respiratory diseases, and harms reproduction [1]. Smoking induced diseases account for nearly one of every five deaths in the United States [2]. Smokers die 13-14 years younger and cost $193 billion annually [3]. In addition, almost 60% of children are still exposed to secondhand smoking which is also a known human carcinogen. To reduce such harmful effects of smoking, there needs to be substantial progress in tobacco control, health education programs, and development of interventions to aid smoking cessation.

Given the adverse impact of smoking on human health, significant research is conducted on development of smoking interventions. Eight (out of 27) divisions at National Institutes of Health (NIH) award research grants for smoking cessation. Of these, National Cancer Institute (NCI) alone awards $100+ million annually in smoking research. Despite extensive efforts, smoking continues to be prevalent. Seventy percent of adult smokers want to quit completely, while 40% try to quit each year - but most quit efforts end in relapse [4]. Each day about 2,000 people become new daily smokers [1].

Most smoking cessation programs achieve low success rate (i.e., less than 10%) because they are unable to intervene at the right moment. Smokers who are trying to quit need to avoid high-risk situations and if they get into a high-risk situation, need an intervention to break their urge. Given the ubiquity of smart phones, such a smoking cessation assistant app can be developed for smart phones that intervenes if a quitter is found to be in a high-risk situation. The

challenge, however, is to automatically identify when the quitter is in a high-risk situation. Considerable amount of research work is focused on identifying the factors (called *antecedents*) that lead to high-risk situations (and, eventually relapse) to design effective interventions [5], [6], [7] .

Identifying antecedents and precipitants of smoking lapse (i.e., an acute condition such as stress that causes a lapse) requires conducting scientific user studies in the natural environment so as to observe the psychological, social, and environmental factors that may be associated with smoking instances [8], [9], [10], [11], [12], [13], [14]. This is done by observing and recording the user's context when smoking occurs for daily smokers or when a smoking lapse occurs in those trying to quit [15]. These studies must have some mechanism of detecting when smoking occurs, so that physical, physiological, psychological, behavioral, social, and environmental contexts before, during, and after a smoking session can be identified. Most current studies on smoking behavior rely on various self-reporting techniques, where subjects are asked to self-report each smoking episode. These methods range from basic pen-paper methods and retrospective recalls, to electronic diary keeping and ecological momentary assessments (EMA) [16], [15], [17], [18]. These methods, in addition to imposing a burden on the study participants, have the limitation of introducing biases when recalling events, forgetting to report, among several others.

Technological methods to detect smoking episodes include the use of external sensors such as carbon monoxide (CO) monitors (piCO$^+$), CReSS Pocket [19], Micro$^+$ [20], instrumented cigarette lighters [21], RespiTrace® [22], video cameras [23]. As discussed in more detail in Chapter 2, each of these methods require manual intervention by subjects or an observer. In recent years, there has been growing interest in the use of body worn sensors (e.g., inertial

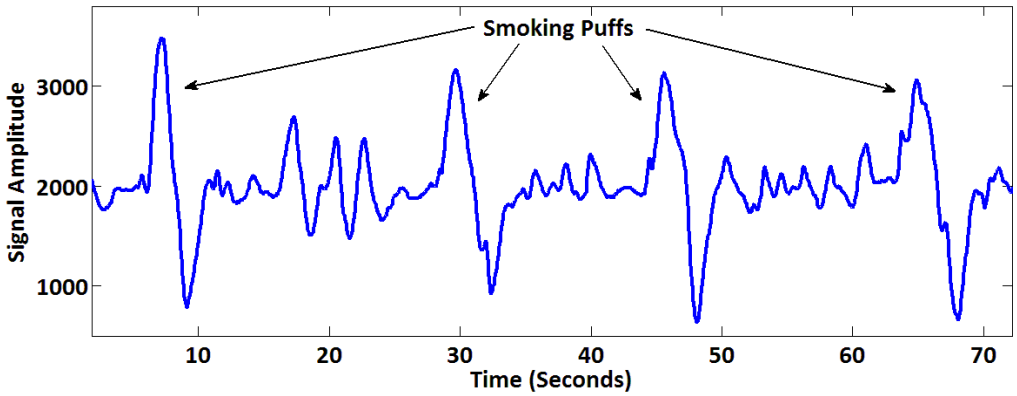

Fig. 1.1: Respiration signals captured during a typical smoking episode.

sensors [24], [25] and respiratory inductive plethysmography (RIP) sensors [26]) to detect smoking in the natural environment that is operator independent.

1.2 Objective

In this thesis, we first investigate the feasibility detecting smoking puff from measurement obtained from RIP sensor. Second, we investigate whether we can improve the accuracy of smoking puff detection by fusing measurements from multiple sensing modalities. In particular, we combine measurements from RIP and inertial sensors to detect cigarette smoking puffs.

We take the first step towards automatically detecting smoking in the natural environment by developing *mPuff*, a model to automatically detect smoking puffs from respiration measurements. mPuff uses respiration measurements collected from a respiration band that the user wears underneath their clothing. Detection of smoking puffs from respiration is feasible because they are associated with deep inhalation and deep exhalation (see Figure 1.1 for an example). It should be noted that a puff lasts only for one respiration cycle. Thus in order to detect puffs we need to find appropriate features that can help discriminate a smoking puff not only from usual respiration cycles, but also from those respiration cycles that may represent speaking, stress, or physical activity such as walking.

In order to achieve better accuracy to classify respiration cycles, we investigate the utility of detecting hand gestures associated with smoking to locate potential puff cycles. During smoking, the hand holding the cigarette moves towards the mouth to put the cigarette between the lips, the hand remains stationary for a short duration at the mouth while inhaling the smoke, and finally the smoke is exhaled usually after the cigarette is removed from the mouth. Wrist worn inertial sensors can be used to detect such gestures. However, trajectory of the hand during a puff events depends on the posture and possible concurrent activity the person is involved in (e.g., standing, sitting, lying, walking, talking, or driving). Rather than tracking the trajectory of the hand we use a 6-axis inertial sensor (3 axis accelerometer and 3-axis gyroscope) to detect the timing when the hand is at mouth. As the exhalation of the smoke occur after the cigarette is removed from the mouth, finding the timing of hand at mouth gesture enables us to locate the respiration cycle that potentially corresponds to a puff. This process reduces number of respiration cycles that are needed to be assessed. Moreover, we find features from the accelerometer and gyroscope measurements during the hand at mouth gesture segments and combine them with respiration features in order to detect puffs.

In the following, we summarize the rest of this dissertation, that includes related works (Chapter 2), a respiration only method for detecting smoking puffs (Chapter 3), composing the respiration measurements with inertial sensor measurements to improve the accuracy of detecting smoking puffs (Chapter 4), and a description of conclusion and future works (Chapter 5).

1.3 Summary of the Dissertation

Chapter 2 presents the current state of art of smoking detection methods. We also describe the works on human state detection (e.g., conversation and stress) from respiration measurements.

In Chapter 3 we describe the feature identification and model development and evaluation of the mpuff system. In order to identify discriminatory features we examine the morphology of respiration cycles during smoking and other confounding events. A smoking puff tends to lengthen the duration of a respiration cycle relative to its neighboring cycles. It also amplifies the degree of inhalation and exhalation (i.e., both directions on y-axis) as compared to the usual level of peak and valley, as well as that compared to the neighboring cycles. We build on these insights to identify 12 new features in respiration measurements that together with 5 previously known features are used in mPuff to detect smoking puffs. A majority of these features are not person specific as they measure relative changes of some basic characteristics of the respiration cycles. For those features that depend on absolute values such as inspiration duration, we normalize it by computing the z-score of the feature values using the person-specific mean and standard deviation, thus accounting for the between-person differences. Therefore, mPuff self-calibrates to each person and does not need to be trained on a person prior to its usage.

In order to develop the mPuff model from the above mentioned features, we use support vector machine (SVM) that is trained over the respiration features. Given that various smoking researchers may need different sensitivities to false positive and false negative rate, the model we use can be customized for a target false positive or false negative rate. Given imperfection in automated models, the smoking research studies may continue collecting self-reports. These self reports, however, may not always be located before the start of a smoking session. They in some cases may be located during or after a smoking session. We propose a semi-supervised support vector model that improves the accuracy of detecting smoking puffs by making use of the self-report markings. Our model can potentially be used as a building block to develop a full-fledged smoking detector

that can identify those smoking episodes that may not have received a self-report marking.

To train mPuff, we collected respiration data during smoking from 10 volunteer daily smokers. During the collection of labeled data, we carefully marked each puff in a smoking session. To ensure generalizability of our model, we also used data sets from major confounding events, e.g., physical activity, conversation, and stress, that may cause similar patterns in respiration as smoking puffs. All these data sets constitute the training and testing data sets of the supervised SVM. These data sets are supplemented by collection of respiration data from 4 volunteers (out of the original 10 volunteers) who wore the sensors for 7 continuous days in their natural environment and self-marked smoking episodes. Together with the labeled data, this data set is used for the training of the semi-supervised model.

Results: We find that smoking puffs can be detected with an accuracy of 91% within a smoking session. When applying the model to confounding events, we obtain an accuracy of 83.85% for the supervised SVM model, which improves to 86.7% by using a semi-supervised model that is able to use a much larger data set from the field. We also find that the accuracy of the classifier increases by more than 10% by using the newly proposed features. We applied mPuff to our data set to observe patterns of smoking behavior. We find that the average duration of a smoking session is 6.62 minutes, a smoking session contains an average of 12 puffs, among several other interesting statistics.

In Chapter 4 we first describe the hand at mouth gesture detection method that is based on wrist worn accelerometer and gyroscope measurements. We identify the segments in the sensor data timeseries to identify where the hand is relatively stable compared to its surrounding regions. We then screen out segments based on their duration and the orientation of hand during those

segments. The candidate segments are then mapped to respiration cycles and we extract features from both the respiration cycle and candidate segments. We propose 2 new respiration features and 12 features on the hand gesture segments. To train the classifier, we collected 32 smoking episodes from 6 participants who wore the respiration and two wrist sensors for a total of 40.3 hours in free living conditions. These 32 smoking episodes include instances where the participants smoked while standing, sitting, walking, and being in a conversation. This is unlike the smoking data collected in mpuff system where the participants smoked while standing and was not involved in confounding activities such as conversation. 10-fold cross validation on the data collected produced 3.7 false positives per hour in non-smoking regions of the data while maintaining a true positive rate of 90.3%. At the same true positive rate we obtain 43.8 false positives per hour if only respiration based features are used. In order to evaluate the generalizability of the classifier we perform leave-one-subject-out cross validation. In this case, we use data from 3 of the 6 participants for each of whom we have at least 5 hours of data. We observe that the classifier fairly generalizable. For the 3 participants the true positive rate is 85%, 74%, and 85% with 1.42, 5.3, and 9.55 false positive per hour respectively.

Finally in Chapter 5 we describe the future work and conclude the dissertation.

Future Applications: Our model opens the opportunity for automated detection of smoking episodes in the natural environment. Since respiration measurements can be used to detect stress [27] as well, which has been found to be a leading predictor of smoking relapse, smoking research can potentially be revolutionized. It has been found that stress levels of abstinent smokers who relapsed rises hours before a lapse [5]. Now, it can be found out what happens in the minutes preceding a smoking lapse. Since several other contexts such as

location, commuting, physical activity, and social interactions can also be detected on a smart phone, rich contextual analysis can be conducted to find true predictors of smoking lapses. Such analysis can then be used to design effective interventions which can be delivered on a mobile phone, when and where smoking urges may occur.

Chapter 2

Background

In this chapter, we review existing smoking cessation treatments and methods for monitoring of smoking events in the field environment targeted at improving smoking cessation research. This will provide an appropriate context for the smoking event detection work presented in this dissertation. To provide an appropriate background from the technological perspective, we describe works on detecting behaviors from respiration measurements that play a role in developing the computational methods proposed in this dissertation.

We note that the mPuff system [28] described in this dissertation, to the best of our knowledge, is one of the first works that presents the potential for automated detection of cigarette smoking in the natural environment. Since the publication of our work there has been growing interest[1] in the detection of smoking using respiration measurements [26] and other modalities (e.g., inertial sensors [24], [25].

2.1 Tobacco Control and Smoking Cessation Programs

Smoking control and cessation strategies have evolved and expanded as our understanding of smoking behavior and its risk has developed over the years. Medical evidence linking smoking with cancer began to accumulate in 1930s and 1940s. In 1950 case-control studies published by Wynder and Graham [29] and Doll and Hill [30] established the association between smoking and lung cancer. These studies further stimulated the research on the health effects of smoking and by late 1950s the causal relationship between smoking and lung cancer was generally accepted [31]. In the subsequent decades smoking was established as the cause nearly 40 diseases such as cancer, heart disease, stroke, lung diseases (including emphysema, bronchitis, and chronic airway obstruction), and

[1] The mPuff work has been cited 18 times in the 2 years since its publication.

diabetes [32], [33], [34]. Subsequent research also established the harmful effects of second hand smoking. Secondhand smoke refers to the involuntary exposure to tobacco smoke of non-smokers. Secondhand smoke has been proved to the cause of cancer and coronary heart disease in non-smokers. It also increases the risk of asthma, respiratory infections, and sudden infant death syndrome [35].

Recognition of the harmful effects of smoking initially lead to efforts organized by both the government and voluntary health organizations to educate the public about the risks of smoking. In next couple of decades, several regulatory measures were taken to add warning labels to cigarette packages, to restrict and eventually ban cigarette advertising in media and smoking in federal facilities and airlines. In the last decade, comprehensive clean indoor air laws were introduced that prohibit smoking indoors in workplaces, restaurants, public transports, hospitals, and educational institutes. This had direct impact on reduction of secondhand smoke and overall consumption of cigarettes [36]. Over the years there has also been steep increase in tobacco excise tax which has lead to higher prices of cigarettes and reduced consumption [37], [38]. Increased awareness about the problems created by smoking and the measures described above has contributed to the decline of smoking prevalence and has encouraged smokers to quit. Among the current adult smokers in US about 68% report that they want to quit completely [39].

However, most quit attempts are unsuccessful. Only 6.2% of smokers successfully abstained from smoking up to a year after quitting [39]. Different pharmacological and behavioral treatments have been developed to assist smokers to quit.

Pharmacological Treatments: Dependence on nicotine as the fundamental reason of sustained smoking behavior has been widely accepted only in the last 30 years [40], [41], [42]. Main purpose of pharmacological drug

treatments is to attenuate nicotine withdrawal syndromes that include anxiety, irritability, depression increased appetite, weight gain, restlessness, and decreased heart rate [10], [43]. Pharmacological treatments include nicotine replacement therapy (NRT) and bupropion (an antidepressant). Nicotine replacement therapy is the most commonly used pharmacological treatment, particularly among heavy smokers [44]. Five NRTs have been approved in the USA and other countries that include slow-acting nicotine patch, and other faster-acting formulations such as nicotine gum (polacrilex), nicotine inhaler, nicotine nasal spray, and nicotine lozenges [45], [42]. The faster-acting NRTs are effective to reduce craving whereas the slow-acting NRT supplies constant and low levels of nicotine which can relieve withdrawal syndromes [46]. These methods are popular in smoking cessation treatments because of their effectiveness, ease of use, availability, and low cost [47]. There have been numerous random control trial studies examining the effectiveness of NRT [48], [49], [50], [51], [52]. Compared to placebo, NRT produces a $1.5 - 3$-fold increase in smoking cessation rates [53], [54]. Although not an NRT per se, use of E-cigarettes or electronic cigarettes has gained popularity in recent years. It is a battery powered device that vaporizes a solution containing nicotine and propylene glycol. It has been marketed as a safer alternative to regular smoking and is used as a smoking cessation aid. However, the effectiveness of their use as a smoking cessation aid has not yet been assessed rigorously [55], [56].

Behavioral Treatments: Face-to-face counseling (in individual or group sessions) [57], [58], [59], [60] and telephone counseling [61] are the most prevalent behavioral treatment efforts in the United States [62], [63]. Other behavioral methods include use of educational and self-help materials in form of booklets or videos (also making them accessible over the internet), and computer tailored interventions that provide specific advice or plans for quitting [64], [63],

[65], [66]. Mobile phone text messaging is also used as a delivery method of behavior interventions in the form of motivational messages and behavioral-change support [67], [68]. It has been shown that combining behavioral and pharmacological treatments additively improves success rates and is currently considered as the gold standard for treatment [63].

Recent developments in research on smoking relapse process indicate that efficacy of smoking cessation programs can be significantly improved if the behavioral interventions are provided at the moment the smokers need help to avert relapse [12]. Research in smoking relapse process focuses on the identification of immediate motives or triggers for smoking lapse and the underlying processes that promote the progression from initial lapse to complete relapse. Evidence from studies based on self-report of smoking suggest that certain situations and contexts, such as negative affect, alcohol and coffee consumption, viewing other people smoking, are likely triggers or antecedents of smoking relapse [69], [70], [71], [72], [73], [74], [75]. However, the use of retrospective self-reports have been shown to be limited in their accuracy due to recall and subjective bias [76], [77], [78]. Ecological momentary assessment (EMA) approach attempts to avoid these problems by collecting real time data in the real world setting [79]. This approach uses devices, such as palm-top computers or smartphones that repeatedly prompts users to report their current experiences and behaviors. EMA is considered the current state of the art method used in smoking cessation and relapse research studies to collect ecologically valid repeated assessment of subjects behavior and experience.

Advent of wearable sensors and sensors on smartphones further enables continuous and unobtrusive collection of such data and reduce the burden on the subjects. Moreover, automated detection of smoking events and subject's experience and context from data collected from these sensors opens up the

opportunity to identify true antecedents of smoking relapse, perform rigorous analysis of influence of these antecedents on smoking relapse, and enable researchers have a better insight into the relapse process which can lead to new and improved smoking cessation treatments. In the next section we describe the current technological methods to detect smoking.

2.2 Technological Methods to Detect Smoking

Technological methods to detect smoking episodes include carbon monoxide (CO) monitors such as piCO$^+$, CReSS Pocket [19], Micro$^+$ [20], RespiTrace® [22], instrumented cigarette lighters [21], image processing [23], and accelerometers or Inertial Measurement Units (IMU). We can categorize the sensors used to detect or monitor smoking into two classes - external sensors and body worn sensors.

External Sensors: piCO$^+$, and Micro$^+$ are handheld devices designed for use as motivational aid in smoking cessation programs. They display the amount of smoke inhalation and carbon monoxide levels in a single breath exhaled, measured through a mouthpiece attached to the devices. They also calculate and display the percentage of carboxyhaemoglobin in the blood, thereby providing visual motivation for the smoker to stop smoking. These device are, however, not intended to be used for automatically detecting smoking in an operator independent fashion.

Clinical Research Support System (CReSS) Pocket/CReSS Micro [19] is a portable device that can be used to acquire the smoking behavioral information in the smoker's natural environment over weeks as they store the data on the device's memory. The subject is asked to insert a cigarette into a holder of CReSS and smoke through a mouthpiece attached to the device. The device then is able to compute several measures of smoking behavior including puffs per cigarette, puff volume, and puff duration, and also the timestamps of cigarette

insertion and removal. All this data can be downloaded later to a computer. Although CReSS has been used in some studies outside of the laboratory settings, it has been mostly in studies by tobacco researchers to establish brand differences [80], [81], [82] by observing the smoking pattern and the degree of tobacco intake. For example, it was observed that with light cigarettes, smokers take more frequent puffs in order to inhale the same amount of tobacco as in a heavy cigarette [81], negating the purpose of making lighter cigarettes. CReSS requires subject's compliance since each time they smoke, they need to smoke through CreSS. Furthermore, it may be embarrassing for the subjects to use it in the natural environment, since the device on their mouth will be visible to others in the vicinity.

Scholl et al. propose the use of sensors embedded in cigarette lighters to detect when the lighters are lit [21]. This is a very simple and effective method to detect the start time of smoking episodes. However, it cannot provide any information of puffing behavior (e.g., duration of smoking episodes and number of puffs per episode) and therefore has limited utility.

If the place of smoking is under the coverage of a video camera, then movement of hands and presence of cigarette in the mouth can be detected by image processing to automatically detect smoking [23]. Use of this method, however, requires installment of video cameras in all locations where a subject may smoke. Alternatively, the subject may be asked to have a portable video camera (e.g., on a smart phone) pointed to them before they smoke, which again requires the involvement of subject each time they smoke.

In summary, each of the above technological methods that involve external sensors require subject compliance and hence are not suited for widespread usage in smoking research in the natural environment. Therefore, smoking researchers continue to rely on self-report method today.

Body Worn Sensors: Respiratory Inductive Plethysmography (RIP) that uses a respiration band worn around the chest and possibly abdomen, wrist worn inertial sensors, and proximity sensors are examples of body worn sensors that have been used in the detection of smoking.

Respitrace® is a device that uses a RIP sensor, such as RespiBand Plus, that measures the chest's expansion as the wearer breathes in and out. The timing of each puff is marked manually by an observer who presses a push button switch when the subject places the cigarette on lips. Authors in [22] make use of these measurements to analyze post-puff breathing patterns in smoking. The use of Respitrace® has been restricted to lab settings to study smoking patterns since it requires manual marking of each puff and is not a portable device.

Over the past two decades there has been extensive research on human activity recognition using inertial sensors, specially accelerometers (see [83], [84] for surveys on activity detection from inertial sensors). Only recently, with the advent of consumer electronic products such as Nitendo Wii remote and smartphones there has been much interest in inertial sensor based hand gesture detection with the aim to facilitate gesture based human computer interaction [85], [86], [87], [88], [89], [90], [87], [91]. Wrist worn smartwathches that include inertial sensors provides the opportunity to detect hand gestures in the natural environment. There has been only a few works on smoking gesture detection from wrist worn inertial sensors. Below we provide a brief description of these works.

A promising recent approach of smoking detection based on hand gesture is presented in [24]. In this work, data from 15 volunteers are collected for a total of 17 smoking episodes. The smoking episodes included smoking alone, in a group while having a conversation and smoking while walking around. The volunteers wore a 9-axis inertial measurement unit (IMU) on the wrist for an average of 2 hours each. The IMU fuses information of a 3-axis accelerometer, a

3-axis gyroscope, and a magnetometer to provide position of the wrist in 3D space. Due to the use of magnetometers the trajectory computed from the IMU is in world co-ordinate system. This creates problem because the gesture trajectory becomes dependent on the orientation of the body. To address this issue, the points in the trajectory are transformed into points in the elbow's frame of reference producing a relative trajectory. Assuming the elbow is stationary during a smoking gesture, the relative trajectory then also becomes invariant to whether the person is smoking while being stationary or walking. The trajectory data is then segmented based on the idea that a gesture starts from a rest position (a low wrist velocity point) and ends at another.

Using the distance from the most recent rest point the whole segment is divided into 3 sub-segments by finding the peaks and troughs in the distance time series. A trough-peak-trough pattern is observed which corresponds to a typical smoking gestures where the hand goes through an ascending stage (trough to peak corresponding to rest position to mouth), then a stationary stage (hand is close to mouth), and finally, peak to trough (back to the rest position). Different duration, velocity and displacement features are computed from the sub-segments. Also different angle features are computed based on the roll and pitch measurements obtained from the orientation data.

The authors show that the use of a classifier (e.g., random forest) to make independent classification of these segments into smoking or not smoking, is outperformed by a linear chain conditional random field (CRF) that jointly classifies the whole sequence of segments. The CRF model is constructed based on the assumption that consecutive gestures are most likely to have the same rather than different labels. The CRF model achieves 10 fold cross validation accuracy of 95.74% where recall and precision are 81% and 91% respectively. However, in leave-one-out evaluation it is observed that for some participants the

recall goes to 0, which implies none of the puffs were correctly recognized. This is due to the fact that these participants had a completely different smoking gestures and temporal properties of these gestures compared to others.

Tang et al. present a wrist accelerometer based smoking puff detection method [25]. Smoking data from 6 participants are collected in a somewhat naturalistic setting. The participants smoked in the patio of a restaurant and the research stuffs annotated the timing of smoking, puff taking, posture of the participants (e.g., sitting, standing, and walking) and concurrent activities (e.g., talking, eating, drinking, and using the phone) performed by the participants while smoking.

This paper highlights the high variability observed in both with-in and between person smoking behavior characteristics such as duration of smoking, duration of puffing, number of puffs per smoking episode, and inter-puff duration. Such variability is most pronounced when a participant is also engaged in concurrent activities. Moreover, the puff signature in the accelerometer data is also observed to vary substantially depending on the body posture, concurrent and interleaved hand activities, and position of the hand at the start and end of a puff. The proposed method to detect smoking puffs involves segmentation of accelerometer data in fixed length sliding windows and feature extraction from the windows. Fifty one features are extracted from each accelerometers placed on each wrist. These features include statistical features (e.g, mean, standard deviation, kurtosis, skew etc.), SNR, and RMS of each window. Additionally, peak-to-peak amplitude, peak rate, correlation and crossing rate between accelerometer axes are computed. F1-score (harmonic mean of recall and precision) of 0.7 is achieved for the detection of puff by applying a Random Forrest classifier. F1-score of 0.75 and 0.40 achieved by individualized model and leave-one-out evaluations respectively, underscores the variability across

participants' smoking behavior. Also, puffs instances that come of the periods while the participant was engaged in concurrent activities were miss-classified most often.

The paper also presents a smoking episode detector that utilizes smoking topological information, namely puffing frequency which computed by dividing the number of puffs by the duration of a smoking episode. A Gamma distribution is fitted to the puffing frequency observed in the training data. In order to detect smoking episodes, puffing frequency of a certain window of time is computed by counting the number of puffs detected by the puff classifier within the window. Score of detecting smoking at a particular time point is computed by a weighted average of scores obtained from the Gamma distribution for the past 1, 4, and 8 minutes. An empirical threshold of 0.3 on this score is used to decide whether a smoking episode is present within the last 8 minutes. This method achieves an F1 score of 0.79.

Both of the above work only focus on hand gestures to detect smoking. As discusses above the presence of confounding events such as drinking and eating and the within and between subject variability of smoking behavior indicate that the gestures lack specificity in the real world. Therefore, a multimodal approach that fuses complementary information from different modalities has the potential to improve the accuracy and reliability of the detectors.

Multimodal Approach: Lopez-Meyer et al. [26] propose a multimodal smoking detector system that combines hand gesture detected by a proximity sensor and respiration measurements. A radio frequency transmitter-receiver proximity sensor is used to capture hand-to-mouth gestures. The transmitter is placed on the wrist while the receiver antenna is attached to the chest area. The receiver is tuned to the same operating frequency of the of the transmitter and hence can detect the proximity of the wrist to the mouth (within 30-17cm). A

hand-to-mouth gesture is detected whenever the amplitude of the received signal rises above and then drops below a certain threshold. Respiration measurements are captured by two respiration bands worn worn around the chest and abdomen. The tidal volume signal computed by averaging the signals obtained from the two respiration bands and airflow signal computed by taking the first derivative of the tidal volume signal are used as respiration features to classify smoking puffs. From each of the signals (i.e., tidal volume, airflow, and proximity signal) 500 samples from the start point of the detected hand gesture is used to construct a feature vector of length 1500. Three additional features - duration, average amplitude, and maximum of the hand gesture signal are used to form the the final feature vector. The classifier is trained using smoking (while sitting and standing) data and data collected from 10 other activities (e.g., sitting, reading, walking, eating, and drinking) in a lab setting. An average F1-score of 81.25% is reported for the leave-one-out evaluation of the SVM classifier.

This method has similarities to the method proposed in this thesis. However, the method presented in our work has several advantages. First, the proximity sensor is susceptible to electromagnetic interference and requires wearing an additional device (the receiver antenna and circuit) limiting the use the sensors in naturalistic settings. Second, our method proposes several respiration features that are interpretable and can be reused in the detection of other activities. Use of fixed number of samples from the raw signals as features does not provide similar interpretability. Third, across the 20 participants' data there were only 4,402 hand gestures and out of them only were 531 puffs. Use of a 1503 length feature vector may have over-fitted the training data and hence is responsible for the low F1 score for some participants.

2.3 Analytical Methods to Infer Human States from Respiration

Respiration sensors have existed for quite some time, most notably for use in sleep studies. With its integration in wearable and wireless sensor suites, its use have been explored for detection of various human behaviors from changes they induce on breathing patterns. These include works for inferring stress [27] and conversation [92]. In [27], respiration measurements are processed to infer physiological and perceived stress. Various features such as inhalation duration, exhalation duration, minute ventilation, and respiration rate are computed from each minute of respiration measurements. Machine learning models are then trained over these features to infer whether the subject is stressed. In [92], features used in [27] are supplemented with some new features such as B-Duration, computed from 30 seconds of respiration measurements to detect if the subject is speaking, listening, or quiet. These states are then composed together in a Hidden Markov Model (HMM) to identify conversation episodes.

While some features identified in the above works can help in detecting smoking puffs, the features and models used in these and other works are not directly applicable to detecting puffs for several reasons. First, robust statistics such as mean and median of various features are used in the detection of stress and conversation since there are several respiration cycles in 30 seconds or 1 minute. For smoking, each puff needs to be identified reliably. Second, there is a pattern of transitions among the speaking, listening, and quiet states that can be leveraged in an HMM to detect conversation episodes, such patterns are not observable in a smoking session. Third, the timing or number of respiration cycles between successive smoking puffs can vary widely among different subjects, and among different smoking session for the same subject, for instance, when smoking in a group or when smoking alone during work hours.

Fourth, the accuracy of detecting a smoking puff may need to be customizable for various use cases. A study on observing smoking behavior may want good accuracy for detecting an entire smoking session, whereas another study on abstinent smokers may want to detect individual puffs, at the cost of a higher false positive rate, since even a single puff can lead to a full relapse. The first puff in such scenarios constitutes the moment of lapse (also called first lapse) and is the main event which is used in the analysis for identifying antecedents and precipitants of smoking relapse. It is critical to be able to obtain the timing of the first lapse, and the entire subject data may need to be removed if the first lapse is not detected [15]. False positives may be acceptable, especially, if the analysis is to be done post-facto, where the goal is to be able to pinpoint the timing of relapse. Each puff may be presented to the subject to identify the one that may correspond to the actual first lapse. None of the above works present a way to customize the model to a given rate of false positive or false negative.

Fifth, none of the above works use a semi-supervised model to use the data collected from the field to improve the accuracy of the model. In a smoking research study, subjects are usually asked to mark each smoking episode in the field. These marks, however, may be before, during, or after the smoking episode. Some smoking episodes may have no marks at all. Hence, the marks provide a label for smoking episodes, but these labels are a noisy source. We develop a semi-supervised approach to make use of these noisy labels to improve the accuracy of our model. In summary, the requirements for the development of a smoking detector are more stringent than other works on detecting psychological and behavioral states from respiration measurements and hence require a new approach to model development.

In summary, the work presented in this dissertation advances the state-of-the-science in both improving the detection of smoking in the field

environment (with significant implications on smoking cessation research) and in detecting human behaviors from respiration measurements.

Chapter 3

mPuff: Detection of Smoking Puffs from Respiration Measurements

In this chapter, we present mPuff - a novel system to automatically detect smoking puffs from respiration measurements We introduce several new features from respiration that can help classify individual respiration cycles into smoking puffs or non-puffs. We then propose supervised and semi-supervised support vector models to detect smoking puffs.

3.1 Data Collection

In this section, we describe the AutoSense sensor suite we used to capture respiration measurements and the data collection procedure for collecting respiration measurements for developing, training, and testing the mPuff model.

3.1.1 The AutoSense Sensor Suite

We use the Autosense sensor suite [93] that includes a Respiratory Inductive Plethysmograph (RIP) band to measure relative lung volume and breathing rate. AutoSense also includes ECG, galvanic skin response and 3-axis acceletometer sensors, but only respiration measurements are used in this work. The sampling rate of RIP is 21.3 samples/sec. RIP uses a conductive thread that is sewn in a zigzag fashion to the elastic band. An alternating current source is applied to the resulting loop of wire, which, in turn, generates a magnetic field that opposes the current whose strength is proportional to the area enclosed by the wire according to Lenz's law. The ratio of the magnetic flux to the current is called self-inductance. Therefore, changes to the chest circumference can be measured by measuring the changes to the self-inductance of the band. The inductance measurement purely depends on the geometry of the band and is not related to the tension in the band. As a result, the measurement is not prone to the trapping of the band and associated artifacts due to changes in tension. The sensors transmit data to an Android mobile phone in real-time over a low-power wireless

Fig. 3.1: Respiratory Inductive Plethysmograph band (in blue color) is worn around the chest area and the wearable AutoSense sensor unit clips to the belt. A 3-axis Acceloerometer, ECG, and Galvanic Skin Response sensors are also included in the same sensor unit. Two coins (a quarter and a penny) are also shown in the picture to indicate the form factor.

link. We use the FieldStream mobile phone software available in [94]. Using the FieldStream software, we obtained the raw data files collected in the phone.

3.1.2 Data Collection for Model Development

To develop a model for smoking detection, we collected data from 10 volunteer participants over 13 individual smoking sessions. Each participant was a daily smoker. They wore the RIP sensor of the Autosense suite in their natural environment and were accompanied by an observer. The observer marked each puff the participant took by pressing a button on the mobile phone that also received the respiration measurements via wireless channel. The timing of the button press was saved. Marking the puff times on the same phone that received the measurements reduced the time lapse between the markings and sample timestamps. In order to get a more precise marking, the data from smoking session was visualized with the markings. The markings of puffs were then adjusted to match each puff, which is visually distinctive due to deep inhalation and exhalation associated with a puff.

Out of the 10 participants, 4 participants wore the sensor suite for 7 days in their natural environment during their awake hours. They were asked to self-report the each time they smoked a cigarette. We modified the interface of data collection software (FieldStream framework) on the phone to facilitate self reporting. Though the participants were asked to mark the smoking sessions as they light up a cigarette, we do not expect that the self reported times exactly corresponds to the beginning of the sessions; they can be anywhere (before, during, or after) in the vicinity of the smoking episode. Visual inspection confirmed this hypothesis. In total, we have 136 self-reports of smoking from these 4 participants.

3.1.3 Data Sets for Model Evaluation

We expect a smoking session to be confused with acute stress, conversation, and physical activity, since they all affect respiration measurements in a similar pattern as smoking. We call these events confounding events. In order to evaluate the model's performance on different confounding events, we use the data collected in our previous user studies [27], [92]. Both of these studies also made use of the AutoSense sensor suite. The first data set is from a study on 21 participants who were exposed to three real-life stressors (e.g. public speaking, mental arithmetic, and cold pressor tasks) in a lab setting [27]; successive stress periods were separated by rest periods. The second data set consists of conversation episodes from 12 participants. The start/end of each conversation episode as well as the start and end of speaking and listening periods in the conversation episode was marked. Total amount of data collected for this set is around 46 hours [92]. In addition to these two data sets, we also collected data from 5 participants that captured different levels of physical activities ranging from running and walking to sitting quietly. This data set consists of 1 hour worth of data.

3.2 Puff Detection

In this section, we describe the development of the mPuff model that classifies each respiration cycle into *smoking puffs* and *non-puffs*. We first run a peak-valley detection algorithm to find the peaks and valleys in each respiration cycle. We use the peak-valley detection method proposed in [92]. Once peaks and valleys are located, features of interest can be computed to use in a classification algorithm. In Section 3.2.1, we describe the features we identify for use in puff detection. We then describe (in Section 3.2.2) the development of a supervised classifier that uses labeled data to detect smoking puffs. This model can be configured for prescribed target false positive or false negative rate. We next describe (in Section 3.2.3) a semi-supervised model that uses self-report labels obtained from field data to improve the accuracy of puff detection.

3.2.1 Feature Computation

Other works that use respiration measurements to infer human states such as [27], [92] make inferences on time windows that are 30 seconds or longer. Hence, these works are able to compute statistics over multiple respiration cycles making them robust to noise and outliers. In contrast, in this work, we need to make accurate decision at the level of an individual cycle since a smoking puff lasts only one respiration cycle. This makes the task of identifying appropriate features more challenging.

We identify 17 distinct features that are computed from the respiration signal. We identify 5 features from existing work and propose 12 new features. The new features are selected through visual inspection of data collected during smoking and other non-smoking episodes such as physical activity, stress and conversation. Computation of the features involves the identification of the respiration cycles, which are composed of an inhalation and an exhalation period.

We now define all features in the following and illustrate three of them in Figure 3.2.

Existing Features: We first describe five features that have previously been proposed for identifying stress and conversation events from respiration [27], [92]. **Inhalation Duration** corresponds to the time elapsed from a valley of a respiration cycle, to the subsequent peak. The amplitude difference in signal values between these points is the maximum expansion of the chest during a respiration cycle (see Figure 3.2). **Exhalation Duration** corresponds to the time duration between a peak and the subsequent valley in a respiration cycle. **IE Ratio** is defined as the ratio of inhalation duration to the exhalation duration in a respiration cycle. **Respiration Duration** is the sum of inhalation and exhalation duration. **Stretch** is defined to be the difference between the maximum (legitimate) amplitude, and the minimum (legitimate) amplitude the signal attains within a respiration cycle (see Figure 3.2).

These features have been shown to be effective in identification of stress and conversation from respiration [27], [92]. As we show in Section 3.3, using these features provides an accuracy of 73.55%, which improves to 86.7% once the new features described below are added.

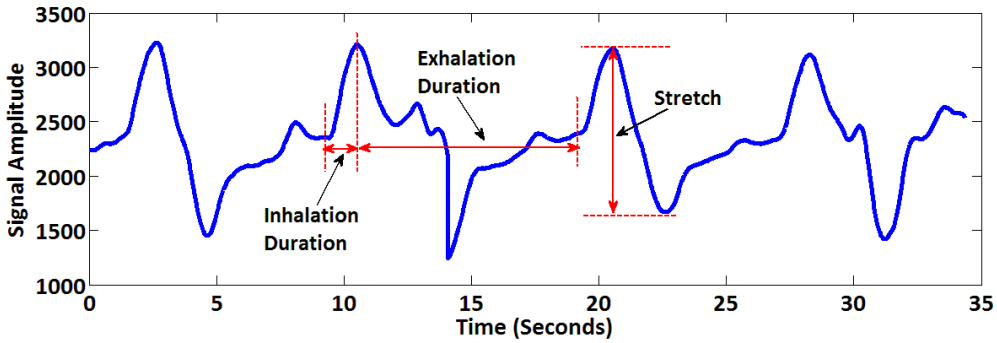

Fig. 3.2: Illustration of three features extracted from respiration cycles.

New Features: Figure 3.3 shows the respiration patterns during smoking and three confounding events (stress, conversation and running). We observe that features such as stretch and expiration duration are distinct during a respiration cycle containing a puff as compared to respiration cycles in speaking, stress, or activity, and hence have discriminatory power. We further observe that the relative change in stretch and exhalation duration from one respiration cycle to the neighboring cycles are higher when there is a puff involved during smoking. On the other hand, we do not see such magnitude of change during running or conversation events. This is because it is quite unusual to take two consecutive puffs without taking any breaths in between. We define the first difference feature and ratio features in order to capture these information concerning relative change.

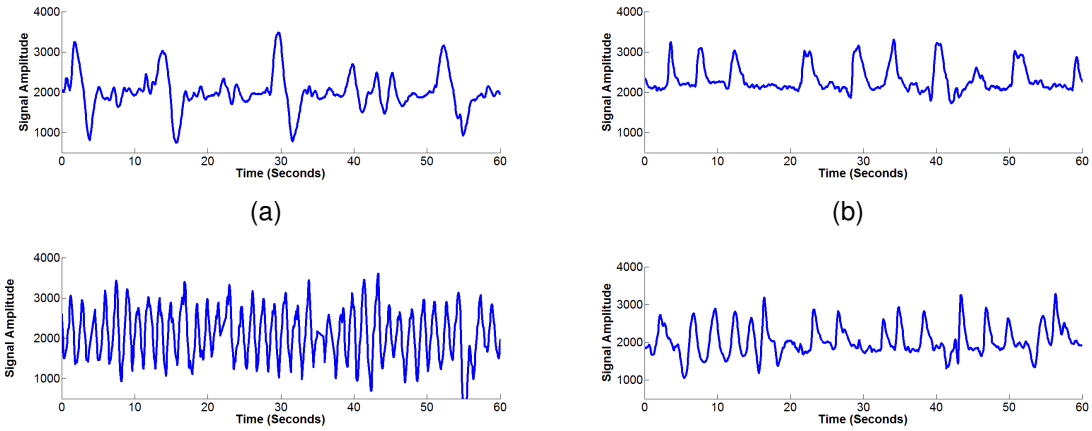

Fig. 3.3: The four figures above show the respiration signal during smoking and three confounding events. We observe that the stretch of a respiration cycle is higher during running and puffs, as compared to speaking and stress. We further observe that unlike during running events, during smoking sessions, the non-puff cycles around the puff do not have as high of a stretch. This simple visual inspection suggests the use of change in stretch relative to its neighboring cycles in discriminating puffs from conversation, stress, and running events. Similar observations can be made for other new features such as relative change in exhalation duration, and upper and lower stretch as described in Section 3.2.1.

We also observe that during a puff, the respiration signals stretch in both upward (called *Upper Stretch*) and downward directions (called *Lower Stretch*), extending the peak amplitude and reducing the valley amplitude respectively, as compared to usual respiration cycles. This suggests that the measurement of relative change in the upper and lower stretch as compared to the running mean of the valley in respiration cycles, can further improve the accuracy of identifying smoking puffs. We now describe the 12 new features.

- **Forward and Backward First Differences** of a feature is derived by computing the first order differences of the feature values from their previous and next feature values respectively. We compute these first order differences for inhalation, exhalation and respiration durations and stretch and use them as features. Altogether, this procedure creates a total of eight new features.

- **Stretch Ratio** of a particular cycle c, is defined as the ratio of its stretch to the average stretch value in a window of five cycles, with the window centered on cycle c. When computing the average, we exclude cycle c. We use a window of five cycles because we never see successive puffs occurring in a window of five cycles, if the window is centered on a puff cycle.

- **Exhalation Ratio** of a particular cycle is similarly computed from the average exhalation duration in a window of five cycles.

- **Upper and Lower Stretch** values are the two features computed from the stretch of each cycle, by splitting it into two parts. The upper stretch magnitude is computed by taking the difference of peak amplitude and running mean value of the valley amplitudes of signal cycles ($ValleyMean$). Similarly, the lower stretch magnitude is computed by taking the absolute

difference of minimum amplitude in a respiration cycle and $ValleyMean$. During the computation of the running mean, $ValleyMean$, any valley amplitude two standard deviations away from the current mean value is discarded in the computation.

In order to visualize the effectiveness of the features, we use the **Fisher's Linear Discriminant method** [95] to project the 17-dimensional data vectors to a single dimension using $y = \mathbf{w}^T x$. The idea is to adjust the components of $\mathbf{w}$ in such a way that the projection maximizes the class separation. The discriminant method maximizes the difference between the projected class means while minimizing the projected class scatter. This method can actually be used as a classifier albeit a weak one as much of the information inherent in the data gets lost in the projection. It, however, provides an easy way to visualize the separation of the classes and separation of the classes in the one dimension does hint to the fact that the features may be promising for classification in the higher dimensional space. Figure 3.4 shows the projected data points and it can be readily observed that there exists good enough separation of the puff and non-puff classes.

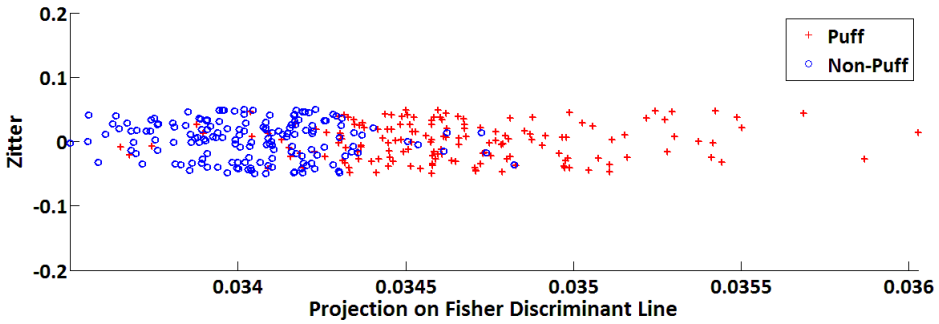

Fig. 3.4: Projection of data on the Fisher's Linear Discriminant Line. The y-axis is used to spread the data points on the line for visual aid. The x-axis presents the value of the projection.

3.2.2 Classifier: Supervised Learning Model

The supervised classifier we use to detect individual puffs is learned using a supervised learning algorithm. The standard classification supervised learning framework is formulated as follows. Given an example set of input observations $X = \mathbf{x}_i \in \mathbb{R}^n | i = 1 \ldots N$, e.g., matrix of n RIP features by N inhalation cycles, with corresponding class labels $Y = y_i \in +1, -1 | i = 1 \ldots N$, e.g., puffs and not puffs, the training algorithm learns a classification function $f_\alpha(\mathbf{x})$, which estimates the class label $\hat{y}$ for a given unlabeled/novel input observation $\mathbf{x}$. The function f is parameterized by the parameter vector α.

Our puff from non-puff classification function is learned using the industry-standard Support Vector Machines (SVM) algorithm, which has been demonstrated to be highly competitive on a great number of problems and tasks, owing to its combination of high learning capacity, i.e., ability to learn highly complex classification functions, with a mathematically rigorous handling of the overfitting/training error trade-off via regularization in the space of kernel functions. Formally, the SVM classification function is defined, using the dual formulation, as

$$\begin{aligned} f_\alpha(\mathbf{x}) &= \operatorname{sign}\left(g_\alpha(\mathbf{x})\right). \\ g_\alpha(\mathbf{x}) &= \sum_{i \in X_{\mathtt{SV}}} \alpha_i K(\mathbf{x}_i, \mathbf{x}) + b. \end{aligned} \tag{3.1}$$

In the above formulation, α_i's and bias constant b are the parameters learned in the course of training, and the set $X_{\mathtt{SV}}$ contains the training observations, called support vectors, that define the boundary separating the two classes. The function $K(\dot{;})$ is the special kernel function that allows SVM to learn highly complex functions f, corresponding to a highly non-linear separation

boundaries. This so-called kernel-trick makes it possible to implicitly transform all observations to a space of much higher dimensionality, called kernel feature space, where difficult problems are simplified. Formally, the kernel function $K(\mathbf{u}, \mathbf{v})$ corresponds to a dot product between the original vectors $\mathbf{u}$ and $\mathbf{v}$ in the kernel feature space, making the explicit transformation unnecessary. One powerful class of kernels is the Radial Basis Functions class: $K(\mathbf{u}, \mathbf{v}) = e^{\gamma\|\mathbf{u}-\mathbf{v}\|^2}$, leading to classification functions capable of classifying very complex datasets.

The training formulation is based on regularized empirical risk minimization, whereby the algorithm minimizes the error on the training observations, while minimizing the L2-norm of the function f, which works to minimize its complexity. The formal definition of the SVM primal learning problem is

$$\underset{\mathbf{w},\mathbf{b}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i}^{N} \max(1 - y_i(\mathbf{w}^T\mathbf{x}_i + b), 0) \tag{3.2}$$

where the primal variables/parameters w and b represent the linear decision boundary. Note that minimizing the L2-norm of f is equivalent to maximizing separation margin. The more useful formulation, however, is the Wolfe dual, which also specifies the implicit transformation of the problem into the kernel feature space.

$$\begin{aligned} &\underset{\alpha}{\operatorname{argmax}} \sum_{i}^{N} \alpha_i - \frac{1}{2}\sum_{i,j}^{N} \alpha_i\alpha_j y_i y_j K(\mathbf{x}_i, \mathbf{x}_j) \\ &0 \leq \alpha_i \leq C, \\ &\sum_{i=1}^{n} \alpha_i y_i = 0. \end{aligned} \tag{3.3}$$

The algorithm requires tuning of the two parameters: 1) the C constant, which directly penalizes the error on the train set, as per equation (3.2), by which it indirectly manages the trade-off between overfiting and training error, and 2) the choice of kernel function, along with any constants in it. We use the RBF kernel, and vary the γ hyper-parameter, in conjunction with the C constant, in order to attain the best performing function f. As per standard practice, we defined a set of candidate C's and a set of candidate γ's, and try all combinations of values from these sets. For both C and γ, the candidate values ranged from 2^{-10} to 2^{10}, increasing in steps of $2^{0.5}$.

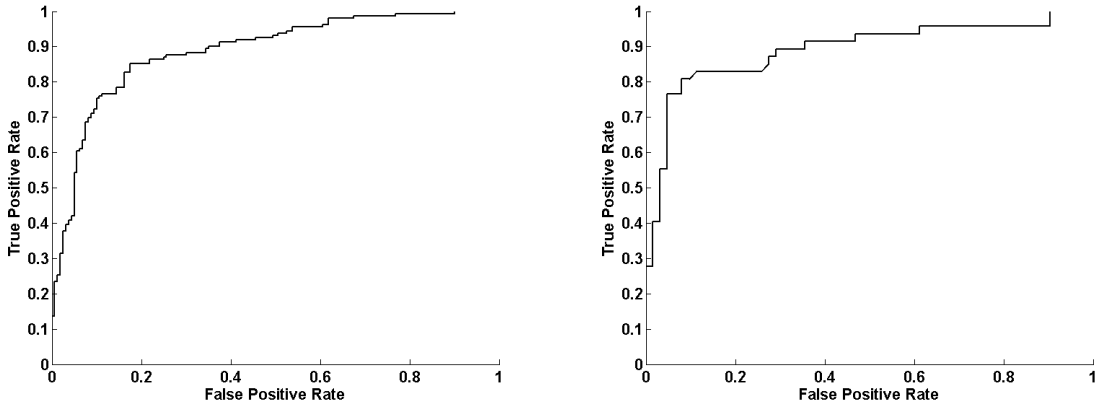

(a) Cross-validation ROC curve (b) Test ROC curve

Fig. 3.5: ROC curves plotted for (a) cross-validation with optimal (C, γ) = (1.415190 , 11.314685), average accuracy of 83.85% and AUC = 0.88, (b) the withheld test set with accuracy = 86.23% and AUC = 0.89. Optimal (C, γ) = (2.829404 , 16.971539).

We follow the standard stratified cross-validation approach to evaluate the learned model, with the number of cross-validation partitions equal to 10. Each fold contains roughly the same proportions of the two types of class labels. For each pair of (C, γ) values we run cross-validation and use the accuracy metric to assess the performance of the model. The optimal pair of (C, γ) is found to be

(1.415190,11.314685). We also split data into training and testing datasets. C and γ values are tuned and the model is trained using the training dataset. The optimal (C, γ) = (2.829404,16.971539). We run the fine tuned trained model on the withheld test dataset. Figure 3.5 contains the cross-validation ROC curve and the test ROC curve showing the performance of the trained model on withheld test dataset.

ROC curve shows the true positive rate against the false positive rate at various SVM score thresholds. In order to build a model with a target true positive or false positive rate, we train an SVM classifier using the entire dataset and select the threshold on the SVM scores that produces the target true positive or false positive rate.

3.2.3 Using Unlabeled Data: Semi-Supervised Model

In this section, we describe a specific semi-supervised model that makes use of unlabeled data obtained from the field data collected from the four participants (see section 3.1.2). Semi-supervised learning has been shown to improve the classification accuracy when there is a scarcity of labeled data but there is a large amount of unlabeled data. We use the Semi-supervised Support Vector Machine (S3VM) model [96] that extends the basic supervised SVM to incorporate unlabeled data.

The basic intuition behind S3VMs is that if we have unlabeled data together with labeled data the decision boundary obtained by the learning algorithm should be such that it separates the labeled data with a maximal margin, while simultaneously maximizing its distance to unlabeled examples. The second part in the model formulation is motivated by the notion that the model should have as little ambiguity in classifying the unlabeled examples as possible, even if there's not assurance that these classifications are correct (given that

there are no labels for them). This intuition is incorporated by modifying the objective function of the basic SVM, and is given by

$$\begin{aligned}\underset{\mathbf{w},\mathbf{b}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i}^{N}\max(1 - y_i(\mathbf{w}^T\mathbf{x}_i + b), 0) \\ + C^*\sum_{j}^{M}\max(1 - |\mathbf{w}^T\mathbf{z}_j + b|, 0)\end{aligned} \tag{3.4}$$

where, $\{\mathbf{z}_j \in \mathbb{R}^n | j = 1 \dots M\}$ is a set of unlabeled input data. We essentially add penalty in the objective function for the unlabeled data points that are too close to the decision boundary, specifically for which $-1 < |\mathbf{w}^T\mathbf{z}_j + b| < 1$, thereby, forcing the decision boundary to go through a low density area of the unlabeled instances. The S3VM experiments were conducted with the SVMlin toolbox [97].

The challenge in developing the semi-supervised learning model is to identify a feature that can connect the self-report to the smoking puffs, knowing that the self-report can be before, during, or after a smoking episode. The new feature we identify is the time distance of the respiration cycle in consideration to the closest self-report timestamp in the field data. As there should exist a marking before, during, or after every smoking session reported, the time distance from each cycle to the nearest self-report time should help the learning algorithm. Note that we do not have actual self-report time for the labeled data, but we can reasonably assume the existence of a hypothetical accurate self-report at the beginning of the carefully labeled smoking sessions (see section 3.1.2).

3.3 Training & Evaluation

In this section, we present the performance of the classifiers for detecting smoking puffs. The training data set for the supervised classifier is comprised of

the instances of puffs cycles and instances of non-puff cycles from the smoking sessions. The other sources of non-puff instances are data from the stress, conversation, and physical activity data sets. These data sets do not include any smoking events. Inclusion of these data sets are required in order to create a robust classifier that should generalize to the natural field environment better than a classifier that uses only the puff and non-puff cycles from the smoking session. Moreover, as stress, conversation and physical activity have been shown to be inferrable from respiration, they form the set of plausible confounding factors in smoking detection. In total, we have 161 puff instances and the same number of non-puff instances. The non-puff instances with equal proportion come from smoking sessions and the 3 confounding data sets.

The training data set for the semi-supervised model set includes the same labeled data set as described above and a large amount of unlabeled data taken from the field data sets obtained from 4 participants. However, not all data from the field data is included as it amounts to 28 hours worth of data. Such a large amount of data proved to be infeasible to run on the SVMlin tool. Experimenting with different amounts of data, we ended up including 10 times the amount of cycles in the labeled data. We ensured that data was included both from the neighborhood of self-report times (± 5 minutes) and far away of from any self reports, thereby ensuring that the unlabeled data included both puffs and non-puffs.

Training and Testing Performance: Tables 3.1 and 3.2 present the performance of the classifiers on labeled data set. Table 3.1 shows the performance of the classifiers on whole data set with 10-fold cross validation. As mentioned above for supervised classification, we use Support Vector Machine (SVM). Also for greater generalizability, we split our labeled data set into training and test sets — one subset contains 66% of the whole data set and other subset

contains the remaining 34%. For this case, the classifier performance on the test data is presented in Table 3.2. Data used in the supervised classifier are carefully labeled. But, when we add the noisy labeled data set from the field, we use the S3VM classifier which is a semi-supervised support vector machine. For the first experiment, when we use all the labeled data, SVM provides 83.85% accuracy; S3VM is able to improve this accuracy to 86.7%.

Table 3.1: Performance of classifiers with 10-fold cross-validation.

Classifier	Accuracy(%)	Precision	TP rate	FP rate
SVM	83.85	0.83	0.85	0.17
S3VM	86.7	0.91	0.81	0.08

In the second experiment, when we apply the trained classifier models on the testing data set, we observe the performances of the classifiers are also similar to the training accuracies (as shown in Table 3.2). In both of these experiments, we get high precision, high recall or high true positive rate. Moreover, if the training set consists of only the non-puffs taken from the smoking session the testing accuracy is 91.43% for the SVM classifier. This indicates that the classifier is quite efficient in detection of puffs and non-puffs in the absence of confounding factors.

Table 3.2: Performance of classifiers on test dataset when data is split into training and test sets.

Classifier	Accuracy(%)	Precision	TP rate	FP rate
SVM	86.23	0.86	0.81	0.097
S3VM	87.27	0.91	0.83	0.08

Figure 3.6 presents the accuracy incurred when the puff detection model is run on the stress, conversation, and physical activity datasets. We note that that

these datasets do not include any smoking sessions and therefore respiration cycles not correctly identified implies it was detected as a puff. We observe that conversation is the most confounding event for the puff detection model. One reason is that during speaking, we tend to take deep breaths at the beginning of the speech and that is sometimes detected as puffs.

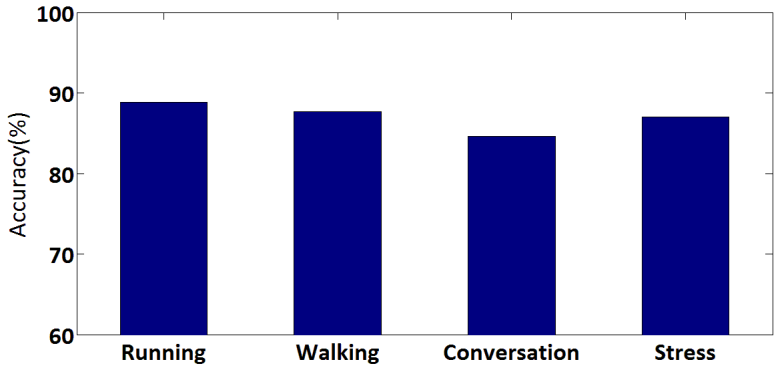

Fig. 3.6: Accuracy of the classifier on the datasets for four confounding events.

Feature Analysis: The performance of the classifiers presented in the Tables 3.1 and 3.2 uses all the 17 features. Here, we present the effect of adding the 12 new features on the performance of the classifier. We partition the set of all the new features into 3 sets, namely, the set of first differences, the set of the stretch and exhalation duration ratio, and the set of Upper and Lower stretches. All the accuracy values reported are for the SVM classifier. The basic set of features comprising only the existing 5 features produces an accuracy of 73.55% on the whole labeled data set. From Table 3.3, we observe that among the three new sets of features, adding the first differences to the basic set improves the accuracy most, reaching up to 81.1% . With this set of features, adding the stretch and exhalation duration ratios increases the performance of the classifier most. With these 2 new sets of features, we obtain an accuracy of 82.7%.

Table 3.3: Effect of new features on the classification accuracy. The classification accuracy obtained using only the existing features is 73.55%. They constitute the basic set of features. Three different sets of features are introduced in this paper. Let S_1 = set of forward and backward first differences of basic features, S_2 = {Stretch ratio, Exhalation ratio} and S_3 = {Upper Stretch, Lower Stretch}. X denotes the inclusion of the set to the basic set of features.

S_1	S_2	S_3	Accuracy(%)
X			**81.1**
	X		76.33
		X	77.86
X	X		**82.7**
X		X	81.32
	X	X	75.79

3.4 Conclusion

As a first step towards building a reliable smoking episode detector, in this work we presented a model to automatically detect smoking puffs in the natural environment from respiration measurement. We achieve 87% accuracy on the detection of puffs even when there exists potential confounding events in the collected data. For the purpose of building the model, we identified 12 new respiration features that are found to be effective compared to the use of only the existing respiration features available in the literature. We also presented a semisupervised moded that improves the accuracy of the model when we provide unlabeled data with self reports that is also collected in the natural environment from participants who are daily smokers.

Chapter 4

Combining Hand Gesture and Respiration to Improve Puff Detection

Although the current puff detection model described in chapter 3 achieves 87% classification accuracy, it can be readily seen that it is bound to produce a large number of false positive and false negatives when it is applied to participant's data collected in a whole day. This is because there are, on average, 14 respiration cycles per minute, which results in more than 8,500 cycles in the usual 10-11 hours of wearing time per day. Then the total number of misclassified cycles amount to more than 1,000 per day. As confounding factors (e.g, physical movement, conversations) in the natural environment significantly alter the respiratory pattern during smoking, we, therefore, believe a more robust detector is required in order to detect smoking episodes in the field. In this chapter, we present a multi-modal method that combines both respiration and hand gestures measurements to detect smoking episodes in the field. First, we describe the label data collection procedure and then the method to fuse the multi-modal measurements.

4.1 Data Collection

In order to collect both respiration and hand gesture data we used the newer version of the AutoSense sensor suit and a pair of wristbands (one on each wrist as participants can smoke with either hand).

Enhanced AutoSense Sensor Suit: It has a thinner packaging and is thus more comfortable to wear. It uses the same chest band to capture the respiration data at 21.33Hz. These samples are transmitted wirelessly using ANT radio to a Sony Ericsson Xperia X8 smart phone at the rate of 28 packets/second. Each of the packets are 8 bytes long and contains 5 samples. The sensors now last more that 10 days on a 750 mAh battery.

Wristband: The participants wore a pair of new wristbands, one on each wrist, that includes a 3-axis accelerometer and a 3-axis gyroscope, a 3-axis magnetometer, and two ambient-light sensors and a temperature sensor. The accelerometer (ADXL335) and gyroscope (ST Microelectronics A3G4250D) will be sampled at 10 Hz at each axis to capture fine-level wrist motion, helping us improve automated detection of smoking by capturing smoking gestures of smoking arms. However, we have turned off the power hungry magnetometer (Honeywell HMC5883L) sensor in order to preserve battery life and have the sensor last 3 days. The two light sensors (Hamamatsu S1087 and S1087-1) capture light energy in different parts of the visible spectrum (320-730nm and 320-1100nm), allowing us to differentiate between outdoor (natural light) and indoor (artificial light) environments. In this work, we use the measurements obtained from the accelerometers and gyroscopes only.

4.1.1 Data Collected

We collected data from 6 volunteer participants who are regular smokers. The participants wore the sensors for a total of 40.3 hours in their natural environment as they went on performing their daily activities. Each time they smoked they were accompanied by an observer who marked each puff by pressing a button on the phone. As it is difficult to mark the precise start and end of hand gesture during smoking, the observer was instructed to mark the puff when the participant held cigarette between the lips and inhaled smoke. From the marking, we therefore, obtained the timing when the hand is at mouth during smoking. The data we obtained contains 32 smoking episodes and 470 puff marking. In 179 instances out of the 470 puff markings, there were data loss due to wireless transmission losses and noise due to physical movement or loosening of the respiration belt. As a result, we obtained 291 puff instances for which we

have respiration and wrist sensor data. The 32 episodes included smoking while standing, sitting, walking, and being in a conversation.

4.2 Detection of Hand-at-mouth Gestures during Puffs

Hand gesture during puffing a cigarette is typically composed of three sub-gestures occurring in the following sequence - hand moving to mouth, hand being at mouth while taking a puff, and hand moving away from the mouth. In this section, we describe a method to locate puff events by detecting the segments in the wrist sensor timeseries that contain hand-at-mouth gestures.

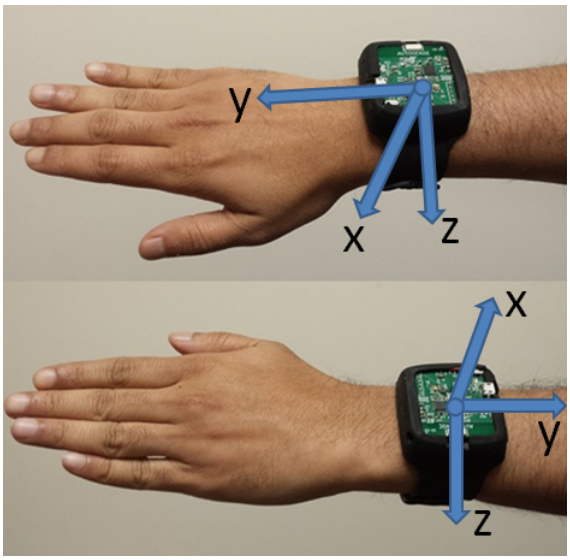

Fig. 4.1: Autosense Wristband.

The wristband contains a 3-axis accelerometer and a 3-axis gyroscope providing us with 6 timeseries. The x, y and z axes of the 3-axis accelerometer and 3-axis gyroscope on the wristband are aligned with each other. The directions of the axes of the wristband sensors are shown in Figure 4.1. When an accelerometer is stationary, any accelerometer axis aligned exactly with the earth's downward gravitational field will result in a measurement of -1g in that axis. So, we observe a positive value on y-axis when the right hand is held in an

upward direction whereas we get a negative value on the y-axis of the left wrist accelerometer. However, when an accelerometer is moving an axis will measure the combination of linear acceleration due to movement and the component of gravity in the direction of that axis. The gyroscope axes, on the other hand, measures the rate of change of rotation or angular velocity about the axes. We use the measurements obtained from both the accelerometers and gyroscope for the detection of hand-at-mouth gestures.

4.2.1 Detection Method

One approach to detect the complete smoking hand gesture from inertial measurements, as proposed in [24], is to track the trajectory of the hand using 3D orientation data obtained from a wrist worn 9-axis inertial measurement unit (IMU). In this work, we propose a simpler method to locate puffs based on the observation that during inhaling smoke from the cigarette the hand is at mouth and for a couple of seconds the hand must remain stationary at that position. We argue that this method of locating puffs by detecting hand-at-mouth gesture is more robust than tracking the hand trajectory which can vary substantially depending on the body posture (i.e., sitting, standing, and walking) and position of the hand as the hand starts moving towards the mouth. In order to detect hand-at-mouth gesture, first, segmenting of the sensor data from both wrists to relatively stationary and non-stationary segments and discarding all non-stationary segments. For this purpose we utilize the gyroscope measurements which, unlike accelerometers, is only affected by movement of the hands and not gravity. However, there will be many such segments where the hand is relatively stationary compared to its surrounding non-stationary segments. We, therefore, screen out segments that are unlikely to contain hand-at-mouth gestures based on their duration and the orientation of the hand during these segments. Orientation of the hand during these relatively stationary segments is

obtained by the roll and pitch angles computed from the accelerometer measurements. Detailed description of detection of stationary segments and the screening procedure is presented below.

In order to find the location of where the hand is relatively stationary, we use the magnitude of the gyroscope axes. Any movement of the hand will manifest as a rotation about one of the gyroscope axes and cause the magnitude value to increase independent of the direction of rotation. When there is very little movement of the hand the gyroscope magnitude will be low since it is not affected by gravity. Therefore, a hand-at-mouth gesture can be detected by finding segment is where the magnitude timeseries attains low values and is preceded and followed by peaks. The first peak is due to the hand moving towards the mouth and the second one is due to the hand moving away from the mouth. However, it should be noted that simply thresholding the magnitude values to locate the segments does not work well in practice. The average magnitude of a hand-at-mouth segment during walking will be higher than that of stationary segments during standing or sitting. This is because, when a person is walking the whole body is moving and there will always be some movement of the hands. Even when the hand is at mouth while taking a puff, there is some movement of the wrists due to taking steps. Therefore, for the hand-at-mouth gestures during walking we expect to find segments that attain low magnitude values only relatively compared to the average magnitude during walking. Moreover, in many instances we observe that the these relatively low amplitude values are sometimes higher than the amplitude of peaks corresponding to the hand movements before and after the hand-at-mouth stationary segments. The amplitude of these peaks depend on the rest position of the hand before and after the puff. While standing or walking usually the hand hangs beside the body and therefore we observe larger peak amplitude. On the other hand, while sitting, the

hands may be resting on the thighs, or on the armrest of a chair. In these cases, the hand is moving a shorter distance to the mouth and hence the peak amplitudes are lower. Also, in many instances the participants were observed to smoke with their hand hanging near the mouth and in these instances the amount of movement is even lower producing the lowest peak amplitudes. Therefore, setting the threshold value as high as the amplitude hand-at-mouth segments during walking will not correctly identify hang-at-mouth segments in many other instances. This necessitates a segmentation method that is adaptive to the current level of movement to find segments that are relatively stationary. The procedure to find relatively stationary segments is described below.

This procedure makes use of a two moving averages based method to detect rise and fall in the gyroscope magnitude timeseries. Such methods are commonly used to compute indicators used by investors in stock markets to identify price rise and fall [98].

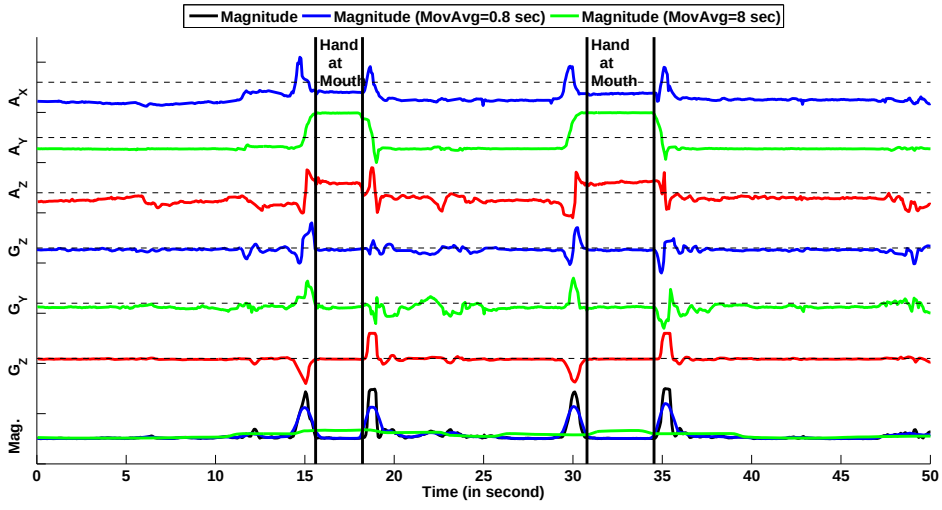

Fig. 4.2: Hand at mouth segment detection. (A_X, A_Y, A_Z) and (G_X, G_Y, G_Z) present the signals of accelerometer and gyroscope x, y, and z axes. The magnitude of the gyroscope and the two moving average signals are shown last.

- Compute a fast (0.8 second window) and a slow (8 second window) moving averages of the gyroscope magnitude. The fast moving average closely follow the dynamic nature of the signal while the slow moving average represents the level of movement in the neighborhood (see Figure 4.2). The window size of the slow moving average is corresponds to the length of a smoking hand gesture length which usually lasts from 3-7 seconds [99]. Because of this window size, during smoking hand-at-mouth gestures the slow moving average is computed over windows that always contain the peaks due to hand moving towards and away from the mouth. Therefore, during these segments the slow moving average attains higher values compared to the fast moving average. The 0.8 second window size for the fast moving average is empirically chosen so that all segments that contain hand-at-mouth gesture during puffs are detected.

- Select the segments where the fast moving average lies below the slow moving average. These segments are demarcated by two consecutive crossing over points of the two moving averages. The first of them corresponds to the location after which the fast moving average moves below the slower one and the second crossing over point corresponds to the location after which the fast moving average rises above the slower one. The selected segments are, therefore, the ones where the magnitude values are relatively lower than the average magnitude of the neighborhood.

The method described above essentially over-generates segments that may contain a hand-at-mouth gesture. Before building a model we screen out segments that have properties substantially different from the segments corresponding to puffs.

Screening: We find the segments that corresponds to hand-at-mouth gestures during puffs by checking whether the segment contains a puff marking.

We call these segments puff segments. We use properties of these segments to discard some of the other segments that are generated. First, we remove the segments that are either too long or too small in duration compared to puff segments. We discard all segments that have duration more that three standard deviation away from the mean duration of segments corresponding to puffs. Second, we compute the mean distances between the fast and slow moving averages for each of the puff segments. Use the minimum of these distances as a threshold (t_m) and discard all segments for which the mean distance is lower than the threshold.

Screening based on Hand Orientation: For each segment we find the orientation of the hand by computing the pitch and roll angles. The pitch and roll angles at a particular orientation indicate the amount of rotation about the x and y axis respectively required to reach the particular orientation from an initial orientation. We assume the in the initial orientation, the hand is kept horizontal with the palm facing down (z-axis is aligned with the gravitational field). Roll and pitch angles can be computed from either accelerometer or gyroscope. However, in the presence of linear acceleration the orientation angles computed from accelerometers are usually less accurate. On the other hand, roll and pitch angles can be computed by integrating the angular velocity measurements obtained from gyroscope. However, a small error in angular velocity measurement gives rise to large integration errors. In order to obtain better accuracy, accelerometer and gyroscope measurements are therefore fused using a Kalman filter [100]. However, since we are interested in computing the orientation when the hand is relatively stationary, relying on only accelerometer measurements are sufficiently accurate for our purpose. For each segment we compute the average of each axis forming the vector (a_x, a_y, a_z). Following the method proposed in [101], we compute the pitch(θ) and roll(ϕ) angles using the following equations:

$$\tan\theta = \left(\frac{-a_y}{-a_z}\right) \tag{4.1}$$

$$\tan\phi = \left(\frac{a_x}{\sqrt{{a_y}^2 + {a_z}^2}}\right) \tag{4.2}$$

Since the direction of x-axis (and y-axis) of the accelerometer on the left and right wrists are opposite to each other, we first negate the x and y axis measurements of the left wrist sensor. In this way, the corresponding the axes point in the same directions in both wrist sensors. By convention the roll and pitch angles are positive when there is clockwise rotation about the y-axis and x-axis respectively. Therefore, when either arm is lifted from the horizontal position pitch angles will have the same negative sign. However, inward (or outward) rotation of the left hand is in the opposite direction for an inward (or outward) rotation of the right hand. Therefore, in order to obtain same sign of the roll angles in both hand for orientations that are only mirror images of each other, we negate the sign of the roll angles obtained for the left hand.

The scatter plot of roll and pitch angles for puff segments (red dots) and other segments (blue dots) is presented in Figure 4.3. We observe that the roll and pitch angles are slightly correlated with each other. For each segments we compute the Mahalanabis distance from the distribution of roll and pitch angles of puff segments. Mahalanabis distance d is computed by, $d = (\mathbf{x} - \mu)S^{-1}(\mathbf{x} - \mu)'$, where $\mathbf{x} = (x_{roll}, x_{pitch})$ is the vector representing the roll and pitch angles of a segment, μ and S are the mean vector and co-variance matrix computed from roll and pitch angles of puff segments. Since for the puff segments the distance

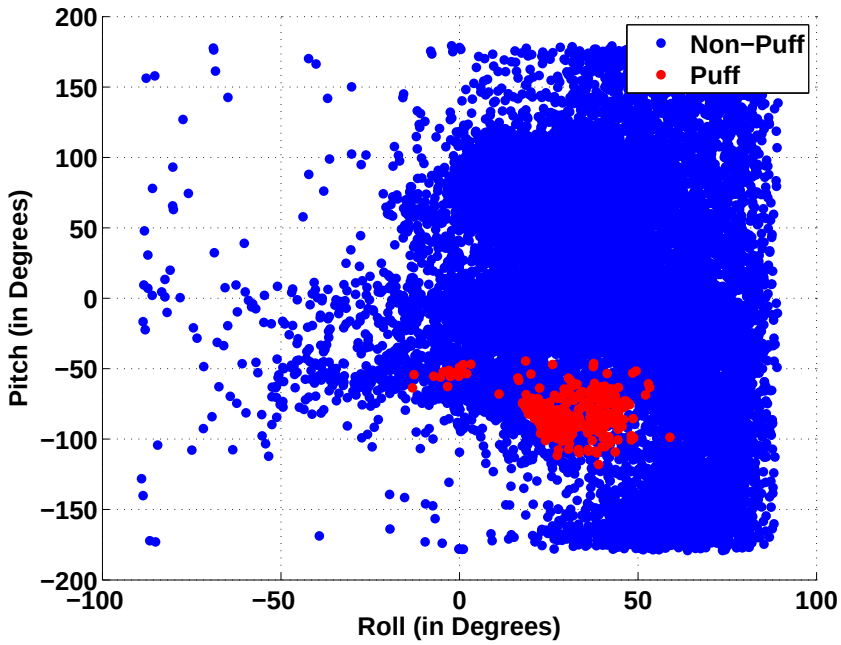

Fig. 4.3: Scatter plot of roll and pitch angles for the puff and non-puff segments.

should be lower than other segments, we find a threshold t_d so that all puff segments are below it. We discard all segments that have distance greater t_d.

4.2.2 Performance of Hand-at-mouth Gesture Detection

As mentioned above, the proposed method to detect relatively stationary segments over-generates potential hand-to-mouth gestures. When we run the method on the 40.3 hours of data that we collected, on average it detects 20.51 and 20.37 segments/min from the left and right wrist sensor respectively. Screening based on duration and the threshold t_m reduces the number of segments of the left and rights wrists by 67.67% and 64.6% respectively. After screening using the orientation of the hands, the total of number of candidate puff segments is reduced to 4721 segments (1.95 segments/min) from both hands. This includes all the hand-at-mouth gesture segments corresponding to the 291 puff events.

4.3 Puff Detection

In this section, we describe the development of the puff classifier based on hand-at-mouth gestures and respiration measurements. First, we find the respiration cycles by computing the peak and valley locations in the respiration signal. Second, we select one respiration cycle for each candidate segment. The respiration signal reaches the peak once the smoke is completely inhaled and exhaling the smoke usually occurs once the cigarette is removed from the mouth. For each candidate hand-at-mouth gesture segments detected, we, therefore, select the first respiration cycle whose peak occurs after the end of the segment. A respiration cycle, however, can be associated with two different candidate segments, one from each hand, whose end times are close to each other. To avoid the situation where the training data contains a puff and a non-puff instance that are both associated with the same respiration cycle but different candidate segments, we only consider the non-smoking regions of the dataset as the source of non-puff instances. We extract features from both the candidate segments and the corresponding respiration cycle. We train a two class SVM classifier to detect puffs using the extracted features.

4.3.1 Features

From each respiration cycle we compute the 17 respiration features presented in 3.2. In addition to these respiration features, we propose to two more features. These features are computed from the rate of change signal obtained by taking the first derivative of the respiration signal. The maximum and minimum values that the rate of change signal attains within a respiration cycle is used as features.

From the candidate hand-at-mouth gesture segments we compute the mean, median, standard deviation, and quartile deviation of magnitude, pitch, and roll. This gives us a total of 12 features.

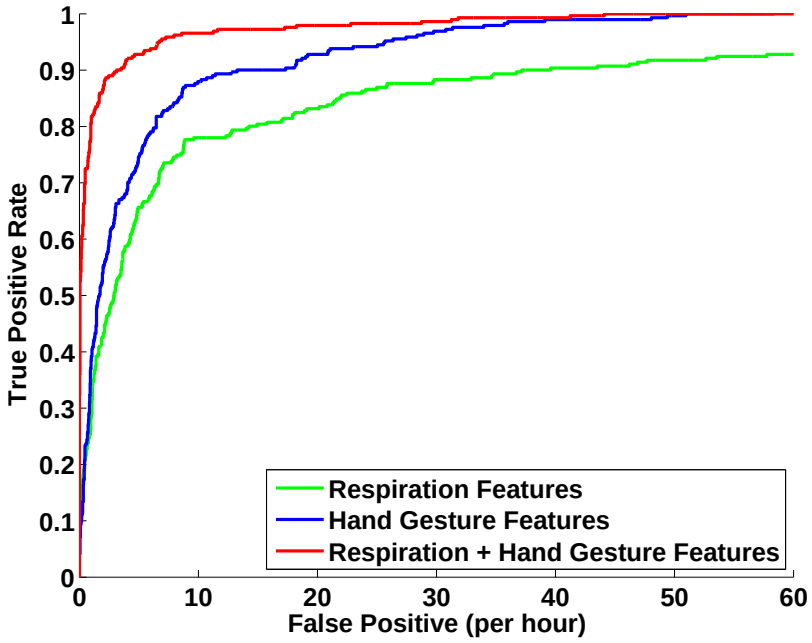

Fig. 4.4: True positive rate vs. false positive per hour of three classifiers based on respiration features, wrist sensors features, and features from both sensors.

4.3.2 Training and Evaluation

The training data for the puff classifier consists of a total of 291 puff instances and 3545 non-puff instances. We build three different classifiers based on only respiration features, only wrist sensor features, and combining the features from both sensors. Figure 4.4 presents the true positive rate of puffs detected versus number of false positive puffs per hour for the classifiers for 10 fold cross validation. As the non-puff instances were taken from regions outside smoking episodes the false positive per hour metric represents the false positives incurred only in those regions. We observe that by combining the features from the two sources we obtain much higher true positive rate without incurring too many false positives. At 90.3% true positive rate there is approximately 3.7 false positives per hour in the non-smoking regions of the data when features from both sensors are used. On the other hand, at 90.3% true positive rate the wrist and

respiration sensors based classifiers attains an average of 30.4 and 43.8 false positives per hour respectively.

Table 4.1: Leave one subject out cross validation results (TPR = True positive rate, FP = False positive).

Participant id	TPR	FP per hour
1	0.85	1.42
2	0.74	5.30
3	0.85	9.55

We further investigate the generalizability of the classifier based on the combined features by performing leave one subject out cross validation. Out of the six participants, we obtained 5 or more hours of data from three participants. We use data from these three participants in this experiment. A puff detection model is learned from 2 of the participant's data and then we evaluate it on the third. We first prepare the training data from the 2 participants by finding the candidate segments. This process also produces the different thresholds used for screening. An SVM classifier is trained using the training data and the threshold on the score of the classifier is set to a value so that is achieves 95% true positive rate on the training data. The true positive rate and false positives per hour metrics on the test data, represents the number of respiration cycles that are correctly or incorrectly classified. If a respiration cycle is associated of two test instances, we consider it to be detected as a puff if either instance is classified as a puff. Table 4.1 presents the cross validation results. We find the true positive rates all participant are high which indicate that our method can generalize to new user. The difference in false positive rate or false detection can be attributed to the difference in the proportion of confounding activities (e.g., conversation and physical activity) present in each participant's data.

4.4 Conclusion

In this chapter, we describe a method to combine hand gesture and respiration measurements for the detection of smoking puffs. We use the accelerometer and gyroscope measurements obtained from wrist worn sensors to first find the candidate segments when the hand is at mouth. The segments are then used to identify respiration cycles that can be potentially puff cycles. A SVM classifier is trained using data collected from 6 participants wearing sensors as they go on their daily lives. The 10-fold cross validation results indicate that combining these two modalities significantly reduces the number of false positives. At 90.3% true positive rate, respiration feature based classifier produces on average 43.8 false positives puff per hours which is reduced to 3.7 false positives per hour when both wrist and respiration features are used. We also perform leave one subject out cross validation which shows our method is able to produce generalizable models.

Chapter 5

Future Work & Conclusion

This dissertation presents, to the best of our knowledge, the first work that attempts to detect cigarette smoking puffs in the natural environment. We explored the feasibility of the use of respiration measurements for the detection of smoking puffs. We presented a model to automatically detect smoking puffs in the natural environment from respiration measurement. We achieve 86.7% accuracy on the detection of puffs even when there exists potential confounding events in the collected data. For the purpose of building the model, we identified 12 new respiration features that are found to be effective compared to the use of only the existing respiration features available in the literature. We also presented a semisupervised moded that improves the accuracy of the model when we provide unlabeled data with self reports that is also collected in the natural environment from participants who are daily smokers.

We also demonstrated that fusion of respiration and wrist worn inertial sensor measurements can further increase the accuracy of detection of puffs. We proposed a method to detect when the hand is at mouth using wrist worn inertial sensors (accelerometer and gyroscopes) and used it to temporally locate respiration cycles that are potentially puffs. Combining features from respiration cycles and segments of accelerometer and gyroscope timeseries corresponding to hand at mouth gesture significantly improves the accuracy of puff detection. Using 10-fold cross validation we observed that at 90.3% true positive rate, a classifier based on only respiration features produces 43.8 false detections per hour which is reduced to 3.7 false detections per hour when we combine features from respiration and wrist sensors measurements. We also performed leave-one-subject-out cross validation to analyze the generalizability of the puff detector. The results indicate that our method is able to generalize to new users.

This work presents the first steps towards the development of a smoking episode detector. As a future work, we plan to develop a smoking episode detector that utilizes the puff detector. We will explore different sequence labeling methods such as conditional random fields for this purpose. As the first step, we plan to conduct a field study to collect more data in the participant's natural environment. We envision several challenges that needs to be addressed in order to develop reliable smoking episode detector. First, as sensors are to be self-worn by the participants methods need to be developed to determine if the sensors are worn correctly. For example, the participants may take off the wrist sensors and when they put them back on they may flip the orientation of the sensors or switch the sensor intended to be worn on one wrist to the other. Second, obtaining fine grained ground truth in the field is also challenging. A smoking topography device can be useful to obtain puff level labels. Use of such a device may result in change in the natural smoking behavior and hand gestures associated with smoking and thereby limiting its utility. We may be able to obtain only coarse level ground truths, such as markings that indicate the beginning or end of an episode in the vicinity of the marking. Third, the episode detector needs to take into account the false detection miss-classification of puffs due to inherent inaccuracy of the puff detector, data loss, or noise. Information at the episode level such as inter-puff duration, number of puffs per episode etc. can be useful to detect episodes. However, being involved in other activities (e.g., drinking or talking) while smoking can cause wide variation in inter-puff durations and number of puffs per episode.

Another direction of future research on smoking detection is to explore the feasibility of using other sensing modalities. As smoking involves exhaling smoke, one can potentially use smart eyeglasses to detect smoke in the surrounding. However, limited battery life of such devices prohibits capturing video or images

continuously. Therefore, a triggering mechanism based on some other sensor needs to be developed that can trigger the smart-glasses to capture images or video.

REFERENCES

[1] U.S. Department of Health and Human Services. (2014) The health consequences of smoking - 50 years of progress: a report of the surgeon general. Atlanta: U.S. Department of Health and Human services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health. [Online]. Available: http://www.cdc.gov/tobacco/data_statistics/sgr/50th-anniversary/index.htm

[2] A. Mokdad, J. Marks, D. Stroup, and J. Gerberding, "Actual causes of death in the united states, 2000," *The Journal of the Americal Medical Association*, vol. 291, no. 10, pp. 1238–1245, 2004.

[3] Centers for Disease Control and Prevention, "Smoking-attributable mortality, years of potential life lost, and productivity losses - United States, 2000 - 2004," *Morbidity and Mortality Weekly Report*, vol. 57, no. 45, pp. 1226–1228, 2008.

[4] U.S. Department of Health and Human Services. (2014) Treating tobacco use and dependence: 2008 update - clinical practice guideline. [Online]. Available: http://www.ahrq.gov/professionals/clinicians-providers/guidelines-recommendations/tobacco/clinicians/update/treating_tobacco_use08.pdf

[5] S. Shiffman and A. Waters, "Negative affect and smoking lapses: A prospective analysis," *Journal of Consulting and Clinical Psychology*, vol. 72, no. 2, pp. 192–201, 2004.

[6] H. Brendryen, P. Kraft, and H. Schaalma, "Looking inside the black box: Using intervention mapping to describe the development of the automated smoking cessation intervention'happy ending'," *The Journal of Smoking Cessation*, vol. 5, no. 1, pp. 29–56, 2010.

[7] K. Matheny and K. Weatherman, "Predictors of smoking cessation and maintenance," *Journal of Clinical Psychology*, vol. 54, no. 2, pp. 223–235, 1998.

[8] N. Hymowitz, M. Sexton, J. Ockene, and G. Grandits, "Baseline factors associated with smoking cessation and relapse," *Preventive Medicine*, vol. 20, no. 5, pp. 590–601, 1991.

[9] K. Doherty, T. Kinnunen, F. Militello, and A. Garvey, "Urges to smoke during the first month of abstinence: relationship to relapse and predictors," *Psychopharmacology*, vol. 119, no. 2, pp. 171–178, 1995.

[10] J. Hughes and D. Hatsukami, "Signs and symptoms of tobacco withdrawal," *Archives of General Psychiatry*, vol. 43, no. 3, pp. 289–294, 1986.

[11] J. Killen and S. Fortmann, “Craving is associated with smoking relapse: Findings from three prospective studies,” *Experimental and Clinical Psychopharmacology*, vol. 5, no. 2, pp. 137–142, 1997.

[12] S. Shiffman, “Reflections on smoking relapse research,” *Drug and Alcohol Review*, vol. 25, no. 1, pp. 15–20, 2006.

[13] G. Swan, M. Ward, and L. Jack, “Abstinence effects as predictors of 28-day relapse in smokers,” *Addictive Behaviors*, vol. 21, no. 4, pp. 481–490, 1996.

[14] M. Stitzer and J. Gross, “Smoking relapse: the role of pharmacological and behavioral factors,” *Progress in Clinical and Biological Research*, vol. 261, pp. 163–184, 1988.

[15] S. Shiffman, J. Paty, M. Gnys, J. Kassel, and M. Hickcox, “First lapses to smoking: Within-subjects analysis of real-time reports,” *Journal of Consulting and Clinical Psychology*, vol. 64, no. 2, pp. 366–379, 1996.

[16] S. Shiffman, D. Scharf, W. Shadel, C. Gwaltney, Q. Dang, S. Paton, and D. Clark, “Analyzing milestones in smoking cessation: illustration in a nicotine patch trial in adult smokers,” *Journal of Consulting and Clinical Psychology*, vol. 74, no. 2, pp. 276–285, 2006.

[17] D. Kalman, “The subjective effects of nicotine: methodological issues, a review of experimental studies, and recommendations for future research,” *Nicotine & Tobacco Research*, vol. 4, no. 1, pp. 25–70, 2002.

[18] H. Ashton, D. Watson, R. Marsh, and J. Sadler, “Puffing frequency and nicotine intake in cigarette smokers,” *The British Medical Journal*, pp. 679–681, 1970.

[19] CReSS. [Online]. Available: http://www.borgwaldt.de/cms/borgwaldt-kc/produkte/rauchmaschinen/geraete-zur-rauchtopographie/cress-pocket.html

[20] Smokerlyzer. [Online]. Available: http://www.bedfont.com/smokerlyzer

[21] P. M. Scholl, N. Kücükyildiz, and K. V. Laerhoven, “When do you light a fire?: Capturing tobacco use with situated, wearable sensors,” in *Proc. ACM Int’l Conf. Ubiquitous Computing (Adjunct)*, 2013.

[22] F. Charles, G. Krautter, and D. Mariner, “Post-puff respiration measures on smokers of different tar yield cigarettes,” *Inhalation Toxicology*, vol. 21, no. 8, pp. 712–718, 2009.

[23] P. Wu, J. Hsieh, J. Cheng, S. Cheng, and S. Tseng, “Human smoking event detection using visual interaction clues,” in *Proc. IEEE Int’l Conf. Pattern Recognition*, 2010.

[24] P. Abhinav, M.-C. Chiu, C. Chadowitz, D. Ganesan, and E. Kalogerakis, "RisQ: Recognizing smoking gestures with inertial sensors on a wristband," in *Proc. ACM Int'l Conf. Mobile System, Applications, and Services (MobiSys)*, 2014.

[25] Q. Tang, D. J. Vidrine, E. Crowder, and S. S. Intille, "Automated detection of puffing and smoking with wrist accelerometers," in *Int'l Conf. on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2014.

[26] P. Lopez-Meyer, S. Tiffany, and E. Sazonov, "Identification of cigarette smoke inhalations from wearable sensor data using a support vector machine classifier," in *Proc. Int'l Conf. IEEE Engineering in Medicine and Biology Society*, 2012.

[27] K. Plarre, A. Raij, S. M. Hossain, A. A. Ali, M. Nakajima, M. Al'absi, E. Ertin, T. Kamarck, S. Kumar, and M. Scott, "Continuous inference of psychological stress from sensory measurements collected in the natural environment," in *Proc. ACM Int'l Conf. Information Processing in Sensor Networks*, 2011.

[28] A. Ali, S. Hossain, K. Hovsepian, M. Rahman, K. Plarre, and S. Kumar, "mpuff: automated detection of cigarette smoking puffs from respiration measurements," in *Proc. ACM Int'l Conf. Information Processing in Sensor Networks*, 2012.

[29] E. L. Wynder and E. A. Graham, "Tobacco smoking as a possible etiologic factor in bronchiogenic carcinoma: a study of six hundred and eighty-four proved cases," *Journal of the American Medical Association*, vol. 143, no. 4, pp. 329–336, 1950.

[30] R. Doll and A. B. Hill, "Smoking and carcinoma of the lung. preliminary report," *British Medical Journal*, vol. 2, pp. 739–748, 1950.

[31] E. L. Wynder, "Tobacco and health: a review of the history and suggestions for public health policy," *Public Health Reports*, vol. 103, no. 1, p. 8, 1988.

[32] R. Doll, "Uncovering the effects of smoking: historical perspective," *Statistical Methods in Medical Research*, vol. 7, no. 2, pp. 87–117, 1998.

[33] P. Vineis, M. Alavanja, P. Buffler, E. Fontham, S. Franceschi, Y.-T. Gao, P. C. Gupta, A. Hackshaw, E. Matos, J. Samet, *et al.*, "Tobacco and cancer: recent epidemiological evidence," *Journal of the National Cancer Institute*, vol. 96, no. 2, pp. 99–106, 2004.

[34] M. Bartal, "Health effects of tobacco use and exposure," *Monaldi Archives for Chest Disease*, vol. 56, no. 6, pp. 545–554, 2001.

[35] U.S. Department of Health and Human Services. (2006) The health consequences of involuntary exposure to tobacco smoke: a report of the

surgeon general. Atlanta: U.S. Department of Health and Human services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health. [Online]. Available: http://www.cdc.gov/tobacco/data_statistics/sgr/2006/index.htm

[36] M. P. Eriksen and R. L. Cerak, "The diffusion and impact of clean indoor air laws," *Annual Review of Public Health*, vol. 29, pp. 171–185, 2008.

[37] J. E. Harris, "Increasing the federal excise tax on cigarettes," *Journal of Health Economics*, vol. 1, no. 2, pp. 117–120, 1982.

[38] M. Grossman, "Health benefits of increases in alcohol and cigarette taxes," *British Journal of Addiction*, vol. 84, no. 10, pp. 1193–1204, 1989.

[39] Centers for Disease Control and Prevention, "Quitting smoking among adults - United States, 2001 - 2010," *Morbidity and Mortality Weekly Report*, vol. 60, no. 44, pp. 1513–1519, 2010.

[40] M. Russell, "Tobacco dependence: is nicotine rewarding or aversive?" *NIDA Research Monograph*, no. 23, pp. 100–122, 1979.

[41] M. J. Jarvis, "Why people smoke," *British Medical Journal*, vol. 328, no. 7434, pp. 277–279, 2004.

[42] I. Mitrouska, I. Bouloukaki, and N. Siafakas, "Pharmacological approaches to smoking cessation," *Pulmonary Pharmacology & Therapeutics*, vol. 20, no. 3, pp. 220–232, 2007.

[43] S. M. Shiffman, "The tobacco withdrawal syndrome," *NIDA Research Monograph*, vol. 23, pp. 158–184, 1979.

[44] A. J. Alberg, J. L. Patnaik, J. W. May, S. C. Hoffman, J. Gitchell, J. Gitchell, and K. J. Helzlsouer, "Nicotine replacement therapy use among a cohort of smokers," *Journal of Addictive Diseases*, vol. 24, no. 1, pp. 101–113, 2005.

[45] T. P. George and S. S. O'Malley, "Current pharmacological treatments for nicotine dependence," *Trends in Pharmacological Sciences*, vol. 25, no. 1, pp. 42–48, 2004.

[46] M. C. Fiore, S. S. Smith, D. E. Jorenby, and T. B. Baker, "The effectiveness of the nicotine patch for smoking cessation: a meta-analysis," *JAMA*, vol. 271, no. 24, pp. 1940–1947, 1994.

[47] S. P. Marlow and J. K. Stoller, "Smoking cessation," *Respiratory Care*, vol. 48, no. 12, pp. 1238–1256, 2003.

[48] A. Hjalmarson, M. Franzon, A. Westin, and O. Wiklund, “Effect of nicotine nasal spray on smoking cessation: a randomized, placebo-controlled, double-blind study,” *Archives of Internal Medicine*, vol. 154, no. 22, pp. 2567–2572, 1994.

[49] P. Tonnesen, J. Norregaard, K. Simonsen, and U. Säwe, “A double-blind trial of a 16-hour transdermal nicotine patch in smoking cessation,” *New England Journal of Medicine*, vol. 325, no. 5, pp. 311–315, 1991.

[50] G. Sutherland, J. Stapleton, M. Russell, M. Jarvis, P. Hajek, M. Belcher, and C. Feyerabend, “Randomised controlled trial of nasal nicotine spray in smoking cessation,” *The Lancet*, vol. 340, no. 8815, pp. 324–329, 1992.

[51] N. G. Schneider, R. Olmstead, F. Nilsson, F. V. Mody, M. Franzon, and K. Doan, “Efficacy of a nicotine inhaler in smoking cessation: a double-blind, placebo-controlled trial,” *Addiction*, vol. 91, no. 9, pp. 1293–1306, 1996.

[52] P. Wu, K. Wilson, P. Dimoulas, and E. J. Mills, “Effectiveness of smoking cessation therapies: a systematic review and meta-analysis,” *BMC Public Health*, vol. 6, no. 1, pp. 1–16, 2006.

[53] G. A. Lillington, C. T. Leonard, and D. P. Sachs, “Smoking cessation: Techniques and benefits,” *Clinics in Chest Medicine*, vol. 21, no. 1, pp. 199–208, 2000.

[54] C. Silagy, T. Lancaster, L. Stead, D. Mant, and G. Fowler, “Nicotine replacement therapy for smoking cessation,” *The Cochrane Database of Systematic Reviews*, no. 3, p. CD000146, 2004.

[55] L. E. Odum, K. A. O’Dell, and J. S. Schepers, “Electronic cigarettes do they have a role in smoking cessation?” *Journal of Pharmacy Practice*, vol. 25, no. 6, pp. 611–614, 2012.

[56] M. B. Siegel, K. L. Tanwar, and K. S. Wood, “Electronic cigarettes as a smoking-cessation tool: results from an online survey,” *American Journal of Preventive Medicine*, vol. 40, no. 4, pp. 472–475, 2011.

[57] T. Lancaster and L. F. Stead, “Individual behavioural counselling for smoking cessation,” *The Cochrane Database of Systematic Reviews*, vol. 2, 2005.

[58] L. F. Stead and T. Lancaster, “Group behaviour therapy programmes for smoking cessation,” *The Cochrane Database of Systematic Reviews*, vol. 2, 2005.

[59] K. A. Perkins, M. D. Marcus, M. D. Levine, D. D’Amico, A. Miller, M. Broge, J. Ashcom, and S. Shiffman, “Cognitive–behavioral therapy to reduce

weight concerns improves smoking cessation outcome in weight-concerned women," *Journal of Consulting and Clinical Psychology*, vol. 69, no. 4, pp. 604–613, 2001.

[60] B. Spring, D. Howe, M. Berendsen, H. McFadden, K. Hitchcock, A. W. Rademaker, and B. Hitsman, "Behavioral intervention to promote smoking cessation and prevent weight gain: a systematic review and meta-analysis," *Addiction*, vol. 104, no. 9, pp. 1472–1486, 2009.

[61] L. F. Stead, R. Perera, and T. Lancaster, "Telephone counselling for smoking cessation," *The Cochrane Database of Systematic Reviews*, vol. 3, 2006.

[62] S. Mottillo, K. B. Filion, P. Bélisle, L. Joseph, A. Gervais, J. O'Loughlin, G. Paradis, R. Pihl, L. Pilote, and S. Rinfret, "Behavioural interventions for smoking cessation: a meta- analysis of randomized controlled trials," *European Heart Journal*, vol. 30, no. 6, pp. 718–730, 2009.

[63] S. Shiffman, S. E. Brockwell, J. L. Pillitteri, and J. G. Gitchell, "Use of smoking-cessation treatments in the united states," *American Journal of Preventive Medicine*, vol. 34, no. 2, pp. 102–111, 2008.

[64] S. T. Walters, J. A. Wright, and R. Shegog, "A review of computer and internet-based interventions for smoking behavior," *Addictive Behaviors*, vol. 31, no. 2, pp. 264–277, 2006.

[65] T. Lancaster and L. F. Stead, "Self-help interventions for smoking cessation," *The Cochrane Database of Systematic Reviews*, vol. 3, no. 3, 2005.

[66] J.-F. Etter, "Comparing the efficacy of two internet-based, computer-tailored smoking cessation programs: a randomized trial," *Journal of Medical Internet Research*, vol. 7, no. 1, 2005.

[67] A. Rodgers, T. Corbett, D. Bramley, T. Riddell, M. Wills, R.-B. Lin, and M. Jones, "Do u smoke after txt? results of a randomised trial of smoking cessation using mobile phone text messaging," *Tobacco Control*, vol. 14, no. 4, pp. 255–261, 2005.

[68] C. Free, R. Knight, S. Robertson, R. Whittaker, P. Edwards, W. Zhou, A. Rodgers, J. Cairns, M. G. Kenward, and I. Roberts, "Smoking cessation support delivered via mobile phone text messaging (txt2stop): a single-blind, randomised trial," *The Lancet*, vol. 378, no. 9785, pp. 49–55, 2011.

[69] S. Shiffman, C. J. Gwaltney, M. H. Balabanis, K. S. Liu, J. A. Paty, J. D. Kassel, M. Hickcox, and M. Gnys, "Immediate antecedents of cigarette smoking: an analysis from ecological momentary assessment," *Journal of Abnormal Psychology*, vol. 111, no. 4, pp. 531–545, 2002.

[70] D. G. Gilbert, J. P. Sharpe, N. V. Ramanaiah, F. R. Detwiler, and A. E. Anderson, "Development of a situation$\times$ trait adaptive response (star) model-based smoking motivation questionnaire," *Personality and Individual Differences*, vol. 29, no. 1, pp. 65–84, 2000.

[71] S. Shiffman, M. Balabanis, J. Fertig, and J. Allen, "Associations between alcohol and tobacco," *Alcohol and Tobacco: from basic science to clinical practice. NIAAA Research Monograph*, vol. 30, pp. 17–36, 1995.

[72] A. C. McKennell, "Smoking motivation factors," *British Journal of Social and Clinical Psychology*, vol. 9, no. 1, pp. 8–22, 1970.

[73] J. D. Lane, "Association of coffee drinking with cigarette smoking in the natural environment," *Experimental and Clinical Psychopharmacology*, vol. 4, no. 4, pp. 409–412, 1996.

[74] T. J. Payne, M. L. Schare, D. J. Levis, and G. Colletti, "Exposure to smoking-relevant cues: Effects on desire to smoke and topographical components of smoking behavior," *Addictive Behaviors*, vol. 16, no. 6, pp. 467–479, 1991.

[75] R. Niaura, D. B. Abrams, M. Pedraza, P. M. Monti, and R. Damaris J, "Smokers' reactions to interpersonal interaction and presentation of smoking cues," *Addictive Behaviors*, vol. 17, no. 6, pp. 557–566, 1992.

[76] S. Shiffman, "Assessing smoking patterns and motives," *Journal of Consulting and Clinical Psychology*, vol. 61, no. 5, pp. 732–742, 1993.

[77] J. C. Tate, J. M. Schmitz, and A. L. Stanton, "A critical review of the reasons for smoking scale," *Journal of Substance Abuse*, vol. 3, no. 4, pp. 441–455, 1991.

[78] S. Shiffman, M. Hufford, M. Hickcox, J. A. Paty, M. Gnys, and J. D. Kassel, "Remember that? a comparison of real-time versus retrospective recall of smoking lapses," *Journal of Consulting and Clinical Psychology*, vol. 65, no. 2, pp. 292–300, 1997.

[79] S. Shiffman, A. A. Stone, and M. R. Hufford, "Ecological momentary assessment," *Annual Review of Clinical Psychology*, vol. 4, pp. 1–32, 2008.

[80] S. Kolonen, J. Tuomisto, P. Puusteinen, and M. Airaksinen, "Puffing behavior during the smoking of a single cigarette in a naturalistic environment," *Pharmacology, biochemistry and behavior*, vol. 41, no. 4, pp. 701–706, 1992.

[81] D. Hammond, G. Fong, K. Cummings, and A. Hyland, "Smoking topography, brand switching, and nicotine delivery: results from an in vivo study," *Cancer Epidemiology Biomarkers & Prevention*, vol. 14, no. 6, pp. 1370–1375, 2005.

[82] J. Veilleux, J. Kassel, A. Heinz, A. Braun, M. Wardle, J. Greenstein, D. Evatt, and M. Conrad, "Predictors and sequelae of smoking topography over the course of a single cigarette in adolescent light smokers," *Journal of Adolescent Health*, vol. 48, no. 2, pp. 176–181, 2011.

[83] O. D. Lara and M. A. Labrador, "A survey on human activity recognition using wearable sensors," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1192–1209, 2013.

[84] A. Avci, S. Bosch, M. Marin-Perianu, R. Marin-Perianu, and P. Havinga, "Activity recognition using inertial sensing for healthcare, wellbeing and sports applications: A survey," in *Proc. Int'l conf. Architecture of Computing Systems (ARCS)*, 2010.

[85] T. Park, J. Lee, I. Hwang, C. Yoo, L. Nachman, and J. Song, "E-gesture: a collaborative architecture for energy-efficient gesture recognition with hand-worn sensor and mobile devices," in *Proc. ACM Conf. Embedded Networked Sensor Systems (SenSys)*, 2011.

[86] G. Raffa, J. Lee, L. Nachman, and J. Song, "Don't slow me down: Bringing energy efficiency to continuous gesture recognition," in *Proc. IEEE Int'l Symp. Wearable Computers (ISWC)*, 2010.

[87] A. Akl, C. Feng, and S. Valaee, "A novel accelerometer-based gesture recognition system," *IEEE Trans. Signal Processing*, vol. 59, no. 12, pp. 6197–6205, 2011.

[88] J. Wu, G. Pan, D. Zhang, G. Qi, and S. Li, "Gesture recognition with a 3-d accelerometer," in *Proc. Int'l Conf. Ubiquitous Intelligence and Computing*, 2009.

[89] H. Junker, O. Amft, P. Lukowicz, and G. Tröster, "Gesture spotting with body-worn inertial sensors to detect user activities," *Pattern Recognition*, vol. 41, no. 6, pp. 2010–2024, 2008.

[90] J. Liu, L. Zhong, J. Wickramasuriya, and V. Vasudevan, "uWave: Accelerometer-based personalized gesture recognition and its applications," *Pervasive and Mobile Computing*, vol. 5, no. 6, pp. 657–675, 2009.

[91] Y. Li, "Protractor: a fast and accurate gesture recognizer," in *Proc. ACM Conf. Human Factors in Computing Systems (CHI)*, 2010.

[92] M. M. Rahman, A. A. Ali, K. Plarre, M. al'Absi, E. Ertin, and S. Kumar, "mConverse: Inferring conversation episodes from respiratory measurements collected in the field," in *Proc. ACM Conf. Wireless Health*, 2011.

[93] E. Ertin, N. Stohs, S. Kumar, A. Raij, M. al'Absi, and S. Shah, "Autosense: Unobtrusively wearable sensor suite for inferencing of onset, causality, and consequences of stress in the field," in *Proc. ACM Conf. Embedded Networked Sensor Systems (SenSys)*, 2011.

[94] Fieldstream. [Online]. Available: http://www.github.com/FieldStream/FieldStream

[95] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[96] X. Zhu and A. Goldberg, *Introduction to Semi-Supervised Learning*. Morgan & Claypool Publishers, 2009.

[97] Svmlin. [Online]. Available: http://vikas.sindhwani.org/svmlin.html

[98] J. Murphy, *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. New York Institute of Finance, 1999.

[99] C. Marian, R. J. O'Connor, M. V. Djordjevic, V. W. Rees, D. K. Hatsukami, and P. G. Shields, "Reconciling human smoking behavior and machine smoking patterns: implications for understanding smoking behavior and the impact on laboratory studies," *Cancer Epidemiology Biomarkers & Prevention*, vol. 18, no. 12, pp. 3305–3320, 2009.

[100] H. J. Luinge and P. H. Veltink, "Measuring orientation of human body segments using miniature gyroscopes and accelerometers," *Medical and Biological Engineering and Computing*, vol. 43, no. 2, pp. 273–282, 2005.

[101] M. Pedley. (2014) Tilt sensing using a three-axis accelerometer. [Online]. Available: http://www.freescale.com/files/sensors/doc/app_note/AN3461.pdf