4-17-2017

Identifying Drug (Cocaine) Intake Events from Acute Physiological Response in the Presence of Free-living Physical Activity

Syed Monowar Hossain

IDENTIFYING DRUG (COCAINE) INTAKE EVENTS FROM ACUTE
PHYSIOLOGICAL RESPONSE IN THE PRESENCE OF FREE-LIVING
PHYSICAL ACTIVITY

by

Syed Monowar Hossain

A Dissertation

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

Major: Computer Science

The University of Memphis

May 2017

ABSTRACT

Hossain, Syed Monowar. Ph.D. The University of Memphis. May, 2017. Identifying Drug (Cocaine) Intake Events from Acute Physiological Response in the Presence of Free-living Physical Activity. Major Professor: Dr. Santosh Kumar.

A variety of health and behavioral states can potentially be inferred from physiological measurements that can now be collected in the free-living environment. The major challenge, however, is to develop computational models for automated detection of health events that can work reliably. In this work, we develop a physiologically-informed model to automatically detect drug (cocaine) use events in the free-living environment of participants from their electrocardiogram (ECG) measurements. The key to reliably detect these events is to incorporate the knowledge of autonomic nervous system (ANS) behavior in the model development so as to decompose the activation effect of cocaine from the natural recovery behavior of the parasympathetic nervous system (after an episode of physical activity). We collected 89 days of data from 9 active drug users in two residential lab environments and 922 days of data from 42 active drug users in the field environment, for a total of 11,283 hours. We developed a model that tracks the natural recovery by the parasympathetic nervous system and then estimates the dampening caused to the recovery by the activation of the sympathetic nervous system due to cocaine. We developed efficient methods to screen and clean the ECG time series data and extract candidate windows to assess for potential drug use. We then applied our model on the recovery segments from these windows. Our model achieves 100% true positive rate while keeping the false positive rate to 0.87/day over (9+ hours/day of) lab data and to 1.13/day over (11+ hours/day of)field data. In order to further improve the sensitivity and specificity of our model, we proposed several new data screening methods. Also we proposed a method to remove the effects of activities that acts as a confounder. We observed the false positive rates of 0.78 and 0.98 per day

when we apply the enhanced model to the lab and field data respectively. Moreover, we observed that the proposed model has high specificity to cocaine. The method also estimates the dosage amount of drug for an event. However, the predicted dosage amount is not reliable for high dosage amounts in free living conditions.

TABLE OF CONTENTS

Content   **Page**

LIST OF TABLES

LIST OF FIGURES

Figure   **Page**

Chapter 1

Introduction

1.1 Background and Motivation

Advances in mobile sensing are enabling a new vision of healthcare (called mobile health or *mHealth*), where consumers can monitor, manage, and improve health and well-being as they go about their daily lives [1]. Wearable, inexpensive sensors allow capture of health relevant data, such as measurements of heart, respiration, physical activity, location, etc. in the natural environment [2,3]. These wearable sensors have become reliable instruments to collect sensor measurements in the field environment, and hence research has shifted from developing wireless sensor platforms to the processing of sensor data to infer health related events such as stress [4], smoking [5], drug use [6], and identify their antecedents and precipitants (i.e., high risk situations). Automated detection of these precipitants on a mobile phone can be used to trigger just-in-time intervention or treatment. For example, momentary exposure to greater physical disorder, social disorder, and drug activity in a neighborhood (as indicated by the NIfETy score [7]) and experiencing craving or stress could constitute a high-risk situation worthy of real-time intervention for a drug user wanting to quit. The first step in building such mHealth control systems for improving human health usually is to collect sensor data in the natural environment and locate the adverse health events in the time-series of sensor data so as to discover the causal role of various contexts in precipitating the adverse health event. Therefore, the development of reliable models for detecting adverse health events from mHealth sensor data in the natural field environment is a critical step in the development of just-in-time mHealth interventions.

Development of such models involve several challenges. First, appropriate sensor(s) are needed that can be used conveniently in the field settings for long

1

enough duration to capture the health events of interest. The sensor should be robust for reliable data collection in the free-living natural environment. In addition, the sensor should be sensitive and specific enough to exhibit a detectable response to the health event of interest. Second, appropriate data collection experiment needs to be designed and conducted to collect the sensor data with appropriate labeling of the times when the health event of interest occurs (to train the model). Third, a robust computational model needs to be developed that is able to detect the event of interest from sensor data in the free-living natural environment.

Drug intake events (e.g., cocaine use), known to acutely excite the autonomic nervous system, can potentially be detected from its response on wearable electrocardiogram (ECG) sensors. Heart Rate (HR) increase of more than 30% (for 16 mg of cocaine) [8, 9] have been observed in lab studies. After (intravenous) cocaine administration, it usually takes 20-60 minutes for the heart rate to recover. Hence, the effect of cocaine use on ECG is quite acute and potentially distinct.

To the best of our knowledge, [6] is the first work to present a computational model for automated detection of cocaine use from ECG. They collect data from six participants in lab setting during cocaine administrations of 8 mg, 16 mg and 32 mg dosage levels. The authors develop a model to classify ECG cycles into drug and baseline classes. Although this method achieves high accuracy (average AUC > 0.9) on clean lab data, there are several challenges in using such a model for detecting cocaine use events in the field setting. For example, as we show in Section 3.3.1, it is not trained to distinguish ECG cycles during drug use from that during physical activity. Second, it is not trained to be invariant to dosage amount and to the modality of administration. Third, since the model classifies each ECG cycle, it is not clear how to use this as a building block

2

to develop a model for detecting the entire drug use event and to distinguish it from physical activity events. Although accelerometry measurements can be used to detect the occurrence of physical activity, drug use events in the field are usually concurrent with physical activity (i.e., subjects are not stationary following cocaine use), and hence acceleromtery alone can't distinguish the two events.

1.2 Problem Statement

In this work, our goal is to design a physiologically-informed model to automatically detect drug events from their acute physiological response, in the presence of various confounders inherent in the free-living lifestyle. The key to reliably detecting drug use event is to incorporate the knowledge of autonomic nervous system (ANS) behavior in the model development. We decompose the activation effect of cocaine from the natural recovery behavior of the parasympathetic nervous system (PNS) that can be observed upon conclusion of a physical activity episode. There is increasing interest in the problem of automatically detecting drug use due to its high societal impact. Illicit drug use, affecting over 150 million people, is a major cause of mortality from fatal overdose and dependence, HIV, Hepatitis B, and Hepatitis C. Other adverse health effects include mental disorder, traffic accidents, suicides, and violence [10]. In drug intake research one of the important goals is to identify the antecedents and precipitants of cocaine intake. Scientific user studies resort to observing and recording user?s context while drug intake occurs. Thus it is imperative that these studies employ some method for detecting drug intake. However, most of the studies on cocaine intake behavior mainly employ self-reporting methods for this purpose. These self-reporting methods range from basic pen-paper methods and retrospective recalls, to electronic diary keeping and ecological momentary assessments(EMA) [11]. These methods, in addition to putting extra burden on users, depend on user compliance, introduce subjective biases etc. that limit their

utility. Therefore, our focus is on the automated detection of drug events in the natural environment that can enable the development of such effective intervention methods.

1.3   Cocaine Event Detection Model

We first designed and conducted three user studies with active drug users — two in the lab and one in the field. In each study, participants wore the AutoSense sensor suite [2] that included an ECG sensor and accelerometers. The lab studies were conducted in residential facilities with 9 drug users (across 89 days). It included free-living lifestyle together with sessions of repeated cocaine administration of various doses under medical supervision. For the field study, 42 drug users wore the sensors for 4 weeks in the field so as to maximize the chances of capturing real-life cocaine use events.

We then develop efficient methods to screen and clean sensor data to handle noise and drift. Next, we develop a data preprocessing stage to identify and locate windows in ECG time series that exhibit physiological response of sufficient magnitude that may result from cocaine use. Activity-free recovery segments from these windows are assessed to determine whether this window is a result of cocaine use. For this purpose, we develop a dynamical system model of the parasympathetic nervous system (PNS) behavior from the heart rate recovery observed upon conclusion of cocaine-free physical activity episodes. In the case of cocaine use, the PNS recovery is dampened due to excitation of the sympathetic nervous system (SNS) from cocaine. The strength of SNS excitation weakens with metabolism of cocaine. Using lab data from cocaine administrations, we develop dynamical system models of both SNS activation and its weakening due to cocaine metabolism, in order to model the cocaine-dampened PNS recovery. We refer to these models collectively as our Autonomous Nervous System (ANS) model. Our ANS model classifies a window

4

into cocaine class if the recovery portion of the window matches that of a cocaine-dampened PNS recovery, and otherwise, if it better matches natural PNS recovery.

We evaluate the performance of our ANS model on both lab and field data. We present some operating points from the ROC curve (see Figure 3.9). For 100% true positive rate, the model keeps the false positive rate to 0.87/day in the lab, even though 10+ candidate windows (due to significant physical activity) are found during 9+ hours/day of sensor wearing in the lab. On 922 days of the field data, we find 27 episodes of cocaine use with good quality sensor data from 13 participants. For these participants, there are 79 confirmed non-cocaine days as established from urine assessment. On these days, there are 1,171 major activity episodes. The model keeps the false alarm rate to 1.13/day, even though 11+ candidate windows are found during 11+ hours/day of sensor wearing in the field.

1.4 Enhancement of the Model

Based on the findings of our initial work, we enhance several modification to the model. First, we improve the sensitivity and specificity of the model. Secondly, we add detection of cocaine dosage as an added functionality to the model.

In order to improve the sensitivity of the model to such noise and increase the number of potentially detectable events, fist we investigate the performance of the model when smaller recovery segment is used. We show that, when the length of the recovery segment is small (30 second), the false positive rate increases up to 10.4 per day. Second, we analyze the performance of the model when we use full recovery portion of the RR signal in the model after removing the effects of confounders, namely activity. To remove the effect of activity, we remove the portion of the RR signal that is affected by intermittent activities of short durations and impute the missing data. This allows us to use the full

5

recovery part of the RR signal in the model. We show that the false positives/day is reduced from 0.87 to 0.74 on the lab data and in case of field data we lower the false positive rate to below 1 per day.

On the other hand, to investigate the specificity of the model, we run the model on the data collected on days where there is self report of only non-cocaine drugs and we find that none of the drugs (e.g., Heroin, Morphine, THC, etc.) are misclassified as cocaine.

As an added functionality of the model, we estimate the dosage amount for each detected drug intake events. First, we analyze the performance of the model to estimate the dosage on lab data. We find that the model estimates the dosage with 27.9% error. Then, we apply the model to estimate the dosage in field data. We see that the error rate is high (101.7%). Our conclusion is, due to the effect of activity and higher intake of dosage than the lab, it is difficult to estimate the dosage. More data with higher dosage in the controlled environment is important to investigate further.

Contributions. This work makes several key contributions. First, it establishes feasibility of collecting good quality physiological sensor data in the field environment from illicit drug users during real-life drug use events. These data provide a first look at physiological response to drug use episodes of up to 600 mg. Second, this work shows that it is indeed possible to develop a model for detection of health events from physiological sensor data collected in the lab settings that generalize to the free-living natural environment, providing a feasibility result for the critical question of lab-to-field generalizability. This is achieved by developing explainable models that model physiological mechanisms and hence enable decomposing the effect of physical activity on the physiological response. Third, automated detection of drug use events in real-time (in a scientific study) can be used to solicit self-reports on the mobile phone to obtain

more information surrounding the context of drug use. Fourth, accurate localization of drug use events in the time series of sensor measurements can enable discovery of antecedents and precipitants that can be detected from sensors in the mobile phone alone (e.g., GPS, accelerometry). Such discoveries can then be used to construct novel just-in-time interventions to predict drug use events and break the urge. Even self-monitoring of these contexts can enhance a patients' awareness of vulnerable contexts and help them avoid such high risk situations [12, 13, 14].

Finally, the model of parasympathetic nervous system (PNS) can be of independent interest in assessing the fitness of individuals. Recovery from physical activity (i.e., the health of PNS) is traditionally used for estimating cardiovascular health in clinical settings [15, 16]. The ANS model developed in this work can be used to obtain stable estimates of PNS recovery in the natural environment. This new measure has a potential to be of similar utility as heart rate variability (HRV) in biofeedback applications for self-improvement of physiological health.

Real-life Usage of the Model. The work presented here is not intended to be used directly by drug users seeking help with cocaine abstinence. Rather, this model is intended to be used in scientific studies of drug use. Using the model, we are able to pinpoint where the cocaine event occurred in time, which cannot be reliably inferred from self-report. Locating cocaine use events in the time series can help identify predictors from other sensor modalities (e.g., location, activity, stress, etc.). During real-life usage to help cocaine abstinence, the system does not detect cocaine use; rather, it predicts potential for lapse and prompts a user to break his/her urge when the predictors are detected. In this phase, the users are not wearing the physiological sensors. They only carry their smartphones.

We further point out that collecting physiological sensor data during real-life

cocaine use episodes is an inherently challenging task due to the context and situation associated with illicit drug use. It took nearly 922 person days (10,449 hours) of sensor wearing in the field to find 27 cocaine use episodes with good quality ECG data. Fortunately, development of a model is a one-time activity. Once a model is published and predictors of cocaine use are discovered and reported (subject of ongoing work), the community does not need to go through this tedious and resource intensive process again. They can use our models and use the predictors we find in our data set in developing cocaine abstinence interventions.

1.5   Organization

Chapter 2 describes some related works and summarizes the key challenges in developing a reliable model for detecting cocaine use in the field. Data collection procedure and statistics of data collected is described in Chapter 3. Data processing steps and model development are also described in this chapter. We describe how we can enhance of our model for better sensitivity and specificity in Chapter 4. In Chapter 5, we describe the limitations of our model. Finally in Chapter 6, we discuss potential generalization of our model in the sensor based detection of other human behaviors.

Chapter 2

Related Works

There exists a substantial body of literature on the measurement of physiological responses to various drugs, such as cocaine, in the laboratory environment, where different doses of cocaine were administered to volunteers and physiological measurements (e.g., heart rate, blood pressure) recorded [8, 9, 17, 18, 19, 20]. These studies focussed on examining the influence of different routes of administration (e.g. smoked, intervenous and intranasal) on pharmacokinetic parameters and drug-induced behavioral and physiological effects of cocaine. Table 2.1 summarizes results from these works on the effect of cocaine on heart rate. We observe that the response time and duration of effect depends highly on the route of administration [21]. In case of smoked or intravenous administration, the onset of action is almost immediate and within 1-5 minutes the effect reaches its peak. In all cases, there is an increase in heart rate (or decrease in RR interval - the interval from one R-peak to the next in the ECG waveform) and blood pressure and both increase with dosage. An increase of 32% and 34% respectively for intervenous cocaine doses of 16 mg and 32 mg was reported in [8, 9, 22].

Table 2.1: Response time and effect duration of cocaine intake on heart rate.

Route of Administration	Onset of Action	Peak Effect	Duration of Action
Smoking	3-5 sec	1-3 mins	5-15 mins
Intravenous	10-60 sec	3-5 min	20-60 min
Intranasal or mucusal	1-5 min	15-20 min	60-90 min
Gastrointestinal	$\leq$ 20 min	$\leq$ 90 min	$\leq$ 180 min

For characterizing physiological response to drug use, [23] used non-linear regression to model the heart rate during cocaine and placebo administration sessions. The placebo session consisted of a bolus intervenous injection and a

9

4-hour continuous infusion of a 0.2N saline solution. It used 40 mg, 60 mg and 80 mg cocaine intervenous doses followed by a 0.2N saline or cocaine infusion that lasted for 4 hours. Heart rate was modeled as the sum of baseline heart rate, an exponentially decaying model of conditioned response effect, and a non-linear drug effect model. The conditioned response effect was used to model the response that was observed a couple of minutes before the injections which also continued a couple of minutes after the injection, even during the placebo sessions. This model, however, cannot be readily used to detect drug use in the field as it does not distinguish from similar effects that are observed during other confounding events in daily life (e.g., physical activity).

In [13,14], the authors describe preliminary investigation of a proposed project called "iHeal", that can potentially detect subject's craving of drugs in their natural environment from sensor data on electrodermal activity, body motion, skin temperature, and, optionally, heart rate. Results of this project, however, are yet to be reported.

To the best of our knowledge, [6] presents the first work on an ECG morphological feature based classifier for detecting cocaine use. As discussed in Section 1.3, the model presented in [6] for classifying each ECG cycle into cocaine use or baseline does not lead to a model for detecting Cocaine use events in the field setting that involves free-living activities.

Key Challenges. In conclusion, detection of cocaine use from physiological measurements collected in the field setting is challenging. We now summarize some of these technical and experimental challenges. First, physiological measurements such as ECG are subject to several sources of noises and quality issues, especially when worn in the mobile environment. They include incorrect placement and poor attachment of electrodes. Second, reliably collecting wearable ECG measurements from illicit drug users during active drug

10

use in their natural environment is challenging due to the scenario and context of its usage. From our multi-year effort to collect ECG data in this population, we observe that in many situations the participant simply chooses not to wear the sensors during drug intake, even though they wear the sensors daily for 11+ hours. This may be due to safety concerns associated with wearing devices in crime-prone neighborhoods, where they may be suspected to be wired by the police. Third, it is very difficult to recruit participants who are active drug users; therefore, in most published scientific studies, the number of participants is usually in single digits. Fourth, obtaining ground truth, i.e., self-reports from participants as they take drug is also difficult, partly due to the same reasons as above. Though urine assessment can indicate drug intake over the previous few days, they do not provide the exact timing of the drug use event. Fifth, dosage amount and the method of administration (in case of cocaine — smoking, intravenous, intranasal or gastrointestinal) has significant effect on the HR response. However, we cannot obtain training data in a lab setting by administering drugs that represent very high dosage levels observed in the field (up to 600 mg). Sixth, as data is collected in unconstrained environments, there are usually confounding factors that can have similar physiological response. For cocaine use detection, the common confounding factors are activity, caffeine intake, and the intake of other drugs. Since we can't ask participants to walk during drug administration in lab, we can't collect lab data that represents drug use mixed with physical activity. Yet, the model developed must be able to distinguish drug use events from physical activity even when they co-occur. The model we present is the first work to handle all of the above issues and detect drug use events that occur in the field setting.

Chapter 3

Cocaine Event Detection Model

In this chapter, we present detection of cocaine intake, a novel system to automatically detect cocaine intake from ECG measurements. We develop a dynamical system model of the parasympathetic nervous system (PNS) behavior from the heart rate recovery observed upon conclusion of cocaine-free physical activity episodes. We also develop dynamical system models of both sympathetic nervous system (SNS) and its weakening due to cocaine metabolism, in order to model the cocaine-dampened PNS recovery. We then evaluate the performance of our model on both lab and field data using generalized likelihood ratio test approach.

3.1   Data Collection

We designed and conducted an in-residence user study with 3 cocaine dependent volunteers at Johns Hopkins University Medical School (termed "JHU Lab Study"). We conducted a second in-residence study with 6 cocaine using volunteers at National Institute on Drug Abuse Intramural Research Program (NIDA IRP) (termed "NIDA Lab Study"). We also conducted a field study with 42 active poly-drug users (for 4 weeks of sensor wearing per user) at NIDA IRP (termed "NIDA Field Study") to collect sensor data in the free-living environment. All studies were conducted upon approval from the Institutional Review Board (IRB) of the respective institutions.

Sensor Suite. We used a wearable wireless sensor suite called AutoSense [2]. AutoSense uses a flexible band worn about the chest to capture respiration data via inductive plethysmography (called RIP). The chest sensor unit also contains two-lead electrocardiograph (ECG), 3-axis accelerometer, temperature sensors (ambient and skin), and galvanic skin response (GSR). The sampling rates for the chest band were 21.33 Hz for RIP, 64 Hz for ECG, 10.67 Hz

for each of the three axes of accelerometers and GSR, and 1 Hz for the two temperature sensors and the battery level. Fig. 3.1 shows how to wear the AutoSense sensor suite. These samples were transmitted wirelessly using ANT

Fig. 3.1: Respiratory Inductive Plethysmograph band (in blue color) is worn around the chest area and wearable AutoSense sensor unit clips to the belt. A 3-axis Accelerometer, ECG, and Galvanic Skin Response sensors are also included in the same sensor unit. Mobile phones are capable of capturing data wirelessly from this sensor.

radio [24] to a Sony Ericsson Xperia X8 smart phone at the rate of 28 packets/second, each of which was 8 bytes long, containing 5 samples. The sensors last more that 10 days on a 750 mAh battery. We used the FieldStream mobile phone software [2] for logging the data. On a full charge, the phone battery lasts over 11 hours.

In the following, we describe the study protocols and the measurements obtained from the three studies.

3.1.1 In-Residence Study Protocols

In the "JHU Lab Study," three non-treatment seeking cocaine dependent volunteers (37-41 years old, 2 males) enrolled in a behavioral pharmacology residential study at Johns Hopkins University. They wore the AutoSense for at least 8 hours daily on the weeks when cocaine self-administration sessions were

scheduled. During a safety session, participants self-administered intravenous cocaine doses of placebo (1 mg), 10 mg, 20 mg and 40 mg (every 30 minutes) via a Patient-Controlled Analgesia (PCA) pump. Safety data were collected only for two participants. During study weeks 1, 3 and 5, participants went through a series of cocaine self-administration sessions. Dose Response session on Monday consisted of three doses 45 minutes apart — placebo (1mg), 20 mg and 40 mg. On Tuesday, Wednesday and Thursday, a *sample* of cocaine dose was self-administered in the morning (double-blind randomized, out of placebo, 20 mg or 40 mg). After 2 hours, they were offered 7 *choices* for either the same morning drug dose or decreasing amounts of money in a sample/choice session. In summary, each participant went through 13 days of cocaine administration (1 safety session, and 3 study weeks of 4 days each). During rest of the awake hours of the day when wearing AutoSense, the participants self reported smoking and craving events as well as some events that may trigger craving (e.g., watching TV, watching movies, or playing video games).

In the "NIDA Lab Study," healthy cocaine users are admitted to a secure residential research unit at NIDA IRP. They undergo baseline assessments on Day -1, receive training on Day 0, and receive single doses of intravenous cocaine (25 mg) on Days 1, 5 and 10. On Days 1, 5, and 10, dried blood spot specimens are collected up to 3 times daily over 1.5 hours. Single oral doses of acetazolamide (15 mg) are given on Days 2-5 and quinine (80 mg) on Days 7-10. Blood, oral fluid, and breath specimens are collected for up to 71 hours, 70 hours, and 22 hours, respectively, after drug administration on Days 1, 4, 5, 9 and 10. Participants wear AutoSense on Days 1, 3, 4, 5, 8, 9 and 10 for up to 12 hours each day. Six participants have completed this lab study.

Data obtained from the in-residence studies are used to develop the model for detecting cocaine use since it consists of clean and carefully labeled

activity-free cocaine use events, as well as various confounding events such as physical activity, games, smoking, etc., that are free from cocaine use.

3.1.2 Field Study Procedure

Methadone-maintained poly-drug (Cocaine, Heroin, etc.) users (different from those in the lab study) were recruited for NIDA field study. Participants wore AutoSense for four one-week periods, in their natural free-living environment. Participants were asked to self-report drug craving and use events in the field by pressing a button on a study smart phone (different from the one that collected ECG sensor data). Whenever they self-reported a drug-use event, they were asked to provide additional information on the timing of the drug use by choosing one of the options provided on the mobile device. The options were: (1) Less than 5 minutes ago, (2) 5-15 minutes ago, (3) 15-30 minutes ago, and (4) More than 30 minutes ago (in which case they were asked to input an estimate of the time of drug use). Urine samples were collected three times weekly (Monday, Wednesday, and Friday) during weeks when participants wore the sensors. Forty two participants have participated in the field study. Data collected in this study are used to understand the challenges in detecting drug use in the field environment and to validate our model for detecting cocaine use in the field setting.

3.1.3 Data Collected

In the "JHU Lab Study," the first participant wore the sensors only during weeks 3 and 5. Twelve days of data is collected from the first participant, 24 days from the second participant, and 22 days from the third participant, making for a total of 58 days. Of these 8, 13, and 13 days respectively are from cocaine administration sessions. An average of 9.55 hours of data was collected per day, for a total of 554 hours of good quality usable data.

In developing the model, we did not use the data corresponding to the 10

mg dose as we do not observe significant physiological response for this dosage level. From the remaining 36 sessions, sensor data is lost for 9 episodes (5 episodes from 20 mg and 4 episodes from 40 mg), due to ECG electrode detachment and/or displacement, leaving 27 episodes (13 episodes from 20 mg and 14 episodes from 40 mg) for use in model development and evaluation.

From the 3 participants who went through the choice session, only 1 chose cocaine, providing 18 instances of active cocaine administration during the choice sessions. We do not use the choice session instances in modeling since cocaine injections were administered only 15 minutes apart, not allowing enough time for the physiology (i.e., heart rate) to recover before re-administration. Further, the effects of multiple doses accumulate, making it quite different from a single cocaine dose response.

In the "NIDA Lab Study," 31 days of data have thus far been collected, of which 14 days had cocaine sessions. Each cocaine day had only one cocaine session of 25 mg. A total of 280 hours (or, 9 hours/day) of good quality ECG data have been collected from these 31 days.

In the "NIDA Field Study," we have thus far collected 922 person days of field data from 42 participants (4 participants did not complete). A total of 10,449 hours of ECG data has been collected (11.33 hours/day, on average). On those days, participants self-reported 211 instances of illicit drug (142 for cocaine) use events. Each self-report has the time of drug use, type of drug (e.g., Cocaine, Heroin, Methamphetamine, opiates, THC, Benzodiazepine), quantity, and how the drug was administered (e.g., smoking, sniffing/snorting, oral, or intravenous). Sometimes, the participants reported the use of multiple types of drug at the same time.

Among the 42 participants, 20 participants actually reported cocaine use. Among the 142 reported cocaine uses, 3 were for 50 mg, 86 for 100 mg, 1 for 150

mg, 34 for 200 mg, 8 for 300 mg, 9 for 400 mg and 1 for 600 mg. For modality of use, smoking was most popular (52 out of 142), followed by intravenous (48), sniffing (39), and oral (3).

We observe several issues with data collection in field. In several instances, data is unavailable due to not wearing the sensor or of not usable quality due to improper or loose contact of ECG electrodes, electrode detachment, loosening of electrical connectors, drying out of gel, noise from physical movement, etc. We adopt a method proposed in [25] for determining acceptability of ECG signals. In addition, a human expert visualized the signals and corrected any error in the labeling produced by the automated algorithm though they are very few in number. Since the participants did not report the exact time of the drug use, we measure the availability and quality of our sensor data for each person day to verify the usability of the collected data. Also, the self-report may not always be accurate. Hence, we visually inspect the ECG, and accelerometer signals together with the drug use self-reports to ascertain whether data in the vicinity of self-report indeed exhibits a drug use response.

Out of the 142 self-reports of cocaine use (from 103 person days), sensors were not worn on 17 days (for 22 reports). Of the remaining, 25 episodes are reported after the participant took off the sensor at night. Of the remaining 95, sensor was taken off during drug use and then put back on in 3 cases. Accelerometry sensor was not working in 6 cases, ECG sensor data was missing around the report in 6 cases, and ECG data quality was unacceptable in 53 cases. Hence, we are left with 27 instances with good quality ECG data. Among the 27 instances of cocaine uses, 14 were for 100 mg, 8 for 200 mg, and 2 for 300 mg and the remaining 3 instance were for 400 mg. These 27 instances come from 320 person days of good quality data, from 13 different participants (out of the 20 who reported cocaine use), where physiological response in the vicinity of self

report of drug use can be observed. On these 320 days, we have a total of 3,631 hours of data.

Urine Reports. Urine samples were collected three times weekly during weeks when participants wore AutoSense. We classify each day using the result from the urine test via the following simple rule. If cocaine is detected from the urine sample collected, one can infer that in the last 4 days from that day, there must be at least one cocaine use. We label all these days as potential cocaine use day. On the other hand, if the urine report is negative, one can assume that the last 24 hours is a cocaine-free day. Also, if a participant has two or more negative urine reports in a row, with no self-report of cocaine use, we can safely infer that they didn't use cocaine at all during those days. However, if the participant self reports cocaine use on a particular day, we mark that day as a potential cocaine day, even if the urine report doesn't reflect it. Using these rules, we identified 385 potential cocaine use days out of the 922 days of data collection.

3.2 Data Processing and Modeling

In this section, we describe the data processing steps and model development. Fig. 3.2 presents an overview of the data processing stages. First, we describe the processing of ECG data in order to obtain an RR interval time series. We also describe the activity detection method that we use to identify segments in the time series that may correspond to physical activity. We then describe the process we use to localize the drug and activity episodes in the RR interval time series. Next, we provide simple rules to screen out some of these segments or windows that obviously do not come from drug episodes. Each of the windows contains two parts. First part corresponds to the excitation of the SNS system up to the point when heart rate reaches it's peak value. We call this the activation response. The second part corresponds to the portion of the window that represents the recovery of the physiology. We extract these two portions and
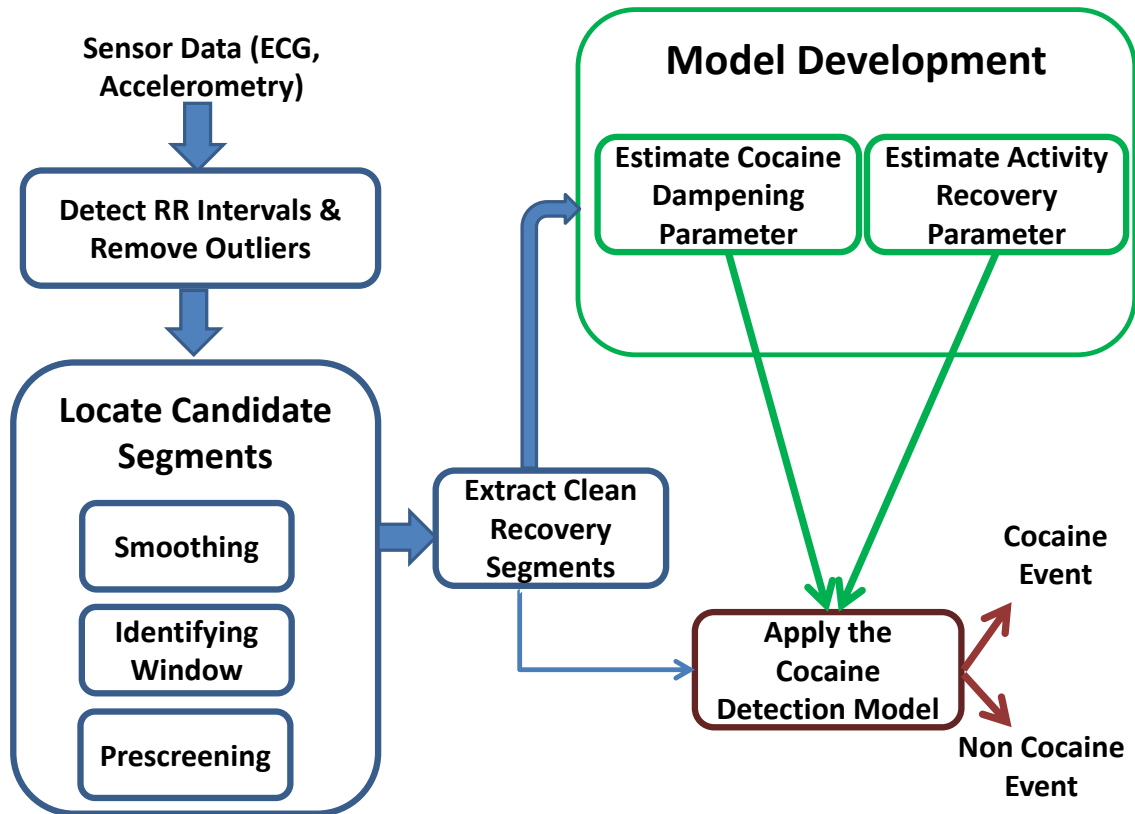
Fig. 3.2: Sensor data are first processed to find the RR interval timeseries and activity. Next, the start and end of cocaine and activity windows are identified and the recovery portions of the cocaine and activity responses are extracted. Two model parameters are learned during the training phase and the trained models are applied on the recovery response portion of the candidate windows extracted from the test data set.

construct models for the RR interval curve corresponding to recovery during cocaine-free activity episodes and activity-free cocaine episodes. Detection of cocaine use episodes makes use of these two models.

3.2.1   RR interval detection

We follow similar preprocessing and RR interval detection method as presented in [6]. We first adopt a method proposed in [25] for determining acceptability of ECG signals. We then apply the Tompkin's algorithm [26] to detect R-peaks. We remove outliers in the time series of RR intervals (time between two

R-peaks) using the outlier removal algorithm presented in [27] and remove the DC offset from each interval to remove the baseline drift.

3.2.2 Activity Detection

In order to determine the presence of physical activity from accelerometery, we use a simple threshold based activity detector using the 3-axis on-body accelerometer (placed on chest). Phone accelerometer data was not used because the phone may not be on the person and thus may miss some physical activity episodes. We make use of existing physical movement detection approach [28, 29, 30] and adapt it to fit our data. As the placement of the accelerometer and the participant population is different from that presented in prior works, we collected training data to determine an appropriate threshold for detecting activity.

We collected labeled data under walking and running (266 minutes), and stationary (183 minutes) states from seven participants while they wore AutoSense. After noise and drift removal, we extract the standard deviation of magnitude, from 10 second windows, which is independent of the orientation of the accelerometers. We scale this signal using the 99^{th} and 1^{st} percentile values of the signal using $m = (m - h)/(max - l)$, where m, h and l are the samples' 99^{th} percentile and 1^{st} percentile values. We find that a threshold of 0.35 is able to distinguish stationary from non-stationary states with an accuracy of 93% in 10-fold cross-validation.

3.2.3 Model Development

To develop a model, we first develop a method to smooth the RR interval time series, identify the windows that exhibit sufficient change so as to result from a potential drug use event, locate the start and end points of such windows, and then extract portions of this window from which the parameters of the parasympathetic nervous system (PNS) and the parameters for dampening

introduced by Cocaine metabolism induced sympathetic nervous system (SNS) activation can be estimated. These steps are described in Section 3.2.3.1. Section 3.2.3.2 presents the rationale and the actual model for detecting drug use events in the selected candidate windows of RR intervals.

3.2.3.1 Candidate Window Selection & Preparation

Fig. 3.3 shows examples of cocaine responses observed during lab administration of cocaine. We observe that the effect of cocaine can last a long time (more than 30 minutes) and the entire response window (which can vary in length as well as intensity of response) must be detected as one single event. Hence, the problem of detecting drug use is a time series pattern recognition problem. The challenge then is to locate candidate windows and identify its boundaries so as to assess it for drug use response. This challenge is compounded by dynamic variations in RR intervals and difficulty in locating the start of cocaine or activity response and the end of recovery.

This problem has similarity with the dynamic fluctuations in stock prices. Therefore, in order to identify the potential windows and to identify the start and end of the window that may indicate a cocaine use event, we use the Moving Average Convergence Divergence method (MACD) [31]. This method is widely used to compute an indicator that investors in stock markets use to identify the stock price rise and fall trends. The indicator is also used as an oscillator indicator that is used to identify when the market moves sideways, i.e., when the price oscillates within a narrow range. MACD thus is more stable than (price) trend following indicators. We make use of this MACD procedure to identify the windows that correspond to 'large' rise and fall trends of the RR intervals. MACD readily provides the start and end of the windows. The MACD procedure makes use of Exponential Moving Averages (EMA), which gives more weight (a constant in this case) to recent values compared to simple moving averages. MACD is

21

computed as follows:

$$\text{MACD Line} = EMA_{slow} - EMA_{fast} \qquad (3.1)$$

$$\text{Signal Line} = \text{EMA of MACD Line} \qquad (3.2)$$

Each of the EMA's takes one parameter — the window size on which the EMA is to be computed. EMA_{slow} is computed on a longer window than that of EMA_{fast}. The crossover points of the MACD line and the Signal line corresponds to the fall and rise points. We learn the parameter (i.e., window size) from the data we collect in the "JHU Lab Study." We mark the start and end of the windows for cocaine data by visual inspection. For activity, we select start and end of the window with the help of accelerometry and visual inspection. In total, we mark 27 cocaine windows and 272 activity windows. The search space for the parameters of the slow and fast moving averages are set to [1, 180] minutes and [2, 90] minutes respectively. The search space for the parameter of EMA_{slow} is same as that of EMA_{fast}. We obtain the parameters that achieve the minimum error in finding the start and end points. The search algorithm chooses crossover points that are nearest to the start and end points marked for each set of parameter values. The error in this decision is computed as the sum of the distance from the crossover points chosen by the algorithm and the start and end points marked via visual inspection.

The input to the MACD process is smoothed RR interval time series that removes high-frequency variations in RR intervals. We use a simple moving average over preceding 10 minutes to smooth the signal. We also learn the optimal window size of this moving average window in this process. Fig. 3.3 shows the MACD, signal line as well as the crossover points selected in this process. The parameters learned for EMA_{slow}, EMA_{fast}, and EMA for Signal

Line are 35 minutes, 4 minutes and 3.67 minutes respectively. These parameters are used to find the windows on the field data too.

      Screening. Next, we use the following two rules to screen out some of the candidate windows before we apply our model. First, we compute average width and height of the windows from the lab data. If the candidate window c is too 'small' compared to the drug windows (i.e., if width or height $< (\text{mean} - 3 * \text{standard deviation})$ of that of width and height of windows corresponding to drug response), it is discarded. In this case, the physiological response is insufficient to be that from drug use. Second, out of all the 10 second accelerometer measurement windows within the first 5 minutes from the start of a candidate window c, if a majority of them are detected as activity we discard the window c. In this case, the activation of heart rate is the result of physical activity.

Fig. 3.3: Illustration of MACD based candidate window selection method. The crossover points, where the MACD and signal lines cross (marked using dotted horizontal lines), indicate the start and end points of the candidate windows.

Extracting recovery episodes. We extract both the recovery and activation part of the windows using the crossover points from the MACD process. For a particular window, the start and end of activation (A_{start} and A_{end} respectively) must occur after the start point of the window. They are defined as the last crossover point above the zero line and first crossover point below the zero line. On the other hand, the start and end of the recovery portion (R_{start} and R_{end} respectively) must occur after the activation part. They are defined as the last crossover point below zero line and first crossover above zero line. Finally, to model the recovery process, we extract the first subsequence of the recovery curve (defined by R_{start} and R_{end}) that is clean, i.e., not affected by activity.

3.2.3.2 Autonomous Nervous System (ANS) Model

To develop a model to detect drug use events, we consider an abstract model describing the interaction of the ANS with the cardiovascular system during physical activity and drug intake. ANS is a control system that regulates heart and respiration rate, cardiac output and constriction of dilation of blood vessels to meet the demands imposed on the body. The ANS is divided into two separate systems — PNS and SNS acting in concert. Roughly speaking, PNS is responsible in conservation of energy by reducing heart, respiration rate and blood pressure providing a dampening effect. In contrast, SNS provides activation by increasing heart rate, blood pressure and cardiac output to meet the demands of physical activity, fear, stress, etc. Generally speaking, SNS works at a faster time scale than PNS and can be excited by many inputs, not all directly observable. In this work, we consider modeling heart rate recovery regulated by PNS with and without drug usage events. We learn a simplified abstract model with few parameters driven by physiology from the lab data to build a statistically optimal detector. The concise model avoids over-fitting and provides parameters readily interpretable as time constants and dosage.
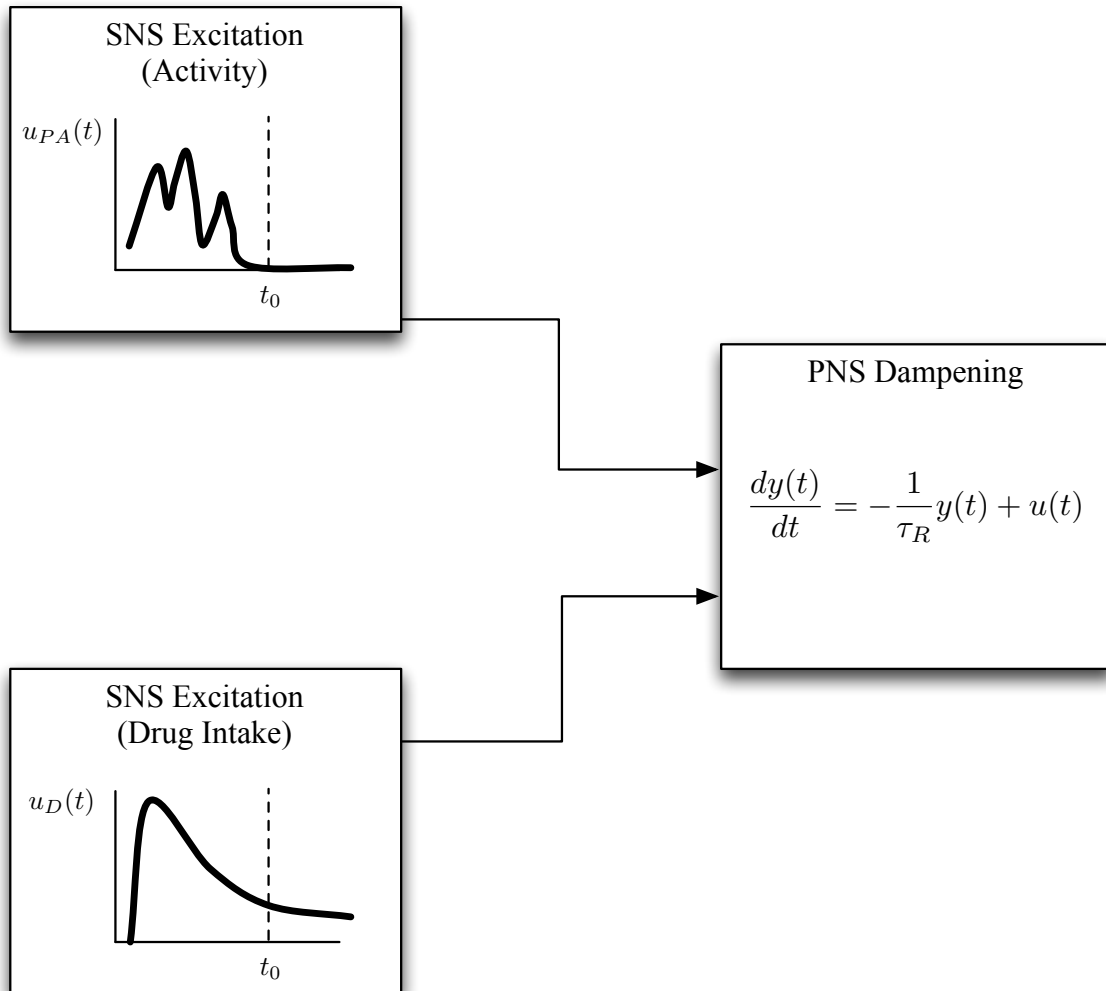
Fig. 3.4: ANS model of heart rate recovery.

We would also like to note that the activation data portion of the response is not used in our modeling. There are several difficulties associated with using this portion of the data. First, during activation there is the combined effect of Vagal withdrawal and SNS activation as well as the metabolism of cocaine. Therefore, in order to model the activation curve, we need to estimate the parameters associated with Vagal withdrawal and SNS activation. However, the activation portion is quite short, not offering sufficient data for robust modeling.

We focus, therefore, on modeling the RR interval time series $r(t)$ (inverse of the instantaneous heart rate) during recovery periods. We assume that the RR

interval time series $r(t)$ can be modeled as noisy samples of a signal $y(t)$ measured from the resting baseline RR interval B.

$$r(t) = B - y(t) + n(t) \qquad (3.3)$$

We assume that the PNS can be modeled by a first-order differential equation:

$$\frac{dy(t)}{dt} = -\frac{1}{\tau_R} y(t) + u(t), \qquad (3.4)$$

where $u(t)$ models the excitatory input from the SNS. In the absence of any excitatory input $u(t) = 0$. In such a case, the solution corresponds to an exponential decay model, with a recovery time constant of τ_R. We model the natural recovery process as

$$y_R(t) = y(t_0) e^{-\frac{(t-t_0)}{\tau_R}}. \qquad (3.5)$$

Such exponential decay models have been previously suggested as models of heart rate recovery from physical exercise [32]. Here, we consider two sources of excitatory input — physical activity and drug intake. Both can elevate the heart rate. We focus on recovery portions where the excitatory input due to activity is absent, so that we need to model only one source of activation. We note that although the activity is missing during our chosen recovery periods, the initial heart rate elevation could be due to both physical activity and drug intake, complicating the detection of drug usage and estimation of dosage. We model the excitation due to drug intake as an exponentially decaying process since its

26

intensity decays due to cocaine metabolism[1] with time-constant τ_D:

$$u_D(t) = u(t_0)e^{-\frac{(t-t_0)}{\tau_D}}.$$

Our ANS model of recovery in the absence of physical activity is given in Figure 3.4. Inserting this input into the differential equation governing dynamics of the PNS, results in the following model of recovery curve under drug intake:

$$y_{DR}(t) = y(t_0)e^{-\frac{(t-t_0)}{\tau_R}} - \frac{u(t_0)}{K}e^{-\frac{(t-t_0)}{\tau_R}} + \frac{u(t_0)}{K}e^{-\frac{(t-t_0)}{\tau_D}}, \tag{3.6}$$

where the constant K is given by $K = (1/\tau_R - 1/\tau_D)$. This represents the rate of recovery by the PNS system when cocaine is still in the bloodstream (i.e., not fully metabolized).

We note that although we build upon established model for modeling the recovery behavior of the parasympathetic nervous system (PNS) from activity (i.e., in the absence of cocaine), but our proposed model introduces two new components to account for the dampening of PNS recovery due to cocaine stimulus and to tease it out from the effect of activity. Prior models of recovery from cocaine (e.g., [23]) consider only the composite recovery of heart rate for cocaine events and hence can't separate out the effects of physical activity. Our proposed ANS model is the first model that considers the dampening effect of cocaine and the decay of dampening due to cocaine metabolism by the body in PNS recovery, and accounts for the recovery from physical activity alone. The first component on the right hand side of (3.6) represents the natural recovery of the heart rate, while the 2nd and 3rd components represent the cocaine dampening and the decay of dampening due to metabolism of cocaine, respectively.

The detection of drug intake events can now be posed as a binary

---

[1] Drug metabolism is usually modeled as an exponentially decaying function since the rate of metabolism is proportional to its remaining concentration.

hypothesis testing problem of identifying the underlying curve from the two models in (3.5) and (3.6) from noise corrupted recovery segments of RR interval time series. We assume the noise on the time samples of $y(t)$ is identical and independently distributed with Gaussian distribution. The nuisance parameters of intercept $y(t_0)$ and initial drug intensity $u(t_0)$ are unknown. Therefore, we follow a generalized likelihood ratio test approach where maximum likelihood estimate of these parameters are used. The detector for drug intake events from RR interval samples $r(t)$ take the form of

$$\frac{\|B - r(t) - y_{RD}(t; \hat{y}(t_0), \hat{u}(t_0))\|^2}{\|B - r(t) - y_R(t; \hat{y}(t_0))\|^2} < \theta$$

where the nuisance parameters are found using least-square fit to the segment under test. The threshold θ has to be chosen to set an appropriate balance between probability of detection and probability of false alarm. For the detected drug episodes the parameter estimate $\hat{u}(t_0)$ is the intensity of remaining drug excitation at the beginning of the recovery segment. If reliable estimates of the drug intake time can be formed, this could be used to extrapolate $u(t)$ to drug-intake time to find an estimate of the drug dosage, which we leave for future work.

The ANS model parameters consist of resting baseline B, recovery time constant τ_R and drug metabolism time constant τ_D. In general, these parameters have to be learned for each individual separately. But, obtaining ground truth data on drug events for each individual is impractical. Therefore, we follow the following unsupervised procedure to construct semi-individual models of ANS activity. For each participant j, we extract segments for recovery from physical activity occurring naturally in their daily life and fit model in (3.5) and learn the

PNS model parameters of the resting baseline B^j and recovery time constant τ_R^j for each subject.

For time-constant determination, the exponential moving averages (EMA) from section 3.2.3.1 cannot be employed for the RR interval signal $y(t)$, since the long time averaging windows used in EMA's will cause an upward bias in the estimate of the recovery time constants. On the other hand, raw RR interval data collected in the field contains significant amount of outliers. These outliers are not well modeled with Gaussian statistics, since they have large amplitudes with a proportion larger than indicated by the tails of the normal distribution. As a result, least square estimates of τ_R and τ_D will be highly sensitive to outliers present in the unsmoothed RR interval data. To have robust estimates of these time constants we follow the robust regression framework introduce by Huber [33]. Specifically, we consider Huber's modified cost-function $\rho(e(t))$,

$$\rho(e) = \begin{cases} 0.5e^2 & \text{for } |e| \leq k \\ k|e| - 0.5k^2 & \text{for } |e| > k \end{cases}, \tag{3.7}$$

where $e(t)$ is the error signal defined as the difference between the data $y(t)$ and the models in (3.5) and (3.6). The constant $k = 1.345\sigma$, with σ is the expected variance of RR interval data. The modified cost function eliminates large bias introduced by non-Gaussian outliers present in RR interval data. Then, for the nine participants with lab data on drug usage we fit the drug recovery model in conjunction with their PNS model of (B^j, τ_R^j) to get an average estimate of drug metabolism constant τ_D minutes (for each dose level) that we use uniformly for every participant in our testing.

3.3   Evaluation

We first present the performance of the ECG cycle based detector presented in [6] on our in-residence lab data. Next, we present estimates of

parameters for our ANS model. To provide a feel for our model, we provide an illustration of our model by showing their fitting with data from a participant. Finally, we show the performance of our ANS model on the data collected in the in-residence studies and the field study.

3.3.1 Performance of Cycle Classifier

To assess the applicability of a model developed in [6] for detecting cocaine usage in free-living conditions, we implemented the classifier in [6] and applied it to the data collected from our in-residence study. As mentioned in Chapter 1, this model classifies each ECG cycle into cocaine and and baseline classes. The in-residence study, in our case, did not include a baseline session, therefore the training data for non-cocaine class consists of data collected from 30 minutes to 2 minutes prior to the first cocaine administration of the day, since subjects are resting for this period. Data only for durations when there was no physical activity were labeled as non-cocaine. This is determined from the accelerometer data. If activity is detected for a particular activity detection window, we discard data up to 5 minutes after that window to make sure the data does not contain any activity. As for the cocaine class, data from the interval (valley -2 minutes) to (valley $+8$ minutes) from each cocaine response window are included. The cocaine response windows are identified by the Moving Average Convergence Divergence (MACD) method described in Section 3.2.3.1. The valley of the window is defined to be the point when the RR interval reaches its minimum value (i.e., where HR peaks). Again, we only admit data that is not effected by activity. Total data cases for cocaine and non-cocaine class for participant #1 amounts to 2,080 cycles each. A support vector machine (SVM) classifier is trained and we obtain AUC = 0.91 for 10-fold cross-validation, which is similar to that reported in [6]. We then apply this classifier on the whole data set

for that participant. To show the applicability of this model on non-baseline data we show the output of the classifier in Fig. 3.5.

Fig. 3.5: The red dots on the top indicate all the cycles marked by the ECG cycle based classifier as cocaine class. The model produces too many false positive outputs rendering it unsuitable for the use in field setting.

We observe that although the classifier performs well for the 20 mg case, it's performance for both 40 mg and placebo (1 mg) cases are not as good. Also, we find that the classifier produces a large number of false positives when there is physical activity since it is not trained for this class. This indicates that not only RR intervals but also other features (hypothesized in [6] to be insensitive to physical activity) are confounded by activity. It remains open whether the model proposed in [6] can be improved by training it with free-living activity data. Another approach to improving it may be to combine this data-driven approach of classifying each ECG cycle with our model-based approach.

3.3.2 Evaluation of the ANS Model

We first estimate parameters of our ANS model and then evaluate its performance on the in-residence and field data.

Parameter Estimates Our ANS model needs two parameters to be estimated — τ_R and τ_D. The $1/\tau_R$ and $1/\tau_D$ values represent the half-time (in minutes) to recover from cocaine-free physical activity and for cocaine to metabolize, respectively. τ_D is estimated from activity-free cocaine episodes, while τ_R is estimated from cocaine-free activity episodes. Table 3.1 presents the data obtained from the in-residence and the field studies. From both in-residence studies, we obtain a total of 41 cocaine episodes from 40 days of data. A total of 497 non-cocaine significant physical activity episodes are found on 49 non-cocaine days.

Since there are several episodes of cocaine-free recovery from activity for each participant, we obtain a person specific estimate of τ_R for each participant (in both lab and field) and use it in testing. Table 3.2 shows estimates for τ_R, which are also used in clinical settings for assessing cardiovascular fitness [15, 16]. We observe that it takes approximately 3 minutes on lab data (and 4 minutes on field data, due to higher intensity of activity episodes) for half-recovery of heart rate in the absence of cocaine.

To estimate τ_D, we need cocaine recovery curves that are free from activity. Since we have activity-free recovery from cocaine for only in-residence participants ($n = 9$), we obtain person-specific estimates for τ_D only for these participants. Table 3.1 shows estimates for τ_D for various dosage amounts. We observe that the half-time to cocaine metabolism (i.e., $1/\tau_D$) increases with increase in dosage. The median half-time for cocaine metabolism is 43 minutes.

When fitting the model to recovery curves of RR intervals, we use person-specific estimates of τ_R, but a person-independent estimate (i.e., median

32

Fig. 3.6: Curve fitting for activity recovery. The four solid green curves represent four raw RR interval timeseries of activity recovery of one participant. Fitting the recovery model (dotted blue curves) we find the estimate of τ_R. Based on this τ_R value we observe that it takes approximately 2.78 minutes for the heart rate to recover to half of the resting baseline for this participant.

over 9 lab participants) of τ_D. Figs. 3.6 and 3.7 show examples of fitting of the models to recovery of RR intervals from an activity episode and from a cocaine use episode. Fig. 3.8 shows an example of fitting an activity recovery model and a cocaine recovery model onto an RR recovery curve from a cocaine use episode. We observe that the recovery rates of the two curves are quite distinguishable, with cocaine recovery model providing a better fit.

Performance on In-residence Data. Although we obtain person-specific estimates for τ_R, we need a person-independent estimate of τ_D for use on the field data. To determine the best estimate of τ_D, we obtain dose-specific estimates of τ_D (for 20 mg, 25 mg, and 40 mg) and dose-independent estimate of

Fig. 3.7: Curve fitting for cocaine recovery. Using the activity recovery model we find the estimate of τ_R from the activity recovery data of this participant. τ_D is estimated from this participant's cocaine response data (solid green curves) by fitting the cocaine recovery model that uses the estimated value of $\tau_R = 0.3576$. In this case the estimated value for τ_D is 0.0282 which gives the half life of cocaine in the blood to be 35.46 minutes.

τ_D (by using all dosage). To determine their suitability, we test their performance on lab data. We test their performance on the dose-specific in-residence dataset (In-residence-Dose) (i.e., on all the cocaine of a particular dose and all non-cocaine episodes of the in-residence studies). This provides an indication of their expected performance if all ocaine episodes are of the same amount. We also test their performance on all lab data (In-residence-All) to assess their invariance to dosage level.

We compute the number of false positives (i.e., the number of non-cocaine episodes identified as cocaine episodes) per day, without misidentifying any cocaine episodes (i.e., true positive rate of 100%). Results appear in Table 3.3.

Fig. 3.8: Fitting of an example cocaine recovery curve to the models. The dotted blue and red curves represent the fitted activity recovery and cocaine recovery models respectively. It is clearly observed that the cocaine recovery model [represented by Equation-3.6] that takes into account both τ_R and τ_D performs better.

We observe that the 25 mg model provides the lowest false positive rate on the In-residence-Dose dataset. But, for the In-residence-All dataset, the lowest false positive/day occurs for τ_D obtained from the 40 mg cocaine data. We, therefore, use this time constant in all of our testing. The ROC for the 40 mg model with the data of all participants combined is presented in Fig. 3.9 with a representative operating point shown.

Performance on Field data. As mentioned in Section 3.1.2, there are 27 cocaine instances from 13 participants for which we have good quality data. These 27 instances come from 25 different days on which the participant both self reported cocaine use and are confirmed with positive urine reports. We also have

Table 3.1: Data statistics for In-residence and Field studies.

Study	Non-cocaine		Cocaine	
	# windows	# days	# windows	# days
In-residence Studies	497	49	41	40
Field Study	740	64	27	25

Table 3.2: Model Parameters ($1/\tau_R$ and $1/\tau_D$) in minutes for different cocaine dose amounts. (L) and (F) refers to the $1/\tau_R$ values computed for the in-residence and field studies respectively. $1/\tau_D$ values for 20mg, 25mg, 40mg, and combined models are presented in the columns labeled $1/\tau_D(20)$, $1/\tau_D(25)$, $1/\tau_D(40)$ and $1/\tau_D(C)$ respectively.

Measures	Range	Median
$1/\tau_R(L)$	2.41-6.75	3.18
$1/\tau_R(F)$	2.00-9.82	4.06
$1/\tau_D(20)$	27.55-75.75	30.30
$1/\tau_D(25)$	35.09-81.97	45.23
$1/\tau_D(40)$	40.16-61.35	51.02
$1/\tau_D(C)$	35.09-81.97	43.29

740 episodes of non-cocaine windows of significant physical activity that come from 64 confirmed non-cocaine days. We obtain person-specific estimates of τ_R. Table 3.3 presents the results of applying the dose-specific and dose-independent models on the field data. An ROC curve for the 40 mg model is presented in Fig. 3.9 with a representative operating point shown for the field.

Table 3.3: Performance of the detector on in-residence and field study data for various estimates of τ_D. We report (Avg) number of False positive/day for a True Positive rate of 100%. L-All is for the dose-independent estimate of τ_D.

Dataset	L-20	L-25	L-40	L-All
In-residence-All	0.91	0.91	0.87	0.91
In-residence-Dose	0.91	0.22	0.65	-
Field	1.13	1.14	1.13	1.14

Fig. 3.10 shows the application of the model on a sample of field data with cocaine use event reported. Our candidate window selection and screening

Fig. 3.9: ROC for detection of drug on the in-residence and field studies using the 40 mg model. The two points on the ROC curves presents two suitable operating points that have similar performances. If we allow for misclassification of drug events (TPR 95%) we can achieve false alarms rate of $< 7\%$.

method identifies several potential cocaine response windows. (These are representatives of the 740 non-cocaine windows). In Fig. 3.10, these windows are represented by dotted vertical lines marking the start and end times of each window. The output of the model for each of these windows is presented using blue or red markers over them. We find that the model detects two cocaine windows. The first window from the left marked with a red dot is the actual cocaine event (identified by the self report). The model identifies one more window as a cocaine response window. It can very well be the case that there was a second drug intake episode that the participant failed to report; there are several instances of repeated drug use episodes in the field.

More generally, from visual inspection of false positive cases, we observe that sometimes the candidate window selection method fails to identify the start of a recovery segment correctly, leading to a false positive. Hence, further reduction in false positives can be achieved by improving the window marking method.

Fig. 3.10: Output of the model on the field data. Here we find that there is one self-report of cocaine use and the participant self reports that he/she had cocaine 80 minutes prior to the self-reporting time (marked by the vertical line). The candidate windows are marked by the dotted lines. The output of the model is presented using small red dots and blue crosses over each window representing cocaine and activity respectively. We see the model makes one error - it wrongly detects one non-cocaine window as a cocaine window.

Chapter 4

Enhancement of the Model

In this chapter we present several enhancement of the model to detect drug intake in free living conditions. First, we discuss several proposed steps to improve the sensitivity and specificity of the model. Secondly, we examine the performance of the model to determine cocaine dosage as an added functionality to the model.

4.1   Improving the Sensitivity of the Model

In Fig. 4.1 provides a summary of the quality of data available from the field study. We observe that out of the 142 self reported cocaine intake events, for 56 events there is no data available. The remaining events are classified into three groups - there were 44 events for which we have acceptable amount and quality of ECG data, 31 events for which data quality was bad or data were missing at places (at least 33% of the segment was missing) and finally there were 11 events for which the accelerometer was turned off for some reason (eg., battery depletion). Out of the 44 events for which data quality was acceptable as well as accelerometer data was available, we in our preliminary work worked with only 27 events. We excluded the remaining 17 events from our analysis because the recovery segment for those events were not 'clean'. A clean recovery segment is defined as a segment, starting from the start of the heart rate recovery, for which there is at least 2 minutes worth of data which is not affected by activity. We observe that in order to obtain such clean recovery segments a large number of events were discarded in the preliminary analysis. In this chapter we relax this criteria of selecting the clean recovery segment and measure the sensitivity of the model.

In order to improve the sensitivity of the model in the presence of such noise and increase the number of potentially detectable events, fist we propose to

39
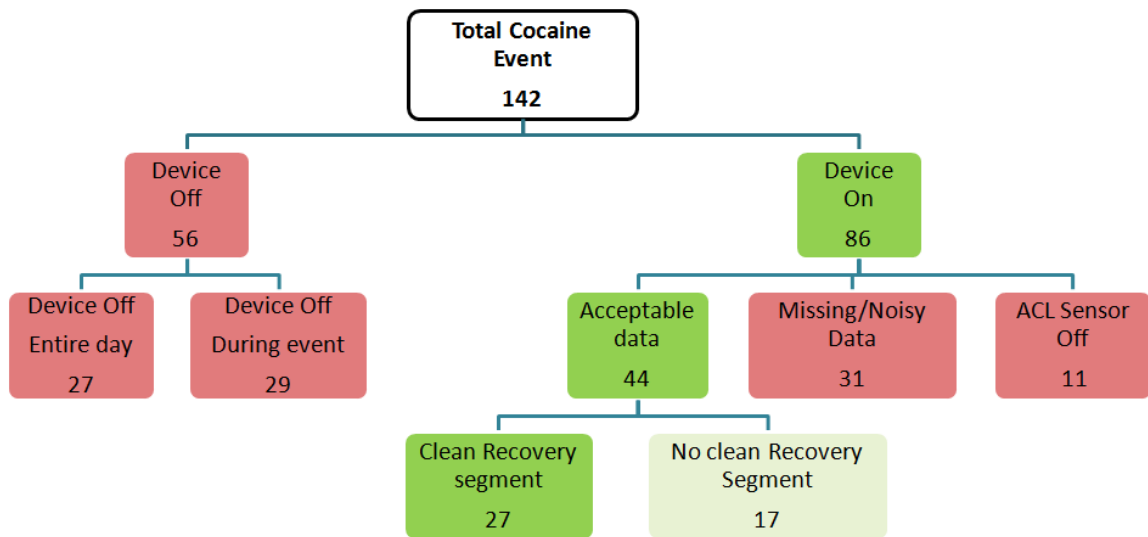
Fig. 4.1: Classification of self-reported cocaine intake events based on data quality. On the left side of the tree we see that in 56 out of the 142 events the sensor/device was turned off and so no data was collected. On the other hand, There were 86 events when the device/sensor was turned on. We in our preliminary work used data collected only for 27 of these events where we obtained 'clean' data for the recovery periods.

use segments that are smaller than the arbitrary 2 minutes threshold. Using smaller windows has the added advantages of potentially having segment with good quality of data as well as having segments that are not affected by activity. Also we get near real time detection which can be important in different intervention applications. We examine the performance of the model with different segment sizes ranging from 30 seconds to 45 minutes. The result of running the model with these window sizes on the data from non-cocaine days is shown in Fig. 4.2. We observe that, the number of false positives per day falls significantly when the window size is more than 15 minutes. However, increasing the window size after that does not seem to change the sensitivity of the model. Therefore, we conclude using very small segments is not very practical. This is likely due to

the fact that the from the lab data we find that the minimum and maximum length of the recovery segments are 17.06 minutes and 56.7 minutes respectively.

Fig. 4.2: Performance of the model with different window sizes ranging from 30seconds to 45 minutes. We observe a significant improvement of the model when the recovery segment size is increased to 15 minutes. We note that in our data, From the lab data we find that the minimum and maximum length of the recovery segments are 17.06 minutes and 56.7 minutes respectively. This is found by visual inspection which agrees with our findings here.

As we find from the above results that smaller window sizes degrades the performance of the model, we will need to work with longer segments to make the model more effective. However, as mentioned before, longer segments are likely to be affected by activity. We examine the sensitivity of the model in the presence of activities by attempting to remove the effects of activity from the heart rate or RR signal. In [34] it is reported that mean time to recovery of heart rate from activity is approximately 19.8 seconds. However, this varies from person to person. We, therefore, propose to use τ_R computed for each person to set the threshold for obtaining clean recovery segments. Instead of removing segments that do not contain 2 minutes worth of data from the start of the recovery (valley

point), we first find the duration of activity t_a within the recovery segment. This is done by putting a threshold on the standard deviation of accelerometer (26.4) magnitude. Once we find the duration of activity, we add to it the τ_R value found for that person. We assume that the heart-rate due to the activity will recover by that time. We then remove data from the recovery segment for the duration of $t_a + \tau_R$ from the start of activity and consider that during this time the data is missing from the recovery segment. We find that, this adds 7 more cocaine events making a total of 34 events. The observe (see table 4.1) that the false positives/day is reduced from 0.87 to 0.74 on the lab data and in case of field data we lower the false positive rate to below 1 per day.

Table 4.1: number of false positive per day on in-residence and field study data.

Method	In-residence	Field
Paper's Method	0.87	1.17
Removing activity from the recovery segment	0.74	0.98

4.2   Improving the Specificity of the Model

We propose to investigate the specificity of the model in detecting cocaine intake. In the field, the participants took a variety of other drugs. There is a total of 38 non-cocaine drug intake events. Similar to cocaine intake instances, participants reported the type of drug (such as Heroine, opiates, THC, Benzodiazepine, Methamphetamine and instances when multiple drugs were taken together) and the approximate time of intake. Some of these drugs have a similar heart rate response as cocaine. Figs. 4.3 and 4.4 show the RR interval signal surrounding the self reports of THC, Heroin, Methamphetamine, Morphine intake respectively. We observe that heroin and THC don't have much effect on heart rate. The recovery rate of RR interval for other drugs are different than that of cocaine.

(a) Cocaine

(b) Morphine
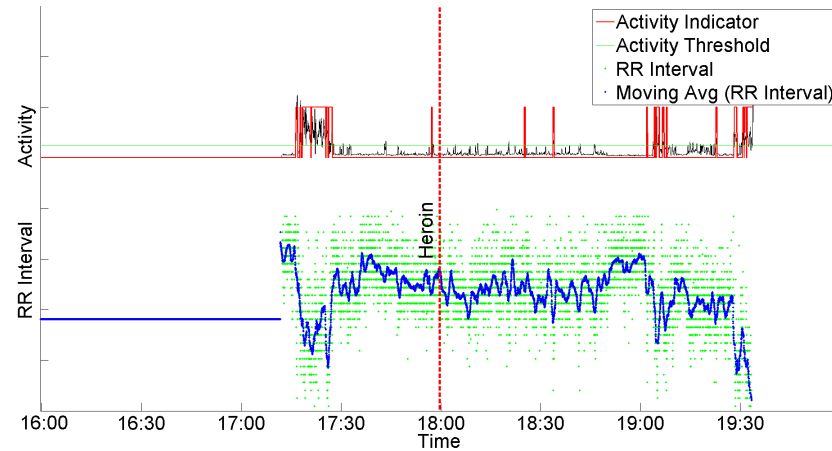
(c) Heroin

Fig. 4.3: Segment with ± 2 hours surrounding the self report. Here we find that there is one self-report of cocaine use(marked by the vertical line). The green dot represents RR interval, the blue line represents moving average of RR interval. Th red line indicates activity detected from activity threshold (black line). Figure (a), (b), (c) show the segments for Cocaine, Morphine and Heroin, respectively.
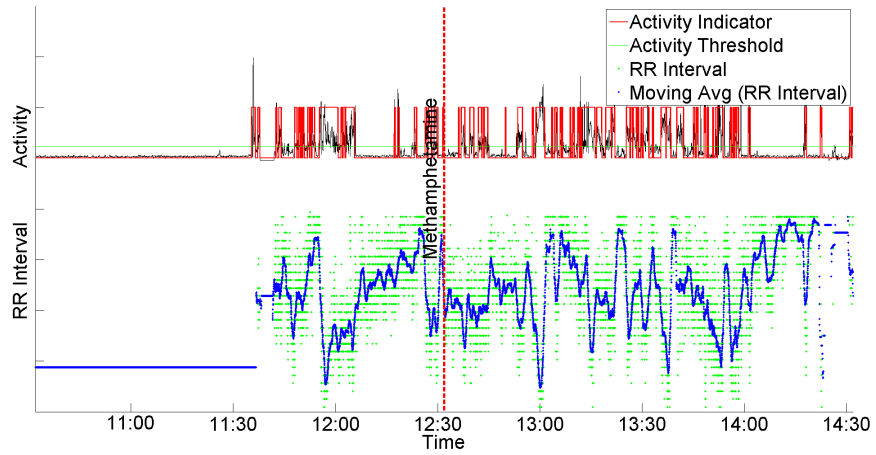
(a) Methamphetamine

(b) THC

(c) THC+Benzodiazepine

Fig. 4.4: Segment with $\pm$ 2 hours surrounding the self report. Here we find that there is one self-report of cocaine use(marked by the vertical line). The green dot represents RR interval, the blue line represents moving average of RR interval. Th red line indicates activity detected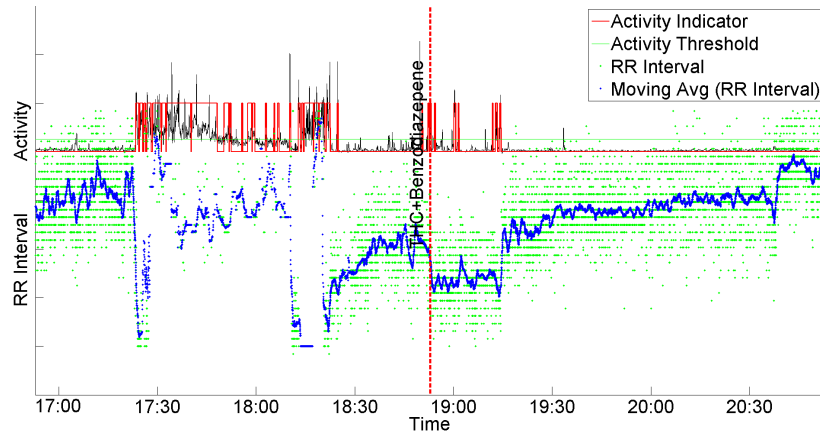 from activity threshold (black line). Figure (a), (b), (c) show the segments for Methamphetamine, THC and THC+Benzodiazepine, respectively.

In order to investigate the specificity of the model, we consider the days where

- There is a self report of drugs other than cocaine

- There is no self report for cocaine

- Urine report for cocaine is negative

We identify 16 such events in total when we consider all the above. For each self report, we consider the window with $\pm$ 2 hours surrounding the self report. For all cases there's no case that is misclassified as cocaine.

4.3 Adding Functionality to the Model - Estimating Dosage

One important functionality that can be added to the model is estimation of dosage for cocaine intake events.

Table 4.2: Cocaine Dosage in In-residence and Field studies.

Location	John Hopkins(Lab)	NIDA IRP (Lab)	NIDA IRP (Field)
# of Participants	3	6	42 (4 drop outs)
# of Episodes (with Dosage)	18 (20 mg) 18 (40 mg)	14 (25mg)	3 (50 mg) 86 (100 mg) 1 (150 mg) 34 (200 mg) 8 (300 mg) 9 (400 mg) 1 (600 mg)

Table 4.2 shows the different dosage of cocaine intake in the different studies. It can be observed that the dosage amount in the field is substantially higher than what is administered in the lab. First, we will attempt to analyze the performance of the model trained using only the lab data whether we can identify the dosage amount. In Fig. 4.5 shows that identifying 20mg and 40mg cocaine intake is likely to be possible due to the difference in recovery rates. Secondly, we

will attempt include field data that are identified using our model to develop a

better regression model for predicting the dosage from the recovery signal.
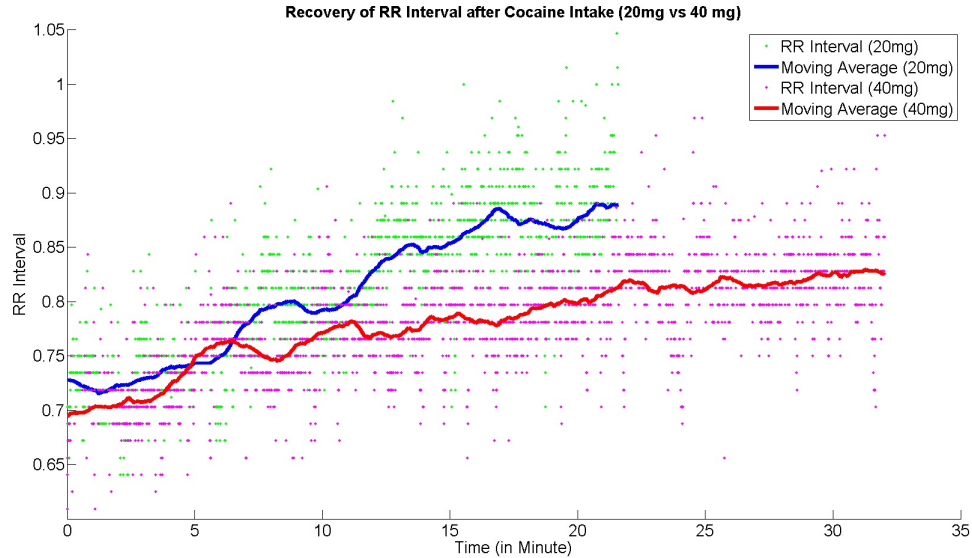
Fig. 4.5: The blue and red lines represent recovery portion of RR interval after 20 mg and 40 mg cocaine intake respectively. The recovery of RR Interval to the base line for 20 mg cocaine intake is faster than 40 mg cocaine intake.

From Table 4.2 we find that the dosage amount reported in the field study

are much higher than what was administered in the lab.

When we run the model on lab data the average error rate 27.9%.

However, average error rate is as high as 101.3% when applied to field data. The

main difficulties to detect dosage in the field are as follows:

- There are intermittent and rigorous activities in the recovery part.

- There's no lab data for higher dosage to understand the relationship

    between low dosage and high dosage.

46
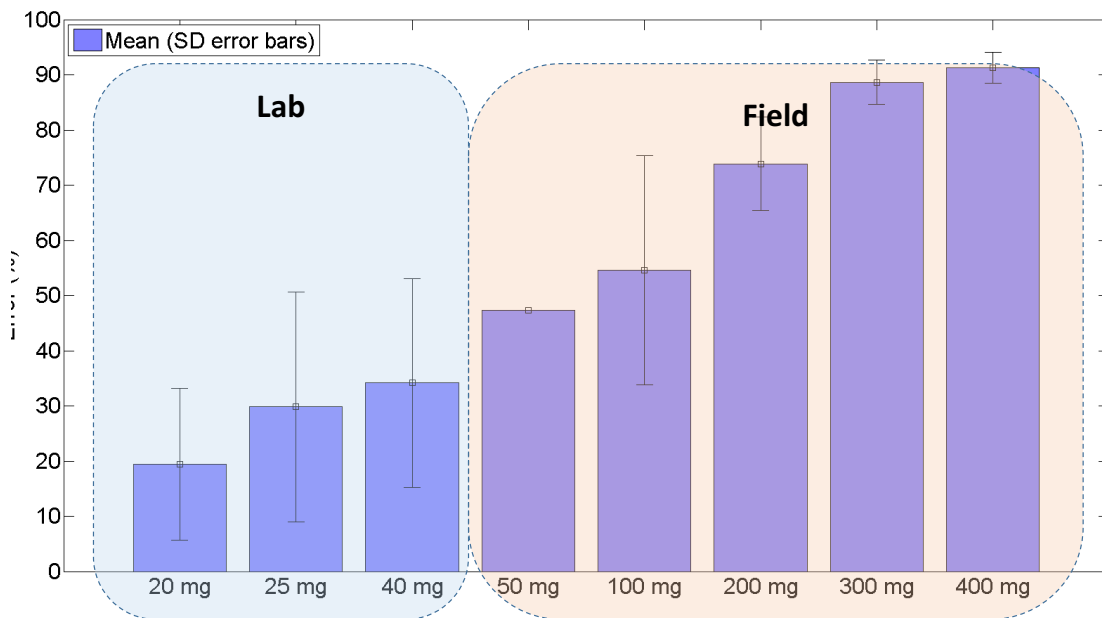
Fig. 4.6: Error rate for various dosage both in the lab and in the field

Fig. 4.7: Intermittent activity in the recovery segment of cocaine segment

Chapter 5

Discussion, Limitations and Future Work

In this chapter, we discuss several limitations of the work that can motivate future research.

First, we present the limitations of the method proposed in chapter 4 that removes affects of activity from the recovery segments. Although the proposed method, enhances the performance of the model, this approach does not consider the initial conditions of the segments that result in after the removal of the segments associated with activity. More research is required that to estimate the initial conditions of each segment. Second, in this work we have developed a cocaine intake detection model that is fairly specific. Development of detection models for other drugs that have similar (or different effects) on the heart rate is another important future research direction. Third, the model requires the computation of baseline heart-rate and recovery rates for each person, as well as cocaine metabolism rate which are computationally expensive to be implemented in a mobile phone. However, for real-time deployment of the model for Just in time interventions, it is imperative to implement such models on mobile devices. In the following, we discuss these issues in more detail and provide directions for future research.

5.1   Removing the Effect of Intermittent Physical Activity

In our approach, we identify 'clean' recovery segment of RR interval, which are not affected by physical activity, and then apply the model on it. If there is any intermittent activity detected during this recovery segment, we discard the data. For this, we first find the duration of activity t_a within the recovery segment. Once we find the duration of activity, we add to it the τ_R value found for that person. We then remove data from the recovery segment for the duration of $t_a + \tau_R$ from the start of activity and consider that during this time the data is missing from the

recovery segment. We then, fit only one exponential function for the whole segment. This method provides larger recovery segment and better false positive rate. However, there is a short coming of this approach. As there are multiple small recovery segments that are created due to discarding of the data, the initial conditions u_t0 for each segment changed. We observe that, as our approach does not estimate the initial conditions, it does not provide good estimation of cocaine dosage. Therefore, an alternative approach that fits multiple exponential functions with proper initial condition may provide better detection accuracy. It should be noted that, the cocaine metabolism rate is common for all of these segments. Applying this approach will lead us to better false positive rate as well as better estimation of dosage.

5.2   Development of Detection Models for other Drugs

The detection model is developed and tested for cocaine intake only. As have observed Cocaine is a stimulant drug. It increases heart-rate and causes hyperactivity. Among other stimulant drugs the most common one is Methmphetamines. Our exponential decay function for cocaine models the ANS to capture the metabolism and excitation of the drug. For similar types of drugs, e.g., Methamphetamines one can use similar decaying model. To train that model and learn its parameters, a similar type of lab study is required.

Other types of drugs include Opiods, Hellucinogens and depressants or sedatives. Opiods are mainly used as painkillers and they are known to slow down the central nervous system. Marijuana, psilocybin mushrooms, and LSD are examples of drugs that cause hallucinations. They can alter the user's perception. The other type of common drug that people use are depressants or sedatives. Though these are not necessarily illicit drugs. More studies are needed to collect physiological data in order to detect these drugs. Further research is also needed to find the confounding factors for these drugs. For example, sleeping causes the

49

heart rate to go down, therefore sleeping events might be confused with drugs such as intake of opioids or marijuana that also decreases the heart rate. Also activity, caffaine intake can be a confounding factor for different drugs that work as stimulants.

5.3 Real time deployment of the model

Finally in this section, we discuss the challenges in deployment of the model.

- Base Line calculation: We see that, the base line of RR interval varies from person to person. Again, for each person it varies from day to day. In our model we calculate baseline by taking 95th percentile of RR interval of each day for each person. This is used as the baseline of RR interval of that day for that person. For the real time model, we cannot do that. Instead, we can start with common base line of RR interval and as the time goes the base line needs to be updated. We pose the dynamic baseline estimation for RR interval as a future work.

- tau_R calculation: To estimate tau_R value for a person, we consider the recovery portion of RR interval of activity for non-cocaine days. We have the luxury of seeing all data. But for real-time algorithm, this tau_R value need to be updated as we process the data stream. Initially, we can start with a default value and it needs to be updated in real time.

- Computation in real time: In the off-line model, we calculate the RR interval at once and to smooth the signal we use one second sliding window with the window size of 10 minutes. For the online version of the algorithm, it is computationally expensive to run all of these at once. We propose to calculate these values with a smaller batch size in real time. It requires careful investigation to find the right batch size so that the model has

sufficient amount of data to calculate RR intervals and do the smoothing operation. Also one has to make sure, it is not computationally expensive. Another option is to send all of the data to the cloud for processing and get back the result. But it suffers from three problems. First, it may be possible that, the user may be out of reach of network. The consequence is, the just-in-time intervention for cocaine intake can't be triggered at the right time. Second, users may not allow to upload their physiological data due to privacy and security concerns. Third, it adds extra cost of internet usage as the system tries to upload as much as 100 megabyte of data per day if they use the system for 14 hours.

Practical real-time deployment also requires research on developing (1) methods that enable context based Just in Time Interventions and self-report prompts and (2) sensors that are cause less burden for the users. Our model detects the cocaine intake event after the body recovers from the cocaine intake i.e. when the heart rate goes to base line. Which means, the cocaine intake event is detected after 30 to 40 minutes from when the cocaine is taken. Early detection may be required to trigger JITI or self reports near the vicinity of cocaine intake. In chapter 4 we show that, if we use smaller windows false positive rate rises. However, from the field study we find that people intake drug in high dosage amounts. This indicates that it might be possible to detect such events using smaller segments as the effect of very high dosage is more acute.

In this work, We use AutoSense Chest Band sensor to collect data both in the controlled and field environments. To get a good quality of ECG data, two ECG electrodes are required to be attached to the body. It requires an effort on the part of the participant to attach it in such a way so that it provides good quality of data. In the NIDA field study, user wore this sensor suite for 4 weeks. Wrist watch based ECG sensor is in under research now. Many commercial devices

such as Apple Watch, Microsoft Band etc. provides coarse grain heart rate measurements too. The MD2K team is also working on developing a new wrist sensor which will provide raw ECG signal. Wrist based sensor can detect the activity pattern and the hand movement of the user which can be used to further improve the cocaine detection model.

Chapter 6

Conclusion

Mobile health can help people improve their health by continuous monitoring of health and behavior on their mobile phone and by delivering timely interventions to motivate healthy lifestyle and abstinence from risky behaviors. Automated detection of various health states and behaviors from sensor data is key to realizing the mobile health vision. Such capabilities can then be used in scientific user studies to identify precipitating contexts that precede undesirable health states and risky behaviors. Automated detection of these predictors can be used to trigger the delivery of timely intervention on the mobile phone as preventative measures. Our work contributes to realization of this mobile health vision by showing that automated detection of drug use is feasible, opening the doors for development of just-in-time interventions.

The ANS model for detecting drug use can itself be improved for both accuracy and generality. First, we use a person-independent estimate of drug metabolism rate. This can be made more person-specific by incorporating demographic information such age, gender, weight, body mass index, etc. that are known to affect the metabolic rate. Second, the model can be improved further by using measurements from other physiological sensors such as respiration, galvanic skin response, electrodermal activity, etc. Third, the model can be expanded to include other illicit drugs. Psychostimulant drugs that may have similar response as Cocaine include Amphetamine and Methamphetamine, Methylphenidate (Ritalin), Methcathinone (an emerging drug under the name of Bath Salt) and MDMA (Ecstasy).

In addition to stimulating further research on mobile health, this work motivates three new sensor data processing issues for future research.

- Generalizable approaches are needed for detecting events of interest from

53

physiological time series data in the presence of numerous confounders encountered in the natural environment. Our work shows the feasibility of doing so for cocaine use that has a pronounced impact on physiology. But, new research is needed to develop explainable (i.e., non black box) methods for other events (e.g., stress, conversation, smoking, eating, etc.) that may not have as pronounced of an impact on physiology.

- Generalizable methods are needed to smooth noisy physiological time series data to see broad trends and to mark the boundary of events of interest in the time series. Our work illustrates an approach to identify the windows of interest in ECG time series for drug use response identification. But, further research is needed to investigate its applicability to other sensing modalities such as respiration, EEG, etc.

- We provide a generative model that tracks the recovery portion of the heart rate response, when physical activity concludes. Such models are succinct (need estimation of only one parameter) and can be used to generate synthetic data for simulation. Again, more generalized generative models are needed that can model the activation portion of the response curve as well so that the entire response curve can be synthetically generated. In addition to generating synthetic data such models can help build more robust detectors by modeling the entire response curves.

REFERENCES

[1] S. Kumar, W. Nilson, M. Pavel, and M. Srivastava, "Mobile health: Revolutionizing healthcare through trans-disciplinary research," *IEEE Computer*, vol. 46, no. 1, pp. 28–35, 2013.

[2] E. Ertin, N. Stohs, S. Kumar, A. Raij, M. al'Absi, T. Kwon, S. Mitra, S. Shah, and J. W. Jeong, "Autosense: Unobtrusively wearable sensor suite for inferencing of onset, causality, and consequences of stress in the field," in *ACM SenSys*, 2011, pp. 274–287.

[3] Z. BioHarness, "Zephyr Bioharness," http://www.zephyr-technology.com/bioharness-bt, Accessed: October 2013.

[4] K. Plarre, A. Raij, S. M. Hossain, A. A. Ali, M. Nakajima, M. al'Absi, T. Kamarck, S. Kumar, M. Scott, D. Siewiorek, A. Smailagic, and L. E. Wittmers, "Continuous inference of psychological stress from sensory measurements collected in the natural environment," in *ACM IPSN*, 2011, pp. 97–108.

[5] A. Ali, S. Hossain, K. Hovsepian, M. Rahman, K. Plarre, and S. Kumar, "mpuff: automated detection of cigarette smoking puffs from respiration measurements," in *ACM IPSN*, 2012, pp. 269–280.

[6] A. Natarajan, A. Parate, E. Gaiser, G. Angarita, R. Malison, B. Marlin, and D. Ganesan, "Detecting cocaine use with wearable electrocardiogram sensors," in *ACM UbiComp*, 2013, pp. 123–132.

[7] C. Furr-Holden, K. Campbell, A. Milam, M. Smart, N. Ialongo, and P. Leaf, "Metric properties of the neighborhood inventory for environmental typology (nifety): An environmental assessment tool for measuring indicators of violence, alcohol, tobacco, and other drug exposures," *Evaluation review*, vol. 34, no. 3, pp. 159–184, 2010.

[8] M. Fischman, C. Schuster, L. Resnekov, J. Shick, N. Krasnegor, W. Fennell, and D. Freedman, "Cardiovascular and subjective effects of intravenous cocaine administration in humans," *Archives of General Psychiatry*, vol. 33, no. 8, pp. 983–989, 1976.

[9] C. Muntaner, K. Kumor, C. Nagoshi, and J. Jaffe, "Intravenous cocaine infusions in humans: Dose responsivity and correlations of cardiovascular vs. subjective effects," *Pharmacology Biochemistry and Behavior*, vol. 34, no. 4, pp. 697–703, 1989.

[10] L. Degenhardt and W. Hall, "Extent of illicit drug use and dependence, and their contribution to the global burden of disease," *The Lancet*, vol. 379, no. 9810, pp. 55–70, 2012.

[11] S. Shiffman, J. A. Paty, M. Gnys, J. A. 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, p. 366, 1996.

[12] K. Johnson, A. Isham, D. Shah, and D. Gustafson, "Potential roles for new communication technologies in treatment of addiction," *Current psychiatry reports*, vol. 13, no. 5, pp. 390–397, 2011.

[13] E. Boyer, R. Fletcher, R. Fay, D. Smelson, D. Ziedonis, and R. Picard, "Preliminary efforts directed toward the detection of craving of illicit substances: the iheal project," *Journal of Medical Toxicology*, vol. 8, no. 1, pp. 1–5, 2012.

[14] E. Boyer, D. Smelson, R. Fletcher, D. Ziedonis, and R. Picard, "Wireless technologies, ubiquitous computing and mobile health: Application to drug abuse treatment and compliance with hiv therapies," *Journal of Medical Toxicology*, vol. 6, no. 2, pp. 212–216, 2010.

[15] C. R. Cole, E. H. Blackstone, F. J. Pashkow, C. E. Snader, and M. S. Lauer, "Heart-rate recovery immediately after exercise as a predictor of mortality," *New England Journal of Medicine*, vol. 341, no. 18, pp. 1351–1357, 1999.

[16] L. Cahalin, R. Arena, V. Labate, F. Bandera, C. Lavie, and M. Guazzi, "Heart rate recovery after the 6 min walk test rather than distance ambulated is a powerful prognostic indicator in heart failure with reduced and preserved ejection fraction: a comparison with cardiopulmonary exercise testing," *European journal of heart failure*, vol. 15, no. 5, pp. 519–527, 2013.

[17] R. Kloner, S. Hale, K. Alker, and S. Rezkalla, "The effects of acute and chronic cocaine use on the heart," *Circulation*, vol. 85, no. 2, pp. 407–419, 1992.

[18] S. Reed, M. Haney, S. Evans, N. Vadhan, E. Rubin, and R. Foltin, "Cardiovascular and subjective effects of repeated smoked cocaine administration in experienced cocaine users," *Drug and alcohol dependence*, vol. 102, no. 1, pp. 102–107, 2009.

[19] E. Cone, "Pharmacokinetics and pharmacodynamics of cocaine," *Journal of Analytical Toxicology*, vol. 19, no. 6, pp. 459–478, 1995.

[20] R. Foltin, A. Ward, M. Haney, C. Hart, and E. Collins, "The effects of escalating doses of smoked cocaine in humans," *Drug and alcohol dependence*, vol. 70, no. 2, pp. 149–157, 2003.

[21] M. Egred and G. Davis, "Cocaine and the heart," *Postgraduate medical journal*, vol. 81, no. 959, pp. 568–571, 2005.

[22] M. Perez-Reyes, S. Di Guiseppi, G. Ondrusek, A. Jeffcoat, and C. Cook, "Free-base cocaine smoking," *Clinical Pharmacology & Therapeutics*, vol. 32, no. 4, pp. 459–465, 1982.

[23] D. Noe and K. Kumor, "A pharmacokinetic–pharmacodynamic model of heart rate during cocaine administration in humans," *Clinical Pharmacology & Therapeutics*, vol. 49, no. 4, pp. 426–432, 1991.

[24] "ANT Radio," http://www.thisisant.com/, Accessed: January 2016.

[25] K. Plarre, A. Raij, S. Guha, M. al'Absi, E. Ertin, and S. Kumar, "Automated detection of sensor detachments for physiological sensing in the wild," in *Wireless Health 2010*. ACM, 2010, pp. 216–217.

[26] J. Pan and W. Tompkins, "A real-time qrs detection algorithm," *Biomedical Engineering, IEEE Transactions on*, vol. 32, no. 3, pp. 230–236, 1985.

[27] G. Berntson, K. Quigley, J. Jang, and S. Boysen, "An approach to artifact identification: application to heart period data," *Psychophysiology*, vol. 27, no. 5, pp. 586–598, 1990.

[28] S. Reddy, J. Burke, D. Estrin, M. Hansen, and M. Srivastava, "Determining transportation mode on mobile phones," in *Wearable Computers, 2008. ISWC 2008. 12th IEEE International Symposium on*, 2008, pp. 25–28.

[29] L. Atallah, B. Lo, R. King, and G. Yang, "Sensor placement for activity detection using wearable accelerometers," in *Body Sensor Networks (BSN)*, 2010, pp. 24–29.

[30] J. Kwapisz, G. Weiss, and S. Moore, "Activity recognition using cell phone accelerometers," *ACM KDD*, vol. 12, no. 2, pp. 74–82, 2011.

[31] J. Murphy, *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. Penguin. com, 1999.

[32] J. Freeman, F. Dewey, D. Hadley, J. Myers, and V. Froelicher, "Autonomic nervous system interaction with the cardiovascular system during exercise," *Progress in cardiovascular diseases*, vol. 48, no. 5, pp. 342–362, 2006.

[33] P. J. Huber *et al.*, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

[34] H. Sarker, M. Tibursky, M. M. Rahman, K. Hovsepian, M. Sharmin, D. H. Epstein, K. Preston, C. D. Furr-Holden, A. Milam, I. Nahum-Shani, M. al'Absi, and S. Kumar, "Finding significant stress episodes in a discontinuous time series of rapidly varying mobile sensor data," in *ACM CHI*, 2016.