A study of time-based features and regularity of manipulation to improve the detection of eating activity periods during free living

---

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Computer Engineering

---

by
Jose Luis Reyes Concha
May 2014

Accepted by:
Dr. Adam W. Hoover, Chair
Dr. Eric R. Muth
Dr. Richard E. Groff

Abstract

This thesis considers the problem of detecting when people eat by tracking their wrist motion. The goal of this work is to automatically detect the start and stop times of these eating activities. It builds upon previous work done by our research group [12] which developed an algorithm for automatically detecting peaks in activities associated with food preparation and clean-up. This peak detector is then used for segmenting data. These segments are then classified as eating or non-eating activities using a naive Bayes classifier based on probabilities obtained from computing different features in each segment. The original work introduced 4 features, all of them based on sensor readings. In this thesis we introduce a set of 3 new features to improve the detection of eating and non-eating activity periods: regularity of manipulation, time since last eating activity and cumulative eating time. We discuss the main concepts behind them, introducing the idea of time-based features. We then we test our new features under the framework developed by our group. Detection including regularity of manipulation, in combination with the original 4 features, achieved an overall accuracy of 79%. The accuracy obtained including time since last eating activity reached 69% and cumulative eating time, an overall accuracy of 64%. Finally, we compare these results to the original work and later discuss and characterize results based on our findings.

Acknowledgments

I would like to thank my advisor Dr. Hoover, for his guidance and advice throughout my Master's program. It has been sincerely, a great pleasure to work with you. My sincere thanks to my commitee members, Dr. Muth and Dr. Groff for their time and effort spent reviewing this thesis. Also, I would like to thank my family for their unconditional support, especially my parents Luz Amelia and Jose Fernando, who have been instrumental in my decision of pursuing my Master's degree.

Table of Contents

Title Page . . . . . **i**

Abstract . . . . . **ii**

Acknowledgments . . . . . **iii**

List of Tables . . . . . **v**

List of Figures . . . . . **vi**

1 Introduction **1**
1.1 Obesity problem 2
1.2 Tools for monitoring energy intake 3
1.3 Previous work of our group 4
1.4 Novelty 5

2 Research Design and Methods **6**
2.1 Overview of algorithm 6
2.2 Data collection 11
2.3 Software tools 12
2.4 New features 13
2.4.1 Regularity of manipulation 13
2.4.2 Time since last eating activity 16
2.4.3 Cumulative eating time 18
2.4.4 Other Features 20
2.5 Evaluation metrics 22

3 Results **24**
3.1 Previous results 24
3.2 Regularity of manipulation 25
3.3 Time since last eating activity 31
3.4 Cumulative eating time 33

4 Conclusions and Future Work **36**

Bibliography **38**

List of Tables

2.1 Average feature values found in training in [12] 11
2.2 FFT Statistics for data sets 1 and 2. 15
2.3 Statistical factors for ranges in Table 2.2. 16
2.4 Feature set. 22

3.1 Results for original 4 features with both sets combined. 24
3.2 Statistics for regularity of manipulation. 25
3.3 Results for original 4 features and regularity of manipulation. 26
3.4 Results for original 4 features and time since last EA. 31
3.5 Results for original 4 features and cumulative eating time. 33

List of Figures

2.1 Raw data for accelerometer and gyroscope sensors over a 2 minute period. 7
2.2 Smoothed data for accelerometer and gyroscope sensors for data in Figure 2.1. . . . 8
2.3 Sum of acceleration for an entire day. Arrows indicate start and end of meals/snacks. 8
2.4 Segmentation for Figure 2.3. Arrows indicate peaks, segments are considered as full periods between peaks. 9
2.5 Software tool developed, used for displaying feature statistics. 12
2.6 Manipulation in an eating activity segment with 7590 samples. 14
2.7 Manipulation in a non-eating activity segment with 4264 samples. 14
2.8 Distributions for regularity of manipulation between $0.0\overline{3}$Hz and 0.1Hz (set 1). . . . 17
2.9 Distributions for regularity of manipulation between $0.0\overline{3}$Hz and 0.1Hz (set 2). . . . 17
2.10 Cumulative measurements for each minute since last EA from data set 1 (averages). 19
2.11 Cumulative measurements for each minute since last EA from data set 2 (averages). 19
2.12 Cumulative eating time averages (set 1). 21
2.13 Cumulative eating time averages (set 2). 21

3.1 Comparison between original 4 features (top) and the combination of them including regularity of manipulation (bottom). Purple horizontal bars indicate EA detection. . 27
3.2 Smoothed manipulation data for an EA segment on the right tail of the distribution in Figure 2.9 of over 7400 samples. 28
3.3 Smoothed manipulation data for a non-EA segment on the left tail of the distribution in Figure 2.9 of 4500 samples. 28
3.4 Smoothed manipulation data for an EA segment that overlaps with the non-EA distribution in Figure 2.9 of roughly 6200 samples. 29
3.5 Smoothed manipulation data for an non-EA segment that overlaps with the EA distribution in Figure 2.9 of 4500 samples. 29
3.6 Original smoothed data for a segment in Figure 3.4. Arrows emphasize vigorous motions. 30
3.7 Original smoothed data for a segment in Figure 3.5. Arrows emphasize vigorous motions. 30
3.8 Classification period for original 4 features over a period of 8 and a half hours. . . . 32
3.9 Classification period for 5 features including time since last EA over a period of 8 and a half hours. Note the many fewer false positives compared to Figure 3.8. 32
3.10 Comparison between classification outputs for original 4 features (top) and 5 features including time since last EA (bottom) over a 30 minute period. The false positive just prior to the actual meal makes the feature of time since last EA a strong inhibitor for immediately subsequent data, causing the actual meal to be mis-classified. 33
3.11 Classification period for original 4 features over a period of 6 and a half hours. . . . 34
3.12 Classification period for 5 features including cumulative eating time over a period of 6 and a half hours. The true positive at the beginning of the segment causes the classifier to suppress immediately subsequent EA detections. 34

3.13 Comparison between classification outputs for original 4 features (top) and 5 features including cumulative eating time (bottom) over a 1 hour period. The false positive prior to the actual meal makes the feature cumulative eating time a strong inhibitor for immediately subsequent data, causing the actual meal to be mis-classified. 35

Chapter 1

Introduction

This thesis considers the problem of detecting when people eat by tracking their wrist motion. People generally eat between 3 to 5 meals and snacks per day [7, 18]. The goal of this work is to automatically detect the start and stop times of these eating activities. To do this, we are investigating methods that use a wrist-worn device that resembles a watch. The device contains accelerometers and gyroscopes that can continuously track the motion of the wrist. We are developing methods to analyze these signals with the goal of detecting the start and stop times of eating activities.

The motivating health concern is obesity. In the United States, more than one-third are obese [8], and obesity-related health issues are an increasing concern. Determining when people eat, and how much they eat, is crucial for studying behaviors related to obesity. The tool we envision could be used to monitor people to develop a calendar-like record of their eating activities. This data could be used to search for problematic eating patterns such as binge eating disorder, night eating syndrome, and day-to-day disparities such as overeating on weekends [18]. The tool could also be used to prompt people to complete a written record of what was eaten each day, commonly called a food diary.

Current methods for monitoring eating activities include food diaries, 24-hour recalls, and food frequency questionnaires. These tools rely heavily upon memory and manual effort. Because they are manual, they are typically only used for short periods of time ranging from 1-7 days. They also suffer from underreporting and underestimation bias, meaning that they provide incomplete records of eating activities.

The advantage of the tool we are developing is that it automates the process. The work considered in this thesis covers only part of the envisioned tool, specifically the detection of when people are eating. Other work has investigated automated measurements of how much someone eats during an eating activity [1, 10, 19]. Studies have shown that this method provides a measure that correlates with kilocalories consumed [31], and that people prefer using it over a 24-hour recall [31]. However, one limitation of those works was that the device had to be turned on and off by the wearer for every eating activity. Previous work in our group investigated the automatic detection of eating activities, achieving an accuracy of 81% [12]. In this thesis, we explore 3 new features to try to improve the detection accuracy.

1.1 Obesity problem

Obesity is common, serious and costly. In the United States, 34% of the population are obese [8]. Obesity and overweight are the fifth leading risk for global deaths [38]. Worldwide obesity has nearly doubled since 1980, and each year at least 2.8 million adults die as a result of being overweight or obese [38]. Overweight and obesity are linked to more deaths worldwide than underweight [38]. For example, 65% of the world's population live in countries where overweight and obesity kill more people than underweight [38]. Obesity-related conditions include heart disease, stroke, type 2 diabetes and certain types of cancer, some of the leading causes of preventable death [8]. Obesity is costly, the estimated annual medical cost of obesity in the U.S. was $147 billion in 2008 U.S. dollars; the annual medical costs for people who are obese were $1,429 higher than those of normal weight [8].

Obesity treatment commonly targets behavior changes including dietary and exercise activities [20, 25]. Severe cases may be treated with weight-loss medication and weight-loss surgery [20, 25]. According to the World Health Organization, the most important aspect at the individual level, is to limit energy intake [38]. This idea is the motivation of various research topics regarding food and energy intake [9, 11, 13, 19]. Obesity and overweight present a real problem in today's society, and combined with all the facts just described set the base of our group's motivation for this research.

1.2 Tools for monitoring energy intake

Recording the intake of a free-living human is a difficult task. The measurement of intake typically refers to the amount of energy (e.g. kilocalories), but can also include the types of foods ingested, the times of day, the duration of a meal, and the social context. The three most widely used tools are eating diary, food frequency questionnaire, and 24-hour dietary recall [35]. An eating diary generally involves manually logging food ingestion, a 24-hour dietary recall relies on subjects recalling everything they ate and drank during the previous day, and a food frequency questionnaire (FFQ) generally means participants reporting the frequency of consumption and portion size over a defined period of time. Numerous studies [9, 14, 15, 17, 24, 37] have shown that people underreport their food consumption when using one of these methods, typically by 10 to 40% per day. Also, any of these methods requires subjects to invest a significant conscious effort to continuously pay attention to the self-monitoring task. The significance of our eating activity monitor is that it automates the recording of intake, greatly decreasing user burden and subjectivity.

Several experts in the field of dietetics have indicated the need for advance in the tools used for self-monitoring [21, 36]. Two main approaches have been researched: using a camera, and using wearable sensors. In the camera-based approach, a camera takes pictures of foods before and after eating. The amount consumed is evaluated by a trainer observer who compares the images to a database with pictures of portion-varying of the same type of food. A comparable accuracy has been shown between this approach and both weighed and direct visual estimation when pictures were taken by trainer observer in a cafeteria setting. Its accuracy is also comparable to when subjects took the pictures themselves [19]. Although this approach has shown to reduce user burden, it takes approximately 30 minutes per day for the trained researcher to analyze an estimate energy intake from the images. The idea of automated image processing instead of post-image analysis by a reviewer to determine the amount of food ingested has also been suggested [27, 34, 40], and has been demonstrated for small amounts of food that are carefully separated and with a carefully controlled background [39]. Studies have shown that the camera-based approach is preferred over the traditional pen and paper method to record food consumption [6, 33], showing the promise of automated technology-based advances. However, the experimenter burden for the camera method is still very high, and the method is not scalable to potentially millions of users.

The idea behind the wearable sensors approach is that sensors are worn on various parts

of the body to track limb motion and estimate sound and muscle activity (e.g. around the neck), in order to determine when a person is eating [3, 28]. Limb motion patterns associated with eating have been detected using five microelectromechanical devices systems (MEMS) sensors, one on the lower and upper segment of each arm, and one on the back [16]. Ear and neck mounted sensors have been used to detect chewing sounds and swallowing motions [2, 3]. Neck-worn sensors have been used to detect swallowing activities [32]. Swallowing motion has been studied using microphones, as well. Microphones have been used in various locations such as the laryngopharynx and mastoid bone portions of the throat, inside the ear, and ambient outward-directed on the throat [30]. Analysts have shown high reliability for identifying bites, chews and swallows using the signals recorded from these sensors. High reliability was shown in other studies where trained analysts reviewed signal recordings for detecting periods of eating activity between periods of quiet rest, talking, and meals of various portion sizes solely based on sensor data [29]. While all these approaches show advances for automating the self-monitoring task, none of them have been tested on subjects under normal free living conditions; all data was collected in brief controlled laboratory visits. Our group's approach uses wearable sensors mounted at the wrist that resembles a watch. This approach enables more freedom of motion and activities, and has a higher likelihood of social acceptance.

A number of researchers would find use in the calendar of eating activities provided by an eating activity (EA) monitor. For example, measuring energy intake (EI) and time spent eating per day, or per meal, or even at different times of the week could all help identify eating disorders [18]. The data obtained from an EA monitor could be used to compare different eating habits of different demographics [23], or even to study the effect of changes in the environment [22]. For studies requiring detailed nutrition information, the calendar could be used in combination with the already mentioned 24-hour recall or even a food diary to help remind participants when they ate, improving the accuracy of self-reported data. The EA monitor could also prompt the participant in real-time, at the completion of an eating activity, to fill out a food diary or complete other questionnaires.

1.3 Previous work of our group

This thesis builds upon an algorithm previously developed by our group [12]. The paper describes a new method that uses a watch-like configuration of sensors to continuously track wrist motion throughout the day and automatically detect periods of eating [12]. The study was carried

out on 43 subjects and evaluated on a large data set. This research included a complete framework for the in depth study of four different features which in combination with each other contribute to classify an activity as eating or non-eating. The features described in the original algorithm were: manipulation, regularity of wrist roll motion, linear acceleration and amount of wrist roll motion. These features are computed using data obtained from an accelerometer and gyroscope sensors included in a mobile phone device. The framework also includes a segmentation portion which divides the data into segments that are then classified as eating or non-eating based on feature probability using a Bayes classifier. Based on evaluation metrics also developed by our group, an overall accuracy of 81% was achieved [12].

1.4 Novelty

The previous work did not consider features related to time. Manipulation, regularity of wrist roll motion, linear acceleration and amount of wrist roll motion are features that rely only on sensor measurements. In this thesis we consider the hypothesis that time-of-day and time-since-last-meal can have an impact on when a person eats. This information could potentially be used in a classifier to improve the detection of periods of eating. We also consider a new feature related to manipulation. Manipulation is a measure of how much a person rotates their wrist relative to linear motion of the wrist. Previous work found that manipulation is generally higher during eating activities than other activities. However, manipulation does not happen continuously during eating. It tends to occur periodically as a person manipulates food and brings it to their mouth, followed by periods of rest during chewing and swallowing. In this thesis we develop a new feature that measures the periodicity of manipulation over time.

Chapter 2

Research Design and Methods

We first overview the method originally developed by our group [12]. We then describe the new features and software tools developed for this thesis. Finally, we describe the data collected to test the new features and the metrics used to evaluate performance.

2.1 Overview of algorithm

Figure 2.1 shows an example of a 2-minute period of raw accelerometer and gyroscope data recorded at 15Hz while tracking a person during daily free-living. It can be seen that the data is noisy and it is difficult visually to make out the continuity of the sensor readings. Displayed from top to bottom we see data from the accelerometer ($AccX$, $AccY$, $AccZ$) and gyroscope (Yaw, $Pitch$ and $Roll$) sensors obtained from a recording in our data set. A recording typically spanned 8-14 hours, and zoomed-out views will be used in many subsequent figures.

In order to reduce the effects of noise, the first step is to smooth the data. This is done using a Gaussian window as:

$$S_t = \sum_{i=-N}^{0} R_{t+i} \frac{exp(\frac{-t^2}{2R^2})}{\sum_{x=0}^{N} exp(\frac{-(x-N)^2}{2R^2})} \tag{2.1}$$

where R_t is the raw value of sensor reading for each signal and S_t is the smoothed datum at time t. Equation 2.1 uses a Gaussian-weighted window centered on the current measurement, so that only half of a Gaussian distribution is used for smoothing. N represents the window size and R

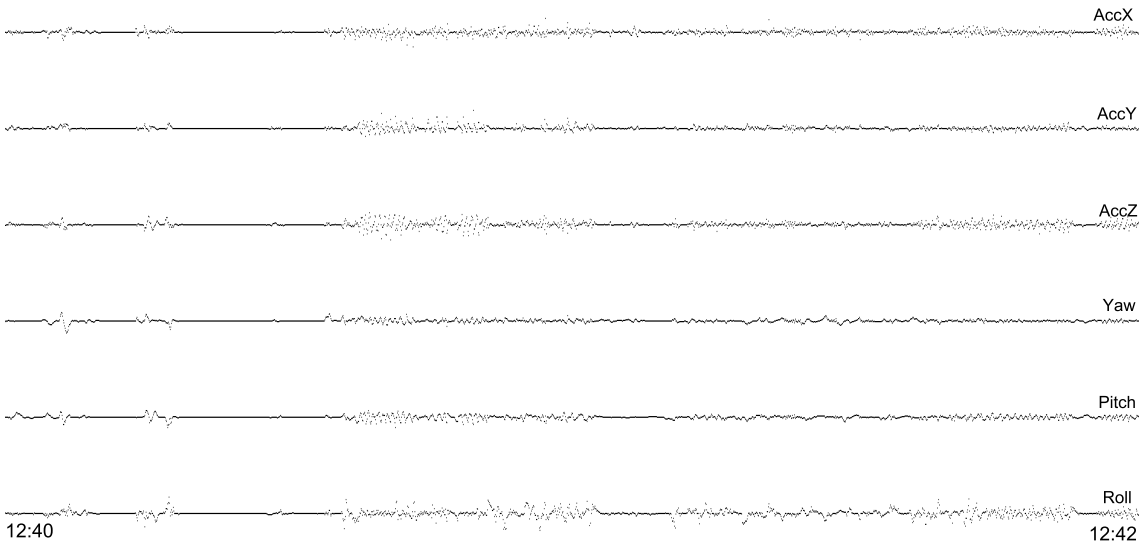

Figure 2.1: Raw data for accelerometer and gyroscope sensors over a 2 minute period.

is the sigma of the Gaussian. The Gaussian kernel is applied independently to the data from each accelerometer and gyroscope axis. Figure 2.2 shows the smoothed signal of each individual axis for the data shown in Figure 2.1.

After smoothing the data, the algorithm calculates a feature called sum of acceleration. Sum of acceleration can be computed as:

$$E_t = \frac{1}{W+1} \sum_{i=t-\frac{W}{2}}^{t+\frac{W}{2}} |S_{x,t}| + |S_{y,t}| + |S_{z,t}| \tag{2.2}$$

where $S_{x,t}$, $S_{y,t}$ and $S_{z,t}$ are the smoothed acceleration readings at time t. W is the window size that is slides across the data. Figure 2.3 shows an example of sum of acceleration for a whole day. The manually marked meal times are indicated by downward arrows. These came from the ground truth, obtained from manual logs kept by a person wearing a wrist motion tracker. It can be seen that the sum of acceleration tends to be high prior to periods of eating, and high again after periods of eating are done. During an eating period, it tends to be low. This is likely due to vigorous activities during the preparation and clean-up of food and materials before and after eating.

The next step is to segment the data. Segmentation is done by automatically detecting peaks by using a custom peak detector based on a hysteresis threshold algorithm. Details are discussed in [12]. This is done by using the feature we just described, sum of acceleration. Like we mentioned

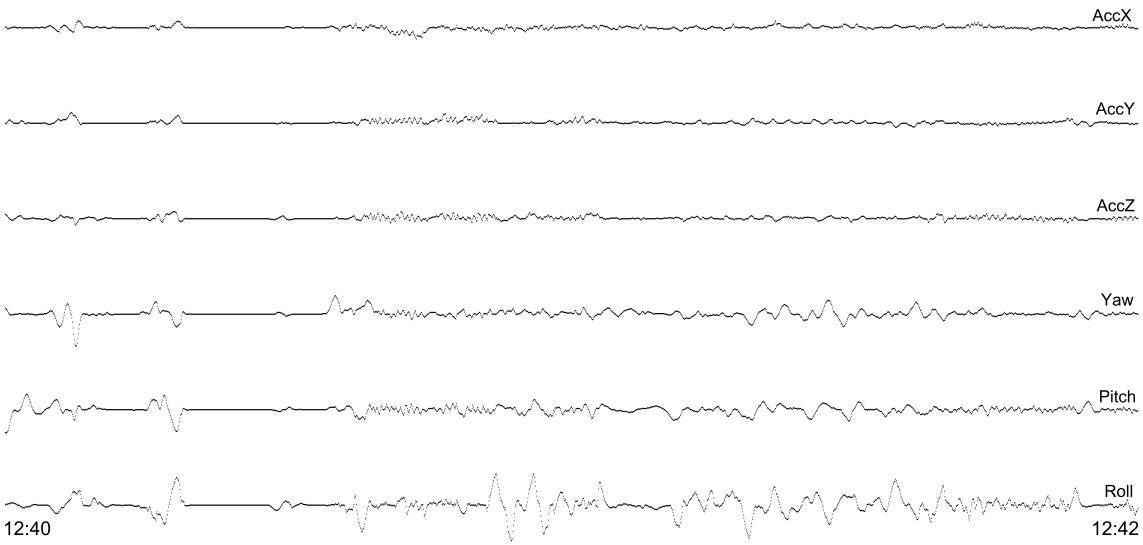

Figure 2.2: Smoothed data for accelerometer and gyroscope sensors for data in Figure 2.1.

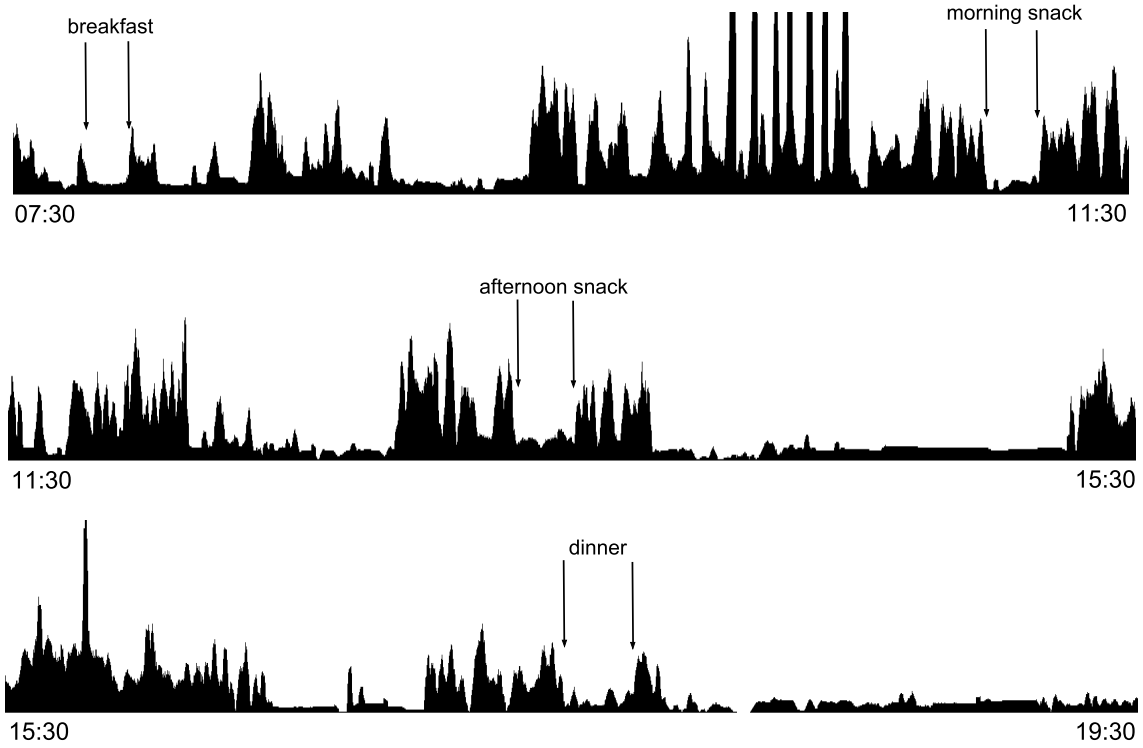

Figure 2.3: Sum of acceleration for an entire day. Arrows indicate start and end of meals/snacks.

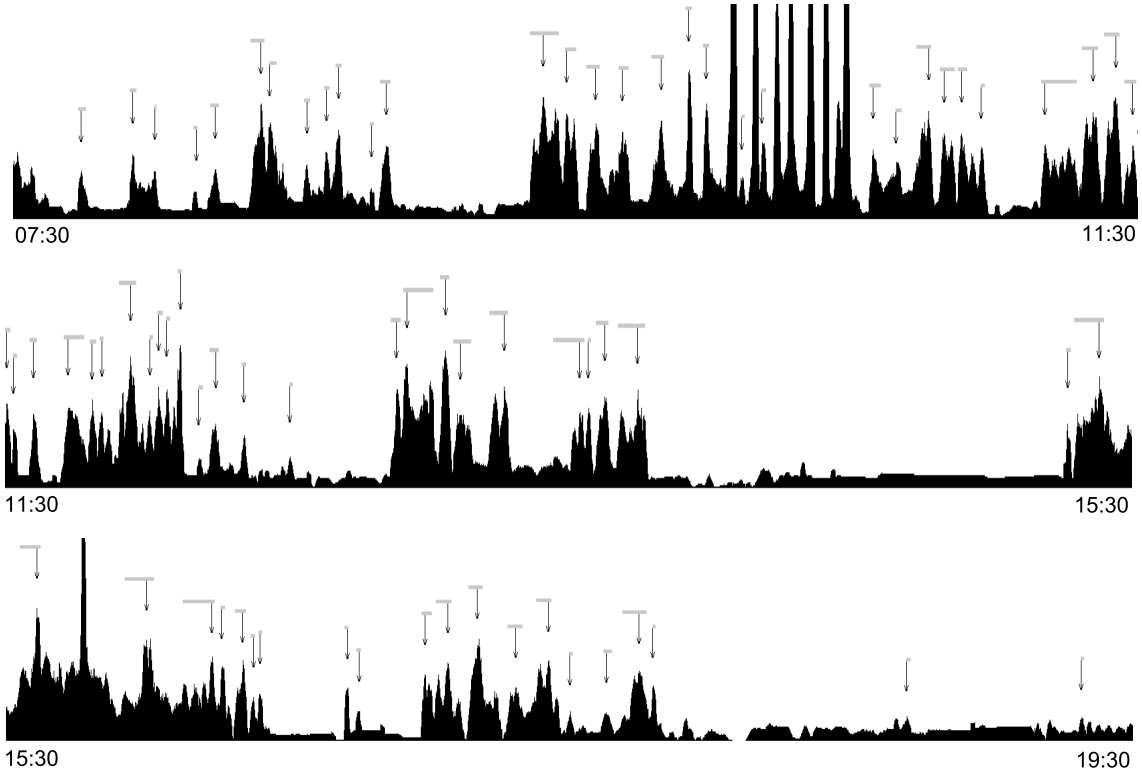

Figure 2.4: Segmentation for Figure 2.3. Arrows indicate peaks, segments are considered as full periods between peaks.

before, sum of acceleration tends to be high prior and after eating activities. This feature is not sufficient for classification but it can be used to derive a segmentation. The indices of the peaks are used to segment the data. Figure 2.4 shows an example of the peaks automatically detected for Figure 2.3. The arrows show the maximum times and the horizontal spans show the widths of the peaks. Segments are taken as all periods of time between the arrows.

The original method used 4 features to classify segments. The first feature is manipulation, and it is computed as:

$$f_{1,w} = \frac{1}{W} \sum^{W} \frac{|S_{\phi,t}| + |S_{\theta,t}| + |S_{\psi,t}|}{|S_{x,t}| + |S_{y,t}| + |S_{z,t}|} \tag{2.3}$$

where W is the span of the segment period, t is the index that iterates across the span, and S is the smoothed datum (ϕ = yaw, θ = pitch, ψ = roll). Manipulation measures the ratio of rotational motion to linear motion. Units for manipulation are $\frac{deg/s}{G}$, since units for the rotational motion are deg/s and G for linear motion. The second feature is linear acceleration. Linear acceleration is

calculated as:

$$f_{2,w} = \frac{1}{W}\sum^{W} |S_{x,t}| + |S_{y,t}| + |S_{z,t}| \tag{2.4}$$

The next feature is the amount of wrist roll motion, and it is computed as:

$$f_{3,w} = \frac{1}{W}\sum^{W} |S_{\phi,t} - \frac{1}{W}\sum^{W} S_{\phi,t}| \tag{2.5}$$

And lastly, regularity of wrist roll motion. Regularity of wrist roll motion is calculated as:

$$f_{4,t} = \frac{1}{W}\int_{W} 1\forall t \in [|S_{\phi,t}| > 10^0 \ldots t + 8sec] \tag{2.6}$$

This feature takes on a value between 0 and 1, and it represents the percentage of time that the wrist is in roll motion. The calculation includes the time the wrist roll is at least 10 deg/s, plus a period of 8 seconds after each occurrence of wrist roll motion falling below 10 deg/s [12].

These features are used in a naive Bayesian classifier for determining eating and non-eating activities based on the segmented data. This approach to classification is to assign the most probable class $c_i \in C$, given feature values $f_1, f_2, \ldots, f_N$. Using the naive assumption of independence of features, the classification problem can be written as:

$$c_i = \underset{C}{\text{argmax}} P(c_i) \prod_j P(f_i|c_i) \tag{2.7}$$

For our problem only 2 classes will be represented, eating and non-eating. Probabilities are initialized to $P(c_i) = 0.5$ and each feature contributes to the classifier by computing its normal distribution:

$$P(f_j|c_i) = \frac{1}{\sqrt{2\pi\sigma_{i,j}^2}} exp(-\frac{(f_i - \mu_{i,j})^2}{2\sigma_{i,j}^2}) \tag{2.8}$$

where σ represents the standard deviation and μ the mean of the distribution. So in order to classify an unknown segment, mean and standard deviations are needed. For our group's previous work, these statistics were found for each feature and are shown in Table 2.1 and were used in Equation 2.8.

Notation	Feature	Eating		Non-Eating	
		Mean	Std. Dev.	Mean	Std. Dev
f_1	Manipulation ((deg/s)/G)	791	214	395	239
f_2	Acceleration (G)	0.039	0.014	0.054	0.066
f_3	Roll motion (deg/s)	9.1	4.27	6.8	6.3
f_4	Roll regularity (%time)	0.58	0.14	0.37	0.26

Table 2.1: Average feature values found in training in [12]

2.2 Data collection

This work uses 2 data sets collected in previous work [12]. Data for the study was collected using an iPhone 4 [4]. This device was chosen because it is programmable, equipped with an accelerometer and gyroscope, and has enough memory (16GB) and battery (1420 mAh) to record data continuously for an entire day. The phone was inside a pouch and worn around the wrist, and it recorded data continuously throughout the day. Two different sets of recordings were obtained. Set 1 was obtained using a preliminary version of the software that was developed by our group to work with the iPhone which required users to manually log start and end of meals. A second set was obtained from a new version of the software were users could log start and ends of meals by pressing on a button on the application. A total of 43 recordings were obtained. Set 1 consists of 20 recordings and set 2 consists of 23 recordings for a total of 449 hours of data. The data was recorded at a sampling rate of 15Hz.

For training purposes, data is segmented into eating and non-eating activity segments. For the case of the non-eating activities, these segments are 5 minutes long unless the end of the record is reached or an eating activity is reached before 5 minutes. If data was missing or corrupted data was filled with zeroes for all axes if this period was under 5 seconds. If it was longer, data was discarded. Eating activity segments simply span for the duration of an entire meal/snack. Each of the EA and non-EA segments were used to compute each of the features previously discussed, in order to be analyzed and find characteristics that separate EAs from non-EAs. For set 1, the total number of meal segments is 35 and 1997 non-meal segments. Set 2 contains a total of 81 meal segments and 3201 non-meal segments.

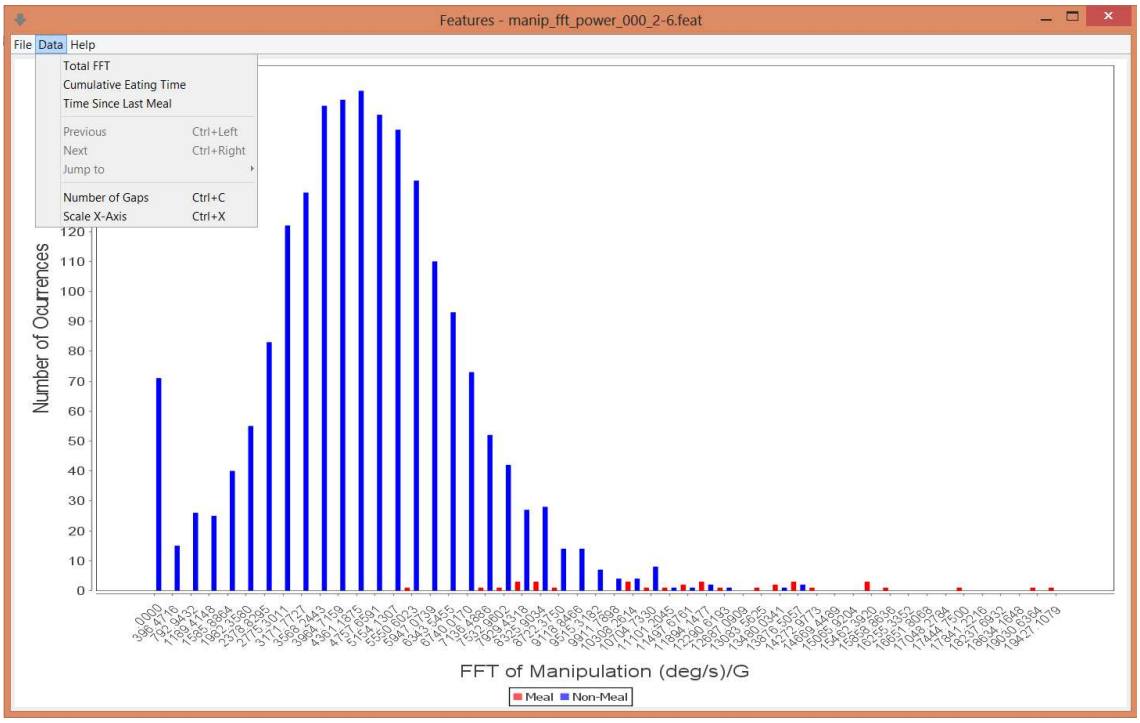

Figure 2.5: Software tool developed, used for displaying feature statistics.

2.3 Software tools

For the work in this thesis, a custom tool was designed in Java (version 7 [26]) to display and analyze different features. The variety of graphing and scientific libraries was the key to choose this language. A screenshot of the tool can be visualized in Figure 2.5. Another small routine was developed in C that computed a variety of different potential features. This software saves the output to a file which is later displayed using the Java tool that was just described. A third tool, called PhoneView, was developed previously by our group and was used in this thesis to display the data, check ground truths, visually go through segmented data and perform the classification process.

2.4 New features

2.4.1 Regularity of manipulation

Figure 2.6 shows manipulation over time in an eating activity segment. The segment spans 7590 samples (506 seconds at 15Hz). Figure 2.7 shows a non-eating segment of 4264 samples (about 284 seconds). Peaks in the EA segment appear to be more regular and about the same magnitude with values around $5000\frac{deg/s}{G}$, whereas for the non-EA they tend to be more irregular and of different sample intervals. This brings up the idea that the peaks in the EA could represent bites. Trying to capture this regularity of peaks is the main idea behind regularity of manipulation. It is not unusual for people to take bites between two to eight times a minute, so it seems reasonable to analyze this regularity based on this concept. Regularity of manipulation uses a Fast Fourier Transform (FFT) approach to do this. An FFT provides us with a way to find the regularity (periodicity) of bites between certain frequency ranges. It reveals periodicities in input data as well as the power of any periodic components. The FFT is an algorithm that computes a Discrete Fourier Transform (DFT), and it is widely used because it takes significantly less computations than a DFT. Basically, a DFT converts the sampled function from its original domain, in our case the time domain to the frequency domain. The Discrete Fourier Transform is computed as:

$$F_n = \sum_{k=0}^{N-1} f_k * e^{\frac{-2\pi i}{N} nk} \tag{2.9}$$

where N is the sequence of numbers, $k = 0, 1 \ldots N-1$, and f represents the input signal amplitude at time k. The FFT reduces the number of computations that are performed by the DFT. The classical FFT uses powers of 2 for input length, but our problem requires more flexibility on this matter, since the EA and non-EA segments are of variable length. The FFT algorithm used in this paper is called Bluestein's FFT algorithm [5], which provides the flexibility we need.

Looking at figures 2.6 and 2.7, it can be seen that the signals are noisy, so in order to take advantage of the FFT, data is heavily smoothed. We applied another Gaussian-weighted window, using Equation 2.1, with the kernel centered at the current measurement, but using a full Gaussian distribution. The kernel for this part uses a N size of 225, which correspond to 15 seconds and an R value of $37.\overline{6}$. Fifteen seconds seemed to be enough to smooth the data without losing information, and R is obtained from covering 99.7% of the area below the Gaussian (centered at the mean). In

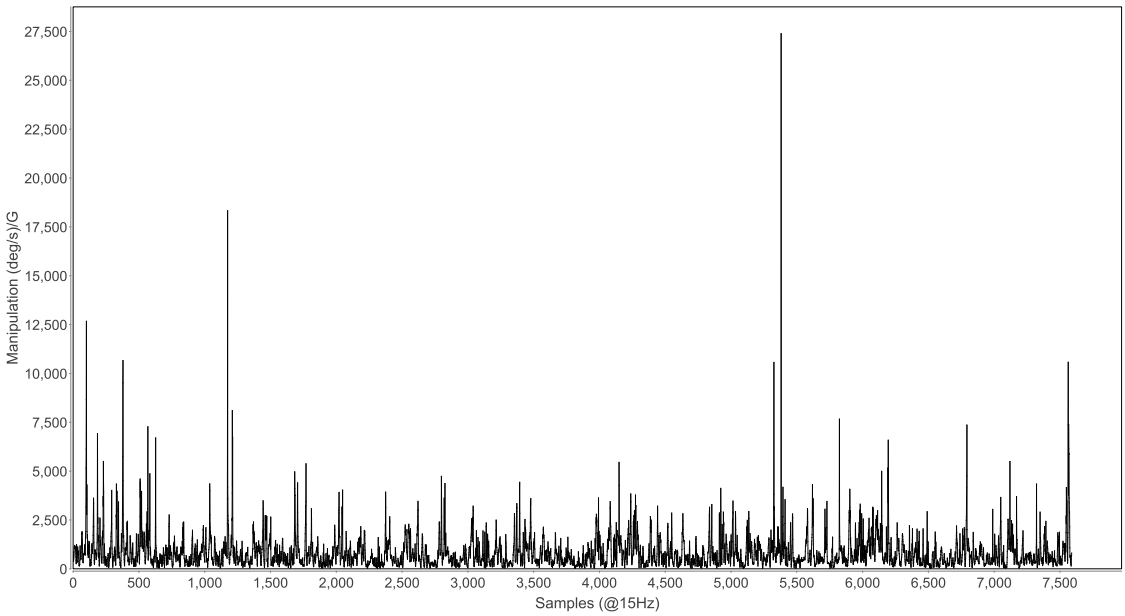

Figure 2.6: Manipulation in an eating activity segment with 7590 samples.

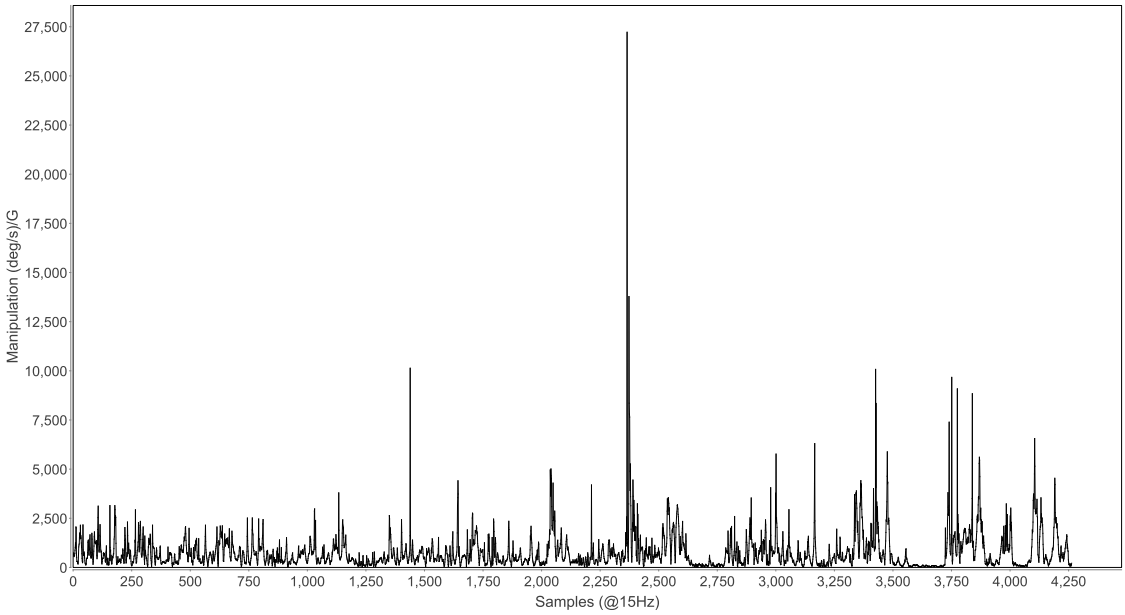

Figure 2.7: Manipulation in a non-eating activity segment with 4264 samples.

	Set 1				Set 2			
	Meal		Non-Meal		Meal		Non-Meal	
Range	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev
$\frac{2}{60}Hz$-$\frac{6}{60}Hz$	12090.3	3423.2	4656.7	2083.6	12709.9	5310.2	4425.1	2383.6
$\frac{4}{60}Hz$-$\frac{6}{60}Hz$	4009.1	1228.6	1681.8	802.36	4392.6	1804	1573.7	885.7
$\frac{4}{60}Hz$-$\frac{8}{60}Hz$	5708.9	1653.7	2363.5	1090.4	6052.2	2491.7	2222.8	1214.5
$\frac{6}{60}Hz$-$\frac{8}{60}Hz$	2676.3	832.94	1030.2	509.09	1659.6	733.47	649.16	370.73

Table 2.2: FFT Statistics for data sets 1 and 2.

order to do this we need to consider at least three standard deviations to each side of the distribution, for a total of 6 standard deviations.

Like mentioned before, we want to analyze a few different frequency boundaries between 2 times per minute and 8 times per minute. A few different range of frequencies between $\frac{2}{60}$Hz and $\frac{8}{60}$Hz were tested and are shown in Table 2.2. These statistics come from computing the FFT for all eating and non-eating segments, and calculating the sum of the magnitude between a range of frequencies:

$$f_{RoM} = \frac{1}{t}\sum_{i=LFB}^{HFB}\sqrt{r_k^2 + i_k^2} \tag{2.10}$$

where time is indicated by t, and the limits of the sum, LFB and HFB, are the lower and higher frequency boundaries, r and i stand for the real and complex parts of the FFT respectively, and k is index of the bin.

We want to pick the range that has eating and non-eating distributions overlapping with each other the least possible. We know from Table 2.2 that the mean for eating activity columns is higher than for non-eating. In other words we are trying to find the means to be as far apart from each other, and standard deviations to be relatively small. We need a measure of how well a given range of frequencies separates eating segments from non-eating segments. This factor is calculated for EAs as:

$$X_{EA} = -\frac{\mu_{nonEA} - \mu_{EA}}{\sigma_{EA}} \tag{2.11}$$

and for non-EAs:

$$X_{nonEA} = \frac{\mu_{EA} - \mu_{nonEA}}{\sigma_{nonEA}} \tag{2.12}$$

	Set 1		Set 2	
Range	X_{EA}	X_{nonEA}	X_{EA}	X_{nonEA}
$\frac{2}{60}Hz$-$\frac{6}{60}Hz$	2.1715	3.5676	1.5602	3.4757
$\frac{4}{60}Hz$-$\frac{6}{60}Hz$	1.8943	2.9006	1.5626	3.1827
$\frac{4}{60}Hz$-$\frac{8}{60}Hz$	2.023	3.068	1.5369	3.1531
$\frac{6}{60}Hz$-$\frac{8}{60}Hz$	1.9763	3.2335	1.3776	2.7255

Table 2.3: Statistical factors for ranges in Table 2.2.

where μ is the mean, and σ is the standard deviation. We want to pick the factors, denoted by X, to be as high as possible since that would mean that distributions are more isolated from each other. Results for these computations are shown in Table 2.3 and represent the output for each of the ranges from Table 2.2. Based on these results, we picked the range $[\frac{2}{60}Hz, \frac{6}{60}Hz]$ since it has higher factors overall, taking both data sets into account.

After the integral is computed, the output is normalized by total time in seconds, because eating activity segments are usually longer than the 5 minutes segments defined for non-eating periods. The distributions for the selected frequency ranges for both sets are shown in Figure 2.8 and Figure 2.9. In order to make the figures clearer the values for eating segments are magnified by a factor of 25. This is done because the ratio of non-eating segments to eating segments is over 50:1. Final units for regularity of manipulation are $\frac{deg/s^3}{G}$. In other words, regularity of manipulation measures manipulation between frequency bands.

Now that we have chosen values for our regularity of manipulation feature, we can compute probabilities that will contribute to the classification process by using normal distribution probabilities, for both EAs and non-EAs, using Equation 2.8.

2.4.2 Time since last eating activity

The idea behind the time since last EA feature is to explore time-based components. Basically, this approach implies that after a person has eaten, it is very unlikely for that person to eat again in the immediate future. As time passes, the probability of eating monotonically increases since a person needs to eat. Let t_{last} = end time of the last segment classified as EA. Let t = middle time of the unknown segment currently being classified. We then calculate time since last EA as:

$$f = t - t_{last} \tag{2.13}$$

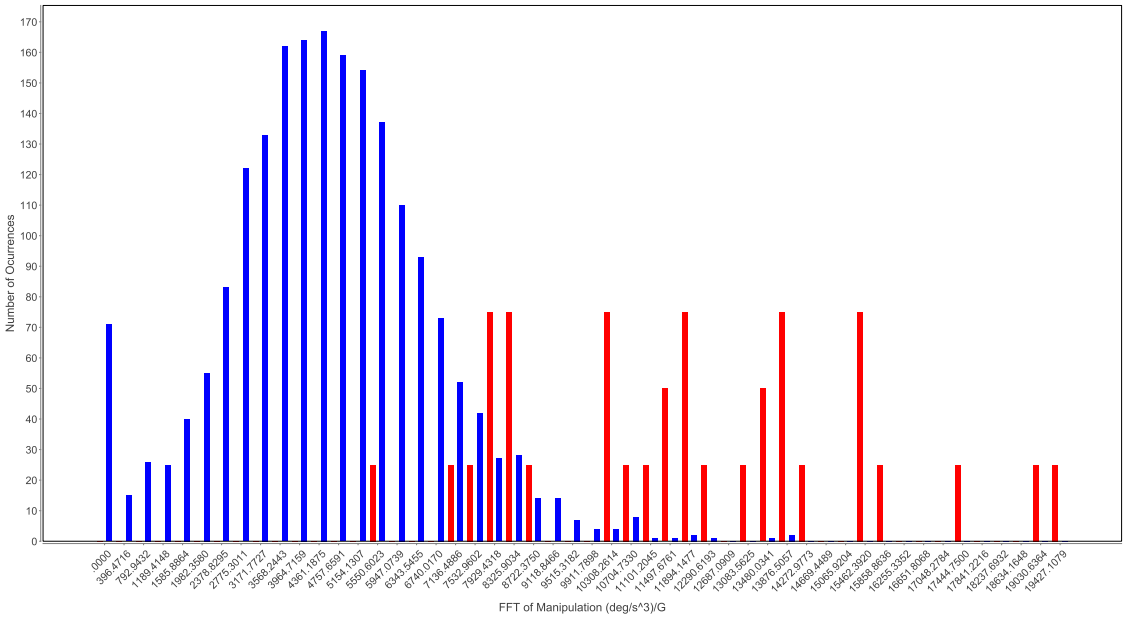

Figure 2.8: Distributions for regularity of manipulation between $0.0\overline{3}$Hz and 0.1Hz (set 1).

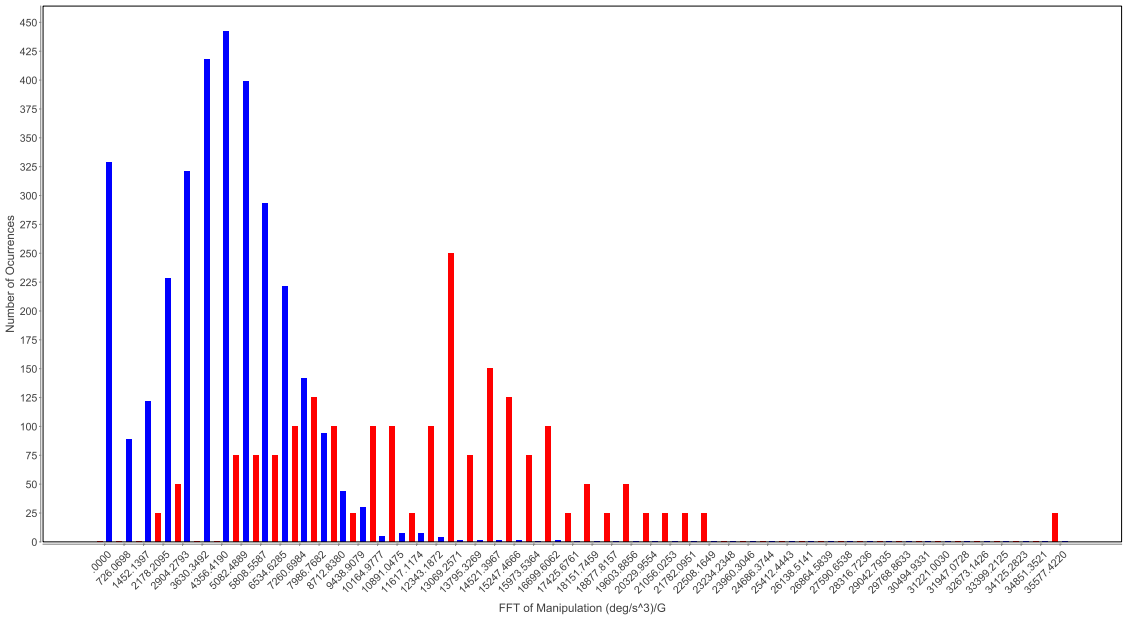

Figure 2.9: Distributions for regularity of manipulation between $0.0\overline{3}$Hz and 0.1Hz (set 2).

where f is the feature value. In other words, this feature represents the amount of time in minutes since the last meal. A Bayesian classifier requires probability distributions for each possible class (see equations 2.7-2.8). In the case of a time-based feature, it is possible to calculate a distribution of times between when 2 events occurred. For example, it is possible to calculate the average time between meals. But for the opposite class, the absence of a meal, it is nonsensical to calculate a distribution; i.e. what is the average time between all non-meal times. In other words, it does not make sense to compute time since last non-EA. We therefore instead calculate the probability that the segment is not a meal as one minus the probability that it is a meal.

To calculate the probability of a meal, we trained our classifier by calculating a cumulative distribution function (CDF) of time since last EA. The CDF takes the range 0 to 1. We trained the CDF using all available data. Figures 2.10 and 2.11 show the values found. Probabilities in Figure 2.10 start at 0 and end at 0.917, and probabilities in Figure 2.11 range from 0 to 0.988. Each point on the x-axis is an amount of time since the last EA in minutes, and each point on the y-axis was the frequency of having observed a meal in the training set that occurred that number of minutes prior to the current time. For an unknown segment, the probabilities in equation 2.8 are calculated as $p(f|EA)$ = CDF and $p(f|nonEA)$ = 1-CDF. If no EA has occurred prior to the current time, we then set our EA and non-EA probabilities to 0.5 since the feature is undefined.

2.4.3 Cumulative eating time

As its name suggests, cumulative eating time is also a time-based feature. Basically, it takes on the idea that a person spends a certain amount of time eating in a day. Studies have shown that on average a person spends around 1.1 hours a day eating and drinking [7]. Every person will have accumulated a certain amount of time in minutes, since they first started recording. Let χ be computed as:

$$\chi_t = \begin{cases} 1 & \text{if } t \in EA \\ 0 & \text{otherwise} \end{cases} \tag{2.14}$$

Then, at time t, cumulative eating time f, can be calculated as:

$$f_t = \sum_{i=0}^{t} \chi_i \tag{2.15}$$

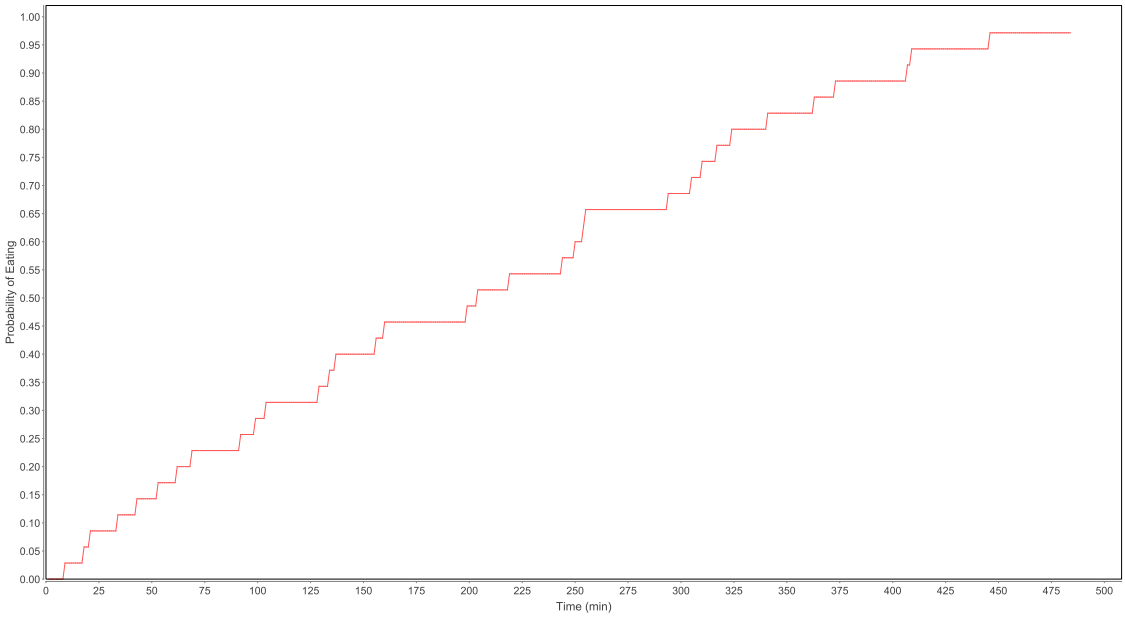

Figure 2.10: Cumulative measurements for each minute since last EA from data set 1 (averages).

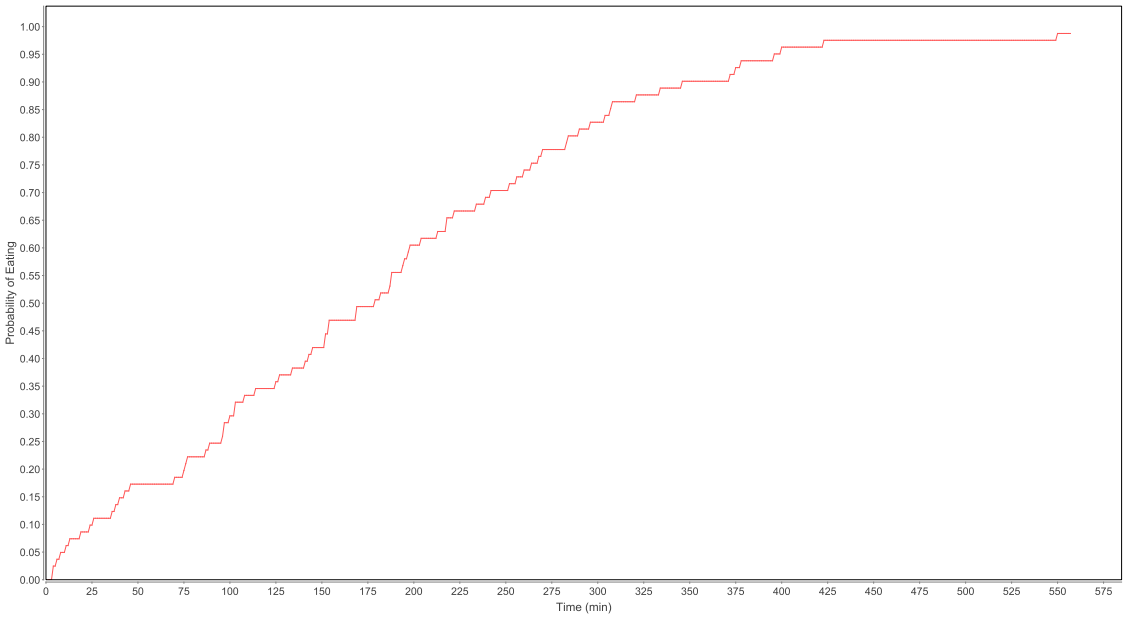

Figure 2.11: Cumulative measurements for each minute since last EA from data set 2 (averages).

where f is the value of the feature. For example, a person started breakfast 15 minutes into their recording and ate for 10 minutes. For every minute after that, that person has accumulated 10 minutes of eating time until they eat again, and then that amount of time will of course, increase. In other words, cumulative eating time, simply computes total time spent eating and drinking in a single day.

As with the time-based feature described in the last section, this one also has the problem that it is nonsensical to define the distribution of times involving non-events. We therefore calculate the probability that the segment is not a meal as one minus the probability that it is a meal. To calculate the probability of a meal, we trained our classifier by calculating a cumulative distribution function (CDF) of time eating for the day. The range of the CDF is 0 to the maximum amount of minutes monitored. We trained the CDF using all available data. Figures 2.12 and Figure 2.13 show the values found for each data set. Each point on the x-axis is an amount of time since the recording started in minutes, and each point on the y-axis was the total minutes a person had spent eating in the training set up to that time. The CDF was calculated independently for each person and then the average and standard deviation at each total time was calculated. Figures 2.12 and Figure 2.13 show the average as the middle line and the +-1 standard deviation as the top and bottom lines.

For an unknown segment, the probabilities in equation 2.8 are calculated by looking up values for μ and σ in the CDF shown in figures 2.12 and 2.13. The x-axis on the CDF is found as f (equation 2.15) considering the unknown segment as an EA, and μ_{cdf} and σ_{cdf} are the values on the y-axis. Using those values, $p(f|EA)$ is given by Equation 2.8 and $p(f|nonEA) = 1$ - $p(f|EA)$.

2.4.4 Other Features

Some other features were analyzed and are shown in Table 2.4. Each feature in this feature set was applied individually to each axis of the accelerometer and gyroscope signals, except for the case of correlation which uses a combination of 2 inputs at a time. A total of 63 other features were tested and analyzed, but after initial testing were not pursued further.

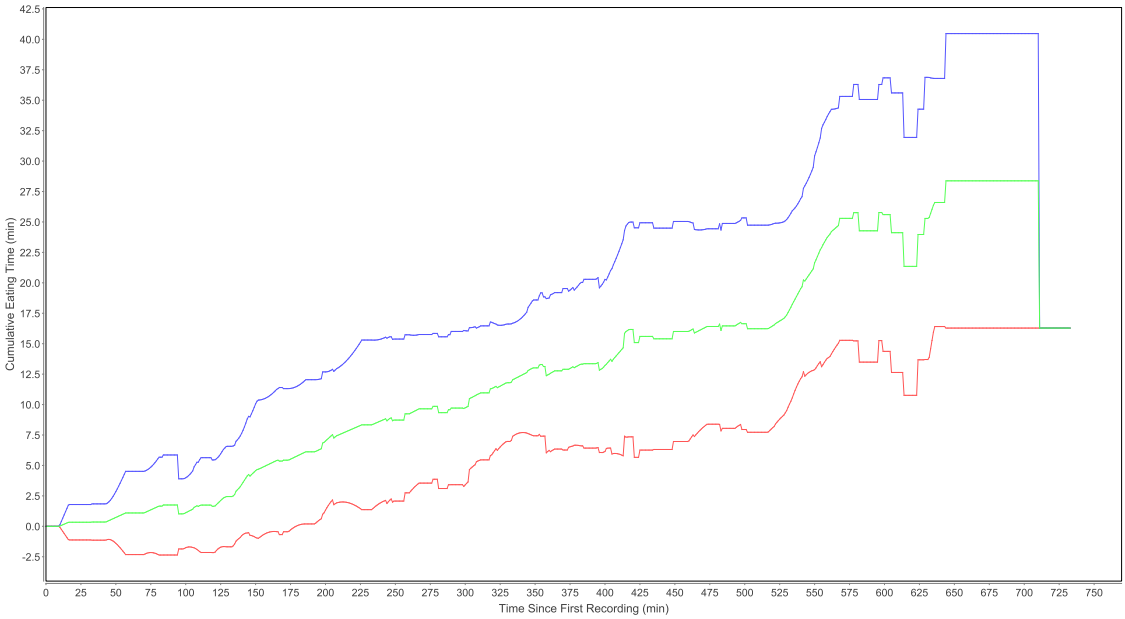

Figure 2.12: Cumulative eating time averages (set 1).

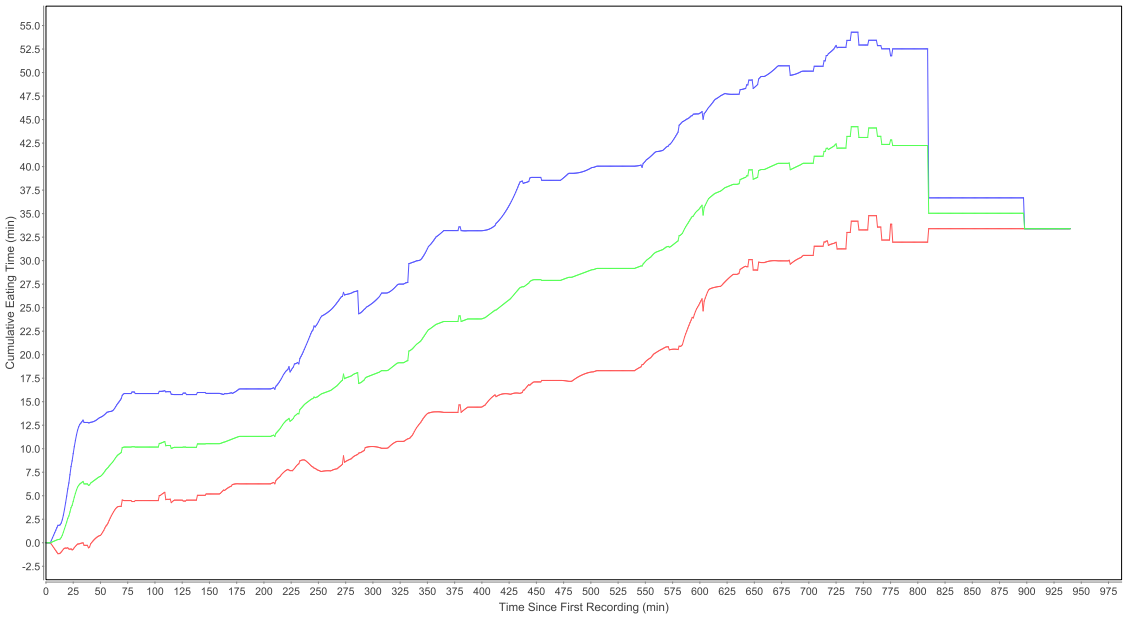

Figure 2.13: Cumulative eating time averages (set 2).

Feature	Equation
Mean	$\mu = \frac{1}{N}\sum_{i=1}^{N} x_i$
Mean of absolute value	$MAV = \frac{1}{N}\sum_{i=1}^{N} \lvert x_i \rvert$
Correlation	$Corr(X,Y) = \frac{\sum_{i=1}^{N}(x_i-\mu_X)(y_i-\mu_Y)}{\sqrt{\sum_{i=1}^{N}(x_i-\mu_X)^2 \sum_{i=1}^{N}(y_i-\mu_Y)^2}}$
Median	$Median\{X\}$
Power	$P = \frac{1}{N}\sum_{i=1}^{N} x_i^2$
Root mean square	$RMS = \sqrt{\frac{1}{N}\sum_{i=1}^{N} x_i^2}$
Standard deviation	$\sigma = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i-\mu)^2}$
Variance	$\sigma^2 = \frac{1}{N-1}\sum_{i=1}^{N}(x_i-\mu)^2$
Zero Crossings	ZC is incremented by 1 if $[x_i > 0$ and $x_{i+1} < 0]$ or $[x_i < 0$ and $x_{i+1} > 0]$

Table 2.4: Feature set.

2.5 Evaluation metrics

The primary metric used to evaluate the classifier is by the total amount of time correctly classified segments. Its accuracy can be computed as:

$$accuracy = \frac{TP * 20 + TN}{(TP + FN) * 20 + (TN + FP)} \tag{2.16}$$

where true positives (TP) were counted as the number of seconds of time that were labeled as eating in the manual logs and actually classified as an eating activity. False positives (FP) were counted as the number of seconds of time that were labeled as non-eating in the manual logs and classified as eating. True negatives (TN) and false negatives (FN) were computed in a similar way by comparing manual logs and activities that were classified as non-eating. The factor of 20 in Equation 2.16 weights true positives to true negatives at a ratio of 20:1. This is done because eating activities occur much less frequently than non-eating activities [12]. In addition, accuracy for eating and non

eating segments were also computed, and they were calculated as follows:

$$accuracy_{EA} = \frac{TP}{TP + FN} \tag{2.17}$$

and

$$accuracy_{nEA} = \frac{TN}{TN + FP} \tag{2.18}$$

An important aspect to consider is that the classifying process was completed individually in sets 1 and 2, but our evaluation metrics include overall results for both sets combined.

Chapter 3

Results

3.1 Previous results

Results for the original four features (described in Section 2.1) are shown first. Table 3.1 shows totals in seconds for TPs, FPs, TNs, and FNs. It also displays overall accuracies, which are based on equations 2.16, 2.17 and 2.18. As seen in Figure 3.1 the overall accuracy achieved by our group's previous work was 81%. Each of our new features was tested in combination with the original four in order to compare results against the previous outcome.

Totals (in s)				Accuracy		
TP	TN	FP	FN	Eating	Non-eating	Overall
65053	1254650	280915	15074	82%	81%	81%

Table 3.1: Results for original 4 features with both sets combined.

3.2 Regularity of manipulation

Statistics for regularity of manipulation computed in the previous chapter are shown on Table 3.2, and represent final values used for Equation 2.8. Results for the regularity of manipulation feature in combination with the original four can be seen in Table 3.3. The table shows totals in seconds for TPs, TNs, FPs and FNs. In addition, it displays accuracy for EAs and non-EAs detection, as well as the overall accuracy. Clearly seen from the table, overall accuracy dropped from 81% to 79%. Including regularity of manipulation, the classifier performs better at detecting eating activities, improving the detection from 82% to 85%. On the other hand the non-eating detection decreased from 81% to 72%. If we explore deep into these results we realize that there are a couple of reasons why the detection did not improve for non-EA activities. Firstly, if we recall figures 2.8 and 2.9 we can see that the distribution for the non-EA is a Gaussian-like distribution, where as for the EAs is not at all clear. Figure 2.8 suggests that the distribution is perhaps, bimodal. In any case, standard deviations for EA segments are much larger than for non-EA segments, as seen in Table 3.2. This fact causes the EA distributions to overlap on the left tail with the right tail of the non-EA distributions, so the classifier detects less true negatives and more false positives. Although, the number of segments of FPs was reduced, the amount of time of FP detection was greatly increased from 280915 seconds to 437175 seconds. An example of this issue is shown in Figure 3.1. This figure shows at the top the output for over four and a half hours after the classification process is completed for the original four features. The bottom displays the output after classification including regularity of manipulation. The blue horizontal bars indicate detected EA segments. This is just case where clearly the time of EA detected was greatly increased due to the combination of the increase in FPs and decrease in TNs.

Another reason for the poorer performance of this feature, is that the FFT is not capturing all the information we would like to. Recall Figure 2.9, data was obtained by computing values for the FFT at each individual segment throughout the entire data set. FFTs for different individual

	Eating		Non-eating	
Set	Mean	Std. Deviation	Mean	Std. Deviation
1	12090	3423	4657	2084
2	12710	5310	4425	2384

Table 3.2: Statistics for regularity of manipulation.

Totals(s)				Accuracy		
TP	TN	FP	FN	Eating	Non-eating	Overall
67988	1099451	437175	11697	85%	72%	79%

Table 3.3: Results for original 4 features and regularity of manipulation.

segments are shown in four figures. Two of these segments come from the EA distribution, and the other two from the non-EA distribution. Figures 3.2 and 3.3 show the extreme cases. Figure 3.2 displays the smoothed output of manipulation for a segment of the EA distribution that does not overlap with the non-EA distribution, and Figure 3.3 shows the output for a segment in a section of the non-EA distribution that does not overlap with the EA distribution. The EA segment spans for over 8 minutes, and the non-EA segment spans for 5 minutes. From these outputs we can conclude that the FFT should perform really well on the extreme cases, since the data shows significant differences between them. The EA segment shows a more regular pattern with peaks of relative equal magnitude. On the other hand, it is not clear that there is any kind of periodicity on the non-EA activity. Now lets visualize an EA and a non-EA that is obtained from a section where both distributions overlap, somewhere between 8000 and 8500 $\frac{deg/s}{G}$ in Figure 2.9. Figure 3.5 shows the smoothed manipulation for a non-EA segment, and we can now see that the segment has more regularity compared to the previous figure, Figure 3.3. Figure 3.4 shows the output for an EA segment. We can still see that the output is somewhat regular, so we must explore deeper into these results. Figure 3.6 shows a segment of smoothed original data from Figure 3.4, and Figure 3.7 shows a segment of smoothed original data for Figure 3.5. We can see that data differs significantly from each other. Figure 3.6 shows clear cycles of motion between 8 and 11 seconds, especially in the *roll* signal, where Figure 3.7 does not show that kind of regularity. Based on this analysis we can assume that the FFT is capturing the cyclical periodicity, but we do not clearly see a sinusoidal periodicity which we believe its causing the classifier to decrease accuracy.

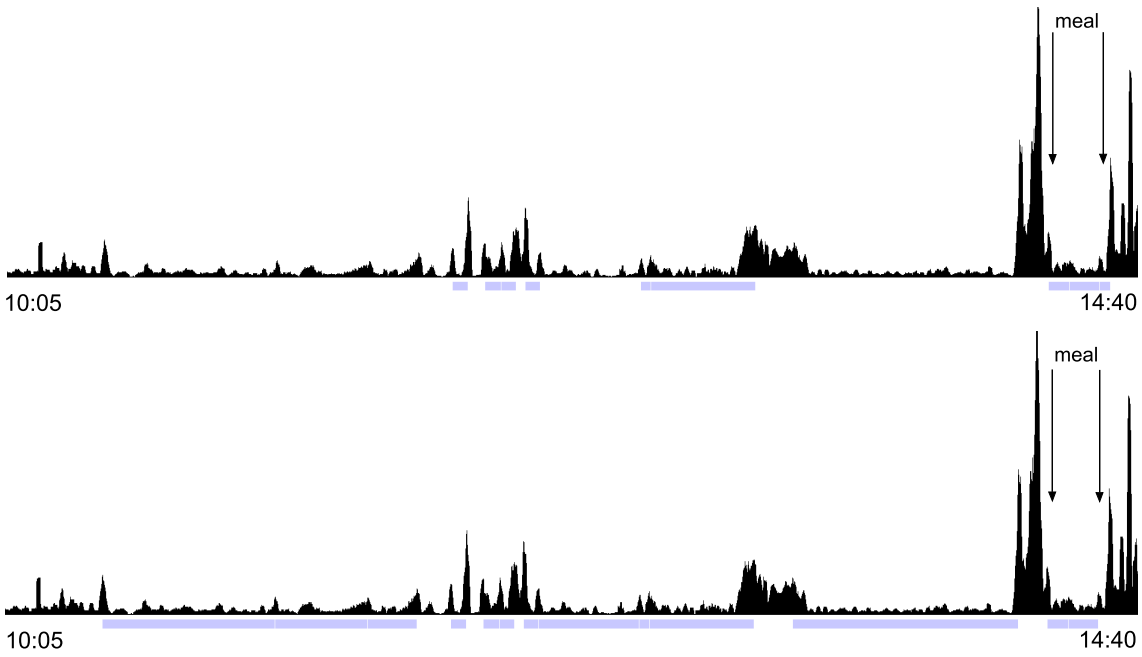

Figure 3.1: Comparison between original 4 features (top) and the combination of them including regularity of manipulation (bottom). Purple horizontal bars indicate EA detection.

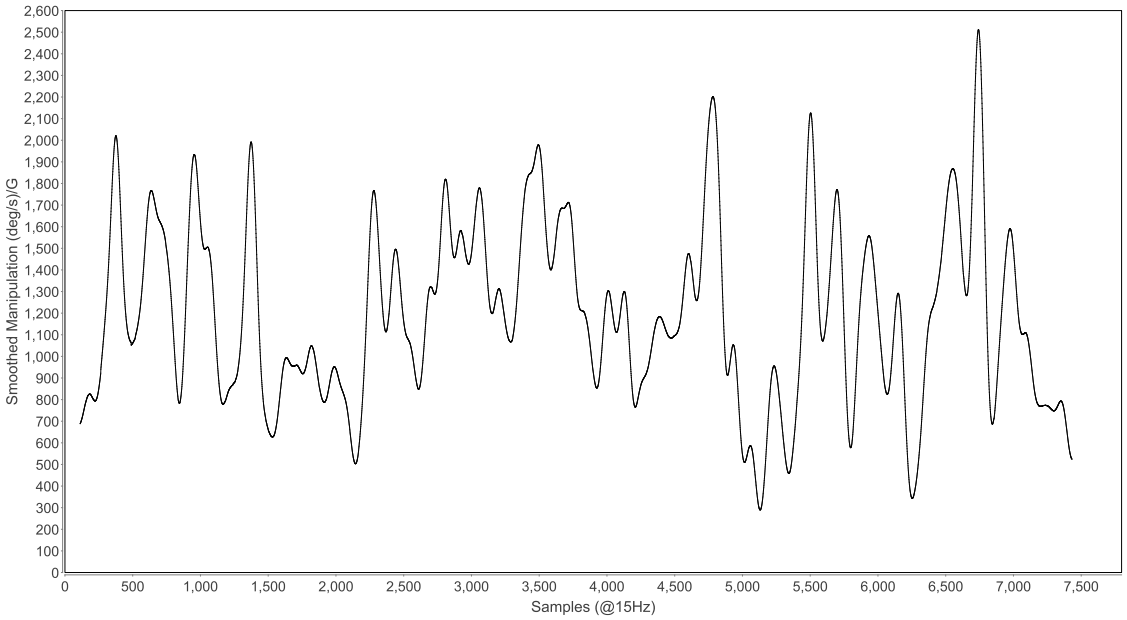

Figure 3.2: Smoothed manipulation data for an EA segment on the right tail of the distribution in Figure 2.9 of over 7400 samples.

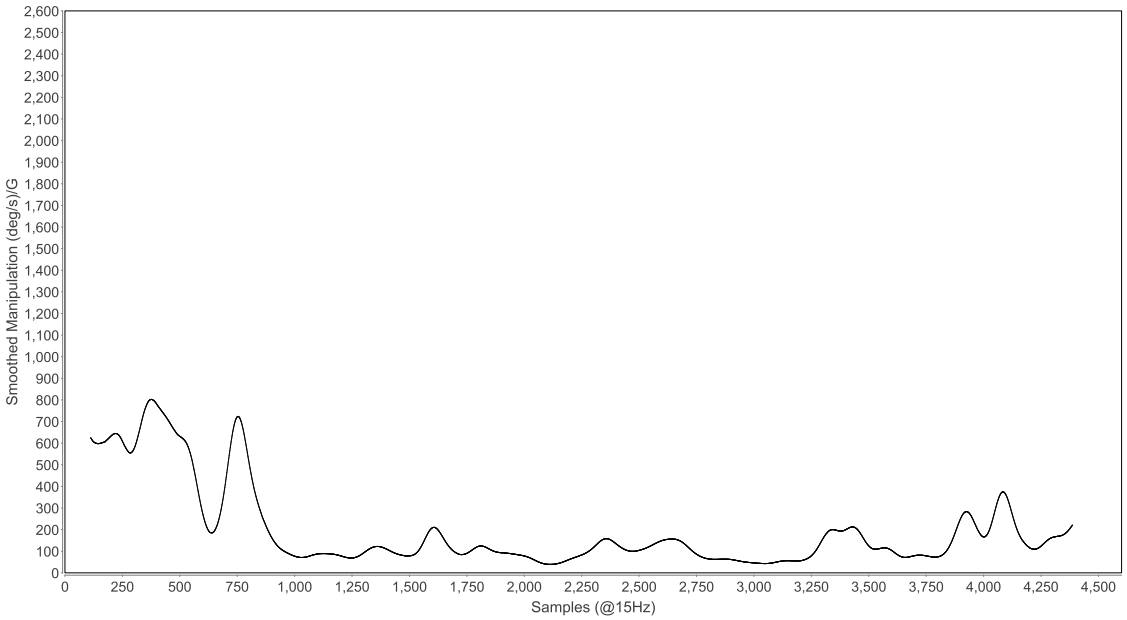

Figure 3.3: Smoothed manipulation data for a non-EA segment on the left tail of the distribution in Figure 2.9 of 4500 samples.

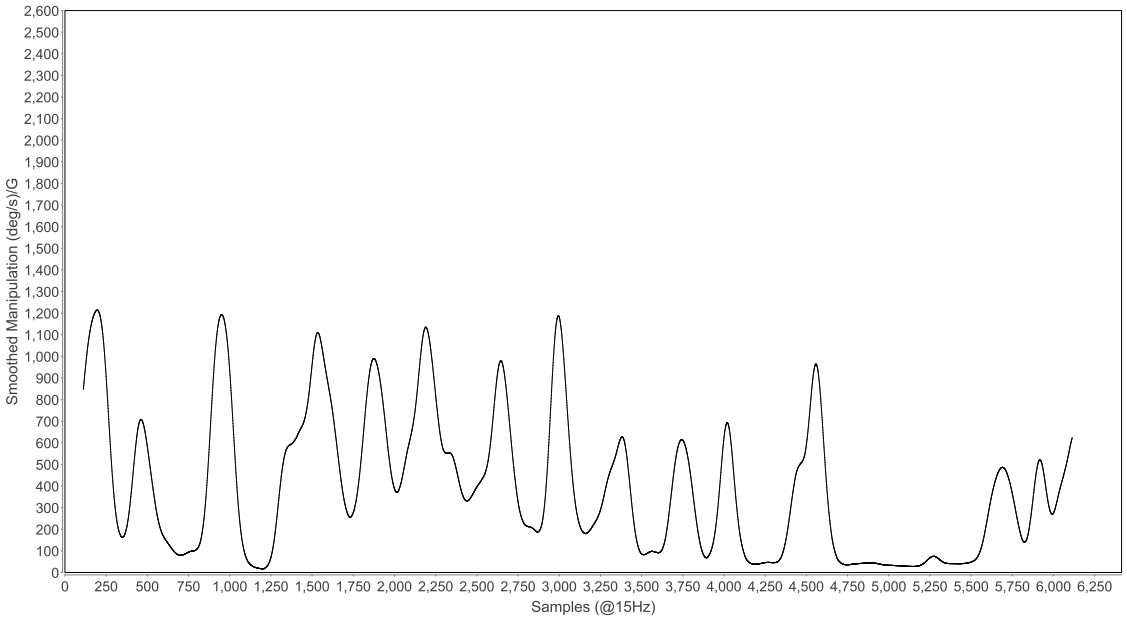

Figure 3.4: Smoothed manipulation data for an EA segment that overlaps with the non-EA distribution in Figure 2.9 of roughly 6200 samples.

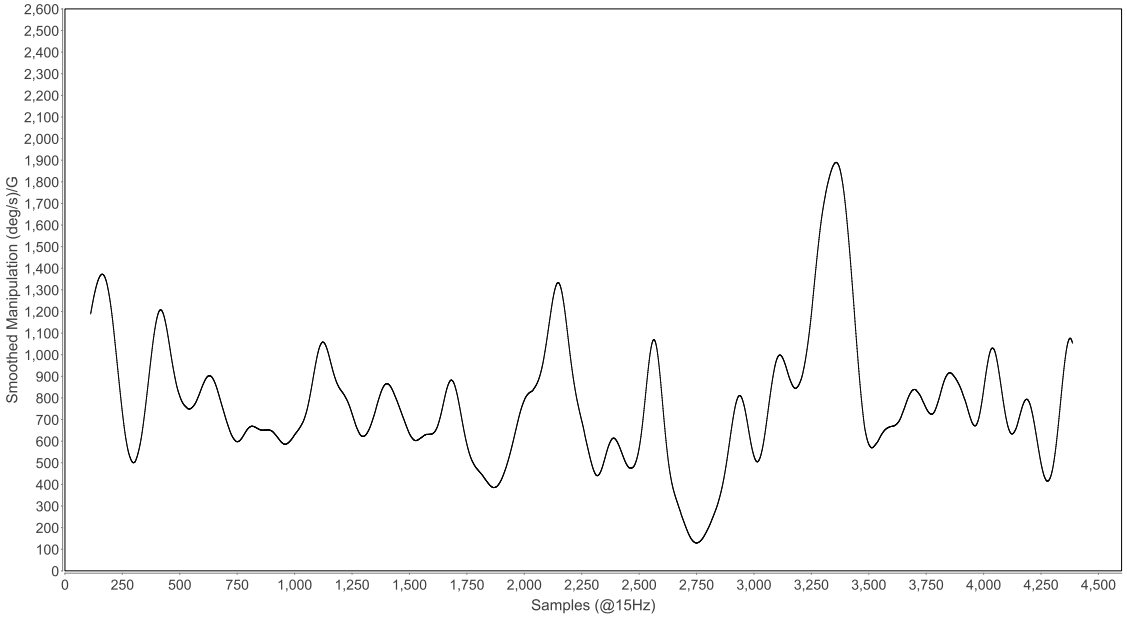

Figure 3.5: Smoothed manipulation data for an non-EA segment that overlaps with the EA distribution in Figure 2.9 of 4500 samples.

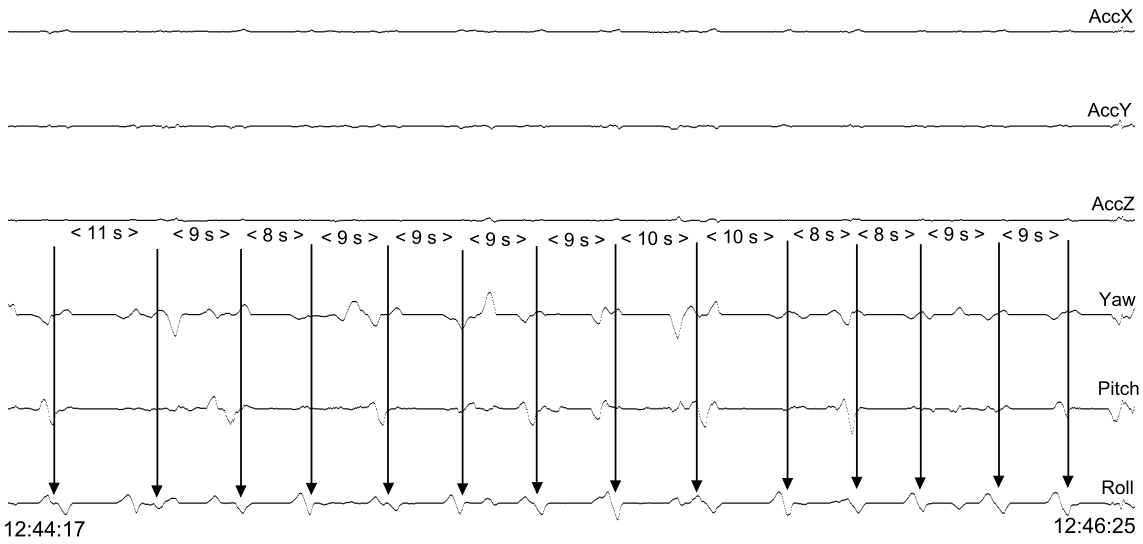

Figure 3.6: Original smoothed data for a segment in Figure 3.4. Arrows emphasize vigorous motions.

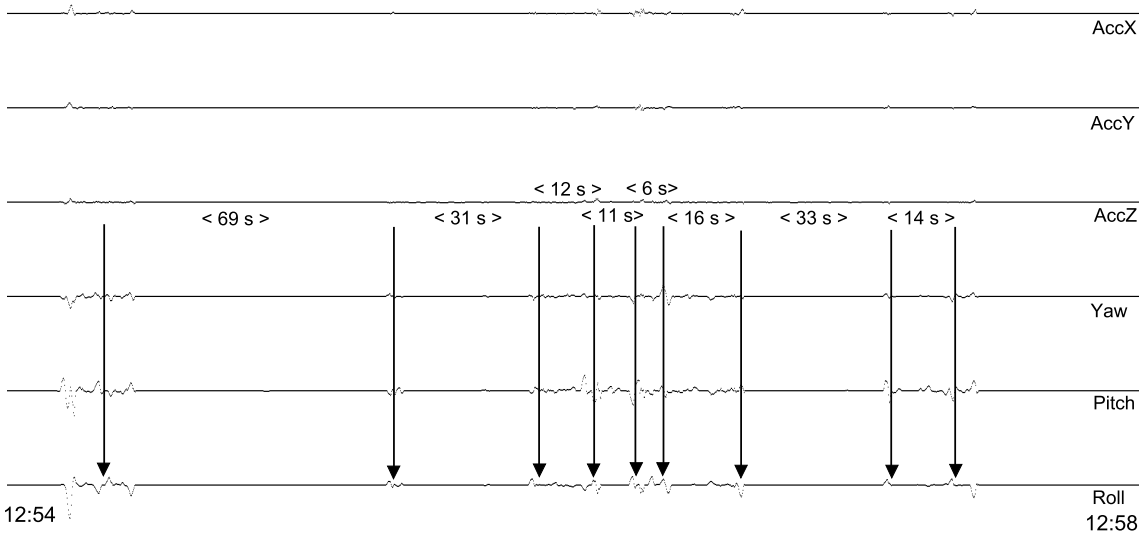

Figure 3.7: Original smoothed data for a segment in Figure 3.5. Arrows emphasize vigorous motions.

3.3 Time since last eating activity

Numerical results for he time since last EA feature are shown in Table 3.4. The results show the classifier's output for five features, the four original in addition to time since last EA. From these results we can see that the eating detection dramatically decreased from 82% to 46%, and the non-eating detection improved from 81 to 93% for an overall accuracy of 69%. It is evident that this feature is detecting non-EAs considerably well. This is because once a meal has been detected, probabilities will be extremely low for some time after the detection, contributing to detect more TNs, and suppress FPs. Figures 3.8 and 3.9 serve as example for this concept. Figure 3.8 outputs the original four features, and Figure 3.9 displays the output with time since last EA included. The purple horizontal bars underneath each segment indicate EA segments detected by the classifier. Arrows indicate start and end of meals based on ground truth. The classifier successfully suppressed 10 FPs in this particular segment, most of these came from FPs that were near EAs.

We can also notice from overall results that the EA detection significantly decreased. Recall the cumulative distribution function seen in figures 2.10 and 2.11. As a person continues to not eat, the probability of a person eating in the next minute keeps increasing, so the reason for the accuracy loss for eating segments, is that as time passes probability of eating increases at such a point that it outweighs the other features forcing the classifier to detect it as an EA segment. This idea is reflected in Figure 3.10. It compares the output between, again, the original four at the top and at the bottom the output includes time since last EA. The bottom figure clearly shows that the EA segment was suppressed when a recent EA was detected, after including our time since last EA feature.

Totals (in s)				Accuracy		
TP	TN	FP	FN	Eating	Non-eating	Overall
36649	1433574	103052	43036	46%	93%	69%

Table 3.4: Results for original 4 features and time since last EA.

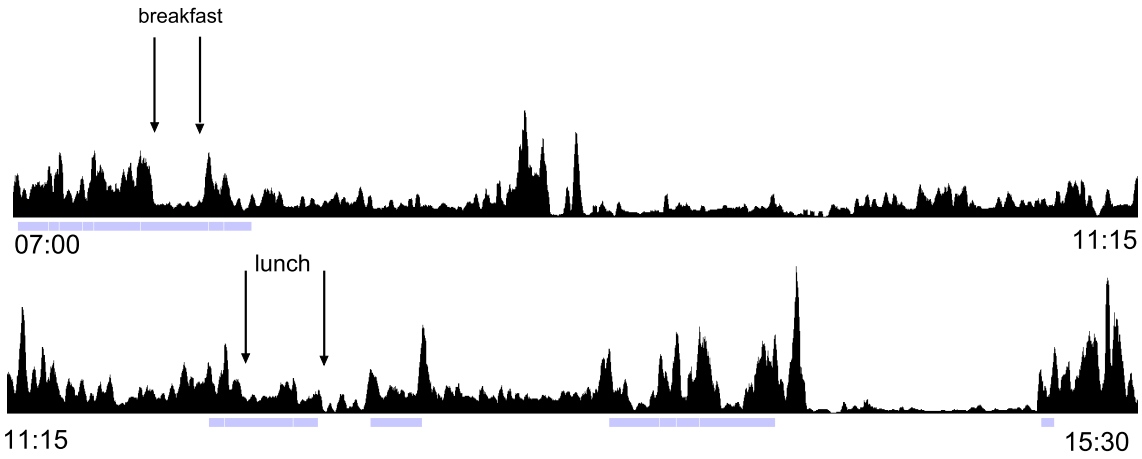

Figure 3.8: Classification period for original 4 features over a period of 8 and a half hours.

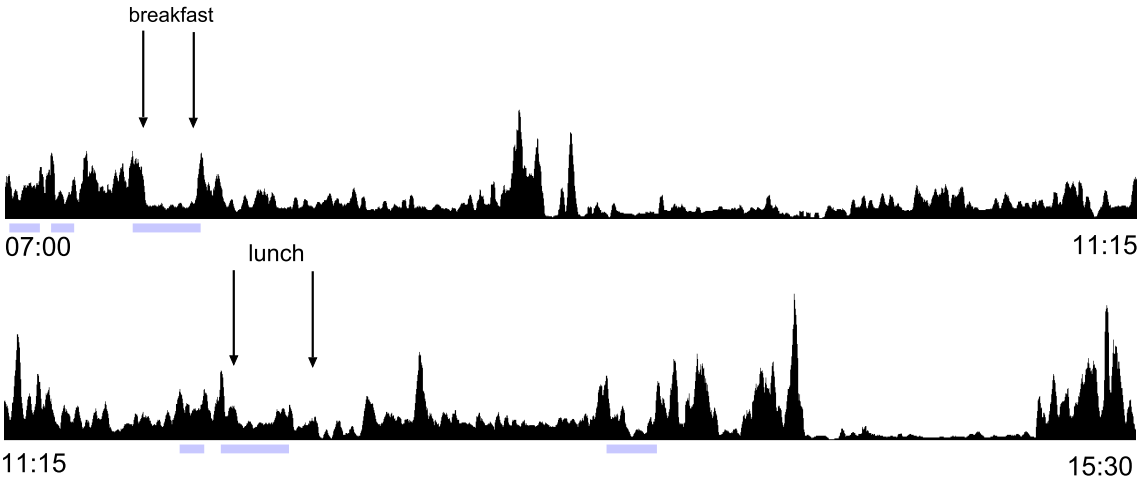

Figure 3.9: Classification period for 5 features including time since last EA over a period of 8 and a half hours. Note the many fewer false positives compared to Figure 3.8.

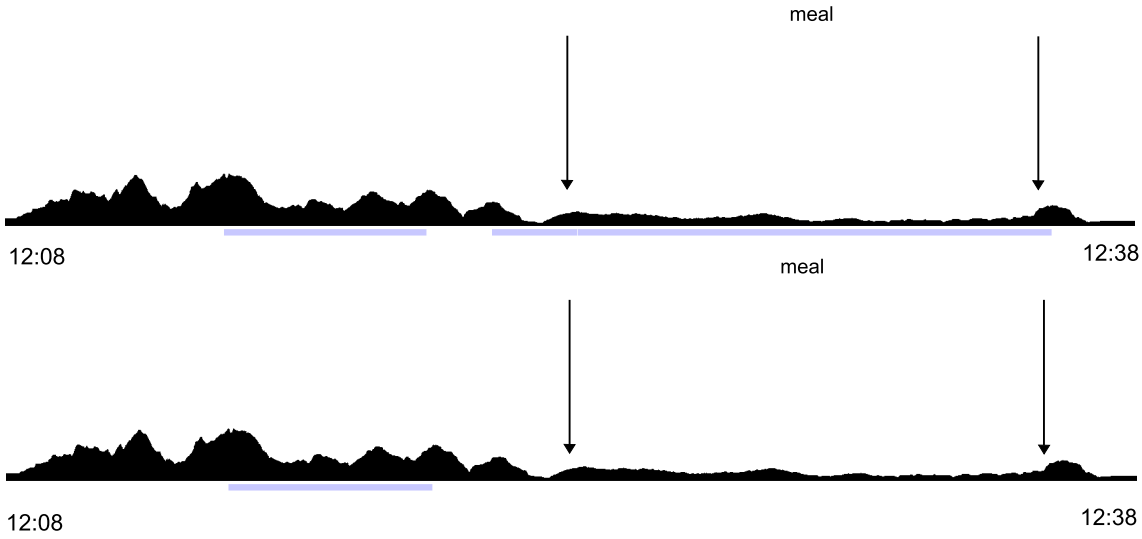

Figure 3.10: Comparison between classification outputs for original 4 features (top) and 5 features including time since last EA (bottom) over a 30 minute period. The false positive just prior to the actual meal makes the feature of time since last EA a strong inhibitor for immediately subsequent data, causing the actual meal to be mis-classified.

3.4 Cumulative eating time

Results for the cumulative eating time feature are displayed in Table 3.5. The outcome shows the results for the classifier, once again, of the original four features plus cumulative eating time. From these results we clearly notice that eating activities accuracy dramatically dropped to 31%, and the non-EA detection significantly increased to 98%. The reason behind the high non-EA accuracy is very similar to the reason behind the time since last EA feature. If a person eats, probabilities for EA activities will be very low for segments after the detected EA. Figures 3.11 and 3.12 enforces this premise. Figure 3.11 shows a classifier output for a six and a half hour period where only the four original features are displayed, and Figure 3.12 displays it for five features including cumulative eating time. It is evident that the classifier has suppressed all FPs. Another way to look at this is that right after *lunch*, "it is not time to eat, yet".

Totals (in s)				Accuracy		
TP	TN	FP	FN	Eating	Non-eating	Overall
24971	1502861	33765	54714	31%	98%	64%

Table 3.5: Results for original 4 features and cumulative eating time.

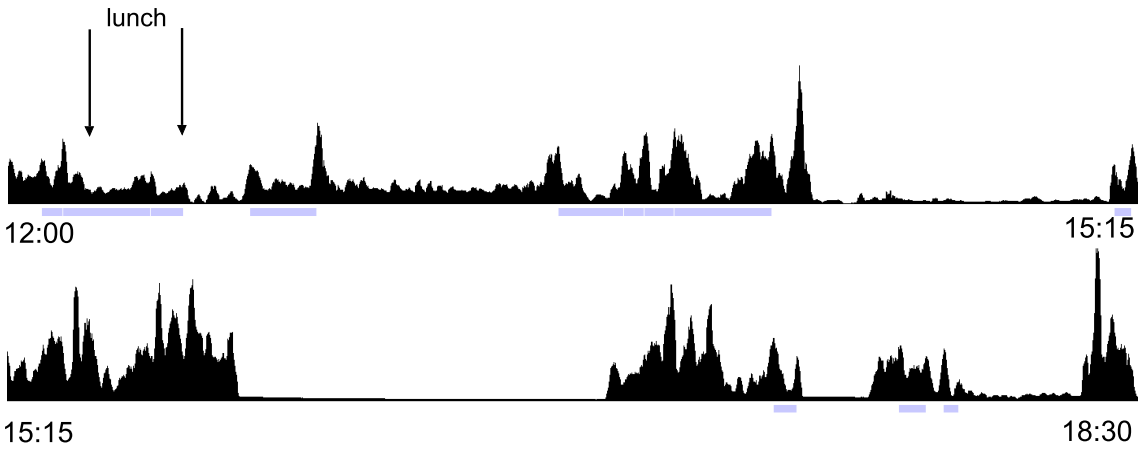

Figure 3.11: Classification period for original 4 features over a period of 6 and a half hours.

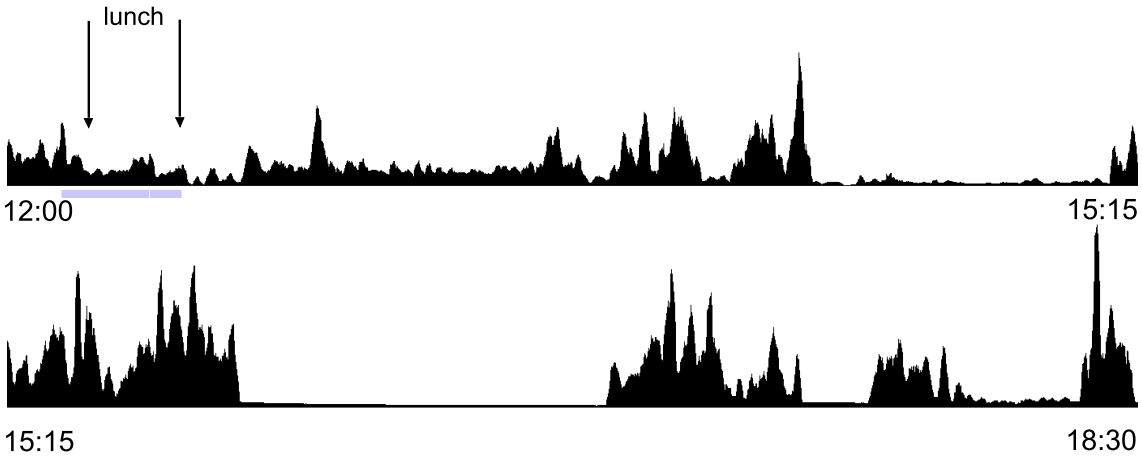

Figure 3.12: Classification period for 5 features including cumulative eating time over a period of 6 and a half hours. The true positive at the beginning of the segment causes the classifier to suppress immediately subsequent EA detections.

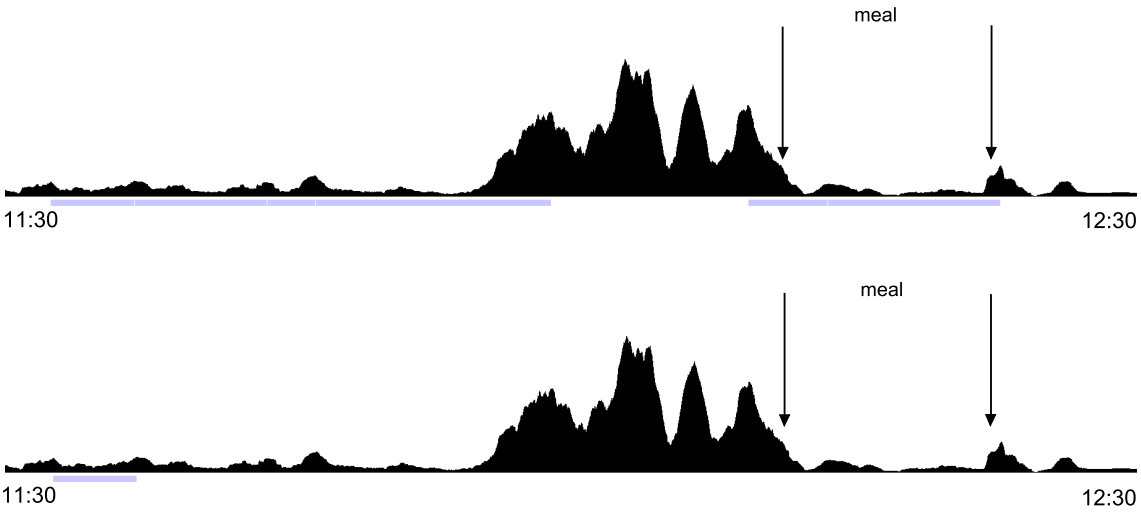

Figure 3.13: Comparison between classification outputs for original 4 features (top) and 5 features including cumulative eating time (bottom) over a 1 hour period. The false positive prior to the actual meal makes the feature cumulative eating time a strong inhibitor for immediately subsequent data, causing the actual meal to be mis-classified.

We also mentioned that the EA detection using this feature significantly decreased. If an EA is detected, as time passes the probability from the distribution for cumulative eating time at each minute thereafter will start increasing until its contribution to the classifying process will outweigh the rest of the features, forcing the detection of an EA. This is a very similar concept to time since last EA. Figure 3.13 shows a comparison between original features at the top, and five features including cumulative eating time for a period of an hour at the bottom. Even though the figure showing just the four features did not suppress all of the FP, it still managed to detect the correct EA segment. On the other hand, the bottom segment detected a FP at the beginning of the period and failed to detect the actual *meal* segment because it had just detected an EA.

Chapter 4

Conclusions and Future Work

In this thesis we have explored 3 different features that contributed to a classification process for detecting eating and non-eating activities in free-living by tracking wrist motion in addition to time-based activities. Regularity of manipulation, time since last eating activity, and cumulative eating time are features that had the potential to be great candidates for the classification process. Regularity of manipulation measured how regular the ratio of rotational motion to linear motion between certain frequency ranges. We found that this feature captures the regularity but that the regularity is not purely sinusoidal and thus may be better captured through the use of a wavelet transform or other method that captures non-sinusoidal periodicity. As for the time-based features, we imagined that time would be significant factor to consider; it turned out that both of the time-based features acted as *clocks*, forcing an eating activity detection even though the individual has not eaten yet.

For future work, we suggest that a more specific feature than regularity of manipulation needs to be modeled, one that can capture not only cyclical activities, but also other irregularities since the data suggests that there is not a consistent pattern that characterizes eating activities. As for the time-based features, we have realized that the features act like probabilistic counters that force the classification of a segment as eating activity based only upon the previously detected eating activity or total eating activity classified up to that point. Based on our findings using time-based features in real-time could be used under certain conditions, since the non-EA detection works significantly well. It would be ideal to come up with a process that could backtrack and re-classify segments when it detects a segment in the future with significant indication of being a meal segment

compared to a recently classified eating activity segment. This suggests the use of Hidden Markov models. A Hidden Markov model could search all possible classifications of all segments throughout the day, combining the time constraints with the other features, suggesting it to be an excellent candidate for future studies.

Bibliography

[1] R. Almaghrabi, G. Villalobos, F. Pouladzadeh, and S. Shirmohammadi, "A novel method for measuring nutrition intake based on food image," in *2012 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, May 2012, pp. 366–370.

[2] O. Amft and G. Tröster, "Recognition of dietary activity events using on-body sensors," *Artificial Intelligence in Medicine*, vol. 42, no. 2, pp. 121–136, 2008.

[3] O. Amft and G. Tröster, "On-body sensing solutions for automatic dietary monitoring," *IEEE Pervasive Computing*, vol. 8, no. 2, pp. 62–70, 2009.

[4] Apple Inc., http://support.apple.com/kb/SP587/, 1 Infinite Loop, Cupertino, CA 95014, 2014.

[5] L. Bluestein, "A linear filtering approach to the computation of discrete fourier transform," *IEEE Transactions on Audio and Electroacoustics*, vol. 18, no. 4, pp. 451 – 455, December 1970.

[6] C. Boushey, D. Kerr, J. Wright, K. Lutes, D. Ebert, and E. Delp, "Use of technology in childrens dietary assessment," *European Journal of Clinical Nutrition*, vol. 63, no. 1, pp. S50–S57, February 2009.

[7] Bureau of Labor Statistics, "American time use survey," http://www.bls.gov/tus/, U.S. Department of Labor, retrieved April 8, 2014.

[8] Centers for Disease Control and Prevention, "Adult obesity facts," http://www.cdc.gov/obesity/data/adult.html, retrieved April 9, 2014.

[9] C. Champagne, G. Bray, J. Delany, A. Kurtz, J. Monteiro, E. Tucker, and J. Volaufova, "Energy intake and energy expenditure: a controlled study comparing dietitians and non-dietitians," *Journal of the American Dietetic Association*, vol. 120, no. 10, pp. 1428–1432, 2002.

[10] Y. Dong, A. Hoover, and E. Muth, "A device for detecting and counting bites of food taken by a person during eating," in *IEEE Int'l Conf on Bioinformatics and Biomedicine*, 2009, pp. 265–268.

[11] Y. Dong, A. Hoover, E. Muth, and J. Scisco, "A new method for measuring meal intake in humans via automated wrist motion tracking," *Applied Psychophysiology and Biofeedback*, vol. 37, no. 3, pp. 205–215, 2012.

[12] Y. Dong, J. Scisco, M. Wilson, E. Muth, and A. Hoover, "Detecting periods of eating during free living by tracking wrist motion," *IEEE Journal of Biomedical and Health Informatics*, 2013.

[13] J. Fisher, B. Rolls, and L. Birch, "Children's bite size and intake of an entre are greater with large portions than with age-appropriate or self-selected portions," *The American Journal of Clinical Nutrition*, vol. 77, no. 5, pp. 1164–1170, May 2003.

[14] K. Glanz, J. Brug, and P. van Assema, "Are awareness of dietary food intake and actual fat consumption associated?a dutch-american comparison," *European Journal of Clinical Nutrition*, vol. 51, no. 8, pp. 542–547, August 1997.

[15] S. Jonnalagadda, D. Mitchell, H. Smiciklas-Wright, K. Meaker, N. V. Heel, W. Karmally, A. Ershow, and P. Kris-Etherton, "Accuracy of energy intake data estimated by a multiplepass, 24-hour dietary recall technique," *Journal of the American Dietetic Association*, vol. 100, no. 3, pp. 303–311, 2000.

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

[17] S. Lichtman, K. Pisarska, E. Berman, M. Pestone, H. Dowling, E. Offenbacher, H. Weisel, S. Heshka, D. Matthews, and S. Heymsfield, "Discrepancy between self-reported and actual caloric intake and exercise in obese subjects," *The New England Journal of Medicine*, vol. 327, no. 27, pp. 1893–1898, December 1992.

[18] Y. Ma, E. Bertone, E. Stanek, G. Reed, J. Hebert, N. Cohen, P. Merriam, and I. Ockene, "Association between eating patterns and obesity in a free-living us adult population," *American Journal of Epidemiology*, vol. 158, no. 1, pp. 85–92, 2003.

[19] C. Martin, H. Han, S. Coulon, H. Allen, C. Champagne, and S. Anton, "A novel method to remotely measure food intake of free-living people in real-time: The remote food photography method (rfpm)," *British Journal of Nutrition*, vol. 101, no. 3, pp. 446–456, 2009.

[20] Mayo Clinic, "Obesity," http://www.mayoclinic.org/diseases-conditions/obesity/basics/treatment/con-20014834, retrieved April 9, 2014.

[21] B. McCabe-Sellers, "Advancing the art and science of dietary assessment through technology," *Journal of the American Dietetic Association*, vol. 110, no. 1, pp. 52–54, 2010.

[22] K. McConaby, H. Smiciklas-Wright, L. Birch, D. Mitchell, and M. Picciano, "Food portions are positively related to energy intake and body weight in early childhood," *The Journal of Pediatrics*, vol. 140, pp. 340–347, 2002.

[23] D. Mozaffarian, T. Hao, E. Rimm, W. Willett, and F. Hu, "Changes in diet and lifestyle and long-term weight gain in women and men," *The New England Journal of Medicine*, vol. 364, pp. 2392–2404, 2011.

[24] L. Muhlheim, D. Allison, S. Heshka, and S. Heymsfield, "Do unsuccessful dieters intentionally underreport food intake?" *International Journal of Eating Disorders*, vol. 24, no. 3, pp. 259–266, 1998.

[25] National Heart, Lung, and Blood Institute, "How are overweight and obesity treated?" http://www.nhlbi.nih.gov/health/health-topics/topics/obe/treatment.html, U.S. Department of Health and Human Services, retrieved April 9, 2014.

[26] Oracle Corporation, http://www.oracle.com/us/technologies/java/standard-edition/overview/index.html, 500 Oracle Parkway, Redwood Shores, CA 94065, 2014.

[27] Y. Saeki and F. Takeda, "Proposal of food intake measurement system in medical use and its discussion of practical capability," in *Knowledge-Based Intelligent Information and Engineering Systems*, vol. 3683, 2005, pp. 1266–1273.

[28] E. Sazonov and E. Schuckers, "The energetics of obesity: A review: Monitoring energy intake and energy expenditure in humans," *IEEE Engineering in Medicine and Biology Magazine*, vol. 29, no. 1, pp. 31–35, 2010.

[29] E. Sazonov, S. Schuckers, P. Lopez-Meyer, O. Makeyev, E. Melanson, M. Neuman, and J. Hill, "Toward objective monitoring of ingestive behavior in free-living population," *Obesity*, vol. 17, no. 10, pp. 1971–1975, 2009.

[30] E. Sazonov, S. Schuckers, P. Lopez-Meyer, O. Makeyev, N. Sazonova, E. Melanson, and M. Neuman, "Non-invasive monitoring of chewing and swallowing for objective quantification of ingestive behavior," *Physiological Measurement*, vol. 29, no. 5, p. 525, 2008.

[31] J. Scisco, "Sources of variance in bite count," PhD dissertation, Psychology Department, Clemson University, SC., 2012.

[32] E. Sejdic, C. Steele, and T. Chau, "Segmentation of dual-axis swallowing accelerometry signals in healthy subjects with analysis of anthropometric effects on duration of swallowing activities," *IEEE Transactions on Biomedical Engineering*, vol. 56, no. 4, pp. 1090–1097, April 2009.

[33] B. Six, T. Schap, F. Zhu, A. Mariappan, M. Bosch, E. Delp, D. Ebert, D. Kerr, and C. Boushey, "Evidence-based development of a mobile telephone food record," *Journal of the American Dietetic Association*, vol. 110, no. 1, pp. 74–79, January 2010.

[34] F. Takeda, K. Kumada, and M. Takara, "Dish extraction method with neural network for food intake measuring system on medical use," in *2003 IEEE International Symposium on Computational Intelligence for Measurement Systems and Applications (CIMSA)*, July 2003, pp. 56–59.

[35] F. Thompson and A. Subar, *Dietary Assessment Methodology*, 2nd ed. "Academic Press/Elsevier", 2008.

[36] F. Thompson, A. Subar, C. Loria, J. Reedy, and T. Baranowski, "Need for technological innovation in dietary assessment," *Journal of the American Dietetic Association*, vol. 110, no. 1, pp. 48–51, 2010.

[37] J. Tooze, A. Subar, F. Thompson, R. Troiano, A. Schatzkin, and V. Kipnis, "Psychosocial predictors of energy underreporting in a large doubly labeled water study," *American Journal of Clinical Nutrition*, vol. 79, no. 5, pp. 795–804, 2004.

[38] World Health Organization, "Obesity and overweight," http://www.who.int/mediacentre/factsheets/fs311/en/, retrieved April 9, 2014.

[39] F. Zhu, M. Bosch, W. Insoo, K. SungYe, C. Boushey, D. Ebert, and E. Delp, "The use of mobile devices in aiding dietary assessment and evaluation," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 4, pp. 756–766, August 2010.

[40] F. Zhu, A. Mariappan, C. Boushey, D. Kerr, K. Lutes, D. Ebert, and E. Delp, "Technology-assisted dietary assessment," in *SPIE: Computational Imaging VI*.