A Feasibility Study of Wrist-Worn Accelerometer Based Detection of Smoking Habits

Philipp M. Scholl and Kristof van Laerhoven
Embedded Sensing Systems, ESS
Technische Universität Darmstadt
Email: scholl,kristof@ess.tu-darmstadt.de

***Abstract*—Cigarette smoking is one of the major causes of lung cancer, and has been linked to a large amount of other cancer types and diseases. Smoking cessation, the only mean to avoid these serious risks, is hindered by the ease to ignore these risks in day-to-day life. In this paper we present a feasibility study with smokers wearing an accelerometer device on their wrist over the course of a week to detect their smoking habits based on detecting typical gestures carried out while smoking a cigarette. We provide a basic detection method that identifies when the user is smoking, with the goal of building a system that provides an individualized risk estimation to increase awareness and motivate smoke cessation. Our basic method detects typical smoking gestures with a precision of 51.2% and shows a user-specific recall of over 70% - creating evidence that an unobtrusive wrist-watch-like sensor can detect smoking.**

***Keywords*- activity recognition; accelerometer; smoking; long-term sensing; ambulatory monitoring**

I. Introduction

Persuasive technologies have an unprecedented ability to monitor all aspects of health and lifestyle, and as such can equip users with novel technological tools to improve on self-monitoring and self-discipline. The increasing set of commercially available products that provide feedback, analysis and visualization of the user's fitness activities [1], sleep patterns [2], or sedentary episodes [3], indicates that there is a large interest in such technology. Such technologies have furthermore been adopted as key components in diagnosis and analysis studies for clinical and psychiatric studies as well [4].

Smoking has been called the single biggest preventable cause of death by the World Health Organization [5] with tobacco claiming millions of lives a year, most of them in developing countries. The report also states that much of the disease and premature mortality caused by tobacco may be considered as side-effects of the disease of addiction. Tobacco dependence itself is classified in the International Classification of Diseases (ICD-10) as a disease [6].

This paper focuses on the automatic detection and long-term capturing of the user's smoking episodes throughout the day, using data from a light-weight and inconspicuous sensor device worn on the wrist. We identify typical accelerometer patterns resulting from the smokers' cigarette-to-mouth gestures, show a basic method to quantify the similarity between these patterns and provide evidence that smoking can be detected with such a sensor, in a similar manner Sazonov et. al. [7] have shown with a different technology.

The remainder of this paper is structured as follows: The next section will provide a description of the proof-of-concept study on data from four regularly smoking participant that were recorded for about a week each, and give more details on the sensor device used. Section III will then discuss the results of said study, focusing on data quality and detection feasibility. A final section will then sum up the paper's contributions and mention future directions for this research.

II. Method & Analysis

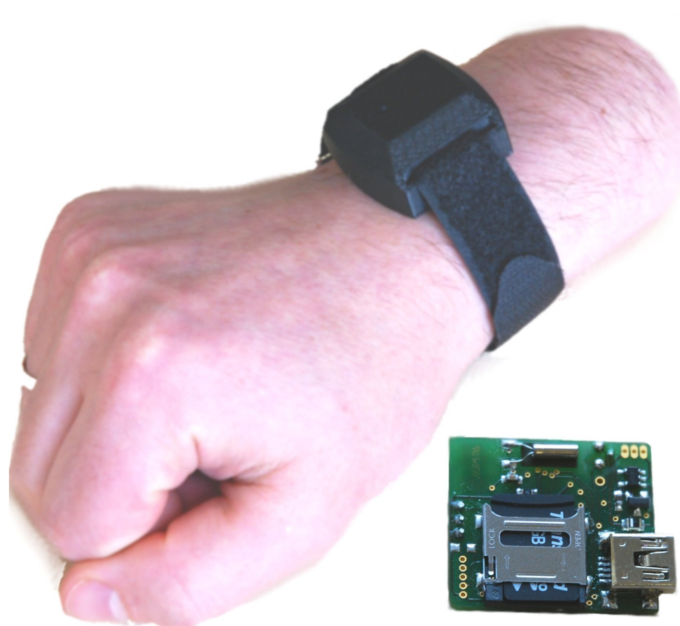

Fig. 1: The Hedgehog platform packaged and worn on the wrist and bare PCB. Participants have been asked to wear these in 12-hour wake-phase to monitor their smoking habits.

For this feasibility study we asked regular four smokers (aged 26 to 40, 2 male, 2 female) to wear an accelerometer-based sensor device (see Fig. 1) on their wrist over the course of about one week. The monitoring time, number of total recorded data samples, covered timespan and some basic statistics on manual labeling can be found in Table I. The participants have been asked to double-tap the wrist-worn

TABLE I: Summary description of collected data for the study participants, number of total smoking gesture patterns (number of "hard", "fair" and "perfect" samples in brackets), the mean duration of those gestures and the number of total accelerometer sample points in the data. Note that the last figure can be misleading as sample points are only recorded when subsequent measured values changed (using run-length compression), not representing the equidistant sampling points.

participant	timespan	#patterns	duration	#samples
0 (male)	8 days	35 (2|7|26)	4.6min	3.08M
2 (male)	5 days	28 (1|6|21)	6.8min	3.57M
1 (female)	5 days	34 (11|8|15)	8.1min	4.58M
3 (female)	5 days	19 (10|3|6)	8.7min	1.53M

sensor in order to get a marker in the data each time a cigarette has been smoked, which was used in the manual labeling phase. Participants were asked to strap the sensors on the respective dominant hand's wrist (all participant were right-handed) because we assume that this is the hand most often used to hold cigarettes. Furthermore, we asked the participant to wear the sensor over the course of the whole day and only take it off while sleeping, which gives a lot of background data to compare our detection algorithm to.

A. Data Acquisition

The data has been collected with wrist-worn "Hedgehog" sensing platform prototypes. Their design is based around a PIC18F microcontroller, which contains furthermore an ADXL345 acceleration sensor and a µSD-card. The acceleration data is continuously sampled at 100Hz and written to a FAT32 filesystem on the µSD-card in a compressed format. The range of the accelerometer has been set to $\pm 4g$ in order to have enough accuracy for human movement. The actual data can be retrieved by accessing the µSD-card via a standard USB mass-storage interface. The 180mAh battery included in the package can power the system for a total run-time of at least 7 days without the need for recharging.

B. Preprocessing

Before explaining how we labeled the data, we would like to build an intuitive understanding of the typical wrist postures of consuming a cigarette and how these facts can help in automatically recognizing this gesture. First of all an average smoker consumes a cigarette in $4\text{-}8\,\text{min}$, which gives us a fixed time-window during which we need to detect several similar wrist postures [7]. Second, for a limited amount of time, depending on the smoking style, the hand/wrist is near the mouth in a fixed angle, while inhaling the cigarette and after that in a different position. And third, a similar wrist posture is given while lighting the cigarette. Based on that we can already say that we are looking for certain repetitive posture sequences of the wrist in a fixed time window of $4\text{-}8\,\text{min}$.

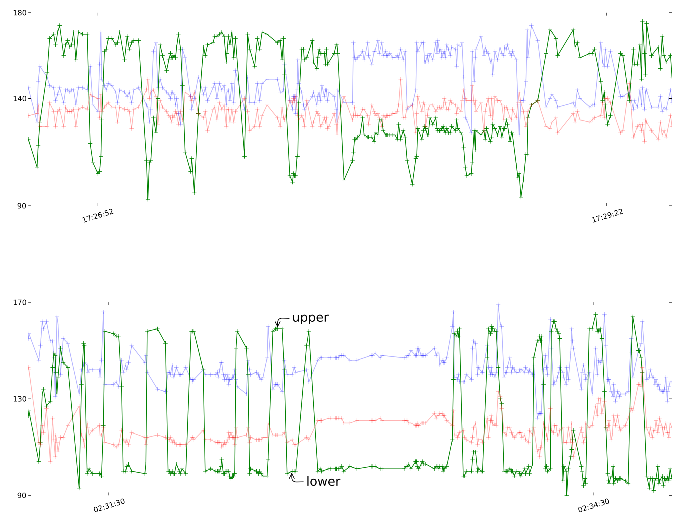

Fig. 2: Two sample patterns in a time series visualization, with the sensor being worn in two different orientations, and with the upper pattern also showing the influence of the sensor strap not being worn tightly enough. Annotations in the bottom pattern show the two fixation states of the smoking gesture.

These postures are directly proportional to the inclination of the wrist. To measure inclination from an accelerometer, i.e. the direction of earth's gravitational pull, the measurements have to be done when the accelerometer is quasi-static. In our case we can safely assume that there will be no constant acceleration of the wrist, but only short peaks of change due to a movement of the arm. Thus, in order to extract posture data from the acceleration signal, we use a low-pass filter on the sensor readings at an empirically determined cut-off frequency of 5Hz and therefore assume that as long as the readings change at a lower frequency we can extract the wrists inclination.

C. Manual Labelling

For further analysis, we labeled the raw accelerometer data with the help of the double-tap indicators - we asked the participants to strongly tap the sensor twice before or during smoking to mark when a cigarette has been smoked. By visually inspecting the raw data (as depicted in Fig. 2 and 3) for the double-taps, we extracted the patterns of smoking gestures following these markers. We then subdivided those extracted patterns into three classes, i.e. perfect, fair and hard. Note that these classes represent the authors' confidence whether the participant was smoking and the subjective similarity to other posture patterns. While perfect means full confidence in a pattern having emerged from having a cigarette, fair means partial confidence due to noisy data and very limited number of repetitions and hard means that there might have been smoking but data is too noisy, the number of repetitions is limited or the pattern is highly different. Fig. 3 shows a pattern in each class for every participant.

TABLE II: The four mean and variance samples of two states in the cigarette-to-mouth gesture. "Lower"-state describes when the hand is not near the mouth, while "upper"-state describes when the hand is near the mouth. The Gaussians are displayed for each participant and the identified Gaussian combined for all participants. Data values represent accelerometer value with an $8\mathrm{bit}$ resolution and a range of $\pm 4g$.

participant	lower	upper
0 (male)	$100.15(2.99)$	$158.13(6.19)$
1 (female)	$100.85(2.41)$	$141.45(8.02)$
2 (male)	$101.35(3.16)$	$158.99(6.59)$
3 (female)	$97.59(3.47)$	$145.23(8.89)$
	$\mathbf{99.84(3.45)}$	$\mathbf{150.95(10.84)}$

D. Feature Selection

By further examining the data we identified the following challenges, which an automatic classification algorithm based on wrist-worn accelerometer data needs to tackle:

- different smoking styles
- non-fixed sensor position
- superposition with other activities

Sample signals of those challenges can be seen in Fig. 3. Different smoking styles exhibit itself in how participants hold theirs cigarettes, the timespan between inhaling and if standing or sitting. All the shown "perfect" (first row, Fig. 3) samples have been recorded while the participant was standing, but there are also several samples where participants have been sitting (second row, third column) which also exhibits the repetitive movement pattern on the Y-axis but with a different range. This is because, while standing, the arm of a smoker rests pointing downwards, while when sitting, the arm rests in a $45°$ angle. Also, people tend to change the hand with which they smoke, making it impossible to record any useful data when the cigarette is not in the dominant hand. An example of this effect can be seen in the bottom column of Fig. 2. Another problem when trying to detect those postures is the position of the sensor on the wrist and if it is firmly attached. Fig. 2 shows two different patterns of the smoking activity where the sensor has been worn with two different orientations, you can see there that the fixation period (the period where values stay around a single value for a longer time) are mirrored. Superposition with other activities present a further challenge, for example some participants prefer to smoke while walking, which cannot be filtered easily.

From these consideration we can deduce a basic classifier. The cigarette-to-mouth gesture can be split into two postures, which we call the "upper" and "lower" posture, that is when the hand is near the mouth and when the hand is near the waist. These two postures mainly influence the Y-Axis of the accelerometer data, and are highlighted in Fig. 2. We selected the mean and variance, i.e. the Gaussian, of these two postures as the feature to classify by, and manually extracted those from one "perfect" patterns of each participant. The numerical result of this can be found in Table II. We then combined these four Gaussians into two single cross-participant Gaussians representing the "upper" and "lower" posture, which are the features we are looking for during automatic classification.

E. Feature Extraction

We now want to find occurrences of the "upper" and "lower" Gaussians in the complete dataset. For this we employed a simple iterative algorithm based on an adaptive window. We continuously calculate the mean and the variance of a buffer of variable length. Accelerometer data is added element-wise to this buffer, until the calculated variance is greater than half of one of the pre-determined Gaussians variance values. In which case we record the deviation between the calculated mean and pre-determined Gaussian mean, empty the buffer and continue with the rest of the accelerometer data.

Applying this algorithm for both the "upper" and "lower" Gaussians, results in two lists of deviations between the accelerometer data and pre-determined Gaussians.

F. Classification

To automatically find the cigarette-to-mouth smoking gestures, we combined the two previously introduced list of Gaussian deviations. Summing up those lists over a fixed time window of roughly $5.4\mathrm{s}$, i.e. the mean length of two subsequent cigarette-to-mouth gestures [7], results in the similarity score we used to identify the gesture. After applying an empirically determined threshold to this summed list, we were able to identify time windows where participants had a cigarette. Because participants also tended to change the hand which hold the cigarette, we furthermore merged identified windows which were separated but do not span more than $4 - 8\mathrm{min}$, the mean time it takes to consume a cigarette.

G. Evaluation

Table III shows the results of this automatic classification compared to our manual labeling. What is visible there is the precision ratio of the classification, i.e. how many automatic classifications match our manual labeling and how many do not, as well as the hit-ratio, which describes the number of matches of automatic classifications in each class of manually labelled data. In total 116 episodes of cigarette-smoking have been monitored, of which for all but one participant more than 55% could be identified with automatic classification.

This is a promising result, since this is achieved by straightforward thresholding and by Gaussian modeling of the "upper" and "lower" posture states before and after the cigarette-to-mouth gestures. Further analysis of actual gesture data, as well as higher-level models of sequences of posture changes might in combination with this method attain better classification results. The proposed algorithm is however both fast and has a small footprint, so that it could be implemented on the sensor and act in an on-line fashion, i.e., on the streaming sensor data.

TABLE III: Accuracy results of our basic Gaussian classifier. Positives are calculated as the ratio between total number of automatically identified occurrences and the ones which matched the manual labelled ground-truth (true positives) and ones which did not match (false positives). The hit-ratio is the number of matches between manually labelled occurrences and automatically identified occurrences.

	positives (precision)		hit-ratio (recall)		
	true	false	hard	fair	perfect
0	56.4%	43.6%	0.0%	14.3%	48.6%
1	61.8%	38.2%	63.6%	50.0%	73.3%
2	69.2%	30.8%	100.%	16.7%	76.2%
3	17.4%	82.6%	30.%	0%	16.7%
	51.2%	48.8%	48.8%	20.2%	53.7%

H. Summary

To summarize the proposed algorithm, our goal is to get from raw accelerometer data to a list of timestamps that mark the start and end of a cigarette-to-mouth gesture, for this we:

1) low-pass filtered accelerometer data with a cut-off frequency of 5Hz.
2) split data into regions of varying length where its variance is below the ones from our "upper"- and "lower"-Gaussians.
3) calculate the deviation of the mean of our regions to the mean of the "upper"- and "lower"-Gaussians.
4) sum up the deviations with a fixed-time window of roughly 5.4s.
5) record the timestamps when the sum of deviations rises and falls below a threshold.
6) merge the rise and fall time if it occurs within a time frame of $4 - 8\text{min}$.

The comparison of manual and automatically labeled smoking data can be found in Table III.

III. Discussion

Several things should be noted when interpreting the results presented in Table III. First of all, we are working with an approximate ground truth, which we gathered by letting the participants double-tap the sensor prior to, during, or at the end of having a cigarette. We then manually labelled the timespan in which we could identify a pattern which we deemed to result from a cigarette-to-mouth gesture. While the probability that this gives us a wrong label is low (since we have been looking for repeating patterns in the whole dataset) the probability that we missed a similar pattern is inevitably higher. Often the participants simply forgot to double-tap the sensor, or the pattern is just not similar enough to the ones we identified beforehand. The true number of false positives is thus likely more optimistic than reported here, as the classifier indeed identified the cigarette-to-mouth gesture correctly but our manual labeling was too conservative.

Furthermore, for this feasibility study, we concentrated on a single frequently occurring cigarette-to-mouth gesture. While the accelerometer pattern that results from this, is prominent in the data for all participants, it does show an interesting variation over different days. At most times the Y-axis of the 3d-accelerometer is influenced the most, while X- and Z-axis are quasi-static, but this is only the case when the sensors is worn tight on the wrist (compare rows of Fig. 3). When it is worn loose, also the X- and Z-axis show a high amount of change similar to the pattern seen on the Y-axis, this phenomena can be seen in the upper row of Fig. 2. This also shows in the data after the participant gets up and re-attaches the sensor in the morning, when the whole dataset shows then a different "smoking"-pattern. This also explains the low number of recall and precision in the dataset of participant 3, which tended to wear the sensor in a loose way that made it harder to recognize our identified pattern with our basic classifier.

When interpreting the recall values, it should be cross-checked with the number of total occurrences of the specific class as given in Table I. Since the number of occurrences of each class is varying for each participant, this number might be unfair for comparison between participants.

The basic classifier we introduced in this paper is based on a number of assumptions regarding the cigarette-to-mouth gesture, which could hold only in specific cases. We assumed that the participants were smoking while standing still and moving their dominant hand between their mouth and a lower position. This is of course only one specific gesture smokers tend to exhibit, others for example might prefer to smoke while moving or walking, which would also result in a different accelerometer pattern, which the algorithm proposed in this paper does not check for. Another assumption that this classifier builds on is that a cigarette is usually smoked in a time-frame of $4 - 8\text{min}$, which is reasonable since a usual cigarette burns down in 10min. Certain cigarettes or cigars might however cause different smoking times.

It is finally important to stress that the dataset for this study is a very realistic one. It was recorded in an unobtrusive manner with the participants reporting being unaware of wearing the sensor for most of the time. Furthermore, participants wore the sensor during their entire wake-phase which gives a large amount of background data to assess the possible confusion with other activities, for example eating or drinking. Those might exhibit similar hand-to-mouth gestures and postures, which could explain the rather high false-positive rate. Because of the way we obtained the ground-truth data, we are unable to assert this. However, compared to a study under laboratory-condition our data can be expected to be highly realistic, since we used an unobtrusive sensor that has been worn through the course of several days.

IV. Conclusion and Future Work

The final conclusions we can draw from the results in Table III is that in this feasibility study we were able to detect a smoking-specific gesture from wrist-worn accelerometer data with a precision of 51.2%. It should be noted that these results

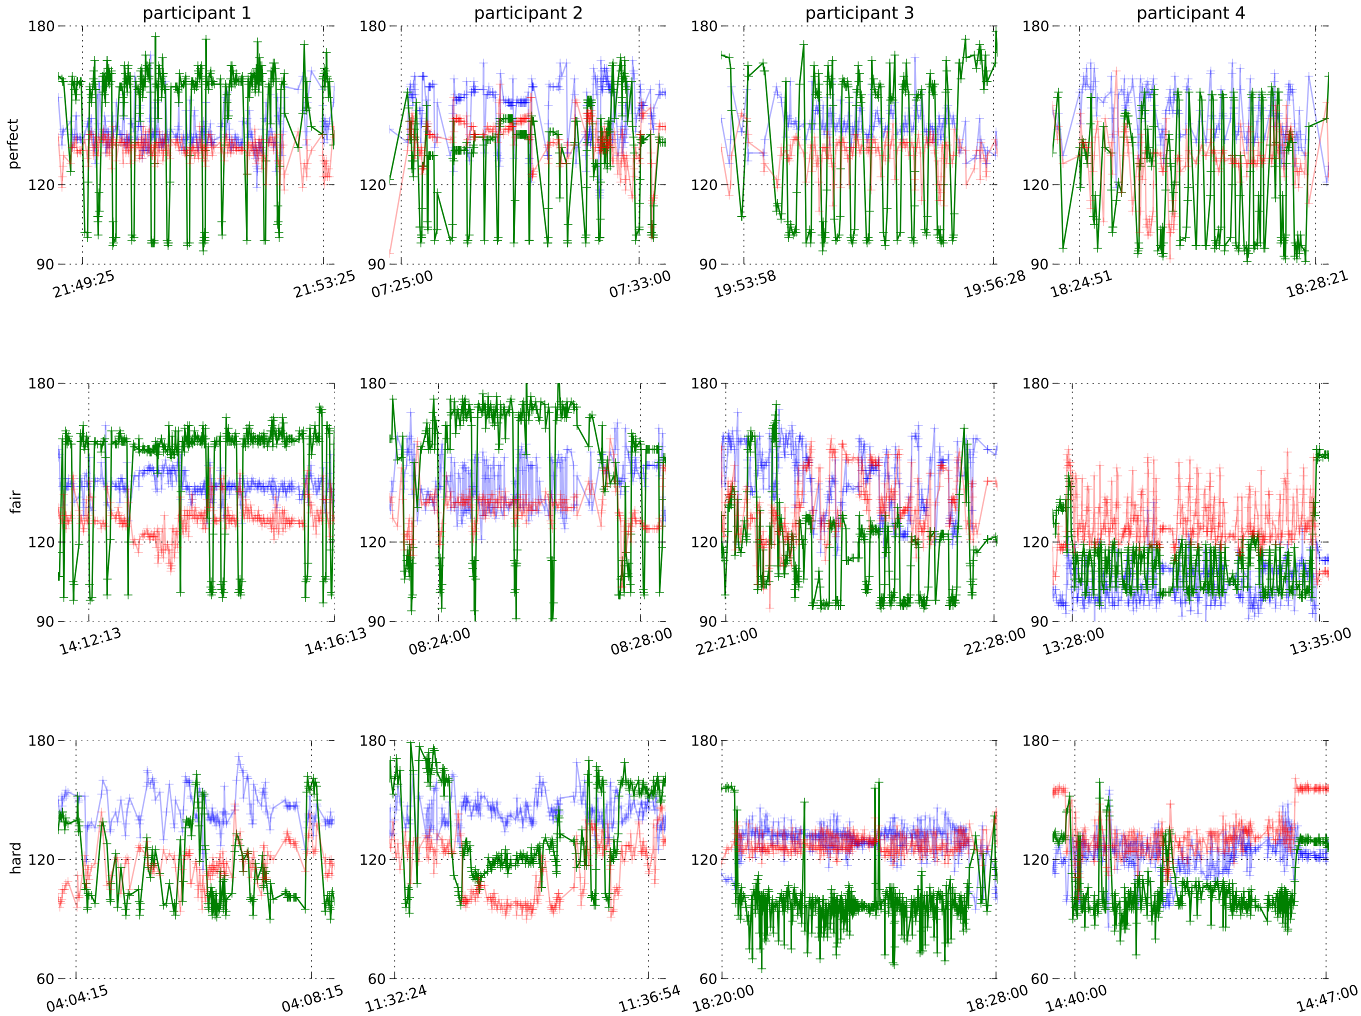

Fig. 3: Raw X-Y-Z accelerometer data (red,green.blue), for each participant (columns) and each class (row). The pattern for smoking while standing is clearly visible in the top row and is most prominently visible on the Y-axis, while X- and Z-axis are rather static. Strong variations can be seen in the "fair"-patterns, where for participant 3 and 4 an influence on all axis can be observed and for participant 1 and 2 a break of the repetition, probably due to flipping the cigarette to a different hand.

are achieved with a very basic classifier that only searches for the occurrences of two Gaussians and with the assumption of a $4-8\mathrm{min}$ duration of smoking a cigarette. The specific recall rates of over 70% for two participants indicates that the classifier is biased to the smoking style of those participants - hinting towards the fact that a classifier trained for specific users will perform better. The classifier used here can be seen as a very basic Gaussian Mixture classifier [8] [9], and we believe that these kinds of classifiers have enough flexibility to incorporate the number of assumption we are able to make about smoking gestures.

In this paper we only used one axis of the wrist-worn sensor. While this axis was the one that was most influenced by the smoking gestures most of the time, we have also found examples where patterns could be seen on all three available axes (possibly due to the sensor not being attached firmly). Future work on smoker classification therefore needs to include all available axes. Adding additional sensor like gyroscopes and compass to build a full Attitude-and-Heading Reference system might also prove useful in detecting gestures based on wrist-postures. Furthermore, another property of the cigarette-to-mouth gesture is that it usually results in the wrist travelling in height. This difference can be measured with a high-sensitivity barometer and should give a pretty detailed view of the gesture.

In order to assess the possibility of detecting smoking with a wrist-worn accelerometer a study with more participants needs to be conducted. This should give enough insight into different smoking styles, smoking behaviour and gestures to detect smoking. For a further study another way to get the exact number of consumed cigarettes needs to be found. Instead of letting participants write down the exact times when they had a cigarette, one could think of a lighter that contains a real-time clock which counts how often it has been lit during the day. Solely used for lighting cigarettes this would enable unobtrusive monitoring of smoking behaviour too. This is also true for the system introduced by Sazonov et.al. who have used a wrist-worn sensor based on rfid-measurements, which results in better recognition rates. Compared to our wrist-worn accelerometer approach this only allows to track smoked cigarettes without any other activities.

Furthermore it should be noted that the ground truth presented in this paper is only approximate, but still gives evidence that detecting typical smoking gestures and therefore consumed cigarettes per day with an unobtrusive wrist-worn inertial sensor is feasible, which can potentially lead to a system that increases self awareness and helps in smoking cessation. This can be as simple as quantifying the number of smoked cigarettes, or having an automated ubiquitous system [10] that makes it harder for the user to ignore the risks and motivate healthier behaviour.

References

[1] A. Teller and J. Stivoric, "The bodymedia platform: continuous body intelligence," in *Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*. ACM, 2004, pp. 114–115.

[2] E. Choe, S. Consolvo, N. Watson, and J. Kientz, "Opportunities for computing technologies to support healthy sleep behaviors," in *Proceedings of the 2011 annual conference on Human factors in computing systems*. ACM, 2011, pp. 3053–3062.

[3] S. Patterson, D. Krantz, L. Montgomery, P. Deuster, S. HEDGES, and L. NEBEL, "Automated physical activity monitoring: Validation and comparison with physiological and self-report measures," *Psychophysiology*, vol. 30, no. 3, pp. 296–305, 1993.

[4] M. Teicher, "Actigraphy and motion analysis: new tools for psychiatry," *Harvard Review of Psychiatry*, 1995.

[5] W. H. Organization, "Mortality country fact sheet, 30 may 2006," *Available on: http://www.wpro.who.int/ (Accessed on 20 February 2012)*, 2006.

[6] *The ICD-10 classification of mental and behavioural disorders: diagnostic criteria for research*. World Health Organization, 1993.

[7] E. Sazonov, K. Metcalfe, P. Lopez-Meyer, and S. Tiffany, "Rf hand gesture sensor for monitoring of cigarette smoking," in *Sensing Technology (ICST), 2011 Fifth International Conference on*. IEEE, 2011, pp. 426–430.

[8] T. Hastie and R. Tibshirani, "Discriminant analysis by gaussian mixtures," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 155–176, 1996.

[9] C. Fraley and A. Raftery, "How many clusters? which clustering method? answers via model-based cluster analysis," *The computer journal*, vol. 41, no. 8, pp. 578–588, 1998.

[10] M. Ayabe, Y. Okuda, V. Lehdonvirta, and E. Tokunaga, "Effecting lifestyle changes through ubiquitous feedback systems," *Proceedings of PerGames*, 2007.