Monitoring Smoking Behaviour Using a Wearable Acoustic Sensor

Iñigo Torres Uribe Echebarría, Syed Anas Imtiaz, Mingxu Peng and Esther Rodriguez-Villegas

Abstract— **Smoking is a cause of multiple health problems resulting in diseases which can also be fatal. It is well known that smoking has long-term impact on the health of an individual as well. While a number of studies have looked at the impact of smoking on health and its economic impacts, most of these rely on input from smokers in the form of questionnaires and surveys. Long-term monitoring of smoking habits and behaviour is thus not possible because of the lack of means to do so. This paper proposes the use of a wearable device to monitor breathing signals of subjects. It is shown that the acoustic properties of a smoking breath are different from a non-smoking breath. To encapsulate these differences, several features from a breath segment are extracted and used with a simple classifier to automatically identify smoking breaths. The proposed algorithm detected smoking and non-smoking breaths with acerage accuracy of 66% and 99% respectively.**

I. Introduction

Smoking is a well documented cause of multiple health problems. An improved understanding and increased awareness of its dangers have reduced smoking prevalence in the UK from 46% of the population in 1974 to 19% in 2013 [1]. However, despite this progress, smoking remains the number one cause of preventable death with over 100,000 deaths in UK [2] each year, and 6 million [3] worldwide, as a result of smoking-related illnesses. Further, there are economic costs associated with smoking as well. The National Health Service (NHS) in the UK spends 2 billion each year on smoking-related medical issues [4].

The long-term effects of smoking have been studied in numerous ways by analyzing its impact on the health of individuals. The most common method of recording smoking behaviour is by using targetted surveys and questionnaires that are carried out by researchers or doctors. However, these have proven to be an ineffective method of data collection with subjects constantly under-reporting their tobacco consumption [5]. As an illustrative example, in one study, subjects reported smoking only 55% of their pre-study daily consumption when logging data into a PDA [6]. Another popular method to monitor smoking invovles analysing the expired air carbon monoxide (EACO). This provides a non-intrusive, low cost, and immediate evaluation of carbon monoxide levels which correlate highly with smoking. However, carbon monoxide has a very short half-life, and thus measurements must be taken frequently (and hence are burdensome) [7]. Other techniques include blood sampling and urine analysis [8] and provide detailed information about many smoking-related indicators (like nicotine or carboxyhemoglobin). However, these are expensive and must be undertaken at a clinic or hospital making them unfeasible for continuous long-term monitoring.

Even though a direct connection between smoking and a wide range of cancers has been established, it has been difficult to assess to what extent smoking affects cancer incidence, as well as survival prospects. This is because current methods of monitoring provide limited and subjective data related to smoking [9]. A much deeper understanding of the effects of smoking would be achievable by continually tracking smoking from disease diagnosis until death or survival at and old age [10]. To make this possible, many techniques to detect and monitor smoking behaviour are being explored explored by different researchers. These include the use of kinetic sensors to measure movements of the arm [11] or chest [12], [13] during respiration, smoke-through devices [14], and video processing [15]. However, while most of these do provide extensive information about smoking cycle, they are not feasible for long-term use. This is because monitoring over long period of time is only possible by creating wearable devices that can be used by subjects without requiring adjustments to their lifestyle.

Based on the above literature, the need to monitor smoking behaviour and its effect on human health in more detail is clearly evident. This paper looks at the feasiblity of using a wearable acoustic sensor to achieve this monitoring. The hypethesis behind this work is that the acoustic properties of a breath will be different at the time of smoking. By being able to characterise these differences, it should be possible to detect the time periods when a person is smoking. Section II describes the different features that have been extracted to highlight differences between smoking and non-smoking breaths. These features are then used to develop an algorithm for automatic classification of smoking breaths. The results of this algorithm are detailed in Section III followed by a discussion of its limitations and potential for future work in Section IV.

II. Material and Methods

A. Data Acquistion

In order to develop the algorithm, acoustic signals were recorded from two different subjects at different times. This recording was performed using a wearable device that has been designed to monitor breathing [16]. It consists of a microphone and is attached at the suprasternal notch of the subjects using a double-sided adhesive.. A total of six recordings were obtained from the two subjects each lasting the duration of smoking one cigarette. At the time

I. T. U. Echebarría, S. A. Imtiaz, M. Peng and E. Rodriguez-Villegas are with the Wearable Technologies Lab, Electrical and Electronic Engineering Department, Imperial College London, United Kingdom. Email: ({it812,anas.imtiaz,ming.peng11,e.rodriguez}@imperial.ac.uk).

of recording, the subjects were also asked to annotate the smoking periods.

B. Algorithm Overview

Fig. 1 shows a high-level overview of the proposed smoking behaviour detection algorithm. It consists of three main stages. The first stage is *pre-processing* where basic information about a breath such as its start time and length are determined. The next stage is *feature extraction* where the characteristics determined at the previous stage are used to extract five different features. These features are carefully chosen to represent relevant traits of smoking behaviour, such as longer expirations or a reduced breathing rate. The last stage is of *classification* where the different features are combined to determine if the breath being processed corresponds to the subject smoking. The output of the algorithm is a classification of individual breaths as smoking or non-smoking.

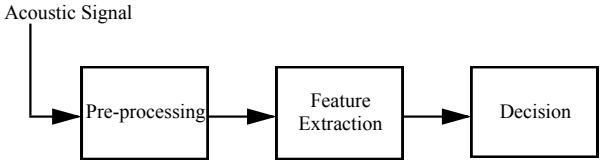


Fig. 1: *An overview of the smoking breath classification algorithm.*

C. Pre-processing

1) Breath Start Detection: The start of a breath is detected by following the method described in [17] which uses a series of filters and a segmentation algorithm to detect a valid breath segment. This is important since all of feature extraction at the next stage is performed on individual breath segments and hence accurate detection of these segments is critical to the overall performance of the algorithm proposed in this paper. The method described in [17] has been successfully validated in a clinical study for a different application [16].

2) Envelope Extraction: The envelope of the signal is extracted by taking the absolute value of the analytic representation of the original signal that is obtained by computing the Hilbert transform. It is more useful than the original representation of the signal because it gives a cleaner view of signal characteristics such as peak amplitudes and rise times. The envelope is used later to extract four other features in the next stage.

3) Breath Length Detection: The start of a breath is already computed earlier, hence to get its length, the end point of the breath needs to be determined. The end of a breath was obtained by using a method based on the instantaneous gradient of the envelope signal. While a person is breathing, the instantaneous gradient of the acoustic signal was found to be larger gradually falling off towards the end of each breath. This gradient signal was first low pass filtered to get a smoother representation of the gradient. The amplitude of this signal was then tracked such that the end of breath was marked if it fell below a certain threshold (after rising). The threshold itself was established by analysing data from multiple breath samples. Fig. 2 shows an example where the end of a breath is established using the instantaneous gradient approach.

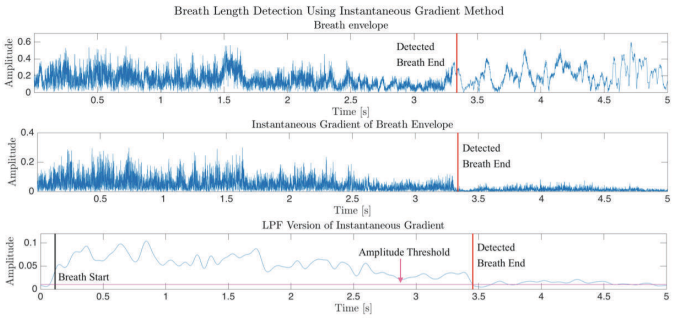


Fig. 2: *Detecting end of a breath segment by thresholding the instantaneous gradient of the acoustic signal.*

D. Feature Extraction

Fig. 3 illustrates the different features which are used in the decision stage in order to determine whether or not a breath is a smoking inhalation or expiration. Note that the breath lengths obtained in the pre-processing stage are also used as an independent feature. The calculation of the rest of the features (except breath separation) is more involved and needs as inputs the signal envelope, the breath start times and the breath lengths.

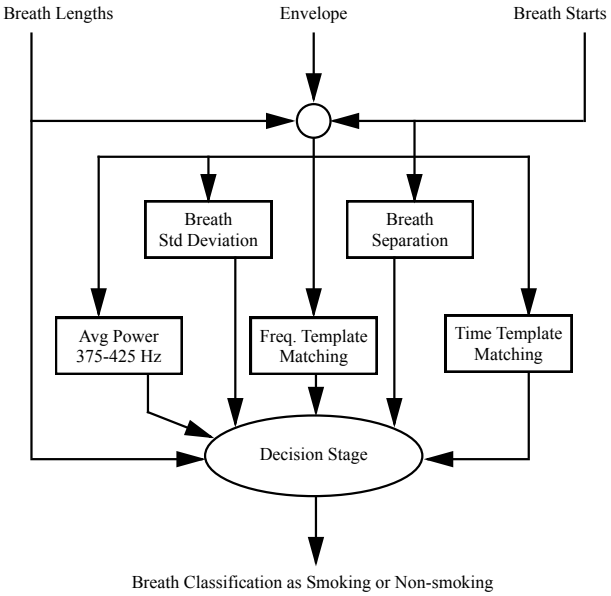


Fig. 3: *An overview of the feature extraction stage where different features are calculated to be fed into the decision stage.*

1) Time-Domain Template Matching: It is anticipated that the shape of the breathing signal will differ when a person is smoking. Hence, by comparing the breathing signals with

the shape of waveform expected during smoking, it should be possible to classify a breath as smoking or non-smoking.

Of the six available recordings, two were used to study the temporal shape of the breathing signal while smoking. By listening to the breaths and qualitatively analyzing their characteristics, a set of the most representative smoking breaths was selected. From this set, the most common waveform shapes were saved as the templates to compare against new data for classification. Normalized cross-correlation of the new breath against the saved templates is performed to determine the extent of similarity between the two. It was found that the template matching method is generally better at detecting smoking expirations than inhalations, as expirations have a more repeatable shape.

A cross-correlation coefficient close to 1 indicates that the template and the analyzed breath are similar, while a lower value implies that the two sequences are uncorrelated. However, since the breath start detection algorithm does not always identify as the breath start the same position as in a breath template, the maximum normalized cross-correlation value is determined for each breath at several different lags.

2) Frequency-Domain Template Matching: An analysis of the power spectral density (PSD) of the smoking and non-smoking breaths revealed that there were small but consistent differences in their general shapes. In order to use these differences, a frequency-domain template matching method was used similar to the time-domain template matching method. The normalized cross-correlation between the PSDs of a saved smoking-breath template and that of each breath showed higher correlation with the smoking breaths. The mean correlation coefficient in case of smoking breaths was 0.6, as opposed to 0.3 for non-smoking breaths.

3) Average Power Between 375Hz and 425Hz: The average signal power around 400 Hz was found to be consistently higher in smoking breaths compared to the non-smoking breaths. It was observed that this rise in power generally starts at 375 Hz and falls off again at 425 Hz. Hence, this feature can be used to distinguish between the two breath types. The larger the average power, the greater the likelihood of the analyzed breath being a smoking breath. The average power within this frequency band must be normalised with respect to the total power to perform a fair comparison. This is because in certain cases, for example, a loud breath, the power at all frequencies will be greater than a quiet breath, and hence its average power across the 375 Hz to 425 Hz range will be higher regardless of whether it is a smoking or a non-smoking breath.

4) Breath Separation: As part of studying smoking behaviour, it was found that smokers tend to keep smoke in their mouth for a longer time than the typical expiration cycle. As a result the expiration time during smoking increases and the time interval between the start of two consecutive breaths is longer. This interval, termed as breath separation, is used as one of the characteristic features. The average separation between smoking breaths was found to be just under 6 seconds, which is significantly higher than the non-smoking breaths (less than 2.5 seconds). Note that the due to this, the breathing rate also tends to decrease at the time of smoking but it is difficult to compute the instantaneous rate hence the use of separation between breaths, particularly when classifying individual breaths, is more appropriate.

5) Breath Standard Deviation: Some of the recordings in the test dataset consisted of speech activity during both smoking and non-smoking breaths. It is important to study the effect that speech will have on smoking activity. To account for this, the standard deviation of the amplitude within each breath sample is computed. The standard deviation for smoking breaths was found to be low even for those breaths that were surrounded by speech activity. A relatively high standard deviation is hence a clear indicator of behaviour other than smoking.

E. Decision

Once all the features for a breath are calculated a smoking score is computed as follows.

$$SmokingScore = \frac{L \times S \times P}{\sigma} \times (TDTM + FDTM) \quad (1)$$

In this equation, L is the length of a breath in seconds, S is the separation between breaths, P is the normalized power in 375-425 Hz band, σ is the standard deviation within a breath while $TDTM$ and $FDTM$ are the outcomes of time-domain and frequency-domain template matching respectively. If the smoking score is above a fixed threshold, then the breath is classified as smoking, otherwise it is marked as a non-smoking breath.

III. Results

The algorithm described in the previous section examines each breath and classifies it as either smoking or non-smoking. Its performance is thus assessed based on its ability to detect smoking breaths accurately. For this, the algorithm output was compared against the manually marked breaths across all six recordings. The algorithm was further tested by using a long recording without any smoking activity to test its ability of rejecting non-smoking breaths.

The performance of the algorithm is characterised by its *sensitivity*, which represents the fraction of smoking breaths correctly detected, and its *specificity* which refers to the fraction of breaths correctly rejected as non-smoking.

$$Sensitivity = \frac{TP}{TP + FN} \quad (2)$$

$$Specificity = \frac{TN}{TN + FP} \quad (3)$$

In the equations above, a detection is considered to be a true positive (TP) when it is correctly identified by the algorithm as a smoking breath and true negative (TN) when correctly rejected by the algorithm as a non-smoking breath. A false positive (FP) represents incorrect classification as smoking breath while a wrong identification of non-smoking breath is considered a false negative (FN).

Table I shows the performance of the algorithm on each recording. It can be seen that the algorithm achieves a very high level of specificity i.e. it is highly capable of rejecting non-smoking breaths correctly. It also achieves a sensitivity of over 70% in four of the six recordings with smoking breaths. The smoking breath detection ratio is, however, particularly low in recording number 5 which consists of loud speech that is causing difficulties in identification. Further, in recording 7 there are no smoking breaths, and only 2 out of 548 breaths are misclassified in this case with 546 correct detections of non-smoking breaths.

TABLE I: *Performance of the smoking breath detection algorithm on seven different recordings.*

No	B_t	B_s	TP	FN	TN	FP	Sen(%)	Spe(%)
1	333	19	14	5	310	4	73.7	98.7
2	186	19	11	8	166	1	57.9	99.4
3	216	4	4	0	211	1	100.0	99.5
4	95	17	12	5	75	3	70.6	96.2
5	155	7	1	6	147	1	14.3	99.3
6	78	13	10	3	62	3	76.9	95.4
7	548	0	0	0	546	2	-	99.6

B_t - number of breaths in a recording.
B_s - number of breaths manually marked as smoking breaths.

IV. Discussion

In this paper, an algorithm for automatic identification of *smoking breaths* is presented. The algorithm uses acoustic signals from a wearable device that is attached at the suprasternal notch. These signals are analysed to detect individual breaths that are subsequently classified into smoking or non-smoking breaths. Using 7 different recordings of breathing signals, the algorithm was able to identify 52 out of total 79 smoking breaths. Overall, the results demonstrate that the novel approach being proposed in this paper is certainly feasible to detect smoking breaths with room for further improvements. These imporvements can be achieved by looking at other methods to reject speech and other motion artefacts as well as exploring other features with better discriminatory abilities.

The work presented in this paper is critical towards realising a complete smoking monitoring system. By having the ability to detect smoking breaths, more details about a smoking session can be extracted. For example, by being able to detect the number of smoking breaths in a given interval throughout the day, information about the quantity of cigarettes smoked can be extracted. Currently, this relies on questionnaires only and any improvement on the existing method would be substantial. It is also possible to determine the total time spent smoking with details such as the inhalation depth and the time during which smoke is held in the mouth. These details, amongst others, can help provide new insights into the effects of smoking at different stages of disease progression. It can also help to monitor certain treatments where smoking can be detrimental by reducing chances of survival [18].

A complete smoking monitoring system would help obtain structured and unbiased data, providing objective evidence of smoking behaviour while conducing large-scale research studies. However, crucial to the success of such a system is the requirement of a highly accurate smoking detection algorithm that works well with signals acquired using a wearable device. Such a device needs to be non-invasive and low-cost without interfering with daily lifestyle of people otherwise they would be opposed to using it over long periods of time. The signal acquisition device used in this paper already fulfills these criteria and has been successfully used in other applications [16], [19]. Although more work is needed to create an algorithm that provides accurate detection of smoking breaths, the work presented in this paper confirms that the approach based on acoustic signals from a wearable sensor is viable as well as practical.

References

[1] Office for National Statistics. (2017) Adult smoking habits in Great Britain: 2014. [Online]. Available: https://www.ons.gov.uk/.

[2] ASH: Action on Smoking and Health. (2017) Smoking and Disease. [Online]. Available: https://www.ash.org.uk/.

[3] World Health Organization. (2017) WHO report on the global tobacco epidemic, 2011: Warning about the dangers of tobacco. [Online]. Available: https://www.who.int/.

[4] ASH: Action on Smoking and Health. (2017) The economics of tobacco. [Online]. Available: http://ash.org.uk/information-and-resources/fact-sheets/the-economics-of-tobacco/.

[5] Coultas *et al.*, "Discrepancies between self-reported and validated cigarette smoking in a community survey of new mexico hispanics," *Am Rev Respir Dis*, vol. 137, no. 4, pp. 810–814, 1988.

[6] Warthen and Tiffany, "Evaluation of cue reactivity in the natural environment of smokers using ecological momentary assessment," *Experimental and Clinical Psychopharmacology*, vol. 17, no. 2, p. 7077, 2009.

[7] Rashid *et al.*, "The pitfalls of measuring changes in smoking habits," *Tropical Doctor*, vol. 30, no. 3, pp. 160–161, 2000.

[8] Benowitz *et al.*, "Carcinogen exposure during short-term switching from regular to "light" cigarettes," *Cancer Epidemiology and Prevention Biomarkers*, vol. 14, no. 6, pp. 1376–1383, 2005.

[9] Warren *et al.*, "Smoking at diagnosis and survival in cancer patients," *International Journal of Cancer*, vol. 132, no. 2, pp. 401–410, 2013.

[10] Gritz, Dresler, and Sarna, "Smoking, the missing drug interaction in clinical trials: Ignoring the obvious," *Cancer Epidemiology and Prevention Biomarkers*, vol. 14, no. 10, pp. 2287–2293, 2005.

[11] Raiff *et al.*, "Laboratory validation of inertial body sensors to detect cigarette smoking arm movements," *Electronics*, vol. 3, no. 1, pp. 87–110, 2014.

[12] Sazonov, Lopez-Meyer, and Tiffany, "A wearable sensor system for monitoring cigarette smoking," *Journal of Studies on Alcohol and Drugs*, vol. 74, no. 6, pp. 956–964, 2013.

[13] Ali *et al.*, "mpuff: Automated detection of cigarette smoking puffs from respiration measurements," in *ACM IPSN*, Beijing, 2012.

[14] Kolonen *et al.*, "Puffing behavior during the smoking of a single cigarette in a naturalistic environment," *Pharmacology Biochemistry and Behavior*, vol. 41, no. 4, p. 701706, 1992.

[15] Stitt and Kozlowski, "A system for automatic quantification of cigarette smoking behavior," in *IEEE EMBC*, New York, August 2006, pp. 4771–4774.

[16] Rodriguez-Villegas, Chen, Radcliffe, and Duncan, "A pilot study of a wearable apnoea detection device," *BMJ Open*, vol. 4, no. 10, 2014.

[17] Corbishley and Rodriguez-Villegas, "Breathing detection: Towards a miniaturized, wearable, battery-operated monitoring system," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 1, pp. 196–204, 2008.

[18] Browman *et al.*, "Influence of cigarette smoking on the efficacy of radiation therapy in head and neck cancer," *New England Journal of Medicine*, vol. 328, no. 3, pp. 159–163, 1993.

[19] Chen, Imtiaz, Aguilar-Pelaez, and Rodriguez-Villegas, "Algorithm for heart rate extraction in a novel wearable acoustic sensor," *IET Healthcare Technology Letters*, vol. 2, pp. 28–33, 2015.