

A Bottom-up method Towards the Automatic and Objective Monitoring of Smoking Behavior In-the-wild using Wrist-mounted Inertial Sensors

Athanasios Kirmizis, Konstantinos Kyritsis and Anastasios Delopoulos

Abstract—The consumption of tobacco has reached global epidemic proportions and is characterized as the leading cause of death and illness. Among the different ways of consuming tobacco (e.g., smokeless, cigars), smoking cigarettes is the most widespread. In this paper, we present a two-step, bottom-up algorithm towards the automatic and objective monitoring of cigarette-based, smoking behavior during the day, using the 3D acceleration and orientation velocity measurements from a commercial smartwatch. In the first step, our algorithm performs the detection of individual smoking gestures (i.e., puffs) using an artificial neural network with both convolutional and recurrent layers. In the second step, we make use of the detected puff density to achieve the temporal localization of smoking sessions that occur throughout the day. In the experimental section we provide extended evaluation regarding each step of the proposed algorithm, using our publicly-available, realistic Smoking Event Detection (SED) and Free-living Smoking Event Detection (SED-FL) datasets recorded under semi-controlled and free-living conditions, respectively. In particular, leave-one-subject-out (LOSO) experiments reveal an F1-score of 0.863 for the detection of puffs and an F1-score/Jaccard index equal to 0.878/0.604 towards the temporal localization of smoking sessions during the day. Finally, to gain further insight, we also compare the puff detection part of our algorithm with a similar approach found in the recent literature.

I. INTRODUCTION

According to the World Health Organization (WHO), smoking is the leading public health problem worldwide, resulting in millions of preventable deaths each year and is responsible for a number of serious chronic diseases (e.g., hypertension, atherosclerosis, cancer) [1].

Globally, male and female smokers have their life expectancy reduced by 13.2 and 14.5 years, respectively [2]. Moreover, at least half of all smokers worldwide die prematurely from smoking [1]. It is important to emphasize that smoking is not only harmful to smokers themselves, but it is also a major risk factor for passive smokers [3].

The modernization of societies has ignited a recent trend that promotes a lifestyle in which smoking is considered an outdated habit. More and more people are taking up a sport, or are beginning to follow a healthy eating regime [4]. An important role for the engagement of people with these healthy habits plays the technology that is constantly evolving and gets integrated into everyday life. The rapid growth of portable and wearable devices has brought with it a great increase in applications that help people develop and maintain a healthy lifestyle. From tracking meals and calories to measuring physical activity or sleep, the applications now available to users are multiple, easy to use, and unobtrusive

All authors are with the Multimedia Understanding Group, Information Processing Laboratory, Aristotle University of Thessaloniki, Greece

[5]. However, the objective monitoring of smoking behavior is still an open research problem. Research shows [6] that tailored feedback to the smoker can greatly facilitate the reduction or even permanent cessation of smoking.

Several works exist in the literature that approach the problem of smoking behavior monitoring using body-worn sensors [7]. The work presented in [8] suggests a method that combines the data from a wrist-mounted inertial measurement unit (IMU) sensor and a chest-worn respiratory inductive plethysmography (RIP) sensor towards the detection of smoking gestures. Evaluation using data from 6 daily smokers reveals a recall of 0.97. It should be mentioned, however, that devices as bulky as the RIP sensor are too obtrusive for the user to properly simulate the normal smoking behavior. The work of M. Shoaib *et al.* [9] proposes a two-step algorithm towards the detection of smoking events. In the first stage, the data are crudely classified, while at the second stage a rule-based correction of the first-stage classification is applied. The classifiers tested by the authors are random forest (RF), decision tree (DT) and support vector machine (SVM). For each classifier, a total of 36 features are extracted from the 3D accelerometer and gyroscope measurements. According to the authors, the second step of their algorithm corrects up to 50% of the misclassified samples. Evaluation is performed using their dataset of 11 participants with a total duration of 45 hours, where the authors achieved an F1-score of 0.83-0.94.

In our work, we propose the use of a smoking behavior model that is based on two fundamental components: a) the *puff* (also referred to as smoking gesture in the literature), defined as the series of hand movements that bring an active cigarette to the mouth with the purpose of smoking it and then back to rest, and b) the *smoking session*, defined as the act of consuming a cigarette. In particular, we model smoking behavior as a series of smoking sessions that occur during the day. Subsequently, each smoking session is modeled as a series of puffs. Figure 1 illustrates the adopted smoking behavior model. Furthermore, we suggest a two-step, bottom-up method towards the objective and automatic monitoring of smoking behavior using all-day, free-living IMU recordings from an off-the-shelf smartwatch. In the first step, we use an artificial neural network (ANN) with convolutional and recurrent layers to detect puffs during a smoking session. In the second step, we use the distribution of the detected puffs to localize the smoking sessions throughout the day.



Fig. 1: The proposed smoking behavior model. In this example, a subject performed six smoking sessions (blue) during the course of the day (grey), with each session containing a number of puffs (dark red).

II. DETECTION OF PUFFS

A. Data pre-processing

Let $\mathbf{s}(t) = [\mathbf{a}_x(t), \mathbf{a}_y(t), \mathbf{a}_z(t), \mathbf{g}_x(t), \mathbf{g}_y(t), \mathbf{g}_z(t)]^T$ represent the vector that contains the 3D acceleration and orientation velocity measurements for a moment t . Then, a complete recording of d_{tot} seconds can be represented by the $M \times 6$ signal $\mathbf{R} = [\mathbf{s}(1), \mathbf{s}(2), \dots, \mathbf{s}(M)]^T$, where $M = d_{tot} \cdot f_s$ is the length of the recording in samples and f_s is the sampling frequency in Hz.

Smoking cigarettes is a process that can be completed by using either hand (right or left) or, in some cases, a combination of both. In order to achieve uniformity among data from different participants, we consider the right hand as the reference and transform all left-handed smoking sessions using the *hand mirroring* process proposed by Kyritsis *et al.* [10]. Particularly, all recordings that are collected with the participant wearing the smartwatch on the left wrist \mathbf{R}_l , are transformed into \mathbf{R}_r by changing the direction of the first, fifth and sixth channels (i.e., \mathbf{a}_x , \mathbf{g}_y and \mathbf{g}_z) of \mathbf{R}_l .

Furthermore, accelerometer measurements also include the influence of the Earth's gravitational field. To attenuate this undesirable effect, a high-pass finite impulse response (FIR) filter is applied to each of the acceleration streams (i.e., the first, second, and third channels of \mathbf{R}), independently. Experimentally, we obtained satisfactory results with a cut-off frequency of 1 Hz and a filter length equal to 512 samples (which corresponds to $512/f_s$ seconds).

B. Training the puff detection model

Given a recording \mathbf{R} that corresponds to a smoking session, we extract training examples using a sliding window. More specifically, the sliding window has a length w_l that corresponds to 4.5 seconds ($4.5 f_s$ samples) and a step w_s that corresponds to 0.5 seconds ($0.5 f_s$ samples). We selected a window length equal to 4.5 seconds as it approximates the median puff duration in the SED dataset (Table I). Each extracted window \mathbf{W}_i has dimensions $(4.5 f_s) \times 6$.

TABLE I: Information regarding the SED and SED-FL datasets. Statistics were calculated using the raw data.

Dataset	SED		SED-FL	
	Smoking	Puffs	In-the-wild	Smoking
Number of instances	20	276	10	39
Mean (sec)	485.14	4.86	28202.14	525.33
Std (sec)	197.32	1.47	13484.42	301.81
Median (sec)	484.26	4.75	25919.27	462.80
Total (sec)	9702.88	1341.18	282021.38	20487.84
Total (hours)	2.69	0.37	78.33	5.69
Participants	11		7	

In order to train the network, each window \mathbf{W}_i needs to be associated with a label y_i that would indicate if the window corresponds to a puff or not ($y_i = \pm 1$, respectively). We use the following formula to perform the labeling process:

$$y_i = \begin{cases} +1 & \text{if } t_j^{gt} - \epsilon \leq t_i^{\mathbf{W}} \leq t_j^{gt} + \epsilon \\ -1 & \text{otherwise} \end{cases} \quad (1)$$

where t_j^{gt} is the moment at which the j -th puff ends (hand has returned to rest) according to ground truth (GT) and $t_i^{\mathbf{W}}$ is the timestamp associated with the right end of the i -th extracted window. Moreover, we select ϵ to be equal to 1.5 seconds. Figure 2 showcases the window labeling process.




Fig. 2: The proposed window labeling method.

The next step is to artificially augment the training set by simulating different positions of the smartwatch, that may occur *involuntarily* while wearing it, with respect to the subject's wrist. Specifically, we draw two numbers from a normal distribution with a mean and standard deviation equal to 0 and 10, respectively. These two numbers represent the *angles* that the smartwatch has rotated around the x (parallel to the subject's arm) and z (perpendicular to the screen of the smartwatch) axes. The transformation for each window \mathbf{W}_i is selected to be one of the following (with equal probability): a) rotation around x , b) rotation around z , c) rotation around x and then around z , or d) rotation around z and then around x . The motivation behind the augmentation step was the significant increase in the performance reported in [10].

The proposed model is a tuned-down version of the renowned VGG architecture [11]. In particular, our network includes a convolutional and a recurrent part. The convolutional part contains three 1D convolutional layers, with each of the first two followed by a max pooling layer with a decimation factor of 2. The convolutional layers have 32, 64 and 128 filters, with a size of 5, 3 and 3, respectively. All convolutional layers use a unary stride and the rectified linear unit (ReLU) as the non-linearity. The recurrent part of the network consists of a single long-short-term-memory (LSTM) layer with 128 cells and the sigmoid function as the activation of the recurrent steps. The output of the LSTM is propagated to a fully connected layer with a single neuron and the sigmoid activation function. In order to avoid overfitting, we apply dropout to the inputs of the

fully connected layer with a probability of 50%. The network minimizes the binary cross-entropy loss with the RMSProp optimizer, and uses a learning rate of 10^{-3} , a batch size of 32 and a number of 10 epochs. In a compact notation, the network can be written as Conv(32×5)-Pool(2)-Conv(64×3)-Pool(2)-Conv(128×3)-LSTM(128)-FC(1), where Conv(32×5) represents a convolutional layer with 32 filters and a filter size of 5, Pool(2) is pooling layer with a decimation factor of 2, LSTM(128) is an LSTM layer with 128 hidden cells and FC(1) is a fully connected layer with a single neuron.

C. Puff detection

By forwarding windows from a recording \mathbf{R} to the trained puff detection network (Section II-B), we obtain the predictions vector \mathbf{p} with length N . Essentially p_i is the probability that the i -th window \mathbf{W}_i is a puff and N represents the total number of extracted windows of length w_l and step w_s .

Puff detection is achieved by initially performing a local maxima search in \mathbf{p} , with a minimum distance between successive peaks equal to 10 samples. The next step is to discard peaks that are associated with a probability p_i that is lower than a threshold λ_p set to 0.8. Both the minimum distance between peaks and λ_p were selected by experimenting with a small part of the SED dataset. As a result, we obtain the set of detected puffs, $\mathcal{F} = \{f_1, \dots, f_K\}$, where f_i is the timestamp of i -th detected puff and K the total number of detected puffs. The process is illustrated in Figure 3.




Fig. 3: Detection of puffs given the probability vector \mathbf{p} (blue line). The ground truth puff durations (black line), local maxima peaks (red dots) and the λ_p threshold (light-red dashed line) are also depicted. In the figure’s example, the left-most and right-most peaks are rejected as they are below the threshold λ_p .

III. TEMPORAL LOCALIZATION OF SMOKING SESSIONS

The second step of the proposed algorithm aims at the temporal localization of smoking sessions that occur during a day. In our early experiments we observed that in all-day recordings the density of puffs is increased during a smoking session and reduced everywhere else. As a result, in the second step of our algorithm we take advantage of this observation and attempt to group the detected puffs into smoking session clusters using the density-based spatial clustering of applications with noise (DBSCAN) [12] algorithm.

More specifically, let \mathbf{R}' be an all-day, in-the-wild recording with dimensions $M' \times 6$, where $M' \gg M$. Next, we use the trained puff detection model (Section II-B) to produce the set of puff detection estimates \mathcal{F}' . Subsequently, we apply

clustering using DBSCAN on the set \mathcal{F}' using a minimum distance between clusters that corresponds to 250 seconds (as this is the minimum distance between consecutive smoking sessions in the SED-FL dataset) and a minimum number of points per cluster set to 4.

Each cluster that DBSCAN produces is then associated with the first and last timestamps of the detected puffs that belong to that specific cluster. This pair of timestamps corresponds to the start and end moments of a smoking session. Formally, the final output of the algorithm is the set $\mathcal{G} = \{C_1, \dots, C_L\} = \{[t_1^s, t_1^e], \dots, [t_L^s, t_L^e]\}$, where $[t_i^s, t_i^e]$ represents the start and end timestamps of the i -th detected smoking session. An example depicting the temporal localization of smoking sessions can be found in Figure 4.




Fig. 4: Figure depicting an example of how the smoking session clusters (red line) are formed using the set of detected puffs \mathcal{F}' (blue dots). The ground truth smoking session durations (as annotated by the participants are also depicted (black line).

IV. EXPERIMENTS AND EVALUATION

A. Datasets

In order to fine-tune and evaluate our method we collected two datasets. The SED dataset was captured in semi-controlled environments (e.g., private residences or cafes) and contains a single smoking session per recording. On the other hand, the SED-FL dataset was captured under in-the-wild conditions and contains all-day recordings that include smoking sessions and other daily activities (e.g., working, eating). Inertial data were collected using a Mobvoi TicWatch E smartwatch at a sampling rate f_s equal to 50 Hz.

The SED dataset consists of 11 subjects performing 20 smoking sessions, with a total duration of 2.69 hours. The SED-FL dataset consists of 10 all-day sessions from 7 subjects, with a total duration of 78.3 hours (Table I). Three of the subjects participate in both datasets. It should be emphasized that we asked from the subjects to smoke naturally; as a result, they were free to engage in a discussion or perform additional activities (two instances are depicted in Figure 5). All subjects were already smokers and signed an informed consent prior to their participation. In order to label the data in SED, we recorded the smoking sessions using the camera from a typical smartphone. To produce the GT for the all-day, in-the-wild sessions of SED-FL, a smartwatch application was developed that enabled subjects to easily note the start and end timestamps of their smoking sessions. It is worth noting that both datasets deal with the consumption of tobacco using *cigarettes*; no electronic cigarettes (also known as vaping devices), pipes or heated tobacco products were used. Both datasets are publicly available at <https://mug.ee.auth.gr/smoking-event-detection/>.



Fig. 5: Two participants from the SED dataset.

B. Experiments

We conducted two series of experiments. In the first experiment (EX-I), we evaluate the puff detection performance using the SED dataset. Moreover, we compare the performance of the proposed puff detection approach with the method proposed in [9]. For the second experiment (EX-II), we evaluated the smoking session temporal localization performance using the SED-FL dataset. Both EX-I and EX-II are performed in a leave-one-subject-out (LOSO) fashion.

C. Evaluation

In order to measure the puff detection performance (EX-I), we apply the strict evaluation scheme presented in [10]. An example of the evaluation scheme is presented in Figure 3. Essentially: a) only the first detected puff within the duration of a GT interval is considered as a true positive (TP), all subsequent ones count as false positives (FP), b) GT intervals without detections count as false negative (FN) and c) predictions outside GT intervals are considered as FP. It should be noted that the evaluation scheme of [10] cannot calculate true negatives (TN). However, at a window level we can effectively measure TP/FP/FN and TN; i.e., by comparing the label y_i of each extracted window W_i with the GT. As a result, we can calculate the weighted accuracy metric, defined as $\frac{TP \cdot w + TN}{(TP + FN) \cdot w + FP + TN}$, using a weight w equal to 7.27 (total time spend during smoking sessions divided by the total time spend during puffs).

Regarding EX-II, a detected smoking session is considered a TP if its middle timestamp (calculated as $\frac{t_i^s + t_i^e}{2}$) is within the duration of a GT interval; in any other case is considered a FP. In addition, GT intervals without detections are considered as FN. Figure 4 illustrates the aforementioned evaluation scheme. Similar to EX-I, we also calculated the weighted accuracy for EX-II using a weight equal to 13.76. Finally, we calculated the Jaccard Index (JI), defined as $\frac{|A \cap B|}{|A \cup B|}$, where A and B are the intervals of the true and the predicted smoking sessions, respectively.

V. RESULTS

The obtained results showcase the high potential of our approach; both towards the detection of individual puffs (upper part of Table II), as well as for the temporal localization of smoking events in-the-wild (lower part of Table

TABLE II: EX-I-II results.

Experiment	Algorithm	W. Acc	Prec	Rec	F1-score	JI
EX-I	[9] with RF	0.894	0.834	0.840	0.837	N/A
	[9] with SVM	0.876	0.836	0.815	0.825	N/A
	[9] with DT	0.730	0.478	0.960	0.638	N/A
	Proposed	0.915	0.921	0.811	0.863	N/A
EX-II	Proposed	0.968	0.837	0.923	0.878	0.604

II). More specifically, regarding EX-I, the proposed approach achieves a weighted accuracy of 0.915 and an F1-score of 0.863 using the stricter evaluation scheme of [10] (against 0.894 and 0.837 obtained by [9]). Concerning EX-II, our approach achieves an F1-score/weighted accuracy/JI equal to 0.878/0.968/0.604 which indicates that smoking sessions can be effectively detected under in-the-wild conditions.

VI. CONCLUSIONS

In this paper we present a two-step, bottom-up method towards the in-the-wild monitoring of smoking behavior. LOSO experimental results using our realistic SED and SED-FL datasets reveal the high potential of our approach towards the detection of puffs and the localization of smoking sessions during the day, under in-the-wild conditions.

VII. ACKNOWLEDGMENTS

The work leading to these results has received funding from the EU Commission under Grant Agreement No. 965231, the REBECCA H2020 project (<https://rebeccaproject.eu/>).

REFERENCES

- [1] W. H. Organization *et al.*, *WHO report on the global tobacco epidemic, 2017: monitoring tobacco use and prevention policies*. World Health Organization, 2017.
- [2] C. for Disease Control, P. (CDC, *et al.*, “Annual smoking-attributable mortality, years of potential life lost, and economic costs—united states, 1995–1999,” *MMWR. Morbidity and mortality weekly report*, vol. 51, no. 14, pp. 300–303, 2002.
- [3] R. Otsuka *et al.*, “Acute effects of passive smoking on the coronary circulation in healthy young adults,” *Jama*, vol. 286, no. 4, 2001.
- [4] G. W. Institute, “Global wellness economy monitor,” 2018.
- [5] J. H. West *et al.*, “There’s an app for that: content analysis of paid health and fitness apps,” *Journal of medical Internet research*, vol. 14, no. 3, p. e72, 2012.
- [6] T. Lancaster and L. F. Stead, “Self-help interventions for smoking cessation,” *Cochrane database of systematic reviews*, no. 3, 2005.
- [7] M. H. Imtiaz *et al.*, “Wearable sensors for monitoring of cigarette smoking in free-living: A systematic review,” *Sensors*, 2019.
- [8] N. Saleheen *et al.*, “puffmarker: a multi-sensor approach for pinpointing the timing of first lapse in smoking cessation,” in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2015, pp. 999–1010.
- [9] M. Shoaib *et al.*, “A hierarchical lazy smoking detection algorithm using smartwatch sensors,” in *2016 IEEE 18th International Conference on e-Health Networking, Applications and Services (Healthcom)*. IEEE, 2016, pp. 1–6.
- [10] K. Kyritsis, C. Diou, and A. Delopoulos, “A data driven end-to-end approach for in-the-wild monitoring of eating behavior using smartwatches,” *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 1, pp. 22–34, 2020.
- [11] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [12] M. Ester *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise.” in *Kdd*, vol. 96, no. 34, 1996, pp. 226–231.