Behavioural Smoking Identification via Hand-Movement Dynamics

Mohamed Benouis
Computer Science
M'sila University
M'sila, Algeria
mohamed.benouis@univ-msila.dz

Maryam Abo-Tabik
Computing and Mathematics
Manchester Metropolitan University
Manchester, UK
maryam.a.abo-tabik@stu.mmu.ac.uk

Yael Benn
Psychology
Manchester Metropolitan University
Manchester, UK
y.benn@mmu.ac.uk

Olivia Salmon
Psychology
Manchester Metropolitan University
Manchester, UK

Alex Barret-Chapman
Psychology
Manchester Metropolitan University
Manchester, UK

Nicholas Costen
Computing and Mathematics
Manchester Metropolitan University
Manchester, UK
n.costen@mmu.ac.uk

***Abstract*—Smoking is a commonly observed habit worldwide, and is a major cause of disease leading to death. Many techniques have been established in medical and psychological science to help people quit smoking. However, the existing systems are complex, and usually expensive. Recently, wearable sensors and mobile application have become an alternative method of improving health. We propose a human behavioural classification based on the use of a one-dimensional local binary pattern (LBP), combined with a Probabilistic Neural Net (PNN) to differentiate smoking from other movements as measured from a wearable device. Human activity signals were recorded from two sets of 6 and 11 participants, using a smart phones equipped with an accelerometer and gyroscope mounted on a wrist module. The data combined structured and naturalistic scenarios. The proposed architecture was compared to previously studied machine learning algorithms and found to out-perform them, exhibiting ceiling level performance.**

***Index Terms*—smoking behaviour, hand movements, acceleration data, machine learning, 1-D Local Binary Patterns.**

I. Introduction

Human activity recognition has received much attention as it is considered one of the most efficient methods for improving health-care by monitoring and prompting improvement in well-being [1]. Some notable systems include a component that detects human gestures and automatically distinguishes between complex activity in real life situations. These systems recognize human motions based on sensor values or self-reported data, store them as behavioural features that are adapted to the available resources (memory, and battery life), and then use the data to automatically predict human activity. However, although such applications have been able to detect a few human activities, difficulties arise when dealing with natural human activity, which has higher levels of complexity. As such, methods for automatically identifying and predicting unique hand movement gestures among a large dataset from sensors are crucially required.

In this paper, we assume that the important information to be detected can be learned from the users' accelerometers data, and propose a method for classifying human activity such as smoking and non-smoking by applying a well-known machine-learning techniques. In the proposed method, we first extract feature values from sequential sensor data, and then train a detector using important features, with training labels of the user's movements. Using the 1D-local binary pattern extraction [2], allows the different movements that are recorded by the user's Smartphone to be classified as either smoking or non-smoking, independent of the presence of simultaneous, confounding behaviours such as eating. Furthermore, because this method uses the feature vectors from sensors instead of self-report as input, the method has the advantage that the model can be easily trained and classify the activity data with high accuracy. To evaluate the proposed method, the data set collected from the University of Twente [3] was used. As a result, for all of the sensor data, the accuracy reached 100%; the accuracy on the more difficult dataset collected for this was lower (at 85%), but still higher than that for other techniques.

A. Related work

Work related to the present study includes approaches that employ the inertial measurements units (IMU) of Smartphones to detect smoking. Several studies in the literature focused on detecting daily human activities [4] in general, and more specifically smoking detection was shown to be feasible when using sensor data acquired from custom [5, 6] or commercial [4, 7] wrist-worn devices. In this work, we focus on detecting smoking gestures using accelerometer data.

Ecological Momentary Assessment (EMA) techniques use mobile devices to continuously collect data about experiences and behaviours in the natural environment [8]. As there is lower likelihood of memory errors or other biases, EMAs have enabled the collection of ecologically valid data. Many studies have used EMAs to model smoking behaviour and/or predict smoking lapses [3, 9, 10, 11, 12, 13, 14] resulting in important insights into smoking behaviour patterns.

For example, two studies using EMA by Businelle et al [9, 10], have been able to identify imminent risk for smoking lapse among smokers seeking smoking cessation treatment. Estimation of the risk of smoking lapse was significantly improved by weighting six risk variables (i.e., urge to smoke, stress, recent alcohol consumption, interaction with other smokers, cessation motivation, and cigarette availability) [10]. In a follow-up [9], EMA was used to deliver tailored messages to individuals to help them to maintain abstinence from smoking.

Most recently, a study by Schick et al [5] has evaluated the effectiveness of a smart-phone application which integrates time and space (measured via GPS) with EMA data. The collected data was then sent to the user's General Practitioner (GP) to review and agree a quitting plan with the smoker. The time and places that the smokers are most likely to smoke was also predicted, in order to implement better delivery of the support messages. Although the results demonstrate that sending timely messages helped smokers quit smoking and prevented relapse, the specific application still relied on self-reporting, and still involved face-to-face meeting with the GP.

In fact, to date all studies using EMA have relied on self-reporting of smoking events by participants. Self-reporting is a problematic weak point, as the reliability of information collected is still at the mercy of users, who may forget or prefer to not report smoking events [15] for reasons such as self-enhancement. In addition, monitoring behaviour is a known behaviour change technique [16], and as such monitoring smoking events is likely to influence the smoker's behaviour, hence directly challenging the ecological validity assumed in self-reported EMA. Therefore, we suggest that reliability of self-reporting style EMA could be dramatically improved if mobile apps could use EMA without relying on user input, but instead detect smoking behaviour automatically, and combine it with environmental data in order to make lapse- prediction and help smokers. In addition, it will allow for a more effective evaluation of the intervention component, without the 'contamination' of behaviour monitoring.

Automatic detection of smoking behaviour presents its own challenges. When measuring smoking related movement, wearing the accelerometer on the non-dominant arm is likely to have great impact. Furthermore, being able to uniquely identify smoking related movement from other similar movement such as eating or drinking is not a trivial task.

Sazonov et al [11] have proposed a method based on the use of wearable sensors to detect and characterize cigarette smoke inhalations through monitoring of breathing and hand-to-mouth gestures. Results indicated that smoking results in a unique breathing pattern that is highly correlated with hand-to-mouth cigarette gestures and as such, the combined signals could be effectively used to detect smoking behaviour. However, the wearable sensors included not only the hand gesture sensor (attached to the wrist) but also a respiration sensor, which is a thoracic and abdominal respiration bands of the respiratory plethysmograph, as well as a chest-mounted antenna that combined chest movement with signals from the hand gesture sensor. This makes it a less suitable solution for using in a naturalistic environment.

Others have managed to make good predictions of smoking behaviour using just arm-movement data [7, 12, 17] including commercially available smart watches [3, 13, 14, 18]. Perhaps the best prediction rate were made by Parate et al [7] who designed a low-power wristband device, which contains a 9-axis inertial measurement unit, fusing information from an accelerometer, gyroscope, and compass to provide 3D orientation of the wrist. This gave a smoking detection accuracy of 95.7%. However, their analysis partially relied on an elbow-worn sensor, while those using commercially available smart-watches generally report lower accuracy (e.g., Alharbi and Farrahi [19]: 92–96%, Cole et al. [18]: 85%–95%). Even when a similar algorithm to Parate et al was used with a smart watch [20], precision was limited to 86% and recall to 71%.

B. Algorithmic issues

Most human activity classification approaches seek to extract the relevant information associated with a particular time (t) (i.e., an action event). However, the action definitions usually do not specify the duration time as well as the discriminate motion features needed to distinguish between them and other events. In other words, the size of the window for each action differs and this is considered one the most serious issues facing the times series segmentation field.

Following the approach of Shoaib et al. [3], this study investigates the efficiency of local binary patterns as a single layer to minimize the confusion between smoking and other activities. A three-stage analysis pipeline was used, where raw motion data are isolated from other information, and a feature extraction procedure applied to extract the relevant information. These features are fed to three classifiers, KNN, SSRC and PNN, to identify the hand movements which corresponding with hand-to-mouth gestures.

Motion features are described with LBP histograms, which look for patterns of values occurring in short as well as in long periods. The optimal window size has not been resolved and the most of studies have used manual selection of windows size. Here, different window sizes were considered and and a window size of 18 samples (180 ms) was adopted.

II. METHODS

Two data sets were used; one collected for the study, and one used by experimenters previously. This ensured that two forms of data were considered; the new dataset had a large, balanced set of smoking gestures (many combining smoking and other activities), with the participant also performing a range of similar, non-smoking actions. The older set was more naturalistic, with data collected over longer periods and with "genuine", rather than prompted or mimed smoking actions.

A. Data collection procedure

Full ethical approval was granted from the University ethics committee. Data was collected from 6 healthy smokers (at least 5 cigarettes a day for at least 6 months; 3 females, none smoked e-cigarettes), aged between 18-55 years (mean 33.6, SD 11.3). Following informed consent, participants performed 16 activities, with an Huawei P smartphone attached to their dominant arm's wrist using an armband. A mobile application was installed on the phone, and facilitated the recording and labelling of data. Before each activity, the researcher selected the appropriate options on the app interface.

Data was collected from 9 sensors (see Table I), and was stored with the appropriate labels on an internal SQLite database. The set of activities included 4 short activities lasting 30 seconds (fast drinking - i.e., from a bottle while standing; hair brushing, teeth brushing, and opening an umbrella while standing still) while the other 12 activities (drinking while sitting, eating with knife and fork, eating snack from bag like crisp, sitting, walking, standing still, reclining, writing, typing, reading, making a phone call, and driving using a simulator) lasted 60 sec. All activities were performed once without smoking, and then together with mimicking either smoking or interacting with a mobile phone or both.

For most actions, participants were given the choice to perform the activity while sitting or standing in order to simulate their normal behaviour. The sitting activities were performed once on a chair with handles, and once without handles. The driving activity was also performed twice: once driving with one hand and then with two hands. In addition, participant were asked to act-out opening cigarette pack, rolling a cigarette, and lighting a cigarette several times in between the other activities. These actions were not fixed in time, allowing participants to perform these naturally. Given that the experiment took place in-doors, no actual smoking took place; the smoking actions were mimed, using real tobacco products as props.

Data group name	Description
Participant code	Unique participant ID
Creation date	Time stamp (DD-MM-YYYY, HH:MM:SS)
Motion sensor data	Accelerometer Gyroscope Linear acceleration Orientation Rotation vector.
Environmental data	Magnetic field Light level Ambient temperature Relative humidity
Labels	Type of activity Smoking / not smoking Interacting / not interacting with their personal mobile phone Cigarette type (traditional or e-cigarette) Chair type (no arm / two arm chair, also used for indicating one / two hand driving)

TABLE I: Categories of parameters collected by the app.

B. Twente Dataset

The Shoaib et al [13] dataset has 45 hours of 50Hz wrist-mounted accelerometer data of smoking and other similar activities such as eating, and drinking coffee or tea. Out of these 45 hours, the smoking activity was performed for 16.86 hours in various forms including smoking while standing (SmkSTD); smoking while sitting (SmkSIT); smoking while partaking in a group conversation (SmkG); smoking while walking (SmkW). Non-smoking activities included drinking while standing (DrinkSTD); drinking while sitting (DrinkSIT); standing (STD); sitting (SIT); walking (WALK) and eating (Eat). Each activity was performed several times by each participant on various days over a period of three months. Usually, the participants smoked 1-4 cigarettes in a day, and performed the non-smoking eating and drinking activities on different days according to their availability. The participants were divided into three groups, as shown in Table II, on the basis of which activities they performed.

Scenario	Participants	Activities performed
Senario 1	1–11	SmkSTD, SmkSIT, DrinkSTD, DrinSIT, Eat
Scenario 2	1–8	SmkG, SmkSTD, SmkSIT, DrinkSTD, DrinkSIT, Eat
Scenario 3	1–3	SmkW, SmkG, SmkSTD, SmkSIT, DrinkSTD, DrinkSIT, Eat

TABLE II: Activities performed by sub-groups of the Twente participants, identified as "Scenarios".

Examples of the acceleration magnitudes obtained for participants undertaking the "smoking while standing" (SmkSTD) activity are given in Figure 1. It demonstrates the potential variation which can be captured by appropriate processing.

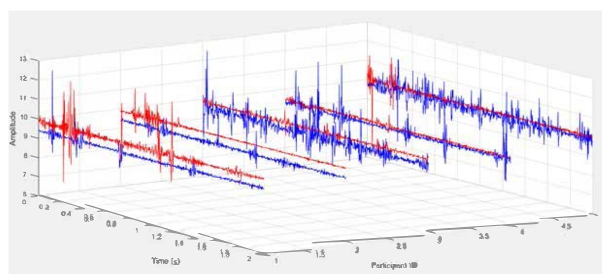

Fig. 1: Acceleration vector magnitude persons of 10 people undertaking the "smoking while standing" (SmkSTD) activity. Each person is a separate line on the axis "Participant ID".

C. Data normalisation and feature extraction

The phone used for data-collection did not generate observations with consistent intervals; data was timestamped to the start of the second in which they were made, but timestamps varied in the number of observations recorded. The data was thus re-sampled, adding additional copies of observations (and also removing excess ones), on the following basis:

```
c = 0, k = 1, newsamples = sample(1)
```

```
while k < #samples
  if c < sr
    k = k + 1
    if k < #samples - 1
      c = c + 1;
      if timestamp(k) == timestamp(k+1)
        concat newsamples, samples(k)
      else
        for c = c : sr
          concat newsamples, samples(k)
        end-for
      end-if
    end-if
  else
    if timestamp(k) != timestamp(k+1)
      c = 0
    end-if
    k = k + 1
  end-if
end-while
```

where `sr` was a tunable parameter, set here to 10Hz, and `concat` adds `samples(k)` to the end of `newsamples`. The signal was then normalised to give a standard range of 1.

An important aspect of describing an action is to quantify its information content. Typically, the statistical properties of the variation inside the signal are used. Recently, techniques based on LBPs have been spreading widely [21, 22]. The success of LBP in signal description reflects their discriminative power, the computational simplicity of the operator, and their robustness to monotonic amplitude values changes caused by artefacts or noise due to the intra/inter-person variability.

The 1D-LBP transformation of the signal point S_c is

$$S^{1D-LBP} = \sum_{i=0}^{m-1} sig(P_i - S_c)2^i \quad (1)$$

where

$$S(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Hence, the transformation codes obtained describe the behaviour across the full extent of the sequence. The histogram of these codes provides the feature vector, which is then fed to a classifier to identify the sequence. Processing is thus not dependent on the overall length of the sequence. The mean and standard deviation of each sample was then used as the features passed to the classifier. Following examination of classification performance, it was found that $m = 18$ gave optimal results for both the MMU and Twente datasets.

D. Classification techniques

Three representative techniques was investigated:

1) k Nearest Neighbours (kNN) - a standard, non-parametric machine learning technique. All of the training samples are retained, with their labels. A distance metric (in this case, sum of squares) is used to find the similarity between a test sample and each training sample. The test sample is assigned the modal label of the k training samples with the smallest distances.
2) Sparse Supervised Representation Classifier (SSRC) - a recently-developed technique for processing data with un-balanced, non-linear class distributions [23]. Class label information is used during the modelling phase to deal with uncontrolled data sets. Both the one-norm of the observation code and a two-norm of the representation error are minimized; each class linearly represents the test sample in its subspace. The test sample is assigned to the class with the lowest representation error.
3) Probabilistic Neural Network (PNN) - a standard, non-parametric classification technique [24] based on the feed-forward neural network. The probability distribution of each class is approximated by a Parzen window. Given a test sample, the likelihood for each class if estimated and Bayes' rule used to allocate the sample to the class with the highest posterior probability.

The evaluation was performed using a ten-fold cross-validation procedure. The samples were processed and three classifiers implemented and tested using the Matlab2017 "Statistics and Machine Learning Toolbox".

III. Results

A. MMU dataset

The system performance was evaluated using the true positive (TP) rate, comparing the categories of smoking and non-smoking, pooled across other concurrent activities. While the KNN was unable to identify either smoking non-smoking cases (see Table III), the PNN shows a ceiling level of performance on both, with the SSRC at an intermediate level.

	KNN	SSRC	PNN
Smoking	0.50	1.00	1.00
Non-smoking	0.14	0.90	1.00

TABLE III: TP rates on the MMU dataset.

The detection of the concurrent activities was also considered, obtaining values separately for them (see Table IV)). Again, the PNN performance is at, or close to, ceiling with SSRC and KNN lower. Interestingly, the hit rates for the non-smoking condition is consistently lower than that for smoking. This may reflect a greater variability of the actions when they are not constrained by the need to manipulate a cigarette.

The data was also classified using the KNN while varying m. This altered the time-period considered for each classification. The effects of varying m between 8 and 53 were explored for smoking detection under Scenario 3, clustered into sets of 4 to allow measurement of distributions. The results are shown in Figure 2; optimal performance is seen with $m = 18$. Since the data is sampled at 10Hz, this corresponds to 1800ms.

B. Twente dataset

The assessment of the Twente dataset was divided between the three Scenarios, as the activities performed varied between participants. In Scenario 1, all participants are considered,

Condition		KNN	SSRC	PNN
Smoking	Calling	0.90	0.87	0.90
	Drinking	0.50	0.70	0.80
	Driving	0.87	1.00	1.00
	Lying down	0.90	0.90	0.87
	Sitting	0.50	0.87	1.00
	Standing	0.90	0.80	1.00
	Walking	0.50	0.70	0.80
Non-smoking	Calling	0.90	0.87	0.90
	Drinking	0.50	0.70	0.80
	Driving	0.25	0.87	1.00
	Lying down	0.90	0.90	0.87
	Sitting	1.00	1.00	1.00
	Standing	0.50	0.60	0.70
	Walking	0.50	0.70	0.80

TABLE IV: TP rates by activity on the MMU dataset.

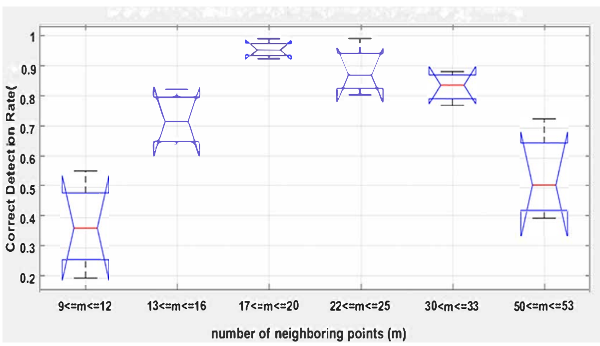

Fig. 2: Effects (mean, standard deviation and range) on classification accuracy of varying the number of items entering the 1-D LBP calculation on the MMU dataset.

but smoking while walking and smoking while in group conversation are ignored, as not all participants performed these two variations of smoking. In Scenario 2, the first eight participants are considered, as all of them performed smoking while in a group conversation. However, in this scenario, the smoking while walking is not considered for the first three participants. Finally, in Scenario 3, only the first three participants are considered, and all of their activities. These participants are notable as only they performed smoking while walking and walking without smoking.

Table V reports the hit rates for the three basic activities of smoking, eating and drinking. As can be seen, a ceiling level accuracy (100%) was obtained from the PNN; SSRC came close to it, with KNN significantly lower. It should be noted that higher levels of performance are seen when the number of participants is reduced.

To make the discriminations more challenging and realistic, a constant level of white noise was added to the acceleration data. The results of this manipulation are shown in Table VI; although the pattern has not changed, the hit rate for KNN and SSRC is reduced. Across the scenarios, it should be observed that it is relatively easy to recognize smoking while standing and while walking, but relatively difficult to detect the other activities. Such a pattern was found by Sohaib et al.,[3, 13], who reported particular difficulty with 'smoking

Scenario		KNN	SSRC	PNN
Scenario 1	Smoke	0.45	0.90	1.00
	Eat	0.36	0.45	1.00
	Drink	0.63	0.63	1.00
Scenario 2	Smoke	0.50	0.87	1.00
	Eat	0.36	1.00	1.00
	Drink	0.50	0.87	1.00
Scenario 3	Smoke	1.00	1.00	1.00
	Eat	0.36	1.00	1.00
	Drink	0.63	1.00	1.00

TABLE V: TP rates on the Twente dataset, no noise.

while sitting, which was on occasion very similar to drinking tea or coffee. This made it difficult to recognize. In contrast, the current study achieved 100% identification in all scenarios, when using the PNN with 1D-LBP.

Scenario		KNN	SSRC	PNN
Scenario 1	Smoke	0.27	0.90	1.00
	Eat	0.36	0.54	1.00
	Drink	0.63	0.67	1.00
Scenario 2	Smoke	0.31	0.91	1.00
	Eat	0.12	1.00	1.00
	Drink	0.63	0.63	1.00
Scenario 3	Smoke	0.33	1.00	1.00
	Eat	0.33	1.00	1.00
	Drink	0.33	1.00	1.00

TABLE VI: TP rates on the Twente dataset, added noise.

IV. CONCLUSIONS

This study has shown that the combination of 1D-LBP and a PNN, both applied to the acceleration magnitude recorded from a wrist-mounted smart-phone, can be used to distinguish between smoking and other similar activities at a ceiling level of performance. This applied to both structured and more naturalistic datasets, also making it possible to identify particular concurrent or separate non-smoking activities with a somewhat lower level of accuracy. This latter result is important since it enables the detection of behaviours which precede the particular act of cigarette smoking, hence allowing both analysis of the context of smoking and also of real-time cuing of participants to not smoke.

As this study did not seek to find the optimal sequence length for smoking detection, the whole of each recorded sub-sequence was used as input. It does however appear that approximately 180ms is appropriate. Future work should be able to detect smoking events within longer sequences and measure error between times of actual and detected smoking events. As smoking actions are likely to vary over shorter periods than other quasi periodic hand to mouth movements, a wider range of LBP distribution measures should be investigated.

Wearable technology potentially provides a more reliable way to identify smoking episodes than approaches that rely on traditional self-report. The overall aim of this project is to provide longer-term records of the times at which the smoking occurs, which can then be used as labels in a machine learning system describing behaviours antecedent to smoking, and collected from smart-phones that can be held and carried in more variable and natural locations. This will in turn provide real-time indicators of the antecedents of smoking, without the requirement to wear a smart-watch in a particular location (especially the need to wear the watch on the hand used for smoking). This will ensure that the recording process is minimally invasive for the participants. Thus there is no need to use the smart-watch data about hand-movement to provide real-time information about activities. The next stage of this research will thus centre on the collection of an extended, naturalistic dataset, via multiple sensors. This will allow combined analysis of the short-term smoking actions, as here, and also the longer-term behavioural and environmental circumstances under which individuals smoke.

References

[1] O. D. Lara and M. A. Labrador, "A survey on human activity recognition using wearable sensors." *IEEE Communications Surveys and Tutorials*, vol. 15, no. 3, pp. 1192–1209, 2013.

[2] M. Pietikäinen, A. Hadid, G. Zhao, and T. Ahonen, *Computer vision using local binary patterns*. Springer, 2011.

[3] M. Shoaib, H. Scholten, P. J. M. Havinga, and O. D. Incel, "A hierarchical lazy smoking detection algorithm using smartwatch sensors," in *Proc. e-Health Networking, Applications and Services*, 2016, pp. 1–6.

[4] K. Ouchi and M. Doi, "Indoor-outdoor activity recognition by a smartphone," in *Proc. ACM Conference on Ubiquitous Computing*, 2012.

[5] R. S. Schick, T. W. Kelsey, J. Marston, K. Samson, and G. W. Humphris, "Mapmysmoke: feasibility of a new quit cigarette smoking mobile phone application using integrated geo-positioning technology, and motivational messaging within a primary care setting," *Pilot and Feasibility Studies*, vol. 4, no. 1, p. 19, 2018.

[6] A. Parate, M.-C. Chiu, C. Chadowitz, D. Ganesan, and E. Kalogerakis, "Risq: Recognizing smoking gestures with inertial sensors on a wristband," in *Proc Mobile Systems, Applications, and Services*, 2014.

[7] A. Parate and D. Ganesan, "Detecting eating and smoking behaviors using smartwatches," in *Mobile Health*, 2017, pp. 175–201.

[8] S. Shiffman, A. A. Stone, and M. R. Hufford, "Ecological momentary assessment," *Annual Review of Clinical Psychology*, vol. 4, pp. 1–32, 2008.

[9] M. S. Businelle, P. Ma, D. E. Kendzor, S. G. Frank, D. J. Vidrine, and D. W. Wetter, "An ecological momentary intervention for smoking cessation: evaluation of feasibility and effectiveness," *Journal of Medical Internet Research*, vol. 18, no. 12, 2016.

[10] M. S. Businelle, P. Ma, D. E. Kendzor, L. R. Reitzel, M. Chen, C. Y. Lam, I. Bernstein, and D. W. Wetter, "Predicting quit attempts among homeless smokers seeking cessation treatment: an ecological momentary assessment study," *Nicotine and Tobacco Research*, vol. 16, no. 10, pp. 1371–1378, 2014.

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

[12] Q. Tang, D. J. Vidrine, E. Crowder, and S. S. Intille, "Automated detection of puffing and smoking with wrist accelerometers," in *Proc. Pervasive Computing Technologies for Healthcare*, 2014.

[13] M. Shoaib, O. D. Incel, H. Scholten, and P. Havinga, "Smokesense: Online activity recognition framework on smartwatches," in *Mobile Computing, Applications, and Services*, 2018, pp. 106–124.

[14] C. Maramis, V. Kilintzis, P. Scholl, and I. Chouvarda, "Objective smoking: Towards smoking detection using smartwatch sensors," in *Precision Medicine Powered by pHealth and Connected Health*, 2018, pp. 211–215.

[15] T. L. Webb, B. P. Chang, and Y. Benn, "'the ostrich problem': motivated avoidance or rejection of information about goal progress," *Social and Personality Psychology Compass*, vol. 7, no. 11, pp. 794–807, 2013.

[16] S. Michie, S. Ashford, F. F. Sniehotta, S. U. Dombrowski, A. Bishop, and D. P. French, "A refined taxonomy of behaviour change techniques to help people change their physical activity and healthy eating behaviours: the calore taxonomy," *Psychology & Health*, vol. 26, no. 11, pp. 1479–1498, 2011.

[17] J. Lu, J. Wang, X. Zheng, C. Karmakar, and S. Rajasegarar, "Detection of smoking events from confounding activities of daily living," in *Proc. Australasian Computer Science Week Multiconference*, 2019.

[18] C. A. Cole, D. Anshari, V. Lambert, J. F. Thrasher, and H. Valafar, "Detecting smoking events using accelerometer data collected via smartwatch technology: Validation study," *JMIR mHealth and uHealth*, vol. 5, no. 12, 2017.

[19] F. Alharbi and K. Farrahi, "A convolutional neural network for smoking activity recognition," in *Proc. e-Health Networking, Applications and Services*, 2018.

[20] A. L. Skinner, C. J. Stone, H. Doughty, and M. R. Munafò, "Stopwatch: The preliminary evaluation of a smartwatch-based system for passive detection of cigarette smoking," *Nicotine and Tobacco Research*, vol. 21, no. 2, pp. 257–261, 2018.

[21] T. Ojala and M. Pietikäinen, "Unsupervised texture segmentation using feature distributions," *Pattern Recognition*, vol. 32, no. 3, pp. 477–486, 1999.

[22] T. Ojala, M. Pietikainen, and T. Maenpaa, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 7, pp. 971–987, 2002.

[23] T. Shu, B. Zhang, and Y. Y. Tang, "Sparse supervised representation-based classifier for uncontrolled and imbalanced classification," *IEEE Transactions on Neural Networks and Learning Systems*, 2018.

[24] D. F. Specht, "Probabilistic neural networks," *Neural Networks*, vol. 3, no. 1, pp. 109–118, 1990.