Detecting Face Touching with Dynamic Time Warping on Smartwatches: A Preliminary Study

Yu-Peng Chen
Department of CISE, University of Florida, USA

Chen Bai
Department of Health Outcomes and Biomedical Informatics, College of Medicine, University of Florida, USA

Adam Wolach
Department of Aging and Geriatric Research, College of Medicine, University of Florida, USA

Mamoun T. Mardini
Department of Health Outcomes and Biomedical Informatics, College of Medicine, University of Florida, USA

Lisa Anthony
Department of CISE, University of Florida, USA

ABSTRACT

Respiratory diseases such as the novel coronavirus (COVID-19) can be transmitted through people's face-touching behaviors. One of the official recommendations for protecting ourselves from such viruses is to avoid touching our eyes, nose, or mouth with unwashed hands. However, prior work has found that people touch their face 23 times per hour on average without realizing it. Therefore, in this Late-Breaking Work, we explore a possible approach to help users avoid touching their face in daily life by alerting them through a smartwatch application every time a face-touching behavior occurs. We selected 10 everyday activities including several that should be easy to distinguish from face touching and several that should be more challenging. We recruited 10 participants and asked them to perform each activity repeatedly for 3 minutes at their own pace while wearing a Samsung smartwatch. Based on the collected accelerometer data, we used dynamic time warping (DTW) to distinguish between the two groups of activities (i.e., face-touching and non-face-touching), which is a method well-suited for small datasets. Our findings show that the DTW-based classifier is capable of classifying the activities into two groups with high accuracy (i.e., 99.07% for the user-dependent scenario). We demonstrated that smartwatches have the potential to detect face-touching behaviors with the proposed methodology. Future work can explore other classification approaches, collect larger datasets, and consider other sensors to increase the robustness of our results.

CCS CONCEPTS

• **Human-centered computing**; • **Mobile devices**; • **Applied computing**; • **Consumer health**;

KEYWORDS

Human Activity Recognition, Dynamic Time Warping, Smartwatch, Accelerometer

ICMI '21 Companion, October 18–22, 2021, Montréal, QC, Canada

ACM ISBN 978-1-4503-8471-1/21/10.
https://doi.org/10.1145/3461615.3485433

ACM Reference Format:
Yu-Peng Chen, Chen Bai, Adam Wolach, Mamoun T. Mardini, and Lisa Anthony. 2021. Detecting Face Touching with Dynamic Time Warping on Smartwatches: A Preliminary Study. In *Companion Publication of the 2021 International Conference on Multimodal Interaction (ICMI '21 Companion), October 18–22, 2021, Montréal, QC, Canada.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3461615.3485433

1 INTRODUCTION

COVID-19 has been a global challenge in many ways [19]. To prevent another global pandemic, it is essential to find solutions to slow the spread of this virus, as well as respiratory diseases caused by other similar viruses. A person can contract COVID-19 by touching a contaminated surface or object and then touching their eyes, nose, or mouth [4]. One way to protect ourselves from the virus is to avoid touching our face with unwashed hands [4]. However, prior work has found that people touch their face 23 times per hour on average without realizing it [12]. A potential approach to reducing such intrinsic behavior is to deliver real-time intervention every time people try to touch their face. We propose to detect users' face-touching behaviors and alert them of such behaviors via smartwatches.

One of the reasons for utilizing a smartwatch is its increasing pervasiveness. Worldwide shipments of wristbands and smartwatches reached more than 170 million units in 2020 [9]. This popularity provides a scalable platform that can be used by a large share of the population, which in this application can mitigate the spread of respiratory diseases more effectively. A smartwatch also has computational and sensory capabilities that are sufficient for achieving our goal, with the help of its triaxial accelerometer [22]. Davoudi et al. [6] showed that a smartwatch can be used as an alternative to a research-grade activity monitor for accelerometer-based human activity recognition. Therefore, we trained a DTW-based classifier on accelerometer data collected from a Samsung Galaxy Watch to detect face touching. More specifically, we asked a total of 10 participants to perform 10 everyday activities, including 6 non-face-touching activities and 4 face-touching activities, repeatedly for 3 minutes at their own pace while wearing the smartwatch on their dominant wrist.

We tested the DTW-based classifier in a user-dependent scenario, in which training and test sets are drawn from the same user, and in a user-independent scenario, in which a set of users are used for training while an additional user is used for testing. Our results showed that recognizing each activity individually (i.e., multiclass classification) is challenging (i.e., 92.48% for the user-dependent scenario and 55.1% for the user-independent scenario). However, our DTW-based classifier can distinguish between face-touching and non-face-touching activities in general with satisfactory accuracies (i.e., 99.07% for the user-dependent scenario and 85.13% for the user-independent scenario). In the user-dependent scenario, the very high recognition accuracy indicates the potential for using a DTW-based classifier in developing a personalized face-touching detection application. For the user-independent scenario, the DTW-based classifier was able to achieve a satisfactory accuracy with fewer training samples than machine learning-based (ML-based) classifiers [24]. This makes the DTW-based classifier potentially more suitable for mobile devices that have limited memory space to store a relatively large ML model. Also, compared with ML-based classifiers that require the process of feature extraction and model training, a DTW-based template matcher does not extract features from training templates but matches point-paths, which makes it more suitable for rapid prototyping [23, 25]. This is especially important in the early stage of developing a new methodology because it helps explore and realize new concepts more easily and quickly. We demonstrated that smartwatches have the potential to detect face touching using the DTW algorithm, which in turn can mitigate the spread of COVID-19 as well as other respiratory diseases.

2 RELATED WORK

2.1 Face Touching Detection with Wearables

Face touching was shown to be a frequent habit by Kwok et al. [12]. This behavior can result in the spread of coronavirus since unwashed hands are considered a common vector for the transmission of respiratory diseases [14]. Wearable devices that can alert the user if they touch their face (e.g., the Immutouch wristband [7] or NASA's PULSE pendant necklace [18]) have been introduced as a response to the COVID-19 outbreak. While these new devices can help slow the spread of the virus, they impose additional cost to users while only providing one specific functionality, which could hinder these solutions from reaching a large share of the population. As an alternative solution, we propose to detect face touching using commercially-available smartwatches for two main reasons: (1) the increasing popularity of such wearable devices [9] provides a platform that is accessible to a large share of the population, and (2) the computational and sensory capabilities of a smartwatch are sufficient to identify the user's current activity [24]. One prior research project [5] also investigated the feasibility of detecting face touching using a smartwatch, but with a very limited number of users (only 3) and prompted face touching behaviors only (e.g., touching their nose and cheek). Their Random Forest (RF) model achieved an accuracy of 92%. Our work goes beyond this exploratory study by recruiting more participants (i.e., 10) and asking them to simulate actual daily activities.

2.2 Human Activity Recognition

Prior work has explored algorithms for human activity recognition using smartphones [2, 11] and wearables [6, 24]. Kwapisz et al. [11] collected accelerometer data from 29 users while they performed 6 daily activities with a smartphone in their pockets. The authors showed that most activities can be recognized with accuracies above 90% using ML models. Anguita et al. [2] had 30 participants perform a similar set of everyday activities while wearing a waist-mounted smartphone and achieved an overall accuracy of 96% using a multiclass SVM. While this prior work on smartphone-based activity recognition has shown promising results, only certain behaviors can be recognized with these approaches: that is, smartphones placed in the pocket may not be as effective for recognizing hand-motion activities. This limitation could be addressed by using wrist-mounted wearables such as smartwatches [6, 24]. Weiss et al. [24] asked 17 participants to perform non-hand-oriented activities (e.g., climbing stairs) and hand-oriented activities (e.g., drinking) while they wore a smartwatch on their dominant wrist and had a smartphone in their pocket. Using the Random Forest algorithm, the authors found that smartwatch-based classifiers generally performed better than smartphone-based classifiers. Furthermore, Davoudi et al. [6] trained ML-based classifiers on accelerometer data collected from a Samsung Gear S smartwatch and achieved similar accuracies in recognizing 15 daily activities as compared to a research-grade ActiGraph GT3X+ activity monitor. Their results showed that smartwatches can be reliable enough for real-world applications.

This prior work has demonstrated that ML-based methods can recognize hand-based activities with high accuracy. However, ML-based methods usually require a relatively large pre-trained model to achieve satisfactory results [10], which may not be suitable for mobile devices with limited memory space. Another possible method for hand-based activity recognition is DTW, an algorithm that measures the similarity between two time series that may vary in length [3]. DTW has been shown to achieve comparable performance in recognizing hand gestures with significantly fewer training samples [1, 13, 20]. Liu et al. [13] designed *uWave*, a personalized system that uses DTW for accelerometer-based gesture recognition. This user-dependent system requires only one training template for each gesture and achieved a competitive accuracy of 98.6% on mobile devices. Additionally, Akl and Valaee [1] asked 7 participants to perform a set of 18 gestures using the Wiimote controller. The authors demonstrated that an accelerometer-based hand gesture recognition system employing DTW can achieve an almost perfect user-dependent accuracy of 99.79% and a satisfactory user-independent accuracy of 96.89%. We go beyond this prior work by asking our participants to simulate actual daily activities while wearing an everyday mobile device (i.e., a Samsung Galaxy Watch) instead of only performing simple gesture patterns (e.g., a circle) with a dedicated device (e.g., Wiimote). To our knowledge, no prior work has used DTW for accelerometer-based face-touching detection on smartwatches. Our goal is to examine the potential for using a DTW-based classifier trained with limited data to detect face touching on resource-constrained devices such as smartwatches.

3 METHOD

To examine the performance of a DTW-based classifier in recognizing face-touching activities based on smartwatch data, we (1) built a smartwatch app for hand-motion data collection on the Samsung Galaxy Watch, (2) asked 10 participants to perform 10

Table 1: Physical activities performed by participants.

Group	Activity	Description
Non-face-touching	Using a mobile phone	Messaging using social media (no phone calls)
	Lying flat on the back	Simulating sleeping or napping
	Computer tasks	Typing a document and navigate websites
	Writing	Writing on a piece of paper
	Leisurely walk	Walking at leisurely pace
	Moving items from one location to another	Simulating moving a light item from one place to another
Face-touching	Repeated face touching	Wiping nose, or other gestures on the face
	Eating and drinking	Eating a snack and drinking water
	Simulated smoking	Simulating the act of smoking a cigarette
	Adjusting eyeglasses	Adjusting eyeglasses placed on the face

everyday physical activities while wearing the smartwatch, (3) implemented the DTW-based classifier, and (4) tested the classifier on the collected data in the user-dependent and user-independent scenarios.

3.1 Data Collection

We implemented a monitoring application that records accelerometer data at a rate of 30Hz on the smartwatch using Tizen Studio [21]. We recruited 10 participants, including 5 females and 5 males who ranged in age from 20 to 83 (M = 47.7, SD = 27.7), from the local community by distributing flyers and word-of-mouth advertising. We excluded individuals with cognitive impairment or medical conditions that would prevent them from performing the activities in our experiment. To simulate face-touching behaviors users might exhibit in everyday life, we selected 10 everyday activities, including 6 activities that should be easy to distinguish from face touching, such as writing, and 4 activities that should be more difficult to distinguish from face touching, such as adjusting eyeglasses (Table 1). We accumulated ten streams of data for each of the ten participants, with each stream representing a 3-minute task consisting of consecutive repetitions of an activity. Therefore, our dataset contains 100 time series that are each 3 minutes long.

The study procedure was approved by our university's Institutional Review Board. We conducted laboratory-based data collection sessions. Before each session, the participant provided written informed consent. We then asked participants to perform each activity (Table 1) repeatedly for 3 minutes at their own pace. Since people typically write with their dominant hand, participants were asked to wear the smartwatch on their dominant wrist and perform each activity using their dominant hand. For each participant, we randomized the order of the tasks to avoid potential order effects caused by fatigue. Participants received a $25 gift card as a compensation at the end of the session. Since the sessions were in person, we took COVID-19 precautions as follows: we took the temperature of all participants and study coordinators prior to starting the session, we wiped down all equipment with disinfecting wipes after each session, and all participants and study coordinators wore masks and gloves and practiced social distancing during the session.

3.2 The DTW-Based Classifier

DTW [3] is a dynamic programming algorithm that is capable of measuring the similarity between two temporal sequences of different lengths. Berndt and Clifford [3] provide a more detailed explanation in their work. In our experiment, we collected accelerometer data for the three spatial dimensions. Therefore, a sample represents a repetition of an activity consisting of three time series, each corresponding to the three axes. We calculate the matching between two samples A and B using the following equation.

$$DTW\,(A,B) = \sqrt{DTW(A_x,B_x)^2 + DTW\left(A_y,B_y\right)^2 + DTW(A_z,B_z)^2}$$

In this equation, we used the open-source Python package DTAIDistance [16] to calculate an overall DTW distance between the two samples based on the three DTW distance values in the x, y, and z directions. To classify a testing sample, we first calculated the DTW distance between the testing sample and all training templates individually, and then assigned to the testing sample the label of the training template that resulted in the minimum DTW distance.

3.3 Analysis

Our overall goal was to classify the data samples into one of the two groups (i.e., face-touching and non-face-touching). We also wanted to explore how well DTW could distinguish each individual activity from each other. Therefore, we formulated the problem in two ways: **binary classification** and **multiclass classification**. For our analysis, we applied data preprocessing techniques to each 3-minute time series to (1) capture a stream of cyclical activity repetitions without the irregular noise at the start and end of the stream, and then (2) divide the stream into smaller segments, each representing a repetition of an activity, to be used for classification. First, we examined the visualization of the time series to identify the extent of irregular noise at the start and end of each task caused by the starting and ending of the task by the user; based on this observation, we removed the first 20 and last 5 seconds from all samples. Next, we conducted preliminary experiments to select the optimum number of segments (e.g., 1, 12, 18, 20, 30, 36), and found that dividing each time series into 12 segments of roughly 15 seconds in length resulted in the best recognition performance. Our resulting dataset contained a total of 1200 samples = 10 (participants) × 10 (activities) × 12 (segments).

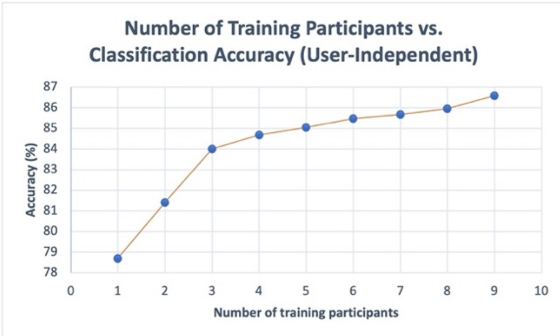

Figure 1: The relationship between the number of training participants per activity type and classification accuracy in the user-independent scenario. For each activity type, a single template was selected from each training participant.

After preprocessing the data, we first aimed at training and testing our DTW-based classifier with all possible combinations of training and test samples in both user-dependent and user-independent scenarios, based on prior work [23]. In the **user-dependent** scenario, the classifier was trained and tested on a specific user (i.e., best-case accuracy) using cross-validation. For each participant, for each activity type, T samples were randomly selected as the training data and one additional sample was selected for testing. We repeated this process 100 times for each value of T (e.g., 1 to 11) and calculated the average of the results for each participant. The results from the 1.1×10^5 classification tests were averaged into an overall classification accuracy. For the **user-independent** scenario, we also used cross-validation to examine how well a classifier can generalize to activities performed by users whose data are not included in the training set. We selected one participant's data to use for testing and randomly selected P other participants' data to use for training. For each activity type, T samples were randomly selected from each training participant while one sample was selected from the testing participant. We repeated the process 10 times for each P (e.g., 1 to 9) and 10 times for each T (e.g., 1 to 11). The results from the 9.9×10^5 classification tests were averaged into an overall classification accuracy. Furthermore, to understand whether face-touching detection would be feasible on resource-constrained devices, we specifically looked at the case of minimum training samples in both scenarios. We only ran this additional experiment for the binary classification problem since our primary goal for the study was to classify the data samples into one of the two groups (i.e., face-touching and non-face-touching). To **train with only a single template** for each activity type, the value of T was set to 1 in the user-dependent scenario while the values of both P and T were set to 1 in the user-independent scenario.

4 RESULTS

On the binary classification problem, when training and testing with all possible combinations of training and testing samples, the DTW-based classifier delivered satisfactory accuracies of 99.07% (user-dependent) and 85.13% (user-independent). The results also showed that non-face-touching activities were recognized with higher accuracies than face-touching activities. When training with only a single template for each activity type, the classifier was able to maintain a satisfactory accuracy of 97.37% in the user-dependent scenario. For the user-independent scenario, however, the classifier only achieved an unsatisfactory accuracy of 78.69%. Therefore, we increased the number of training participants per activity type by varying P from 1 to 9 while keeping T as 1. As Figure 1 shows, our DTW-based classifier can achieve an acceptable accuracy of more than 85% when the number of training participants is increased to five or more per activity type (while T is held static at 1). Interested readers can find in the supplementary material the relationship between the number of training samples per activity type and classification accuracy in the user-dependent scenario.

In the multiclass classification problem, the DTW-based classifier achieved accuracies of 92.48% (user-dependent) and 55.1% (user-independent). Table 2 shows the confusion matrix in the user-dependent scenario. The "using mobile phone" and "writing" tasks were recognized with accuracies over 99%. However, "adjusting eyeglasses" was recognized with the lowest accuracy of 78.36% and was mainly confused with "repeated face touching". Regarding the face-touching (red text) and non-face-touching (black text) categories shown in Table 2, **there were more within-category confusions than between-category confusions**, especially among the activities in the face-touching category. The main between-category confusions were related to "moving items". User-independent tests revealed the same trend for the between-category confusions. Interested readers can find the user-independent results in the supplementary material.

5 DISCUSSION

For the binary classification problem (i.e., face-touching vs. non-face-touching), in the user-dependent scenario, our classifier achieved an accuracy of 97.37% when training with only a single template for each activity type. Although Liu et al. [13] achieved a slightly better accuracy of 98.6% using DTW for accelerometer-based gesture recognition under the same data constraint, some of this performance advantage is a result of the authors improving the accuracy of their DTW-based recognizer (from 93.5% to 98.6%) with the help of a method called template adaptation [15]. Without using template adaptation, our DTW-based classifier still delivered a competitive accuracy (i.e., 97.37%) with a minimum amount of training data. We were able to achieve such performance by formulating the problem as a binary classification problem instead of a multiclass classification problem. Even though we lose some granularity by only classifying the activity into these two broad groups, the approach can still be useful because performing any one of the face-touching activities in Table 1 is close enough to face touching that the system should still alert the user of the potential face-touching risk. This result showed that a DTW-based classifier has the potential to provide a personalized face-touching detection service on resource-constrained devices such as smartwatches. The system can allow users to input their own templates for each activity type without the need to re-train the model. In the user-independent scenario, more than one training participant per activity type was required for our DTW-based classifier to achieve an acceptable accuracy of more than 85%. As Figure 1 shows, five

Table 2: Confusion matrix of individual activity recognition in the user-dependent scenario. Face-touching activities are colored in red. The matrix elements highlighted with deep blue indicate correct classifications while the elements highlighted with light blue indicate relatively high rates of confusions (i.e., larger than 1.00%).

Actual Class \ Predicted Class	Mobile Phone	Lying Flat	Computer Tasks	Writing	Leisure Walk	Moving Items	Repeated Face Touching	Eating and Drinking	Simulated Smoking	Adjust Eye-glasses
Mobile Phone	99.09%	0.00%	0.55%	0.00%	0.00%	0.00%	0.00%	0.00%	0.36%	0.00%
Lying Flat	0.00%	98.64%	0.27%	0.00%	0.00%	0.00%	0.18%	0.00%	0.73%	0.18%
Computer Tasks	2.00%	0.00%	96.45%	1.45%	0.00%	0.09%	0.00%	0.00%	0.00%	0.00%
Writing	0.00%	0.00%	0.18%	99.82%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Leisure Walk	0.82%	0.00%	0.00%	0.09%	94.36%	4.00%	0.64%	0.09%	0.00%	0.00%
Moving Items	0.18%	0.00%	0.45%	3.45%	0.00%	94.18%	0.00%	1.45%	0.27%	0.00%
Repeated Face Touching	0.00%	0.27%	0.00%	0.00%	0.00%	0.55%	83.91%	3.27%	3.45%	8.55%
Eating and Drinking	0.00%	0.00%	0.00%	0.09%	0.00%	1.55%	2.18%	90.91%	2.73%	2.55%
Simulated Smoking	0.45%	0.09%	0.55%	0.18%	0.00%	0.64%	2.91%	1.45%	89.09%	4.64%
Adjust Eyeglasses	0.18%	0.55%	0.00%	0.00%	0.00%	0.09%	10.91%	3.45%	6.45%	78.36%

or more training participants per activity type, each contributing a single template to the training set, were required to make the classifier more generalizable. These results indicated the need for more memory space to store a total of 50 or more templates. However, this classifier can still be suitable for mobile devices with limited space since the required memory space for a total of 50 templates is less than 1 MB. The findings from our study also indicate a potential for using smartwatches to achieve health goals such as quitting smoking in which the target behavior (i.e., smoking) is similar to face touching. Other possible scenarios in which our findings could potentially be helpful include the treatment of hair-pulling disorder [8] and the treatment of nail-biting disorder [17].

In the multiclass classification problem, there were more within-category than between-category confusions, especially for the face-touching activities (Table 2), which explains the DTW-based classifier's better performance in binary classification than in multi-class classification. Furthermore, "moving items" in the non-face-touching group caused relatively high rates of confusion with the activities in the face-touching group. This can be explained by examining the actual steps of the "moving items" activity: picking up (i.e., lifting) a light item, moving it to an adjacent spot, and putting down the item. The hand movement involved in this activity is relatively similar to face-touching activities, which consist of lifting the hand, moving the hand toward the face, and touching the face. The only difference was the direction (axis) in which the hand moved when performing the tasks. A possible way to resolve this issue could be placing a higher weight on the z-direction acceleration during classification in order to emphasize the unique face-touching feature of moving the hand toward the face. Also, note that we did not try to differentiate between types of face-touching activities (e.g., "repeated face touching" and "simulated smoking") for two reasons: (1) our primary goal was to classify the data samples into one of the two groups (i.e., face-touching and non-face-touching), and (2) all of the face-touching activities (Table 1) are potentially risky regarding the spread of COVID-19.

6 LIMITATIONS AND FUTURE WORK

Though we have shown the potential for using DTW to classify face-touching behaviors from smartphone accelerometer data, there are several limitations to the scope of this late-breaking work. First, we asked our participants to wear the smartwatch on their dominant wrist since we included a writing task, but many people prefer wearing watches on their non-dominant hand. Future studies should supplement our dataset with data collected from non-dominant hand activities. Second, we collected data in a laboratory setting, but users' behaviors as well as external factors may differ in a real-world setting. Also, we did not collect gyroscope data, which has been shown to increase the classification accuracy [6]. Future studies can consider collecting gyroscope data in addition to accelerometer data from a larger number of users in more diverse settings. Third, we focused on investigating the feasibility of using DTW to detect face touching as an alternative to ML-based methods that require more data samples. It could be beneficial to compare the quantitative results of both methods regarding classification accuracy and required resources in future studies. Finally, our classification experiments were only conducted in an offline manner. Future work should implement real-time face-touching detection on smartwatches and study the efficacy in reducing users' face-touching behaviors.

ACKNOWLEDGMENTS

We thank Emmie Palmer and Megan Siobhan Roberts for assisting us with running the study. This project was supported by the UF Informatics Institute (UFII) COVID-19 Response SEED Program.

The content is solely the responsibility of the authors and does not necessarily represent the official views of the UF Informatics Institute (UFII).

REFERENCES

[1] Ahmad Akl and Shahrokh Valaee. 2010. Accelerometer-based gesture recognition via dynamic-time warping, affinity propagation, compressive sensing. In 2010 IEEE International Conference on Acoustics, Speech and Signal Processing, 2270–2273. https://doi.org/10.1109/ICASSP.2010.5495895

[2] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. 2013. A public domain dataset for human activity recognition using smartphones. In Esann, 3.

[3] Donald J Berndt and James Clifford. 1994. Using dynamic time warping to find patterns in time series. In KDD workshop, 359–370.

[4] Centers for Disease Control and Prevention. 2020. Prevent Getting Sick. Retrieved March 8, 2021 from https://www.cdc.gov/coronavirus/2019-ncov/prevent-getting-sick/index.html

[5] Xiang "Anthony" Chen. 2020. FaceOff: Detecting Face Touching with a Wrist-Worn Accelerometer. CoRR abs/2008.0. Retrieved from https://arxiv.org/abs/2008.01769

[6] Anis Davoudi, Amal Asiri Wanigatunga, Matin Kheirkhahan, Duane Benjamin Corbett, Tonatiuh Mendoza, Manoj Battula, Sanjay Ranka, Roger Benton Fillingim, Todd Matthew Manini, and Parisa Rashidi. 2019. Accuracy of Samsung Gear S Smartwatch for Activity Recognition: Validation Study. JMIR Mhealth Uhealth 7, 2: e11270. https://doi.org/10.2196/11270

[7] Avrelle Harrington. 2020. Ready to Finally Break Your Bad Habit? Introducing the New Immutouch Habit Band. Retrieved March 26, 2021 from https://immutouch.posthaven.com/ready-to-finally-break-your-bad-habit-introducing-the-new-immutouch-habit-band

[8] Joseph A Himle, Deborah Bybee, Lisa A O'Donnell, Addie Weaver, Sarah Vlnka, Daniel T DeSena, and Jessica M Rimer. 2018. Awareness enhancing and monitoring device plus habit reversal in the treatment of trichotillomania: An open feasibility trial. Journal of obsessive-compulsive and related disorders 16: 14–20.

[9] International Data Corporation. 2021. Consumer Enthusiasm for Wearable Devices Drives the Market to 28.4% Growth in 2020, According to IDC. Retrieved March 22, 2021 from https://www.idc.com/getdoc.jsp?containerId=prUS47534521

[10] Michael I Jordan and Tom M Mitchell. 2015. Machine learning: Trends, perspectives, and prospects. Science 349, 6245: 255–260.

[11] Jennifer R Kwapisz, Gary M Weiss, and Samuel A Moore. 2011. Activity Recognition Using Cell Phone Accelerometers. SIGKDD Explor. Newsl. 12, 2: 74–82. https://doi.org/10.1145/1964897.1964918

[12] Yen Lee Angela Kwok, Jan Gralton, and Mary-Louise McLaws. 2015. Face touching: a frequent habit that has implications for hand hygiene. American journal of infection control 43, 2: 112–114.

[13] Jiayang Liu, Lin Zhong, Jehan Wickramasuriya, and Venu Vasudevan. 2009. UWave: Accelerometer-Based Personalized Gesture Recognition and Its Applications. Pervasive Mob. Comput. 5, 6: 657–675. https://doi.org/10.1016/j.pmcj.2009.07.007

[14] A E Macias, A la Torre, S Moreno-Espinosa, P E Leal, M T Bourlon, and G M Ruiz-Palacios. 2009. Controlling the novel A (H1N1) influenza virus: don't touch your face! Journal of Hospital Infection 73, 3: 280–281.

[15] F R McInnes, M A Jack, and J Laver. 1989. Template adaptation in an isolated word-recognition system. IEE Proceedings I (Communications, Speech and Vision) 136, 2: 119–126.

[16] Wannes Meert, Kilian Hendrickx, and Toon Van Craenendonck. 2020. wannesm/dtaidistance v2.0.0. https://doi.org/10.5281/zenodo.3981067

[17] Steffen Moritz, Andras Treszl, and Michael Rufer. 2011. A Randomized Controlled Trial of a Novel Self-Help Technique for Impulse Control Disorders: A Study on Nail-Biting. Behavior Modification 35, 5: 468–485. https://doi.org/10.1177/0145445511409395

[18] NASA Jet Propulsion Laboratory. 2020. PULSE: A Pendant to Warn You When You Touch Your Face. Retrieved March 26, 2021 from https://medeng.jpl.nasa.gov/covid-19/pulse/

[19] Maria Nicola, Zaid Alsafi, Catrin Sohrabi, Ahmed Kerwan, Ahmed Al-Jabir, Christos Iosifidis, Maliha Agha, and Riaz Agha. 2020. The socio-economic implications of the coronavirus pandemic (COVID-19): A review. International journal of surgery (London, England) 78: 185–193. https://doi.org/10.1016/j.ijsu.2020.04.018

[20] Jaime Ruiz and Yang Li. 2011. DoubleFlip: A Motion Gesture Delimiter for Mobile Interaction. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11), 2717–2720. https://doi.org/10.1145/1978942.1979341

[21] Tizen Developers. An open source, standards-based software platform for multiple device categories. Retrieved April 12, 2021 from https://developer.tizen.org/

[22] Richard P Troiano, James J McClain, Robert J Brychta, and Kong Y Chen. 2014. Evolution of accelerometer methods for physical activity research. British journal of sports medicine 48, 13: 1019–1023. https://doi.org/10.1136/bjsports-2014-093546

[23] Radu-Daniel Vatavu, Lisa Anthony, and Jacob O Wobbrock. 2012. Gestures as Point Clouds: A $P Recognizer for User Interface Prototypes. In Proceedings of the 14th ACM International Conference on Multimodal Interaction (ICMI '12), 273–280. https://doi.org/10.1145/2388676.2388732

[24] Gary M. Weiss, Jessica L. Timko, Catherine M. Gallagher, Kenichi Yoneda, and Andrew J. Schreiber. 2016. Smartwatch-based activity recognition: A machine learning approach. In 2016 IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI), 426–429. https://doi.org/10.1109/BHI.2016.7455925

[25] Jacob O Wobbrock, Andrew D Wilson, and Yang Li. 2007. Gestures without Libraries, Toolkits or Training: A $1 Recognizer for User Interface Prototypes. In Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology (UIST '07), 159–168. https://doi.org/10.1145/1294211.1294238