# COUGHTRIGGER: EARBUDS IMU BASED COUGH DETECTION ACTIVATOR USING AN ENERGY-EFFICIENT SENSITIVITY-PRIORITIZED TIME SERIES CLASSIFIER

Shibo Zhang<sup>\*</sup>, Ebrahim Nemati<sup>†</sup>, Minh Dinh<sup>‡</sup>, Nathan Folkman<sup>‡</sup>, Tousif Ahmed<sup>†</sup>, Mahbubur Rahman<sup>†</sup>, Jilong Kuang<sup>†</sup>, Nabil Alshurafa<sup>\*</sup>, Alex Gao<sup>†</sup>

\* Northwestern University, Chicago, IL, USA
<sup>†</sup> Samsung Research America, Mountain View, CA, USA
<sup>‡</sup> Samsung Design Innovation Center, San Francisco, CA, USA

# ABSTRACT

Persistent coughs are a major symptom of respiratory-related diseases. Increasing research attention has been paid to detecting coughs using wearables, especially during the COVID-19 pandemic. Among all types of sensors utilized, microphone is most widely used to detect coughs. However, the intense power consumption needed to process audio signals hinders continuous audio-based cough detection on battery-limited commercial wearable products, such as earbuds. We present CoughTrigger, which utilizes a lowerpower sensor, an inertial measurement unit (IMU), in earbuds as a cough detection activator to trigger a higher-power sensor for audio processing and classification. It is able to run all-the-time as a standby service with minimal battery consumption and trigger the audio-based cough detection when a candidate cough is detected from IMU. Besides, the use of IMU brings the benefit of improved specificity of cough detection. Experiments are conducted on 45 subjects and our IMU-based model achieved 0.77 AUC score under leave one subject out evaluation. We also validated its effectiveness on free-living data and through on-device implementation.

*Index Terms*— Cough Detection Activation, Sensitivityprioritized Classification, Multi-Center Classifier, Template Matching, Earbuds, Wearables

# 1. INTRODUCTION

Persistent coughs can be a sign of serious lung diseases, such as Chronic Obstructive Pulmonary Disease (COPD), asthma, pneumothorax, atelectasis, bronchitis, lung cancer, and COVID-19. Since a cough can be an important indicator of respiratory disease, reliable automated detection of coughs using everyday wearable devices is especially desirable. In recent decades, wearable devices such as smartphones and earbuds are becoming prevalent in our daily life. A body of work has recently emerged, with wearable sensors showing promise in automatically detecting coughs [1–9]. Earbuds-based devices have been shown to be effective in cough detection [7–9].



**Fig. 1**: Earbuds embedded with IMU and 3-axis accelerometer cough data (blue: x-axis, green: y-axis, red: z-axis)

Audio-based sensing has shown promise in detecting coughs on device, but requires higher battery consumption and introduces privacy concerns [1-3, 5, 6, 8]. Battery consumption is a major concern, particularly among commercial wearables, such as earbuds, where power is a limited resource. For example, Samsung Galaxy Buds2 has a 61 mAh battery that supports up to 7.5 hours of play time per charge [10], and adding extra features will cause a drop in battery life, and ultimately utility and user-interest. Some studies explored the use of an inertial measurement unit (IMU), due to its low battery footprint and computational load, to detect coughs [7, 9]. It was soon realized that using a traditional classifier (XGBoost) to detect coughs could only yield a 47% sensitivity and 54% F1-Score [7]. The large number of confounding head movements made it challenging to distinguish between a cough and a non-cough.

Audio and IMU sensing have complementary characteristics: audio can better distinguish between coughs and noncoughs, while IMU enables battery-efficiency. Given that the majority of the time is often spent by people not coughing, the low-power IMU can be used to trigger the audio sensing pipeline when a candidate cough is detected. To do this, the IMU must yield high sensitivity, to ensure all coughs are ultimately passed onto the audio sensing. Since traditional classification methods yield low sensitivity, we propose a novel multi-center template matching algorithm to achieve high sensitivity in the IMU data. This algorithm is then used in a two-stage pipeline, where an always-on IMU triggers audio processing, only when needed, to reliably confirm the detection of a cough.



Fig. 2: Illustration of sensitivity-prioritized classification.

In this work, we focus on the IMU-based sensitivityprioritized cough detection problem and propose CoughTrigger, an IMU-based activator that utilizes a novel multi-center classifier to trigger audio processing and cough detection. CoughTrigger can further alleviate the privacy issue because it does not require constant collecting and processing audio data. We summarize the contributions of this work as follows: (1) A battery-efficient dual IMU-audio cough detection framework using earbuds; we define the requirements for the first IMU stage of the pipeline and formalize it as a sensitivity-prioritized classification problem. (2) We propose CoughTrigger, an IMU-based cough detection activator based on a novel template matching method - Multi-Center Classifier and present superior experimental results compared with baselines. (3) We implement our proposed method on a commercial device and prove the feasibility and effectiveness of CoughTrigger with on-device result.

#### 2. METHODOLOGY

#### 2.1. Sensitivity-prioritized Classification

Unlike a traditional binary classification problem where the positive and negative samples are distinguishable either by human perception or machine learning techniques in a transformed feature space, the IMU cough data and non-cough ones are not fully separable due to an overlapping region in feature space, as illustrated in Figure 2. To guarantee the positive samples (coughs) can be detected to trigger audio-based cough detection in the next stage, we formalize the problem as a *sensitivity-prioritized* classification task, which means higher sensitivity is prioritized over specificity, and approach this problem with a novel template matching algorithm.

#### 2.2. Multi-Center Classifier (MCC)

Zhang *et al.* proposed a template matching algorithm [9], called Multi-Centroid Classifier, which aims at iteratively creating an increasing number of clusters, each of which has its own centroid and radius and all together cover all the positive samples while include as few negative samples as possible. When a satisfying accuracy is achieved, the derived centroids and radii will be used in the test set as templates and thresholds to classify positive samples from negative ones. The method has shown merits in accuracy, inference speed, and model size. In this work, we adopt the idea, make sub-

Algorithm 1: Training Multi-Center Classifier
<b>input</b> : Positive and negative training samples and
stop criterion H
<b>output:</b> K templates $C_k$ , $(k = 1, 2,, K)$ and K
thresholds
Initialize number of clusters $K = 1$ ;
Assign all the positive samples to cluster $R_1$ ;
Randomly select seed center $C_1$ from positive
samples of cluster $R_1$ ;
Do discrepancy-based clustering with $C_1$ to obtain
total cost L, updated center $C_1$ and threshold;
while total cost $L > stop$ criterion $H$ do
Select $R_t$ with the highest cost $t = \arg \max_i L_i$ ;
From the positive samples of $R_t$ , select the
sample which brings the largest cost increase as
the new seed centroid;
Do discrepancy-based clustering using $K + 1$
centers $C_1, \ldots, C_{K+1}$ to obtain total cost $L$ ,
updated $K + 1$ centers, and $K + 1$ thresholds;
Calculate cost $L_k$ for each of the $K + 1$ clusters;
$L = \sum_{k=1}^{K+1} L_k;$
K = K + 1;
end

stantial modifications to the training phase, and propose an enhanced version, Multi-Center Classifier, on which we build *CoughTrigger*.

**Training Phase** We present the new training procedure in Algorithm 1. The original MCC comprises three major parts: discrepancy cost, discrepancy-based clustering, and cluster averaging. The discrepancy cost is used to measure and compare the purity of each cluster. Discrepancy-based clustering assign positive samples into clusters. Cluster averaging derives one template and one threshold for each cluster. As the original method is sensitive to regional density in the feature space, which prevents the derivation of a pseudo-optimal clustering result, we make substantial modifications to the discrepancy-based clustering and cluster averaging. Also, we change the way we select new seeds for better convergence. Below are the three modifications introduced:

(1) We update the averaging step, making it easier to converge on noisy datasets. When deriving the centroid for each cluster, instead of averaging operation, we select as center the positive sample which has the minimum  $Dist_{max}$  ( $Dist_{max}$  is the max distance between the current sample and any other positive sample).

(2) By modifying the clustering step, we make the algorithm less sensitive to regional density in the feature space. When clustering, we require each sample to be assigned to only one cluster instead of allowing samples assigned to multiple clusters. This modification brings about better convergence. We will showcase the effect with a toy dataset in Section 4.1. (3) When increasing the number of clusters, instead of adding one random centroid seed, we select the sample that brings the greatest increase to the cost function.

**Inference Phase** The inference steps are unchanged [9]. The distance between a test sample and each template is calculated and compared against the threshold. If the distance is smaller than the threshold, then the test sample is predicted as positive, otherwise it is negative.

# 3. EXPERIMENT

# 3.1. Data Collection

We generated an earbud-based cough dataset from 45 participants (15 with lung disease, 22 male,  $41.4 \pm 10.7$  years old) with approval by the Institutional Review Board. In the experiments, one earbud was worn by a participant to collect IMU data at 50 Hz and audio at 16 KHz.

In-lab Experiment To evaluate the sensitivity of our model we collected eight cough sessions, where we asked each participant to cough continuously for a period of time with a short pause between every two coughs. Cough sessions include five stationary and three non-stationary periods. Stationary periods comprised: coughing while seated (30s), coughing while lying down (30s), coughing while listening to music from an earbud (30s), coughing with background fan noise (30s), and coughing with background music/TV noise (30s). Non-stationary periods comprised: coughing while performing yoga in a quiet environment (45s), coughing while performing yoga in a noisy environment (45s), and coughing while walking (1 min). To evaluate the specificity, non-cough activities that involve signals that could resemble cough motion were collected including: eating (30s), drinking (30s), laughing (30s), scripted speech (1 min), throat clearing (30s), free head motion while talking (30s), and one additional bystander coughing session (30s). On average, 10.5 coughs are captured in each cough session for each participant. We aim at detecting each single cough in every session while preventing false alarms.

**Free-living Experiment** Enrolled participants also took part in a free-living experiment, where they were asked to cough naturally in the morning and in the afternoon for one week. At each time, data were collected for coughs while seated for 30s and while walking for 30s.

Annotation We annotated every single cough using audio. Due to the limitations of data transmission from earbud to data logging app, there is a random time drift of up to about 400 ms between the audio and IMU data. We resolved the asynchronization by observing the data and identifying peaks in the x-axis of accelerometer, which corresponds to the most probable motion of the wearer while coughing. We labeled 45 participants for the in-lab experiment and the 15 participants from the lung disease cohort in the free-living experiment.

### 3.2. Model Development

**Data Preprocessing** 3-axis accelerometer data are preprocessed using a moving average filter with a window size as 10 samples and a Butterworth high-pass filter ( $\omega_c = 3\pi$ ). Positive samples are extracted from seated and lying down cough sessions using a 0.4s window centered around the annotated IMU cough moment. To increase the variety of cough data during training, we applied three time series augmentation methods, namely jittering, scaling, and magnitude warping [11] on cough data enlarging the positive sample size by four times. For negative samples, we segmented the non-cough session accelerometer data using a 0.4s sliding window with 0.1s stride size. Then we randomly subsampled four times the positive samples from non-cough sessions as negative samples, in order to balance class ratio.

**Model Training** To expedite the training process, we trained one MCC model with stop precision criterion as 0.8 for each participant in the training set using multi-variate DTW distance measure, then aggregated the models of each training participant together and ran on all the training participants. We used a greedy algorithm to rank the templates and select the top K templates based on template importance measured by how many new positive samples are hit by each template.

**Model Testing** We applied MCC with K templates on the accelerometer data of the test participant's cough and noncough sessions, with 0.4s window size and 0.02s stride size. Then, we aggregated and merged all the predicted cough windows to determine the final cough event prediction. In our case, the capability of manually adjusting the trade-off between sensitivity and specificity is desirable. A benefit of MCC is that we have two ways to adjust it: to choose the number of top K templates used and to adjust the thresholds of templates. When more (less) templates are used, or when we increase (decrease) the thresholds of templates, it leans towards a higher sensitivity (specificity).

**Evaluation Method** Leave one subject out cross validation (LOSOCV) is adopted. For cough sessions, we compared the predicted cough segments against the ground truth coughs. True positives are defined as predicted cough segments intersected by ground truth coughs. We calculated the sensitivity of cough sessions as a ratio of true positives to ground truth coughs. We used sample-level specificity to test how well it can specify non-cough events. The model is validated on free-living data in a similar manner.

### 4. RESULT

### 4.1. Improvement of MCC

We used a synthetic dataset generated from Gaussian distribution to validate the improved MCC, as in Fig. 3. The original MCC achieved 0.57 test accuracy with 10 centroids. Using improved MCC, 0.99 test accuracy was obtained with 8 cen-



Fig. 3: Illustration of Multi-Center Classifier improvement.

ters. The first reason is that the original method allows one sample to belong to multiple clusters, increasing the possibility of highly overlapped clusters. Instead, the new method can separate clusters towards different directions. Second, the new clustering averaging method makes MCC less sensitive to regional density.

# 4.2. CoughTrigger Results

**Choice of Input Data** We compared all the combinations of the 3-axis accelerometer data, including three using one axis, three using two axes, and one using three axes. The best result is achieved with both the x- and y-axis. That aligns with the IMU direction in Fig. 1 as z-axis only contributes to sensing motion in the left/right direction, which is not as useful in detecting coughs. We tried with only 3-axis gyroscope and combining both accelerometer and gyroscope data but observed lower accuracy.

**In-lab Experiment Result** Under LOSOCV, when training MCC, on average 14.2 templates were generated for each participant in each fold. When using top five templates, we achieved 90% sensitivity on average for the five stationary cough sessions, and an average sensitivity of 86% on all cough sessions. The average specificity is 53% across all non-cough sessions. When the number of adopted templates ranges from 1 to 30, we achieved 0.77 AUC with all the 15 sessions. We present the ROC curve of *CoughTrigger* in Fig. 4 Left. We emphasize that our specificity evaluation was designed for the worst case scenario with a variety of activities. In real life, as for most of the time the wearer is stationary, the specificity is expected to be higher.

**Free-living Experiment Result** We trained with the in-lab data from all 45 participants and tested on the free-living data from the 15 participants in the disease cohort. When we applied top 10 templates, we received an average sensitivity of 87% during the stationary periods, and an average sensitivity of 82% on both seated and walking cough periods. Although we only have cough sessions in the free-living setting, we found a 55% average specificity across both sessions for all the participants. When we adjusted the thresholds of each templates, we received 0.80 AUC for stationary coughs and 0.73 AUC for stationary and walking sessions.

Baseline Methods Since there is no existing out-of-box



**Fig. 4**: Left: In-lab experiment ROC curve; Right: Battery life comparison.

sensitivity-prioritized classification method, we investigated ways that may apply, including one-class classifier and adjusting the decision boundary of a traditional classifier. Because one-class classifier only utilizes the positive class and models the distribution of positive samples, the overlapping of two classes should not interfere with modeling of the positive class. We tested the popular OC-SVM which relies on identifying the smallest hypersphere consisting of all the positive samples [12, 13]. We utilized the same preprocessing steps and concatenated x- and y-axis into one vector of size 40 as input. After adjusting the hyperparameters in a large range with different kernels (linear, RBF, and Sigmoid), we only received 0.51 AUC, which is no better than random guess. For a traditional classifier, we implemented a 3-layer NN with 20, 10, and 5 neurons in each layer and we observed the same result as OC-SVM.

**On-device Implementation** We implemented *CoughTrigger* on Samsung Galaxy Buds2. Fig. 4 Right shows the battery life of base firmware without cough detection, *CoughTrigger* using 10 and 20 templates integrated in the base firmware, and an integrated audio-based cough detection method [8]. We see that the base firmware without cough detection has around 18 hours battery life, and integrating *CoughTrigger* (10 or 20 templates) makes no significant change in the battery life. The IMU-based cough detection implementation consumes less battery than an audio-based method, which shows the feasibility of leveraging *CoughTrigger* as a first-stage filter to reduce battery consumption.

### 5. CONCLUSION

In this work, we introduce a battery-efficient earbuds-based two-stage IMU-audio cough detection framework and formalize the IMU first stage as a sensitivity-prioritized classification problem. We propose using a novel multi-center classifier as a first-stage cough detection activator. By conducting in-lab, free-living, and on-device experiments, we demonstrate the feasibility and effectiveness of our proposed method. Our proposed sensitivity-prioritized template matching algorithm can be adopted as a plug-and-play module for other sensor-fusion wearable applications.

# 6. REFERENCES

- Eric C. Larson, TienJui Lee, Sean Liu, Margaret Rosenfeld, and Shwetak N. Patel, "Accurate and privacy preserving cough sensing using a low-cost microphone," in *Proceedings of the 13th International Conference on Ubiquitous Computing*, New York, NY, USA, 2011, UbiComp '11, p. 375–384, Association for Computing Machinery.
- [2] Xiao Sun, Zongqing Lu, Wenjie Hu, and Guohong Cao, "Symdetector: Detecting sound-related respiratory symptoms using smartphones," in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, New York, NY, USA, 2015, UbiComp '15, p. 97–108, Association for Computing Machinery.
- [3] Xuhai Xu, Ebrahim Nemati, Korosh Vatanparvar, Viswam Nathan, Tousif Ahmed, Md Mahbubur Rahman, Daniel McCaffrey, Jilong Kuang, and Jun Alex Gao, "Listen2cough: Leveraging end-to-end deep learning cough detection model to enhance lung health assessment using passively sensed audio," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 1, Mar. 2021.
- [4] Yanzhi Ren, Chen Wang, Jie Yang, and Yingying Chen, "Fine-grained sleep monitoring: Hearing your breathing with smartphones," in 2015 IEEE Conference on Computer Communications (INFOCOM). IEEE, 2015, pp. 1194–1202.
- [5] Jesús Monge-Álvarez, Carlos Hoyos-Barceló, Paul Lesso, and Pablo Casaseca-de-la Higuera, "Robust detection of audio-cough events using local hu moments," *IEEE Journal of Biomedical and Health Informatics*, vol. 23, no. 1, pp. 184–196, 2019.
- [6] Daniyal Liaqat, Salaar Liaqat, Jun Lin Chen, Tina Sedaghat, Moshe Gabel, Frank Rudzicz, and Eyal de Lara, "Coughwatch: Real-world cough detection using smartwatches," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 8333–8337.
- [7] Tobias Röddiger, Michael Beigl, Michael Hefenbrock, Daniel Wolffram, and Erik Pescara, *Detecting Episodes of Increased Cough Using Kinetic Earables*, p. 111–115, Association for Computing Machinery, New York, NY, USA, 2021.
- [8] Ebrahim Nemati, Shibo Zhang, Tousif Ahmed, Md Mahbubur Rahman, Jilong Kuang, and Alex Gao, "Coughbuddy: Multi-modal cough event detection using earbuds platform," in 2021 IEEE 17th International

Conference on Wearable and Implantable Body Sensor Networks (BSN). IEEE, 2021, pp. 1–4.

- [9] Shibo Zhang, Ebrahim Nemati, Tousif Ahmed, Md Mahbubur Rahman, Jilong Kuang, and Alex Gao, "A novel multi-centroid template matching algorithm and its application to cough detection," *arXiv preprint arXiv:2109.00630*, 2021.
- [10] Samsung, "All about your samsung earbuds' battery life," [Online]. Available: https: //www.samsung.com/us/support/answer/ ANS00078081/, Accessed: 2021-09-23.
- [11] Terry T. Um, Franz M. J. Pfister, Daniel Pichler, Satoshi Endo, Muriel Lang, Sandra Hirche, Urban Fietzek, and Dana Kulić, "Data augmentation of wearable sensor data for parkinson's disease monitoring using convolutional neural networks," in *Proceedings of the 19th* ACM International Conference on Multimodal Interaction, New York, NY, USA, 2017, ICMI 2017, pp. 216– 220, ACM.
- [12] Yunqiang Chen, Xiang Sean Zhou, and T.S. Huang, "One-class svm for learning in image retrieval," in *Proceedings 2001 International Conference on Image Processing (Cat. No.01CH37205)*, 2001, vol. 1, pp. 34–37 vol.1.
- [13] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson, "Estimating the support of a high-dimensional distribution," *Neural computation*, vol. 13, no. 7, pp. 1443–1471, 2001.