



Transfer Learning to Detect COVID-19 Coughs with Incremental Addition of Patient Coughs to Healthy People's Cough Detection Models

Sudip Vhaduri¹(✉), Seungyeon Paik¹, and Jessica E. Huber^{2,3}

¹ Computer and Information Technology Department, Purdue University,
West Lafayette, IN 47907, USA

{svhaduri,paiks}@purdue.edu

² Speech, Language, and Hearing Sciences Department, Purdue University,
West Lafayette, IN 47907, USA

jhuber@purdue.edu

³ Communicative Disorders and Sciences Department, University at Buffalo, Buffalo,
NY 14214, USA

jehuber@buffalo.edu

Abstract. Millions of people have died worldwide from COVID-19. In addition to its high death toll, COVID-19 has led to unbearable suffering for individuals and a huge global burden to the healthcare sector. Therefore, researchers have been trying to develop tools to detect symptoms of this human-transmissible disease remotely to control its rapid spread. Coughing is one of the common symptoms that researchers have been trying to detect objectively from smartphone microphone-sensing. While most of the approaches to detect and track cough symptoms rely on machine learning models developed from a large amount of patient data, this is not possible at the early stage of an outbreak. In this work, we present an incremental transfer learning approach that leverages the relationship between healthy peoples' coughs and COVID-19 patients' coughs to detect COVID-19 coughs with reasonable accuracy using a pre-trained healthy cough detection model and a relatively small set of patient coughs, reducing the need for large patient dataset to train the model. This type of model can be a game changer in detecting the onset of a novel respiratory virus.

Keywords: transfer learning · COVID-19 · Cough detection

1 Introduction

1.1 Motivation

Coronavirus and other human-transmissible respiratory viruses have become prevalent and have led to human suffering and a large number of deaths in recent times. According to the World Health Organization (WHO), the novel coronavirus SARS-CoV-2 (COVID-19) has so far caused a total of over 771 million infections and over 6.9 million deaths globally [8]. Even after the development of a vaccine, over 300 thousand infections and 1.5 thousand deaths occur a day [8]. Additionally, COVID-19 created a heavy economic burden on the health sectors, e.g., the United States incurred a total of \$163.4 billion in direct medical expenses during the pandemic [40]. Early onset detection can help prevent the rapid spread and its adverse consequences. But, traditional diagnosis approaches are slow and require resources, such as viral tests (based on samples from the nose and mouth) or antibody tests [3], chest X-ray or spirometry tests [4], blood tests, pulse oximetry, and sputum tests [1, 7]. These resources are not readily available in peoples' homes or at healthcare access points, such as primary care or urgent care. Therefore, there is a need for an approach that can be easily deployed to quickly detect the onset and control disease spread.

1.2 Related Work

Researchers have been trying to develop tools/systems to objectively detect and remotely report typical symptoms of respiratory diseases, such as coughing. Many of these techniques require the use of wearable technology. For example, researchers have detected coughing with 0.82 accuracy using smartwatch accelerometers and audio recordings [31], and 0.94–0.95 sensitivity using ECG, thermistor, chest belt, accelerometer, contact microphone, audio microphone [23] and chest sensor [11]. Some researchers proposed a respiratory monitoring system using a wearable patch sensor [24] and a wearable radio-frequency (RF) cough monitoring system [26]. On the other hand, a group of researchers proposed a COVID-19 symptom tracker utilizing a headset-like sensor [45].

However, we have found that people's adherence to wearables drops significantly over time, compared to smartphone adherence [9, 28]. Therefore, some researchers have been trying to detect objective symptoms, such as coughing, using smartphone data [12, 13, 20, 48, 56]. One team of researchers proposed machine learning-based COVID-19 cough, breath, and speech detection using smartphone recording files [37, 38]. They achieved up to 0.93 area under the curve using the k-nearest neighbor classifiers. Other researchers developed frameworks to diagnose COVID-19 using a smartphone app and built-in sensors in the smartphone [27, 30].

However, a major limitation of all these existing cough detection models is the underlying assumption of the availability of a good amount of relevant data, which is not always possible [13]. For example, during the early stage of a new outbreak, there is not much data to develop a reasonably good traditional

machine learning model due to the need for a large volume of data. But coughs from healthy people and patients have similarities, which can be utilized to detect COVID-19 coughs using a healthy cough detection model and a relatively small set of coughs from COVID-19 patients, reducing the need for large amounts of COVID-19 patient coughs. This kind of model would be invaluable in the detection of a new novel respiratory disease.

1.3 Contribution

The main contribution of this work is to present a novel approach leveraging similarities between healthy people's coughs and COVID-19 patients' coughs to incrementally transfer healthy cough detection models to COVID-19 cough detection models with smaller batches or folds of COVID-19 coughs. Compared to large data-driven traditional modeling approaches, incremental transfer learning approaches can help detect the onset of a novel respiratory virus early utilizing a relatively small set of cough samples obtained from the first few people infected and a pre-trained healthy people's cough detection model to control the spread of the novel respiratory disease to minimize adverse consequences.

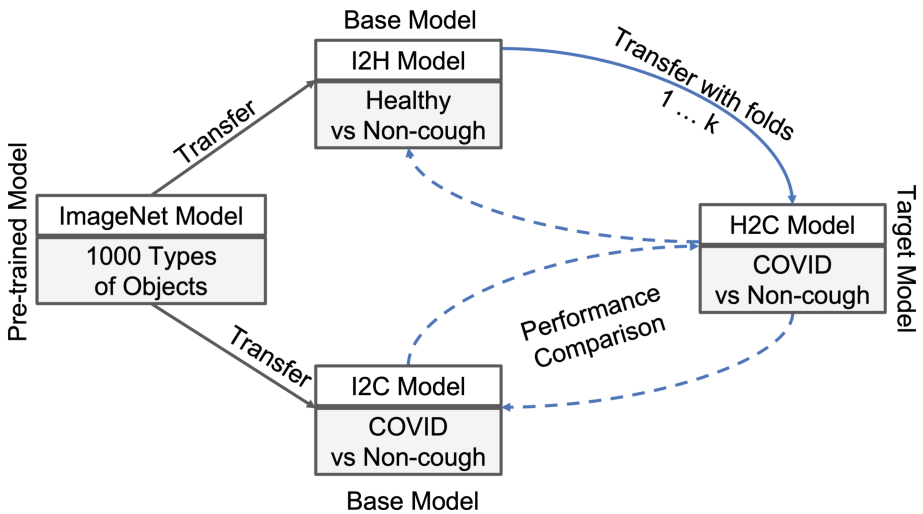


Fig. 1. Proposed modeling approach

2 Methods

In this section, we first present our modeling approach. Next, we discuss cough and non-cough audio recordings obtained from different sources, data processing steps, and neural network architectures and parameters used in this work.

2.1 Modeling Approaches

In Fig. 1, we present our modeling scheme. We start with a pre-trained model [44] developed to detect 1000 objects, i.e., classes from the ImageNet dataset [6]. Next, we use transfer learning to develop two base models and one target model to detect coughs obtained from healthy people and COVID-19 patients using transfer learning.

1. ImageNet to Healthy (**I2H**) model: We transfer our ImageNet 1000 object detection model to a binary model detecting healthy cough versus non-cough. Throughout this manuscript, we name this model as “ImageNet to healthy”, i.e., **I2H** model. This is one of the two base models. This model will be later used to develop the target COVID-19 cough detection model incrementally.
2. ImageNet to COVID-19 (**I2C**) model: We transfer the ImageNet 1000 object detection model to a binary model that detects COVID-19 cough versus non-cough, which we refer to as “ImageNet to COVID-19”, i.e., **I2C** model in this manuscript. This is the second base model, which will be used as a reference model when comparing the performance of our target COVID-19 cough detection models.
3. Healthy to COVID-19 (**H2C**) model: We transfer the healthy cough detection model, i.e., **I2H** model, incrementally by adding smaller batches, i.e., folds of COVID-19 coughs to detect COVID-19 coughs. This target model is named “Healthy to COVID-19”, i.e., **H2C** model.

Our ultimate goal is to utilize the **I2H** model and smaller batches/folds of COVID-19 coughs to develop a target **H2C** model that achieves close performance to the **I2C** base model to investigate the capability and feasibility of incremental transfer learning. The incremental fold addition is the core of our transfer learning approach. Model performance is analyzed to show how the target models improve over time and to determine the minimum number of samples required to identify disease coughs accurately. There may not be a base model for a new disease for the performance comparison, such as the **I2C** model. However, our approach is to demonstrate that incremental transfer learning can be a pathway to get a valid model even when we have few samples of the new disease-specific patient data.

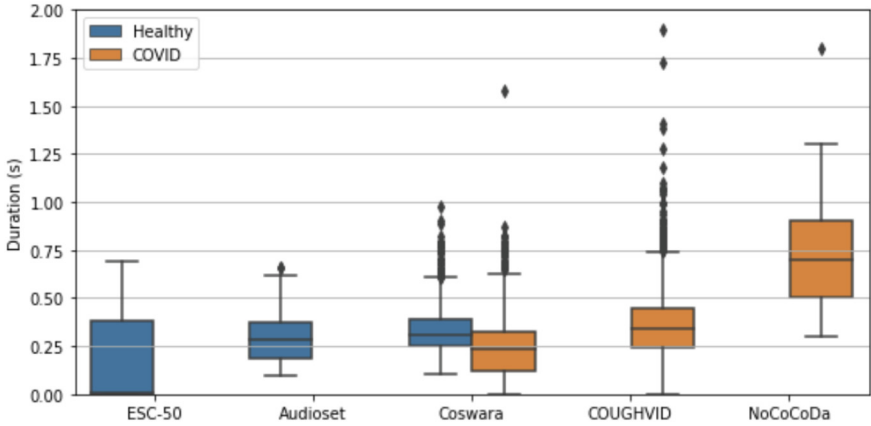


Fig. 2. Distribution of cough duration

2.2 Datasets

This section briefly describes six major datasets used in our experiments, including three COVID-19, two healthy cough datasets, and an image dataset.

1. ImageNet Dataset [6]: This is a publicly available image dataset with a thousand object classes and millions of images. The dataset contains various types of objects types, including geese, balloons, and fruits. We use the ImageNet dataset to develop the pre-trained “ImageNet” model.
2. Coswara Dataset [41]: This crowdsourced dataset consists of breathing, coughing, and voice sound recordings from healthy people and COVID-19 patients. The sampling rate is in the range of 47.82 ± 0.83 kHz. In this work, we use the voluntary coughs obtained from 274 COVID-19-positive patients.
3. COUGHVID Dataset [36]: This crowdsourced dataset contains more than 25,000 cough recordings from COVID-19-positive patients or asymptomatic people of varying ages, genders, and regions. The sampling rate of this dataset is around 44.1 kHz. We use the voluntary coughs obtained from 719 COVID-19-positive patients or asymptomatic people.
4. NoCoCoDa Dataset [17]: This dataset contains natural cough recordings of 13 COVID-19-positive patients collected from public media interviews. Audio recordings are collected at a 44.1 kHz sampling rate.
5. ESC-50 Dataset [39]: The environmental sound classification (i.e., ESC-50) dataset consists of audio recordings from five categories of sounds (i.e., animal, natural soundscapes, human sounds, interior sounds, and exterior noises) with 10 types of sounds per category recorded at a rate of 44.1 kHz. There are 40 audio recordings per type (2,000 recordings in total). We use the voluntary cough recordings from five healthy subjects as the cough class, i.e.,

healthy coughs. The remaining 49 sound types are used to create the non-cough class when developing the binary cough versus non-cough models presented in Sect. 2.1.

6. AudioSet Dataset [25]: This dataset contains a wide range of 632 sound classes obtained from YouTube videos, where samples are recorded at 16 kHz and 44.1 kHz. In this work, we use voluntary coughs from 88 healthy subjects.

In this work, we combined all COVID-19 data to create one COVID-19 patient cough dataset ($n = 1006$ patients). Similarly, we combined all healthy cough datasets to develop another cough dataset ($n = 83$ healthy people). While combining data from different datasets, we keep the subject information so that they can be utilized later to create mutually exclusive splits or folds among the training, validation, and test sets. All non-cough data is obtained from the ESC-50 dataset.

2.3 Data Processing

This section presents our cough ground-truth label collection approach from long audio recordings, followed by additional processing steps, including finding an optimal window size, padding, and feature extraction. Finally, we present our cross-validation approach with incremental training to develop **H2C** models from **I2H** model which requires relatively fewer COVID-19 coughs.

Data Cleaning and Ground-Truth Label Collection. Most of the audio data used in this work come from crowd-sourced datasets. Therefore, we first perform a rigorous cleaning process to remove unwanted parts, including quiet, speech, and noises using different audio signal processing libraries and tools, including the Audacity toolbox [2].

After the initial cleaning, we performed data segmentation to extract the ground-truth cough labels from long audio recordings with multiple coughs utilizing the Audacity toolbox. Adapted from our previous work [46, 74], we automated the process by developing an energy threshold-based audio segmentation followed by a phase classification approach.

We obtained 144 and 252 healthy coughs from the ESC-50 and AudioSet datasets. Similarly, we obtained 1892, 2690, and 73 COVID-19 coughs from the Coswara, COUGHVID, and NoCoCoDa datasets, respectively. In total, we used 396 healthy coughs and 4655 COVID-19 coughs.

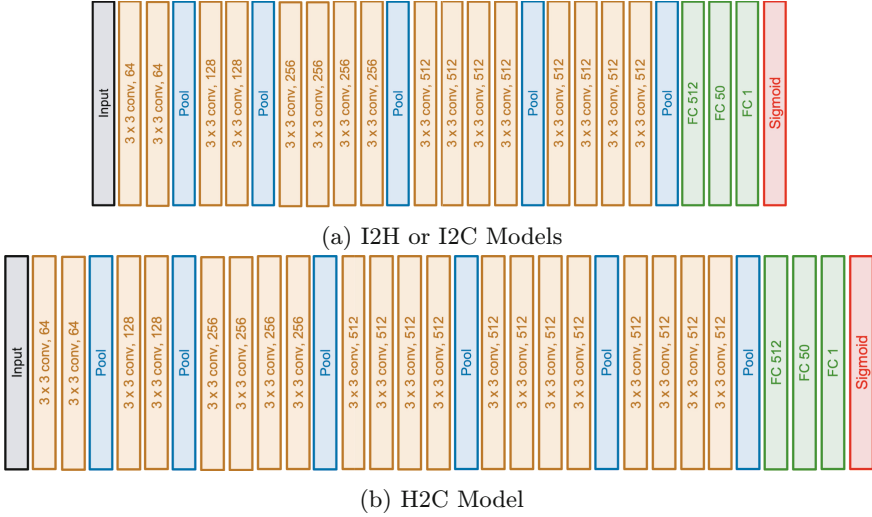


Fig. 3. Optimal Architecture of different models

Feature Extraction. Since we collected audio recordings from different sources, we first changed all sampling frequencies to 44.1 kHz frequency before additional processing.

Next, we determined a suitable window size before computing features. In Fig. 2, we use boxplots to represent the distribution of cough duration obtained from various datasets. In the figure, we find that 75th percentile of Coswara healthy cough duration, i.e., 0.3917s, is a suitable choice for window size since most of the other healthy and patient coughs, except the NoCoCoDa coughs, have a duration shorter than that. Compared to other datasets, the NoCoCoDa dataset has relatively fewer samples.

We perform padding (i.e., add 0s) at the end of coughs shorter than the window size (i.e., 0.3917s). In the case of longer cough and non-cough audio recordings, we truncate the parts longer than the window size. Finally, we compute the log mel-spectrogram [34] (logmel) and Mel-frequency cepstral coefficient (mfcc) image features from each cough and non-cough audio recording using the Python Librosa library [33].

Training, Validation, and Test Split. First, we uniformly split the cough data into 10 mutually exclusive folds based on subjects. For class balancing, we select the same number of non-cough samples as we have cough samples for the 10 folds. Thereby, we maintain the same number of cough (either healthy or COVID-19) and non-cough instances when training, validating, and testing binary cough (either healthy or COVID-19) versus non-cough detection models presented in Sect. 2.1. Next, we follow a “leave-2-fold-out” for validation and “leave-2-fold-out” for the test approach, where we use the remaining six folds

for training while developing and validating/testing a **I2H** or **I2C** model from the pre-trained ImageNet Model. We follow a rotational approach, where we perform this mutually exclusive 6-2-2 train-validation-test split 10 times to develop 10 different models. While developing **H2C** models, we follow an incremental training approach, where we add the six training folds of COVID-19 data to the **H2C** models one-by-one. We perform this incremental training for one of the 10 6-2-2 splits. During the incremental training, we distribute COVID-19 subjects in 10 folds using a snaking approach, where we first sort the subjects in descending order of cough counts. Then, we put the top 10 subjects into 10 folds. Next, 10 subjects are put into 10 folds. This way, we maximize the number of cough samples in 10 folds. We end up with 100 subjects in each fold with 450 random COVID-19 cough samples.

2.4 Neural Network Architectures and Parameters

In Fig. 3a, we present the optimal architecture of the two models (i.e., **I2H** or **I2C** models) transferred from the pre-trained ImageNet model developed with VGG19 [44]. In Fig. 3b, we present the optimal architecture of the **H2C** model transferred from the **I2H** model by adding folds of COVID-19 coughs. In the figure, the sequence and meaning of different parameters in each layer are kernel size and the number of feature maps, i.e., nodes in each layer. The pre-trained ImageNet VGG19 model had fully connected (FC) three layers with 4096, 4096, and 1000 nodes, followed by a softmax layer to classify 1000 objects. We changed the last four layers with FC 512, 50, and 1 node, followed by a sigmoid layer for binary classification of healthy cough versus non-cough (**I2H** model) or COVID-19 cough versus non-cough (**I2C** model). Compared to the **I2H** or **I2C** model, we add four additional convolutional layers and a pooling layer in the case of **H2C** models.

We used TensorFlow and Keras libraries to develop our models. We used the ReLU activation function in the hidden layers and the sigmoid decision function in the final layer for every model. For the loss function, we used binary cross-entropy. We tried Adam and RMSprop optimizer and found Adam achieves 53% higher accuracy than the RMSprop. We also tried batch sizes 16 and 32 and found that batch size 16 is more accurate and has a shorter execution time. The input size used in this work is (320,320). We considered a range of learning rates, including 0.00001, 0.0001, and 0.0005, and epochs ranging from 30 to 200. We found a learning rate of 0.0001 works better for the **I2H** and **I2C** models. Similarly, we found a learning rate of 0.00001 is a good compromise for the **H2C** models. We developed models on Purdue University’s Gilbreth GPU server with 8 GPU nodes and 16 cores per node [5]. Each node has 192 GB memory, 100 Gbps Infiniband interconnects, and 2 P100 GPUs.




Fig. 4. **I2H** and **I2C** models tested on different coughs (H-H refers to the case when **I2H** models are tested on healthy coughs; similarly, H-C and C-C refer to the cases when **I2H** and **I2C** models are tested on COVID-19 coughs)

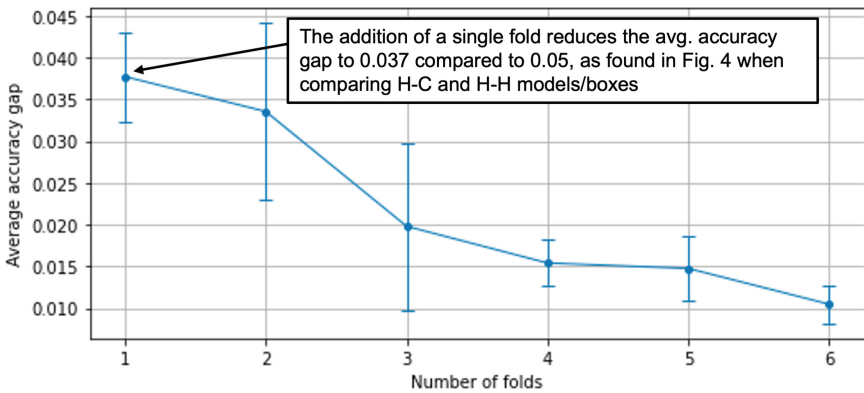


Fig. 5. Performance of **H2C** models developed using transfer learning incrementally by adding folds of COVID-19 coughs to **I2H** models

3 Results

Since we use the mutually exclusive folds with the same number of cough and non-cough instances in each fold (Sect. 2.3) when developing binary models, our classes are always balanced, and we use classification accuracy when comparing the performance of different models.

In Fig. 4, we present the performance of the two base models, i.e., **I2H** and **I2C** models when tested on different cough datasets. First, we analyze the performance of the **I2H** base models, i.e., healthy cough detection models transferred from the pre-trained ImageNet models using healthy people’s coughs. The **I2H** base models are tested on the healthy test coughs, and the findings are presented using the Healthy-Healthy or H-H box in the figure. We observe average accuracy values of 0.93 ± 0.02 (logmel) and 0.92 ± 0.05 (mfcc) with median values of 0.93 (logmel) and 0.93 (mfcc).

Next, we analyze the performance of the second base model, i.e., **I2C** model (COVID-19 cough detection models transferred from the pre-trained ImageNet models using COVID-19 patient coughs). The **I2C** model is tested on the COVID-19 test coughs, and the findings are presented using the COVID-COVID or C-C box in the figure. We observe average accuracy values of 0.96 ± 0.02 (logmel) and 0.95 ± 0.01 (mfcc) with median values of 0.96 (logmel) and 0.95 (mfcc).

Next, we analyze the performance of the **I2H** models (i.e., base models to detect healthy people’s coughs) when tested on COVID-19 test coughs and present the findings using the Healthy-COVID or H-C box in the figure. We observe average accuracy values of 0.88 ± 0.07 (logmel) and 0.86 ± 0.03 (mfcc) with median values of 0.91 (logmel) and 0.86 (mfcc). In the case of logmel feature, the drop in average accuracy is 0.05 and 0.08 when compared with the findings in the Healthy-Healthy (H-H) and COVID-COVID (C-C) boxes, respectively. In the case of mfcc feature, the drop in average accuracy is 0.06 and 0.09 when compared with the findings in the Healthy-Healthy (H-H) and COVID-COVID (C-C) boxes, respectively. To improve the accuracy values of the H-C box (i.e., **I2H** models tested on COVID-19 test coughs), we incrementally developed **H2C** models transferring the **I2H** models by adding small amounts of COVID-19 coughs in folds. Since logmel features outperform the mfcc features across all measures, we consider the logmel features in the next analysis.

In Fig. 5, we present the accuracy gap of **H2C** models with respect to the average accuracy of the base **I2C** models by varying amounts of COVID-19 coughs (i.e., number of folds) included in the training set. In the figure, we observe that with the addition of only one fold of COVID-19 coughs, the average accuracy gap drops to 0.037.

The average accuracy gap drops to 0.02 by adding two more folds of COVID-19 coughs to the base **I2H** models. Thereby, with the addition of three folds of COVID-19 coughs with the healthy cough detection models, i.e., **I2H** models, we can achieve a performance close to that of the base COVID-19 detection model performance.

As we continue adding more folds of COVID-19 coughs to the base **I2H** models, we witness a drop in accuracy gap, and after adding all six folds of COVID-19 patient coughs, the accuracy gap drops to 0.01. Additionally, adding more folds makes the error bar tighter, reflecting more consistent accuracy values. Thereby, using this incremental transfer learning approach, we can develop COVID-19 detection models, such as **H2C** models, from the base healthy people’s cough detection models, i.e., **I2H** models and smaller amounts of COVID-19 coughs to

achieve very close performance to that of the base COVID-19 detection models, i.e., **I2C** models.

4 Conclusion and Discussion

This work attempts to utilize the power of transfer learning and similarities between two types of coughs, i.e., healthy and COVID-19 coughs, to incrementally develop new models requiring a relatively small set of patient coughs to achieve similar performance to that of the COVID-19 cough detection models trained from a bigger COVID-19 cough dataset. Our findings show the promise of utilizing healthy cough detection models to detect COVID-19 coughs after training with relatively fewer patient coughs.

This model can be useful to detect an early-onset novel respiratory virus with a smaller amount of relevant data. However, before generalizing the findings to similar or other problems, extended studies with a diverse population, diseases, and stages will be needed. While image feature-based transfer learning has been adopted in this feasibility work, in the future, other types of data, e.g., acoustic signals, can be utilized to adapt transfer learning models and can be compared with this feasibility work. This work and findings will also impact other domains of predictive modeling, including place of importance discovery [59, 65, 67, 70, 71], health condition monitoring [14, 42, 43, 57, 66, 72, 73] and well-being tracking [21, 29, 47, 50, 52, 55, 58, 60, 61], securing a user's cyber-physical space [10, 15, 16, 18, 19, 22, 32, 35, 49, 51, 53, 54, 62–64, 68, 69], as it presents the feasibility of developing predictive models with relatively small datasets to alternate the traditional approaches requiring large-scale datasets.

Acknowledgement. The authors would like to thank the Clifford B. Kinley Trust for funding this research.

References

1. Asthma: Steps in testing and diagnosis - Mayo Clinic. <https://mayoclinic.org/3vPs3J7>. Accessed Jan 2023
2. Audacity: Free, open source, cross-platform audio software. <https://www.audacityteam.org/>. Accessed Jan 2023
3. CDC: COVID-19 Testing. <https://bit.ly/3nRjYOM>. Accessed Jan 2023
4. COPD Symptoms and Diagnosis | American Lung Association. <https://bit.ly/3hefi2f>. Accessed Jan 2023
5. Gilbreth. <https://www.rcac.purdue.edu/compute/gilbreth>. Accessed Jan 2023
6. ImageNet. <https://image-net.org/download.php/>. Accessed Jan 2023
7. Pneumonia | Disease or Condition of the Week | CDC. <https://bit.ly/35aOE7Q>. Accessed Jan 2023
8. WHO Coronavirus (COVID-19) Dashboard. <https://covid19.who.int/>. Accessed Jan 2023
9. Aditi Pai: Survey: One third of wearable device owners stopped using them within six months. <https://bit.ly/3yjuzrC>. Accessed Jan 2023

10. Al Amin, M.T., Barua, S., Vhaduri, S., Rahman, A.: Load aware broadcast in mobile ad hoc networks. In: IEEE International Conference on Communications (ICC) (2009)
11. Amoh, J., Odame, K.: Deepcough: a deep convolutional neural network in a wearable cough detection system. In: 2015 IEEE Biomedical Circuits and Systems Conference (BioCAS), pp. 1–4. IEEE (2015)
12. Cai, J., Vhaduri, S., Luo, X.: Discovering covid-19 coughing and breathing patterns from unlabeled data using contrastive learning with varying pre-training domains. In: INTERSPEECH (2023)
13. Chang, Y., Jing, X., Ren, Z., Schuller, B.W.: Covnet: a transfer learning framework for automatic covid-19 detection from crowd-sourced cough sounds. *Front. Digit. Health* **3**, 799067 (2022)
14. Chen, C.Y., Vhaduri, S., Poellabauer, C.: Estimating sleep duration from temporal factors, daily activities, and smartphone use. In: IEEE Computer Society Computers, Software, and Applications Conference (COMPSAC) (2020)
15. Cheung, W., Vhaduri, S.: Context-dependent implicit authentication for wearable device users. In: IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC) (2020)
16. Cheung, W., Vhaduri, S.: Continuous authentication of wearable device users from heart rate, gait, and breathing data. In: IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechanics (BioRob) (2020)
17. Cohen-McFarlane, M., Goubran, R., Knoefel, F.: Novel coronavirus cough database: Nococoda. *IEEE Access* **8**, 154087–154094 (2020)
18. Dibbo, S.V.: SoK: model inversion attack landscape: taxonomy, challenges, and future roadmap. In: 2023 IEEE 36th Computer Security Foundations Symposium (CSF), pp. 439–456. IEEE (2023)
19. Dibbo, S.V., Cheung, W., Vhaduri, S.: On-phone CNN model-based implicit authentication to secure IoT wearables. In: EAI International Conference on Safety and Security in Internet of Things (SaSeIoT) (2021)
20. Dibbo, S.V., Kim, Y., Vhaduri, S.: Effect of noise on generic cough models. In: IEEE International Conference on Wearable and Implantable Body Sensor Networks (BSN) (2021)
21. Dibbo, S.V., Kim, Y., Vhaduri, S., Poellabauer, C.: Visualizing college students' geo-temporal context-varying significant phone call patterns. In: 2021 IEEE 9th International Conference on Healthcare Informatics (ICHI), pp. 381–385. IEEE (2021)
22. Dibbo, S.V., Moore, J.S., Kenyon, G.T., Teti, M.A.: Lcanets++: robust audio classification using multi-layer neural networks with lateral competition. *arXiv preprint [arXiv:2308.12882](https://arxiv.org/abs/2308.12882)* (2023)
23. Drugman, T., et al.: Objective study of sensor relevance for automatic cough detection. *IEEE J. Biomed. Health Inform.* **17**(3), 699–707 (2013)
24. Elfaramawy, T., Fall, C.L., Arab, S., Morissette, M., Lellouche, F., Gosselin, B.: A wireless respiratory monitoring system using a wearable patch sensor network. *IEEE Sens. J.* **19**(2), 650–657 (2018)
25. Gemmeke, J.F., et al.: Audio set: an ontology and human-labeled dataset for audio events. In: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 776–780. IEEE (2017)
26. Hui, X., Zhou, J., Sharma, P., Conroy, T.B., Zhang, Z., Kan, E.C.: Wearable RF near-field cough monitoring by frequency-time deep learning. *IEEE Trans. Biomed. Circuits Syst.* **15**(4), 756–764 (2021)

27. Imran, A., et al.: Ai4covid-19: AI enabled preliminary diagnosis for covid-19 from cough samples via an app. *Inform. Med. Unlocked* **20**, 100378 (2020)
28. Kamei, T., Kanamori, T., Yamamoto, Y., Edirippulige, S.: The use of wearable devices in chronic disease management to enhance adherence and improve telehealth outcomes: a systematic review and meta-analysis. *J. Telemed. Telecare* **28**(5), 342–359 (2022)
29. Kim, Y., Vhaduri, S., Poellabauer, C.: Understanding college students' phone call behaviors towards a sustainable mobile health and wellbeing solution. In: *International Conference on Systems Engineering* (2020)
30. Laguarda, J., Hueto, F., Subirana, B.: Covid-19 artificial intelligence diagnosis using only cough recordings. *IEEE Open J. Eng. Med. Biol.* **1**, 275–281 (2020)
31. Liaqat, D., et al.: Coughwatch: real-world cough detection using smartwatches. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8333–8337. IEEE (2021)
32. Lien, C.W., Vhaduri, S.: Challenges and opportunities of biometric user authentication in the age of IoT: a survey. *ACM Comput. Surv.* **55**(12), 1–39 (2023)
33. McFee, B., et al.: Librosa: audio and music signal analysis in python. In: *Proceedings of the 14th Python in Science Conference*, vol. 8, pp. 18–25 (2015)
34. Meng, H., Yan, T., Yuan, F., Wei, H.: Speech emotion recognition from 3D log-mel spectrograms with deep learning network. *IEEE Access* **7**, 125868–125881 (2019)
35. Muratyan, A., Cheung, W., Dibbo, S.V., Vhaduri, S.: Opportunistic multi-modal user authentication for health-tracking IoT wearables. In: *EAI International Conference on Safety and Security in Internet of Things (SaSeIoT)* (2021)
36. Orlandic, L., Teijeiro, T., Atienza, D.: The coughvid crowdsourcing dataset, a corpus for the study of large-scale cough analysis algorithms. *Sci. Data* **8**(1), 1–10 (2021)
37. Pahar, M., Klopper, M., Warren, R., Niesler, T.: Covid-19 cough classification using machine learning and global smartphone recordings. *Comput. Biol. Med.* **135**, 104572 (2021)
38. Pahar, M., Niesler, T.: Machine learning based covid-19 detection from smartphone recordings: cough, breath and speech. *arXiv pre-print* (2021)
39. Piczak, K.J.: ESC: dataset for environmental sound classification. In: *Proceedings of the 23rd ACM International Conference on Multimedia*, pp. 1015–1018 (2015)
40. Richards, F., et al.: Economic burden of covid-19: a systematic review. *ClinicoEconomics Outcomes Res. CEOR* **14**, 293 (2022)
41. Sharma, N., et al.: Coswara—a database of breathing, cough, and voice sounds for covid-19 diagnosis. *arXiv preprint* [arXiv:2005.10548](https://arxiv.org/abs/2005.10548) (2020)
42. Sharmin, M., et al.: Visualization of time-series sensor data to inform the design of just-in-time adaptive stress interventions. In: *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pp. 505–516 (2015)
43. Simhadri, S., Vhaduri, S.: Understanding user trust in different recommenders and smartphone applications. In: *EAI International Conference on Wireless Mobile Communication and Healthcare (MobiHealth)* (2022)
44. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint* [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
45. Stojanović, R., Škraba, A., Lutovac, B.: A headset like wearable device to track covid-19 symptoms. In: *2020 9th Mediterranean Conference on Embedded Computing (MECO)*, pp. 1–4. IEEE (2020)

46. Vhaduri, S.: Nocturnal cough and snore detection using smartphones in presence of multiple background-noises. In: ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS) (2020)
47. Vhaduri, S., Ali, A., Sharmin, M., Hovsepian, K., Kumar, S.: Estimating drivers' stress from GPS traces. In: International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI) (2014)
48. Vhaduri, S., Brunschwiler, T.: Towards automatic cough and snore detection. In: IEEE International Conference on Healthcare Informatics (ICHI) (2019)
49. Vhaduri, S., Cheung, W., Dibbo, S.V.: Bag of on-phone ANNs to secure IoT objects using wearable and smartphone biometrics. *IEEE Trans. Dependable Secure Comput.* **20**(3), 1–12 (2023)
50. Vhaduri, S., Cho, J., Meng, K.: Predicting unreliable response patterns in smartphone health surveys: a case study with the mood survey. *Smart Health J.* **28**, 100398 (2023)
51. Vhaduri, S., Dibbo, S.V., Chen, C.Y.: Predicting a user's demographic identity from leaked samples of health-tracking wearables and understanding associated risks. In: 2022 IEEE 10th International Conference on Healthcare Informatics (ICHI). IEEE (2022)
52. Vhaduri, S., Dibbo, S.V., Chen, C.Y., Poellabauer, C.: Predicting next call duration: a future direction to promote mental health in the age of lockdown. In: IEEE Computer Society Computers, Software, and Applications Conference (COMPSAC) (2021)
53. Vhaduri, S., Dibbo, S.V., Cheung, W.: HIAuth: a hierarchical implicit authentication system for IoT wearables using multiple biometrics. *IEEE Access* **9**, 116395–116406 (2021)
54. Vhaduri, S., Dibbo, S.V., Cheung, W.: Implicit IoT authentication using on-phone ANN models and breathing data. *Internet Things* **24** (2023)
55. Vhaduri, S., Dibbo, S.V., Kim, Y.: Deriving college students' phone call patterns to improve student life. *IEEE Access* **9**, 96453–96465 (2021)
56. Vhaduri, S., Dibbo, S.V., Kim, Y.: Environment knowledge-driven generic models to detect coughs from audio recordings. *IEEE Open J. Eng. Med. Biol.* **4**, 1–12 (2023)
57. Vhaduri, S., Munch, A., Poellabauer, C.: Assessing health trends of college students using smartphones. In: IEEE Healthcare Innovation Point-of-Care Technologies Conference (HI-POCT) (2016)
58. Vhaduri, S., Poellabauer, C.: Design and implementation of a remotely configurable and manageable well-being study. In: EAI SWIT-Health (2015)
59. Vhaduri, S., Poellabauer, C.: Cooperative discovery of personal places from location traces. In: International Conference on Computer Communication and Networks (ICCCN) (2016)
60. Vhaduri, S., Poellabauer, C.: Human factors in the design of longitudinal smartphone-based wellness surveys. In: IEEE International Conference on Healthcare Informatics (ICHI) (2016)
61. Vhaduri, S., Poellabauer, C.: Design factors of longitudinal smartphone-based health surveys. *J. Healthc. Inform. Res.* **1**(1), 52–91 (2017)
62. Vhaduri, S., Poellabauer, C.: Towards reliable wearable-user identification. In: 2017 IEEE International Conference on Healthcare Informatics (ICHI) (2017)
63. Vhaduri, S., Poellabauer, C.: Wearable device user authentication using physiological and behavioral metrics. In: IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC) (2017)

64. Vhaduri, S., Poellabauer, C.: Biometric-based wearable user authentication during sedentary and non-sedentary periods. *International Workshop on Security and Privacy for the Internet-of-Things (IoTSec)* (2018)
65. Vhaduri, S., Poellabauer, C.: Hierarchical cooperative discovery of personal places from location traces. *IEEE Trans. Mob. Comput.* **17**(8), 1865–1878 (2018)
66. Vhaduri, S., Poellabauer, C.: Impact of different pre-sleep phone use patterns on sleep quality. In: *IEEE International Conference on Wearable and Implantable Body Sensor Networks (BSN)* (2018)
67. Vhaduri, S., Poellabauer, C.: Opportunistic discovery of personal places using smartphone and fitness tracker data. In: *IEEE International Conference on Healthcare Informatics (ICHI)* (2018)
68. Vhaduri, S., Poellabauer, C.: Multi-modal biometric-based implicit authentication of wearable device users. *IEEE Trans. Inf. Forensics Secur.* **14**(12), 3116–3125 (2019)
69. Vhaduri, S., Poellabauer, C.: Summary: Multi-modal Biometric-based Implicit Authentication of Wearable Device Users. *arXiv preprint [arXiv:1907.06563](https://arxiv.org/abs/1907.06563)* (2019)
70. Vhaduri, S., Poellabauer, C.: Opportunistic discovery of personal places using multi-source sensor data. *IEEE Trans. Big Data* **7**(2), 383–396 (2021)
71. Vhaduri, S., Poellabauer, C., Striegel, A., Lizardo, O., Hachen, D.: Discovering places of interest using sensor data from smartphones and wearables. In: *IEEE Ubiquitous Intelligence & Computing (UIC)* (2017)
72. Vhaduri, S., Prioleau, T.: Adherence to personal health devices: a case study in diabetes management. In: *EAI International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)* (2020)
73. Vhaduri, S., Simhadri, S.: Understanding user concerns and choice of app architectures in designing audio-based mHealth apps. *Smart Health J.* **26**, 100341 (2022)
74. Vhaduri, S., Van Kessel, T., Ko, B., Wood, D., Wang, S., Brunschwiler, T.: Nocturnal cough and snore detection in noisy environments using smartphone-microphones. In: *IEEE International Conference on Healthcare Informatics (ICHI)* (2019)