

# Audio-Based Deep Learning Frameworks for Detecting COVID-19

Dat Ngo<sup>1</sup>, Lam Pham<sup>2</sup>, Truong Hoang<sup>3</sup>, Şefki Kolozali<sup>1</sup>, Delaram Jarchi<sup>1</sup>

1. School of Computer Science and Electronic Engineering, University of Essex, United Kingdom.  
{dn22678, sefki.kolozali, delaram.jarchi}@essex.ac.uk
2. Center for Digital Safety & Security, Austrian Institute of Technology, Austria.  
lam.pham@ait.ac.at
3. Solution Technology Unit Department, FPT Software Company Limited, Vietnam.  
hoangvantruong1369@gmail.com

**Abstract**—This paper evaluates a wide range of audio-based deep learning frameworks applied to the breathing, cough, and speech sounds for detecting COVID-19. In general, the audio recording inputs are transformed into low-level spectrogram features, then they are fed into pre-trained deep learning models to extract high-level embedding features. Next, the dimension of these high-level embedding features are reduced before fine-tuning using Light Gradient Boosting Machine (LightGBM) as a back-end classification. Our experiments on the Second DiCOVA Challenge achieved the highest Area Under the Curve (AUC), F1 score, sensitivity score, and specificity score of 89.03%, 64.41%, 63.33%, and 95.13%, respectively. Based on these scores, our method outperforms the state-of-the-art systems, and improves the challenge baseline by 4.33%, 6.00% and 8.33% in terms of AUC, F1 score and sensitivity score, respectively.

**Index Terms**—low-level spectrogram feature, high-level embedding feature, pre-trained model, convolutional neural network.

## I. INTRODUCTION

The COVID-19 pandemic is now continuing to spread around the world, with more than 300 million confirmed cases, and causing more than five million deaths across almost 200 countries [1]. Although many countries are in relentless efforts of launching their vaccination programs, the number of global daily new cases is still increasing as COVID-19 restrictions are being eased in many countries and people can travel worldwide. To deal with the rapid spread of infection across populations, many countries have to conduct massive daily tests for their citizens to control the virus spreads. Particularly, the molecular testing approaches, namely the reverse transcription polymerase chain reaction test (RT-PCR) [2] and the rapid antigen test (RAT) [3], are now widely applied as primary testing methodologies in most countries. However, these methodologies present various limitations as their procedure of sample collection violates physical distancing in many countries that the ordering online lateral test at home is limited. Additionally, analysing and receiving results require high-cost, chemical equipment, and involving labour-intensive tasks. As a result, there is an urgent need for non-invasive, scalable, and cost-effective tool to detect infected individuals in a decentralized manner. As COVID-19 disease is related

to primary symptoms such as fever, sore throat, cough, chest pain, etc. Therefore, many researchers and practitioners are motivated to use acoustic-related signals such as breathing, cough, speech from human respiratory system to early detect COVID-19 and non-COVID-19 individuals [4]. Indeed, calls for development of diagnostic tools were announced in the Interspeech 2021 as a special session titled ‘Diagnostics of COVID-19 using Acoustics (DiCOVA) Challenge’ as well as in ICASSP 2022 as a calling paper, referred to as the First DiCOVA and the Second DiCOVA Challenges, respectively.

In this paper, we aim to explore all types of human respiratory sounds: breathing, cough, speech, provided by the recent Second DiCOVA Challenge, and then propose a robust framework for detecting COVID-19. Our contributions include: (1) to conduct extensive experiments and pinpoint the most effective approach related to the extraction of well-represented features for each type of breathing, cough, speech sound input; (2) to evaluate how oversampling on positive samples and dimension reduction on represented features affect the system performance; and (3) to demonstrate that our proposed framework is reliable and robust for detecting COVID-19 which outperforms the state-of-the-art systems and is potential for a real-life application.

## II. DATASET AND TASKS DEFINED

In this paper, we evaluate the dataset derived from the Second DiCOVA Challenge [5]. This dataset provides audio recordings of three different sound categories: breathing, cough, speech which were collected from both COVID-19 positive and negative patients in the age group of 15 to 45 years old. For each sound category, there are a total of 1436 acoustic samples which are then divided into the Development set (965 subjects) and the Test set (471 subjects) for training and testing processes, respectively. It is of note that Development set in each categories has recordings of 172 individuals positive to COVID-19, while there is a higher number of recordings from 793 individual negative to COVID-19, presenting an imbalanced dataset. To compare our proposed systems with the state-of-the-art, we follow the Second DiCOVA Challenge.

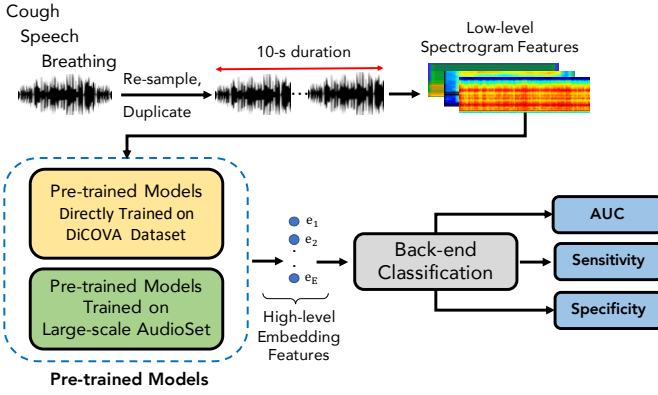


Fig. 1. The high-level architecture of the proposed frameworks

Particularly, we aim to detect COVID-19 positive subjects by exploring only breathing sound, only cough sound, only speech sound, and using all sound categories, which match the Track-1, Track-2, Track-3, and Track-4, respectively in the Second DiCOVA Challenge. Regarding the evaluation metrics, we also obey the Second DiCOVA Challenge which uses An Area Under the Curve (AUC) and the specificity (SPEC.) & sensitivity (SEN.). As SPEC. has to be equal or larger than 95.13% which is an essential requirement of the challenge, a decision threshold from 0 to 1 with a step size of 0.0001 is evaluated to obtain SEN. when SPEC. value is satisfied. As a result, we report AUC and SEN. in this paper (i.e. The SPEC. is always equal or greater than 95.13% to meet the challenge’s requirement as presented in the leaderboard of the challenge [6]).

### III. THE PROPOSED FRAMEWORK

To explore DiCOVA dataset, we firstly present a high-level architecture of proposed framework as shown in Figure 1. Generally, an entire framework proposed is separated into three main steps: Low-level spectrogram feature extraction, high-level embedding feature extraction, and back-end classification.

#### A. The low-level spectrogram feature extraction

At the first step shown in the upper stream in Figure 1, the raw audio recordings are re-sampled to 16.000 Hz, duplicated to make sure that all audio recordings have an equal duration of 10 seconds. These audio recordings are then transformed into spectrograms where both temporal and spectral features are presented.

#### B. The high-level embedding feature extraction

The low-level spectrogram features are then fed into pre-trained deep learning models to extract embeddings (e.g. vectors), referred to as high-level features. The pre-trained deep learning models used for extracting high-level embedding features come from two different approaches: **(I)** pre-trained models directly trained on DiCOVA dataset and **(II)** pre-trained models trained with the large-scale AudioSet dataset [7].

TABLE I  
CONFIGURATION OF LENET BASED ARCHITECTURE

| Lenet architecture layers                                     | Output                    |
|---|---------------------------|
| Input layer (spectrogram patch of $128 \times 154 \times 1$ ) |                           |
| BN - Conv [3×3] @ 32 - ReLU - BN - AP [2×2] - Dr (20%)        | $64 \times 77 \times 32$  |
| BN - Conv [3×3] @ 64 - ReLU - BN - AP [2×2] - Dr (25%)        | $32 \times 38 \times 64$  |
| BN - Conv [3×3] @ 128 - ReLU - BN - AP [2×2] - Dr (30%)       | $16 \times 19 \times 128$ |
| BN - Conv [3×3] @ 256 - ReLU - BN - GAP - Dr (35%)            | 256                       |
| FC - ReLU - Dr (40%)  | 1024                      |
| FC - ReLU - Dr (40%)  | 1024                      |
| FC - Softmax  | C=2                       |

**(I) In the first approach**, we construct deep learning models and directly train these models on DiCOVA dataset mentioned in Section II with Cross-entropy loss and Adam method [8] for optimization. After the training process, we reuse these models to extract the feature map at the global pooling layer, which is considered as the high-level embedding features mentioned. As the performance of high-level embedding features depends on the low-level spectrograms and the model architectures in this approach, we therefore base on our previous work [9]–[13] to evaluate: three types of spectrograms such as log-Mel [14], Gammatonegram [15], and Scalogram [16] that are proven effective in representing respiratory-related sounds; and a wide range of benchmark deep learning models from low footprint models such as Lenet based network [17] to high-complex architectures such as Xception [18], InceptionV3 [19], etc.

To identify which spectrogram features are well-represented for breathing, cough, or speech audio inputs, we evaluate three proposed spectrograms using a Lenet based network architecture as shown in Table I. To make sure that the spectrograms that are fed into the Lenet based model have the same size, we use the same setting parameters with window size = 2048, hop size = 1024, and filter number = 128; to generate the same spectrogram of  $128 \times 154$ . Regarding the Lenet based architecture as shown in Table I, it includes the convolutional layer (Conv [kernel size]), batch normalization (BN) [20], rectified linear units (ReLU) [21], average pooling (AP), global average pooling (GAP), dropout [22] (Dr(percentage)), fully connected (FC) and Softmax layer. Since certain low-level spectrograms are suitable for various types of audio inputs (i.e. breathing, cough, or speech), we therefore evaluate these spectrograms with different benchmark neural network architectures of VGG16 [23], VGG19 [23], MobileNetv1 [24], ResNet50 [25], Xception [18], InceptionV3 [19], and DenseNet121 [26] to achieve the best framework configuration (i.e. which low-level spectrogram and training network architecture). As evaluated network architectures are reused from the Keras library [27], the final fully connected layer of these networks is modified from 1000 (i.e. the number of classes of image object in ImageNet dataset) to 2 which matches the number of classes in DiCOVA dataset. It is worth noting that we only reuse the network architectures from the Keras library instead of using available weights trained with ImageNet dataset, all trainable parameters of these networks are initialized with mean and variance set to 0 and 0.1, respectively.

**(II) In the second approach**, we leverage three available pre-trained models trained with the large-scale AudioSet

TABLE II  
COMPARISON OF LOW-LEVEL SPECTROGRAMS (AUC/SEN.)

| Spec           | Cough              | Speech             | Breathing          | All                |
|----------------|--------------------|--------------------|--------------------|--------------------|
| GAM            | 70.11/10.00        | 73.41/21.67        | 73.29/15.00        | 77.38/25.00        |
| <b>log-MEL</b> | <b>71.39/16.66</b> | <b>73.15/18.33</b> | <b>74.31/18.33</b> | <b>80.11/33.33</b> |
| Scalogram      | 63.71/8.33         | 68.68/10.00        | 78.20/20.00        | 69.83/18.33        |

dataset in advance: PANN [28], OpenL3 [29], TRILL [30]. We evaluate whether these up-stream pre-trained models are beneficial for the down-stream task of detecting COVID-19 positive in DiCOVA dataset. While PANN [28] and OpenL3 [29] leverage VGGish architecture and cross-entropy loss to train the large-scale Audioset dataset, TRILL [30] is based on Resnet and tripless loss. Additionally, while PANN [28] was trained using input signal with 10-second duration, both TRILL [30] and OpenL3 [29] analyse short time duration of 1 second. As using different network architectures, loss functions, as well as analysing different audio durations, these proposed pre-trained models may differently perform on three types of audio inputs (i.e. breathing, cough, and speech) from DiCOVA dataset.

Regarding the high-level embedding features extracted from these pre-trained models, while OpenL3 [29] and PANN [28] extract the feature map at the global pooling layer, TRILL [30] extracts the feature map at the final layer which proves effective for different down-stream tasks mentioned in [30]. Notably, since PANN pre-trained model works on 10-second input duration from AudioSet dataset that matches the input duration of our proposed system in Figure 1, only one embedding feature (e.g. one vector) is extracted from each 10-second audio recordings of the input. Meanwhile, as OpenL3 and TRILL pre-trained models were trained with 1-second audio segment of AudioSet dataset, therefore, multiple embeddings are obtained when we feed one 10-second DiCOVA audio sample into these two pre-trained models. Therefore, an average of these multiple embeddings across the time dimension is computed to obtain one embedding feature which represents for each 10-second duration audio input. Additionally, only low-level log-Mel spectrogram is used in this approach as these three pre-trained models explore this type of spectrogram for training on the large-scale AudioSet dataset.

As we present two different approaches of using pre-trained models for extracting high-level embedding features, we now refer to two main frameworks as: (I) three low-level spectrograms (log-Mel, Gammatonegram, Scalogram), pre-trained models directly trained on DiCOVA dataset, LightGBM back-end classification; and (II) low-level log-Mel spectrogram, pre-trained models trained on the large-scale AudioSet, and LightGBM back-end classification.

### C. The back-end classification

In this paper, we use Light Gradient Boosting Machine (LightGBM) [31] as the final back-end classification model to fine-tune high-level embedding features. The LightGBM is implemented by using an available toolkit [31] and the parameters are set as: `learning_rate = 0.02`, `objective = 'binary'`, `metric = 'auc'`, `subsample = 0.68`, `colsample_bytree = 0.28`, `early_stopping_rounds = 1000`, `num_iterations = 10000`,

TABLE III  
COMPARISON OF PRE-TRAINED DEEP LEARNING NETWORK ARCHITECTURES DIRECTLY TRAINED ON DiCOVA DATASET (AUC/SEN.)

| Input           | Cough              | Speech             | Breathing          | All                |
|-----------------|--------------------|--------------------|--------------------|--------------------|
| VGG16           | 70.05/10.00        | 75.64/10.00        | 71.75/16.66        | 78.17/30.00        |
| VGG19           | 63.07/5.00         | 67.81/8.33         | 70.28/15.00        | 68.24/6.67         |
| MobileNetv1     | 58.92/5.00         | 66.86/16.67        | 67.76/13.33        | 70.73/25.00        |
| ResNet50        | 66.67/10.00        | 70.07/21.67        | 69.39/18.33        | 77.35/23.33        |
| <b>Xception</b> | <b>71.21/31.66</b> | <b>70.91/28.33</b> | <b>72.75/28.33</b> | <b>80.81/43.33</b> |
| InceptionV3     | 61.97/8.33         | 74.39/30.00        | 67.56/18.33        | 78.83/30.00        |
| DenseNet121     | 68.12/23.33        | 75.55/20.00        | 72.63/18.33        | 80.10/31.66        |

`subsample_freq = 1`. As the Track-4 in the Second DiCOVA Challenge suggests to use all audio input data (i.e. breathing, cough, speech), high-level embedding features extracted from different types of audio inputs are concatenated before feeding into the back-end LightGBM for classification.

## IV. EXPERIMENTS AND RESULTS

### A. Performance comparison of frameworks (I): three low-level spectrograms, pre-trained models directly trained on DiCOVA dataset, and LightGBM back-end classification

We firstly evaluate how low-level spectrograms affect the performance in frameworks (I). As the results are shown in Table II, log-Mel performs better than Gammatonegram and Scalogram with the highest scores of 71.39/16.66 in cough while still witness the second highest scores of 73.15/18.33 and 74.31/18.33 in speech and breathing, respectively. As a result, we only use log-Mel spectrogram for evaluating network architectures.

As Table III shows, it can be seen that Xception outperforms the other network architectures. However, the high-complex Xception model only slightly improve, compared to the low footprint Lenet based architecture. Notably, when high-level embedding features extracted from breathing, cough, and speech inputs by Xception model are concatenated to detect COVID-19 (i.e. The Track-4 in the Second DiCOVA Challenge), it helps to improve significantly the performance, recording a score of 80.81/43.33 compared with 71.21/31.66, 70.91/28.33, 72.75/28.33 for only cough, speech, and breathing, respectively.

### B. Performance comparison of frameworks (II): low-level log-Mel spectrogram, pre-trained models trained on AudioSet dataset, and LightGBM back-end classification

As Table IV shows, it can be seen that different pre-trained models works well on different types of sound input. Particularly, the best scores of 79.82/48.33 are associated with TRILL for cough sound input. PANN achieves the best scores of 84.05/55.00 for breathing. Meanwhile, speech sound input presents the best scores of 86.65/61.67 with OpenL3. Notably, using pre-trained models with AudioSet dataset outperforms the frameworks (I) analysed in Section IV-A in terms of detecting COVID-19 with only breathing, cough, or speech.

Comparing the performance of various models among audio inputs, breathing and speech show potential to detect COVID-19 rather than cough. However, when we concatenate high-level embedding features extracted from breathing,

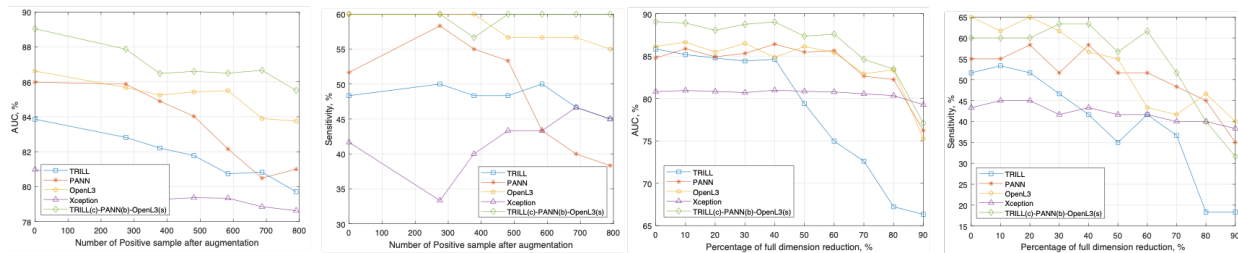


Fig. 2. Influence of over-sampling on positive samples and dimensionality reduction using high-level features within Track-4 DiCOVA Challenge

TABLE IV  
COMPARISON OF PRE-TRAINED MODELS TRAINED ON AUDIOSET DATASET (AUC/SEN.)

| Pre-trained model | Cough              | Speech             | Breathing          | All                |
|-------------------|--------------------|--------------------|--------------------|--------------------|
| PANN              | 70.49/26.66        | 79.33/33.33        | <b>84.05/55.00</b> | 84.77/55.00        |
| TRILL             | <b>79.82/48.33</b> | 80.50/46.67        | 76.97/33.33        | 85.83/51.67        |
| OpenL3            | 78.42/38.33        | <b>86.65/61.67</b> | 82.40/50.00        | <b>86.14/65.00</b> |

cough, and speech for Track-4 of the challenge, the performance from TRILL shows a significant improvement to 85.83/51.67. Meanwhile, OpenL3 demonstrates the highest scores of 86.14/65.00 compared to the others in this track.

### C. Influence of over-sampling on positive samples and dimensionality reduction of high-level features

As the proposed frameworks with a concatenation of high-level features prove effective to detect COVID-19 in Track-4 of the Second DiCOVA Challenge, we further conduct experiments to evaluate how over-sampling on positive samples and reducing dimension over high-level features affect the performance. In these experiments, we concatenate (1) three TRILL based high-level features, (2) three PANN based high-level features, (3) three OpenL3 based high-level feature, (4) three Xception based high-level features, and (5) (TRILL(c)-PANN(b)-OpenL3(s)) high-level features for cough, breathing, and speech, respectively. We conduct the combination among high-level features (TRILL(c)-PANN(b)-OpenL3(s)) as each feature shows effective for each audio data input as shown in Table IV.

To identify less significant dimensions of high-level features, we firstly compute an average of each dimension from all features associated with positive and negative COVID-19 independently. Then, two vectors representing for two groups of COVID-19 positive and negative features are obtained. Based on the absolute difference of these two vectors from each dimension (i.e. the lower absolute difference is related to the less significant dimension), the dimension of the feature set has been reduced from 10% to 90%. Note that the dimension of high-level features is reduced before concatenating. To over-sample the positive cases, we apply SVM-SMOTE [32] on the high-level embedding features, then increase the the positive samples from double to five times.

As Figure 2 shows, oversampling on positive samples reduces the AUC scores. This method only helps to improve SEN. scores with PANN (double of positive samples) and Xception (five times of positive samples). Meanwhile, when

we reduce the dimension of high-level embedding features from 20 to 40%, it almost helps to improve both AUC and SEN. scores. As a result, we can achieve the best scores: 80.32/48.33 with TRILL (at the reducing dimension of 20%) for cough input only, 84.05/55.00 with PANN for breathing input only (no reduction), 86.65/61.67 with OpenL3 for speech only (no reduction), and 89.03/63.33 with (TRILL(c)-PANN(b)-OpenL3(s)) (at the dimensionality reduction of 40%) for all audio inputs.

### D. Performance comparison among different back-end classifiers

We further compare the LightGBM with other machine learning models implemented by using Scikit-Learn toolkit [33] on of Support Vector Machine (SVM with C=1.0, Kernel='RBF'), Random Forest (RF with Max Depth of Tree = 20, Number of Trees = 100), Multi-layer perceptron (MLP with 4096 nodes, Adam optimization, Max iter = 200, Learning rate = 0.001, Entropy Loss), and Linear Regression (LR with C=1.0). Noting that we keep TRILL based embedding (at the dimensionality reduction of 20%) for cough, PANN and OpenL3 based embeddings (with no reduction) for breathing and speech, respectively. Similarly, we continue to apply the best concatenation of (TRILL(c)-PANN(b)-OpenL3(s)) (at the dimensionality reduction of 40%) for all audio inputs. As a result, the comparison in Table VI again shows that LightGBM model still achieves the best scores and outperforms other models.

### E. Performance comparison across the top-10 systems submitted for the Second DiCOVA Challenge

As performance comparison to the state-of-the-art systems [6] is shown in Table V, we can achieve the top-6 in Track-1 with breathing input only and the top-3 in Track-2 with cough input only. Notably, our proposed systems outperform the state-of-the-art, achieve the top-1 in both Track-3 and Track-4 with speech input only and all audio inputs, respectively.

To evaluate the best score of the LightGBM with the best-selected embedding features for cough, speech, breathing and all types of audio inputs, we conducted ten times of running the experiments on a different randomly chosen Test set. Next, we calculate and achieve an average confidence interval (CI) of [0.7628, 0.8352] in cough, [0.8369, 0.8981] in speech, [0.8133, 0.8785] in breathing, and [0.8580, 0.9153] in all audio inputs. All of them matches with the mentioned AUC scores

TABLE V  
PERFORMANCE COMPARISON WITH THE STATE-OF-THE-ART SYSTEMS ON TEST SET

| System            | Cough        |              |              |              | Speech       |              |              |              | Breathing    |              |              |              | All          |              |              |              |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                   | AUC          | SEN.         | Precision    | F1 score     | AUC          | SEN.         | Precision    | F1 score     | AUC          | SEN.         | Precision    | F1 score     | AUC          | SEN.         | Precision    | F1 score     |
| 1st system        | 81.97        | 36.67        | 52.38        | 43.14        | 85.21        | 45.00        | 57.45        | 50.47        | 87.18        | 48.33        | 59.18        | 53.21        | 88.44        | 58.33        | 63.64        | 60.87        |
| 2nd system        | 81.21        | 48.33        | 59.18        | 53.21        | 84.73        | 38.33        | 53.49        | 44.66        | 86.72        | 40.00        | 54.55        | 46.16        | 87.26        | 63.33        | 65.52        | 64.41        |
| 3rd system        | 80.12        | 35.00        | 51.22        | 41.58        | 84.55        | 51.67        | 60.78        | 55.86        | 86.41        | 45.00        | 57.45        | 50.47        | 86.87        | 58.33        | 63.64        | 60.87        |
| 4th system        | 79.06        | 35.00        | 51.22        | 41.58        | 84.26        | 43.33        | 56.52        | 49.05        | 85.77        | 41.67        | 55.56        | 47.62        | 85.79        | 33.33        | 50.00        | 40.00        |
| 5th system        | 77.85        | 46.67        | 58.33        | 51.85        | 84.04        | 48.33        | 59.18        | 53.21        | 84.50        | 31.67        | 48.72        | 38.39        | 85.37        | 60.00        | 64.29        | 62.07        |
| 6th system        | 77.60        | 33.33        | 50.00        | 40.00        | 82.98        | 43.33        | 56.52        | 49.05        | 82.16        | 41.67        | 55.56        | 47.62        | 84.70        | 55.00        | 62.26        | 58.41        |
| 7th system        | 76.98        | 25.00        | 42.86        | 31.58        | 82.83        | 38.33        | 53.49        | 44.66        | 82.05        | 41.67        | 55.56        | 47.62        | 84.26        | 38.33        | 53.49        | 44.66        |
| 8th system        | 76.36        | 30.00        | 47.37        | 36.74        | 81.95        | 48.33        | 59.18        | 53.21        | 80.84        | 40.00        | 54.55        | 46.16        | 83.78        | 46.67        | 58.33        | 51.85        |
| 9th system        | 75.95        | 28.33        | 45.95        | 35.05        | 81.86        | 31.67        | 48.72        | 38.39        | 80.55        | 30.00        | 47.37        | 36.74        | 80.51        | 40.00        | 54.55        | 46.16        |
| 10th system       | 75.71        | 35.00        | 51.22        | 41.58        | 80.92        | 45.00        | 57.45        | 50.47        | 80.35        | 36.67        | 52.38        | 43.14        | 74.15        | 35.00        | 51.22        | 41.58        |
| Baseline system   | 74.89        | 36.67        | 52.38        | 43.14        | 84.26        | 43.33        | 56.52        | 49.05        | 84.50        | 31.67        | 48.72        | 38.39        | 84.70        | 55.00        | 62.26        | 58.41        |
| <b>Our system</b> | <b>80.32</b> | <b>48.33</b> | <b>59.18</b> | <b>53.21</b> | <b>86.65</b> | <b>61.67</b> | <b>64.91</b> | <b>63.25</b> | <b>84.05</b> | <b>55.00</b> | <b>62.26</b> | <b>58.41</b> | <b>89.03</b> | <b>63.33</b> | <b>65.52</b> | <b>64.41</b> |

TABLE VI  
PERFORMANCE COMPARISON ACROSS DIFFERENT BACK-END CLASSIFICATION MODELS ON TEST SET (AUC/SEN.)

| Model    | Cough              | Speech             | Breathing          | All                |
|----------|--------------------|--------------------|--------------------|--------------------|
| LR       | 76.22/35.00        | 80.90/40.00        | 78.61/43.33        | 77.29/43.33        |
| SVM      | 78.40/41.66        | 85.93/53.33        | 82.25/51.66        | 85.73/55.00        |
| RF       | 78.78/43.33        | 81.06/48.33        | 79.06/38.33        | 85.34/51.66        |
| MLP      | 74.80/41.66        | 81.53/56.66        | 80.89/41.66        | 83.75/46.67        |
| LightGBM | <b>80.32/48.33</b> | <b>86.65/61.67</b> | <b>84.05/55.00</b> | <b>89.03/63.33</b> |

of 80.32, 86.65, 84.05, and 89.03 for cough, speech, breathing, and all, respectively.

## V. CONCLUSION

This paper has presented an exploration on how to extract effectively well-represented features for breathing, cough, speech sound input via pre-trained models. By conducting extensive experiments, we achieve a robust framework for detecting COVID-19 compared with the state-of-the-art systems. This is demonstrated by the rank of our performance with top-1 in both Track-3 and Track-4, top-3 in Track-2 and finally top-6 in Track-1 in the Second DiCOVA Challenge. Our best AUC score of 89.03%, F1 score of 64.41%, sensitivity score of 66.33% from the Track-4 demonstrate the potential of detecting COVID-19 through the respiratory-related sounds.

## REFERENCES

- [1] B. News. (2022-01-18) Covid map: Coronavirus cases, deaths, vaccinations by country. <https://www.bbc.com/news/world-51235105>.
- [2] V. M. Corman and et al., "Detection of 2019 novel coronavirus (2019-ncov) by real-time rt-pcr," *Eurosurveillance*, vol. 25, no. 3, pp. 23–30, 2020.
- [3] R. W. Peeling, P. L. Olliaro, D. I. Boeras, and N. Fongwen, "Scaling up covid-19 rapid antigen tests: promises and challenges," *The Lancet infectious diseases*, vol. 21, no. 9, pp. e290–e295, 2021.
- [4] J. Andreu-Perez and et al., "A generic deep learning based cough analysis system from clinically validated samples for point-of-need covid-19 test and severity levels," *IEEE Transactions on Services Computing*, 2021.
- [5] P. Muguli and et al., "Dicova challenge: Dataset, task, and baseline system for covid-19 diagnosis using acoustics," *preprint arXiv:2103.09148*, 2021.
- [6] T. S. D. Challenge. (2021-10-02) Leaderboard results. <https://competitions.codalab.org/competitions/34801results>.
- [7] J. F. Gemmeke and et al., "Audio set: An ontology and human-labeled dataset for audio events," in *Proc. ICASSP*, 2017, pp. 776–780.
- [8] P. K. Diederik and B. Jimmy, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2015.
- [9] L. Pham, D. Ngo, T. Hoang, A. Schindler, and I. McLoughlin, "An ensemble of deep learning frameworks applied for predicting respiratory anomalies," *preprint arXiv:2201.03054*, 2022.
- [10] D. Ngo and et al., "Deep learning framework applied for predicting anomaly of respiratory sounds," in *Proc. ISEE*, 2021, pp. 42–47.

- [11] L. Pham and et al., "Cnn-moe based framework for classification of respiratory anomalies and lung disease detection," *IEEE journal of biomedical and health informatics*, vol. 25, no. 8, pp. 2938–2947, 2021.
- [12] L. Pham, H. Phan, A. Schindler, R. King, A. Mertins, and I. McLoughlin, "Inception-based network and multi-spectrogram ensemble applied to predict respiratory anomalies and lung diseases," in *Proc. EMBC*, 2021, pp. 253–256.
- [13] L. Pham and et al., "Robust deep learning framework for predicting respiratory anomalies and diseases," in *Proc. EMBC*, 2020, pp. 164–167.
- [14] B. McFee and et al., "librosa: Audio and music signal analysis in python," in *The 14th Python in Science Conference*, 2015, pp. 18–25.
- [15] W. Ellis, "Gammatone-like spectrogram," <http://www.ee.columbia.edu/dpwe/resources/matlab/gammatonegram>, 2009.
- [16] L. Jonathan and O. Sofia, "Higher-order properties of analytic wavelets," *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 146–160, 2008.
- [17] Y. LeCun, P. Haffner, L. Bottou, and Y. Bengio, "Object recognition with gradient-based learning," in *Shape, contour and grouping in computer vision*. Springer, 1999, pp. 319–345.
- [18] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. CVPR*, 2017, pp. 1251–1258.
- [19] C. Szegedy and et al., "Rethinking the inception architecture for computer vision," in *Proc. CVPR*, 2016, pp. 2818–2826.
- [20] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. PMLR*, 2015, pp. 448–456.
- [21] V. Nair and E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proc. ICML*, 2010.
- [22] N. Srivastava and et al., "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [23] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *preprint arXiv:1409.1556*, 2014.
- [24] A. G. Howard and et al., "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *preprint arXiv:1704.04861*, 2017.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.
- [26] G. Huang and et al., "Densely connected convolutional networks," in *Proc. CVPR*, 2017, pp. 4700–4708.
- [27] F. Chollet, "Keras library," <https://keras.io>, 2015.
- [28] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "Panns: Large-scale pretrained audio neural networks for audio pattern recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [29] J. Cramer and et al., "Look, listen, and learn more: Design choices for deep audio embeddings," in *Proc. ICASSP*, 2019, pp. 3852–3856.
- [30] J. Shor and et al., "Towards learning a universal non-semantic representation of speech," *preprint arXiv:2002.12764*, 2020.
- [31] K. Guolin and et al., "Lightgbm: A highly efficient gradient boosting decision tree," *Advances in neural information processing systems*, vol. 30, pp. 3146–3154, 2017.
- [32] N. Hien, E. Cooper, and K. Kamei, "Borderline over-sampling for imbalanced data classification," *International Journal of Knowledge Engineering and Soft Data Paradigms*, vol. 3, no. 1, pp. 4–21, 2011.
- [33] F. Pedregosa and et al., "Scikit-learn: Machine learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.