

## Detecting COVID-19 from breathing and coughing sounds using deep neural networks

Mina A. Nessiem, Mostafa M. Mohamed, Harry Coppock, Alexander Gaskell, Björn W. Schuller

### Angaben zur Veröffentlichung / Publication details:

Nessiem, Mina A., Mostafa M. Mohamed, Harry Coppock, Alexander Gaskell, and Björn W. Schuller. 2021. "Detecting COVID-19 from breathing and coughing sounds using deep neural networks." In *34th International Symposium on Computer-Based Medical Systems (CBMS 2021), Aveiro, Portugal, 7-9 June 2021*, edited by José Luís Oliveira, Agma Traina, Paolo Soda, Joel Arrais, and Bridget Kane, 183–88. New York, NY: IEEE.  
<https://doi.org/10.1109/cbms52027.2021.00069>.



# Detecting COVID-19 from Breathing and Coughing Sounds using Deep Neural Networks

Mina A. Nessiem<sup>\*§</sup>, Mostafa M. Mohamed<sup>\*§</sup>, Harry Coppock<sup>†</sup>, Alexander Gaskell<sup>†</sup>, and Björn W. Schuller<sup>\*†</sup>

<sup>\*</sup>Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Germany

<sup>†</sup>GLAM – Group on Language, Audio, & Music, Imperial College London, UK

mina.nessiem@informatik.uni-a.de, mostafa.mohamed@student.uni-a.de

**Abstract**—The COVID-19 pandemic has affected the world unevenly; while industrial economies have been able to produce the tests necessary to track the spread of the virus and mostly avoided complete lockdowns, developing countries have faced issues with testing capacity. In this paper, we explore the usage of deep learning models as a ubiquitous, low-cost, pre-testing method for detecting COVID-19 from audio recordings of breathing or coughing taken with mobile devices or via the web. We adapt an ensemble of Convolutional Neural Networks that utilise raw breathing and coughing audio and spectrograms to classify if a speaker is infected with COVID-19 or not. The different models are obtained via automatic hyperparameter tuning using Bayesian Optimisation combined with HyperBand. The proposed method outperforms a traditional baseline approach by a large margin. Ultimately, it achieves an Unweighted Average Recall (UAR) of 74.9 %, or an Area Under ROC Curve (AUC) of 80.7 % by ensembling neural networks, considering the best test set result across breathing and coughing in a strictly subject independent manner. In isolation, breathing sounds thereby appear slightly better suited than coughing ones (76.1 % vs 73.7 % UAR).

**Index Terms**—COVID-19, Speech Analysis, Deep Learning, Ensemble Models, Convolutional Neural Networks

## I. INTRODUCTION

The COVID-19 pandemic has forced the global community to reckon with a shortage of adequate testing capacity all over the globe [1]. The Foundation for Innovative New Diagnostics (FIND) tracker [2] shows that the pandemic has exacerbated global economic inequalities. Developed countries have mature industries available with which testing equipment and materials can be produced locally or the funds to procure the necessary materials from abroad. Meanwhile, developing countries face a distinct lack of testing equipment and materials, having even during the second wave of the pandemic the capacity to only conduct less than 5 000 daily tests per capita, or even 1 000 tests per capita for some developing countries. This lack of access to testing capacity forces countries to grapple with two options: either push through hard or complete lockdown measures in bids to slow down or break the spread of the virus, placing excessive strain on local economies for those workers whose jobs can not be done in a Work from Home (WfH) fashion [3], or allow the virus to pass undetected through their populations in trying to shore up the economies, but while placing further pressure on already under-equipped health systems.

<sup>§</sup>Equal Contribution

Artificial intelligence (AI) and in particular, Deep Neural Networks (DNNs), started to grow in popularity ever since amongst others [4] used them to surpass – by a large margin – previous classical machine learning approaches on the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [5]. Since then, they have come to set the state of the art in a number of different fields and challenges. Among these fields where AI has been finding usage is that of medicine in general [6]–[8], where it has been shown to be effective for cancer screening [9]–[11] and for COVID-19 screening as well [12]–[14]. Of particular interest here, is the usage of AI to help in combating the COVID-19 pandemic, as researchers have already made use of AI to analyse different signals for signs of COVID-19 [15, 16].

It is in this context that it becomes important to consider the usage of everyday tools such as internet-connected mobile phones [17], together with artificial intelligence as tools to detect the infection with COVID-19 as a way of mitigating the necessity for everybody to take tests, or to quarantine at home. In this paper, we introduce a machine learning model that detects, by way of coughing or breathing samples, whether the audio contains traces of COVID-19 or not, which could be used, e. g., as pre-selection filter for more reliable, but more expensive testing methods. The audio stems from a public database of coughing and breathing sounds collected from mobile phones and over the Internet crowd-sourced by the University of Cambridge.

For an overview on the current state-of-play in COVID-19 detection from audio, a short overview is given in [18].

In Section III we introduce our approach, including how we represent the audio as input to the models, the model architecture and construction, the training technique we used, as well as the hyperparameter optimisation process undertaken; we finish the section by introducing the baseline with which we compare our results. In Section IV, we showcase our experiments, beginning by introducing the dataset that serves as the basis for our experiments, the evaluation metrics that we consider, and the results of our approach, as well as finish the section with a discussion of the results. We conclude the paper in Section V.

## II. RELATED WORK

[19] developed models that can classify COVID-19 according to cough sounds, however, they report limited results of Area Under ROC Curve (AUC) 72 %. [20] developed a similar approach with much better results, however, their data size is quite limited (80 different speakers); they collected three different samples from each person, without mentioning that they split the data speaker independently, which might put their results in jeopardy if there is a speaker in both training and test data. [21] used audio data of relatively low quality (collected from YouTube), with phoneme features and SVM models to classify if a person has COVID or not from speech, instead of coughing or breathing sounds.

## III. APPROACH

Our approach utilises raw audio in combination with different spectrogram variations. We implement an independent branch in our model for each of the input formats, where each branch consists of a Convolutional Neural Network (CNN). In the end, we combine the learnt representations from each branch by concatenation and use fully-connected layers to output a final classification label.

### A. Audio representations

Spectrograms are visual representations of audio [22], which are computed by calculating the Short-Time Fourier Transform (STFT). This is done by cutting the audio into (not necessarily overlapping-) chunks and then computing the Discrete Fourier Transform (DFT) for each chunk. The difference between each chunk and the next is defined as the *hop length*  $S$ . The output of the DFT is a decomposition of all the frequency bands contained within an audio chunk, this is either shown on a linear-scale or gets transformed to a log-mel scale. One of the various aspects that are considered when constructing spectrograms is whether they are *narrowband* or *wideband*. Spectrograms can display a high resolution in either the frequency dimension or the time dimension, but not both. Given these spectrogram variations, we make use of a combination of eight different spectrograms that differ in terms of three main parameters:

- 1) The hop length  $S$  of the spectrogram chunks: 64 or 128.
- 2) The chunk length of either  $2S$ , or  $4S$ .
- 3) Mel-scale (128 bins, NFFT= 1024) or Linear-scale.

### B. Models

Our models make use of two building blocks. We will introduce each of the building blocks first, followed by an explanation of the general architecture of the models. We refer to a *temporal feature map* as a real-valued 2-dimensional tensor of dimensions  $T \times F$ , where  $T$  is the length of time dimension and  $F$  is the number of features. According to this definition, the spectrograms and raw audio are both temporal feature maps. The first building block is a *convolution block*, which consists of a 1D convolution layer with  $C$  output channels and a Rectified Linear Unit (ReLU) activation function, followed by a max

pooling layer, then a dropout layer [23]. The dropout layer has the aim of reducing overfitting; we adopt a dropout rate of 0.2. Additionally, a *convolution branch* is a stack of  $N$  convolution blocks. A convolution branch will map a temporal feature map into another temporal feature map, which is more condensed in the time dimension. In other words, the features at the end have high representations, while the time dimension is reduced several times via the pooling layers in the convolution blocks.

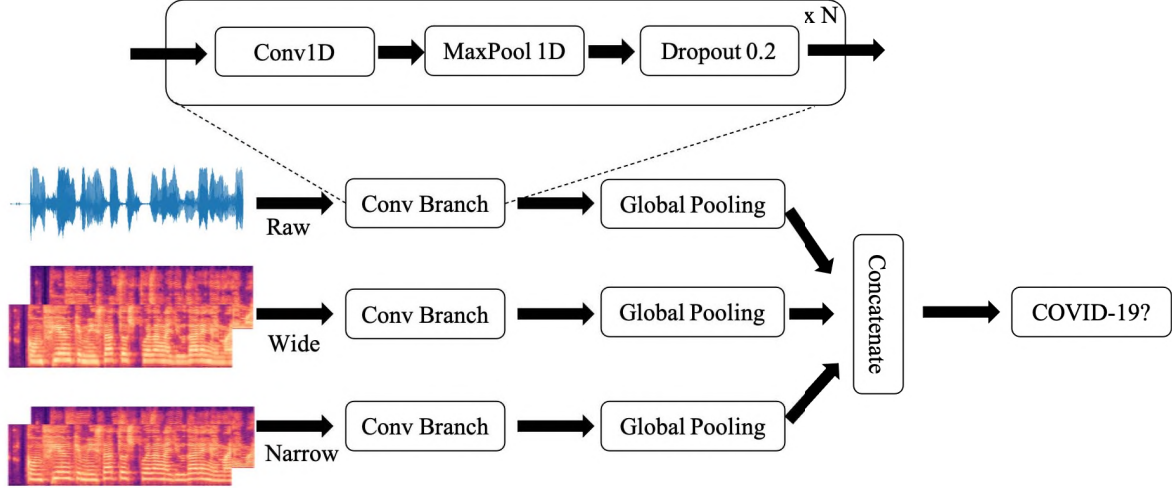
The second building block is a reduction of the time dimension of a temporal feature map, using global pooling layers. We make use of global average pooling and global maximum pooling; then, we concatenate the resulting features from both. For the model, we make use of several representations of the input, inspired by [24] which utilised multi-channel spectrograms. The spectrograms are grouped by their hop-length, because different hop-lengths result in different lengths of the time dimension of the spectrograms. Each group of spectrograms is concatenated along the channels (frequencies) axis, which results in a total of two groups of spectrograms in addition to the raw audio. For each of those three, we construct separate but identical convolution branches, followed by global pooling. After that, the outputs from the different branches are concatenated, and followed by  $F$  fully-connected layers, which include the last layer the outputs the classification labels. The resulting model architectures is shown in Figure 1.

### C. Training and hyperparameter optimisation

We train several variants of the model architecture described in Subsection III-B while splitting the data using nested  $k$ -fold cross validation [25] on the mixture of train and development parts of the data. The variants differ in at least one of two aspects; they are either trained using a different fold or constructed having different hyperparameters. We utilise an ensemble of the variants to reach best results, as ensembling tends to reduce overfitting and hence reaches the best results [26]. We use averaging of the prediction probabilities to ensemble the different models. For all the variants, we employ the Adam optimiser [27] using a mini-batch size of 16 for 100 epochs, then the model with the best validation score on the corresponding hold-out set is saved. The learning rate of Adam  $\alpha$  is considered as a hyperparameter to be tuned. Since the database consists of audio tracks of varying lengths, we pad the examples with 0 values to match the longest example in a given batch. The network parameter optimisation is done by minimising the binary crossentropy loss function, while giving a different weight  $\lambda$  to the positive class. This will direct the model to focus more on getting the positive cases correct, which is helpful in case the positive cases are underrepresented in the data. The final equation of the loss function is given by:

$$\mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) = -\frac{1}{N} \sum_{i=1}^N (\lambda y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i))$$

For hyperparameter tuning, we use Bayesian Optimisation HyperBand (BOHB) [28] by instructing it to reach the best



**Fig. 1:** An illustration of the model’s architecture. Three input formats are used: raw audio, a mixture of spectrograms of hop length  $S = 8$  ms, and  $S = 16$  ms respectively. For each of those,  $N$  convolution blocks are used, followed by a time reduction block, concatenation of features then fully-connected layers for classification.

Condition	Platform	Symptom	# files
COVID	Android	no cough	128
COVID	Android	cough	92
COVID	Web	no cough	46
COVID	Web	cough	16
Total COVID-19 positive			283
Asthma	Android	cough	26
Asthma	Web	cough	16
Healthy	Android	no symptom	282
Healthy	Android	cough	16
Healthy	Web	no symptom	362
Healthy	Web	cough	50
Total COVID-19 negative			752

**TABLE I:** Unaugmented database file statistics.

Unweighted Average Recall (UAR) (explained further in Subsection IV-B) on the validation data, namely by maximising the average of the UAR across the hold out folds in the cross-validation.

#### D. Baseline

We compare our approach against a baseline linear Support Vector Machine (SVM) classifier [29], which is trained on audio functionals extracted with openSMILE [30]. We choose the large feature configuration (6 373 functionals) as introduced in the Interspeech 2016 ComParE challenge [31], which has been the configuration of choice for the Interspeech ComParE challenges ever since. On that, we perform Principal Component Analysis (PCA) [29] to reduce the number of features to 100 in a standard manner. The SVM classifier’s complexity parameter is optimised on a logarithmic scale between  $[10^{-5}, 1]$ , based on the achieved UAR on a development partition. The optimisation yielded a value  $10^{-5}$ . Afterwards, the SVM is fit on the combined training and development data with the optimised complexity value and finally used for inference on the test partition of the data.

## IV. EXPERIMENTS

### A. Database

The experiments were performed on a crowdsourced database [32] that was collected via the “COVID-19 sounds” Android app, as well as through a web form (an iOS app also exists, though it did not contribute to the database used here) by the University of Cambridge [32]. The participants are asked to fill a survey about their demographic information (such as age and location), medical history, as well as symptoms (if any). The app instructed the participants to “breathe deeply five times, cough three times, and read three times a short sentence appearing on screen” [32].<sup>1</sup>

The present database is an excerpt of the total data collected via the app. The database includes the coughing and breathing audio samples, their associated medical condition (COVID, Asthma, or Healthy), whether the submitter suffered from coughing as a symptom or not, and finally which platform the samples were collected through (Android or Web).

The database consists of a total of 1 427 audio files that total 3.93 hours. Of these 1 427 files, 1 034 are original audio samples that total two hours, while the rest are the result of data augmentation [33]. The mean file length for the dataset is 9.96 seconds with a standard deviation of 6.02 seconds. The audio files have a sample rate 16 kHz and 16 bit quantisation. In total, 174 unique participants are included in the database. For the purposes of our experiments, only the non-augmented files are used. The audio files are downsampled to 8 kHz and the different database labels resembling telephone speech quality in this respect and are collapsed into two sets of experiment labels: *covid positive*, which includes all the COVID-19 labels, while all the other label categories are classified as *covid negative*. The summary statistics of the labels of the unaugmented audio files can be found in Table I.

<sup>1</sup>Please refer to <https://www.covid-19-sounds.org/en/app/> for more details.

Model description							Valid.	UAR	AUC	ACC	ACC <sub>+</sub>	ACC <sub>-</sub>
Baseline (openSMILE ComParE features + SVM)							-	59.3	61.0	56.6	67.2	51.3
index	$N$	$C$	$F$	$\lambda$	$\alpha$	Fold						
1	5	77	2	1.582	$4.8 \times 10^{-5}$	1	65.8	71.4	76.0	72.6	68.1	<b>74.8</b>
2	2	308	2	2.093	$5.9 \times 10^{-6}$	1	61.4	66.0	74.8	59.4	85.3	46.6
3	2	235	2	1.805	$2.6 \times 10^{-5}$	1	60.7	64.6	72.8	62.3	71.6	57.7
4	3	466	2	1.859	$7.1 \times 10^{-6}$	1	60.6	63.4	71.1	59.7	74.1	52.6
5	4	73	2	1.917	$9.7 \times 10^{-6}$	2	63.8	53.4	59.6	52.3	56.9	50.0
6	5	151	1	2.176	$1.8 \times 10^{-6}$	2	62.1	57.7	73.1	44.9	<b>95.7</b>	19.7
7	3	486	1	1.678	$8.8 \times 10^{-6}$	2	61.4	60.5	65.3	58.6	66.4	54.7
8	5	435	1	1.351	$3.7 \times 10^{-5}$	2	59.9	69.1	72.9	68.3	71.6	66.7
9	4	445	2	0.748	$2.3 \times 10^{-5}$	3	63.2	58.5	64.8	50.9	81.0	35.9
10	2	215	2	1.218	$7.3 \times 10^{-6}$	3	63.2	52.4	60.4	52.3	52.6	52.1
11	3	486	1	0.882	$4.3 \times 10^{-6}$	3	61.3	52.9	59.3	61.4	27.6	78.2
12	6	424	1	1.328	$2.3 \times 10^{-5}$	3	60.6	70.4	74.8	68.9	75.0	65.8
Ensemble description							Ensembled group					
Best test UAR							-	<b>74.9</b>	80.5	<b>73.1</b>	80.2	69.7
Best test AUC							-	74.5	<b>80.7</b>	72.9	79.3	69.7
Best model per fold							-	70.8	77.3	68.3	78.4	63.2
All							-	70.2	77.6	67.7	77.6	62.8

**TABLE II:** Results of the different models using the introduced performance measures in percentages: validation UAR on hold out set, followed by test set UAR, AUC, ACC, and class-wise (+/-) ACC (recall/sensitivity), respectively. The models are the baseline, several CNNs with different hyperparameters or training folds, and ensembled models.  $N, C, F$  are model hyperparameters (Subsection III-B), and  $\lambda, \alpha$  are training hyperparameters (Subsection III-C). The models used for the ensemble groups are specified by the index column. The best model per fold is specified by the best validation UAR on the hold out set.

Ensemble	UAR	breath UAR	cough UAR
1	71.4 ± 4.9	72.2 ± 7.0	70.9 ± 6.7
8	69.1 ± 4.8	70.6 ± 7.5	67.8 ± 6.8
12	70.6 ± 5.4	68.1 ± 7.2	72.6 ± 6.9
1,2,6,7,8,11	<b>74.9 ± 4.5</b>	<b>76.1 ± 6.4</b>	<b>73.7 ± 7.2</b>
1,2,6,7,8,11,12	74.5 ± 4.4	75.7 ± 6.5	73.2 ± 6.7
1,5,9	70.8 ± 4.8	71.0 ± 6.9	70.7 ± 6.4
1-12	70.2 ± 4.8	71.4 ± 6.5	69.0 ± 7.6

**TABLE III:** An overview of the 95 % confidence intervals of a selection of models for the UAR metric: for the whole dataset and for each modality separately. The selected models consist of the ensemble models and the best test UAR individual models from each fold.

We split the data in a stratified, strictly subject-independent manner into three roughly balanced sets: train, dev(elopment), and test. The train set consists of 464 samples, the dev set consists of 220 samples, and the test set consists of 350 examples. The models in our experiments are trained using  $k$ -fold nested cross validation [25], with  $k = 3$ , wherein each model is trained on two thirds of the mixture of the train and dev sets, with its performance during training evaluated on the remaining hold out part<sup>2</sup>. We saved models with best Unweighted Average Recall (UAR) on the hold out set.

In other words, we merge the train and development sets, and then split the resulting set in a speaker independent manner into three folds. Then, we train different models on different folds, and we test them at the end on the overall ‘final’ test set, which is always strictly held out, also in the case of model fusions. During training, we validate on the hold out fold and never look at the test set. Likewise, all four parts: Fold1, Fold2, Fold3, and the test set are pairwise mutually exclusive in terms of speakers, so an individual speaker appears only in one of these four. Fold1 is what we declare as the dev(elopment)

<sup>2</sup>The split indices are available by the authors for reproducibility.

part for the baseline. Note that, as the splits are generated in a stratified manner, in all splits, the COVID-19/total ratio is close to the 27 % of the COVID-19/total ratio of the whole dataset.

Also, one can see that on some folds, high validation set UAR yields bad test set UAR and vice versa, which shows we are not biasing our training in order to get good test set results.

## B. Evaluation metrics

We adapt several classification metrics:

- Recall COVID-19 positive (ACC<sub>+</sub>), also known as sensitivity, which is the true positive rate.
- Recall COVID-19 negative (ACC<sub>-</sub>), also known as specificity, which is the true negative rate.
- Unweighted Average Recall (UAR): Also known as Unweighted Average Accuracy, is the average of sensitivity and specificity – the standard competition measure in the Interspeech ComParE challenge series. Chance level for two classes resembles 50.0 % UAR.
- Area Under ROC Curve (AUC): The ROC curve plots the true-positive rates against the false-positive rates with varying classification thresholds, and measures the area under the drawn curve. This measures how well a model can distinguish between the positive and negative classes, where a random baseline classifier will get a value of 50.0 %.
- Accuracy (ACC): The ratio of examples that are answered correctly in the evaluation set.

UAR is the preferred measure when compared to Accuracy, as Accuracy will reward trivial classifiers on unbalanced data, where the classifier always just predicts the more frequent class (here, the negative class).

### C. Results and discussion

In Table II, we showcase the following models: The baseline model, trained using openSMILE ComParE features (introduced in Subsection III-D), individually trained models (introduced in Subsection III-B), whose hyperparameters are tuned using BOHB. The range for the parameters examined is  $N$  in [2,6],  $C$  in [64,512],  $F$  in [1,2] (number of fully connected layers at the end, including the final classification layer),  $\alpha$  log sampled from  $[10^{-7}, 10^{-3}]$ , and  $\lambda$  in [0.5,2.2]. Note that, if lambda is less than one, the model will focus on the negative class, which non-intuitively yielded good results in few cases. Finally, we present the ensemble models (introduced in Subsection III-C). The performance of the models is shown using the metrics introduced in Subsection IV-B. The metrics themselves indicate performance on the test set, while the validation column indicates the performance of the individual models on the hold-out set of their particular fold.

As can be seen in Table II, all the variations of ensembled models outperform the baseline by a wide margin on all metrics. The same cannot be said for the individual models trained on one of the three total folds: while all of the models trained on the first fold outperform the baseline for all the metrics, only two to three of the models trained on folds two and three surpass the baseline, depending on the metric. That being said: model six produces the best overall positive class recall score of 95.7 %, while model eleven does the same for the negative class recall, with a score of 78.2 %, with both of them surpassing the respective metric scores of any of the ensembles, yet, at the cost of the respective other class. Still, these results probably justify why both models were included in both of the best performing ensembles.

In terms of ensembles, we show the following configurations: An ensemble that provides the best test set UAR, an ensemble that provides the best test set AUC, an ensemble of the best model per fold (based on the UAR metric) and finally, an ensemble of all the trained models. For UAR, the highest result is 74.9 %, while for AUC, the highest result is 80.7 %. Selecting only the best model per fold and ensembling those together results in performances of 70.8 % UAR and 77.3 % AUC. Finally, ensembling all models together reaches the results of 70.2 % UAR, and 77.6 % AUC.

In order to provide an insight into uncertainty of the results, in Table III, we calculate the 95 % confidence intervals of the UAR metric for deeper inspection for the ensembled models and the best test UAR individual model from each fold. The table shows that breathing modality seems generally better to detect COVID-19 as compared to the coughing modality. However, the results are not significant. Table II shows that the models trained on fold 1 are on average much more capable than their equivalents trained on folds 2 and 3, for UAR, AUC And ACC metrics. The reason for this discrepancy is not clear, it could be due to some random factors related to the model structure or parameter initialisation. Another possible reason is that, the validation scores on the holdout portions for folds 2 and 3, unlike fold 1, are not representative of the test data,

as the models of better validation UAR tend to get worse test UAR and vice versa.

From Table II, it can be seen that it appears indeed possible to use cough and breathing audio samples, fed to neural network models to predict whether a patient has COVID-19 or not, e.g., for a pre-diagnosis to pre-select candidates for more reliable, yet more effort and cost requiring testing. The results shown above should be considered as initial results, given the limited size of the data set of two hours of original samples, which prevents us from using larger, more complex neural networks. The effects of the small size of the dataset can be seen in Table II, which shows that the specific selection of which fold a model is trained on impacts the average predictive strength of the model. This is factor that can be reduced by increasing the data size, and thus allowing for fold splits with more uniform information content. Table III shows that the UAR performances for the breath modality are consistently higher, suggesting that breath audio samples contain more COVID-19 information, or at the very least information that is easier to extract. These results would suggest that future data collection efforts should place a particular focus on collecting more breathing audio recordings, in addition to more validated COVID-19 test results, as well as a more diverse range of control group samples from other respiratory diseases.

### V. CONCLUSION

In this paper, we explored the usage of deep learning models as a way to predict whether someone is infected with COVID-19 based on an audio sample of either their breathing or their coughing. The need for this usage arose from issues with lacking COVID-19 testing capacity in developing countries across the world, as opposed to the abundance of mobile phone quality microphones, but also the general opportunities coming with real-time low cost pre-scanning for selective testing with more reliable approaches. Accordingly, the aim of the models would be to function as a ubiquitous, low-cost pre-testing mechanisms that could help mitigate the demand for COVID-19 lab tests, which are relatively expensive to conduct, as they require access to materials, equipment and manpower that are not equally available around the world.

To this end, we used a subset of a crowdsourced database collected via the University of Cambridge's COVID-19 Speech Android app and web interface. The database contained samples of breathing and coughing recordings, as well as associated demographic information, medical history, and COVID-19 testing status. We illustrated how we pre-processed the database, splitting it in a stratified, strictly subject-independent manner into 3-fold train and development sets, as well as an independent test set. We then showed how we trained a number of individual Convolutional Neural Networks (CNNs), which we then ensembled together in order to produce our predictions. Our proposed models achieved at best a UAR score of 74.9 % and an AUC score of 80.7 % on the held-out speaker independent test partition.

The achieved results suggest that it is indeed possible to detect COVID-19 by way of either breath or cough samples with

an accuracy relevant to use-cases such as pre-selection for more reliable testing, and also possible to use deep learning models to perform this detection. However, the current results are limited by amount of available data, which might prevent the usage of even larger models, which is where deep learning models tend to produce their best results. A future direction for this research would be to collect a larger database with highly validated and more varied control data, including a plethora of other respiratory and further related diseases, which would open the door to even better, but also more tangible results.

## VI. ACKNOWLEDGEMENT

The University of Cambridge does not bear any responsibility for the analysis or interpretation of the data used herein, which represents the own view of the authors.

## REFERENCES

- [1] A. A. for Clinical Chemistry, "Coronavirus Testing Survey," <https://www.aacc.org/science-and-research/covid-19-resources/aacc-covid-19-testing-survey>, 2020.
- [2] Foundation for Innovative New Diagnostics, "SARS-COV-2 Test Tracker," <https://www.finddx.org/covid-19/test-tracker/>, 2020.
- [3] C. Gottlieb, J. Grobovsek, M. Poschke, and F. Saltiel, "Lockdown accounting," *Discussion Paper Series*, no. 13397, 2020.
- [4] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Proc. NIPS*, vol. 25, Lake Tahoe, CA: Curran Associates, Inc., 2012, pp. 1097–1105.
- [5] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision*, vol. 115, pp. 211–252, 2015.
- [6] I. Kononenko, "Machine learning for medical diagnosis: history, state of the art and perspective," *Artificial Intelligence in Medicine*, vol. 23, pp. 89–109, 2001.
- [7] P. Hamet and J. Tremblay, "Artificial intelligence in medicine," *Metabolism*, vol. 69, pp. 36–40, 2017.
- [8] S. Moein, *Medical diagnosis using artificial neural networks*. IGI global, 2014.
- [9] S. M. McKinney, M. Sieniek, V. Godbole, J. Godwin, N. Antropova, H. Ashrafi, T. Back, M. Chesus, C. G. S., A. Darzi, M. Etemadi, F. Garcia-Vicente, F. J. Gilbert, M. Halling-Brown, D. Hassabis, S. Jansen, A. Karthikesalingam, C. J. Kelly, D. King, L. J. R., D. Melnick, H. Mostofi, L. Peng, J. J. Reicher, B. Romera-Paredes, R. Sidebottom, M. Suleyman, D. Tse, K. C. Young, J. De Fauw, and S. Shetty, "International evaluation of an AI system for breast cancer screening," *Nature*, vol. 577, no. 7788, pp. 89–94, 2020.
- [10] D. Ardila, A. P. Kiraly, S. Bharadwaj, B. Choi, J. J. Reicher, L. Peng, D. Tse, M. Etemadi, W. Ye, G. Corrado, D. P. Naidich, and S. Shetty, "End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography," *Nature medicine*, vol. 25, pp. 954–961, 2019.
- [11] W. L. Bi, A. Hosny, S. M. B., M. L. Giger, N. J. Birkbak, A. Mehrtash, T. Allison, O. Arnaout, C. Abbosh, I. F. Dunn, R. H. Mak, R. M. Tamimi, C. M. Tempny, C. Swanton, U. Hoffmann, L. H. Schwartz, R. J. Gillies, R. Y. Huang, and H. J. W. L. Aerts, "Artificial intelligence in cancer imaging: clinical challenges and applications," *CA: a cancer journal for clinicians*, vol. 69, pp. 127–157, 2019.
- [12] B. W. Schuller, D. M. Schuller, K. Qian, J. Liu, H. Zheng, and X. Li, "Covid-19 and computer audition: An overview on what speech & sound analysis could contribute in the SARS-CoV-2 Corona crisis," *arXiv preprint*, no. 2003.11117, 2020.
- [13] S. A. Harmon, T. H. Sanford, S. Xu, E. B. Turkbey, H. Roth, Z. Xu, D. Yang, A. Myronenko, V. Anderson, A. Amalou, M. Blain, M. Kassan, D. Long, N. Varble, S. M. Walker, U. Bagci, A. M. Ierardi, E. Stellato, G. G. Plensich, G. Franceschelli, C. Giraldo, G. Irmici, D. Labella, D. Hammoud, A. Malayeri, E. Jones, R. M. Summers, P. L. Choyke, D. Xu, M. Flores, K. Tamura, H. Obinata, H. Mori, F. Patella, M. Cariati, G. Carrafiello, P. An, B. J. Wood, and B. Turkbey, "Artificial intelligence for the detection of COVID-19 pneumonia on chest CT using multinational datasets," *Nature communications*, vol. 11, pp. 1–7, 2020.
- [14] K. Zhang, X. Liu, J. Shen, Z. Li, Y. Sang, X. Wu, Y. Zha, W. Liang, C. Wang, K. Wang, L. Ye, M. Gao, Z. Zhou, L. Li, J. Wang, Z. Yang, H. Cai, J. Xu, L. Yang, W. Cai, W. Xu, S. Wu, W. Zhang, S. Jiang, L. Zheng, X. Zhang, L. Wang, L. Lu, J. Li, H. Yin, W. Wang, O. Li, C. Zhang, L. Liang, T. Wu, R. Deng, K. Wei, Y. Zhou, T. Chen, J. Y.-N. Lau, M. Fok, J. He, T. Lin, W. Li, and G. Wang, "Clinically Applicable AI System for Accurate Diagnosis, Quantitative Measurements, and Prognosis of COVID-19 Pneumonia Using Computed Tomography," *Cell*, vol. 181, pp. 1423–1433.e11, 2020.
- [15] G. Deshpande and B. Schuller, "An Overview on Audio, Signal, Speech, & Language Processing for COVID-19," *arXiv preprint*, no. 2005.08579, 2020.
- [16] F. Shi, J. Wang, J. Shi, Z. Wu, Q. Wang, Z. Tang, K. He, Y. Shi, and D. Shen, "Review of Artificial Intelligence Techniques in Imaging Data Acquisition, Segmentation and Diagnosis for COVID-19," *IEEE Reviews in Biomedical Engineering*, pp. 1–1, 2020.
- [17] World Bank Group, *World development report 2016: Digital dividends*. The World Bank, 2016.
- [18] K. Qian, B. W. Schuller, and Y. Yamamoto, "Recent Advances in Computer Audition for Diagnosing COVID-19: An Overview," *arXiv preprint*, no. 2012.04650, 2020, 2 pages.
- [19] P. Bagad, A. Dalmia, J. Doshi, A. Nagrani, P. Bhamare, A. Mahale, S. Rane, N. Agarwal, and R. Panicker, "Cough against covid: Evidence of covid-19 signature in cough sounds," 2020.
- [20] A. Hassan, I. Shahin, and M. B. Alsabek, "COVID-19 Detection System using Recurrent Neural Networks," in *International Conference on Communications, Computing, Cybersecurity, and Informatics (CCCI)*, 2020, pp. 1–5.
- [21] K. V. S. Ritwik, S. B. Kalluri, and D. Vijayasenan, "COVID-19 Patient Detection from Telephone Quality Speech Data," *arXiv preprint*, no. 2011.04299, 2020.
- [22] P. C. Loizou, *Speech enhancement: theory and practice*. CRC press, 2013.
- [23] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, pp. 1929–1958, 2014.
- [24] J. Szep and S. Hariri, "Paralinguistic Classification of Mask Wearing by Image Classifiers and Fusion," in *Proc. Interspeech*, Shanghai, China, 2020, pp. 2087–2091.
- [25] G. C. Cawley and N. L. Talbot, "On Over-Fitting in Model Selection and Subsequent Selection Bias in Performance Evaluation," *Journal Machine Learning Research*, vol. 11, p. 2079–2107, 2010.
- [26] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*. Springer, 2001, vol. 1.
- [27] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *3rd International Conference on Learning Representations, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds. San Diego, CA: ICLR, 2015.
- [28] S. Falkner, A. Klein, and F. Hutter, "BOHB: Robust and Efficient Hyperparameter Optimization at Scale," in *Proc. ICML*, J. Dy and A. Krause, Eds., vol. 80. Stockholm, Sweden: MLR Press, 2018, pp. 1437–1446.
- [29] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2006.
- [30] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: the munich versatile and fast open-source audio feature extractor," in *Proc. 18th ACM international conference on Multimedia*. Florence, Italy: ACM, 2010, pp. 1459–1462.
- [31] B. Schuller, S. Steidl, A. Batliner, J. Hirschberg, J. K. Burgoon, A. Baird, A. Elkins, Y. Zhang, E. Coutinho, and K. Evanini, "Computational Paralinguistics Challenge: Deception, Sincerity & Native Language," in *Proc. Interspeech*, San Francisco, CA, 2016, pp. 2001–2005.
- [32] C. Brown, J. Chauhan, A. Grammenos, J. Han, A. Hasthanasombat, D. Spathis, T. Xia, P. Cicuta, and C. Mascolo, "Exploring automatic diagnosis of covid-19 from crowdsourced respiratory sound data," in *Proc. KDD*, virtual conference, 2020, pp. 3474–3484.
- [33] J. Schlüter and T. Grill, "Exploring Data Augmentation for Improved Singing Voice Detection with Neural Networks," in *Proc. 16th International Society for Music Information Retrieval Conference*, Málaga, Spain, 2015, pp. 121–126.