



Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Transfer learning for the efficient detection of COVID-19 from smartphone audio data

Mattia Giovanni Campana^{a,*}, Franca Delmastro^a, Elena Pagani^b^a Institute for Informatics and Telematics of the National Research Council of Italy (IIT-CNR), Pisa, Italy^b Computer Science Department, University of Milano, Milan, Italy

ARTICLE INFO

Article history:

Received 25 November 2022

Received in revised form 18 January 2023

Accepted 25 January 2023

Available online 30 January 2023

Keywords:

Deep audio embeddings

Transfer learning

Deep learning

m-health

COVID-19

ABSTRACT

Disease detection from smartphone data represents an open research challenge in mobile health (m-health) systems. COVID-19 and its respiratory symptoms are an important case study in this area and their early detection is a potential real instrument to counteract the pandemic situation. The efficacy of this solution mainly depends on the performances of AI algorithms applied to the collected data and their possible implementation directly on the users' mobile devices. Considering these issues, and the limited amount of available data, in this paper we present the experimental evaluation of 3 different deep learning models, compared also with hand-crafted features, and of two main approaches of transfer learning in the considered scenario: both feature extraction and fine-tuning. Specifically, we considered VGGish, YAMNET, and L³-Net (including 12 different configurations) evaluated through user-independent experiments on 4 different datasets (13,447 samples in total). Results clearly show the advantages of L³-Net in all the experimental settings as it overcomes the other solutions by 12.3% in terms of Precision-Recall AUC as features extractor, and by 10% when the model is fine-tuned. Moreover, we note that to fine-tune only the fully-connected layers of the pre-trained models generally leads to worse performances, with an average drop of 6.6% with respect to feature extraction. Finally, we evaluate the memory footprints of the different models for their possible applications on commercial mobile devices.

© 2023 Elsevier B.V. All rights reserved.

1. Introduction

COVID-19 pandemic largely changed our lives, not only for the impact of the disease on the entire world population, but also for an emerging perception of human weaknesses to the virus's spread and the dependency on the limitations of national healthcare systems. In the last couple of years, people became more conscious of the risks, and they increased their trust in digital solutions aimed at counteracting and mitigate the pandemic. The digital solutions proposed by researchers are mainly characterized by a massive use of Artificial Intelligence (AI) technologies and big data [1,2]; data derived from clinical analysis (e.g., blood test [3], X-ray and lung CT images [4]), but also from commercial personal and wearable devices collecting physiological and/or audio signals [5].

Mobile health (m-health) systems have the potentiality to become scalable and low-cost solutions for fast screening, aimed at recognizing the onset of new COVID-19 cases, thus possibly preventing new outbreaks. To this aim, they have to rely on smartphone-embedded sensors, especially microphones, to recognize or distinguish the main COVID symptoms,

* Corresponding author.

E-mail addresses: m.campana@iit.cnr.it (M.G. Campana), f.delmastro@iit.cnr.it (F. Delmastro), elena.pagani@unimi.it (E. Pagani).

like cough and breath, from those of either other pathologies or healthy conditions. We investigated previous works in the literature, and we found interesting preliminary works. Schuller et al. [6] have been the first to investigate how the automatic analysis of speech and audio data can contribute to fight the pandemic crisis, presenting the potential of Computer Audition techniques (CA, i.e., computer-based speech and sound analysis) [7]. Subsequently, researchers investigated the effective applicability of those techniques in real scenarios. The main difficulty they encountered was the collection of objective data from large populations. Initial studies focused on small patients' cohorts [8] even though the target of the work would require a huge amount of data that could not be collected rapidly. The authors of [8] presented both a preliminary evaluation of a cough detector system aimed at distinguishing cough signals from noise, and an AI tool for COVID-19 diagnosis based on data collected from 70 subjects in controlled environments. Other works released mobile and web apps to directly collect crowdsourced datasets from the population [9–11]. Some of them focused only on cough audio data, while others introduced also speech signals, trying to improve the classification performances. Respiratory sound samples (e.g., cough and breath) are generally processed by using standard modeling procedures proposed in the CA literature to extract different sets of features (referred as *hand-crafted acoustic features*) [12]. Different Deep Learning (DL) approaches have been proposed [13,14], including the use of pre-trained models for latent features extraction from raw audio signals, also called *deep audio embeddings*, to enrich standard CA features [9].

In our previous paper [15], we investigated the feasibility of using a novel embedding model called *Look, Listen and Learn* (L³-Net) [16] to improve the detection of COVID-19 from respiratory sounds data, including both coughing and breathing recordings. Specifically, we employed a pre-trained version of L³-Net to extract latent features from audio files, thus relying on transfer learning to characterize raw audio samples in a low-dimensional space, which highlights the differences among the data. By comparing L³-Net performances with the other proposed models, we demonstrated that it overcomes the others in several standard metrics, scoring an overall gain of 28.57% of AUC, 23.75% of Precision, and 39.43% in terms of Recall.

In this work, we present a more thorough investigation of possible solutions to efficiently detect COVID-19, comparing the use of hand-crafted acoustic features and two different transfer learning techniques aimed at adapting pre-trained DL models: *feature extraction* and *fine-tuning*. Specifically, we can summarize the novel contributions of this work as follows:

- We evaluate the effectiveness of detecting COVID-19 from respiratory sounds data by using both hand-crafted acoustic features and deep audio embeddings extracted by 3 state-of-the-art pre-trained DL models: VGGish, YAMNET, and L³-Net;
- We study in details L³-Net, evaluating its classification performances based on different model's configurations;
- We compare the two main approaches of transfer learning for model adaptation, i.e., features extraction and fine-tuning, identifying their advantages and drawbacks in our application scenario;
- We perform an extensive evaluation of the considered solutions through a series of user-independent experiments by using 4 different benchmark datasets, which correspond to 8726 data samples more than those we used in our previous work;
- Finally, we study the memory footprint of both transfer learning approaches, thus evaluating their feasibility of being entirely executed on mobile devices.

The remainder of the paper is organized as follows. Section 2 presents the related work. In Section 3, we describe in details the procedure we follow to compare the different approaches for COVID-19 detection, including pre-processing of the raw audio sample, hand-crafted features extraction, and the use of deep audio embeddings models relying on the two transfer learning techniques. Then, Section 4 outlines the experimental setup we adopt to evaluate the classification performance of the considered approaches, and presents the obtained results. In Section 5, we evaluate the memory footprint of the two transfer learning techniques, considering their implementation on mobile devices. Lastly, in Section 6, we draw our conclusions and present some directions for future works.

2. Related work

It is largely recognized in literature that the use of deep learning techniques (namely, Convolutional Neural Networks, CNNs) for solving audio classification problems typically obtains better performances than classical ML methods [17]. However, labeled datasets available for research on the application of ML to automated COVID-19 detection usually involve just a few hundreds of samples. In such cases, *transfer learning* is the main technique for system training so as to achieve good predictions in case of scarce labeled samples [18]. Leveraging transfer learning with CNNs for audio classification problems is studied in some papers. In [19], the usage of both image and sound CNNs is studied. In the former case, data samples are image-based sound representations such as spectrograms. Exhaustive experiments are conducted, which show that both hyper-parameters and the choice of the retrained part of the neural network impact on accuracy and retraining time, but in all the cases transfer learning yields better results than training a CNN from scratch. The sound CNNs (namely VGGish and YAMNet [20]) obtained better results than image CNNs (GoogLeNet, SqueezeNet, ShuffleNet) with all the considered datasets.

Transfer learning leverages neural networks, pre-trained on large sets (millions) of generic labeled objects, as a starting point. Let \mathcal{G} be the generic dataset used to pre-train the network \mathcal{N} , and \mathcal{S} the dataset specific to the considered

classification problem. Two techniques are mainly used for transfer learning: either *feature extraction* or *fine-tuning*. In the former, the first n layers of \mathcal{N} – usually stopping prior the fully-connected layers – receive as input the data in \mathcal{S} , and their outputs are used as features of the objects in \mathcal{S} . Thanks to the pre-training of \mathcal{N} , these should be the relevant features of \mathcal{S} objects. The remaining layers are then trained for the specific classification problem. In the latter, the last (or a few last), most specialized, layer of \mathcal{N} is removed and substituted with a new, non-trained, one; the resulting network \mathcal{N}' is then trained on data from \mathcal{S} (thus avoiding starting from scratch).

The superiority of one technique over the other is not definitively established. Feature extraction is largely believed to avoid overfitting, i.e., excessive specialization of the network in properly classifying data with the same characteristics as those from the training set while poorly behaving with additional data of the same set; this phenomenon may easily occur with small datasets. In the computer vision field [21], the two methods are compared, and an analysis is conducted about where to split \mathcal{N} for an effective feature extraction. The drawn conclusions are that, in general, fine-tuning performs better than feature extraction; however, this benefit is smaller the more \mathcal{G} and \mathcal{S} differ. Furthermore, the benefit depends on the number n of preserved layers and it may drop for greater n due to too few general features and/or complex feature interactions. Similar hypotheses appear in [22] (sec.14.2.1) where an assumption for fine-tuning usage is that the knowledge learned and preserved by \mathcal{N}' is also applicable to \mathcal{S} . In the NLP field, the authors of [23] obtain comparable performances of the two methods in most of the cases, with fine-tuning outperforming feature extraction when \mathcal{G} and \mathcal{S} are aligned, while achieving poor performance in case of significantly different tasks. [24] analyzes the performances of fine-tuning applied to medical image processing. The presented results confirm that layers selection in fine-tuning affects the classification metrics, and an automatic method for layer selection is proposed and validated, which consists in solving an optimization problem. The proposed method shows an improvement for all the considered metrics, but at the cost of higher time spent for training. An analogue research is presented in [25] applied to computer vision.

Going through the specific application domain of respiratory sounds, some solutions exist in the literature, even though it may be difficult to classify these sounds due to their complicated structures and variety of noises [26]. [27] classifies lung sounds via feature extraction applied to pre-trained CNNs. It compares different techniques of feature extraction, considering three classifiers. It obtains the best results with the pre-trained ResNet18 CNN and short-time Fourier transformation as the mechanism for feature extraction, and it improves the state-of-the-art results on the same dataset. In [26], deep learning pre-trained CNNs are used, and transfer learning is leveraged via feature extraction, representing sounds with Mel-spectrograms. CNNs show better results than Support Vector Machine (SVM), and VGG-16 supplies the best results as feature extractor, with lower computation time than SVM; anyway, the best measured accuracy in detecting anomalous sounds is around 86% confirming the difficulty of the considered classification task. In [28], CNNs are used to classify lung diseases using audio classification, with a dataset of 920 samples reportedly challenging for other methods. Comparison is conducted between two policies: (i) performing feature extraction via a pre-trained (VGG-16) CNN whose output is then sent to SVM, or (ii) fine-tuning the pre-trained CNN using spectrograms of lung sounds. Both policies obtain better results than the state-of-the-art for the same dataset, with the former policy slightly outperforming (2.4%) the latter.

In the last couple of years, researchers have explored several audio processing techniques to develop effective and low-cost COVID-19 screening methods based on respiratory data [29], especially derived from smartphone embedded microphones. One of the most promising approach in this domain consists in converting the audio files into a visual representation (e.g., time–frequency spectrogram or Mel-spectrogram) that can be used as input to a CNN model for both features extraction and classification. This category includes, for example, the application AI4COVID proposed in [8] based only on cough recordings. Specifically, they modeled the audio sample as both Mel-spectrogram and MFCC, which are then processed by an ensemble model composed of two CNNs and one SVM to categorize the cough into 4 classes: COVID-19, bronchitis, pertussis, and normal cough. A similar solution has been proposed in [30], where different visual representations of cough recordings (e.g., Mel-spectrogram, Chromagram, and Power-Spectrogram) have been compared to train an end-to-end CNN architecture. Such approaches are particularly interesting because they avoid the features engineering and selection phases in the data processing pipeline, mainly relying on the intrinsic capabilities of DL to automatically model the raw input data. However, due to the scarcity of public COVID-19 respiratory sound data, their training has been performed on small-size datasets, typically composed by a few hundreds of samples. DL models, especially those with complex architectures, tend to overfit in such settings, often providing unreliable results.

By contrast, in [31] a large, crowdsourced, dataset is available. Cough, breathing, and voice recordings of the participants are converted in log-mel spectrograms, which feed VGGish [32], a CNN-based embedding model pre-trained on the large-scale YouTube-8M dataset (approximately 2.6 billion audio/video features), for feature extraction. The obtained feature vectors are then aggregated by a pooling layer; the resulting latent feature vector is passed on to the prediction layers. Thanks to the extension of the available dataset, no cross-validation is used. Instead, participants whose data are included in the training set, are not considered for the testing set. Exhaustive experiments are conducted; the results reveal that the considered dataset is challenging: a maximum ROC-AUC of 0.71 is obtained in average, leveraging all the three data sources mentioned above. By splitting the participants into categories, the best results are achieved for female participants and for over-60 participants. The considered method is also validated by intentionally introducing biases into the data and repeating the experiments. In Section 4, we present the results we obtained by conducting some experiments with the [31] dataset.

Another approach deals with the mentioned DL drawback by using a combination of hand-crafted acoustic features and audio embeddings extracted by pre-trained deep models. Representative of this category is [5], in which the authors




Fig. 1. COVID-19 detection approaches: (1) extraction of handcrafted acoustic features from the audio waveform, which are then classified by a shallow ML model; (2) usage of a pre-trained deep learning model as features extractor, in series with a shallow model to classify the deep audio embeddings; and (3) fine-tuning of the pre-trained deep model for both features extraction and classification.

used a set of acoustic features and a pre-trained DL model to train a shallow ML classifier (e.g., Logistic Regression, LR) to identify COVID-19 subjects from cough and breath audio recordings. Specifically, as deep features extraction model, they employed VGGish, thus taking advantage of Transfer Learning concept to deal with the scarcity of COVID-19 audio data [33].

Our work stands out from the related literature for the type of analysis we conduct in order to identify an effective solution to detect COVID-19 by using commercial mobile devices as low-cost and pervasive diagnostic tools. Indeed, while similar works consider just one or a few approaches of audio classification and evaluate their performance with a limited amount of data (typically only one dataset), here we provide a thorough comparison of several state-of-the-art techniques, including both shallow and deep learning models, by using four large-scale datasets collected from real devices, for a total of 13,448 audio samples. Moreover, to the best of our knowledge, our work is the first investigating the two main approaches of Transfer Learning (i.e., features extraction and fine-tuning) in this specific application scenario, comparing their performance and providing a feasibility study of their implementation on real-world resource-constrained devices.

3. Acoustic features and deep audio embeddings

Fig. 1 shows the flow diagram of the entire data processing we adopt to compare the three main approaches of audio classification, which include: (1) using handcrafted acoustic features to train a shallow ML classifier; (2) using a pre-trained model as features extractor to input a shallow classifier; and, finally, (3) the fine-tuning of a pre-trained model, thus creating a single solution for both audio modeling and classification.

In the first approach, a set of handcrafted acoustic features is firstly extracted from the user's audio sample; then, a Principal Component Analysis (PCA) is applied to reduce the dimensionality of the features vector, which is finally used to input a shallow ML classifier that performs the diagnosis as a binary classification problem, distinguishing between COVID-positive and COVID-negative samples.

As far as transfer learning is concerned, after the signal conversion in Spectrogram-like images, and the extraction of deep audio embeddings, we use a pooling layer that calculates the mean and standard deviation of each dimension of the deep audio embeddings across all the windows. This generates a single feature vector.

Then, based on the classification procedure, we can distinguish the two forms of transfer learning as follows: in (2), we first reduce the dimension of the embeddings vector through PCA to input a shallow ML classifier that provides the final prediction. By contrast, in (3) we train from scratch a series of fully-connected layers to process the input embeddings vector, where the output softmax layer generates the final probability vector used to classify COVID-19 positive and negative predictions.

In the following, we describe in detail both the acoustic features and the pre-trained deep models considered in this work to evaluate the different approaches of audio classification aimed at detecting COVID-19 from respiratory sounds.

3.1. Acoustic features extraction

In order to transform the raw audio sample into a numerical representation manageable by a ML model, legacy approaches rely on different sets of handcrafted acoustic features. To evaluate this approach, we adopt the common

procedure used in similar audio-based medical applications [7]. Firstly, the original audio sample recorded by the user's device microphone is re-sampled to a standard value for audio tasks (e.g., 16 kHz or 22 kHz). Then, we manually extract common audio features related to both the frame (i.e., a chunk of the audio) and the segment (i.e., the entire audio sample) perspectives from the raw audio waveform, including frequency-based, structural, statistical, and temporal characteristics.

Specifically, we characterize the user's audio sample by using the same acoustic features proposed in [9]:

Duration : total length (in seconds) of the audio sample;

Onset : number of pitch onsets (i.e., "events") in the audio file;

Tempo : the beats rate that occur at regular intervals throughout the whole audio signal;

Period : the frequency with the highest amplitude among those obtained from the Fast Fourier Transform (FTT);

RMS Energy : the root-mean-square value of the signal power, i.e., the magnitude of the short-time Fourier transform;

Spectral Centroid : the centroid value of the frame-wise magnitude spectrogram, which can be used to identify percussive and sustained sounds;

Roll-off Frequency : the frequency under which the 85% of the total energy of the frame-wise spectrum is contained;

Zero-crossing rate : the number of times the signal value crosses the zero axes, and it is computed for each frame;

Mel-Frequency Cepstral Coefficients (MFCC) : coefficients that represent the shape of the cosine transformation of the sound logarithmic spectrum, expressed in Mel-bands, which characterize the set of frequency audible by human beings;

Δ -MFCC, Δ^2 -MFCC : the first and second order derivatives of MFCC along the time dimension, respectively.

The total number of acoustic features we extract from the audio sample is 477, including standard statistics (e.g., mean, median, max/min values, and skewness) to describe time-series descriptors for the entire audio signal, i.e., for RMS Energy, Spectral Centroid, Roll-Off Frequency, Zero-crossing rate crossing, MFCC, Δ -MFCC, and Δ^2 -MFCC.

3.2. Pre-trained deep audio embeddings models

In order to evaluate the effectiveness of the transfer learning approaches in our specific application scenario, we take into account 3 state-of-the-art deep learning models for audio classification: YAMNET, VGGish, and L³-Net [34]. These models mainly differ in the network's architecture and training approach.

Specifically, VGGish [32] is a 24-layer-deep network based on the VGG architecture, a popular CNN-based model for image classification [35]. The input audio sample must be firstly re-sampled at 16 kHz and normalized in the range $[-1, +1]$. Then, the audio waveform is framed into sliding windows of 0.96 s with no overlap. Finally, for each frame, the corresponding Mel-spectrogram image is used to input VGGish, which relies on the CNN layers to produce a 128-dimensional embeddings representation of the input image, which can be further classified either by a set of fully-connected layers or a shallow ML model.

YAMNET follows the same idea of VGGish, but it is composed by 28 deep layers and it is based on the architecture of MobileNet [36], a lightweight CNN model especially designed to perform visual applications directly on mobile devices. In this case, the input waveform should be framed through a sliding windows of length 0.96 s and hop of 0.48 s, from which the model extracts a 1024-dimensional embeddings vector.

By contrast, L³-Net is based on a totally different approach. In fact, as we mentioned in Section 2, this model has been designed to learn deep representations by identifying if a video frame (i.e., an image) and an audio segment come from the same video source. This allows to train the model by exploiting the self-supervised approach: since both matched and mismatched image-audio pairs can be automatically generated by extracting the image and audio from the same or different videos, no manual labeling is required to train the model.

The architecture of L³-Net consists of two distinct CNN subnetworks, used to extract different embeddings for the video and audio inputs respectively. To check the correspondence between both embeddings, a fusion network is used, which concatenates the deep representations and uses two fully connected layers as well as a softmax layer for binary classification. After training, the two subnetworks can be used as two distinct models to extract deep audio and video embeddings, respectively.

L³-Net has been further extended in [37], where the authors investigated different design choices of the original model, and how they impact on the performances of downstream audio classifiers. In particular, they studied 3 main aspects of the model: (i) the input representation, considering Linear-frequency power spectrogram (L), Mel-spectrogram with 128 bands (M128), and Mel-spectrogram with 256 bands (M256); (ii) the type of audio samples used during the training, including both environmental sounds (E) and music (M); and (iii) the size of the deep audio embeddings, considering both

Table 1
L³-Net configurations.

Name	Training data	Embedding size	Input type
L3 E 512 L	Environmental sounds	512	Linear
L3 E 512 M128			Mel-128
L3 E 512 M256			Mel-256
L3 E 6144 L		6144	Linear
L3 E 6144 M128			Mel-128
L3 E 6144 M256			Mel-256
L3 M 512 L	Music	512	Linear
L3 M 512 M128			Mel-128
L3 M 512 M256			Mel-256
L3 M 6144 L		6144	Linear
L3 M 6144 M128			Mel-128
L3 M 6144 M256			Mel-256

512 and 6144 dimensions. Based on this study, the authors publicly released OpenL3,¹ a framework that includes all the 12 investigated configurations of L³-Net, which we summarized in Table 1 for the sake of clarity.

In order to compare the different approaches of audio processing for COVID-19 detection, we rely on the pre-trained instances of the aforementioned deep learning models, thus taking advantage of their ability of characterizing audio data samples obtained by their training with a massive amount of data. In particular, for VGGish and YAMNET, we use the pre-trained versions available on *TensorFlow Hub*,² which have been trained by using the YouTube-8M [38] (6.1 million videos) and AudioSet [7] (more than 2 million of data samples), respectively. We also test all the configurations of OpenL3, which have been trained on different subsets of AudioSet.

4. Experimental evaluation

In this section, we describe both the benchmark datasets and the experimental protocol we use to evaluate the performances of the different approaches presented in Section 3. In a first set of experiments, we compare the effectiveness of both handcrafted acoustic features and deep audio embeddings in detecting COVID-19 from respiratory sounds data. In a second set of experiments, we evaluate the feasibility of fine-tuning the considered deep learning models with the available data, thus creating a single component for both features extraction and classification.

4.1. Datasets

In order to perform our experiments, we rely on 4 crowdsourced datasets commonly used in the literature to benchmark AI-based solutions for COVID-19 detection from respiratory sounds, including both breathing and coughing recordings: COSWARA [10] and Coughvid [39] are publicly available, while we obtained the access to two different versions of the COVID-19 Sounds dataset collected by the Cambridge University,³ here called Cambridge KDD [5] and Cambridge npj [31].

The first dataset is part of the COSWARA research project of the Indian Institute of Science (IISc), Bangalore, attempting to build a diagnostic tool for COVID-19 using different audio recordings of individuals, including breathing, cough and speech sounds. Currently, the project is still ongoing and it is continuing the data collection stage through crowdsourcing. Through the use of a web and a mobile application, the researchers ask volunteers to send their health status along with different types of audio recordings: two samples of cough (shallow and heavy), two audios of breath (shallow and deep), two recordings of counting numbers (normal and fast), and the phonation of sustained vowels.

The dataset is freely available on the official Github repository of the project.⁴ We downloaded the whole dataset available on February, 26th 2022, which includes audio recordings collected up to 24th February 2022. In this work, we considered only valid files (i.e., not corrupted or empty recordings) of both breath and cough recordings donated by people without other respiratory diseases (e.g., asthma), including both the available modalities, i.e., shallow/heavy cough, and shallow/deep breath. The final dataset includes a total of 2501 audio samples shared by people who have declared they were positive to COVID-19 at record time (1234 breath samples, and 1267 cough samples), and 860 samples labeled as negative examples (425 breath and 435 cough).

Similarly to COSWARA, the Coughvid corpus has been collected through a web application by a research group of the École Polytechnique Fédérale de Lausanne, Switzerland, between April 1st, 2020 and December 1st, 2020, and it is freely

¹ <https://openl3.readthedocs.io>.

² <https://www.tensorflow.org/hub>.

³ <https://www.covid-19-sounds.org>.

⁴ <https://github.com/iiscleap/Coswara-Data>.




Fig. 2. Number of audio samples in the 4 considered datasets, grouped by their respective labels and modalities.

available as a Zenodo repository.⁵ The authors collected a total of 25,000 crowdsourced samples representing a wide range of participant ages, genders, and geographic locations. Moreover, among the different metadata, each record has been labeled according to the output of a ML classifier specifically trained by the authors to estimate the probability that the input audio sample actually represents a cough recordings. Taking into account only those records with a high cough estimation (i.e., greater than 80%), we selected a total of 547 COVID-positive samples and 5625 negative recordings.

Finally, both Cambridge KDD and Cambridge npj have been collected by the Mobile System Research Lab of the University of Cambridge as part of the ERC EAR research project. Similarly to the previous datasets, they contain respiratory sounds collected by using both web and mobile applications. Cambridge KDD represents the first dataset used by the authors to evaluate the performance of a VGGish-based solution to detect COVID-19 from both breathing and coughing recordings [5], and it contains 282 COVID-positive samples (141 breath and 141 cough) and 660 negative samples (330 breath and 330 cough). More recently, in [31] the same authors have studied how different data selection strategies can overestimate the performances of ML models in detecting COVID-19 from respiratory sounds, and they released a second dataset to avoid potential acoustic bias, selecting only English-speaking participants and balancing the data based on the users' demographics.⁶ In this work, we refer to such corpora as Cambridge npj to distinguish it from the former dataset, and it contains a total of 1468 (734 breath and 734 cough) audio samples donated by COVID-positive subjects and 1504 (752 breath and 752 cough) recordings of healthy people.

Fig. 2 shows the main peculiarities of the four datasets, highlighting the number of audio files available for both negative and positive examples, according to the available ground truth. As we can note, most of the datasets are characterized by a large class imbalance, which can lead to obtain biased learning models that have a poorer predictive accuracy over the minority classes compared to the majority classes [40]. To cope with this issue, we balanced all the datasets by using random under sampling, thus randomly removing samples from the majority class in each corpus. In order to make the experiments reproducible, we share both the name of the files used in our experiments and the code on our Github repository.⁷

4.2. Evaluation protocol

We follow a rigorous protocol for the experiment's evaluation. Specifically, while similar works in the literature employ simple dataset splits that can lead to overestimate the performance of a predictor, we decided to rely on the 5-fold Nested Cross-Validation approach with user-based splits to avoid biasing the results based on patterns of specific users or specific train/validation/test splits of the dataset. The procedure can be summarized as follows. Firstly, we divide the whole dataset into 5 non-overlapping folds, ensuring that all the recordings generated by the same user are not placed in two different folds. In an outer loop, we iterate over the 5 folds, where, for each iteration, one fold is selected as the test set, while the remaining 4 represent the development set. The development set is then used in an inner loop for hyperparameters tuning, where it is further divided in 5 non-overlapping folds to define both training and validation sets.

⁵ <https://doi.org/10.5281/zenodo.4048312>.

⁶ While the access to the raw data requires a formal Data Access Agreement, the code required to produce the balanced dataset is freely available on Github at <https://github.com/cam-mobsys/covid19-sounds-npjDM>.

⁷ <https://github.com/mattiacampana/Transfer-Learning-COVID-19>.

Table 2

Hyperparameters tuned in both the Features Extraction (FE) and Fine-Tuning (FT) experiments. Algorithm DE indicates the classification layers added to the Deep Embeddings models considered in our experiments.

Task	Algorithm	Parameter	Values
FE	SVM	Regularization	$[10^{-3}, \dots, 10^3]$
		Kernel	[rbf, poly, sigmoid]
		Kernel coefficient	$[10^{-3}, \dots, 10]$
		Degree of poly kernel	$[2, \dots, 5]$
	AB	Estimators	[10, 20, 50, 100]
		Learning rate	[1, .5, .1, .05, .01, .001]
LR	Penalty	[1, 12]	
	Regularization	$[10^{-3}, \dots, 10^3]$	
RF	Estimators	[10, 20, 50, 100]	
	Min samples split	[2, 8, 10, 12]	
	Max depth	[10, 30, 50]	
	Split criterion	[entropy, gini]	
PCA	Explained variance	[.6, .65, .7, .75, .8, .85, .9, .95, .99]	
FT	DE	Hidden layers	[1, ..., 5]
		Hidden units	[128, 512, 1024, 2048, 6144]
		Dropout rate	[0, .1, .2, .3, .4]

The second step of the evaluation protocol differs based on the Transfer Learning approach. For the features extraction experiments, where the pre-trained deep embeddings models are simply used as feature extractors, the tuning phase includes: (i) finding the best PCA coefficient, that is, the amount of variance that needs to be explained by the held components; and (ii) finding the best shallow classifier and tuning its parameters by using the *Bayesian optimization algorithm* [41]. In these experiments, we test 4 broadly used ML algorithms: Logistic Regression (LR), Support Vector Machines (SVMs), Random Forest (RF), and AdaBoost (AB).

By contrast, in the set of fine-tuning experiments, we tune the considered deep embedding models by using the *Hyperband* algorithm, which is broadly used to optimize neural network models because it is able to speed up the random search of the parameter spaces through adaptive resource allocation and early-stopping [42]. Specifically, we exploit such optimization algorithm to tune the hyperparameters of the additional classification layers added to the base deep embeddings models, including the number of fully-connected layers, the number of hidden units in each layer, and the percentage of dropout rates to control the over-fitting of the model.

Table 2 summarizes the value spaces of the hyperparameters considered during the tuning phase of both features extraction and fine-tuning experimental settings. In addition, for each model, we also evaluate which type of audio files (i.e., modality) allows us to obtain the best performances among those available in each dataset: *Cough* (C), *Breath* (B), or the combination of the two (CB).

Finally, we evaluate the classification performances over the outer splits by using the Precision–Recall Area Under Curve (PR-AUC), which summarizes the Precision–Recall curve with a wide range of threshold values as a single score, and it is commonly used to evaluate binary classification models with imbalanced datasets [43].

4.3. Features extraction results

In the first set of experiments we compare the quality of different audio features as input of standard shallow classifiers in COVID-19 detection task. Such features include both the hand-crafted acoustic features (HC) presented before, and the audio embeddings extracted by YAMNET, VGGish, and OpenL3.

Fig. 3 summarizes the obtained results with all the considered datasets and the relative audio modalities, indicating as *OpenL3 (best)* the configuration of OpenL3 that presents the best performances. As we can note, the performances of each approach strongly differs for each combination of dataset and audio modality but, in general, OpenL3 is able to obtain the best results in all the experimental settings. Compared with the other deep models, the embeddings of OpenL3 allows to achieve an overall improvement of 12.3% with respect to the deep features generated by YAMNET, and of 7.2% more than those extracted by VGGish. Surprisingly, the lowest improvement obtained by OpenL3 is the one compared with HC (6.7%), which perform better than YAMNET and VGGish in all the experimental settings, except for Cambridge npj (B) and Cambridge npj (CB), where the handcrafted features achieve the same or slightly worse results than the two deep embedding models. According to the obtained results, Cambridge npj represents a challenging dataset, where all the considered solutions achieve lower performances compared with the results achieved in the other experimental settings, ranging from 0.62 and 0.67 of PR-AUC, except for OpenL3 that can reach 0.70 by using only cough samples.

In order to analyze in detail the different solutions, we also report the shallow classifiers and PCA explained variance rate that led to the best results. Specifically, Table 3 shows the details for the Coswara and Coughvid datasets, while Table 4 shows those with Cambridge KDD and Cambridge npj. With the former dataset, the best configuration is different for each modality. HC performs best with Breath recordings, reaching a PR-AUC score of 0.74 by using all the




Fig. 3. Overall results for the features extraction experiments. OpenL3 (best) refers to the model's configuration that has obtained the best result in each dataset/modality.

features (i.e., 99% PCA) and AB as shallow classifier, while it obtains its lowest score (0.68) using the Cough modality. Also YAMNET obtains its best result with the breathing recordings, but showing a significant decrease in terms of PR-AUC, obtaining an overall score of 0.66 with LR. VGGish performs slightly better, reaching 0.68 of PR-AUC using RF and 75% as PCA explained variance. On the other hand, OpenL3 obtains the best performances with all the three modalities, but with different configurations: using only cough samples, the best configuration is L3 E 512 L, which reaches 0.73 PR-AUC by using LR and 0.70 of PCA; L3 E 6144 M128 reaches the best score over the whole dataset (i.e., 0.76 PR-AUC) relying on breathing sounds and AB as classifier; while L3 M 6144 M256 scores 0.75 of PR-AUC by using RF and 85% of PCA with the combination of the two previous modalities.

In the experiments with the Coughvid dataset, all the solutions are characterized by a slight decrease in the performance: YAMNET scores 0.56 of PR-AUC by using SVM, thus representing the worst model; VGGish performs better, reaching 0.59 with LR; HC still obtains better performance compared with the previous two deep models, reaching a PR-AUC of 0.61 with SVM; while OpenL3 overcomes the other solutions in all the tested configurations, starting from a PR-AUC of 0.62 obtained by L3 E 512 L in combination with SVM, up to the best result of 0.68 PR-AUC scored by L3 E 6144 M128 with the simple LR as shallow classifier.

On the other hand, with Cambridge KDD all the considered features models reach higher performances with all the three modalities. The lowest value scored by HC is 0.78 by using AB as classifier and coughing records only; while it reaches its maximum result with the breath modality, obtaining 0.85 of PR-AUC and AB. YAMNET still represents the worst model, obtaining the best results with breath sounds: 0.78 of PR-AUC reached by using SVM as shallow classifier. In this case, VGGish performs better than HC and YAMNET, reaching a PR-AUC of 0.89 with cough recordings and the classifier AB. OpenL3 still greatly overcomes the other solutions, obtaining its lowest score with the Breath modality (i.e., 0.92 of PR-AUC) with LR as classifier and three different configurations (i.e., L3 E 6144 M256, L3 M 6144 M128, and L3 M 6144 M256); while it reaches the best performance with the L3 E 6144 M256 configuration and LR as classifier, obtaining an overall PR-AUC of 0.94.

Finally, in the Cambridge npj experiments, HC, YAMNET, and VGGish obtain very similar results with all the three modalities, while OpenL3 still represents the best model. Specifically, HC shows the best performances with coughing samples only, scoring 0.64 of PR-AUC with SVM and 70% of PCA. YAMNET performs very similarly, reaching 0.63 as PR-AUC as maximum value, but using only breathing samples and the combination of cough and breath modalities. With the same combination of respiratory sounds, VGGish slightly overcomes the previous two solutions, scoring a PR-AUC of 0.65 by using SVM. Also in this case, OpenL3 generates the best features set, which allows to reach a PR-AUC score of 0.70 by using the L3 E 512 L configuration, the SVM classifier, and cough recordings only.

Table 3
Features extraction results with Coswara and Coughvid.

D	Features	Cough			Breath			Both		
		Clf	PCA	PR-AUC	Clf	PCA	PR-AUC	Clf	PCA	PR-AUC
Coswara	HC	AB	.90	.68	AB	.99	.74	SVM	.65	.70
	YAMNET	RF	.99	.56	LR	.99	.66	RF	.99	.61
	VGGish	RF	.70	.61	RF	.75	.68	SVM	.65	.63
	L3 E 512 L	LR	.70	.73	RF	.90	.70	AB	.99	.74
	L3 E 512 M128	LR	.99	.71	RF	.75	.70	AB	.95	.73
	L3 E 512 M256	SVM	.85	.68	AB	.95	.71	AB	.85	.72
	L3 E 6144 L	LR	.95	.65	LR	.99	.69	LR	.99	.68
	L3 E 6144 M128	LR	.85	.72	AB	.99	.76	SVM	.75	.73
	L3 E 6144 M256	LR	.75	.72	RF	.85	.74	AB	.99	.71
	L3 M 512 L	RF	.70	.71	RF	.85	.73	SVM	.85	.70
	L3 M 512 M128	RF	.80	.69	AB	.99	.74	RF	.70	.72
	L3 M 512 M256	LR	.60	.70	RF	.90	.71	AB	.99	.69
	L3 M 6144 L	LR	.75	.70	LR	.99	.70	AB	.95	.72
	L3 M 6144 M128	SVM	.75	.69	LR	.99	.74	AB	.99	.72
	L3 M 6144 M256	SVM	.99	.69	AB	.90	.72	RF	.85	.75
Coughvid	HC	SVM	.80	.61	-	-	-	-	-	-
	YAMNET	SVM	.60	.56	-	-	-	-	-	-
	VGGish	LR	.70	.59	-	-	-	-	-	-
	L3 E 512 L	SVM	.75	.62	-	-	-	-	-	-
	L3 E 512 M128	LR	.99	.67	-	-	-	-	-	-
	L3 E 512 M256	LR	.99	.64	-	-	-	-	-	-
	L3 E 6144 L	LR	.99	.65	-	-	-	-	-	-
	L3 E 6144 M128	LR	.99	.68	-	-	-	-	-	-
	L3 E 6144 M256	LR	.99	.67	-	-	-	-	-	-
	L3 M 512 L	RF	.90	.62	-	-	-	-	-	-
	L3 M 512 M128	LR	.99	.62	-	-	-	-	-	-
	L3 M 512 M256	LR	.75	.60	-	-	-	-	-	-
	L3 M 6144 L	LR	.99	.67	-	-	-	-	-	-
	L3 M 6144 M128	LR	.99	.65	-	-	-	-	-	-
	L3 M 6144 M256	LR	.99	.64	-	-	-	-	-	-

4.4. Fine-tuning results

Instead of using the considered deep audio embedding models as simple features extractors, in this second set of experiments we evaluate the use of transfer learning via fine-tuning to create a single model to both represent the input sound through deep embeddings, and then classify it. In other words, here we start with the base models (YAMNET, VGGish, and OpenL3), pre-trained for generic audio classification tasks, and then we fine-tune (train) only the final fully-connected layers of the neural networks to detect COVID-19 from respiratory sounds data.

Table 5 summarizes the performances of the fine-tuned models, highlighting in bold face the best results for each dataset and modality. With Coswara, both YAMNET and VGGish perform better with the Breath modality, obtaining a PR-AUC score of 0.59 and 0.63, respectively, while OpenL3 performs better with coughing samples and the L3 E 6144 M128 configuration, reaching a PR-AUC score of 0.70. The Coughvid dataset is the only one where OpenL3 performs slightly worse than VGGish. In fact, while VGGish achieves a score of 0.59, OpenL3 reaches at most 0.57 with 3 different configurations (L3 E 512 M256, L3 E 6144 M128, and L3 M 512 M128), and YAMNET obtains 0.54 of PR-AUC.

By contrast, with Cambridge KDD all the models obtain higher results: YAMNET is able to reach a maximum score of 0.74 with the combination of both coughing and breathing modalities; VGGish obtains a better result (0.75) by using only breath recordings; and, similarly, OpenL3 achieves a score of 0.85 with the L3 M 512 L configuration, based on breathing samples only. Finally, with the Cambridge npj dataset, all the models perform better by combining the two modalities: YAMNET achieves a PR-AUC score of 0.61, VGGish performs slightly worse, obtaining at most 0.60, while OpenL3 reaches 0.62 with the L3 M 512 M256 configuration.

Based on the obtained results, it is worth noting that the fine-tuned models perform considerably worse than their use as feature extractors to input a shallow classifier. More specifically, compared with the results obtained in the previous set of experiments, the fine-tuned models have lost, on average, 4%, 6.9%, and 8.8% of PR-AUC, respectively for YAMNET, VGGish, and OpenL3. According to the literature [23,44], the relative performance of the fine-tuning and feature extraction approaches strongly depends on the similarity of the pretraining and target tasks. Therefore, we can conclude that the considered respiratory sounds do not allow an effective fine-tuning of the analyzed models because they considerably differ from their original training data, which actually include heterogeneous sounds extracted from YouTube videos, including, for example, music, animal sounds, and traffic noises [45]. A possible solution for this issue could be the fine-tuning not only of the final classification layers but also part of the convolutional components dedicated to the feature extraction. Such an approach could possibly improve the deep representation of the input audio file; however, it typically requires a considerable amount of data that is not currently available in public COVID-19 audio datasets [24].

Table 4
Features extraction results with Cambridge KDD and Cambridge npj.

D	Features	Cough			Breath			Both		
		Clf	PCA	PR-AUC	Clf	PCA	PR-AUC	Clf	PCA	PR-AUC
Cambridge KDD	HC	AB	.60	.78	AB	.80	.85	RF	.60	.83
	YAMNET	LR	.99	.74	SVM	.99	.78	LR	.95	.73
	VGGish	AB	.99	.89	LR	.65	.85	AB	.95	.86
	L3 E 512 L	LR	.70	.86	LR	.95	.87	RF	.90	.88
	L3 E 512 M128	LR	.60	.85	LR	.85	.88	LR	.80	.86
	L3 E 512 M256	RF	.80	.91	SVM	.85	.88	AB	.99	.89
	L3 E 6144 L	RF	.75	.88	RF	.75	.84	SVM	.99	.86
	L3 E 6144 M128	RF	.70	.77	AB	.85	.83	AB	.90	.81
	L3 E 6144 M256	LR	.90	.94	LR	.80	.92	LR	.85	.93
	L3 M 512 L	LR	.70	.82	RF	.80	.89	SVM	.65	.86
	L3 M 512 M128	SVM	.95	.77	SVM	.60	.81	SVM	.99	.80
	L3 M 512 M256	LR	.95	.86	RF	.95	.87	LR	.90	.89
	L3 M 6144 L	LR	.75	.91	SVM	.85	.89	LR	.70	.88
	L3 M 6144 M128	RF	.90	.93	LR	.85	.92	AB	.95	.91
	L3 M 6144 M256	LR	.65	.92	LR	.90	.92	LR	.85	.93
Cambridge npj	HC	SVM	.70	.64	SVM	.99	.62	LR	.70	.63
	YAMNET	SVM	.60	.62	SVM	.99	.63	SVM	.75	.63
	VGGish	SVM	.70	.62	SVM	.99	.65	SVM	.65	.65
	L3 E 512 L	SVM	.75	.70	SVM	.65	.65	SVM	.65	.63
	L3 E 512 M128	SVM	.99	.67	LR	.99	.64	SVM	.90	.67
	L3 E 512 M256	LR	.99	.65	SVM	.99	.67	SVM	.90	.65
	L3 E 6144 L	LR	.95	.65	SVM	.80	.63	LR	.80	.65
	L3 E 6144 M128	LR	.95	.65	SVM	.75	.62	SVM	.85	.66
	L3 E 6144 M256	LR	.95	.66	LR	.85	.64	SVM	.80	.66
	L3 M 512 L	SVM	.95	.65	SVM	.65	.62	SVM	.90	.64
	L3 M 512 M128	LR	.99	.66	LR	.60	.62	SVM	.80	.64
	L3 M 512 M256	SVM	.65	.66	SVM	.90	.66	SVM	.75	.64
	L3 M 6144 L	SVM	.90	.65	SVM	.65	.62	LR	.99	.64
	L3 M 6144 M128	LR	.95	.68	SVM	.85	.66	SVM	.85	.68
	L3 M 6144 M256	LR	.95	.67	SVM	.90	.63	SVM	.75	.66

Table 5
PR-AUC scores obtained in the fine-tuning experiments.

Features	Coswara			Coughvid	Cambridge KDD			Cambridge npj		
	C	B	CB	C	C	B	CB	C	B	CB
YAMNET	.52	.59	.57	.54	.70	.73	.74	.57	.56	.61
VGGish	.57	.63	.62	.59	.69	.75	.74	.57	.58	.60
L3 E 512 L	.68	.62	.65	.54	.68	.72	.66	.57	.55	.57
L3 E 512 M128	.60	.63	.59	.55	.75	.62	.70	.57	.58	.58
L3 E 512 M256	.63	.58	.67	.57	.75	.70	.66	.60	.57	.61
L3 E 6144 L	.60	.69	.59	.53	.74	.68	.73	.56	.54	.55
L3 E 6144 M128	.70	.67	.62	.57	.47	.70	.70	.59	.55	.57
L3 E 6144 M256	.63	.56	.58	.56	.57	.67	.69	.58	.60	.57
L3 M 512 L	.57	.55	.65	.56	.75	.85	.55	.58	.58	.60
L3 M 512 M128	.58	.57	.58	.57	.72	.75	.76	.58	.61	.61
L3 M 512 M256	.56	.67	.64	.54	.71	.76	.75	.57	.61	.62
L3 M 6144 L	.61	.56	.67	.56	.64	.76	.61	.56	.53	.59
L3 M 6144 M128	.58	.54	.60	.56	.75	.64	.73	.60	.58	.61
L3 M 6144 M256	.68	.66	.63	.56	.81	.59	.72	.57	.57	.60

5. Memory footprint

In order to investigate the feasibility of a COVID-19 detection system embedded on commercial mobile devices, we compare the memory footprint of the two Transfer Learning approaches evaluated in Section 4. In fact, while the computational capabilities of modern smartphones are increasingly comparable with the ones of personal computers and they can even execute complex ML models in just a few milliseconds [46], the amount of usable memory still represents a limited resource to be optimized.

When the deep audio embedding models are used as features extractors, two main components must be considered to estimate the overall memory usage: (i) the size of the deep model used to extract the audio embeddings from the raw input file, and (ii) the size of the shallow model used for the classification phase. Based on the number of parameters, the sizes of the deep audio embeddings models are the following: YAMNET is the smallest model, with just 3.7M parameters;




Fig. 4. Memory footprint of (a) the considered shallow classifiers based on different input sizes (i.e., percentage of PCA explained variance), and (b) the fine-tuned deep learning models. Please, consider that in (a) the overall memory footprint is given by taking into account also the size of the deep embedding models to extract the input features from the raw audio sample.

OpenL3 has 4.7M parameters [37]; while VGGish is the biggest one, with a total of 62M parameters [32], approximately requiring 16 MB, 288 MB, and 18 MB for YAMNET, VGGish, and OpenL3, respectively.

Moreover, the memory footprint of the classification phase is represented by the size of the employed shallow classifier, which is generally affected by the input data dimension. Fig. 4(a) shows the memory size (in MB) of the 4 shallow classifiers considered in our experiments, based on the input size in terms of the percentage of PCA explained variance. As expected, the size of most of the classifiers increases according to the input dimension, except for Random Forest (RF), whose size remains nearly constant at around 1 MB (between 0.98 and 1.11). Logistic Regression (LR), the simplest classifier, is also the one with the lowest memory footprint in all the experiments, starting from less than 1 KB (i.e., 922 Bytes), up to 3.64 KB with 99% of PCA explained variance. On the other hand, AdaBoost (AB) results to be the most demanding model in terms of memory, with an overall size that ranges from just 5.74 MB with 60% of PCA, up to 185.05 MB with the full dimension of the input. Finally, SVM has an intermediate memory footprint among the other classifiers, ranging from 42.6 KB up to 1.88 MB.

On the other hand, when the deep audio models are fine-tuned, the size of the additional fully-connected layers should be considered to estimate the overall memory footprint. Fig. 4(b) shows the average size of the fine-tuned models, highlighting both the size of the original pre-trained models, and the size of the additional layers for classification. We can note that, in general, fine-tuning introduces a non-negligible overhead, due to both the number and density of the classification layers considered in our experiments. More specifically, the overhead introduced by the new layers for YAMNET and VGGish is 5.8M and 5.3M parameters, respectively, which corresponds to an average increment of 33.7 MB for the former model, and 7.37 MB for the latter. On the other hand, the overhead introduced for OpenL3 strongly depends on the size of the deep audio embeddings: for the configurations that generate a total of 512 embeddings, the fine-tuning introduces, on average, 5.6M parameters (+21.8 MB); while the configurations with 6144 embeddings have been augmented by 14.8M parameters, which corresponds, on average, to an overhead of 56.73 MB with respect to the base model.

The obtained results clearly show that in both the Transfer Learning approaches, the most demanding component is represented by the pre-trained deep learning model, especially in the case of VGGish, which is the larger model among those considered in our analysis. Even though the general memory footprint of the considered solutions is manageable on modern smartphones, they can still be prohibitive for mobile devices with less memory resources, such as smartwatches and smart bands.

Therefore, in order to develop novel pervasive m-health applications that can be used as low-cost pre-screening tool for COVID-19, processing the user's data directly on her personal device, the deep learning model should be firstly optimized to be executed on resource-constrained devices. In this regards, several techniques can be employed to optimize different aspects of deep learning models, including the overall model size [47,48]. For example, quantization is a practical and broadly used technique aimed at reducing the size of a pre-trained model by simply lowering the operations precision from 32 bit floats to 16 bit floats or even 8 bit integers [49], thus reducing the overall model's size by 4x at least, with little degradation in model accuracy, thus enabling its execution on devices with limited computational and memory capabilities.

6. Conclusions and future work

In this paper, we investigated the use of different Artificial Intelligence approaches to automatically detect COVID-19 from respiratory sound signals, aimed at developing an effective and low-cost solution to identify new cases, performing the whole data processing directly on the users' mobile devices. Specifically, we compared the performances of the two most common approaches for audio classifications: the use of handcrafted acoustic features, and pre-trained deep learning models.

In particular, we considered the pre-trained instances of 3 state-of-the-art deep learning models, namely VGGish, YAMNET, and OpenL3, evaluating their efficacy under two transfer learning settings: features extraction and fine-tuning. In the first case, the deep models are used as simple feature extractors, characterizing respiratory sounds as deep audio embeddings, which are further classified by shallow machine learning classifiers, such as Logistic Regression, and Support Vector Machines. On the other hand, in the latter approach, the final fully-connected layers of the deep learning models are fine-tuned to adapt the model to the new application domain, thus creating a single model for both data modeling and classification.

Through an extensive evaluation, we compared the performances of the different approaches by using four public datasets and 5-fold nested cross-validation with user-based splits, to avoid biasing the experiments based on specific users or specific train/validation/test splits. The obtained results clearly show the great advantage of OpenL3 in all the experimental settings, improving the other solutions by 12.3% in terms of Precision-Recall AUC as features extractor, and by 10% when the model is fine-tuned. Moreover, we noted that the fine-tuned models generally perform worse than the feature extractors, with an average drop of 6.6% of the performances. Besides, we evaluated the feasibility of the considered solutions of being entirely executed on mobile devices, especially considering their memory constraints. As imagined, the component with the greatest impact on the overall memory footprint is represented by the deep learning models, especially in the fine-tuning experiments, where additional fully-connected layers are introduced to adapt the model to the new application domain, requiring, on average, between 37 MB and about 300 MB for the most demanding model (i.e., VGGish).

Our future work is threefold. Firstly, we would like to investigate different techniques to improve the fine-tuning of the deep learning models. In fact, while we focused on updating only the classification part of the models (i.e., the final fully-connected layers), another possible approach would imply also the partial update of the features extraction component of the model (i.e., the convolutional layers), thus adapting its input representation capabilities to the novel application scenario [24]. Even though fine-tuning of a larger portion of the model will surely require a non-negligible amount of COVID-19 data, we can still use the combination of the available datasets during the draining process as data augmentation approach.

Secondly, we are planning to develop a prototype mobile application to evaluate the performance of the considered approaches for COVID-19 detection in real-world scenarios. To this aim, we need to optimize the models in order to be efficiently executed on memory-constrained devices by testing different optimization approaches, including, for example, *quantization*, which reduces the precision of the mathematical operations performed by the neural network, and *pruning*, where redundant connections among hidden units are removed during the training phase [50].

Finally, we would like to explore Explainable Artificial Intelligence (XAI) techniques aimed at supporting the predictions produced by the deep learning models. In particular, since the audio classification models are generally based on Convolutional Neural Networks, we can exploit the same XAI techniques proposed in the literature for visual applications [51] to describe, for example, which set of frequencies have led the model to classify a given respiratory sound as COVID-positive instead of negative sample, thus providing an additional validation of the classification model. This is a critical requirement in our reference scenario, since it has been proven that the ability of explaining the output of an AI model increases the faith held in the system, thus leading to a widespread adoption.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors shared the link to access both data and code in the manuscript.

Acknowledgments

The authors express their gratitude to Prof. Cecilia Mascolo, Dept. of Computer Science and Technology, University of Cambridge (UK) for sharing the databases of COVID19 sound App described in [5,31]. This work was supported partially by the project EcoCyber included in the Spoke 08 - "Risk Management & Governance" of the Research and Innovation Program PE00000014, "Security and Rights in the CyberSpace (SERICS)", and partially by the project Active Ageing@THE included in the Spoke 3 - "Advanced Technologies, Methods and Materials for Human Health and Wellbeing" of the Research and Innovation Program ECS_00000017, "THE - Tuscany Health Ecosystem", both under the National Recovery and Resilience Plan, Mission 4 "Education and Research" - Component 2 "From Research to Enterprise" - Investments 1.3 and 1.5, funded by the European Union - NextGenerationEU.

References

- [1] J. Chen, K. Li, Z. Zhang, K. Li, P.S. Yu, A survey on applications of artificial intelligence in fighting against COVID-19, *ACM Comput. Surv.* 54 (8) (2021) <http://dx.doi.org/10.1145/3465398>.
- [2] Q.-V. Pham, D.C. Nguyen, T. Huynh-The, W.-J. Hwang, P.N. Pathirana, Artificial intelligence (AI) and big data for coronavirus (COVID-19) pandemic: A survey on the state-of-the-arts, *IEEE Access* 8 (2020) 130820–130839, <http://dx.doi.org/10.1109/ACCESS.2020.3009328>.
- [3] J. Wu, P. Zhang, L. Zhang, W. Meng, J. Li, C. Tong, Y. Li, J. Cai, Z. Yang, J. Zhu, M. Zhao, H. Huang, X. Xie, S. Li, Rapid and accurate identification of COVID-19 infection through machine learning based on clinical available blood test results, 2020, <http://dx.doi.org/10.1101/2020.04.02.20051136>, medRxiv. <https://www.medrxiv.org/content/early/2020/04/06/2020.04.02.20051136.full.pdf>. URL <https://www.medrxiv.org/content/early/2020/04/06/2020.04.02.20051136>.
- [4] O. Gozes, M. Frid-Adar, H. Greenspan, P.D. Browning, H. Zhang, W. Ji, A. Bernheim, E. Siegel, Rapid AI development cycle for the coronavirus (COVID-19) pandemic: Initial results for automated detection & patient monitoring using deep learning CT image analysis, 2020, [arXiv:2003.05037](https://arxiv.org/abs/2003.05037).
- [5] C. Brown, J. Chauhan, A. Grammenos, J. Han, A. Hasthanasombat, D. Spathis, T. Xia, P. Cicuta, C. Mascolo, Exploring automatic diagnosis of COVID-19 from crowdsourced respiratory sound data, in: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ACM, 2020, <http://dx.doi.org/10.1145/3394486.3412865>.
- [6] B.W. Schuller, D.M. Schuller, K. Qian, J. Liu, H. Zheng, X. Li, COVID-19 and computer audition: An overview on what speech & sound analysis could contribute in the SARS-CoV-2 corona crisis, *Front. Digit. Health* 3 (2021) <http://dx.doi.org/10.3389/fgth.2021.564906>, URL <https://www.frontiersin.org/article/10.3389/fgth.2021.564906>.
- [7] K. Qian, X. Li, H. Li, S. Li, W. Li, Z. Ning, S. Yu, L. Hou, G. Tang, J. Lu, F. Li, S. Duan, C. Du, Y. Cheng, Y. Wang, L. Gan, Y. Yamamoto, B.W. Schuller, Computer audition for healthcare: Opportunities and challenges, *Front. Digit. Health* 2 (2020) <http://dx.doi.org/10.3389/fgth.2020.00005>, URL <https://www.frontiersin.org/articles/10.3389/fgth.2020.00005>.
- [8] A. Imran, I. Posokhova, H.N. Qureshi, U. Masood, M.S. Riaz, K. Ali, C.N. John, M.I. Hussain, M. Nabeel, AI4COVID-19: AI enabled preliminary diagnosis for COVID-19 from cough samples via an app, *Inform. Med. Unlocked* 20 (2020) 100378, <http://dx.doi.org/10.1016/j.imu.2020.100378>, URL <https://www.sciencedirect.com/science/article/pii/S2352914820303026>.
- [9] J. Han, C. Brown, J. Chauhan, A. Grammenos, A. Hasthanasombat, D. Spathis, T. Xia, P. Cicuta, C. Mascolo, Exploring automatic COVID-19 diagnosis via voice and symptoms from crowdsourced data, in: *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing*, ICASSP, 2021, pp. 8328–8332, <http://dx.doi.org/10.1109/ICASSP39728.2021.9414576>.
- [10] N. Sharma, P. Krishnan, R. Kumar, S. Ramoji, S.R. Chetupalli, N. R., P.K. Ghosh, S. Ganapathy, Coswara – A database of breathing, cough, and voice sounds for COVID-19 diagnosis, in: *Interspeech 2020*, ISCA, 2020, <http://dx.doi.org/10.21437/Interspeech.2020-2768>.
- [11] B. Subirana, F. Hueto, P. Rajasekaran, J. Laguarda, S. Puig, J. Malvehy, O. Mitja, A. Trilla, C.I. Moreno, J.F.M. Valle, A.E.M. González, B. Vizmanos, S. Sarma, Hi sigma, do I have the coronavirus?: Call for a new artificial intelligence approach to support health care professionals dealing with the COVID-19 pandemic, 2020, [arXiv:2004.06510](https://arxiv.org/abs/2004.06510).
- [12] J. Han, K. Qian, M. Song, Z. Yang, Z. Ren, S. Liu, J. Liu, H. Zheng, W. Ji, T. Koike, et al., An early study on intelligent analysis of speech under COVID-19: Severity, sleep quality, fatigue, and anxiety, 2020, [arXiv preprint arXiv:2005.00096](https://arxiv.org/abs/2005.00096).
- [13] J. Laguarda, F. Hueto, B. Subirana, COVID-19 artificial intelligence diagnosis using only cough recordings, *IEEE Open J. Eng. Med. Biol.* 1 (2020) 275–281, <http://dx.doi.org/10.1109/OJEMB.2020.3026928>.
- [14] E.A. Mohammed, M. Keyhani, A. Sanati-Nezhad, S.H. Hejazi, B.H. Far, An ensemble learning approach to digital corona virus preliminary screening from cough sounds, *Sci. Rep.* 11 (1) (2021) 1–11.
- [15] M.G. Campana, A. Rovati, F. Delmastro, E. Pagani, L3-net deep audio embeddings to improve COVID-19 detection from smartphone data, in: *2022 IEEE International Conference on Smart Computing*, SMARTCOMP, 2022, pp. 100–107, <http://dx.doi.org/10.1109/SMARTCOMP55677.2022.00029>.
- [16] R. Arandjelovic, A. Zisserman, Look, listen and learn, in: *Proceedings of the IEEE International Conference on Computer Vision*, ICCV, 2017.
- [17] S. Sigtia, A. Stark, S. Krstulović, M. Plumbley, Automatic environmental sound recognition: Performance versus computational cost, *IEEE/ACM Trans. Audio Speech Lang. Process.* 24 (11) (2016) 2096–2107, <http://dx.doi.org/10.1109/TASLP.2016.2592698>.
- [18] Y. Bengio, Deep learning of representations for unsupervised and transfer learning, in: I. Guyon, G. Dror, V. Lemaire, G. Taylor, D. Silver (Eds.), *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, in: *Proceedings of Machine Learning Research*, vol. 27, PMLR, Bellevue, Washington, USA, 2012, pp. 17–36, URL <https://proceedings.mlr.press/v27/bengio12a.html>.
- [19] E. Tsalera, A. Papadakis, M. Samarakou, Comparison of pre-trained CNNs for audio classification using transfer learning, *J. Sens. Actuat. Netw.* 10 (4) (2021) 72.
- [20] D. Ellis, YAMNet: A pretrained audio event classifier, 2019, <https://github.com/tensorflow/models/tree/master/research/audioset/yamnet>.
- [21] J. Yosinski, J. Clune, Y. Bengio, H. Lipson, How transferable are features in deep neural networks? *Adv. Neural Inf. Process. Syst.* 27 (2014).
- [22] A. Zhang, Z. Lipton, M. Li, A. Smola, Dive into deep learning, 2005, <http://dx.doi.org/10.48550/arXiv.2106.11342>, URL <https://d2l.ai>.
- [23] M. Peters, S. Ruder, N. Smith, To tune or not to tune? Adapting pretrained representations to diverse tasks, in: *Proc. 4th Workshop on Representation Learning for NLP (Repl4NLP)*, 2019, pp. 7–14, <http://dx.doi.org/10.18653/v1/W19-4302>.
- [24] G. Vrbancić, V. Podgorelec, Transfer learning with adaptive fine-tuning, *IEEE Access* 8 (2020) 196197–196211, <http://dx.doi.org/10.1109/ACCESS.2020.3034343>.
- [25] Y. Guo, H. Shi, A. Kumar, K. Grauman, T. Rosing, R. Feris, SpotTune: Transfer learning through adaptive fine-tuning, in: *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, CVPR, 2019, pp. 4800–4809, <http://dx.doi.org/10.1109/CVPR.2019.00494>.
- [26] Y. Kim, Y. Hyon, S. Jung, S. Lee, G. Yoo, C. Chung, T. Ha, Respiratory sound classification for crackles, wheezes, and rhonchi in the clinical field using deep learning, *Sci. Rep.* 11 (2021) <http://dx.doi.org/10.1038/s41598-021-96724-7>.
- [27] Z. Chen, H. Wang, C.-H. Yeh, X. Liu, Classify respiratory abnormality in lung sounds using STFT and a fine-tuned ResNet18 network, in: *Proc. IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 2022, <http://dx.doi.org/10.1109/BioCAS54905.2022.9948614>.
- [28] F. Demir, A. Sengur, V. Bajaj, Convolutional neural networks based efficient approach for classification of lung diseases, *Health Inf. Sci. Syst.* 8 (4) (2020) <http://dx.doi.org/10.1007/s13755-019-0091-3>.
- [29] G. Deshpande, A. Batliner, B.W. Schuller, AI-based human audio processing for COVID-19: A comprehensive overview, *Pattern Recognit.* 122 (2022) 108289, <http://dx.doi.org/10.1016/j.patcog.2021.108289>, URL <https://www.sciencedirect.com/science/article/pii/S0031320321004696>.
- [30] E.A. Mohammed, M. Keyhani, A. Sanati-Nezhad, S.H. Hejazi, B.H. Far, An ensemble learning approach to digital corona virus preliminary screening from cough sounds, *Sci. Rep.* 11 (1) (2021) 15404, <http://dx.doi.org/10.1038/s41598-021-95042-2>.
- [31] J. Han, T. Xia, D. Spathis, E. Bondareva, C. Brown, J. Chauhan, T. Dang, A. Grammenos, A. Hasthanasombat, A. Floto, P. Cicuta, C. Mascolo, Sounds of COVID-19: exploring realistic performance of audio-based digital testing, *Npj Digit. Med.* 5 (1) (2022) 16, <http://dx.doi.org/10.1038/s41746-021-00553-x>.
- [32] S. Hershey, S. Chaudhuri, D.P.W. Ellis, J.F. Gemmeke, A. Jansen, R.C. Moore, M. Plakal, D. Platt, R.A. Saurous, B. Seybold, M. Slaney, R.J. Weiss, K. Wilson, CNN architectures for large-scale audio classification, in: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing*, ICASSP, 2017, pp. 131–135, <http://dx.doi.org/10.1109/ICASSP.2017.7952132>.

- [33] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, Q. He, A comprehensive survey on transfer learning, *Proc. IEEE* 109 (1) (2021) 43–76, <http://dx.doi.org/10.1109/JPROC.2020.3004555>.
- [34] E. Tsalera, A. Papadakis, M. Samarakou, Comparison of pre-trained CNNs for audio classification using transfer learning, *J. Sens. Actuat. Netw.* 10 (4) (2021) <http://dx.doi.org/10.3390/jsan10040072>, URL <https://www.mdpi.com/2224-2708/10/4/72>.
- [35] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2014, arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
- [36] A.G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, H. Adam, Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017, arXiv preprint [arXiv:1704.04861](https://arxiv.org/abs/1704.04861).
- [37] J. Cramer, H.-H. Wu, J. Salamon, J.P. Bello, Look, listen, and learn more: Design choices for deep audio embeddings, in: ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, 2019, pp. 3852–3856, <http://dx.doi.org/10.1109/ICASSP.2019.8682475>.
- [38] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, S. Vijayanarasimhan, Youtube-8m: A large-scale video classification benchmark, 2016, arXiv preprint [arXiv:1609.08675](https://arxiv.org/abs/1609.08675).
- [39] L. Orlandic, T. Teijeiro, D. Atienza, The COUGHVID crowdsourcing dataset, a corpus for the study of large-scale cough analysis algorithms, *Sci. Data* 8 (1) (2021) 156, <http://dx.doi.org/10.1038/s41597-021-00937-4>.
- [40] W. Zheng, M. Jin, The effects of class imbalance and training data size on classifier learning: An empirical study, *SN Comput. Sci.* 1 (2) (2020) 71, <http://dx.doi.org/10.1007/s42979-020-0074-0>.
- [41] J. Snoek, H. Larochelle, R.P. Adams, Practical Bayesian optimization of machine learning algorithms, in: F. Pereira, C. Burges, L. Bottou, K. Weinberger (Eds.), *Advances in Neural Information Processing Systems*, Vol. 25, Curran Associates, Inc., 2012, URL <https://proceedings.neurips.cc/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf>.
- [42] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, A. Talwalkar, Hyperband: A novel bandit-based approach to hyperparameter optimization, *J. Mach. Learn. Res.* 18 (1) (2017) 6765–6816.
- [43] P. Branco, L. Torgo, R.P. Ribeiro, A survey of predictive modeling on imbalanced domains, *ACM Comput. Surv.* 49 (2) (2016) <http://dx.doi.org/10.1145/2907070>.
- [44] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P.J. Liu, et al., Exploring the limits of transfer learning with a unified text-to-text transformer, *J. Mach. Learn. Res.* 21 (140) (2020) 1–67.
- [45] J.F. Gemmeke, D.P.W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R.C. Moore, M. Plakal, M. Ritter, Audio set: An ontology and human-labeled dataset for audio events, in: *Proc. IEEE ICASSP 2017, New Orleans, LA, 2017*.
- [46] S. Yao, Y. Zhao, H. Shao, S. Liu, D. Liu, L. Su, T. Abdelzaher, FastDeepIoT: Towards understanding and optimizing neural network execution time on mobile and embedded devices, in: *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems, SenSys '18*, Association for Computing Machinery, New York, NY, USA, 2018, pp. 278–291, <http://dx.doi.org/10.1145/3274783.3274840>.
- [47] S. Han, J. Pool, J. Tran, W. Dally, Learning both weights and connections for efficient neural network, in: C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Vol. 28, Curran Associates, Inc., 2015, URL <https://proceedings.neurips.cc/paper/2015/file/ae0eb3eed39d2bcef4622b2499a05fe6-Paper.pdf>.
- [48] S. Han, H. Mao, W.J. Dally, Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding, 2015, arXiv preprint [arXiv:1510.00149](https://arxiv.org/abs/1510.00149).
- [49] A. Kwasniewska, M. Szankin, M. Ozga, J. Wolfe, A. Das, A. Zajac, J. Ruminski, P. Rad, Deep learning optimization for edge devices: Analysis of training quantization parameters, in: *IECON 2019 - 45th Annual Conference of the IEEE Industrial Electronics Society*, Vol. 1, 2019, pp. 96–101, <http://dx.doi.org/10.1109/IECON.2019.8927153>.
- [50] T. Liang, J. Glossner, L. Wang, S. Shi, X. Zhang, Pruning and quantization for deep neural network acceleration: A survey, *Neurocomputing* 461 (2021) 370–403, <http://dx.doi.org/10.1016/j.neucom.2021.07.045>, URL <https://www.sciencedirect.com/science/article/pii/S0925231221010894>.
- [51] B.H. van der Velden, H.J. Kuijff, K.G. Gilhuijs, M.A. Viergeever, Explainable artificial intelligence (XAI) in deep learning-based medical image analysis, *Med. Image Anal.* 79 (2022) 102470, <http://dx.doi.org/10.1016/j.media.2022.102470>, URL <https://www.sciencedirect.com/science/article/pii/S1361841522001177>.