L³-Net Deep Audio Embeddings to Improve COVID-19 Detection from Smartphone Data

Mattia Giovanni Campana*, Andrea Rovati†, Franca Delmastro*, and Elena Pagani†

*Institute for Informatics and Telematics of the National Research Council of Italy (IIT-CNR), Pisa, Italy
†Computer Science Department, University of Milano, Milan, Italy
Email: {m.campana, f.delmastro}@iit.cnr.it, andrea.rovati1@gmail.com, elena.pagani@unimi.it

Abstract—Smartphones and wearable devices, along with Artificial Intelligence, can represent a game-changer in the pandemic control, by implementing low-cost and pervasive solutions to recognize the development of new diseases at their early stages and by potentially avoiding the rise of new outbreaks. Some recent works show promise in detecting diagnostic signals of COVID-19 from voice and coughs by using machine learning and hand-crafted acoustic features. In this paper, we decided to investigate the capabilities of the recently proposed deep embedding model L³-Net to automatically extract meaningful features from raw respiratory audio recordings in order to improve the performances of standard machine learning classifiers in discriminating between COVID-19 positive and negative subjects from smartphone data. We evaluated the proposed model on 3 datasets, comparing the obtained results with those of two reference works. Results show that the combination of L³-Net with hand-crafted features overcomes the performance of the other works of 28.57% in terms of AUC in a set of subject-independent experiments. This result paves the way to further investigation on different deep audio embeddings, also for the automatic detection of different diseases.

Index Terms—Audio embeddings, Deep Learning, m-health, COVID-19

I. INTRODUCTION

COVID-19 pandemic has highlighted the limitations of national healthcare systems in containing the spread of a virus at a large scale. Until effective vaccines were available, countries struggled for more than a year in flattening the pandemic curve by testing the population and isolating infected people, causing, as a side effect, an economical crisis that affected the whole society [1]. Researchers from all over the world have proposed diverse digital solutions to mitigate the pandemic and study its diffusion, most of them characterized by a massive use of Artificial Intelligence (AI) technologies and big data [2], [3]. For example, Machine Learning (ML) classifiers have been successfully employed to identify COVID-19 cases from blood tests [4], while Deep Learning (DL) models achieved incredibly high performance (i.e., 99.6% accuracy [5]) in analyzing chest X-ray and lung Computed Tomography (CT) images, thus supporting medical personnel in rapidly diagnosing positive subjects and providing appropriate medical treatments. AI-based solutions have been also proposed to deal with other aspects of the pandemic, including: estimation of patient mortality and survival rate based on medical annotation, demographic and physiological data [6], [7]; extraction of COVID-19 symptoms from unstruc-

tured data by exploiting Natural Language Processing (NLP) techniques [8]; DL-based video tracking to detect suspicious COVID-19 patients in public places [9].

Another aspect of the pandemic that has been recently investigated is the definition of scalable and low-cost digital solutions for fast screening, aimed at recognizing the onset of new cases and possibly preventing new outbreaks. Specifically, smartphones and mobile health systems (m-health) can represent pervasive instruments for the early detection of COVID-19 by exploiting embedded sensors, with particular attention to microphones and generated audio signals, considering that COVID-19 is a respiratory illness characterized by specific dysfunctions in respiratory physiology, affecting patterns of breathing, speech, and coughing [10].

Schuller et al. [11] firstly investigated 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) [12]. Subsequently, researchers investigated the effective applicability of those techniques in real scenarios. Initial studies focused on small patients' cohorts trying to automatically distinguish between COVID-19 cough and cough sounds related to other pathologies [13]. However, this requires a huge amount of data that could not be collected rapidly. Therefore, [13] presents 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 [14]–[16]. As a first analysis, 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*) [17]. Then, DL-based approaches have been proposed [18], [19], including the use of deep audio embeddings to enrich standard CA features [14].

In this paper, we investigate the feasibility of using the recently proposed *Learn, Listen and Learn* (L³-Net) [20] embedding model to improve the detection of COVID-19. Specifically, we employ 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. In addition, we combine deep embeddings with hand-

crafted acoustic features already recognized in literature so as to further enhance the system performances.

To evaluate the proposed solution, we directly compare it with two reference works: [19] and [14]. We perform a series of subject-independent experiments by using the same reference datasets and we demonstrate that L^3-Net overcomes the reference works in terms of different standard metrics: 28.57% AUC, 23.75% Precision, and 39.43% Recall. Moreover, since we would like to investigate the real feasibility of the proposed solution as a m-health system component, we provide a preliminary evaluation of the complexity of the proposed approach by taking into account the typical memory constraints of personal mobile devices. Specifically, we compare different ML classifiers and DL-based feature extraction models in order to identify the best trade-off between classification performances and model's size.

II. RELATED WORK

In the last couple of years, during the pandemic, researchers have explored several audio processing techniques, already known in the CA field, to develop effective and low-cost COVID-19 screening methods based on respiratory data [21], especially derived from smartphone embedded microphones.

We can classify the proposed methods in 3 main approaches. First, the use of speech and audio analysis to extract hand-crafted features that characterize different aspects of the acoustic signal for classification purpose. This includes, for example, basic frequency-based and temporal features [22], [23], but also sets of features especially designed for voice and paralinguistic applications (e.g., GeMAPS [24] and COM-PARE [25], [26]), which have been successfully employed to detect different diseases in the past, including tuberculo-sis [27], asthma [28], and Parkinson [29]. Alsabek et al. [30] are among the first who studied the relevance of using the Mel-Frequency Cepstral Coefficients (MFCCs) features to detect COVID-19 from both cough and breathing sounds, while Han et al. [14] used both basic features and COMPARE set to detect COVID-19 from voice samples. Moreover, Han et al. [17] compared the use of GeMAPS and COMPARE to analyze speech recordings from COVID-19 patients to categorize their health status from four aspects, including severity of illness, sleep quality, fatigue, and anxiety.

The main drawback of these techniques is that designed features might not be optimal for the classification objective, and they are typically outperformed by DL models [31]. In order to overcome this issue, a second approach has been investigated, consisting 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 Convolutional Neural Network (CNN) model for both features extraction and classification. This category includes, for example, the application AI4COVID proposed by Imran et al. [13] based only on cough recordings. Specifically, they modelled the audio sample as both Mel-spectrogram and MFCC, which are then processed by an ensemble model composed of two CNNs and one Support Vector Machine (SVM) to categorize

the cough into 4 classes: COVID-19, bronchitis, pertussis, and normal cough. A similar solution has been proposed by Mohammed et al. [32], 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 modelling 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.

The third approach, which we can consider as hybrid, 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 [33], in which the authors 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 [34], a CNN-based embedding model trained on the large-scale YouTube-8M dataset (approximately 2.6 billion audio/video features), thus taking advantage of Transfer Learning concept to deal with the shortage of COVID-19 audio data [35].

Given its simplicity and effectiveness, we consider the third approach as the most suitable to implement an early detection system for COVID-19 on mobile devices. For this reason, in this work, we propose an enhancement of the solution presented in [33], investigating the use of the more recent L^3-Net model to extract deep audio embeddings from respiratory sound recordings. Compared with VGGish, L^3-Net processes not only audio data, but also video streams, and it has been designed especially to model the correspondence between the two. In this way, it is able to extract a meaningful set of embeddings, which have been proven to outperform other embedding models in several audio classification tasks [36], [37]. To the best of our knowledge, this is the first attempt of using such a model for the early detection of COVID-19.

III. ACOUSTIC FEATURES AND L^3-NET DEEP AUDIO EMBEDDINGS

In this section, we present the high-level architecture we propose to improve COVID-19 detection from smartphone data. Specifically, Figure 1 shows the flow diagram of the entire data process, that can be summarized in 6 main steps: (i) the audio sample is firstly collected through the device microphone; (ii) we extract several hand-crafted acoustic features already proposed in the CA literature for similar tasks and considered the main standard features; (iii) concurrently, we use L^3-Net deep model to extract deep audio embeddings from the raw audio sample; (iv) acoustic features and deep embeddings are then combined in a single features vector,
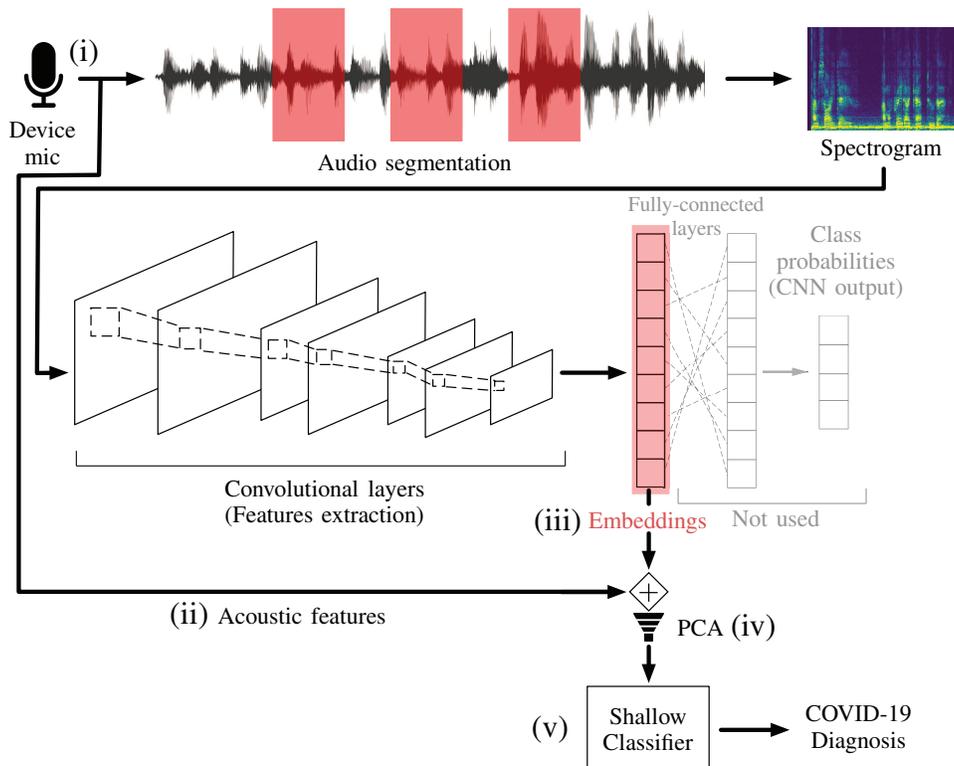
Fig. 1: Flow diagram of the proposed system.

which is further reduced by using Principal Component Analysis (PCA); (v) eventually, the user's audio is classified as potentially positive or negative COVID-19 example by using a shallow ML classifier, such as SVM or LR.

A. Acoustic features extraction

To transform the raw audio sample into a numerical representation manageable by a ML classifier, we use acoustic features and the L^3-Net embeddings, both independently and integrated. In terms of acoustic features, we implement the common approach used in similar audio-based medical applications [12]. Firstly, the audio sample recorded by the user's device microphone is re-sampled to a standard value for audio tasks (e.g., 16kHz or 22kHz). 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, the complete list of acoustic features we consider in this work is presented in Table I and it is the same already used in [14].

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.

B. L^3-Net for deep audio embeddings

Among the hand-crafted features, we use L^3-Net to extract deep latent features from the raw file. As we mentioned in Section II, this model has been designed to learn embeddings by identifying if a video image frame and an audio segment come from the same video. This allows to train the model in a self-supervised way: 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.

L^3-Net architecture consists of two distinct CNN sub-networks to extract different embeddings for the video and audio inputs, respectively. To check the correspondence between both embeddings, a fusion network is used. It concatenates both embeddings and uses two fully connected layers as well as a softmax layer for binary classification. As far as the audio embeddings is concerned, L^3-Net extracts a 512-dimensional features vector from overlapping windows with 1-second length and a 0.1 hop size of Mel-spectrograms images generated with 256 Mel bins. We take the mean and standard deviation of each dimension across all the windows to characterize the entire audio segment as a 1024-dimensional features vector (i.e., 512×2).

As depicted in Figure 1 (step (iii)), we use this model as feature extractor. In other words, we discard the fully-connected layers and final output of the deep model, and keep only the features extraction part: the CNN sub-network that processes the audio and its corresponding embeddings layer.

TABLE I: Hand-crafted acoustic features.

Feature	Description
Duration	total length (in seconds) of the audio sample, after removing possible starting and ending silence
Onset	number of pitch onsets (i.e., "events") in the audio signal
Tempo	rate of beats that occur at regular intervals throughout the entire audio signal
Period	the frequency with the highest amplitude among those obtained from the Fast Fourier Transform (FFT)
RMS Energy	root-mean-square 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. It can be used to identify percussive and sustained sounds [22]
Roll-off Frequency	the frequency under which the 85% of the total energy of the frame-wise spectrum is contained [22]
Zero-crossing rate	the number of times the signal value crosses the zero axe, and it is computed for each frame
MFCC	the shape of the cosine transformation of the sound logarithmic spectrum, expressed in Mel-bands [38]
Δ-MFCC and Δ^2-MFCC	the first and second order derivatives of MFCC along time

Moreover, to cope with the shortage of COVID-19 respiratory sound data, we rely on the OpenL3 [37] model, which has been trained on approximately 2 millions videos contained in the AudioSet dataset [39]. In this way, we follow the Transfer Learning approach, by exploiting the training of L^3-Net on a massive amount of data in a different application domain to take advantage of its ability to characterize audio samples.

C. Features combination and classification

As a final step, we combine the acoustic features and deep audio embeddings (Figure 1, step (iv)), thus obtaining a single representation of the original audio sample composed by a total of 1501 features.

As discussed in Section II, due to the moderate size of the available audio-based COVID-19 datasets (a few thousands of samples in the best case), in order to predict the user's COVID-19 condition we rely on shallow ML classifiers (Figure 1, step (v)), which have been proven to provide excellent results in similar applications, even with a limited amount of training data. Preliminarily, in order to avoid the well-known curse of dimensionality problem that can affect the performance of several classifiers, we use PCA to reduce the dimension of the input samples and to remove possible noisy or redundant features.

IV. EXPERIMENTAL EVALUATION

In order to evaluate the effectiveness of L^3-Net to automatically extract effective latent features for COVID-19 detection, we perform two main sets of experiments by using 3 datasets: COSWARA [15] and Virufy [40] are publicly available, while we obtained the access to the Cambridge dataset [33] through a data transfer agreement between CNR and Cambridge University for research purposes. We then compare the obtained classification performances with the other solutions presented in the literature and detailed in Section II.

In addition, since we are interested in the real implementation of this model in m-health platforms, we provide a preliminary evaluation of the complexity of the proposed approach by comparing different combinations of features sets and shallow classifiers in terms of memory usage, considering the limited resources of personal mobile devices. This allows us also to identify the best candidate solution for the development of a prototype application on real mobile devices.
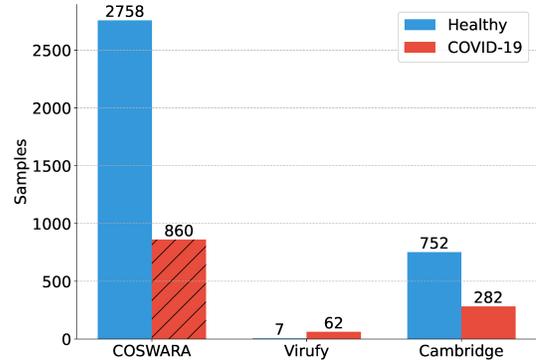
Fig. 2: Number of audio samples in the 3 considered datasets, grouped by their respective labels.

A. Datasets

Figure 2 shows the main peculiarities of the three datasets, highlighting the number of audio samples obtained by negative (*Healthy*) and positive (*COVID-19*) subjects.

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 asked 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 [1]. Similarly to [19], in this work we take into account only cough sounds, whose 2758 have been shared by people who have declared they were healthy and 860 are labelled as COVID-19 positive examples.

Virufy is a no-profit corporation developing AI technology to detect COVID-19 from cough patterns. They publicly released a dataset, collected by 69 voluntary subjects who were visiting an Indian hospital for COVID-19 test [2]. Even

---

[1]https://github.com/iiscleap/Coswara-Data
[2]https://github.com/virufy/virufy-cdf-india-clinical-1

though the number of samples in this dataset is limited (i.e., 69 audio samples, one per person), the labels with which they have been tagged are very accurate because they are based on COVID-19 PCR test results obtained by qualified personnel of the hospital. The total number of samples obtained by healthy subjects is only 7, while the number of COVID-19 cough samples is 62. As we detail in Section IV-B, we use this dataset in combination with COSWARA to compare our proposal with a reference solution based on cough sound recordings [19].

The Cambridge dataset [33] has been collected by the Mobile System Research Lab of the University of Cambridge as part of the ERC EAR research project, which aims at exploiting microphones of mobile devices to collect human body sounds as indicators of disease or disease onsets. Similarly to COSWARA, Cambridge contains respiratory sounds crowdsourced by using both web and mobile applications. It is composed by a total of 1034 audio samples donated by 356 people, who also self-reported their health status related to COVID-19. The dataset is divided in different groups, based on the users' medical condition: positive subjects with/without cough, healthy subjects without any symptoms, healthy subjects with cough, and asthmatic people with/without cough. In Figure 2, we summarize the dataset characteristics, considering as *COVID-19* the 282 samples related to people who have tested positive to the virus (with or without cough), while the other 752 samples are considered as *Healthy*.

B. Evaluation protocol and metrics

In order to compare our proposal with the state-of-the-art, we consider the following works as reference baselines: (i) [33] based on the combination of acoustic features and audio embeddings produced by VGGish model applied to Cambridge dataset; and (ii) [19] based on the ensemble of CNN evaluated by combining COSWARA and Virufy in one single dataset of cough audio samples.

For a fair comparison, we reproduce as much as possible the experiments performed by the reference works. On the one hand, for comparison with [19], we perform a standard binary classification task, i.e., we simply distinguish between positive and negative subjects based on the cough audio samples contained in both COSWARA and Virufy. On the other hand, the comparison with [33] is based on the three different classification tasks defined in the baseline paper and that we detail in the following: **Task 1** (*COVID-positive vs COVID-negative*): distinguishing between people who have declared they tested positive for COVID-19 (*COVID-positive*) and users who have not declared a positive test for COVID-19 with a clean medical history, without symptoms, no smoking, and living where COVID-19 was not prevalent at the recording time; **Task 2** (*COVID-positive with cough vs COVID-negative*): similar to the previous task, but in this case we consider as *COVID-positive* the people who tested positive *and* declared cough as a symptom; and **Task 3** (*COVID-positive with cough vs COVID-negative with asthma and cough*): distinguishing between people who have declared they tested positive for COVID-19

TABLE II: Grid search parameters.

Algorithm	Parameter	Values
SVM	regularization	$[10^{-3}, \ldots, 10^3]$
	kernel	[rbf, poly, sigmoid]
	kernel coefficient	$[10^{-3}, \ldots, 10]$
	degree of poly kernel	$[2, \ldots, 5]$
AB	estimators	[10, 20, 50, 100]
	learning rate	[1, .5, .1, .05, .01, .001]
LR	penalty	[l1, l2]
	regularization	$[10^{-3}, \ldots, 10^3]$
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	[.7, .8, .9. .95, .99]

and reported cough as symptom, and negative subjects with asthma and cough.

Moreover, to avoid bias in the experiments based on patterns of specific users, we adopt the Leave-One-Subject-Out (LOSO) approach, thus ensuring that samples from the same user do not appear in both training and test splits. Specifically, we use a nested cross-validation-like approach as follows. Firstly, in an outer loop, we randomly shuffle the entire dataset for 10 times, based on the users. Then, after each shuffle, we keep 80% of the users as developing set and 20% as test set, and we ensure that the classes in both the sets are always balanced by randomly undersampling the majority class.

The development set is then used in an inner 5-fold cross validation for hyperparameters tuning. This include: (i) selection of the best features to combine with the deep audio embeddings; (ii) finding the best PCA coefficient, that is, the amount of variance that needs to be explained by the held components; and (iii) finding the best ML classifier and fine-tuning its parameters. In these experiments, we test 4 broadly used ML classification algorithms: SVMs, LR, Random Forest (RF), and AdaBoost (AB); and we tune their hyperparameters by performing an exaustive search through grid-search with the parameters value spaces specified in Table II.

As far as the features selection is concerned, we followed the approach used in [33], by testing the following sets of features: (**F1**) deep audio embeddings only; (**F2**) embeddings with *Period*, *Tempo*, and *Duration*; (**F3**) embeddings with all the acoustic features, except Δ-*MFCC* and Δ^2-*MFCC*; and (**F4**) embeddings with all the hand-crafted features. In addition, for the experiments with the Cambridge dataset, we also evaluate which type of audio files (i.e., Modality) allows us to obtain the best performance among those available in the dataset: *Cough*, *Breath*, or the combination of the two.

Finally, we calculate the average classification performances over the outer 10 splits by using 3 standard metrics: Area Under the ROC Curve (*AUC*), which provides an aggregate measure of performance across all possible classification thresholds; *Precision*, which measures the ability of the classifier not to misclassify positive examples; and *Recall* (also

known as *Sensitivity*), which indicates the ability of a classifier to correctly label all the positive samples in the test set.

C. Results

Table III summarizes the classification performances of the proposed solution compared with the reference baselines, high-lighting the best configurations and results (in terms of mean and standard deviation) obtained through the nested cross-validation. Specifically, for the Cambridge dataset, we report for each task the configuration of the best baseline and related metrics, to be compared both with the results we obtained by using the *same* setup, and with our *best* configuration. By contrast, for the experiments with COSWARA+VIRUFY, we use as baseline reference the best configuration reported in [19], that is, the combination of the top 4 audio representations found in their evaluation: Spectrogram, Mel-spectrogram, Power-spectrogram, and MFCC.

In the first set of experiments, we can note that our solution based on L^3-Net is able to usually obtain better results than the baseline, but with different configurations. In the first task, the embeddings generated by L^3-Net allows to obtain the same AUC score and a higher true-positive rate (i.e., $+6,94\%$ in terms of Precision) by using the same Modality (*Cough+Breath*) and features set (*F2*) as the baseline, but by using SVM as shallow classifier instead of LR and less PCA components. In Task 2, our solution shows a higher false-negative rate (i.e., -16.67% in terms of Recall), but overcomes the baseline for both AUC ($+2,4\%$) and Precision ($+15\%$), thus correctly detecting COVID-19 subjects 92% of the time. Surprisingly, the L^3-Net embeddings extracted from audio samples of *Breath* seem more effective than using the *Cough* recordings, making the latter less relevant to distinguish between COVID-positive with cough and COVID-negative subjects in this dataset. Finally, in the last task, our proposal is far better than the baseline in distinguishing between COVID-positive subjects with cough from COVID-negative subjects with asthma and cough, overcoming the reference solution for all the considered metrics: $+10\%$ AUC, $+18.84\%$ Precision, and $+14.49\%$ in terms of Recall.

While the experiments with the Cambridge dataset show the advantage of using L^3-Net over VGGish for COVID-19 detection, the test performed with the COSWARA+Virufy dataset clearly demonstrate the effectiveness of Transfer Learning in our scenario. Our proposal obtains perfect classification performances, considerably overcoming the baseline in all the three evaluation metrics: $+28.57\%$ AUC, $+23.75\%$ Precision, and $+39.43\%$ Recall. This is surely due to the amount of data points contained in the dataset, which enable the shallow classifier to correctly capture the intrinsic patterns among the samples. Moreover, it further motivates our choice of using a pre-trained DL model instead of training it from scratch: using the knowledge learnt during the training with millions of data samples, OpenL3 is able to better characterize the audio data, even though they refer to a different context than the ones used during the training. By contrast, training end-to-end a complex DL model as the one proposed in [19] requires a

considerable amount of annotated data [41], which typically far more exceeds the number of samples contained in the considered datasets.

V. 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 different ML classifiers considered in the conducted experiments so as to find the best trade-off between classification performances and model size.

Figure 3 shows the memory size (in MB) and classification accuracy (AUC) of the 4 shallow classifiers presented in Section IV, taking into account the best Modality and Features sets. Since the input dimension can greatly affect the models sizes, we also show their differences among the considered values of PCA coefficients.

According to the results, we can note that LR, the simplest classifier, is also the one with the lower footprint in all the experiments (i.e., 0.1 MB at most), but it can achieve the best AUC score in two settings (i.e., Cambridge Task 2 and COSWARA+Virufy). On the other hand, AB and RF are generally the most demanding models in terms of memory (up to 11 MB for AB, and approximately 3.05 MB for RF). However, AB obtains the best result in the Cambridge Task 3 experiment, requiring a limited memory by using 0.70 as PCA coefficient (i.e., 0.06 MB). Finally, SVM generally has an average memory footprint compared with the other classifiers, ranging from 0.01 and 0.1 MB (except for the last experiment), and it scores the best result in Cambridge Task 1 with PCA 0.7, requiring only 0.05 MB for an AUC score of 0.80.

The obtained results clearly show that all the considered ML classifiers are viable for being installed on mobile devices, with a low impact on the general memory usage.

VI. Conclusions and future work

In this paper, we investigate the use of the recent embedding model L^3-Net to train shallow classifiers aimed at identifying COVID-19 subjects from cough and breathing audio samples. L^3-Net demonstrated to outperform several Deep Learning solutions in other audio classification tasks, and it can further improve the classification performances in this specific task. To deal with the shortage of public COVID-19 audio data, we employed OpenL3, an instance of L^3-Net pre-trained on approximately 2 millions videos. In this way, applying the Transfer Learning paradigm, we exploited the training of L^3-Net on a massive amount of data, thus taking advantage of its ability to effectively characterize audio data and improve the detection of COVID-19 from respiratory sound samples.

Through an extensive evaluation, employing three public datasets, we evaluated the effectiveness of L^3-Net to automatically extract latent features for COVID-19 detection, comparing its performance with two baseline approaches: the original VGGish-based proposal, and an ensemble of four Convolutional Neural Networks trained from scratch. The obtained results clearly show the great advantage of our proposal over the other solutions, achieving a gain of 10% AUC

TABLE III: Classification results

Dataset	Task	Method	Modality	Features	Classifier	PCA	Mean (± std)		
							AUC	Precision	Recall
Cambridge	1	baseline	Cough + Breath	F2	LR	.95	.80 (.07)	.72 (.06)	**.69 (.11)**
		our (same)	Cough + Breath	F2	LR	.95	.76 (.092)	.69 (.095)	.68 (.158)
		our (best)	Cough + Breath	F2	SVM	.70	.80 (.068)	**.77 (.096)**	.68 (.139)
	2	baseline	Cough	F2	SVM	.90	.82 (.18)	.80 (.16)	**.72 (.23)**
		our (same)	Cough	F2	SVM	.90	.69 (.227)	.74 (.187)	.61 (.276)
		our (best)	Breath	F1	LR	.80	**.84 (.168)**	**.92 (.106)**	.60 (.237)
	3	baseline	Breath	F3	SVM	.70	.80 (.14)	.69 (.20)	.69 (.26)
		our (same)	Breath	F3	SVM	.70	.64 (.254)	.69 (.154)	.66 (.269)
		our (best)	Breath	F1	AB	.70	**.88 (.066)**	**.82 (.152)**	**.79 (.192)**
COSWARA + Virufy		baseline		Top 4	Ensemble CNN	-	.77	.80	.71
		our		F3	LR	.99	**.99 (.001)**	**.99 (.006)**	**.99 (.007)**

(a) Cambridge Task 1 (b) Cambridge Task 2 (c) Cambridge Task 3 (d) COSWARA+Virufy
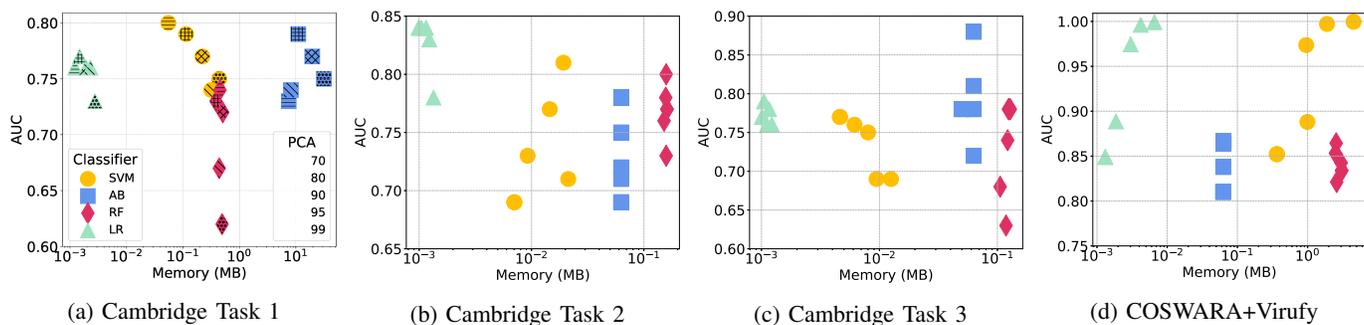
Fig. 3: Memory footprint of the considered shallow classifiers for each value of PCA coefficient in the different experiments.

compared with the former baseline, and 28.57% AUC with respect to the latter. In addition, we also performed a series of experiments to evaluate the trade-off between the classification accuracy and the memory occupancy of 4 shallow classifiers, based on different input size. Support Vector Machines and Logistic Regression performed the best, obtaining a high level of accuracy and, at the same time, requiring only few KB of memory, representing the best candidates to be deployed on mobile devices.

As a future work, we would like to make an extensive comparison of different deep audio embeddings models for COVID-19 detection and, if other datasets are available, for the automatic detection of other important diseases, like Parkinson or post-stroke, in which audio and speech analysis can provide fundamental diagnostic information. Finally, from the algorithmic point of view, we would like to combine different public datasets for fine-tuning OpenL3 on COVID-19 respiratory data, defining a single model combining features extraction and classification tasks.

REFERENCES

[1] N. Akbulaev, I. Mammadov, and V. Aliyev, "Economic impact of covid-19," *Sylwan*, vol. 164, no. 5, 2020.

[2] J. Chen, K. Li, Z. Zhang, K. Li, and P. S. Yu, "A survey on applications of artificial intelligence in fighting against covid-19," *ACM Comput. Surv.*, vol. 54, no. 8, oct 2021. [Online]. Available: https://doi.org/10.1145/3465398

[3] Q.-V. Pham, D. C. Nguyen, T. Huynh-The, W.-J. Hwang, and P. N. Pathirana, "Artificial intelligence (ai) and big data for coronavirus (covid-19) pandemic: A survey on the state-of-the-arts," *IEEE Access*, vol. 8, pp. 130 820–130 839, 2020.

[4] 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, and S. Li, "Rapid and accurate identification of covid-19 infection through machine learning based on clinical available blood test results," *medRxiv*, 2020. [Online]. Available: https://www.medrxiv.org/content/early/2020/04/06/2020.04.02.20051136

[5] O. Gozes, M. Frid-Adar, H. Greenspan, P. D. Browning, H. Zhang, W. Ji, A. Bernheim, and 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.

[6] M. Pourhomayoun and M. Shakibi, "Predicting mortality risk in patients with covid-19 using machine learning to help medical decision-making," *Smart Health*, vol. 20, p. 100178, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352648320300702

[7] B. Xu, B. Gutierrez, S. Mekaru, K. Sewalk, L. Goodwin, A. Loskill, E. L. Cohn, Y. Hswen, S. C. Hill, M. M. Cobo, A. E. Zarebski, S. Li, C.-H. Wu, E. Hulland, J. D. Morgan, L. Wang, K. O'Brien, S. Scarpino, J. S. Brownstein, O. G. Pybus, D. M. Pigott, and M. U. G. Kraemer, "Epidemiological data from the covid-19 outbreak, real-time case information," *Scientific Data*, vol. 7, no. 1, p. 106, Mar 2020. [Online]. Available: https://doi.org/10.1038/s41597-020-0448-0

[8] G. M. Silverman, H. S. Sahoo, N. E. Ingraham, M. Lupei, M. A. Puskarich, M. Usher, J. Dries, R. L. Finzel, E. Murray, J. Sartori *et al.*, "Nlp methods for extraction of symptoms from unstructured data for use

in prognostic covid-19 analytic models," *Journal of Artificial Intelligence Research*, vol. 72, pp. 429–474, 2021.

[9] J. Chen, K. Li, Q. Deng, K. Li, and P. S. Yu, "Distributed deep learning model for intelligent video surveillance systems with edge computing," *IEEE Transactions on Industrial Informatics*, pp. 1–1, 2019.

[10] T. F. Quatieri, T. Talkar, and J. S. Palmer, "A framework for biomarkers of covid-19 based on coordination of speech-production subsystems," *IEEE Open Journal of Engineering in Medicine and Biology*, vol. 1, pp. 203–206, 2020.

[11] 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," *Frontiers in Digital Health*, vol. 3, 2021. [Online]. Available: https://www.frontiersin.org/article/10.3389/fdgth.2021.564906

[12] 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, and B. W. Schuller, "Computer audition for healthcare: Opportunities and challenges," *Frontiers in Digital Health*, vol. 2, 2020. [Online]. Available: https://www.frontiersin.org/article/10.3389/fdgth.2020.00005

[13] A. Imran, I. Posokhova, H. N. Qureshi, U. Masood, M. S. Riaz, K. Ali, C. N. John, M. I. Hussain, and M. Nabeel, "Ai4covid-19: Ai enabled preliminary diagnosis for covid-19 from cough samples via an app," *Informatics in Medicine Unlocked*, vol. 20, p. 100378, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352914820303026

[14] J. Han, C. Brown, J. Chauhan, A. Grammenos, A. Hasthanasombat, D. Spathis, T. Xia, P. Cicuta, and 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.

[15] N. Sharma, P. Krishnan, R. Kumar, S. Ramoji, S. R. Chetupalli, N. R., P. K. Ghosh, and S. Ganapathy, "Coswara — a database of breathing, cough, and voice sounds for covid-19 diagnosis," *Interspeech 2020*, Oct 2020. [Online]. Available: http://dx.doi.org/10.21437/Interspeech.2020-2768

[16] B. Subirana, F. Hueto, P. Rajasekaran, J. Laguarta, S. Puig, J. Malvehy, O. Mitja, A. Trilla, C. I. Moreno, J. F. M. Valle, A. E. M. González, B. Vizmanos, and 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.

[17] 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," *arXiv preprint arXiv:2005.00096*, 2020.

[18] J. Laguarta, F. Hueto, and B. Subirana, "Covid-19 artificial intelligence diagnosis using only cough recordings," *IEEE Open Journal of Engineering in Medicine and Biology*, vol. 1, pp. 275–281, 2020.

[19] E. A. Mohammed, M. Keyhani, A. Sanati-Nezhad, S. H. Hejazi, and B. H. Far, "An ensemble learning approach to digital corona virus preliminary screening from cough sounds," *Scientific Reports*, vol. 11, no. 1, pp. 1–11, 2021.

[20] R. Arandjelovic and A. Zisserman, "Look, listen and learn," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.

[21] G. Deshpande, A. Batliner, and B. W. Schuller, "Ai-based human audio processing for covid-19: A comprehensive overview," *Pattern Recognition*, vol. 122, p. 108289, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0031320321004696

[22] G. Peeters, "A large set of audio features for sound description (similarity and classification) in the cuidado project," *CUIDADO Ist Project Report*, vol. 54, no. 0, pp. 1–25, 2004.

[23] T. Drugman, J. Urbain, and T. Dutoit, "Assessment of audio features for automatic cough detection," in *2011 19th European Signal Processing Conference*, 2011, pp. 1289–1293.

[24] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan, and K. P. Truong, "The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing," *IEEE Transactions on Affective Computing*, vol. 7, no. 2, pp. 190–202, 2016.

[25] B. Schuller, S. Steidl, A. Batliner, A. Vinciarelli, K. Scherer, F. Ringeval, M. Chetouani, F. Weninger, F. Eyben, E. Marchi *et al.*, "The interspeech 2013 computational paralinguistics challenge: Social signals, conflict, emotion, autism," in *Proceedings INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France*, 2013.

[26] B. W. Schuller, A. Batliner, C. Bergler, F. B. Pokorny, J. Krajewski, M. Cychosz, R. Vollmann, S.-D. Roelen, S. Schnieder, E. Bergelson, A. Cristia, A. Seidl, A. S. Warlaumont, L. Yankowitz, E. Nöth, S. Amiriparian, S. Hantke, and M. Schmitt, "The INTERSPEECH 2019 Computational Paralinguistics Challenge: Styrian Dialects, Continuous Sleepiness, Baby Sounds & Orca Activity," in *Proc. Interspeech 2019*, 2019, pp. 2378–2382.

[27] I. D. S. Miranda, A. H. Diacon, and T. R. Niesler, "A comparative study of features for acoustic cough detection using deep architectures*," in *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2019, pp. 2601–2605.

[28] S. Yadav, M. Keerthana, D. Gope, U. Maheswari K., and P. Kumar Ghosh, "Analysis of acoustic features for speech sound based classification of asthmatic and healthy subjects," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6789–6793.

[29] J. S. Almeida, P. P. Rebouças Filho, T. Carneiro, W. Wei, R. Damaševičius, R. Maskeliūnas, and V. H. C. de Albuquerque, "Detecting parkinson's disease with sustained phonation and speech signals using machine learning techniques," *Pattern Recognition Letters*, vol. 125, pp. 55–62, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167865519301163

[30] M. B. Alsabek, I. Shahin, and A. Hassan, "Studying the similarity of covid-19 sounds based on correlation analysis of mfcc," in *2020 International Conference on Communications, Computing, Cybersecurity, and Informatics (CCCI)*, 2020, pp. 1–5.

[31] H. Purwins, B. Li, T. Virtanen, J. Schlüter, S.-Y. Chang, and T. Sainath, "Deep learning for audio signal processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 206–219, 2019.

[32] E. A. Mohammed, M. Keyhani, A. Sanati-Nezhad, S. H. Hejazi, and B. H. Far, "An ensemble learning approach to digital corona virus preliminary screening from cough sounds," *Scientific Reports*, vol. 11, no. 1, p. 15404, Jul 2021. [Online]. Available: https://doi.org/10.1038/s41598-021-95042-2

[33] 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," *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, Aug 2020. [Online]. Available: http://dx.doi.org/10.1145/3394486.3412865

[34] 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, and 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.

[35] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2021.

[36] S. Grollmisch, E. Cano, C. Kehling, and M. Taenzer, "Analyzing the potential of pre-trained embeddings for audio classification tasks," in *2020 28th European Signal Processing Conference (EUSIPCO)*, 2021, pp. 790–794.

[37] J. Cramer, H.-H. Wu, J. Salamon, and 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.

[38] B. J. Mohan and R. B. N., "Speech recognition using mfcc and dtw," in *2014 International Conference on Advances in Electrical Engineering (ICAEE)*, 2014, pp. 1–4.

[39] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 776–780.

[40] G. Chaudhari, X. Jiang, A. Fakhry, A. Han, J. Xiao, S. Shen, and A. Khanzada, "Virufy: Global applicability of crowdsourced and clinical datasets for ai detection of covid-19 from cough," *arXiv preprint arXiv:2011.13320*, 2020.

[41] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015. [Online]. Available: https://doi.org/10.1038/nature14539