



Audio Feature Ranking for Sound-Based COVID-19 Patient Detection

Julia A. Meister¹ , Khuong An Nguyen¹, and Zhiyuan Luo²

¹ University of Brighton, Brighton, East Sussex BN2 4GJ, UK
{J.Meister,K.A.Nguyen}@brighton.ac.uk

² Royal Holloway University of London, Egham, Surrey TW20 0EX, UK
Zhiyuan.Luo@rhul.ac.uk

Abstract. Audio classification using breath and cough samples has recently emerged as a low-cost, non-invasive, and accessible COVID-19 screening method. However, a comprehensive survey shows that no application has been approved for official use at the time of writing, due to the stringent reliability and accuracy requirements of the critical healthcare setting. To support the development of Machine Learning classification models, we performed an extensive comparative investigation and ranking of 15 audio features, including less well-known ones. The results were verified on two independent COVID-19 sound datasets. By using the identified top-performing features, we have increased COVID-19 classification accuracy by up to 17% on the Cambridge dataset and up to 10% on the Coswara dataset compared to the original baseline accuracies without our feature ranking.

Keywords: COVID-19 classification · Audio event engineering · Sound feature ranking

1 Introduction

A widely accessible, non-invasive, low-cost testing mechanism is the number one priority to support test-and-trace in most pandemics. The advent of COVID-19 has abruptly brought respiratory audio classification into the spotlight as a viable alternative for mass pre-screening, needing only a smartphone to record a breath or cough sample [3].

It has long been common knowledge that respiratory diseases physically alter the respiratory environment in a way that often induces audible changes [17]. Consequently, manually listening to lung sounds (auscultation) is a common method to identify and diagnose respiratory disorders. However, many abnormalities only subtly affect auditory cues, making the inherently subjective auscultation process error-prone even when performed by a trained medical professional [2]. To counteract subjectivity, automated audio classification approaches with promising results have become more and more common in recent years [1, 2, 9].

One of the main limiting factors is the lack of ground truth data which may be difficult to obtain, prone to limited population diversity, and requires medical training to label correctly. Because COVID-19 detection is a widespread and

critical problem, multiple universities and research institutions have published COVID-19 audio datasets [3, 19]. This offers a unique opportunity to verify classification solutions on independently collected samples from a diverse population. The datasets have supported the development of a variety of applications with Machine Learning (ML) audio classification. However, at the time of writing, none have yet been officially endorsed for medical usage, largely because of the high accuracy and reliability expectations for such a critical healthcare task.

The paper gives a comprehensive overview of relevant audio features (Sect. 2) and identifies the most indicative ones for COVID-19 (Sect. 3). Finally, the findings are put into the context of existing literature (Sect. 4).

1.1 The Paper’s Contributions

The rigorous feature analysis presented in this paper improves COVID-19 respiratory classification by optimising and holistically evaluating audio signal representations for Machine Learning (ML). The following contributions are made:

- *Audio feature analysis and ranking.* We performed an extensive comparative analysis and ranking of 15 sound features prevalent and less-common in audio classification. The evaluation was carried out on two independent datasets, allowing the findings to be generalised.
- *Highlighted effective features.* We identified sound-based ML features with strong discriminative performance that go against common rules of thumb.
- *Increased COVID-19 detection accuracy.* We improved accuracy up to 17% by incorporating new training features based on our feature ranking.

2 Audio Features Overview

As in any Machine Learning (ML) application, feature engineering is a vital step for COVID-19 cough classification. We provide a detailed overview of 15 audio features from a variety of signal domains (Table 1) before rigorously evaluating their performance.

2.1 Time Domain

Low-level features extracted directly from the signal are in the time domain. They may identify crackling sounds caused by secretions in the throat and lungs [17], and have been previously used for COVID-19 classification [3, 19].

Root Mean Square Energy (RMSE). A measure of the signal’s amplitude over N frames, see Eq. (1). x_n is the average energy per frame [15].

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{n=1}^N x_n^2} \quad (1)$$

Zero-Crossing Rate (ZCR). The signal’s sign change rate (Eq. (2)). x_n is amplitude at frame n of N . $\text{sign}(a)$ returns 1 if $a > 0$, 0 if $a = 0$, and -1 else [15].

$$\text{ZCR} = \frac{1}{2} \times \sum_{n=2}^N |\text{sign}(x_n) - \text{sign}(x_{n-1})| \quad (2)$$

Table 1. *Audio feature selection.* The 15 audio features evaluated in the paper.

Domain	Feature category	Name	Intuition
Time	Signal energy	RMSE	Loudness of the signal
	Waveform	ZCR	Percussive vs tonal
Frequency	Spectral	S-BW	Perceived timbre
	Spectral	S-CENT	‘Brightness’ of a sound
	Spectral	S-CONT	Prevalence of formants
	Spectral	S-FLAT	Similarity to white noise
	Spectral	S-FLUX	Rate of frequency changes
	Spectral	S-ROLL	‘Skewness’ of the energy
Time-frequency	Cepstral	MFCC	Timbre, tone colour/quality
	Cepstral	MFCC- Δ	Velocity of temporal change
	Cepstral	MFCC- Δ^2	Acceleration of temporal change
	Tonal	C-ENS	Pitch
	Tonal	C-CQT	Pitch
	Tonal	C-STFT	Pitch
	Tonal	TN	Pitch & pitch height

2.2 Frequency Domain

To reveal frequency information of digital audio, it is decomposed into its constituent frequencies. This domain may identify abnormal lung sounds caused by an infection by examining the signal’s intensity [17]. A subset has previously been used for COVID-19 detection [3, 19].

Spectral Bandwidth. Equation (3) shows energy concentration, i.e. variance of expected frequency E given energy P_k and frequency f_k in $1 \leq k \leq K$ bands [16].

$$\text{S-BW} = \sqrt{\sum_{k=1}^K (f_k - E)^2 \times P_k} \quad (3)$$

Spectral Centroid. Equation (4) shows the weighted and unweighted sums of spectral magnitudes P_k in the k -th of K subbands. f_k is the corresponding frequency [20].

$$\text{S-CENT} = \frac{\sum_{k=1}^K P_k \times f_k}{\sum_{k=1}^K P_k} \quad (4)$$

Spectral Contrast. Compare spectral peaks P_k and valleys V_k in frequency band k , see Eq. (5). N is the number of frames and $x'_{k,n}$ the FFT vector [7].

$$\text{S-CONT}_k = P_k - V_k = \left(\log \frac{1}{N} \sum_{n=1}^N x'_{k,n} \right) - \left(\log \frac{1}{N} \sum_{n=1}^N x'_{k,N-n+1} \right) \quad (5)$$

Spectral Flatness. Eq. (6) measures similarity to white noise. P_k is the signal’s energy at the k -th frequency band s.t. $1 \leq k \leq K$ [10].

$$\text{S-FLAT} = \frac{(\prod_{k=1}^K P_k)^{\frac{1}{K}}}{\frac{1}{K} \sum_{k=1}^K P_k} \quad (6)$$

Spectral Flux. Equation (7) measures a signal’s energy change between frames. $E_{n,k}$ is the k -th of K Discrete Fourier Transform coefficients in frame n [20].

$$\text{S-FLUX}_n = \sum_{k=1}^K E_{n,k} - E_{n-1,k}^2 \quad (7)$$

Spectral Rolloff. Equation (8) finds frequency f_R s.t. the energy accumulated below is no less than proportion S of total energy. P_k is energy in one of K bands [20].

$$\text{S-ROLL} = \arg \min f_R \in \{1, \dots, K\} \sum_{k=1}^{f_R} P_k \geq S \sum_{k=1}^K P_k \quad (8)$$

2.3 Time-Frequency Domain

This domain shows a signal’s frequency as it varies over time. We consider two types of features: cepstral (timbre or tone colour) and tonal (pitch).

Cepstral Features. Non-linear Mel-frequency Cepstrum (MFC) is ubiquitous in respiratory classification because it explores a signal’s temporal frequency content. It has been previously used for COVID-19 [3, 12].

Mel-Frequency Cepstral Coefficients. Equation (9) shows the signal’s transformation. $s(k)$ is the log energy of the k -th of K coefficients at frame n [3].

$$\text{MFCC}_n = \sum_{k=1}^K s(k) \cos \frac{\pi n(k-0.5)}{K} \quad (9)$$

MFCC- Δ . The first-order derivative of MFCC, velocity, represents temporal change and is often included due to its low extraction cost [4].

MFCC- Δ^2 . The second-order derivative, acceleration, is commonly included because it may improve audio classification [4].

Tonal Features. Based on the human perception of periodic pitch [13]. Two types are considered: chromagram and lattice graph. Secretions are a common consequence of COVID-19 which may alter the pitch of in- and expiration [17].

Chroma Energy Normalised. Chroma abstraction considering short-time statistics within chroma bands. Normalisation makes C-ENS resistant to timbre [13].

Constant-Q Chromagram. Extracted from a time-frequency representation. The constant-Q transform (C-CQT) has a good resolution of low frequencies [8].

Short-Time Fourier Transform Chromagram. The difference to C-CQT is the initial transformation, in this case the Short-time Fourier Transform (STFT) [8].

Tonnetz. A lattice graph of harmonic information. Distances between points become meaningful by encoding pitch as geometric areas [5].

Table 2. *Sample counts of the datasets.* Each coswara participant has ‘shallow’ and ‘deep’ breath (B), cough (C), and breathcough (BC) recordings.

Label	Cambridge			Coswara-deep			Cos.-shallow		
	B	C	BC	B	C	BC	B	C	BC
COVID-19	111	111	111	81	81	81	81	81	81
Healthy	194	194	194	1074	1074	1074	1074	1074	1074
Σ	305	305	305	1155	1155	1155	1155	1155	1155

3 Experimental Method and Results

The 15 investigated features range from prevalent to traditionally excluded from audio classification. They were ranked based on the empirical results analysis of two independent datasets. We assume that patterns repeated across both datasets are likely inherent to the COVID-19 respiratory recordings.

3.1 Research Questions

Three research questions were formulated to inform the experimental design and results analysis. Each is focused on improving COVID-19 audio classification.

- What are the most predictive audio features for Machine Learning?
- Are the feature rankings comparable across independent datasets?
- How much does the performance accuracy of Machine Learning models improve by using the most dominant features?

3.2 The Datasets

Two parallel independent datasets were considered throughout the paper to indicate whether feature rankings were likely generally applicable: the *Cambridge* and *Coswara* COVID-19 audio datasets. The sample counts are shown in Table 2.

Introduced in [3], the *Cambridge dataset* is a collection of healthy and COVID-positive cough and breath recordings. The data we used is a curated set of 48kHz WAV file samples, collected April-May 2020. Additionally, the Indian Institute of Science has collected shallow and deep breath and cough recordings in the *Coswara dataset* [19]. Compatible samples from April-December 2020 were considered. For consistency, we filtered for COVID-positive and healthy participants.

3.3 Feature Engineering

Cleaning the audio data was especially important because the recording devices and environments were not controlled. The pre-processing steps were carried out with the Python-toolkit *librosa*, and included trimming the leading/trailing silences and normalising the amplitude to $(-1, 1)$.




Fig. 1. *Sample lengths pre- and post-processing.* We trim leading and trailing silences (60 dB, empirically identified). Lengths were reduced by 1–3 s.

We evaluated 15 audio features from three signal domains (Sect. 2). To standardise feature dimensions for Machine Learning (ML) models regardless of sample length (1–30 s, Fig. 1), seven summary statistics were calculated to describe the feature distribution across frames: (i) minimum, (ii) maximum, (iii) mean, (iv) median, (v) variance, (vi) 1st quartile, and (vii) 3rd quartile. Only a small subset of features was considered for evaluation and ranking at a time to avoid overfitting (812 features total, Table 3).

Table 3. *Feature dimensions.* 812 features were considered. 7 Summary statistics were taken across frames to ensure consistent dimensions (sample length 1–30 s). To reduce overfitting risk, feature subsets were considered at a time for ranking.


Dimension	Features (min, max, mean, median, var, Q_1 , and Q_3)
(1×7)	RMSE, ZCR, S-BW, S-CENT, S-FLAT, S-FLUX, S-ROLL
(6×7)	TN
(7×7)	S-CONT
(12×7)	C-ENS, C-CQT, C-STFT
(20×7)	MFCC, MFCC- Δ , MFCC- Δ^2

3.4 Results Description and Analysis


We identified the most informative features by evaluating two datasets in parallel. We propose that recurring predictive patterns are likely independent of the dataset, and should be strongly considered for future ML COVID-19 classification applications. Features were analysed in the following configurations:

- The Cambridge, Coswara-deep, and Coswara-shallow datasets.
- Breath (B), Cough (C), and BreathCough (BC) feature vectors. The latter is a concatenation of the previous two feature vectors, i.e. double the size.
- 5 models, selected for the variety in which they partition the label space: AdaBoost-Random Forest (ADA), K-Nearest Neighbours (KNN), Logistic Regression (LR), Random Forest (RF), and Support Vector Machine (SVM).

5-fold Cross-Validation ensured reliable results. We selected 3 metrics to compare the features’ efficiency: *Receiver Operating Characteristic* (ROC), *Precision*



(a) Mean ROC over 5-fold CV (positive: COVID). AUC is ‘Area Under Curve’.



(b) Mean PR over 5-fold CV (positive: COVID). AP is ‘Average Precision’.

Fig. 2. *BreathCough* results. Even though the ROC-curves look similar across datasets, the PR-curves reveal that Cambridge performed better overall. We also identified SVM and RF as the top-performing models. In PR-curves, the unskilled classifier corresponds to the dataset’s positive label ratio.

(P), and *Recall* (R). PR curves are well suited to imbalanced data by omitting true negatives, counteracting ROC’s optimism [18]. The mean over folds was a suitable indicator because the performance values passed the normality test [6].

Feature Categories. An overview of full feature vectors showed promising results, as most models outperformed their no-skill equivalent in ROC and PR-curves (Fig. 2). SVM and RF outperformed their counterparts across BC, B and C. Even though the two datasets had similarly shaped ROC curves, Cambridge had the best Average Precision (AP), and illustrates ROC’s optimism on imbalanced datasets. An influential factor in Coswara’s lower overall accuracies was the greater imbalance of COVID samples (Coswara 13:1 vs Cambridge 2:1, Table 2). Nonetheless, Coswara-trained models performed significantly better than their unskilled classifier counterparts (13–38% vs 7% AP, Fig. 2b).

BC signal domain results confirmed SVM and RF as the best performing models (Table 4). Considering SVM’s BC ROC-AUC across all datasets, we note that the 4 feature categories were broadly ranked in increasing predictive efficiency (Cambridge, Coswara-deep, Coswara-shallow): *time domain* (79%, 64%, 56%), *tonal* (83%, 73%, 69%), *spectral* (85%, 74%, 72%), and *cepstral* (87%,

Table 4. *BreathCough* 5-fold CV ROC-AUC results as mean(std). SVM and RF achieved the highest accuracies across most domains. The feature categories were ranked in the following increasing order: time, tonal, spectral, cepstral.

Dataset	Category	ADA	KNN	LR	RF	SVM
Cambridge	Time	67.17(.04)	77.96(.07)	76.01(.07)	78.21(.05)	78.78(.07)
	Spectral	87.09(.04)	85.34(.05)	84.17(.06)	87.15(.05)	84.84(.07)
	Cepstral	83.84(.05)	85.56(.07)	83.27(.06)	87.82(.07)	87.15(.06)
	Tonal	84.74(.09)	81.04(.05)	81.44(.04)	81.11(.07)	82.59(.07)
Coswara-deep	Time	55.65(.07)	62.34(.02)	54.21(.09)	64.65(.05)	63.94(.07)
	Spectral	65.77(.07)	68.18(.04)	72.03(.05)	71.76(.06)	74.46(.06)
	Cepstral	70.83(.06)	71.03(.03)	75.01(.05)	77.55(.06)	75.62(.08)
	Tonal	69.29(.06)	66.27(.02)	68.02(.03)	72.32(.06)	72.98(.03)
Coswara-shallow	Time	61.63(.04)	55.05(.06)	56.16(.09)	54.27(.07)	55.90(.09)
	Spectral	66.69(.04)	61.02(.05)	69.85(.05)	69.15(.05)	72.32(.04)
	Cepstral	63.13(.09)	68.35(.04)	65.83(.03)	71.79(.06)	70.62(.04)
	Tonal	58.37(.08)	63.98(.05)	65.21(.08)	67.17(.08)	68.81(.08)

76%, 71%). Spectral and cepstral categories achieved similarly high accuracies. Interestingly, the same ranking was prevalent for all 5 ML models, leading to the conclusion that the cepstral and spectral feature categories encode particularly informative COVID-19 data from breath and cough signals. A Repeated Measures ANOVA test [6] confirms that the feature domains lead to statistically significant differences in ROC score for all three datasets ($p < 0.02$).

Individual Features. We start with the best-performing SVM classifier before broadening to include all models to identify general predictive efficiency patterns. The results forming the basis of our analysis are available in Table 5. A Repeated Measures ANOVA test [6] verifies that the sample type leads to statistically significant differences in ROC score across all datasets ($p < 0.05$).

The majority of the 15 features significantly outperformed random guesses for COVID-19 classification across all datasets and sample types. The lowest accuracies were achieved by Coswara-shallow, matching previous findings. Similarities between Cambridge and Coswara-deep were underlined by sample types: BC achieved the highest mean ROC-AUC scores on average, whereas Coswara-shallow was split evenly between B and C. However, given all considered features, the Coswara-shallow dataset still showed its highest accuracy on BC samples since cepstral and tonal features were the most influential overall. MFCC (cepstral), S-CONT (spectral), and C-ENS/C-STFT (tonal) were the highest-scoring features in their categories, whereas the time domain was more variable.

Lastly, we note a surprising trend for MFCC. A prevalent rule of thumb suggests 12–13 coefficients for audio classification [3, 7, 19]. However, Fig. 3 shows that higher-order features provided discriminative information for COVID-19

Table 5. 5-fold CV ROC-AUC as mean(std). The majority of features showed the most accurate results on the BreathCough (BC) vector. Feature categories were ranked in increasing accuracy: time domain, tonal, spectral, and cepstral.

(a) SVM, Cambridge data.					(b) SVM, Coswara-deep data.		
		B	C	BC	B	C	BC
All	All	85.86(.07)	85.80(.05)	87.68(.06)	76.79(.04)	70.85(.06)	77.15(.05)
Time	All	72.77(.04)	74.90(.08)	78.78(.07)	61.80(.04)	58.58(.06)	63.94(.07)
	RMSE	72.28(.05)	76.45(.08)	77.88(.08)	55.89(.10)	61.14(.07)	61.81(.07)
	ZCR	64.59(.08)	69.73(.06)	71.40(.06)	64.68(.03)	59.45(.13)	64.60(.04)
Spectral	All	85.28(.06)	84.03(.07)	84.84(.07)	76.34(.05)	66.74(.05)	74.46(.06)
	S-BW	69.24(.08)	71.57(.04)	75.4(.08)	61.63(.07)	63.51(.05)	65.46(.04)
	S-CENT	73.45(.08)	70.06(.08)	78.07(.07)	68.53(.06)	59.91(.06)	71.95(.05)
	S-CONT	86.14(.06)	84.03(.08)	85.98(.08)	74.89(.05)	63.42(.08)	73.57(.09)
	S-FLAT	74.22(.07)	75.44(.05)	75.87(.06)	61.77(.08)	59.86(.06)	61.14(.03)
	S-FLUX	79.70(.08)	77.14(.06)	82.08(.06)	63.79(.06)	62.76(.07)	67.20(.04)
	S-ROLL	70.70(.07)	67.22(.04)	71.22(.06)	65.35(.05)	63.16(.05)	67.58(.08)
Cepstral	All	86.25(.06)	83.98(.06)	87.15(.06)	74.57(.03)	70.15(.09)	75.62(.08)
	MFCC	86.56(.04)	83.25(.05)	87.68(.04)	74.24(.03)	70.74(.01)	75.38(.05)
	MFCC- Δ	84.21(.04)	79.67(.08)	85.54(.08)	64.85(.07)	68.90(.05)	68.99(.04)
	MFCC- Δ^2	84.25(.09)	78.29(.07)	85.24(.09)	66.65(.08)	67.72(.06)	70.72(.07)
Tonal	All	79.69(.07)	78.06(.07)	82.59(.07)	71.74(.05)	64.06(.06)	72.98(.03)
	C-CQT	76.29(.06)	71.12(.09)	77.30(.06)	67.87(.04)	62.78(.07)	61.50(.05)
	C-ENS	77.56(.07)	72.11(.07)	83.50(.03)	70.03(.07)	65.14(.03)	65.96(.05)
	C-STFT	77.57(.05)	72.65(.03)	77.78(.07)	67.01(.05)	61.80(.08)	68.19(.10)
	TN	74.28(.04)	70.85(.04)	77.57(.05)	60.90(.04)	62.84(.02)	61.33(.03)

on par with (Coswara-deep) or significantly outperforming (Cambridge) lower orders. This phenomenon was most noticeable in BC/B vectors and MFCC features. Since higher-order features contain information about details such as pitch and tone quality [11], we extrapolate that timber is highly relevant to COVID-19.

Discussion. Our extensive analysis, comparison, and ranking of 15 features has found recurring patterns of predictive efficiency for COVID-19 audio classification across independent datasets. There was a distinct category ranking consistent across models, sample types, and datasets (increasing): time domain, tonal, spectral, and cepstral. Contrary to the intuitive expectation, some ‘complex’ categories provided less discriminative information than ‘simpler’ ones (e.g. tonal/spectral features). However, this is justified when considering that tonal features describe pitch and so are more suited to tasks with melodic content.

The ranking underlines the significance of frequency-based features by elevating the spectral and cepstral categories describing timbral aspects and tone quality/colour. We have also shown that the common guideline to use only the first 13 MFCC features [3, 7, 19] was not applicable to COVID-19. Indeed, the




Fig. 3. Normalised ROC-AUC of MFCC and derivatives for BreathCough (BC), Breath (B), and Cough (C) vectors. Contrary to a common rule of thumb [3,7,19], 13+ features provided significant discriminatory data, and showed that timbral information is especially relevant to COVID-19 classification.

higher-order (timbre) features’ predictive efficiency provided significantly more discriminatory information, especially for the BC and B feature vectors.

Taking a step back from the individual features, we note that the most prevailing pattern across all previous descriptions was that the concatenated BC feature vector outperformed the individual B and C vectors in most cases.

Given our insights, we compare our results to the published baselines, summarised in Table 6. The evaluated models were of similar type and complexity; The major difference was our introduction of new training features. We observe that our improved feature vectors significantly outperformed both the Cambridge and Coswara baseline accuracies by 10–17%, validating our feature selection.

Table 6. Comparison to dataset papers’ 5-fold CV baseline results. The most comparable configurations are shown (feature processing and classification model).

Origin	Dataset	Sample	Model	ROC-AUC	Precision	Recall
This paper	Cambridge	BC	SVM	87.68 (.06)	87.61(.07)	81.39(.07)
[3]	Cambridge	BC	LR	71.00(.08)	69.00(.09)	66.00(.14)
This paper	Cos-deep	BC	SVM	77.15 (.05)	76.7(.05)	53.09(.03)
[12]	Cos-unknown	C	RF	67.45(—)	—	—

4 Related Work

During in- and exhalation, air travelling through the respiratory tract undergoes turbulence and produces sounds. Consequently, any physical changes to the airways or lungs (e.g. caused by diseases such as COVID-19) also alter the produced respiratory sounds [17]. Even though listening and evaluating lung sounds manually is inherently subjective, medical professionals have long used this technique to non-invasively diagnose a wide variety of respiratory diseases [2].

The popularisation of digital signal processing techniques and Machine Learning (ML) have made the automatic classification of respiratory sounds possible as a less subjective, low-cost, and patient-friendly (pre-)screening method. A literature review of existing implementations shows that ML can reliably pick up on subtle cues in audio signals for a variety of diseases.

Smartwatches and wearable devices have made audio monitoring for healthcare purposes feasible. Nguyen et al. apply a dynamically activated respiratory event detection mechanism to detect cough and sneeze events non-intrusively [14]. [1] presents classifiers distinguishing between asthma and pneumonia in pediatric patients. Lastly, an image classification solution with comparable results is developed in [2], using spectrograms as the input.

One of the first COVID-19 audio datasets containing breath and cough samples was presented in [3]. Using standard ML and audio processing techniques, the authors report 71% ROC accuracy for COVID classification. [12] and [19] consider further recording types such as vowel intonation and sequence counting, achieving 67% and 66% accuracy with ML models respectively.

5 Conclusion and Future Work

Our extensive comparative analysis of 15 audio features has provided significant insight into Machine Learning (ML) feature selection for COVID-19 respiratory audio classification and addressed the research questions laid out in Sect. 3.1. Primarily, we identified the most informative feature characteristics and verified their ranking across two independent datasets. Since the two feature rankings showed considerable overlap, we conclude that the features' relative salience was likely inherent to the respiratory signals rather than the evaluated datasets.

Throughout our analysis, a number of informative audio features were newly incorporated in the context of COVID-19 classification. In combination with our feature ranking, we achieved 88% and 77% accuracy on the Cambridge and Coswara datasets. Since the complexity of the signal processing and ML models is comparable to the baselines, the increase of up to 17% and 10% respectively was a consequence of our feature selection. Our established feature ranking could benefit future sound-based COVID-19 classification applications.

This paper provides a starting point for the holistic evaluation of respiratory audio features for COVID-19 classification. Considerations that could be addressed in future work are a comprehensive strategy to regularise different sample lengths, and to identify the most informative audio features for complex architectures such as Deep Learning neural networks.

Although sound-based COVID-19 detection was the primary purpose of this research, many other respiratory diseases and disorders could benefit from the development and improvement of automatic audio detection systems for diagnosis, treatment, and management. Therefore, the approach described in this paper could be generalised for the detection of other respiratory diseases.

Acknowledgements. We would like to thank Chris Watkins for the stimulating discussions, and University of Cambridge for access to the COVID-19 sound dataset. This research is funded by University of Brighton’s Connected Futures, Radical Futures’ initiatives, and Santander’s Global Challenges Research grant.

References

1. Amrulloh, Y., Abeyratne, U., Swarnkar, V., Triasih, R.: Cough sound analysis for pneumonia and asthma classification in pediatric population. In: 2015 6th International Conference on Intelligent Systems, Modelling and Simulation, pp. 127–131. IEEE (2015)
2. Aykanat, M., Kılıç, Ö., Kurt, B., Saryal, S.: Classification of lung sounds using convolutional neural networks. *EURASIP J. Image Video Process.* **2017**(1), 1–9 (2017). <https://doi.org/10.1186/s13640-017-0213-2>
3. Brown, C., et al.: 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, pp. 3474–3484 (2020)
4. Hossan, M.A., Memon, S., Gregory, M.A.: A novel approach for MFCC feature extraction. In: 2010 4th International Conference on Signal Processing and Communication Systems, pp. 1–5. IEEE (2010)
5. Humphrey, E.J., Cho, T., Bello, J.P.: Learning a robust Tonnetz-space transform for automatic chord recognition. In: 2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 453–456. IEEE (2012)
6. Iantovics, L.B.: Black-box-based mathematical modelling of machine intelligence measuring. *Mathematics* **9**(6), 681 (2021)
7. Jiang, D.N., Lu, L., Zhang, H.J., Tao, J.H., Cai, L.H.: Music type classification by spectral contrast feature. In: Proceedings, IEEE International Conference on Multimedia and Expo, vol. 1, pp. 113–116. IEEE (2002)
8. Korzeniowski, F., Widmer, G.: Feature learning for chord recognition: the deep chroma extractor. In: Proceedings of the 17th ISMIR Conference, pp. 37–43. International Society for Music Information Retrieval (ISMIR), New York (2016)
9. Laguarda, J., Hueto, F., Subirana, B.: COVID-19 artificial intelligence diagnosis using only cough recordings. *IEEE Open J. Eng. Med. Biol.* **1**, 275–281 (2020)
10. Madhu, N.: Note on measures for spectral flatness. *Electron. Lett.* **45**(23), 1195–1196 (2009)
11. Mitrović, D., Zeppelzauer, M., Breiteneder, C.: Chapter 3 - features for content-based audio retrieval. In: *Advances in Computers*, vol. 78, pp. 71–150. Elsevier (2010)
12. Muguli, A., et al.: DiCOVA challenge: dataset, task, and baseline system for COVID-19 diagnosis using acoustics. arXiv preprint [arXiv:2103.09148](https://arxiv.org/abs/2103.09148) (2021)
13. Müller, M., Kurth, F., Clausen, M.: Audio matching via chroma-based statistical features. In: *ISMIR*, vol. 2005, p. 6 (2005)

14. Nguyen, K.A., Luo, Z.: Cover your cough: detection of respiratory events with confidence using a smartwatch. In: *Conformal and Probabilistic Prediction and Applications*, pp. 114–131. PMLR (2018)
15. Panagiotakis, C., Tziritas, G.: A speech/music discriminator based on RMS and zero-crossings. *IEEE Trans. Multimedia* **7**(1), 155–166 (2005)
16. Peeters, G., Giordano, B.L., Susini, P., Misdariis, N., McAdams, S.: The timbre toolbox: extracting audio descriptors from musical signals. *J. Acoust. Soc. Am.* **130**(5), 2902–2916 (2011)
17. Rizal, A., Hidayat, R., Nugroho, H.A.: Signal domain in respiratory sound analysis: methods, application and future development. *J. Comput. Sci.* **11**(10), 1005 (2015)
18. Saito, T., Rehmsmeier, M.: The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS ONE* **10**(3), e0118432 (2015)
19. Sharma, N., et al.: Coswara-a database of breathing, cough, and voice sounds for COVID-19 diagnosis. arXiv preprint [arXiv:2005.10548](https://arxiv.org/abs/2005.10548) (2020)
20. Stolar, M.N., Lech, M., Stolar, S.J., Allen, N.B.: Detection of adolescent depression from speech using optimised spectral roll-off parameters. *Biomed. J.* **2**, 10 (2018)