

Influence of Event Duration on Automatic Wheeze Classification

Bruno M. Rocha
Univ Coimbra
CISUC
DEI
Coimbra, Portugal
bmrocha@dei.uc.pt

Diogo Pessoa
Univ Coimbra
CISUC
DEI
Coimbra, Portugal
dpessoa@dei.uc.pt

Alda Marques
Univ Aveiro
Lab3R
ESSUA, iBiMED
Aveiro, Portugal
amarques@ua.pt

Paulo Carvalho
Univ Coimbra
CISUC
DEI
Coimbra, Portugal
carvalho@dei.uc.pt

Rui Pedro Paiva
Univ Coimbra
CISUC
DEI
Coimbra, Portugal
ruipedro@dei.uc.pt

Abstract—Patients with respiratory conditions typically exhibit adventitious respiratory sounds, such as wheezes. Wheeze events have variable duration. In this work we studied the influence of event duration on wheeze classification, namely how the creation of the non-wheeze class affected the classifiers’ performance. First, we evaluated several classifiers on an open access respiratory sound database, with the best one reaching sensitivity and specificity values of 98% and 95%, respectively. Then, by changing one parameter in the design of the non-wheeze class, i.e., event duration, the best classifier only reached sensitivity and specificity values of 55% and 76%, respectively. These results demonstrate the importance of experimental design on the assessment of wheeze classification algorithms’ performance.

Keywords—Performance evaluation; Classification; Audio and acoustic processing and analysis

I. INTRODUCTION

Respiratory conditions are leading causes of morbidity and mortality worldwide [1] and generate a significant burden for public health systems and society [2]. Early diagnosis and routine monitoring of patients with respiratory conditions are important for timely interventions; thus, considerable research efforts have been focused on these issues in the last couple of decades [3].

Respiratory sounds are a simple, objective, and noninvasive marker to assess patients’ respiratory condition [4]. In clinical practice they are commonly assessed with pulmonary auscultation using a stethoscope. The expert clinician is trained to listen to and to recognize pathological findings, such as the presence of adventitious sounds (e.g. crackles, wheezes), commonly in the anterior and posterior chest of the patient [5]. Conventional auscultation is commonly applied but has some drawbacks, such as the inability to provide continuous monitoring and its inter-listener variability. Automated respiratory sound analysis could potentially overcome these limitations.

Respiratory sounds are divided into normal and abnormal categories. Normal respiratory sounds are nonmusical sounds produced from breathing and heard over the trachea and chest wall [6]. They show different acoustic properties, such as duration, pitch, and sound quality depending on subjects’ characteristics, subjects’ position, respiratory flow, and place of recording [7], [8]. On the other hand, adventitious sounds are abnormal sounds that are superimposed on the normal

respiratory sounds [6], [9]. Adventitious sounds can be categorized in two main types, continuous and discontinuous [10]. The nomenclature preconized by the standardization of lung sound nomenclature taskforce [11], i.e., continuous adventitious sounds are referred to as wheezes, while discontinuous adventitious sounds are called crackles, will be followed in this study.

Crackles are explosive, short, discontinuous, and nonmusical adventitious sounds that are attributed to the sudden opening and closing of abnormally closed airways [12]. In contrast, wheezes are musical respiratory sounds usually longer than 80-100 ms and with frequencies ranging from less than 100 Hz to 1000 Hz, with harmonics that occasionally exceed 1000 Hz [13]. Wheezes occur when there is a flow limitation and they can be produced by all mechanisms that reduce airway caliber. Clinically, they can be defined by their frequency (mono- or polyphonic), intensity, number, duration, and position in the respiratory cycle (inspiratory or expiratory), gravity influence, and respiratory maneuvers [12]. Wheezes have been used for diagnostic purposes in several respiratory conditions in children (e.g., bronchiolitis) and in adults (e.g., chronic obstructive pulmonary disease (COPD)) [12].

In this work we propose several machine learning methods for the classification of wheeze events against other types of events, randomly selected. The main objectives of this study were to: a) develop several wheeze classification models (traditional and deep learning approaches) and b) study the influence of the randomly generated events on the classification performance.

The article is organized as follows: in section II, we provide a general overview of several state-of-the-art algorithms that have been used in similar works; in section III, we present the database, as well as showing all the methods used in the different experiments; in section IV, the obtained results are displayed; and, lastly, we discuss the results and give a conclusion in section V.

II. RELATED WORK

With the onset of pattern recognition, several features and machine learning approaches have been proposed in the literature to develop automatic systems for wheeze classification.

TABLE I: Summary of selected works

Ref	Data	Classes	Best Results
[14]	Prt: NA; Rec: NA; Src: Priv	2: Wheezes and Normal	Acc: 96%
[15]	Prt: NA ; Rec: 28 Src: R.A.L.E.	2: Wheezes and Normal	Acc: 85%; Sens: 86% ; Spec: 82%
[16]	Prt: 284; Rec: 500; Src: Priv	3: Wheezes, Crackles, and Normal	Wheeze AUC: 86%
[17]	Prt: 30; Rec: 870; Src: Priv	2: Wheezes and Normal	Acc: 94%; Prec: 95%; Sens: 94%; Spec: 94%
[18]	Prt: 12; Rec: 24; Src: Priv	2: Wheezes and Normal	Acc: 98%; Sens: 91%; Spec: 99%; MCC: 93%
[19]	Prt: 1630; Rec: 17930; Src: Priv	2: Healthy and Pathological	Acc: 86%; Prec: 86%; Sens: 86%; Spec: 86%
[9]	Prt: 15; Rec: 15; Src: R.A.L.E.	7, including Wheezes, Crackles, and Normal	Acc: 96%; Wheeze Prec: 98%; Wheeze Sens: 100%
[20]	Prt: 126; Rec: 920; Src: RSD	3: Wheezes, Crackles, and Normal	Wheeze Sens: 79%; Spec: 91%
[21]	Prt: 126; Rec: 920; Src: RSD	3: Wheezes, Crackles, and Normal	Wheeze Sens: 52%; Spec: 52%
[22]	Prt: 126; Rec: 920; Src: RSD	4: Wheezes, Crackles, Wheezes+Crackles, and Normal	Wheeze Sens: 70%; Spec: 74%
[23]	Prt: NA; Rec: 240; Src: R.A.L.E. and RSD	2: Wheezes and Normal	Acc: 99%; Sens: 96%; Spec: 99%
[24]	Prt: 126; Rec: 920; Src: RSD	4: Wheezes, Crackles, Wheezes+Crackles, and Normal	Acc: 50%; Wheeze Sens: 65%; Spec: 63%
[25]	Prt: 126; Rec: 920; Src: RSD	4: Wheezes, Crackles, Wheezes+Crackles, and Normal	Acc: 66%; Wheeze Prec: 55%; Wheeze Sens: 43%; Spec: 78%

Ref: reference; Prt: participants; Rec: recordings; Src: source; Priv: private; AUC: area under the Receiver Operating Characteristic curve; Acc: accuracy; Prec: precision; Sens: sensitivity; Spec: specificity; MCC: Matthews correlation coefficient; NA: not available

Pramono et al. [7] identified the most common features and machine learning algorithms employed in the literature. These include Mel-frequency cepstral coefficients (MFCCs) [26], spectral features [27], entropy [28], wavelet coefficients [29], support vector machines (SVM) [17], artificial neural networks (ANN) [16], rule-based models [30] and logistic regression models [18].

Over the years, several authors have reported excellent results on the wheeze classification task. A summary of selected works can be found in Table I. However, one crucial problem of this field has been its reliance on small or private data collections. Besides, public repositories that have been commonly used in the literature (e.g., R.A.L.E. [31]) were designed for teaching, typically include a small number of adventitious respiratory sounds, and usually do not contain environmental noise. Therefore, we chose to perform the evaluation on the Respiratory Sound Dataset (RSD), the largest publicly available respiratory sound database, which is described in section III.

III. MATERIALS AND METHODS

A. Database

The Respiratory Sound Database (RSD) is a publicly available database with 920 audio files containing a total of 5.5 hours of recordings acquired from 126 participants of all ages [32]. The database main characteristics are described in Table II. The database contains audio samples that were collected independently by two research teams in two different countries and it is a challenging database, since the recordings contain several types of noises, background sounds and different sampling frequencies. It contains annotations of 1898 wheezes, which are found in 341 audio files. The training set contains 1173 wheezes distributed among 203 files, while the testing set includes 725 wheezes distributed among 138 files.

B. Random Event Generation

To establish a comparison with other methods that were tested on this database, we mimicked “Experiment 2” of [32] but focusing only on wheeze classification. As detailed in that experiment, we generated random events using a custom script. To simultaneously guarantee variation and reproducibility, the seed for the random number generator changed for each file but was predetermined. In this work, we employed two approaches

to generate the random events. Both approaches generated the same number of events, as our goal was to study the impact of event duration on the algorithms’ performance. To keep the dataset classes’ proportion similar to the original (40% wheezes, 60% other), we established a rule imposing a limit of one event per 5 s, i.e., the script generated at most four random events in a 20 s file, depending on the number of annotated wheezes. A total of 2910 random events were generated, 1781 in the training set and 1129 in the test set. In the first approach, following [32], events with a fixed duration of 150 ms were generated. In the second approach, we visually inspected the distribution of the annotated wheezes’ durations and found that a Burr distribution [33] provided a good fit. The Burr distribution used to generate the events had probability density function:

$$f(x | a, c, k) = \frac{\frac{kc}{\alpha} \left(\frac{x}{\alpha}\right)^{c-1}}{\left(1 + \left(\frac{x}{\alpha}\right)^c\right)^{k+1}}, x > 0; \alpha > 0; c > 0; k > 0 \quad (1)$$


with $\alpha = 0.2266$, $c = 4.1906$ and $k = 0.3029$.

All the random events were then generated with a duration belonging to this distribution, ranging between 100 ms and 2 s.

TABLE II: Demographic information of database

Number of recordings	920
Sampling frequency (number of recordings)	4 kHz (90); 10 kHz (6); 44.1 kHz (824)
Bits per sample	16
Average recording duration	21.5 s
Number of participants	126: 77 adults, 49 children
Diagnosis	COPD (64); Healthy (26); URTI (14); Bronchiectasis (7); Bronchiolitis (6); Pneumonia (6); LRTI (2); Asthma (1)
Sex	79 males, 46 females (NA: 1)
Age (mean ± standard deviation)	43.0 ± 32.2 years (NA: 1)
Age of adult participants	67.6 ± 11.6 years (NA: 1)
Age of child participants	4.8 ± 4.6 years
BMI of adult participants	27.2 ± 5.4 kg m ² (NA: 2)
Weight of child participants	21.4 ± 17.2 kg (NA: 5)
Height of child participants	104.7 ± 30.8 cm (NA: 7)

COPD: Chronic Obstructive Pulmonary Disease; LRTI: Lower Respiratory Tract Infection; NA: Not Available; URTI: Upper Respiratory Tract Infection.



(a) Wheeze Event

(b) Random Event

Fig. 1: Sample padded spectrograms (Left: Wheeze Event; Right: Random Event).

C. Preprocessing

The audio files in RSD were recorded with different sampling rates. Therefore, we resampled every recording at 4000 Hz, the lowest sampling rate in the database. As the signal of interest was below 2000 Hz [13], this was considered a good resolution for Fourier analysis.

D. Spectrogram

To generate the time-frequency representation of the audio events, the Short-Time Fourier Transform (STFT) was used. The STFT is one of the most used tools in audio analysis and processing since it describes the evolution of the frequency components over time. The STFT representation (F) of a given discrete signal is given by [25]:

$$F(n, \omega) = \sum_{i=-\infty}^{\infty} i\omega(n-i)e^{-j\omega} \quad (2)$$

where $\omega(i)$ is a window function centered at the instant n .

Since the database events had a wide range of durations, a maximum time for event was defined. Using the typical duration of an expiratory cycle, the maximum size per event was established as 2 seconds, and smaller events were centered and zero-padded. The database also contained annotated events longer than 2 s (77 events). For these cases, only the 2 first seconds were considered, as these bigger annotations were regarded as less precise. This stresses the need for a future review and correction of the database annotations.

Spectrograms with a 128 ms hamming window and 75% of overlap were generated. The spectrograms were normalized between 0 and 1. The spectrogram images used by the CNN had 59 by 257 pixels. In Figure 1, it is possible to observe an example of two representative events.

TABLE III: Features description.

Feature	Description
Spectral Centroid	Center of mass of the spectral distribution
Spectral Spread	Variance of the spectral distribution
Zero-crossing Rate	Waveform sign-change rate
Spectral Entropy	Estimation of the complexity of the spectrum
Spectral Flatness	Estimation of the noisiness of a spectrum
Spectral Roughness	Estimation of the sensory dissonance
Spectral Irregularity	Estimation of the spectral peaks' variability
Spectral Flux	Euclidean distance between the spectrum of successive frames
Spectral Brightness	Amount of energy above {100,200,400,800} Hz
Spectral Rolloff	Frequency such that {95,75,25,5}% of the total energy is contained below it
MFCC	13 Mel-frequency cepstral coefficients
Delta-MFCC	1st-order temporal differentiation of the MFCCs
Chromagram Centroid	Tonal centroid
Chromagram Peak	Tonal peak
Pitch	Fundamental frequency estimation
Voicing	Presence of fundamental frequency
Inharmonicity	Partials non-multiple of fundamental frequency

E. Feature Extraction

In this step, 47 features were extracted from each frame of the spectrogram using the MIR Toolbox [34]. Most of these features were used in [18] for wheeze classification and are commonly used in the field of music information research. Table III provides a small description of all the features. Please note that the selected settings were based on previous research and it was not possible to validate all the frame sizes and other extraction parameters for this study. For each event, five statistics of each feature were calculated: mean, standard deviation, median, minimum value, and maximum value. Therefore, the total number of features fed to the classifiers was 235.

F. Classifiers

To classify the samples as wheezes or other we used several machine learning algorithms: linear discriminant analysis (LDA), Linear SVM (SVMlin), Gaussian SVM (SVMrbf), boosted trees (Boost), and a CNN (using the spectrogram images as input). All the classifiers were trained 10 times with different seeds and their parameters were optimized on a validation set containing 25% of the training set. The models with the best parameters were then applied to the testing set. Bayes optimization was used to optimize the following parameters of each traditional machine learning algorithm: delta and gamma for LDA; box constraint and kernel scale for SVM; method, number of learning cycles (NLearn), and learning rate (LearnRate) for Boost. The best parameters for the best run of each system are listed in Table IV.

To develop the CNN model, a base architecture was defined and the parameters of each layer were then optimized using a grid search approach. The optimized parameters were the following ones: Convolution size; Number of convolutional filters; Dropout rate; Pooling size; Size of the fully connected layer. The models were trained with a maximum number of 15 epochs, a mini batch size of 128, Adam optimization (adaptive moment estimation), and 10% of the training dataset was used for validation during the training phase. The best models' architectures for both experiments are described in Table V.

A system consisting of all the features and a logistic regression model was used as a baseline for comparison.

G. Evaluation Metrics

To evaluate the algorithms' performance, we used the following measures:

$$Accuracy(Acc) = \frac{(TP + TN)}{(TP + TN + FP + FN)} \quad (3)$$

$$Precision(Prec) = \frac{TP}{(TP + FP)} \quad (4)$$

$$Sensitivity(Sens) = \frac{TP}{(TP + FN)} \quad (5)$$

$$F1Score(F1) = \frac{(2 \times Prec \times Sens)}{(Prec + Sens)} \quad (6)$$

$$Specificity(Spec) = \frac{TN}{(TN + FP)} \quad (7)$$

$$MatthewsCorrCoef(MCC) = \frac{((TP \times TN) - (FP \times FN))}{\sqrt{((TP + FP)(TP + FN)(TN + FP)(TN + FN))}} \quad (8)$$

where TP (True Positives) are wheeze events that are correctly classified; TN (True Negatives) are other events that are correctly classified; FP (False Positives) are other events that are incorrectly classified as wheezes; FN (False Negatives) are wheeze events that are incorrectly classified as other.

IV. EVALUATION

In this section, we analyze the performance of the algorithms on two experiments corresponding to the two approaches for the generation of random events. We focus especially on MCC, as this is the only evaluation measure that simultaneously considers TP, TN, FP, and FN. To analyze statistical significance, we performed the right-tailed Wilcoxon sign rank test, for a significance level of $\alpha = 1\%$, against the baseline. To adjust for multiple comparisons, we applied Bonferroni correction to the alpha value. Table VI provides details about the training and testing sets.

TABLE IV: Classifiers' hyperparameters (best models).

Algorithm	Parameters (Fixed Durations)	Parameters (Variable Durations)
LDA	Delta: 0.00088684; Gamma: 0.001	Delta: 0.0104; Gamma: 0.0405
SVMlin	BoxConstraint: 931.133; KernelScale: 425.5836	BoxConstraint: 9.6309; KernelScale: 923.6736
SVMrbf	BoxConstraint: 86.0544; KernelScale: 990.1481	BoxConstraint: 27.2645; KernelScale: 999.2519
Boost	Method: Logitboost; NLearn: 186; LearnRate: 0.2582	Method: LogitBoost; NLearn: 264; LearnRate: 0.6745

TABLE V: CNN's architecture.

Layer	Activations (Fixed Durations)	Parameters (Fixed Durations)	Activations (Variable Durations)	Parameters (Variable Durations)
Input	257 x 59 x 1	-	257 x 59 x 1	-
2D Convolution	251 x 53 x 64	Convolution Size-7; Stride-1; Filters-64	253 x 55 x 32	Convolution Size-5; Stride-1; Filters-32
Batch normalization	251 x 53 x 64	-	253 x 55 x 32	-
ReLU	251 x 53 x 64	-	253 x 55 x 32	-
Max Pooling	250 x 52 x 64	Pooling Size-2; Stride-1	250 x 55 x 32	Pooling Size-4; Stride-1
Dropout	250 x 52 x 64	50%	250 x 55 x 32	50%
Fully Connected	1 x 1 x 10	Size-10	1 x 1 x 20	Size-20
Dropout	1 x 1 x 10	50%	1 x 1 x 20	50%
Fully Connected	1 x 1 x 2	Size-2	1 x 1 x 2	Size-2
Softmax	1 x 1 x 2	-	1 x 1 x 2	-

A. Fixed Durations

Table VII displays the results achieved by all the combinations of classifiers and feature sets on the fixed durations (FD) test set. Figure 2 shows the MCC results with the corresponding error bars.

As mentioned in section II, the best system's wheeze Sens and Spec in [32] were 78.6% and 90.5%, respectively. With Sens and Spec of 79% and 97.9%, respectively, our baseline's performance was on par with the state-of-the-art. In addition, it achieved better results than the other traditional classifiers (MCC of 80.3%). Nevertheless, the CNN achieved even better results, reaching an MCC of 91.8%. Considering the MCC, only the CNN performed significantly better than the baseline.

B. Variable Durations

Table VIII displays the results achieved by all the combinations of classifiers and feature sets on the variable durations (VD) test set. Figure 2 shows the MCC results with the corresponding error bars.

The decline in performance is quite noticeable. The best classifier in this experiment, Boost, reached an MCC of 31.7%. However, none of the optimized algorithms significantly outperformed the baseline (MCC of 30.6%).

C. Variance in performance

The substantial variance in performance between experiments might indicate that the generation of the random events with fixed durations introduced considerable bias. Figure 2

TABLE VI: Training and test dataset classes distribution.

	Training set	Testing set	Total
Number of wheezes	1173	725	1898
Number of random events	1781	1129	2910
Total	2954	1854	4808

TABLE VII: Performance results on FD test set (* statistically significant difference ($\alpha = 0.01$)) (mean \pm standard deviation)

	Accuracy	Precision	Sensitivity	F1	Specificity	MCC
Baseline	90.5 \pm 0	96 \pm 0	79 \pm 0	86.7 \pm 0	97.9 \pm 0	80.3 \pm 0
LDA	89.1 \pm 0.4	94.2 \pm 0.5	76.8 \pm 1.3	84.6 \pm 0.7	97 \pm 0.3	77.2 \pm 0.8
SVMlin	87.7 \pm 2.5	91.1 \pm 7.4	76.7 \pm 3.2	83 \pm 2.5	94.7 \pm 5.3	74.3 \pm 5.1
SVMrbf	88.4 \pm 0.5	90.2 \pm 0.8	78.8 \pm 1.8	84.1 \pm 0.9	94.5 \pm 0.6	75.4 \pm 1.1
Boost	90.6 \pm 0.7	91.7 \pm 2.4	83.6 \pm 1.6 *	87.4 \pm 0.8	95 \pm 1.6	80.1 \pm 1.4
CNN	96 \pm 2.2 *	91.9 \pm 4.3	98.1 \pm 5.1 *	94.7 \pm 2.8 *	95.1 \pm 2.4	91.8 \pm 4.4 *

TABLE VIII: Performance results on VD test set (* statistically significant difference ($\alpha = 0.01$)) (mean \pm standard deviation)

	Accuracy	Precision	Sensitivity	F1	Specificity	MCC
Baseline	67.2 \pm 0	58.4 \pm 0	55.9 \pm 0	57.1 \pm 0	74.5 \pm 0	30.6 \pm 0
LDA	66.2 \pm 0.2	57.8 \pm 0.4	50.2 \pm 1.0	53.7 \pm 0.6	76.5 \pm 0.6 *	27.5 \pm 0.5
SVMlin	66.7 \pm 0.2	60.5 \pm 0.7 *	42.7 \pm 1.2	50.1 \pm 0.7	82.1 \pm 0.9 *	27.1 \pm 0.5
SVMrbf	62.9 \pm 0.7	53.2 \pm 1.0	43.6 \pm 1.3	47.9 \pm 1.1	75.4 \pm 0.8	19.8 \pm 1.4
Boost	67.8 \pm 1.3	59.6 \pm 2.1	55.3 \pm 4.3	57.2 \pm 2.4	75.9 \pm 3.1	31.7 \pm 2.7
CNN	60.3 \pm 5.2	67.5 \pm 20.4	51 \pm 4.9	55.9 \pm 7.4	74.7 \pm 5.4	24.4 \pm 4.5

shows the MCC results for each experiment with the corresponding error bars. The classifiers could be implicitly learning to identify the duration of the events. To analyze this bias, we can plot the histogram of event durations of the errors the best classifiers made. The best run of the Boost model yielded 124 FN. As displayed in Figure 3a, all the misclassified wheezes had short durations of less than 325 ms. When classifying wheezes against random events of variable duration, the distribution of FN durations was quite different, encompassing all durations, as shown in Figure 3b. When using the CNN to perform the same classification, the effect of implicitly learning the duration of events was even more noticeable, as can be seen when comparing Figure 3c and Figure 3d. All the misclassified wheezes in the fixed durations test set belonged to the same histogram bin, comprising durations between 125 ms and 175 ms. When considering the variable durations, the CNN model’s FN distribution also covered the full range of durations. To improve visualization, all the annotated wheezes longer than 2 s were equalized and assigned to the same histogram bin.

Figure 4 illustrates how the same algorithm attributed different classes to the same events depending on how the non-wheeze class was designed. On the FD experiment, all annotated wheezes in this sound file were correctly classified, while on the VD experiment, two out of six annotated wheezes were incorrectly classified. Through visual inspection of the spectrogram alone, it is not possible to understand why the 3rd and 6th wheezes were correctly classified while the 4th and 5th wheezes were not, as they seem to show similar spectrogram signatures. This figure also reveals potential annotation inaccuracies, as the spectrogram suggests a possible wheeze that was not annotated between 19.4 s and 20 s.

V. DISCUSSION AND CONCLUSION

In this work, we proposed and evaluated several wheeze classification systems, and demonstrated how event duration can

have a significant impact on automatic wheeze classification. As the algorithms’ performance presented in the previous section show, methods that seem to achieve promising results can fail if we change the way the non-wheeze class is designed. This can happen even if the dataset where the systems are evaluated does not change. It is important to consider how data are used to train, validate, and test a classification algorithm. Our goal was not to question state-of-the-art results in wheeze classification, but to emphasize the critical role of experimental design in the evaluation of the algorithms’ performance. As asserted by Sturm [35], when the performance of an artificial system appears to support the claim that the system is addressing a complex human task (e.g., classify sounds as wheezes), the default position (null hypothesis) should be that the system is







Fig. 2: MCC values for 10 runs in both experiments (mean and standard deviation).




(a) FD test set (best run – Boost).



(b) VD test set (best run – Boost).



(c) FD test set (best run – CNN).



(d) VD test set (best run – CNN).

Fig. 3: Histograms of wheezes and FN durations.




Fig. 4: Classification example on VD test set (file: 133_3p4_Tc_mc_AKGC417L)

not actually addressing the problem it appears to be solving; in this case, a system that claims to distinguish a wheeze from a random event is implicitly learning an irrelevant characteristic of the dataset, i.e., event duration. To determine if a system is relevant, we need to understand the extent to which the characteristics it is extracting from the signal are confounded with the ground truth, i.e., the expert's annotations. Besides, investigations like this are only possible if algorithm evaluation is performed on open access datasets.

Nevertheless, it is important to reiterate the limitations of this database, which may have influenced the performance of the evaluated systems. As pointed out in [32], these include the lack of healthy adult participants and the absence of gold standard annotations, i.e., annotations from multiple annotators. A future update of the database should also check for possible errors.

Wheeze classification is a complex task that is not yet solved, despite the claims made in the literature. It may be particularly hard when algorithms are evaluated on challenging datasets, such as the RSD. Though significant work has been developed to classify wheezes, none has been widely accepted [36]. While CNNs have become state-of-the-art solutions in several tasks [9], a simple CNN with no preprocessing was not enough to tackle this problem. Therefore, accelerating the development of machine learning algorithms is critical to the future of respiratory sounds analysis.

ACKNOWLEDGMENT

This work was partially supported by: Fundação para a Ciência e Tecnologia (FCT) PhD scholarship SFRH/BD/135686/2018; WELMO project, co-funded by the Horizon 2020 Framework Programme of the European Union under grant agreement 825572; Fundo Europeu de Desenvolvimento Regional (FEDER) through Programa Operacional Competitividade e Internacionalização (COMPETE) and FCT under the project UID/BIM/04501/2013 and POCI-01-0145-FEDER-007628—iBiMED.

REFERENCES

- [1] WHO, "The top 10 causes of death," 2018.
- [2] G. J. Gibson, R. Loddenkemper, B. Lundbäck, and Y. Sibille, "Respiratory health and disease in Europe: The new European Lung White Book," *European Respiratory Journal*, vol. 42, no. 3, pp. 559–563, 2013.
- [3] A. Marques, A. Oliveira, and C. Jácome, "Computerized adventitious respiratory sounds as outcome measures for respiratory therapy: A systematic review," *Respiratory Care*, vol. 59, no. 5, pp. 765–776, 2014.
- [4] C. Jácome and A. Marques, "Computerized Respiratory Sounds in Patients with COPD: A Systematic Review," *COPD: Journal of Chronic Obstructive Pulmonary Disease*, vol. 12, pp. 104–112, January 2015.
- [5] S. Fleming, A. Pluddemann, J. Wolstenholme, C. Price, C. Heneghan, and M. Thompson, "Diagnostic Technology: Automated lung sound analysis for asthma," Tech. Rep. July, 2011.
- [6] A. R. Sovijärvi, F. Dalmasso, J. Vanderschoot, L. P. Malmberg, G. Righini, and S. A. Stoneman, "Definition of terms for applications of respiratory sounds," *European Respiratory Review*, vol. 10, no. 77, pp. 597–610, 2000.
- [7] R. X. A. Pramono, S. Bowyer, and E. Rodriguez-Villegas, "Automatic adventitious respiratory sound analysis: A systematic review," *PLoS ONE*, vol. 12, no. 5, pp. 1–43, 2017.
- [8] A. Oliveira and A. Marques, "Respiratory sounds in healthy people: A systematic review," *Respiratory Medicine*, vol. 108, pp. 550–570, apr 2014.

- [9] D. Bardou, K. Zhang, and S. M. Ahmad, "Lung sounds classification using convolutional neural networks," *Artificial Intelligence in Medicine*, vol. 88, pp. 58–69, June 2018.
- [10] L. J. Hadjileontiadis and Z. M. K. Moussavi, *Current Techniques for Breath Sound Analysis*, pp. 139–177. Cham: Springer International Publishing, 2018.
- [11] H. Pasterkamp, P. L. Brand, M. Everard, L. Garcia-Marcos, H. Melbye, and K. N. Priftis, "Towards the standardisation of lung sound nomenclature," *European Respiratory Journal*, vol. 47, no. 3, pp. 724–732, 2016.
- [12] A. Marques and A. Oliveira, *Normal Versus Adventitious Respiratory Sounds*, ch. 10, pp. 181–206. Cham: Springer International Publishing, 2018.
- [13] A. Bohadana, G. Izbicki, and S. S. Kraman, "Fundamentals of Lung Auscultation," *New England Journal of Medicine*, vol. 370, pp. 744–751, February 2014.
- [14] K. E. Forkheim, D. Scuse, and H. Pasterkamp, "Comparison of neural network models for wheeze detection," *IEEE WESCANEX Communications, Power, and Computing*, vol. 1, no. 95, pp. 214–219, 1995.
- [15] R. Riella, P. Nohama, and J. Maia, "Method for automatic detection of wheezing in lung sounds," *Brazilian Journal of Medical and Biological Research*, vol. 42, pp. 674–684, July 2009.
- [16] D. Chamberlain, R. Kodgule, D. Ganelin, V. Miglani, and R. R. Fletcher, "Application of semi-supervised deep learning to lung sound analysis," *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pp. 804–807, 2016.
- [17] M. Lozano, J. A. Fiz, and R. Jané, "Automatic Differentiation of Normal and Continuous Adventitious Respiratory Sounds Using Ensemble Empirical Mode Decomposition and Instantaneous Frequency," *IEEE Journal of Biomedical and Health Informatics*, vol. 20, no. 2, pp. 486–497, 2016.
- [18] L. Mendes, I. M. Vogiatzis, E. Perantoni, E. Kaimakamis, I. Chouvarda, N. Maglaveras, V. Tsara, C. Teixeira, P. Carvalho, J. Henriques, and R. P. Paiva, "Detection of wheezes using their signature in the spectrogram space and musical features," *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pp. 5581–5584, 2015.
- [19] M. Aykanat, Ö. Kılıç, B. Kurt, and S. Saryal, "Classification of lung sounds using convolutional neural networks," *EURASIP Journal on Image and Video Processing*, vol. 2017, p. 65, December 2017.
- [20] G. Serbes, S. Ulukaya, and Y. P. Kahya, *An Automated Lung Sound Preprocessing and Classification System Based On Spectral Analysis Methods*, vol. 66, pp. 45–49. Singapore: Springer Singapore, 2018.
- [21] N. Jakovljević and T. Lončar-Turukalo, *Hidden Markov Model Based Respiratory Sound Classification*, pp. 39–43. Singapore: Springer Singapore, 2018.
- [22] K. Kochetov, E. Putin, M. Balashov, A. Filchenkov, and A. Shalyto, "Noise Masking Recurrent Neural Network for Respiratory Sound Classification," in *International Conference on Artificial Neural Networks*, 2018.
- [23] H. Chen, X. Yuan, J. Li, Z. Pei, and X. Zheng, "Automatic multi-level in-exhale segmentation and enhanced generalized s-transform for wheezing detection," *Computer Methods and Programs in Biomedicine*, vol. 178, pp. 163 – 173, 2019.
- [24] S. Ntalampiras and I. Potamitis, "Classification of Sounds Indicative of Respiratory Diseases," in *Engineering Applications of Neural Networks*, no. PART 1, pp. 93–103, Springer International Publishing, 2019.
- [25] F. Demir, A. Sengur, and V. Bajaj, "Convolutional neural networks based efficient approach for classification of lung diseases," *Health Information Science and Systems*, vol. 8, no. 1, 2020.
- [26] N. Nakamura, M. Yamashita, and S. Matsunaga, "Detection of patients considering observation frequency of continuous and discontinuous adventitious sounds in lung sounds," *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, vol. 2016-October, pp. 3457–3460, 2016.
- [27] P. Bokov, B. Mahut, P. Flaud, and C. Delclaux, "Wheezing recognition algorithm using recordings of respiratory sounds at the mouth in a pediatric population," *Computers in Biology and Medicine*, vol. 70, pp. 40–50, 2016.
- [28] X. Liu, W. Ser, J. Zhang, and D. Y. T. Goh, "Detection of adventitious lung sounds using entropy features and a 2-D threshold setting," *2015 10th International Conference on Information, Communications and Signal Processing, ICICS 2015*, 2016.
- [29] S. Ulukaya, G. Serbes, and Y. P. Kahya, "Overcomplete discrete wavelet transform based respiratory sound discrimination with feature and deci-

- sion level fusion,” *Biomedical Signal Processing and Control*, vol. 38, pp. 322–336, 2017.
- [30] S. A. Taplidou and L. J. Hadjileontiadis, “Wheeze detection based on time-frequency analysis of breath sounds,” *Computers in Biology and Medicine*, vol. 37, no. 8, pp. 1073 – 1083, 2007.
- [31] D. Owens, “R.A.L.E. Lung Sounds 3.0,” *CIN: Computers, Informatics, Nursing*, vol. 5, no. 3, 2002.
- [32] B. M. Rocha, D. Filos, L. Mendes, G. Serbes, S. Ulukaya, Y. P. Kahya, N. Jakovljević, T. Lončar-Turukalo, I. M. Vogiatzis, E. Perantoni, E. Kaimakamis, P. Natsiavas, A. Oliveira, C. Jácome, A. Marques, N. Maglaveras, R. P. Paiva, I. Chouvarda, and P. Carvalho, “An open access database for the evaluation of respiratory sound classification algorithms,” *Physiological Measurement*, vol. 40, no. 3, 2019.
- [33] I. W. Burr, “Cumulative Frequency Functions,” *The Annals of Mathematical Statistics*, vol. 13, no. 2, pp. 215–232, 1942.
- [34] O. Lartillot and P. Toivainen, “Mir in matlab (II): A toolbox for musical feature extraction from audio,” *Proceedings of the 8th International Conference on Music Information Retrieval, ISMIR 2007*, pp. 127–130, 2007.
- [35] B. L. Sturm, “A simple method to determine if a music information retrieval system is a ‘horse’,” *IEEE Transactions on Multimedia*, vol. 16, no. 6, pp. 1636–1644, 2014.
- [36] A. Marques and C. Jácome, “Future Prospects for Respiratory Sound Research,” in *Breath Sounds*, pp. 291–304, Cham: Springer International Publishing, 2018.