UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

PRACTICAL COUGH DETECTION IN PRESENCE OF BACKGROUND NOISE AND PRELIMINARY DIFFERENTIAL DIAGNOSIS FROM COUGH SOUND USING ARTIFICIAL INTELLIGENCE

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE IN ELECTRICAL AND COMPUTER ENGINEERING

BY

CHARLES N. JOHN

Norman, Oklahoma

2020

PRACTICAL COUGH DETECTION IN PRESENCE OF BACKGROUND NOISE AND PRELIMINARY DIFFERENTIAL DIAGNOSIS FROM COUGH SOUND USING ARTIFICIAL INTELLIGENCE

A THESIS APPROVED FOR THE

SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY

Dr. Ali Imran, Chair

Dr. Thordur Runolfsson

Dr. Samuel Cheng

Acknowledgments

Firstly, I would like to thank my Elder Brother for His love and for all the help He has done until now.

Secondly, I would like to express my deepest appreciation to my committee chair, Professor Ali Imran, who has the attitude and the substance of a genius: he continually and convincingly conveyed a spirit of adventure in regards to research and scholarship, and an excitement and unwavering enthusiasm in regard to teaching. His passion for the medical field inspired me to appreciate the importance of research in medicine which led me to choose a practical thesis topic on this critical issue.

Thirdly, I would like to thank the rest of the committee, Dr. Thordur Ronulfsson for being kind and approachable, and Dr. Samuel Cheng for being very approachable and for taking the time to review the manuscript and attending my thesis defense.

To all my friends and colleagues at AI4Networks Research Center who supported and helped me during my stay here. I am so thankful and proud to be a part of this group.

Lastly, I would like to dedicate this work to the families who lost their loved ones during this COVID-19 pandemic. Hopefully, we would all be prepared with all the technology we have to face this kind of pandemic if it ever hits again.

Table of Contents

Table of Tables

Table of Figures

Acronyms Used

ConvNet/CNN	– Convolutional Neural Networks
ReLU/relu	– Rectified Linear Unit
ResNet	– Residual Networks
WHO	– World Health Organization
XGBoost	– eXtreme Gradient Boosting

Abstract

Cough is one of the most common symptoms for many of the diseases. Physicians have been using characteristics of cough for preliminary diagnosis of certain respiratory diseases for ages. But the methods have been subjective and often depend on self-reported history and description of cough by the patients. Recently, with advent of the omnipresent recording devices and advances in machine learning capabilities, many studies have attempted to partially fill the gap. These studies have approached the problem objectively to create devices like cough monitors, cough counters, and partial automatic cough detection using machine learning. There is still a huge gap that exists in detecting and diagnosing the cough in a practical way. This study is an attempt to contribute towards filling this gap. We propose and analyze a machine learning based method to automatically detect cough in presence of background noise. After successful cough detection, we investigate the possibility of preliminary differential diagnosis by distinguishing the cough associated with Asthma, Bronchitis, Bronchiolitis, Pertussis patients and healthy people.

As more training data could be collected for cough and non-cough sounds, it allowed us to leverage the potential of powerful deep architecture like ResNet for the cough detection part. For the diagnosis part of the work, not much data was available. In this case the preliminary results show that XGBoost performed better than CNN and ResNet architectures. While the cough detection part of the study offers mature results, lot more cough sound data for the examined diseases is needed before generalizable conclusions can be drawn from the diagnosis results observed in this study.

CHAPTER 1

Introduction

1.1. Introduction and Background

We live in the era where artificial intelligence is very widely used technology in many fields. It already affected the way we live and drastically changed this world we live-in in the past decade. Artificial intelligence has a potential to revolutionize our lives for good. Medicine is one of those fields that is being revolutionized but has a lot more potential to leverage AI in its field. AI has a very huge potential to leverage that data to make useful predictions on the image data especially after the introduction of the Convolutional Neural Networks.

One of the medium that can be leveraged for finding the symptoms of a disease is cough. Many doctors that are researching on the cough sounds believe that a cough carries valuable information to at least preliminarily diagnose the underlying disease. The premise is that the underlying disease determines the physical character of the cough [1]. According to WHO, Pertussis is one of the diseases that has taken the lives of 89,000 people in the year of 2008 alone while infecting 151, 074 globally.

The premise is, if the lungs are affected then the sound of the cough should be affected as well, but the human ear may not capable of differentiating the coughs of various diseases. This is where the AI comes into play to detect the complex patterns or latent features or distinctive features in an signal/cough.

In this work five kinds of cough sounds are used Asthma, Bronchitis, Bronchiolitis, Normal and Pertussis with the size of 37, 96, 35, 247 and 130 respectively. RESNET and XGBoost models are used for detection and the diagnosis which are expounded in the following sections. The precision and the recall of the model on this dataset is high and is useful for medical purposes assuming the model performs at least the same, if not better with the large data.

This model could be then deployed on an App and it can be accessible for anyone with a smartphone to detect and diagnose a cough.

1.2. Previous Works

Cough has been used for diagnosing by physician for centuries. Recently machine learning has been leveraged to do the same as well for Pertussis, wet-cough and dry-coughs. For automatic cough segmentation, Martinek et al. [3] extracted many time, frequency, and entropy features and used a decision tree to discriminate between voluntary cough sounds and speech. They used data from 20 subjects, with 46 coughs from each subject, and reported median sensitivity and specificity values of 100% and 95% respectively. However, their method is subject-dependent since the subjects are required to cough at the beginning of each recording to obtain individual cough signal patterns. Tracey et al. [5] developed an algorithm for cough detection to monitor patient recovery from tuberculosis. They extracted MFCCs from the audio signals of 10 test subjects. These were used to detect coughs with a combination of artificial neural networks (ANN) and support vector machine (SVM) classifiers achieving an overall sensitivity of 81%. In both these methods, the total number of coughs in the dataset were not reported. Barry et al. [4] used linear predictive coding (LPC) coefficients with a probabilistic neural network (PNN) classifier to create an automatic cough counting tool called Hull Automatic Cough Counter (HACC). This successfully discriminated between a cough and non-cough events from 33 subjects with a sensitivity of 80% and specificity of 96%. Swarnkar et al. [6] used other spectral features such as formant frequencies, kurtosis, and B-score together with MFCC features for cough detection. These were fed into a neural network resulting in a sensitivity of 93% and a specificity of 94% for a test dataset consisting of 342 coughs from 3 subjects only. Matos et al. [8] extracted thirteen MFCCs which were classified using a Hidden Markov Model (HMM). Their test dataset consisted of 2155 coughs from 9 subjects and their method resulted in 82% sensitivity for cough detection. The total number of false detections was not reported, however, the average false positives per hour were 7 with the high variance between subjects. Amrulloh et al. [7] used ANN classification to develop a cough detector using a non–contact recording system for pediatric wards achieving sensitivity and specificity values of 93% and 98% respectively using over 1400 cough sounds from 14 subjects. Liu et al. [9] used Gammatone Cepstral Coefficient (GMCC) features with SVM classification of 903 coughs from 4 subjects resulting in sensitivity and specificity of 91% and 95% respectively. Larson et al. [11] presented a method of cough detection using the built-in microphone of a mobile phone for data collection. Their algorithm, which was not implemented on the phone, used random forest classification with a maximum for 500 decision trees achieving 92% sensitivity on over 2500 cough sounds from 17 subjects. Lucio et al. [10] extracted 79 MFCC and Fast Fourier Transform (FFT) coefficients and used k-Nearest Neighbor (kNN) for classification. From a dataset acquired from 50 individuals, their algorithm achieved a sensitivity of 87% in classifying 411 cough sounds with a specificity of 84%. All these methods work to identify cough segments from audio recordings but do not classify them.

Several algorithms for automatic cough classification have been published to identify various cough types. Chatrzarrin et al. [12] studied the different phases of dry and wet coughs and found the second phase of dry coughs to have lower energy compared to wet coughs. They also noted that, during this phase, most of the signal power is contained between 0-750 Hz in case of wet coughs and 1500-2250 Hz in case of dry coughs. Using a simple thresholding method, they successfully identified 14 wet and dry coughs with 100% accuracy. Swarnkar et al. [13] used a Logistic Regression Model (LRM) based classifier to discriminate between dry and wet coughs from pediatric patients with different respiratory illnesses. They used several features including B–score, non-gaussianity, formant frequencies, kurtosis, zero-crossing rate, and MFCCs. For a test database with 117 coughs from 18 subjects, they reported sensitivities of 84% and 76% for detecting wet and dry coughs respectively. Kosasih et al. [14] developed an algorithm for the automatic diagnosis of childhood pneumonia by assessing cough sounds and crackles. They used MFCCs, non–gaussianity index, and wavelet features with an LRM classifier to differentiate between pneumonia and non–pneumonia cough sounds. Their method achieved sensitivity and specificity of 81% and 50% respectively for a total of 375 cough samples from 25 subjects. Specific to pertussis coughs, Parker et al. [15] studied the performance of three different classifiers for their classification. They used audio files of pertussis cough sounds available on the internet to create a dataset consisting of 16 non-pertussis cough signals and 31 pertussis cough signals. From this data, the cough sounds were then manually isolated and divided into three parts for which 13 MFCC features and the energy level were extracted. These features were subsequently classified using an ANN, a random forest classifier, and a kNN classifier. For each of these classifiers, the false-positive error was 7%, 12%, and 25% respectively while the false negative error was 8%, 0%, and 0% respectively[16].

1.3. Contributions

The works that are presented above contribute on similar lines in terms of diagnosing and detecting the cough. The novel contributions of this work are:

1. Detecting cough in presence of background noise. It is not only the latent features of cough that has value for the medical diagnosis using machine learning, but knowing intensity and temporal pattern of cough also has value for diagnosis by doctors. The cough detection algorithm developed in this thesis when deployed via a smart phone or wearable device can help gather this valuable data.
2. Preliminary differential diagnosis of screening of five respiratory diseases. Other studies have focused on Pertussis, wet and dry cough. No existing study has have focused on differential diagnosis while considered mixed cough sounds from Asthma, Bronchiolitis, Bronchitis, Pertussis patients and healthy people.

1.4. Limitations

a. Majority of effort during this study was consumed by data collection and pre-processing. While we managed to collect reasonable amount of training data to train and develop a robust cough detection algorithm in presence of various type of background and environmental noises, the data that could be collected for training and testing the diagnosis algorithm is still too small. Therefore, while the 2nd part of this work investigates first feasibility study of possible preliminary diagnosis from cough sound for Asthma, Bronchiolitis, Bronchitis and Pertussis, its generalization capability remains an open question till more data becomes available.

b. ResNet while being powerful and proven tool, is a very deep model with millions of parameters to be tuned during the training process. Given the small amount of training data available for the diagnosis part of this study, it is quite possible that it has memorized instead of learning. More data and investigation into interpretability in future would be needed to address this limitation.

1.5. Articles currently under review for Publication

C. Bales Et al., *"Can Machine Learning Be Used to Recognize and Diagnose Coughs?"* (Submitted)

- The major contribution I have done for this work was data collection, pre-processing, and detection part of the algorithm

Charles John Et al., *"Practical and Deployable Cough Detection and Differential Diagnosis of Screening Using Artificial Intelligence"* (In preparation)

- My contribution in this paper was data preprocessing, collecting more data, applying machine learning and deep learning techniques for detection and the diagnosis of the coughs

1.6. Organization

The rest of this thesis is structured as follows: Data Generation and Preprocessing methodology is presented in Chapter II and Chapter III respectively. Chapter IV presents the method for Detection and Diagnosis model which includes ResNet, CNN, and XGBoost models used in this thesis. Performance Analysis and Final Results are presented in Chapter V. Chapter VI concludes the thesis and Future Work is presented in Chapter VII.

CHAPTER 2

Data Generation and Preprocessing

2.1. Data Generation

The data for the detection algorithm was collected through various mediums, via the AI4Lyf app, social media, and used online datasets ESC-50, FSDKaggle 2018 and 2019. All the data was in the wav form and it was converted to spectrograms.

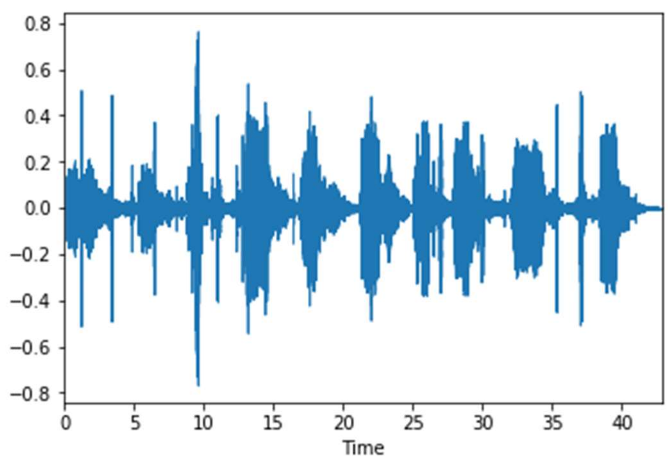

Figure 2.1.: Wave form of a sound

The data consists of cough sounds, male voices, singing, female voices, music, speech, random noises, and the 50 categories from ESC-50 dataset. All these sounds are of 44 kHz converted to spectrogram using the librosa library in python. Here are the steps involved to convert a wave form representation to spectrogram representation:

- The Fourier Transform is applied to the sound and because human hearing range is concentrated in very small frequencies and amplitude ranges, a small adjustment is made to transform both the y-axis (frequency) to log scale, and the "color" axis (amplitude) to Decibels
- In contrast to Hz scale, where the difference between 500 and 1000 Hz is obvious, whereas the difference between 7500 and 8000 Hz is barely noticeable
- This representation is called Spectrogram, as shown in Fig. 2 below

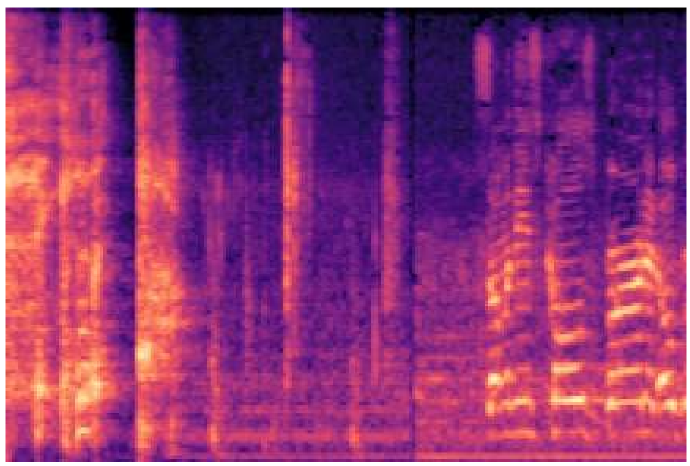

Figure 2.2.: Spectrogram of the sound file

As there the data for the cough sounds was less, several data augmentation techniques are used to generate huge number of data to train the neural network. Noise injection, shifting time, changing speed, changing pitch and the combination of these are used to increase the dataset size. The number has increased from 1,470 to 14,700 for cough and to 29,080 for the non-cough class. The majority of the data for the diagnosis is obtained from the YouTube. No data augmentation is used for diagnosis dataset.

2.2. Data Preprocessing

The size of the detection training dataset is approximately 35,000 images with the size of 640 x 480 each. The test data is separated in a separate folder, carefully examining to see that there is no same sample or augmented image split between the training and the test sets. The diagnosis dataset is of the same size but the size of the image is reduced to 432 x 288 during the runtime. All the different types of diseases are in their respective folders. The raw data obtained were listened physically and individually and cut them into two second events each. It is done for all the cough and the non-cough samples.

All the images before they are fed to the model, are converted to greyscale images to reduce the computational costs and the complexity of the model and training. A greyscale image is represented by a pixel matrix. Each pixel in the image is presented by one integer from 0 – 255. The training dataset is further divided into validation and training sets.

For the diagnosis part of this work, the image data is reshaped into one-directional array then it is given to the machine learning model, which is XGBoost here.

This is the procedure followed for most of the data collection and pre-processing:

1. Downloaded the audio file from the internet sources of the particular diseases – some from the social media website, some from the online open source websites.

2. Listened to each sound file and extracted the cough sound by cutting that part from the audio file using an audio cutter. This was done for all the audio sound files extracting three second files from each sound file
3. Each sound then is preprocessed with a python program to make it precisely three seconds each as most of the cough sounds are less than three seconds by adding silence to the file
4. All the sound files are converted to mono and the output was given as 44kHz with .wav as the format of the sound file
5. This then was fed to a python program that converts the sound file into a spectrogram image by using the package 'librosa'
6. All the images are stored in a folder which is then fed to the machine learning/deep learning algorithm
7. Inside the algorithm, all the images are resized to 432 x 288 images and was reshaped to greyscale image for the sake of reducing complexity and computational costs

Until here, the process is the same for all the algorithms used. For XGBoost there are a few extra steps to follow, the input cannot be a matrix, so the input has to be flattened to one directional array. The training dataset is then split to validation and test dataset and is fed to the algorithm.

CHAPTER 3

Detection and Diagnosis

3.1. Residual Networks

One of the problems that very deep models face is the problem of vanishing/exploding gradients[1,9], which hamper convergence from the beginning. This problem, however, has been largely addressed by normalized initialization [23, 9, 37, 13] and intermediate normalization layers [16], which enable networks with tens of layers to start converging for stochastic gradient descent (SGD) with backpropagation [22]. When the deeper networks' depth increases the accuracy degrades rapidly, this problem is addressed by Deep Residual Learning framework.

The core of these networks is introducing a 'identity shortcut connection' that skips one or more layers. This block is called Residual Block, as shown in the figure below. This identity mapping does not have any parameters and is just there to add the output from the previous layer to the layer ahead. However, sometimes x and F(x) may not have the same dimension. A convolution operation typically shrinks the spatial resolution of an image, e.g. a 3x3 convolution on a 32 x 32 image results in a 30 x 30 image. The identity mapping is multiplied by a linear projection W to expand the channels of shortcut to match the residual. So the input x and F(x) are added as input to the next layer. The Skip Connections between layers add the outputs from previous layers to the outputs of stacked layers. This results in the ability to train much deeper networks than the previous ones.

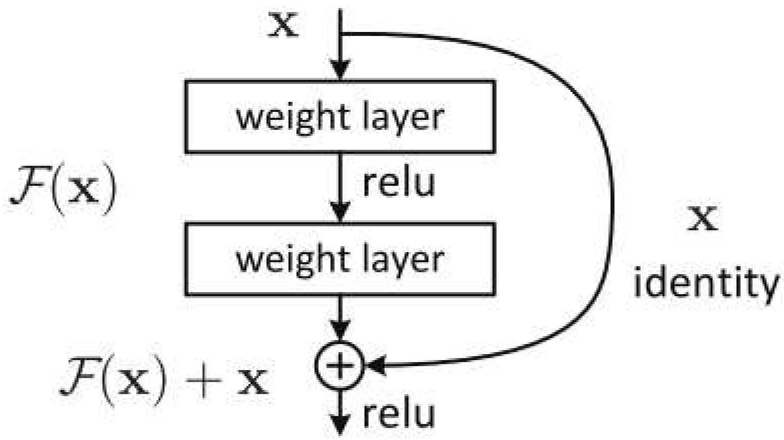

Figure 3.1.: Residual Block

This architecture is used for detection of cough and non-coughs sounds. The input given was 432 x 288. The model used is ResNet50 which consists of 5 stages each with a convolution and Identity block. Each convolution block has 3 convolution layers and each identity block also has 3 convolution layers. The ResNet-50 has over 23 million trainable parameters.

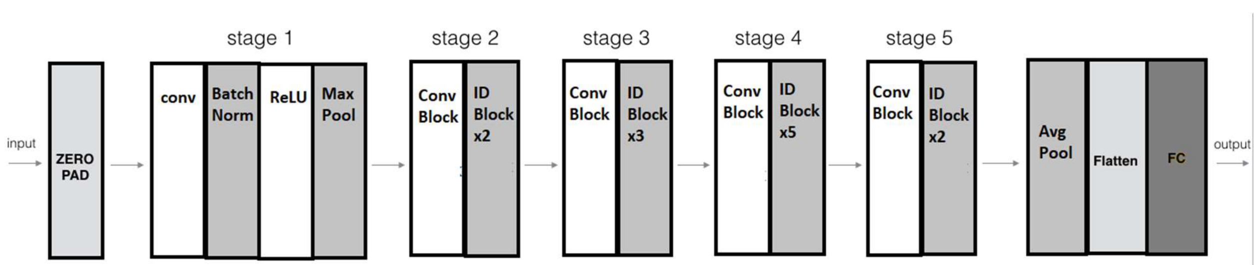

Figure 3.2.: ResNet 50 Model

There are two fully connected layers with the drop out of 50% and with the SoftMax activation function.

The input is then given to this model which outputs the binary value, whether a sound file is a cough or not. There are a total of 30 epochs and the learning rate was 0.001 and with the batch size of 32 while using binary_crossentropy as the loss function and adam as the optimizer.

3.2. XGBoost

One of the models used for diagnosis in this work is XGBoost. It is one of the dominating algorithms in the field of machine learning for tabular data. XGBoost stands for eXtreme Gradient Boosting. XGBoost is a decision-tree-based ensemble Machine Learning algorithm that uses a gradient boosting framework. When it comes to the small tabular data, XGBoost is dominant in this field of machine learning. Gradient boosting as the name suggests it is Gradient blended with Boosting.

In boosting, the individual models are not built on completely random subsets of data and features but sequentially by adding more weight on instances with wrong predictions and high errors. The general idea behind this is that instances, which are hard to predict correctly is focused during learning, so that the model learns from the past mistakes. When the model is trained on each ensemble on a subset of the training set, it is called Stochastic Gradient Boosting, which can help improve generalizability of the model.

The gradient is used to minimize a loss function. In each round of training, the weak learner is built and its predictions are compared to the correct outcome that we expect. The distance between prediction and true value represents the error rate of the model. These errors can now be used to calculate the gradient. The gradient is the derivative of the loss function - so it describes the steepness of the error function. The gradient can be used to find the direction in which to change the model parameters in order to reduce the error in the next round of training by “descending the gradient”. Here, the image data is flattened to make it a single directional array and then is fed to the XGBoost algorithm.

The flattened input is given to this model which outputs the categorical value, if it is one of the four diseases or normal. The model then was trained and fitted and the model is tested on the test set that gave an average of 97% accuracy, mode details in the results section 4.1.

3.3. Convolutional Neural Networks

Convolutional Neural Networks (CNN) is heavily used in the field of Computer Vision. As with every neural network it has many hidden layers and these typically consist of convolutional layers, pooling layers, fully connected layers, and normalization layers. Here instead of using the normal activation functions as in multi-layer perceptron models, CNN uses convolution and pooling functions as activation functions.

In Convolutional Layer, a weight matrix is defined and it extracts certain features from the image. Let’s say, we have initialized the weight of a matrix 3*3, this weight matrix will run across the image such that all the pixels are at least covered once, to give a convolved output. This weight can perform different functions like, extracting edges, corners, colors, etc. The weights are learnt such that the loss function is minimized. In the convolutional layered network, one layer might extract the basic features and the others might extract more complex features from the image. The stride which is the movement of the weight matrix on the image can be determined as a hyperparameter. As the stride value increases the size of the image decreases. Padding reduces the problem of decreased image size.

In Pooling Layer, the number of the trainable parameters are reduced. This is used mostly when we have very large images and wanted to reduce the size of the images. Pooling is done independently on each depth dimension, therefore the depth of the image is not changed. The most commonly used one is called max pooling.

The output layer gives the output in the form of a class. To generate a final output there has to be a fully connected layer to generate an output equal to the number of classes. The output has a loss function like categorical cross-entropy, to compute the error in prediction. Once the forward pass is complete the backpropagation begins to update the weight and biases for error and loss reduction.

The input is then given to this model which outputs the categorical value, whether if it is one of the four diseases or normal and the batch size for training was 8, the learning rate is set to 0.001 and the maximum number of epochs is 15, while using categorical_crossentropy as the loss function and adam as the optimizer.

CHAPTER 4

Performance Analysis

4.1. Results

In this section, the results of the models of this work are presented. For the detection of cough, the accuracy of the model 0.99 and the loss is 0.03. The number of epochs were 8, early stopping is used to avoid overfitting. The figures are shown in fig. 5 and 6.

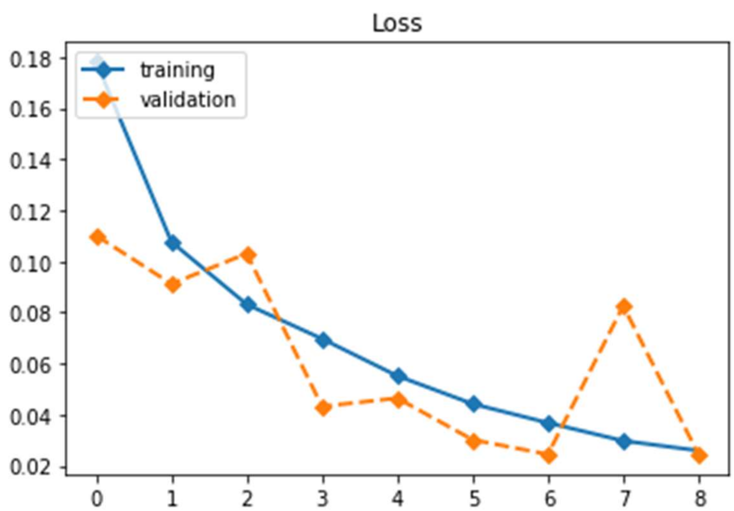

Figure 4.1.: Loss of ResNet 50 Model

The loss and accuracy for the detection model are useful in the real time scenarios, which is yet to be tested practically. This model performs better than CNN, which has an accuracy of 90.08%.

For the diagnosis models the CNN, Resnet and XGBoost almost performed the same and gave the same accuracy with XGBoost having a slight more accuracy than the rest. Repeatedly the results are able to be replicated without a lot of margin.

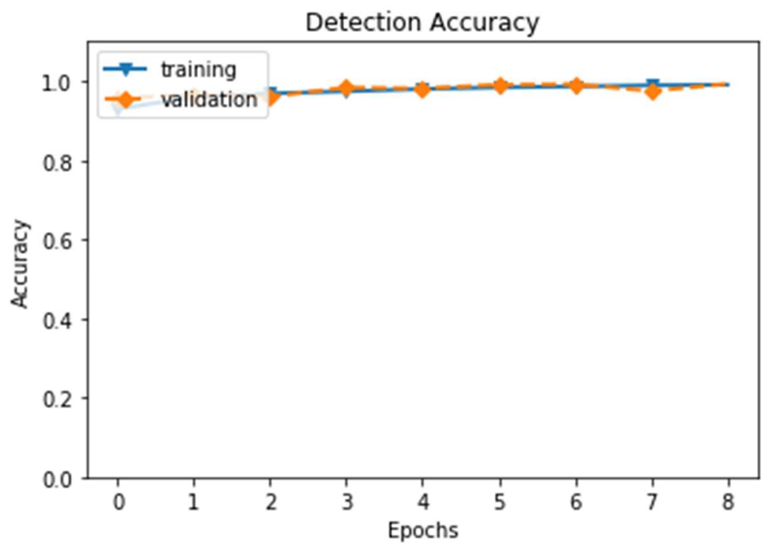

Figure 4.2.: Accuracy of ResNet 50 Model

XGBoost performed better than the neural network models CNN and ResNet see Fig. 8, Fig. 9 and Table 1. The formulas for calculating the metrics are mentioned below, the Support indicates the number of the test set. One of the reasons might be that the data required for the deep networks is very large and this dataset is relatively very small when compared and XGBoost performs well with the smaller data. Getting more data might increase the prediction probability of the neural networks.

	Precision	**Recall**	**F-1 Score**	**Support**
Bronchitis	0.95	0.95	0.95	35
Normal	0.98	0.98	0.98	99
Pertussis	0.93	1.00	0.96	52
Bronchiolitis	1.00	0.86	0.92	14
Asthma	1.00	0.87	0.93	15

Table 4.1.: Classification Report for XGBoost

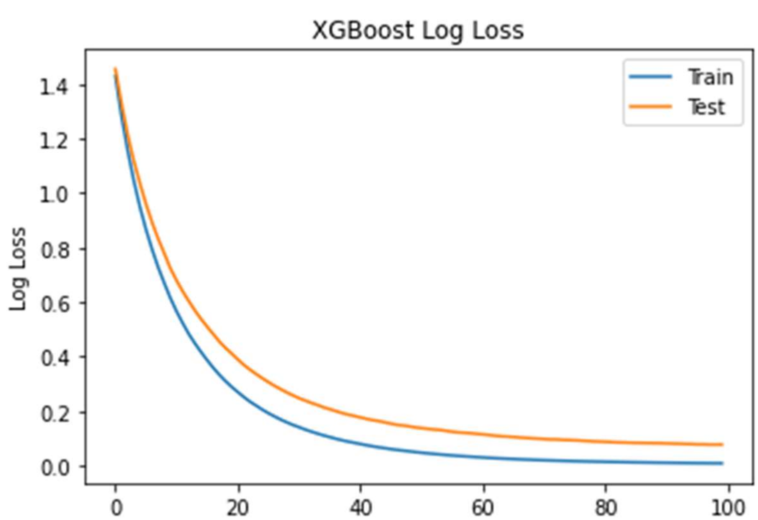

Figure 4.3.: XGBoost Log Loss

The test loss function of the XGBoost performed well when compared to the training set, with a test log loss of 0.2% and classification error of 0.045. This performance is good for the testing data and there is a very high chance for improvement by training on more data using the ResNet.

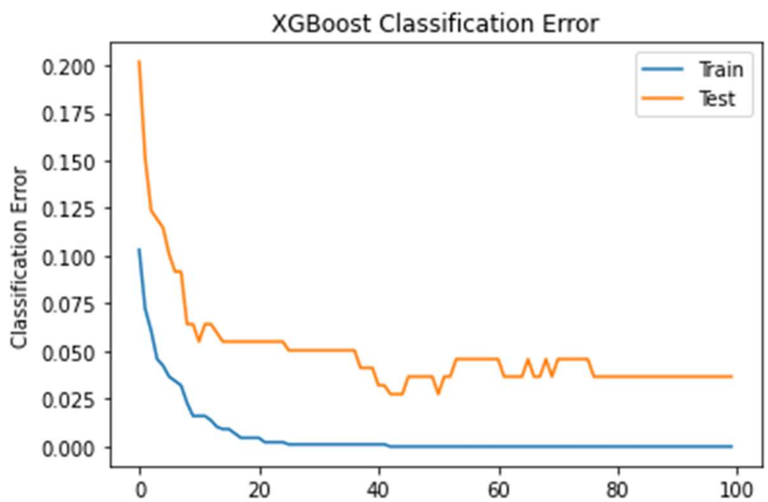

Figure 4.4.: XGBoost Classification Error

These are the formulas to calculate the metrics:

$$Precision = \frac{True\ Positive}{\text{True Positive + False Positive}}$$

$$Recall = \frac{True\ Positive}{\text{True Positive + False Negative}}$$

$$F1 \ = 2 * \frac{\text{Precision} * \text{Recall}}{\text{Precision + Recall}}$$

CHAPTER 5

Conclusion and Future Work

5.1. Conclusion

Cough detection is one of the most needed step in order to diagnose a disease. By using ResNet, this model are able to achieve an accuracy of 99% which outperforms many of the models that are in existence [17],[18] using just a cough sound without extra equipment to be worn.

The diagnosis part of this work has good results also in diagnosing different kinds of coughs with an average accuracy of 97% which makes this work unique and is unlike existing solutions in literature. For this solution to be deployable, it has to be tested on more data and has to be proven to maintain the same accuracy that is now with this small dataset. For now, this solution cannot be relied to be deployed as no clinical trials are done so far.

5.2. Future Work

For future works there is a more need of medical labeled data to carry on this work. Furthermore, the ResNet model can be applied even for diagnosis.

This work can be continued with other techniques for other diseases or other techniques for the same diseases, using speech or breath sounds. It can also be combined with other diseases are disabilities that can help people with speech or hearing disabilities. This work can be further carried on for counting cough sounds, which give a lot of information on the severity of a disease that can be used to measure the health of a person.

References

1. J. Korpáš, J. Sadloňová, and M. Vrabec, "Analysis of the cough sound: an overview," Pulmonary Pharmacology, vol. 9, no. 5-6, pp. 261–268, 1996

2. Immunization, Vaccines and Biologicals – Pertussis. Accessed on May 6, 2020. [Online]. Available: https://www.who.int/immunization/diseases/pertussis/en/

3. Martinek J, Tatar M, Javorka M. Distinction between voluntary cough sound and speech in volunteers by spectral and complexity analysis. J Physiol Pharmacol. 2008; 59(6):433–440

4. Barry SJ, Dane AD, Morice AH, Walmsley AD. The automatic recognition and counting of cough. Cough. 2006

5. Tracey BH, Comina G, Larson S, Bravard M, Lopez JW, Gilman RH. Cough detection algorithm for monitoring patient recovery from pulmonary tuberculosis. In: IEEE EMBC; 2011. p. 6017–6020.

6. Swarnkar V, Abeyratne UR, Amrulloh Y, Hukins C, Triasih R, Setyati A. Neural network based algorithm for automatic identification of cough sounds. In: IEEE EMBC; 2013. p. 1764–1767.

7. Amrulloh YA, Abeyratne UR, Swarnkar V, Triasih R, Setyati A. Automatic cough segmentation from non-contact sound recordings in pediatric wards. Biomed Signal Process Control. 2015; 21:126–136

8. Matos S, Birring SS, Pavord ID, Evans H. Detection of cough signals in continuous audio recordings using Hidden Markov Models. IEEE Trans Biomed Eng. 2006; 53(6):1078–1083

9. Liu JM, You M, Li GZ, Wang Z, Xu X, Qiu Z, et al. Cough signal recognition with Gammatone Cepstral Coefficients. In: 2013 IEEE China Summit and International Conference on Signal and Information Processing (ChinaSIP). IEEE; 2013. p. 160–164.

10. Lucio C, Teixeira C, Henriques J, de Carvalho P, Paiva RP. Voluntary cough detection by internal sound analysis. In: Biomedical Engineering and Informatics (BMEI), 2014 7th International Conference on; 2014. p. 405–409.

11. Larson EC, Lee T, Liu S, Rosenfeld M, Patel SN. Accurate and privacy preserving cough sensing using a low-cost microphone. In: 13th International Conference on Ubiquitous Computing, UbiComp'11 and the Co-located Workshops. Beijing, China; 2011. p. 375–384.

12. Chatrzarrin H, Arcelus A, Goubran R, Knoefel F. Feature extraction for the differentiation of dry and wet cough sounds. In: Medical Measurements and Applications Proceedings (MeMeA), 2011 IEEE International Workshop on. IEEE; 2011. p. 162–166.

13. Swarnkar V, Abeyratne U, Chang A, Amrulloh Y, Setyati A, Triasih R. Automatic Identification of Wet and Dry Cough in Pediatric Patients with Respiratory Diseases. Ann Biomed Eng. 2013; 41(5):1016–1028

14. Kosasih K, Abeyratne UR, Swarnkar V, Triasih R. Wavelet Augmented Cough Analysis for Rapid Childhood Pneumonia Diagnosis. IEEE Trans Biomed Eng. 2014; 62(4):1185–1194

15. Parker D, Picone J, Harati A, Lu S, Jenkyns MH, Polgreen PM. Detecting paroxysmal coughing from pertussis cases using voice recognition technology. PloS one. 2013; 8(12):e82971

16. Renard Xaviero Adhi Pramono*, Syed Anas Imtiaz, Esther Rodriguez-Villegas, "A Cough-Based Algorithm for Automatic Diagnosis of Pertussis"

17. Teyhouee A., Osgood N.D. (2019) Cough Detection Using Hidden Markov Models. In: Thomson R., Bisgin H., Dancy C., Hyder A. (eds) Social, Cultural, and Behavioral Modeling. SBP-BRiMS 2019. Lecture Notes in Computer Science, vol 11549. Springer, Cham

18. J. Monge-Álvarez, C. Hoyos-Barceló, L. M. San-José-Revuelta and P. Casaseca-de-la-Higuera, "A Machine Hearing System for Robust Cough Detection Based on a High-Level Representation of Band-Specific Audio Features," in *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 8, pp. 2319-2330, Aug. 2019

19. Apple Watch wristband sensor claims to detect potassium in your blood — without needles. Accessed on May 6, 2020. [Online]. Available: https://www.theverge.com/2018/3/12/17109036/apple-watch-kardiaband-alivecor-ekg-reader-hyperkalemia-health-artificial-intelligence