

Automated Classification of Adult Cough Recordings into Wet vs. Dry with Applications to Pulmonary Disease Assessment

by

Jiayao Clara Lu

Supervised by [Shwetak N. Patel (Chair) and Elliot Saba]

A senior thesis submitted in partial fulfillment of
the requirements for the degree of

Bachelor of Science
With Departmental Honors

Computer Science & Engineering

University of Washington

December 2017

Presentation of work given on _____ December 7th, 2017 _____

Thesis and presentation approved by _____ Prof. Shwetak Patel _____

Date _____ December 13th, 2017 _____

University of Washington

Abstract

Automated Classification of Adult Cough Recordings into Wet vs. Dry with Applications
to Pulmonary Disease Assessment

by

Jiayao Clara Lu

Chair of the Supervisory Committee:

Professor Shwetak N. Patel

Paul G. Allen School of Computer Science & Engineering

Coughing is a common symptom of pulmonary ailments, and serves as a valuable measure when quantifying pulmonary health. We propose the research and development of a set of techniques to enable physicians to investigate pulmonary health through cough sounds. A variety of signal processing and machine learning approaches are proposed, each with various performance and usability tradeoffs, and we perform a careful analysis of the benefits and drawbacks of each separate approach, opening up new avenues of pulmonary health research in a variety of applications. We propose the application of these techniques toward the classification of coughs.

This thesis is specifically focused on our process and findings when we applied our techniques to distinguish between adult wet cough and dry cough sounds.

TABLE OF CONTENTS

	Page
Chapter 1: Introduction	1
Chapter 2: Background and Related Work	3
Chapter 3: Method	7
3.1 Data Acquisition and Pre-processing	7
3.2 Feature Extraction	8
3.3 Model Training and Testing	12
Chapter 4: Results	13
Chapter 5: Discussion	16
5.1 Lessons Learned	16
5.2 Future Work	17
Chapter 6: Conclusion	19
Chapter 7: Acknowledgement	21
Bibliography	22

Chapter 1

INTRODUCTION

Pulmonary diseases afflict over 1 billion people worldwide, especially those in less developed countries [3, 11]. Four out of the top ten causes of death globally are due to pulmonary disease [11]. Coughing is one of the earliest and most common symptoms across various pulmonary diseases, such as asthma, tuberculosis, pneumonia, and bronchitis [6]. Symptom tracking, including cough tracking, is a crucial part of the health care assessment process [7, 4]. Traditionally, clinicians rely on caregivers or patients themselves to track and report symptom history, or request patients to stay at the facility for hours for cough monitoring, while having personnel to manually annotate. The former approach can be burdening or even impossible, especially for patients suffering chronic pulmonary disease; the latter is time and cost inefficient. Also, due to time constraints of each hospital visit, cough characteristics that can be directly observed by a clinician during an appointment are very limited. Furthermore, the most substantive coughing may occur at unpredictable times, which may not be captured during a visit. Hence, technology that supports continuous, automatic cough analysis can reduce the cost and burden of manual tracking, increase the patient’s quality of life, and provide a more accurate and complete symptom history for health care assessment.

Previous attempts to automate cough monitoring primarily focused on cough frequency [9, 4, 5]. Beyond cough frequency, other characteristics of cough sounds, such as “dryness” and “wetness”, can also reveal valuable information about the state of the respiratory tract for health care assessment. In fact, dryness and wetness of the cough is commonly referenced during subjective assessments conducted by health care professionals globally [8]. A wet (or productive) cough refers to a cough that contains acoustic indication of mucus in the airway; likewise, a dry (or unproductive) cough refers to one that lacks identifiable wetness. The

presence or absence of mucus in the respiratory tract allows health care providers to quickly rule out large classes of diseases during the diagnostic process. Furthermore, once a certain disease is diagnosed, this dryness and wetness information can indicate the progression of the condition and effectiveness of current treatment.

While there have been a couple of preliminary studies in automated classification of cough sounds into dry vs. wet [2, 10], previous work either had limited training and testing dataset sizes, limited patient demographic diversity or limited model performance, as measured by sensitivity and specificity. The limited number of preliminary studies we observed as well as the limitations in prior works highlight the inherent barriers of studying audio-based cough classification. Specifically high quality data is difficult to obtain. For a dataset to be useful, it must be sufficiently large with a sufficient amount of patient diversity. Patient diversity is a necessary criterion of the dataset to ensure the algorithm doesn't learn extraneous features beyond the wetness and dryness of coughs. For example, if all patients in the "wet" group are males, while all patients in the "dry" group are females, then the algorithm may learn to distinguish wet and dry coughs based on the pitch of the coughs due to population differences in vocal tract size. Furthermore, supervised training models require high quality labels. This requires many hours of health care experts' time to listen and manually annotate types of recorded coughs. Even if these barriers are overcome, knowing what relevant features embedded in a cough can be utilized for training is difficult.

In this thesis we will discuss our research process and findings when we applied our techniques to distinguish between adult wet cough and dry cough sounds. We hope that this can provide insights for future researchers who are interested in conducting similar cough sound classification works.

Chapter 2

BACKGROUND AND RELATED WORK

There exists a wide range of cough sound variations across different pulmonary conditions [6]. Therefore, there has been rising research interest in understanding what can be inferred about a given pulmonary condition from cough sounds, with growth potential in clinical applications, particularly the screening or diagnostic process. A cough consists of the following three phases in the given order: inspiration, compression, and expulsion, as shown in Figure 2.1 [6]. Since cough sounds mainly capture the expulsion phase, most cough recording studies restrict their analysis to the last phase only, which still contains broad varieties of cough sounds, as visualized in Figure 2.2 [6], and Figure 2.3 [1]. Our work will also focus on the cough expulsion phase, as it is the only phase with significant audible data.

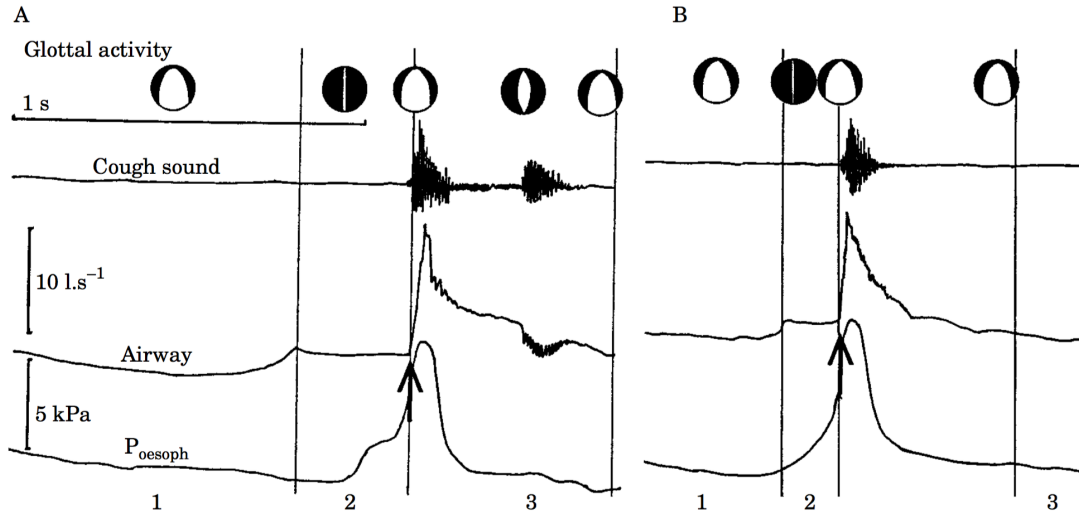


Figure 2.1: From Korpás et al. [6], showing glottal position, recorded sound, airflow and pressure within the esophagus for two example coughs, segmented into three time regions.

Figure 2.2: The variation of sound recordings of the expulsive segment of a cough, from Korpás et al. [6]

During our literature review, we encountered two preliminary studies on dry-wet cough classification with varying approaches and performance.

Chatrzarrin et al. examined cough sound signals from both the temporal and spectral domains [2]. They proposed two features extracted from cough sounds: the number of peaks of the energy envelope of the cough signal, and the power ratio of two frequency bands of the cough signal. Cough data points containing both feature values were then plotted on a two dimensional graph, and a clear clustering pattern separating dry and wet coughs was observed. However, the dataset only contained eight highly dry and eight highly wet cough recordings, undermining the generalizability of the results.

Swarnkar et al. analyzed cough sound signals by training a Logistic Regression Model (LRM) with a set of eight complex mathematical features, including Kurtosis (Kurt) and Mel-frequency cepstral coefficients (MFCC) [10]. The performance of the trained model in terms of sensitivity and specificity were around 55% and 93%, respectively. While the size of the dataset is significantly larger than that of Chatrzarrin et al., the data were mainly

obtained from pediatric patients, with the majority being less than 5 years old. Our work will differ here in that the patients represented in our data will consist primarily of adults.

One observation from this study that is worth mentioning is that the cough recordings were labeled by two health care experts independently, and the agreement rate between the two annotation sets was only 80%. This rate highlights the inherent difficulty in cough dry-wet classification; even human experts cannot reach a consensus in a nontrivial fraction of cases. The commonly accepted resolution for this problem is to only utilize data points for which the majority of experts agree on the labeling, which will be adopted in our work as well.

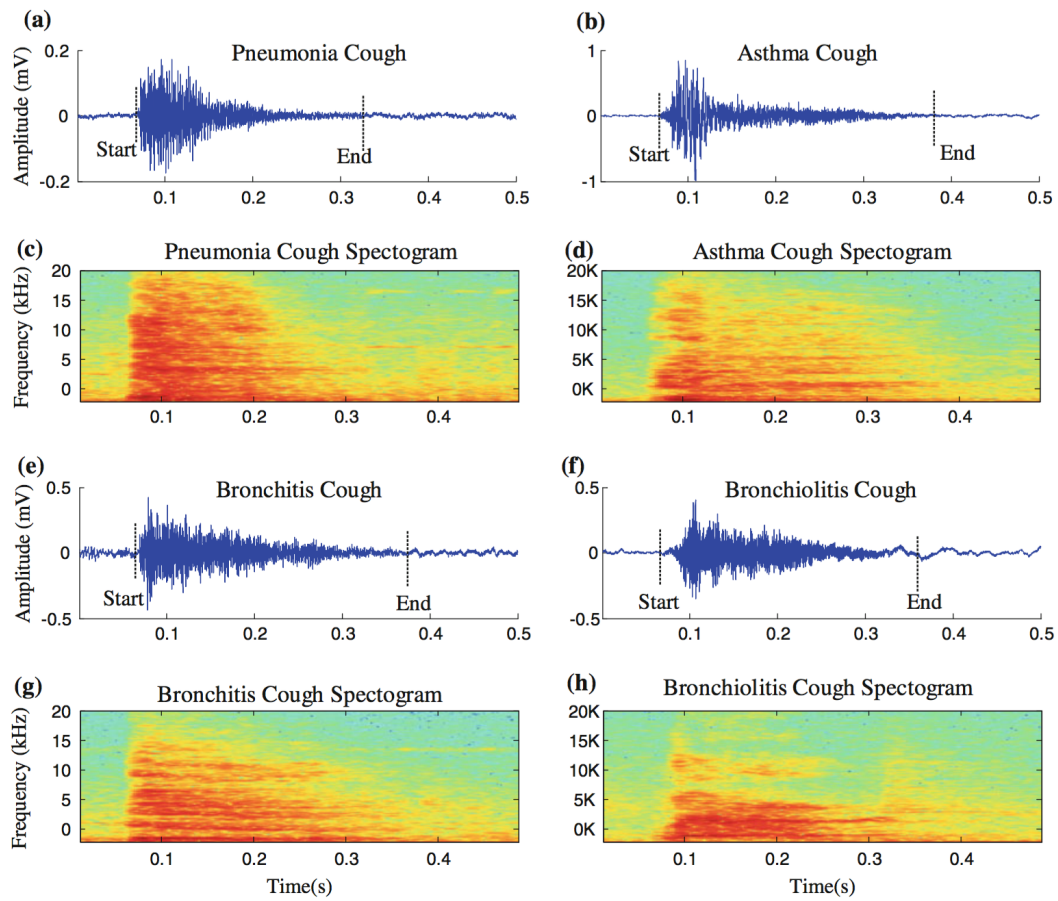


Figure 2.3: Abeyratne et al. display the variation of coughs from patients with different pulmonary ailments [1].

Chapter 3

METHOD

We propose to utilize machine learning and signal processing as our primary techniques to build a classifier to distinguish dry and wet coughs. Our work is divided into three stages: 1) data acquisition and pre-processing, 2) feature extraction, and 3) model training and testing. We will provide more details of each of these three stages in the following sections.

3.1 Data Acquisition and Pre-processing

We collaborated with the international tuberculosis community through the Gates Foundation to obtain thousands of cough recordings from South Africa, which has a large population infected with tuberculosis. Having wet coughs is one of the signature symptoms of tuberculosis. Patients with tuberculosis is therefore an ideal source for us to obtain wet cough recordings. Around 90 subjects were recruited, with healthy controls and tuberculosis patients evenly distributed. Subjects were each placed in the "TARDIS box"-like station where they were provided with some entertainment options such as TV, and their coughs were recorded by a non-contact microphone. Each subject contributed an one-hour-long cough recording.

A group of trained annotators segmented out each cough from the recordings into its own audio file and labeled the cough type as either dry or wet. As mentioned in the previous chapter, some coughs can be ambiguous in terms of dry-wet classification even to trained personnel. For the purposes of this analysis, we only included cough recordings with labels that were agreed upon by the majority of the annotators in our training and testing datasets.

3.2 *Feature Extraction*

We started our exploration for features that can be potentially useful in dry-wet classification by first visually examining the waveform and spectrogram of three highly dry and highly wet cough signals chosen at random. The waveforms and spectrograms are shown in Figures 3.1 and 3.2. Next, we asked questions such as “given our daily experience and these visual representations of dry and wet cough signals, what are the clues that help us humans tell dry and wet coughs apart?” and “how can we quantify these ‘clues’ with signal processing techniques to train machines to perform the same classification?” All of the feature extractions in our work were implemented in Python with a few helpful opensource libraries, including Numpy, Scipy, and PySoundFile. In this section, we will explain each of the features we explored, and some of our hypotheses on how they should behave differently for dry and wet coughs.

3.2.1 Cough Duration

Signal energy is a useful mathematical measure of how much signal is present over time, which we can use by obtaining the positive signal energy envelope to visualize the change of energy levels over time. The peaks of the envelope correspond to high sound energy levels. Under a controlled environment, such as a room with stable ambient noise, these peaks therefore further correspond to the sounds generated by the rapid air expulsion during a subject’s cough, with high probability.

Cough duration is calculated by first locating the time point with the highest energy level, and then extending a time window along the energy envelope on both sides of the peak until it falls below a certain threshold. A set of heuristics incorporates nearby peaks together into a single cough event, resulting in a single duration containing a single peak or multiple neighboring peaks.

As we can observe from the waveforms in the Figure 3.1, the durations of the wet cough samples seem to be slightly longer than those of the dry cough samples. This also coincides with our daily experience in that dry coughs tend to be short and require less effort, while

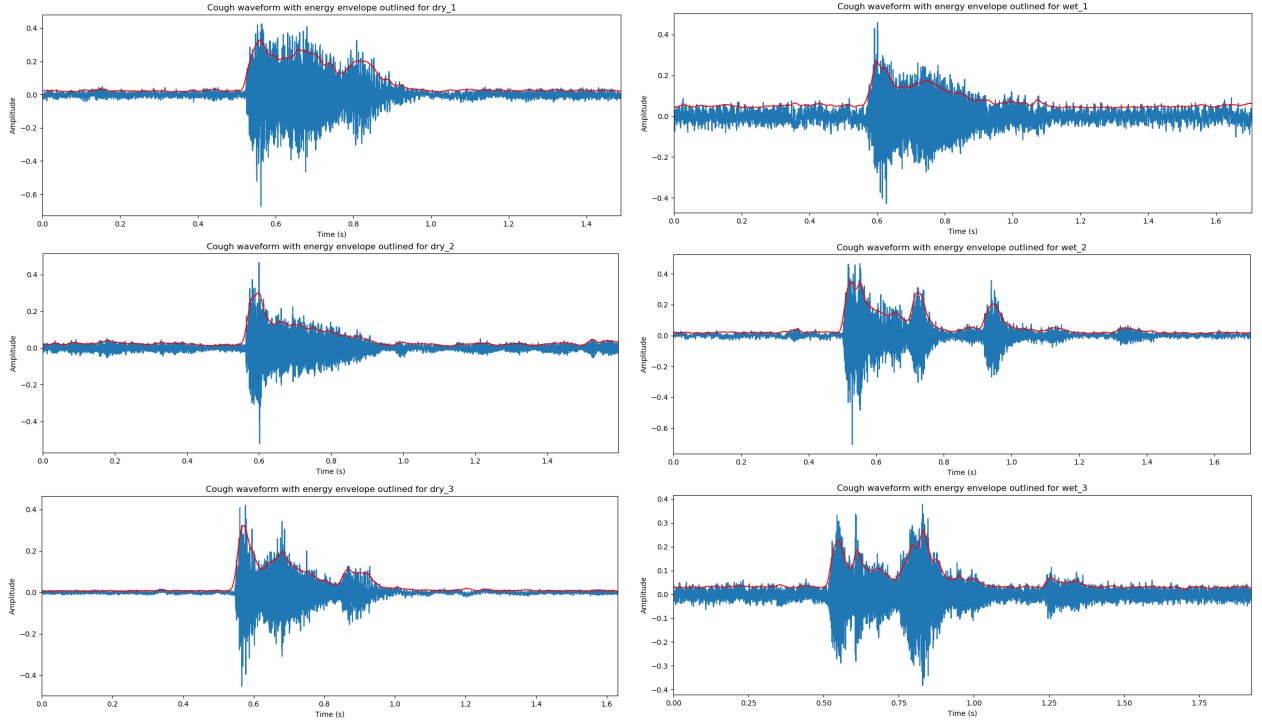


Figure 3.1: The waveforms of cough signals with upper energy envelopes outlined in red present information in the time domain. The x-axis represents time with the y-axis representing amplitude. The left column consists of three highly dry cough signals, and the right column consists of three highly wet cough signals.

wet coughs tend to be longer and require more effort as we try to clear the mucus in the airway. Therefore, we hypothesized that the duration of a wet cough is longer than that of a dry cough.

3.2.2 Peak Frequency

Almost all of the sounds we hear in the real world are not composed of a single frequency and amplitude, but multiple frequencies and amplitudes added together simultaneously. This is the same for cough sounds, as they are generated by vibrations of different parts of the respiratory system at different frequencies and energy levels during the rapid air expulsion.

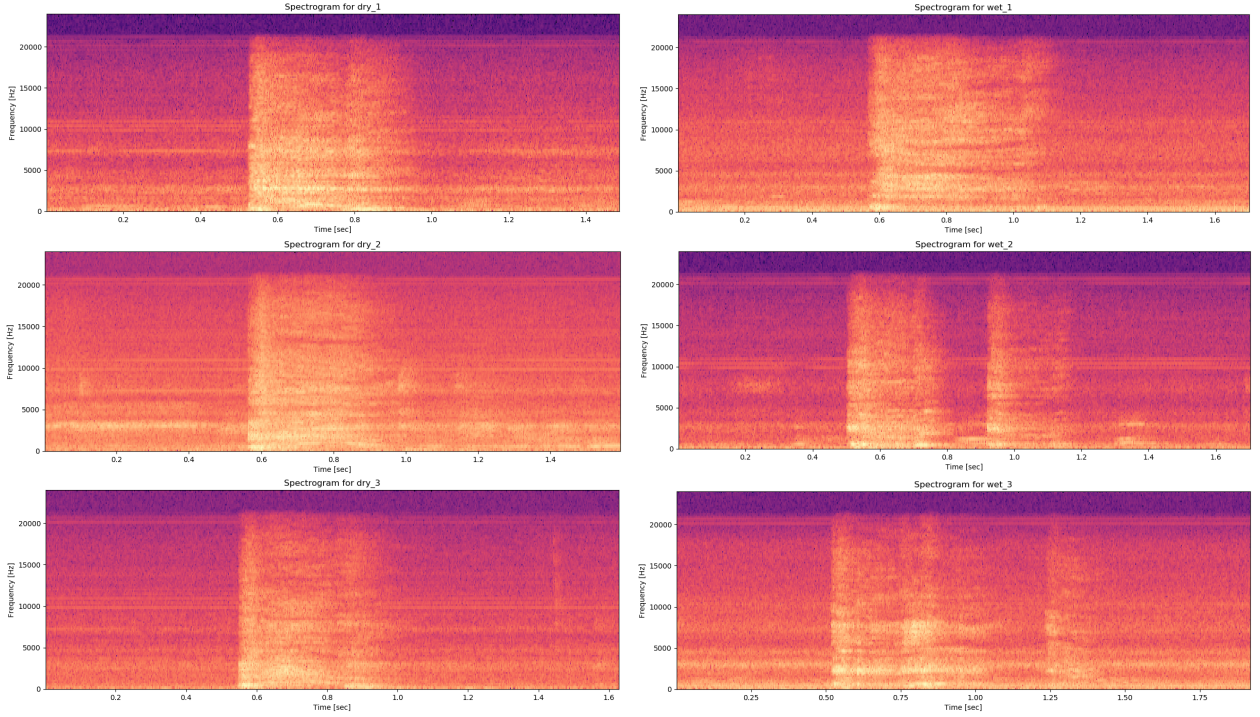


Figure 3.2: The spectrograms of cough signals present informations in both time and frequency domain. The left column consists of three highly dry cough signals, and the right column consists of three highly wet cough signals.

That explains why our cough sound signals (visualized in Figure 3.1 in blue) are “fuzzy” instead of a smooth curve. We can examine the frequencies and corresponding amplitudes’ components of the cough sound in the frequency domain through the Fourier Transform technique.

We define the “peak frequency” to be the single frequency with the highest magnitude (energy). Note this is not the highest frequency, as the name might intuitively suggest. For the same reasons we described in previous subsection, under a controlled environment, we expect that among all the sound components of a cough signal, the sound with the highest amplitude (loudest sound) is from the cough. Because from our daily experience, dry coughs tend to have higher pitch, while wet coughs tend to be more dull, we were interested to see

whether peak frequency would be useful in the dry-wet classification.

3.2.3 Mean Frequency

With similar motivations as the peak frequency, we define the “mean frequency” to be average of the frequencies of a cough signal weighted by their amplitudes, and propose it to be one of feature candidates to be explored.

3.2.4 Mean Energy Levels of the Low, Middle, and High Frequency Bands

The spectrogram is a common tool used to visualize the spectrum of frequencies of sounds, and how energy levels corresponding to each frequency changes over time. The x-axis represents time, the y-axis represents frequency, and the color brightness represents the energy level intensity associated with a specific frequency at a given time (the brighter the color, the stronger the energy is). As we can observe in Figure 3.2, the energy distribution patterns are different between dry and wet coughs. In particular, high frequencies tend to have more energy in dry coughs. Thus, we decided to further explore the mean energy levels of the low, middle, and high frequency “bands” (to be more specific, the bottom, middle, and top thirds of the frequencies spectrum) as three separate features.

3.2.5 Spectral Centroid

The spectral centroid, which describes the “center point” of the sound energy distribution across the frequency spectrum, is a common measure in signal processing associated with the “brightness” of a sound. It provides similar information as the mean frequency, but we included both of them in our feature set, and will explore which one is more informative for our task.

3.2.6 Standard Deviation of Spectral Centroid

For this feature, we are analyzing small chunks of the cough signal at a time, instead of the entire cough all at once. We are particularly interested to see how much the spectral centroid changes over time, given the observation that the color blocks in the dry cough spectrograms seem to be more stable over time. This can be quantified by the standard deviation of spectral centroid.

3.2.7 Average Difference of Short-Time Fourier Transformation

Similar to the regular Fourier Transform, the Short-time Fourier Transform (STFT) shows the frequency components of a sound signal and corresponding amplitudes. However, instead of processing the entire sound signal all at once, STFT processes the sound signal by taking the Fourier Transform of smaller sliding windows separately, which can provide us a better picture of the changes of the energy level corresponding to each frequency over time. Similar to the motivations for exploring the standard deviation of the spectral centroid above, we are interested in how “stable” the energy levels are across time.

3.3 *Model Training and Testing*

We will partition our dataset into a training set (80%) and a testing set (20%), with the same proportion of dry and wet data points, followed by extracting all the features mentioned above from both of the datasets. We will then feed the normalized feature vectors extracted from the training data and corresponding labels into several machine learning models implemented by scikit-learn, a Python library. We specifically explore Random Forest, K-Nearest Neighbors, and several linear models, and compare their performances with cross validation using the testing dataset. Once we narrow down the model with the best performance, we will further examine what feature subset will produce the best result.

Chapter 4

RESULTS

Due to the difficulty of acquiring a large amount of data initially, we conducted a preliminary evaluation with the small dataset we had that were manually segmented and labeled by a group of trained annotators. This initial dataset was obtained from 3 female and 3 male adult tuberculosis patients, with ages between 21 to 54 years. The dataset contains 291 cough events in total, out of which 52 were labeled as wet, 151 were labeled as dry, and the rest were ambiguous (and therefore ruled out from our training and testing sets). The results we obtained when we fed our initial dataset into various machine learning model candidates are shown in Table 4.1. Each classifier’s performance is evaluated by the classifier’s accuracy, sensitivity (true positive rate), and specificity (true negative rate) with respect to the testing dataset.

<i>Model</i>	<i>Results</i>
Random Forest	Sensitivity: 90.00% Specificity: 66.67% Accuracy: 83.33%
K Neighbors	Sensitivity: 77.42% Specificity: 36.36 Accuracy: 66.67%
Support Vector Machines (SVC)	Sensitivity: 75.61% Specificity: 100% Accuracy: 76.19%
Ridge Regression	Sensitivity: 75.61% Specificity: 100% Accuracy: 76.19%
Linear Regression	Sensitivity: 75.61% Specificity: 100% Accuracy: 76.19%
Logistic Regression	Sensitivity: 76.92% Specificity: 66.66% Accuracy: 76.19%

Table 4.1: Performances of different machine learning models.

Since the Random Forest model yielded the best overall performance, we decided to further explore the optimal subset of features with it. We enumerated all possible nonempty subsets of the nine features we extracted, and the results are listed in Table 4.2.

<i>Size of subset</i>	<i>Optimal feature combination</i>	<i>Results</i>
2	mean frequency, peak frequency	Sensitivity: 87.50% Specificity: 70% Accuracy: 83.33%
3	cough duration, mean frequency, peak frequency	Sensitivity: 88.24% Specificity: 87.50% Accuracy: 88.10%
4	cough duration, mean frequency, peak frequency, standard deviation of spectral centroid	Sensitivity: 93.55% Specificity: 81.82% Accuracy: 90.48%
5	cough duration, mean frequency, peak frequency, mean energy of middle frequency band, average difference of STFT	Sensitivity: 96.55% Specificity: 76.92% Accuracy: 90.47%
6	cough duration, mean frequency, peak frequency, mean energy of middle frequency band, spectral centroid, average difference of STFT	Sensitivity: 93.94% Specificity: 100% Accuracy: 95.24%
7	cough duration, mean frequency, peak frequency, mean energy of high frequency band, spectral centroid, average difference of STFT, standard deviation of spectral centroid	Sensitivity: 96.43% Specificity: 71.43% Accuracy: 88.10%
8	all nine features except peak frequency	Sensitivity: 83.33% Specificity: 83.33% Accuracy: 83.33%
9	all nine features	Sensitivity: 90% Specificity: 66.67% Accuracy: 83.33%

Table 4.2: The most optimal combination of features for each size (2) of subset, and their performances.

Chapter 5

DISCUSSION

5.1 *Lessons Learned*

In this thesis we proposed a method to automatically classify cough sounds into wet vs. dry. To the best of our knowledge, this is the first attempt to automate classification of adult dry and wet cough sounds. We examined a novel set of nine features for the the application of distinguishing wet and dry coughs. These features have the advantage of being backed by straightforward human intuitions, as compared to the complex mathematical features proposed in previous work [10], without sacrificing the quality of our model. In fact, the classifier we trained on our initial dataset yielded a more promising result compared to those obtained in previous works, as we can see in Table 5.1.

<i>Studies</i>	<i>Results</i>	<i>Dataset</i>	<i>Method</i>
Chatrzarrin et al. [2]	Sensitivity: 100% Specificity: 100%	8 (unclear age)	two features: number of peaks of the energy envelope, power ratio of two frequency bands
Swarnkar et al. [10]	Sensitivity: 55% Specificity: 93%	536 (from children)	a set of nine complex mathematical features with an LRM
Our work	Sensitivity: 93.94% Specificity: 100%	291 (from adults)	a novel set of features with a Random Forest Model

Table 5.1: Comparison between our work and priors work on similar wet-dry cough classification.

From the results shown in Table 4.2, we learned that the subset of features we examined

that yielded optimal results using the Random Forest Model (RFM) contains:

- 1) cough duration
- 2) mean frequency
- 3) peak frequency
- 4) mean energy of middle frequency band
- 5) spectral centroid
- 6) average difference of STFT

Moreover, using only the mean and peak frequency features with RFM already yielded a sensitivity of 87.50% and specificity of 70%, a surprisingly promising result given the little computation required. Such few features provides easier implementation with a “reasonable” performance. For resource-constrained devices in technologically and medically underdeveloped communities, this can be particularly useful.

5.2 Future Work

- 1) Due the time constraints of this research project, we were not able to segment out single cough events from all of the cough recordings we had access to. This is why we presented results built on only a small portion of it (291 cough events from 6 subjects). We would like to further validate the performance of the feature set and model on a larger dataset when all cough segmentations are complete.
- 2) Since our samples were all collected from South Africa, validating our model on a second dataset collected from a different location in the future can ensure that the dry-wet classifier we trained is not biased based on certain location-dependent factors.

- 3) Since we only examined a set of models whose theory we sufficiently understand, future work can explore other machine learning models as well. One particularly interesting direction is the application of recent deep learning methods to our problem. We did not attempt this due to the small amount of data we had relative to what is typically needed to train a nontrivial-sized network.
- 6) It is an interesting question to ask what are the specificity-sensitivity tradeoffs in using these features and models. Depending on the medical context, one may be more important than the other.
- 5) Future work can also focus on the deployment of this classifier to portable devices, such as mobile phones or Amazon Echoes, that can be used in medical screening or monitoring.
- 6) Finally, we are interested to explore automated classification for directly distinguishing between tuberculosis vs. non-tuberculosis using our dataset and the features we examined, which can be especially impactful for tuberculosis screening in communities with poor medical resources.

Chapter 6

CONCLUSION

To our best knowledge, our work is the first attempt in automated classification of adult cough sounds into dry vs. wet. We acquired one-hour-long cough recordings from 90 South African subjects, with healthy controls and TB patients evenly distributed. The initial dataset we used for study was segmented and annotated by a team of trained annotators. It contained 291 cough events in total, out of which 52 were labeled wet, 151 were labeled dry, and the rest were unclear due to the lack of consensus. Only coughs with dry and wet labels were utilized in our training and testing. We then extracted nine features from those cough events, and fed them into various machine learning models, out of which RFM yielded the best result. Finally, we explored the performances of all subsets of the features. We learned that RFM achieved the best performance with the following features: 1) cough duration, 2) mean frequency, 3) peak frequency, 4) mean energy of middle frequency band, 5) spectral centroid, and 6) average difference of STFT. For this model and feature combination with our initial dataset containing cough samples from 6 subjects, we achieved sensitivity 93.94% and specificity 100%. Future work can focus on:

- 1) validation of our model on a larger dataset
- 2) validation of our model on a dataset collected outside South Africa
- 3) exploration of different machine models and deep learning technology
- 4) exploration of specificity-sensitivity tradeoffs in various medical contexts
- 5) deployment to portable devices in actual medical screening and monitoring

6) extension to direct classification of tuberculosis vs. non-tuberculosis

Chapter 7

ACKNOWLEDGEMENT

I would like give my sincere thanks to Elliot Saba for his mentorship and guidance on this project, and kind encouragement and support both inside and outside the academic environment in the past year. Thank you professor Shwetak Patel for giving me the opportunity to work on such an amazing project combining my interests in both computer science and health care, and for access to great resources in the lab. Finally, I want to thank every wonderful member of the UW Ubiquitous Computing lab for their friendship and inspiration.

BIBLIOGRAPHY

- [1] Udantha R. Abeyratne, Vinayak Swarnkar, Amalia Setyati, and Rina Triasih. Cough Sound Analysis Can Rapidly Diagnose Childhood Pneumonia. *Annals of Biomedical Engineering*, 41(11):2448–2462, November 2013.
- [2] H. Chatrarrin, A. Arcelus, R. Goubran, and F. Knoefel. Feature extraction for the differentiation of dry and wet cough sounds. In *2011 IEEE International Symposium on Medical Measurements and Applications*, pages 162–166, May 2011.
- [3] Donald Enarson. *Respiratory diseases in the world: realities of today - opportunities for tomorrow*. European Respiratory Society, Sheffield, 2013. OCLC: 868996973.
- [4] J. Y. Hsu, R. A. Stone, R. B. Logan-Sinclair, M. Worsdell, C. M. Busst, and K. F. Chung. Coughing frequency in patients with persistent cough: assessment using a 24 hour ambulatory recorder. *European Respiratory Journal*, 7(7):1246–1253, July 1994.
- [5] Richard S. Irwin, Michael H. Baumann, Donald C. Bolser, Louis-Philippe Boulet, Sidney S. Braman, Christopher E. Brightling, Kevin K. Brown, Brendan J. Canning, Anne B. Chang, Peter V. Dicpinigaitis, Ron Eccles, W. Brendle Glomb, Larry B. Goldstein, LeRoy M. Graham, Frederick E. Hargreave, Paul A. Kvale, Sandra Zelman Lewis, F. Dennis McCool, Douglas C. McCrory, Udaya B. S. Prakash, Melvin R. Pratter, Mark J. Rosen, Edward Schulman, John Jay Shannon, Carol Smith Hammond, Susan M. Tarlo, and American College of Chest Physicians (ACCP). Diagnosis and management of cough executive summary: ACCP evidence-based clinical practice guidelines. *Chest*, 129(1 Suppl):1S–23S, January 2006.
- [6] J. Korpáš, J. Sadloňová, and M. Vrabec. Analysis of the cough sound: an overview. *Pulmonary pharmacology*, 9(5):261–268, 1996.
- [7] Robert G. Loudon and Linda C. Brown. Cough Frequency in Patients with Respiratory Disease. *American Review of Respiratory Disease*, 96(6):1137–1143, December 1967.
- [8] Jaclyn A Smith, H Louise Ashurst, Sandy Jack, Ashley A Woodcock, and John E Earis. The description of cough sounds by healthcare professionals. 2006.
- [9] Arietta Spinou and Surinder S. Birring. An update on measurement and monitoring of cough: what are the important study endpoints? *Journal of Thoracic Disease*, 6(7):S728–S734, October 2014.

- [10] Vinayak Swarnkar, Udantha R. Abeyratne, Anne B. Chang, Yusuf A. Amrulloh, Amalia Setyati, and Rina Triasih. Automatic Identification of Wet and Dry Cough in Pediatric Patients with Respiratory Diseases. **Annals of Biomedical Engineering**, 41(5):1016–1028, May 2013.
- [11] World Health Organization. The top 10 causes of death, January 2017.