Measuring Tuberculosis Treatment Response: Cough Analysis and Presentation

By Tara Watson, ECE '16

Cough analysis can be used in resource-limited areas to determine, without requiring a laboratory, if patients are recovering from tuberculosis. The focus of the Purple Team's project is to use an accelerometer and microphone to collect patient data that is analyzed using an algorithm to detect cough rate of patients.

Background

The Importance of Cough Rate Detection

Tuberculosis (TB) is one of the world's deadliest infectious diseases. Basic TB is treatable, and has been for many years; however, recently multi-drug resistant (MDR) strains have emerged. These new forms of TB are more difficult to diagnose and treat due to their drug resistance.

The Purple Team's goal is to develop a low-cost ambulatory cough analysis system for areas of the world that do not have access to typical methods of treatment. This system will help identify if patients are responding to tuberculosis treatment sooner, thereby allowing for an earlier diagnosis of MDR TB if applicable. The project can be broken down into two main parts – collecting the data and analyzing and displaying it. Emily Gill wrote about how to collect the data, and this Tech Note discusses how to analyze and display it.

Identifying a Cough

An example of a cough is seen in Figure 1. The initial spike in energy indicates the beginning of the cough, and the cough continued throughout the duration of time the graph was being recorded. For our purposes, cough count refers to the number of coughs per hour and a coughing episode or bout is defined as one or more cough events closely spaced in time. An epoch is a series of coughs where the end of one event is separated from the beginning of the next by less than two seconds.

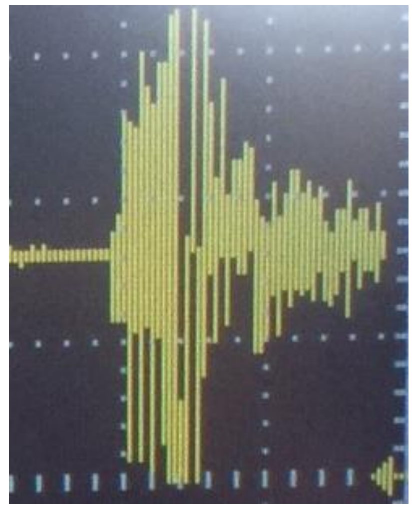

Figure 1. Cough on an oscilloscope

How We Analyze and Display Cough

Detecting an Event

To detect cough events, we are using an accelerometer and microphone to collect patient data that is analyzed using an algorithm to detect cough rate of patients. The initial step of the algorithm is event detection. An event is detected by checking the input signal to see if it has the characteristics of a cough (Figure 1). Qualitatively, this looks like a quick increase in signal energy above an estimate of the ambient noise in the signal. This estimate is time-variant, meaning that it changes over time to take into account different environments the patients encounter as they wear this device throughout their daily life.

First, the baseline ambient noise level needs to be determined. An estimate of the ambient noise is calculated by finding the average minimum amount of energy that occurs over a 20 second window. That 20 second window is always moving so that it is centered on the point that it is calculating the baseline for. In quiet areas of the signal, a fixed value is multiplied with the signal to estimate the median noise (Larson et. al., 2012).

After these calculations, the input signal is compared with the baseline to find the initial set of cough events. These events are found where the signal exceeds the ambient noise estimate by a fixed threshold (Larson et. al., 2012). The start and end samples as well as the peak energy are also determined for each event. These start and end times underestimate event durations for loud events, or events where the peak is greater than the 9 dB threshold. In those cases, the start and end times are adjusted such that the signal energy either stops decreasing or it reaches the ambient noise threshold. If this results in two events overlapping, those events are combined. Finally, if the median energy of the initial onset of the cough is greater than a threshold of 6 dB, the onset is considered to be that of an event that needs to be classified (Larson et. al., 2012).

Now the algorithm has to classify the event. To do that, the event is broken up into 32ms long frames where acoustic features of each frame are calculated (Larson et. al., 2012). These features are analyzed and fed into a classifier offline to determine if the frame contains a cough. The individual decisions for each frame are combined to make an overall decision about the characteristic of the event. To make these individual decisions, the classifier outputs of first third of the frames in an event that appears to be a cough are averaged (Larson et. al., 2012). This approach helps differentiate between pure cough and a mix of cough and other vocalizations.

Classifying an Event
In order to classify an event as “cough” or “not cough”, it is compared against an existing database of signals. To create the database, researchers manually reviewed cough data files. They classified each file as representing either “cough,” “not cough,” or “unclear.” An example of an “unclear” event is one that contains combinations of coughs and other vocalizations such as a grunt or

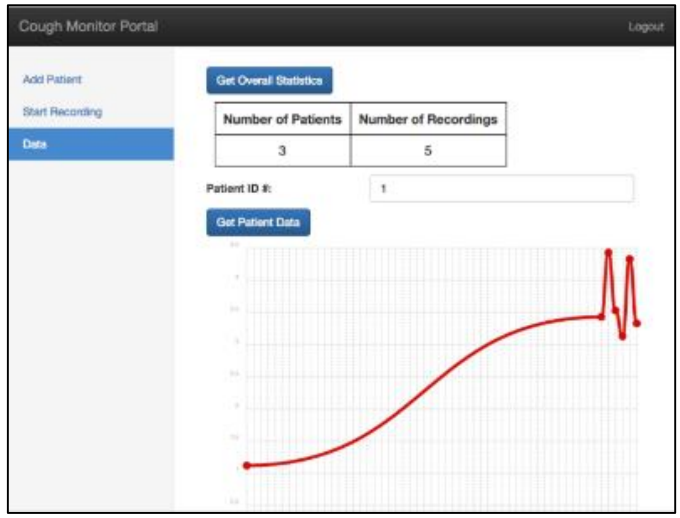

Figure 2. Web interface for clinician use in patient monitoring

speech, while a “not cough” event is one where there is only noise picked up from the environment. There is an opportunity for misclassification here due to human error, however this should not have been a frequent occurrence. Since misclassification should rarely have occurred, a large dataset was used so that a misclassified event would appear as an outlier.

Displaying the Data
The final step is to display the collected patient data in a meaningful way. To do this, the patient data is aggregated after they wear the device for a predetermined period of time (usually in increments of 1 hour). Aggregating the data allows the algorithm to process all of the recorded events at once and can determine how many of them represent cough events. The number of cough events over the number of hours the device was worn is equal to the cough rate. If the cough rates for each patient are graphed over time, it can be shown if the patient is responding to treatment. Figure 2, for example, shows that the average cough rate of the patient is *increasing* over time, and therefore we know that the patient is actually getting worse.

Conclusion
The Purple Team used this algorithm and display method in combination with a small hardware device to accurately detect cough. Our project was tested in Peru with doctors and nurses that work with patients with TB and MDR-TB. The medical staff and engineers

determined that the algorithm was sufficient for analyzing data from the microphone; however a new one should be developed to analyze accelerometer data. The clinician interface was edited to remove patient data that is not relevant to the study, thus making the interface more secure. Finally, the Purple Team discussed future opportunities for the device, including use for monitoring other respiratory diseases and developing medicines to treat those diseases."

References

Larson, S., Comina, G., Gilman, R. H., Tracey, B. H., Bravard, M., & López, J. W. (2012). Validation of an Automated Cough Detection Algorithm for Tracking Recovery of Pulmonary Tuberculosis Patients. *PLOS One , 7* (10). doi: 10.1371/journal.pone.0046229

World Health Organization. (2015). *2015 Global Tuberculosis Report.* Geneva: World Health Organization. Retrieved from http://www.who.int/tb/publications/global_report/en/