Using Deep Learning to Categorize Abnormal Respiratory Sounds

Tarkar, Ishani (School: duPont Manual High School)

The purpose of this project was to use machine learning to accurately categorize abnormal lung sounds. Hundreds of millions of people die globally due to respiratory diseases that would be treatable in the early stages. It was hoped that this project would result in a working product with a high accuracy, which could eventually be used in a medical setting. Four separate machine learning models were used to create the final product. Based on the answer to two questions (if the lung sound is inspiratory or expiratory, and if a stethoscope is required to hear the noise), one of the four models was chosen. Then, the file names and their prediction were printed to a CSV file, which was used to determine how accurate the model was. It was found that a similar model had been created where only 2 of the respiratory sounds were classified, with a 76% accuracy rate. Therefore, it was hypothesized that the finished product would have a 90% accuracy rate. The model had a 64.29% overall accuracy, which is significantly higher than the 36.5% accuracy of pulmonologists in the classification of respiratory audio. The results showed that one (wheeze) out of the five classes had a high inaccuracy incorrectly classifying the sounds, while another (stridor) had a high percentage of incorrect classifications (compared to correct classifications). One possible explanation for this was the small amount of audio in specific classes, resulting in underfitting in at least one of the models.