Do Masks Affect Social Interaction At All?

Justin Donovan and James Haddad
MECE E3018: Mechanical Engineering Lab I
December 30th, 2020

***Abstract*—The COVID-19 pandemic has had drastic effects on the lives of many human beings internationally; there are now enforced protocols declared by the government and state that people must follow in order to keep their loved ones and people around them safe. One of these enforced protocols involves wearing a face mask in a public setting. The most popular complaint when wearing a face mask is the miscommunication that comes along with it. Is there an actual effect on communication, and if so, do the frequencies that humans speak at change this effect? The goal of this experiment is to investigate if there is a change in the energy of the sound waves leaving our mouths when wearing a mask and not wearing a mask, and whether or not the pitch that a human speaks at has an effect on this. A phone was used to record the sound that a passive buzzer made, both with a mask on the buzzer, and without a mask on the buzzer, as it mimicked the range of frequencies that the average human speaks at. The RMS of the observed sound waves would be taken in order to compare the amplitudes for each frequency's trials. A two-tailed hypothesis test was used in order to determine if there was a significant statistical difference between the data sets, and there was. It was concluded with 95% confidence that there is a noticeable difference between the amplitudes for the mask and no mask scenarios. A linear regression could then be applied to the RMS values in order to observe the direct correlation pitch has on the amount of energy a sound wave retains after traveling through the barrier. This direct correlation implies that as the frequencies increase, the amount of energy lost throughout the propagation of the mask increases as well.**

I. Introduction

In the year 2019, a virus conventionally known as COVID-19 plagued the earth. This virus has no physical boundaries or limitations, striking every corner of the earth. Although modern medicine has not found a way to completely overcome the virus, there are many preventative measures that society uses to reduce the spread of COVID-19. These preventative measures consist of things ranging from social distancing (remaining six feet apart), sanitizing, limiting capacities of indoor places, moving school to online platforms, reducing indoor dining and other forms of nightlife, and most importantly wearing masks.

One byproduct of these preventative measures, that may be relatively intuitive, is the limitations placed on social interaction. Because social interaction is the main culprit for the spread of the virus, these limitations reduce human contact at the expense of our ability to socialize. Although safety is undeniably the most important focus regarding the virus, our aim was to negate some of the negative effects that these safety measures have on social interaction while still keeping the health and well being of society as the main priority.

This work was completed for the MECE E3018 Laboratory I course

One of the preventative measures, wearing a mask, is used to protect others from respiratory droplets that might contain the virus. Masks are worn over the nasal and oral cavities, which are the sources of the respiratory droplets. It is indisputable that these masks are necessary to assist in limiting the spread of the virus; however, they do have a direct impact on ones ability to verbally communicate with others. Masks have become a part of every day life for the majority of Americans, and will continue to be a greater part of peoples lives until we have the proper means to overcome this virus. Below is a graph of mask usage in the United States as of July 14th 2020, and public restrictions and government mandates to wear masks became even more prevalent in society in the months that followed July. [1]

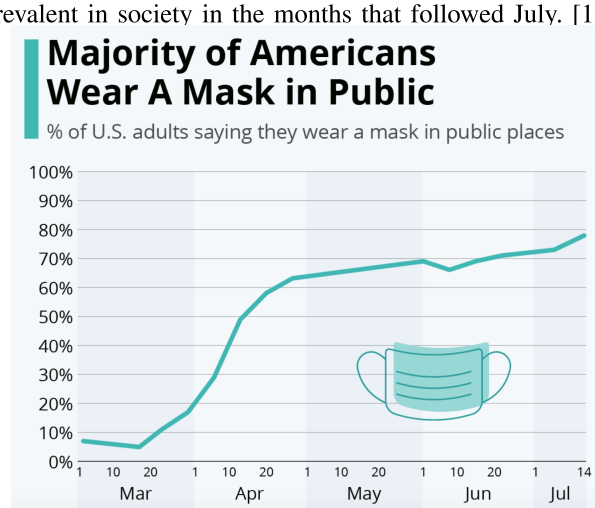

Fig. 1. The trend of mask usage in the United States as of March 1st 2020.[1]

Our research identifies two main impacts that masks have on social interactions: taking away our ability to read lips or facial cues, and muffling the actual sound of the speaker. When considering face to face communication, reading lips is an extremely important aspect of speech processing. The absence of lip reading results in the decline of the ability to process speech in interactions. Presbycusis [2], otherwise known as age related hearing loss, gets worse with age. Therefore, people at older ages are affected disproportionately more from the covering of facial cues. This is because they tend to rely on on the facial cues to understand the verbal message more than someone who is more equip to perceive the sound. We initially thought about transparent masks, but they seem to pose more problems than solutions for reasons such as, the material, sweat, fog and difficulty breathing.

Our experiment aims to test pitches in the range of common human speech (100 Hz - 250 Hz)[3] and determine which pitch is muffled the least by the mask. When sound waves are emitted from the speakers mouth, they have to to penetrate the fabric of the mask to reach their receiver, resulting in dissipated energy. Therefore the sound waves that reach the receiver are dampened. We want to find out if there is an optimal pitch to speak at, or one that will result in the least energy dissipated. Our set up is designed to mirror the exact relationship that humans have with the mask. We developed two scenarios, one where a mask covers the buzzer forcing the sound waves to travel through the mask, and one where the sound waves travel directly to the receiver with no obstructions. A phone will be six inches from the buzzer in both scenarios to record the sound.

Low frequency has less attenuation through matter, therefore the sound wave's energy dissipates less as it penetrates matter. This is based on the concept that longer wavelength sound waves have more profound and less frequent vibrations on their barrier (the mask), allowing sound to propagate through the barrier more effectively [4]. Therefore, as the pitch gets higher, the energy will dissipate more through the medium, demonstrating an inverse relationship between frequency of the sound wave approaching the barrier and energy of the wave on the other side of the barrier. In the case of a face mask being worn, the mask can act almost like a high frequency filtration system.

Therefore, it is hypothesized by this experimental lab group that pitch and frequency have a direct effect on the permeability of sound waves through a face mask, and that a lower pitch on the human spectrum of speech is more audible than a higher pitch, when passing through a mask.

II. Methods

A. Theoretical Background

Sound waves travel as sinusoidal waves. They can be classified as longitudinal waves, which mean that they're classified as some sort of disturbance in the molecules of the medium they travel through. In the case of sound, the air molecules are disturbed. Since the molecules are disturbed, sound can be classified as energy. The medium of which a sound wave travels through may differ, as it can propagate throughout solids, liquids, and gases. The human ear acts as the sensor that receives the sound wave for humans to process. Louder sounds push harder against the eardrum, while quiet sounds push softer against the ear drum.

As a sound wave approaches a barrier, it undergoes reflection, absorption and transmission. When the wave reflects off the surface of the boundary, only a portion of it reflects, reducing the percentage and energy of the wave that travels through the actual boundary. The percentage of the wave that remains in its initial direction of motion is partially absorbed by the new medium, causing the resulting percentage of the wave that continues through the boundary to have even less energy. Because the resulting wave has less energy, the sound wave is dampened and has a smaller amplitude. This may be visualised with the assistance of the following picture [5].

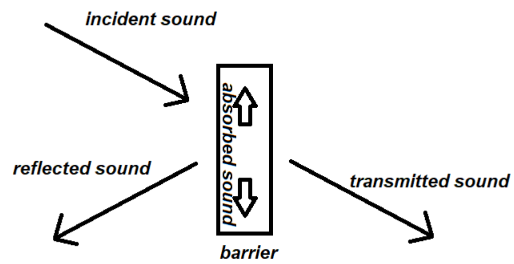

Fig. 2. Sound waves through barriers

The equation of a wave can be represented as a simple sinusoidal function, given by Equation 1, where A is the amplitude of the wave, t is time, and ϕ is the phase shift. The angular frequency ω_n can be represented by 2πf and the wave number k can be represented by $\frac{2\pi}{\lambda}$. The amplitude A can be shown in Figure 3, where A can be both positive and negative. In terms of data manipulation, this isn't a suitable form to manage the data with. We must take the RMS of the sound wave's trial in order to compare the amplitudes due to this periodic trend. In this experiment, we will be examining how the amplitude A changes as the frequency f changes with the mask acting as the barrier that the sound wave propagates throughout.

$$x(t) = Asin(kx + \omega_n t + \phi) \quad (1)$$

As the pitches vary, we note that the wave number k and the angular frequency ω_n change. Our goal is to determine how this impacts the amplitude as the wave travels through the medium by testing our apparatus at varying frequencies.

B. Apparatus

To find the best pitch range, we looked at every pitch in the spectrum of common human speech. We then hard coded the pitches in the buzzer through the Arduino UNO. Next, we set the buzzer six inches away from a phone that records the audio as an m4a file. We set up each apparatus as shown in each image below. The distance between the buzzer and the phone don't necessarily matter due to the distance being constant throughout the whole experiment. The buzzer stand was taped down to the cardboard bottom and fixed with no room to move. The phone stand was also consistently taped down and didn't move throughout the experiment.

It is important to note that in order to emulate a real world scenario, we wrapped the mask entirely around the buzzer. We did this to make sure that no noise reached the receiver without passing through the mask. We also set the receiver six inches from the buzzer because the receiver did not have enough sensitivity to record the buzzer accurately from a further distance. We made sure that the mask did not actually touch the buzzer, because this would alter the sound waves, making the mask also vibrate and propagate sound waves. Therefore, with the mask wrapped around the buzzer but not

Fig. 3. Experimental apparatus without a mask

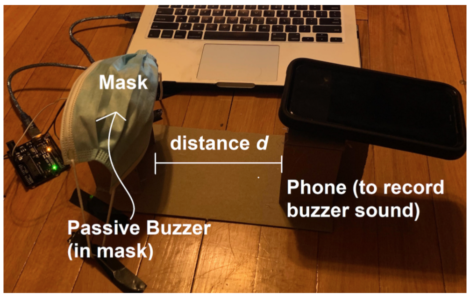

Fig. 4. Experimental apparatus with a mask

in direct contact with it, we replicated how masks interact with humans as accurately as we could.

The sensor used in an iPhone X is called a MEMS microphone. It's a very small non-electret condenser microphone that takes up very little space and power relative to the rest of the phone. MEMS stands for micro-electromechanical systems and this type of microphone is the most common one used in many mobile phones [6]. The way this microphone works is that it detects the vibrations that sound makes via a microscopic diaphragm which then causes variations in the capacitance of the MEMS. This data is then converted into a form that the smart phone is able to read and utilize. Despite knowing the type of sensor that the smart phone uses, there currently isn't any literature nor public information on the biases that come along with it as well as the specific model.

The data was then stored in the phone as an mp4 file and uploaded to MATLAB Audio Toolbox. This is an add on to MATLAB and is used for things such as speech analysis, audio processing, and acoustic measurement.

C. Approach

We began by graphing the audio data set that we recorded and physically examining it. We then truncated the data set into a one second sample with the least amount of outside interference as well as abnormalities between the start of recording and the end. This was done for every trial recording.

We then made the one second sample of the audio with a mask obstructing the noise emitted and a one second sample of the audio without any masks. Each one second sample contains approximately 10,000 data points that record both positive and negative voltages and performed a root mean square on the data in each one second sample. We took a one second sample with the mask and a one second sample without the mask and performed a root mean square on each sample. This is because the values periodically spike from negative to positive and the magnitude of the volume is is recorded at all points in the perceived sound wave. We did this for all of the trials in a given frequency and compared the mask and no mask trials with a two tail t-test in order to validate that they are different, despite the values being physically different when observed numerically. We then found the difference in the root mean squares for these samples and recorded it as the volume lost by the dampening from the mask. We repeated this process four times at each frequency. Once we recorded all of our data we took the average volume lost at each frequency and performed a linear regression to determine if there is a trend between volume lost and pitch of the sound emitted.

D. Data Processing

The 95% confidence intervals for the mask trials and no mask trials under the same frequencies were calculated using equations 3 and 4, where ν is the degrees of freedom ν = N-1. The $t_{a/2,\nu}$ value can be found by a standard t-distribution chart. α will always be a constant of 0.05 under the assumption of a 95% confidence.

$$S_x = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i - \overline{x})^2} \tag{2}$$

where n is the number of trials

$$S_{\bar{x}} = \frac{S_x}{\sqrt{N}} \tag{3}$$

$$\bar{x} - t_{a/2,\nu} S_{\bar{x}} \leq x' \leq \bar{x} + t_{a/2,\nu} S_{\bar{x}} \tag{4}$$

E. Linear Regression

When calculating the equation for the best fit line we note that the equation of our line is.

$$y = a_0 + a_1 x \tag{5}$$

We began by finding the value for Δ, which we see as

$$\Delta = N\sum_{i=1}^{N} x_i^2 - [\sum_{i=1}^{N} x_i]^2 \tag{6}$$

where N is the number of trials. We then use this value to calculate a_0 and a_1 which we see as

$$a_0 = \frac{(\sum_{i=1}^{N} x_i^2 \sum_{i=1}^{N} y_i - \sum_{i=1}^{N} x_i \sum_{i=1}^{N} x_i y_i)}{\Delta} \tag{7}$$

and

$$a_1 = \frac{(N\sum_{i=1}^{N} x_i y_i - \sum_{i=1}^{N} x_i \sum_{i=1}^{N} y_i)}{\Delta} \tag{8}$$

Now that we have our best fit lines, we now need to find our best fit lines to show the upper and lower bounds of our linear regression line. To do this we first need to calculate the variance of fitting parameters or S_y^2 by

$$S_y^2 = \frac{1}{N-2} * \sum_{i=1}^{N}(y_i - [a_0 + a_i x_i])^2 \tag{9}$$

We then use this thus to get S_{a0}^2 and S_{a1}^2 which we calculate with

$$S_{a0}^2 = \frac{\sum_{i=1}^{N} x_i^2}{\Delta} S_y^2 \tag{10}$$

and

$$S_{a1}^2 = \frac{N}{\Delta} S_y^2 \tag{11}$$

And we can use these to calculate an estimation for the true intercept of the true mean regression line and slope as

$$a_0^{true} = a_0 \pm t_{\frac{\alpha}{2},N-2} S_{a_0} \tag{12}$$

and

$$a_1^{true} = a_1 \pm t_{\frac{\alpha}{2},N-2} S_{a_1} \tag{13}$$

where $\alpha = 0.05$ because this gets us the region of confidence that our best fit falls into with 95% confidence. We can put together lower and upper bounds by using the smaller a_0 in conjunction with the smaller a_1 to make the lower bound, and the larger values to make our upper bounds. This forms our 95% confidence region.

The linear regression was applied to percentage of sound energy lost when comparing the masked trials to the no mask trials. The no mask trials were treated as the initial sound energy sent to the sensor with no dampening. The masked trials had dampening. The percentage was calculated by subtracting the masked trials from the no masked trials and dividing that quantity by original no masked trials' values, as shown by equation 14. This is all with respect to the average RMS values in each respective frequency. In order to truly check to see whether or not the points comfortably lie within the confidence intervals, error bars could be applied to these points. The error bars can be applied by calculating the standard deviation of the 5 points that were found from equation 14. The standard deviation can be calculated based off of equation 2. These error bars that are calculated will show how likely the points generated on the linear regression graph will fall within the confidence interval.

$$A_{RMS} = \frac{A_{nomask} - A_{mask}}{A_{nomask}} \tag{14}$$

Please note that all equations used in this Methods section were provided during Dr. Yesilevskiy's lectures given for the MECE E3018 Lab I course.

III. Results

The goal of this experiment was to test whether or not wearing a mask affects communication, and if there was a certain frequency interval that this effect dominates in. This experimental lab group hypothesized that the mask indeed affects communication by hindering it, and that lower frequency sound waves can propagate throughout the face mask in an efficient manner compared to higher frequency sound waves. This effectively means that a face mask may act as some sort of higher frequency filter.

In order to address this hypothesis, the amplitudes of the sound waves were measured and compared in the mask and no mask trials. The trials ranged from 100, 140, 180, 220, and 250 Hertz, where there were 4 trials for each mask and no mask case. The raw data was converted into a more useful format by taking the RMS values of the multiple data points per trial. Figure 5 shows the average RMS values of each trial in each frequency set compared to one another. It can be visually seen that the no mask amplitudes are visually higher than the mask values, which indicates that some of the original sound energy was dissipated as the sound wave propagated throughout the polypropylene medium. This supports the experimental lab group's hypothesis. This hypothesis is further supported by Figure 6. Here, the corresponding mask

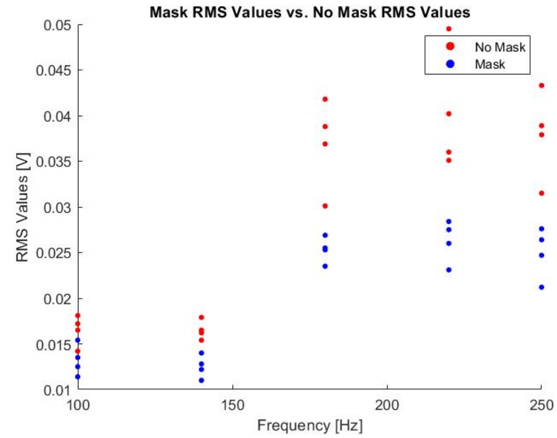

Fig. 5. RMS values of all of the no mask trials and mask trials are plotted on a scatterplot and compared to one another so that the difference between the amplitudes are visually seen.

TABLE I
Sound energy lost calculated by comparing average RMS values to one another in mask and no mask scenarios

Frequency	% Lost With vs Without Mask
100 Hz	0.20000
140 Hz	0.22424
180 Hz	0.31436
220 Hz	0.33831
250 Hz	0.36412

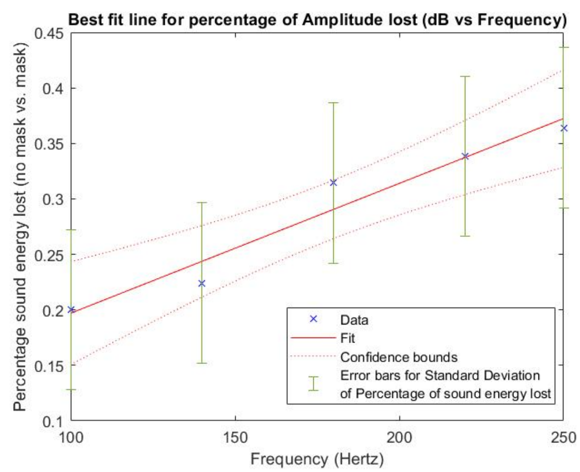

Fig. 6. The amount of amplitude lost (mask compared to no mask) is shown with the calculated 95% confidence intervals and linear regression best fit line.

TABLE II

SLOPE, Y-INTERCEPT, AND 95% CONFIDENCE INTERVALS FROM THE GRAPH IN FIGURE 6

Slope:	0.0012
Y-Intercept:	0.0800
Confidence Interval	± 4.894E-4

and no mask trials are compared to one another respectively. The mask RMS value is subtracted from the no mask RMS value, and this quantity is then divided by the no mask RMS value, as shown by equation 14. This leads to the quantity at hand representing the percentage of the amplitude lost before and after the mask. These values are explicitly shown in Table 1. The confidence bounds are then graphed within a 95% confidence interval. The error bars calculated from the standard deviation of these 5 points are shown as green. Here, it's seen that the points at the end are more likely to fall within the confidence interval, but as the points in the middle of the graph steer away from the linear regression line, the error balls fall out of the confidence intervals drastically. This is shown in the middle 3 points. Despite there only being 5 points on this graph, each point represents 8 total trials used (4 mask and 4 no mask), where there are hundreds of thousands of data points from each trial being averaged and compared. Despite the error bars being very small, they represent the 95% confidence interval of each point.

IV. DISCUSSION & CONCLUSION

A. *Interpretation of Results*

From our data, we can say with 95% confidence that the sound emitted through a mask is quieter than the sound emitted without the presence of a mask, at all frequencies of common human speech. This is proven from the results of our two tailed hypothesis test on the average RMS values of samples with the mask versus without the mask at all frequencies. We can then say that when the frequency of the sound wave propagating throughout the mask increases, the percent lost of sound energy is increased. Thus, as the frequency of the sound wave increases, the noise becomes increasingly dampened as it passes through the mask. This is best shown through the trajectory of the linear regression of the data that compares frequency to the percent of the sound lost when we compare waves through a mask and through no mask. We can see in the confidence bounds that the entire region has the same trajectory as the linear regression. We therefore were able to successfully provide evidence for our hypothesis that the lower frequencies are more successful at penetrating the mask. Despite the error bars falling out of the confidence intervals towards the middle frequencies of the given range we chose, the multiple two tailed tests and the linear regression provide a sufficient amount of evidence within a 95% confidence that the experimental lab group's hypothesis is proven.

The implications of this result may have a significant impact on industries and occupations that primarily rely on verbal communication. If this experiment's results provide a conclusive model for most types of masks, then this would mean that it would be in a majority of organizations' best interests to develop or invest in an alternative method for protection that offers the same amount of functionality that a face mask provides without hindering communication . For example, developing a different type of safe yet non-dampening material or even designing a new shape of face mask that does not necessarily dampen incoming sound waves at a significant rate. This could potentially lead to increased efficiency in communication at workplaces, because workers and employees would have to spend less time reiterating sentences.

One shortcoming of the results that we found is that the data was initially given with positive and negative values, and we had to perform a root mean square on the data to obtain the average volume of the wave emitted. Because each recording had to be recorded and converted to an mp4 file and then read into MATLAB separately, we had trouble recording a large amount of trials at each frequency. We determined the number of trials that we used, eight at each frequency (four with the mask and four without the mask), appropriate, because each once second trial recorded 10,000 sound bytes. Therefore, from this we could see a difference between when a mask is used and a mask is not. This is clearly evident even before any data processing from Figure 5. We also were able to see a clear trend between pitch and amplitude of the sound that is aligned with our hypothesis. Thus, we kept our number of trials at eight for each frequency. If we were to improve on this data, we would get more data points at each frequency by either taking more trials, or finding a more efficient or better sensor.

We chose not to include the inherent bias uncertainty in the microphone of the iPhone, because we could not find a widespread consensus for its value. We also found that the resolution that we recorded the audio in was to four decimal points, and then deemed any bias uncertainty produced by the iPhone negligible. We did however get the standard deviation of the loss in pitch at all five frequencies and use these as our error bars when plotting. We also calculated error bars that contain our linear regression, forming our error region with 95% uncertainty. With the uncertainty that we included, we could still confirm our hypothesis with 95% confidence

Although we were able to find many shortcomings for our apparatus, we believed that our set up successfully tested the motivation that led us to this experimental setup. We were able to trim and clean each trial to remove any parts of the recording that had background noise and cut each trial down to the cleanest one second recording possible. We also believe that we accurately modeled how the mask interacts with the oral cavity, and replicated it perfectly. We did so by fully covering the buzzer so that all of the sound emitted from the apparatus passed through the mask. We also were able to successfully keep the buzzer from touching the mask. Although we had trouble replicating the real world example of humans socially distanced at six feet apart, we were able to use our set up to successfully identify the optimal pitch in the range of common human speech: 100 Hz.

We initially attempted to use a big sound sensor driven by an Arduino UNO, but, the sensor was not sensitive enough. We were able to see very small perturbations in the detected noise, but it wasn't precise enough to see a significant difference nor apply data processing methods. This resulted in alternatively using an Apple iPhone Model X as the sensor. Due to the acoustic environment and the sensitivity of the iPhone's microphone at a distance, we had to resort to an approximation of the experiment where the phone was six inches from the buzzer. One issue with this is that the approximation might not be applicable to some scenarios, because in real world scenarios people must be six feet from each other. A more sensitive sensor in a significantly more isolated environment would be able to handle this large distance, but the current tools available could not handle such a feat. In hindsight, however, the distance does not necessarily matter with regards to this experiment as long as it is fixed with no chances of movement. One possible solution to this would be to either find a more suitable sensor for the experiment, or find a better environment to record in. A more suitable environment may include one that is completely quiet, where no outside interference can be picked up from sensor. A real world application of such an environment may include a studio recording room.

Since masks are located right next to the mouth, there are times when the moisture content can be at a significantly high rate due to saliva and sweat. This increase in moisture content may have some sort of effect on the dampening of incoming sound waves as well. Therefore, one way we could make this experiment more applicable to a real world scenario is by also increasing the moisture content of the mask at each pitch and comparing it to the sound emitted when there is no mask present at all.

Although we were able to clean each trial and get the most consistent one second clip of audio, we still are unable to guarantee that there was no outside noise recorded. Because the room was not a perfect acoustic environment and there might have been some outside noise, the results could have been altered. We believe that the outside factors such as noise and acoustics were negligible; however, if we were to improve on this experiment we would repeat our trials in a perfectly silent and sound proof room.

Another shortcoming of the experiment is only using a polypropylene mask. Many people nowadays use many types of different masks, ranging from cloth masks to N95 masks and everything in between. We came to the conclusion that masks do hinder communication and act as a high frequency filtration system, but this only applies for a polypropylene mask in our case. This is corroborated with some data from other literature investigating the dampening of sound waves as they propagate throughout some sort of mask as a medium.

According to one study by the Illinois Augmented Listening Library research team at the University of Illinois, face masks generally muffle high frequency sounds on the spectrum that humans communicate on, and this was proven to be true for a large range of different mask materials. It was proven that polypropylene masks tend to serve as one of the best masks for speech, whereas the more heavy-duty and reliable N-95 masks and heavier masks are typically worse for communication. [7]. The literature proposes for the ranges of frequencies to be heard to be extremely high relative to the range we have chosen, where the dampening is very evident at the extreme frequencies. This still supports the conclusion that we found, even though they used a more sensitive microphone and isolated environment.

Another way our finding aligns with literature we used to determine our hypothesis is in its parallels to the technology that goes into SONAR. The word SONAR stands for sound navigation ranging. This technique is used to measure the depth of the ocean and then-some. Low frequency active sonar uses pulses of low frequency sound waves that are shot beneath the device into the water.[8] Low frequency waves are used due to the high wavelength that's associated with it. Whenever the pulses hit any sort of object or even the bottom of the ocean, they are reflected back onto the surface where a sensor can use the received information to determine how far an object is. This phenomenon is based off of how dolphins and bats utilize echolocation. SONAR is proof that low frequency waves have less attenuation throughout matter, therefore the sound wave's energy dissipates less as it propagates throughout a medium.

In summary, the main research question for this experiment was whether or not masks muffle sound when humans communicate with each other, while the secondary question for this experiment was if the pitch humans spoke at had any sort of effect on the dampening of the outgoing sound waves. In order to answer these questions, a passive buzzer that was programmed to produce noise at certain frequencies was used to produce noise that was detected by an iPhone sensor. We recorded that the noise from the buzzer was louder when there was no mask covering the buzzer; we also found that as frequency increases, more sound energy is lost throughout the medium of the mask. This experiment was successful in highlighting the effect frequency has on the propagation of sound throughout a medium. In order to truly investigate this topic, future experiments should be designed where different factors including a different frequency range, moisture content, type of mask, etc. are tested in the interest of reducing error and uncertainty.

REFERENCES

[1] F. Richter, "Majority of americans wear a mask in public." https://www.statista.com/chart/22347/share-of-us-adults-saying-they-wear-a-mask-when-in-public-places/.

[2] "Hearing loss: A common problem for older adults." https://www.nia.nih.gov/health/hearing-loss-common-problem-older-adults.

[3] I. Titze, "Principles of voice production, prentice hall." Prentice Hall.

[4] R. P., "Low versus high radio spectrum." High Tech Forum.

[5] "The science of soundproofing." https://www.soundstop.co.uk/science-of-soundproofing.

[6] A. Fox, "What is a mems microphone? (micro-electro-mechanical systems)." https://mynewmicrophone.com/what-is-a-mems-micro-electro-mechanical-systems-microphone.

[7] S. A. Corey RM, Jones U, "Acoustic effects of medical, cloth, and transparent face masks on speech signals." The Journal of the Acoustical Society of America 148: 2371–2375.

[8] H. JA, "Anthropogenic and natural sources of ambient noise in the ocean." https://doi.org/10.3354/meps08353.