HEARTWARE

we listen

Johnathan Chen	EE/M&T
Jason Schwartz	CMPE
Jared Winograd	CMPE
Matthew Wolfman	EE

Advisor: James Weimer
ESE Senior Design Team 1

Table of Contents

II. Executive Summary

HeartWare keeps athletes safe and empowered to perform by guarding their heart health and communicating issues they would otherwise overlook. By combining a wearable heart-sound monitor with real-time, advanced machine learning technology, HeartWare alerts the wearer or coaches/trainers of cardiac problems that are plaguing young athletes the most, including sudden cardiac arrest and arrhythmias.

HeartWare uses a piezoelectric sensor to record the heart sound cleanly and without picking up noise from outside sources. This sensor is connected to a custom-designed PCB that uses an on board WiFi chip to stream the audio directly to our cloud platform. Once there, the audio is separated into chunks and analyzed via our convolutional neural net in the cloud. If there is any sign of an issue with the heart, the user can get a notification to seek medical attention.

This technology, while focused on the heart for this project, has the potential to help diagnose conditions that affect many other parts of the body. It makes personal health more accessible to the user, and can simplify and lower the cost of high-quality, wearable medical technologies. The potential for this product is huge, as it demonstrates the complete, integrated pathway from picking up sound, streaming it, and analyzing it via machine learning in real time.

III. Overview of Project

Sudden cardiac death is the most common medical cause of death in athletes, afflicting 1/40,000 to 1/ 80,000 athletes per year (Harmon, 2014; Wasfy, 2016). Especially at the upper levels, athletes are required to push their physical limits, sometimes undergoing greater than 40 hours a week in high-intensity training and competition (Jacobs, 2015). Because athletes are accustomed to exertion, however, they can habituate to pain, ignoring symptoms that may normally warrant more serious medical attention. Thus, coaches and trainers need new technologies that can encourage high performance in their athletes while ensuring their safety.

On a wider scale, it is interesting to note that of the many biometric data currently collected, little is recorded for one of the most central functions of our health. We are building a continuous, wearable smart stethoscope that analyzes and reports what it listens to. The idea stems from a simple observation: cardiologists cannot follow their patients (or anyone) around 24/7, but there are many specific points in which monitoring would be useful. By building a profile of what the user’s heart should sound like in its healthy state (see Appendix A), we can predict the need for medical attention earlier on, bringing health care to patients as early as possible.

Our project distinguishes itself in that it is wearable and continuously monitoring. Specifically, we start with the problem of heart murmurs during intense physical activity, which can be precursors to sudden cardiac death. The only solution is to constantly listen to the heart, monitoring for any potential sudden changes that may need the attention of medical professionals.

This use case is only the beginning, but it is one that showcases the immediate importance of such a product. A device that acts as an always-on, low-resolution nurse for the heart can be adapted to any variety of diseases that are able to be picked up by current non-smart stethoscopes. These include heart, lung, digestive, and other diseases and issues that can be diagnosed via sound. Senior citizens who cannot afford to bring someone in or regularly go to the doctor can now utilize more immediate analysis which will alert them of serious issues that may necessitate the visit to the doctor. Nurses who wish to ensure healthy physical activity can require students wear these during Physical Education. The use cases continue while the core product stays the same: a stethoscope that is always listening and analyzing the heart and body.

IV. Technical Description

Hardware

Our final hardware product is a mounted printed circuit board. Significant engineering work went into developing different circuits on the PCB and the PCB itself. The sensor is mounted on the PCB, and the PCB has analog circuitry that interfaces the sensor to the ADC on the microcontroller. The microcontroller (with the help of associated digital circuitry and firmware) then reads the values and sends them to the cloud via WiFi. The PCB+sensor combo is mounted on the chest using a strap and laser-cut mount that was designed in Solidworks. Each of these components of the hardware will be explained in greater detail below.

The sensor itself is a piezoelectric sensing element. There is a flexible piezo-membrane inside the sensor that bends as the nub at the top of the sensor is pressed. This pressing induces the piezoelectric effect that generates voltage proportional to how much the sensor is bending. In the fall semester we spent a lot of time developing circuitry to amplify and capture the voltage created by the piezoelectric effect. However, it was very difficult to design an amplifier to capture all of the low frequency vibrations we wanted. We spent the beginning of the spring semester experimenting with different sensor and amplification architectures, and in the end, we chose a sensor package that had some amount of pre-amplification optimized for low frequencies built in. This sensor is the TE Connectivity CM-01B. The schematic and frequency characteristic bode plot diagrams can be seen in Appendix B.

The next stage of the PCB is the analog circuitry to interface the sensor to the ADC. There were a few different design goals and constraints for this aspect of the design. The first is that the chosen microprocessor, the ESP8266, has a 10-bit ADC with an input range of 0-1V. The sensor itself output a signal that was centered around 0V, and ranged from -150mV to 150mV when mounted. Furthermore, the signal picked up all noises traveling along the skin, which included both low frequency and high frequency noise outside of the desired 20-200Hz range for the PCG. Thus, an analog circuit that acted as a level shift, amplifier, and filter was designed. The circuit and other aspects of the analog design work can be seen in Appendix C. The level shift component involved using resistor dividers throughout the signal to ensure that it was above ground. The amplifier component was to expand the voltage range of the output from about 250-300mV to 600mV to span more of the ADC's input range without risking going to high and damaging the ADC input port. We discovered that because of the limited 10-bit resolution and noise, it would not suffice to have the signal range only 300mV. A single supply operational amplifier was used to prevent the need for a negative power rail on the device. The filter was designed as a first-order bandpass to eliminate high and low frequency noise, with cutoffs of about 19Hz and 312Hz to help the ADC pick up more of the desired range of 20-200Hz. We experimented with higher-order filtering in hardware but ultimately determined that the bulk of filtering should be done in software where it was cleaner and easier to adjust cutoffs. This filtering, high order butterworth low and high pass filters, were implemented in python on the cloud where they were applied in real time on the live streaming data and displayed alongside the unfiltered data in our cloud visualization platform. The filter and amplifier circuits were analyzed and tested by using a Matlab analysis of the signals captured. We created a dashboard with 6 elements–the time and frequency representations of the signal directly taken from the sensor, the time and frequency representations of the signal taken from the sensor through the amplification and filtering circuitry, and the time and frequency representations of a signal from the PhysioNet challenge that we used as a control. The goal of the analysis was to make the signal coming through the analog circuitry look as much like the PhysioNet data in both the time and frequency domains. While we were very successful in getting data that looked great in the time domain, which is the live-streaming demo we showed to the judges at senior design demo day and at the SEAS competition, the frequency profile of the captured sound still does not look exactly the same, which results in lower accuracies when running on our own captured data vs. the data trained from PhysioNet. If we were to continue with this project by getting an IRB approval for human testing, we could collect a few hundred samples from patients labeled with and without heart conditions from our own sensor and train our neural net classifier from our own data, eliminating the need for the PhysioNet data.

The PCB is a 4 layer-board that comprises the analog and digital circuitry, and is powered by a 9V battery. For the ESE demo day, the analog circuitry was a separate perf-board taped to the back of the PCB that took the sensor as an input, connected to the digital circuitry on the PCB as an output, and shared power lines with the digital circuitry. For the SEAS competition, we created a second spin of the PCB that had all of the analog components built-in. The PCB was designed in Altium and was manufactured in the US. This is further discussed in the self-learning section of the report. The microprocessor on board is the ESP8266, which is a chip with onboard Wifi. The digital circuitry on the board is meant to setup, power, and run this chip. The full schematic for the PCB, and screenshots of the PCB layout, and pictures of the manufactured PCB can be seen in Appendix C.

The sensor is mounted directly on top of the ESP8266 chip with double sided tape. It is held well in placed by the wires that attach it to the PCB, the tape, and the acrylic mounting plate which has a hole cut out exactly to the diameter of the sensor. The board as four half-circle screw holes meant for 2-56 screws that fix the position of the board and clamp it to the mount. The mount was designed in SolidWorks and lasercut in Penn's Rapid Prototyping Laboratory. The design has belt loop holes for the chest strap to run through to press the device against the user's chest. This worked well, but the device still picked up the PCG better when the user directly pressed the device against their chest than when the strap was tightened as much as possible. Further work to productize this device might require an additional strap or some other component to add pressing force against the chest. The chest strap can be taken on and off by the user but it is slightly easier to have a second party help. The chest strap holds the 9V battery in the back. Using power calculations we estimate the battery life of the device to be about 9.5 hours, which is greater than the 4 hour design specification we got from user research. The weight of the device is 84g, slightly heavier than the 60g design spec. Further work would include optimizing battery selection to hit both targets. The mount design and fully assembled hardware can be seen in Appendix D.

Classifier & Preprocessing

In order to classify heartbeat sounds, we first transformed them into images, and then fed them into a convolutional neural network, which is discussed later. Our preprocessing techniques for creating an image from our raw audio recordings of the heartbeat are based off of a top-place finisher[1] in the 2016 PhysioNet heartbeat classification challenge[2]. First, the time series signal is segmented into shorter segments that each begin with an S_1 heart sound and last for a length of T = 3 seconds. This segmentation is done with a logistic regression hidden semi-Markov model by incorporating information about the expected duration of each state (S_1, Systole, S_2, and Diastole).

Each audio segment is then transformed from a one-dimensional time series signal into a two-dimensional heatmap that captures the frequency-time distribution of the signal's energy. Mel Frequency Cepstral Coefficients (MFCCs) are used for this transformation. MFCCs are commonly used as features in speech recognition systems, making them a powerful tool in effectively characterizing heartbeats. The high-level steps of producing the heatmap are listed below. More details are included in the research paper linked earlier.

1. Use a sliding window of length 0.025s and displacement 0.01s to apply the operations in the remaining steps to overlapping portions of the signal.
2. Computer the Discrete Fourier transform for each window.
3. Apply a filterbank of 6 triangular bandpass filters and sum the energies of each filter together.
4. Apply a Discrete Cosine Transform on the summed energies to undo the correlation from the overlapping windows.
5. Represent all 6 uncorrelated energies with grayscale pixels, and then stack them on top of each other using a colormap. These stacked energies are the MFCC's.
6. Finally, align all 6x1 MFCC's horizontally to produce an 300x6 image, which now contains the MFCC's of each of the 300 time windows in a temporal sequence. An example of this final image is included in Appendix H. This final image becomes the input to the neural net.

We designed a convolutional neural net (CNN) based on the same top-place finisher in the challenge. The machine learning classifier was built to work with the 2D sound images that incorporate both time and frequency information. Our original model that just used the

[1] https://arxiv.org/pdf/1707.04642.pdf
[2] https://physionet.org/challenge/2016/

time-series data was much less accurate than our final solution (described below). This is because frequency information is an important part of audio. Out final classifier architecture was a relatively standard classification network and can be found in Appendix G. The input to the network is a 6x300 grayscale image (ie 6x300x1). First, the image is run through a convolutional layer with a 2x20 kernel, followed by batch normalization, ReLU activation, and a max pooling layer with a 1x20 kernel. This reduces the image to a 60x6 size, but there are now 64 channels. It is then run through another convolutional/norm/ReLU/pool layer, which reduces the image size to 30x6 (still 64 channels). Finally, there are two fully-connected layers, with 1024 and 512 input channels respectively. A dropout rate of 50% is applied to both of these fully-connected layers to ensure the neural net does not overfit the training data and underperform on the test data. Ultimately, a log softmax is used on the two output channels to get the probability of the heart sound being “Normal” and the probability it is “Abnormal.”

Using the dataset of thousands of sounds from the PhysioNet 2016 challenge, we were able to generate around 65,000 training images, train the CNN, and ultimately recorded a training accuracy of more than 99%. We used standard cross-entropy loss to train the network, with an adam optimizer and a learning rate of 1.5e-4. Slight variations in the architecture and hyperparameters were made to optimize the network, and we selected the parameters as described above. Training was done on an NVIDIA GPU to accelerate the process- this allowed us to train for 40 full iterations, and we then selected the iteration with the highest test accuracy.

A validation image set, which consisted of images that were not used to train the net, was used to confirm to determine test accuracy, or accuracy on “new” samples. The validation set was split 50% normal, 50% abnormal recordings. We recorded a test accuracy of 79%, which was on par with top finishers in the PhysioNet 2016 challenge.

Our preprocessing and classifier system was hosted as an interactive cloud platform on Amazon Web Services. The webapp consists of a server written in Node.js that listens for data that is streamed over Wi-Fi from our hardware system. The data are periodically recorded from the ADC at a sampling rate of 1000 Hz and sent as the body of a POST request to the server. The server then stores this data into an mLab database. The server also hosts a homepage that fetches the live heartbeat data from the database and displays it in real-time (Chart.js is used for this visualization). Simultaneously, the homepage also calls a Python program to filter the raw heartbeat data with a Butterworth bandpass filter and then display this real-time, filtered data. Upon a user button-press, the server will call a MATLAB program to use all the filtered data collected thus far to generate a set of frequency heatmaps and classify the heartbeat as normal or abnormal. The classification result is then displayed on the web application. A screenshot of the cloud platform is included in Appendix I.

V. Self-learning

A.

On the hardware side, there was a lot of self-learning about how sensors work. We decided to run with an analog piezoelectric sensor architecture for the spring semester after seeing promising results for outside noise-suppression in our comparison of microphone types we demoed at the end of the fall semester. Both Johnathan and Matt hadn’t take semiconductor/sensor physics beyond ESE218, and have minimal mechanical knowledge, so a lot of learning occurred in this area. One great resource was the piezosensor kit that we bought that had about 10 different types of sensors to try, with documentation of all of them and a mini-textbook about how the technology actually worked. Another great resource was Penn professor Jay Zemel who helped guide Johnathan and Matt towards understand how traditional stethoscopes work and how piezosensors can actually be mounted and used.

Although Matt and Johnathan had previous experience designing and ordering PCBs, there was learning that occurred on this front as well. Being in an environment that was somewhat time and cost-sensitive forced us to learn about the differences of manufacturing in China and the US, and about the process of getting many different quotes at once and choosing the right fab. We also learned to a more professional standard how to generate Gerber files and communicate with the fab effectively after they caught some errors with our initial output package. Hal Paver helped us learn a lot in this area.

On the software side, Jared taught himself a number of web tools and mathematics. The preprocessing techniques that were implemented from the PhysioNet paper involve complex math and signal processing, including applications of Fourier transforms, discrete cosine transforms, and filter banks. Jared also had to learn about server programming in Node.js, writing HTTP clients on embedded devices, working with cloud instances on Amazon Web Services, and a number of techniques to integrate MATLAB and Python scripts for preprocessing and filtering with a standard web server.

While Jason has taken many machine learning classes, the skill of reading literature, designing a novel CNN, and collecting results is complicated. Whenever machine learning is involved, using classroom knowledge on a real application always has challenges. We needed to become even more familiar with Pytorch, a python library designed specifically for building neural networks. In addition, learning new concepts related to neural nets was a challenge. Ideas like dropout (which prevents overfitting), different loss functions, and normalization, were briefly covered in classes, but not in enough detail for this project, so we needed to study these more in depth.

B.

On the hardware side, classes that were most helpful were ESE215 and ESE319 for the analog circuitry, ESE218 for understanding how the different sensors worked, ESE319 and ESE292 for creating the circuit board, ESE292 for creating the mechanical mount, ESE224 and BE521 for signal processing to make sure the signal was being picked up well. Other knowledge that was useful was Johnathan’s experience with PCBs from Penn Electric Racing and Tesla and Matt’s experience with microphones from working on Cortana at Microsoft.

On the software side, CIS519 and CIS680 were the most helpful in developing the CNN, as those classes specifically cover neural nets and reading existing literature. ESE350 was also helpful in learning how to create a complete, integrated product with a microcontroller, reading analog values, and streaming via WiFi. Additionally, CIS197 was very useful in learning how to serve static and dynamic content on a server, as well as how to host one remotely.

VI. Ethical and Professional Responsibilities

Most people don't know that heart murmurs exist in 40-45% of children and 10% of adults at some point in their lives.[3] Sudden Cardiac Death (SCD) in children is very rare in the US, but claims .6 to 6.2 deaths per 100,000 children in the US per year, 20-25% of which occur during sports.[4] With about 75 million children in the US,[5] this means that over 4,000 children die of unknown heart reasons each year, about 1,000 during physical activity, and that up to about 33 million children in the US with some level of heart murmur that could potentially put them at risk. While most heart murmurs and arrhythmias are benign, we believe that all children and parents should have easier access to cardiac knowledge for better understanding and more holistic risk assessment of engaging in high stress physical activities. Globally, with about 2 billion children, one could project these statistics with a 27x multiplier (US children to world children), but due to the differing nature of healthcare in other countries we are hesitant to do so.

Economically, the home healthcare market is projected to be worth $517.23 billion by 2025.[6] More specifically, the fitness tracker market is expected to reach $62.13 billion by 2023.[7] There is a clear market desire for healthcare outside of the hospital. Much of these purchases will be driven by desire for heart-monitoring features such as on the new Apple watch series 4, and we expect to be able to capture much of the higher end of that market focused on more advanced heart monitoring features. In terms of replacement cost, an initial doctor's assessment (before getting any kind of echocardiogram or stress test) can cost between $200 and $400.[8] While it is true that often insurance covers this cost to the patient, the cost still exists, and it's possible that if our device replaced much of the time needed in an initial diagnosis (or replaced it all together) by presenting an accurate, temporal, and automatically analyzed heart log, that health insurers would subsidize the cost of the device to ultimately save money spent on doctors' time.

Potential ethical issues fall into the categories of false negatives and data privacy. In terms of false negatives, we will not market our device to indicate that it is a complete replacement for a doctor. Patients are encouraged to visit the doctor yearly for a checkup, and should have their chest listened to by a stethoscope while they are there in case the device is missing something. In terms of data privacy, a concern could arise if using acquired patient data to further train our machine learning classification model. If we decide to do this, we will do it in an anonymized way such that training data is not identifiable to a specific patient.

VII. Meetings

---

[3] https://www.aurorahealthcare.org/patients-visitors/blog/5-vital-answers-you-should-know-heart-murmurs
[4] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3017912/
[5] https://www.childtrends.org/indicators/number-of-children
[6] https://www.statista.com/statistics/678737/total-number-of-children-worldwide/
[7] https://www.grandviewresearch.com/press-release/global-home-healthcare-market
[8] https://www.angieslist.com/articles/how-be-heart-smart-cardiology-costs.htm

At the beginning of the first semester, we met with many different professors to help identify what our project would be. Once we settled on a project, we met with our advisor James Weimer about once every 1-2 weeks on Fridays in his office. We gave Professor Weimer updates about our progress in these meetings and he would give us great advice (and sometimes hardware) to point us in the correct direction. Our team had a weekly meeting without our advisor scheduled for Thursday morning each week during the spring semester where we would set goals for the next week while sitting in the MBA cafe of Huntsman. We often worked together in Detkin and would discuss our next steps there or over our Facebook Messenger in between our official weekly team meetings. The weekly team meetings were essential to making sure we were still on target for our milestones and helped us keep the big picture on our minds rather than just the specific subtasks we happened to be caught up with at the time.

VIII. Proposed schedule with milestones

Below are our spring semester milestones that were created at the end of the fall semester. We updated this table below to include a column to indicate the degree of completion for each milestone.

Table 1: Spring semester milestones

What	Why	When	Who	**Completed?**
I. v2 microphone fully integrated	Use a technology that in inherently less susceptible to noise, and tune our own amplifier and filter	Before Spring Break	Matt and Johnathan	Yes, the v2 microphone with a piezoresistive sensor was fully integrated.
II. Wearable form factor and begin testing	To make a product rather than just tech, need signal to see if classifier will work in the real world	Middle of the semester and improvement throughout	Matt and Johnathan	Yes, our strap-like design was tested on human bodies.
III. More external data + analyzing our own to train neural net	Total product integration and to make our product robust	Beginning of the semester for external, own analysis after I.	All	No. Instead, we chose to improve preprocessing on our model's input data.

IVa. Pare down NN and implement in C to run on microcontroller	Total product integration, difficult but more completely integrated	Middle of the semester	Jason and Jared	No, the neural net has remained on an AWS EC2 server for maximal computation power and ease of configuration.
IVb. Develop mobile application and Bluetooth streaming to run existing NN	Total product integration, easier to implement but less completely integrated	Towards the end of the semester (after IVb)	Jason and Jared	Yes, but our design pivoted to a web application that incorporates Wi-Fi streaming instead of Bluetooth.

IX. Discussion of teamwork

A.

Our team coordinated work through our team's Facebook Messenger group and through weekly team meetings. We also had a team Slack Channel with our advisor in it which we used to give updates and schedule meetings with him that were roughly every 1-2 weeks.

Our basic team structure was that Johnathan and Matt were the hardware team, and Jared and Jason were the software team. There was much collaboration within and between these two-person teams, and there were very few tasks that were solely accomplished by one person. Generally, Johnathan drove the digital and PCB development, and Matt worked with the sensor and analog circuitry. Jared worked on streaming data from the board and our cloud platform, and Jason worked on the machine learning model that ran in the cloud.

B.

Not applicable, we are not ISD.

X. Budget and justification

A low COGS for our device was a specific design goal of ours throughout this project. The final device has a COGS of $60, mostly driven by a sensor, WiFi radio, printed circuit board, and the elastic chest strap. See Appendix F for the COGS breakdown. However, our budget for the project is very different than the simple COGS of one device. We learned and experimented with many technologies along the way.

Last semester, we were unsure if we would use a capacitive or piezoelectric microphone for our project. We also spent ~$350 on traditional stethoscopes and capacitive microphones that we tested with throughout the year. Some of this was offset by the budget we won by applying and becoming finalists for the Rothberg Catalyzer in the fall. We also spent ~$300 on various piezosensors and a piezosensing kit that came with a guide to the technology as well as many different types of sensors to try. This was invaluable for us to learn about how the technology worked and helped us settle on the exact sensor we would eventually use.

Another important aspect of our budgeting was getting PCBs fabricated. We spent ~$1000 fabricating the 4-layer boards in the US. We tried to get boards made in China for ~$150 but the team there was very unresponsive, and the boards came in one and a half weeks late (after demo day) and were not created correctly. We then spent another ~$1000 to get our final version of the board made for the SEAS Competition after ESE demo day. Both times we got PCB spins we spent about ~$250 on electronics parts so that we could assemble 5-10 boards. We discovered that because of the way PCB manufacturing works, it is the same cost to have 1 or 10 boards made since they all fit on the same standard sized panel.

XI. Standards and compliance

The most important standard we comply to is the FDA's regulation of medical devices. We aim to be classified as a Class II medical device, the category under which most medical devices fall. While manual stethoscopes are covered in Class I, we believe the Class II category is necessary for or device because it includes an electronic component that communicates with a WiFi network. While the hazards are minimal, having a Wi-Fi chip against one's chest always poses some safety risks. The Apples Watch Series 4 is also considered a Class II device due to its EKG and fall detection capabilities, and it is actually exempt from full FDA approval[9]. This is promising because, like the Apple Watch, we do not believe the Class II label would pose a huge threat for our ability to bring our device to market. We will be careful in the labeling of our device to make sure users do not think it is a full replacement for a doctor. This should lower any risk the FDA would associate with our device, and would almost certainly make it a Class II device. We will make sure we stay in compliance with hardware standards such as the IEEE 1801-2015 low power device spec[10], but this is mostly taken care of for us because we are using off-the-shelf electronics components whose manufacturers have already done the work of adhering to these standards.

In addition to the FDA's medical device regulation, we are also complying to the HIPAA data privacy standard. Our platform will be storing user's heartbeat data in the cloud, so making sure that customer's data is secure and private is a top priority. Moreover, we designed our product around the IEEE 11073-10441 standard for cardiovascular activity monitors, which ensures that there is a normative definition of communication established between personal

---

[9]https://www.theverge.com/2018/9/13/17855006/apple-watch-series-4-ekg-fda-approved-vs-cleared-meaning-safe
[10] https://standards.ieee.org/standard/1801-2015.html

telehealth cardiovascular fitness and activity monitor devices and compute engines (e.g. cloud servers in our case) in a manner that enables plug-and-play interoperability[11]. Finally, we chose not to pursue the course of action to obtain an IRB approval. Should the project have gone far as collecting data from human subjects and to train our machine learning model with, then IRB approval would be required. However, this was not necessary given our readjusted goals this semester.

XII. Work done since last semester

At the end of last semester we had demo versions of two different microphone technologies in order to capture the PCG, and a basic neural net classifier in the cloud that had 75%. We also had some basic capabilities to log data to an SD card and transmit data over bluetooth. However, we were nowhere near a finished, working, actual product. There was much work that had to be done in terms of packaging everything up so that the signal could follow the chain from the sensor, to the microcontroller, and into the cloud and through the machine learning model. To make this possible, a PCB was designed to put all the hardware in one place, with an onboard WiFi radio to send the data to the cloud. This semester the backend systems that received the data were also upgraded to allow real-time data streaming from the device, and to run the software filtering and machine learning classification in the background in real time as well.

In addition to this systems-level work, we made each piece of our project better this semester. We spent a significant amount of time testing different piezosensor architectures and finally choosing the right sensor. Then all of the preprocessing work in analog circuitry, to magnify the desired signal frequencies and move the signal into the correct range for the ADC to capture it well, was designed and built this semester. Significant time was also invested in using an existing audio recorder to capture sounds and analyze and compare them in Matlab against the PCG sounds we found from the PhysioNet challenge, experimenting with different filtering and amplification combinations for preprocessing. Software filtering was designed to match the real world recorded signals to be as close as possible to those as found from PhysioNet in both the temporal and frequency domains.

On the machine learning side, one of the main accomplishments of this semester was to update the classifier model from a simple neural net that gave 75% accuracy to a more sophisticated model that was tailored to our application and would provide better results. This required reading the literature, identifying a desired architecture (we used a model adapted from *Recognizing Abnormal Heart Sounds Using Deep Learning,* citation 1*)*, implementing the model, and integrating it into the cloud platform.

XIII. Discussion and conclusion

This past year of ESE Senior Design has been an amazing ride. We were honored to represent ESE as Finalists at the M&T Summit and receive an Honorable Mention at the SEAS

[11] https://standards.ieee.org/standard/11073-10441-2013.html

Competition this semester. However, we are even more proud of all of the engineering work we did over the past year to make this product work. When deciding what project our team wanted to pursue, we wanted to make something that doctors would actually find useful, and we wanted to work on a project that would involve serious engineering work that spanned our skills across all levels of ESE–from working with sensors and analog circuitry to PCBs and microcontrollers to cloud technologies, signal processing, and machine learning. We were successful on both fronts. While each of these pieces was its own technological challenge, we discovered that the systems-level integration of getting everything working together at the level at which data could be live streamed with AI analytics in real-time was incredibly more challenging than we had anticipated.

In summary, we successfully created a wearable device that picks up the PCG sound directly from your heart while filtering out unwanted noise, digitizes the signal, sounds the signal to the cloud, and a cloud platform that runs filtering, segmentation, and machine learning classification on the real time data to indicate whether the wearer is experiencing a heart murmur or not.

Potential challenges that might need to be met if this project were to move further towards productization would occur after obtaining an IRB approval to test on human patients. The machine learning models would need to be retrained to data exclusively gathered by our device from labeled patients, or the analog circuitry would need to be redesigned so that the frequency profile of the captured PCG more closely matches those from the PhysioNet dataset that the models were originally trained on. This would bring the classification accuracy even higher. In addition to classification accuracy, for real world productization we would need to do more work to ensure that both are false-positive and false-negative rates would be appropriate for each specific use case the product could be used for. This would be accomplished through extensive user and stakeholder research, and might require further classifier parameter tweaks. The battery selection would also need to be optimized to achieve both our battery life (4 hours) and product weight (60g) goals.

In terms of general life lessons learned, one very important one was how to effectively manage a team of peers with no manager or team captain by subdividing tasks and holding each other accountable in weekly meetings. Another important skill that we developed specifically in the second semester was learning how to leverage our resources and learn more efficiently. For example, while the trial and error of different sensors we endeavoured in the fall helped us move in the right direction for our project, meeting with Jay Zemel and learning from him the physics behind these devices made us much more efficient in our search for the right sensing architecture. Another example is when our advisor James Weimer spoke from his experience about what kinds of chips are popular for IoT medical projects, and was able to both show us the microcontroller we would want to use and also secured one for us on the spot from one of his colleagues. One more example is when the PCB fab was having difficulty understanding exactly what we wanted on the board design, Hal Paver showed us an important trick of using a third-party Gerber viewer to ensure that Altium wasn't misrepresenting anything because we did the actual design in the same program (it turned out it was). The overarching lesson here is that there is no substitute for experience, and that asking the right questions of your colleagues and mentors results in a significantly more efficient usage of time in engineering work. Thus, we exit our undergrad years of Penn with a sense of pride in all that we have learned and accomplished in this project and throughout the past four years, but also with the soft skills and deep sense of humility required to learn from and work with our colleagues and superiors in the world outside.

XIV. Business Analysis

For Professor Babin's EAS 546, HeartWare was used to perform this business analysis with the help of Angelica Du, Angel Fan, Olivia Huang, & Palmer Paul

Value Proposition & Stakeholders

HeartWare keeps athletes safe and empowered to perform by guarding their heart health and communicating issues they would otherwise overlook. By combining a wearable heart-sound monitor with real-time analytics and advanced machine learning technology, HeartWare reduces the risk of cardiac problems that are plaguing young athletes the most, including sudden cardiac arrest and arrhythmias.

HeartWare is a heart sound monitor patch that adheres to the user's chest and transmits audio data directly to the cloud via WiFi. Machine learning algorithms then analyze the audio for signs of impending problems and, if necessary, an alert is sent instantly to a relevant trainer, coach, or staff member who can take action or acquire professional medical attention. In addition, coaches and training staff are provided periodic performance insights based on heart audio and heart rate information so they can see how athletes are performing relative to their baseline and to the other players on the field.

HeartWare follows a subscription revenue model where the user makes an initial purchase of $160 for the device and first year of subscription. For every subsequent quarter, the user pays a $40 subscription fee (per device) for access to our online analytics and alert platform. Over time, our subscription model (recurring revenue) will become a larger and higher-margin portion of our total revenue each period. We will become profitable by 2023 and our total revenue in 2024, our 5th year in operation, will be $6.5 million.

Our total addressable market (TAM) includes all competitive high-school and collegiate athletes across the nation who may or may not be predisposed to cardiac illness, of which there are close to 500,000 in the NCAA alone. The cardiac monitoring and cardiac rhythm management industry is growing rapidly, expected to reach $26B by 2024 (Market Research Engine, 2019). The capital market is also favorable; investment in this sector has demonstrated annual growths between 80% and 120% from the past few years, according to Thakur (2017).

Our competitors include firms selling chest-strap heart rate monitors for general use, such as Polar and TICKR X (Stables 2019), as well as companies developing wearable EKG devices for use in medical settings, such as CardioNet, iRhythm, and Eko. However, customers are dissatisfied with these current solutions, citing common issues like physical discomfort, tedious manual analytics, and a lack actionable information and a user-friendly alert system as reasons they would be willing to switch to a better device.

Therefore, our strategy involves filling this market gap by promoting 1) comfortable wearability, 2) advanced auscultation metrics and 3) real-time analytics and alert platforms in a single device. Our marketing message to our primary stakeholders - the trainers, coaches, and athletes who have grown disillusioned with a saturated market of look-alike heart rate monitors -- will be that we are a medical-grade device that delivers advanced, critical insights in real time, but is suitable for use by active individuals outside of the hospital space. Indeed, “We are *not* just another heart rate monitor.” To reach individuals in our entry market - athletic programs, we will first directly contact college sports programs with histories of cardiac emergencies and attend sports medicine conferences to make personal connections within the athletic training community. We will also organize follow-ups after delivery, to provide instruction, foster relationships with teams, and gather and incorporate feedback rapidly.

Our management team is a driven group of bio, computer, systems, and software engineers, with experience in key areas for this precise type of project, including physiology, medical device development, and machine learning. In addition, we will bring in advisors with biomedical expertise, including a small scientific advisory board of 3-5 cardiologists, sports medicine doctors and medical device designers to help us shape our technology, make decisions on clinical trial/FDA approval, and raise our company credibility. We hope to build a strong advisory relationship with our VC partners, from whom we can garner critical financial expertise as well.

HeartWare is seeking an angel funding round of $300,000 by the end of 2020. Our first year will be focused on R&D, obtaining Class II medical device approval from the FDA and beginning to implement this direct sales strategy. After producing initial revenues of $1 million in the first two years, we plan to procure a Series A investment of $500,000 by the end of 2021 to expand our R&D and sales outreach, access more schools, and capture a larger portion of our TAM with superior technology and design. Between years 2022-2023, we plan to increase our sales team to 5 regional senior managers (1 located in each of the main athletic conferences), who will each handle up to 3 commissioned sales agents local to their conference region.

In 2024, we plan to exit via acquisition by an established medical device firm with a special interest in entering the cardiac monitoring industry. Seed investors can expect an ROI of 50%, while Series A investors can expect an ROI of 35%. We believe HeartWare will save athletes’ lives and keep them healthy, and that our founding team is well equipped to deliver a product that solves this problem and addresses the problems that coaches, trainers, and athletes face.

Product and Company Overview

HeartWare is a thin, flower-shaped electronic stethoscope, to be worn wirelessly as a patch on the user’s chest (Figure 2). It rests above the heart, uses an amplifier and microphone to record the user’s heartbeat audio, and then transmits the recording to the user’s phone via WiFi, where it is uploaded automatically to the HeartWare analytics platform in the cloud. Once there, machine learning algorithms identify any potential problems with the heartbeat sound and notify the user immediately if they should seek further medical attention.

HeartWare combines continuous monitoring and advanced analytics into a comfortable, wearable device, to deliver real-time heart sound data instantly to the user or related persons in the absence of medical professionals. Our product is unique in that it is designed to be worn during intense exercise, which means it can capture heart murmurs and arrhythmias that may be precursors to exercise-induced sudden cardiac arrests. We are pricing our product at $160 for the device and first year subscription combined, and a $40 / quarter subscription fee thereafter.

Our machine learning model for detecting heart abnormalities is going to be a trade secret. Additionally, we will make abundant use of copyright to protect the source code for the app, training the machine learning model, the website, etc. We will also use copyright to protect any user documentation/manuals and will trademark the name "HeartWare."

We also see a potential secondary source for revenue in licensing out the analytics platform to other wearable device OEMs. We envision the platform becoming a one-stop shop for people to check on their overall wellbeing and making it possible for other wearable devices plug into our platform is critical to achieving that vision.

HeartWare is a C-corporation located in Philadelphia, PA. We chose Philadelphia given the founders' affiliations with the University of Pennsylvania and established connections in the area. HeartWare decided to form a C-corp to more easily raise multiple rounds of financing and make avail of the more attractive taxation regulations for outside investors.

Industry and Market Analysis

Wearable Monitoring Device Industry is Growing

HeartWare falls under the industry of 'cardiac monitoring and rhythm management devices,' as well as that of 'diagnostic and monitoring wearable devices.' The 'cardiac monitoring and rhythm management devices' industry is progressing at a compound annual growth rate (CAGR) of 4% (MarketWatch, 2019), having garnered high demand given the growing population of aging individuals suffering from cardiac conditions and expansion of new technology for monitoring heart rate, electrical signals, and auscultations. With the benefits of mobility, flexibility, and connectivity making wearable devices increasingly attractive for consumers, more and more medical device firms are investing in developing wearable versions of traditional cardiac health tools. As such, the capital market for global diagnostic and monitoring wearable devices has also been favorable, with investment in this sector demonstrating annual growths between 80% and 120% from the past few years (Ritu Thakur, 2017).

Though the overall industry is growing rapidly and the investor climate is active, there are high barriers to entry that must be considered. Primarily, customer affordability poses a barrier for new entrants due to the high cost of incorporating sensitive, accurate detection systems, large data storage capacities, and reliable battery lives into cardiac devices, and the pressure to price competitively. Federal insurance coverage further contributes to this barrier, as reimbursement programs for affected target groups (the elderly and disabled, etc) are usually inadequate for wearables, making them less financially accessible for users. Because purchases are such a large percentage of the customer's own cost, the force of customer bargaining power is a potential threat. Other barriers to entry include a tedious regulatory structure, with cardiac auscultation monitors requiring FDA clearance (Class II, which requires premarket notification, and the 510(k) application), and tied-up intellectual property; for example, some proprietary rights for heart auscultation monitoring algorithms are currently held by Motorola. To penetrate and achieve dominance in this industry, we will need to tailor our strategy around the several non-scale cost disadvantages (in the form of proprietary rights) and high R&D capital requirements for new firms.

9 Million Athletes in TAM

As reported by Tatkare (2016), the market for cardiac monitoring and rhythm management is expected to reach $32B by 2022. The total global revenues for cardiac rhythm monitoring market segment surpassed $14 million in 2015 and is expected to exceed $18 million by 2020 with a CAGR of 5.4% (Elder, 2015). An important trend driving this market growth is the emphasis on technology-driven improvement, with medical device firms emphasizing more convenient data transfer capabilities, higher-accuracy detection systems, longer battery life, and more advanced features to gain market attention. HeartWare intends to expound on this trend by delivering high-quality analytics technology at a competitive price.

The target market for HeartWare includes ~9 million competitive high-school and collegiate level athletes (NCAA, 2019). We selected this market after researching how universities invest in, profit upon, and protect their athletic programs, and their decision-making processes when it sports medicine equipment spending. NCAA FBS teams, a very specific subset of our target market, collectively had $8 billion in expenses in 2017, out of which over $100 million went to medical expenses and over $1.75 billion were spent on facilities and equipment (Knight Commission, 2017).

At this time, the dominant heart monitoring devices used by athletic programs are traditional chest-strap heart rate monitors, which garnered a global revenue of $439 million in 2015 and is expected to reach $1.3 billion by 2020 (Wearable Devices in Healthcare, 2015). Since we are offering a professional cardiac monitoring device in a wearable form, our market will draw from small subsets of both of two market segments: cardiac rhythm monitoring and wearables. If we are able to get a 15% market share from the two markets segments combined, our expected global annual revenue would become $198 million.

HeartWare Addresses Needs of Coaches & Athletes

HeartWare's initial market will be collegiate and high school athletics programs. Sudden cardiac death is the most common medical cause of death in athletes, afflicting 1/40,000 to 1/80,000 athletes per year (Harmon, 2014; Wasfy, 2016). Especially at the upper levels, athletes are required to push their physical limits, sometimes undergoing greater than 40 hours a week in high-intensity training and competition (Jacobs, 2015). Because athletes are accustomed to exertion, however, they can habituate to pain, ignoring symptoms that may normally warrant more serious medical attention. Coaches and trainers are responsible for encouraging high performance in their athletes while ensuring their safety.

HeartWare will address this need by providing coaches and trainers with actionable analytics in real-time so that players who are experiencing cardiac distress can get medical attention faster. The product's continuous heart sound monitoring feature will allow staff to instantly access their athlete's heart data, decrease the need for professional visits, and ultimately reduce the chances of missing signs of heart issues, which athletes may be more likely to ignore. Value can be measured quantitatively, in the number of warnings caught, as well as qualitatively, in consumer satisfaction/feedback.

Based on our primary research, the decision-makers in collegiate athletic programs are the coaches for the individual teams or an athletics director that oversees several teams. In most programs, coaches have the sole power to decide which products to purchase for their teams. They either purchase the devices themselves directly, or they place a purchase order through a standardized university system. In either case, once the coaches decide to purchase our product, they typically have the power to directly complete and manage the purchase. Ultimately, the coaches (who mandate the athletes wear the technology) benefit from the performance insights and the safety alerts, so they are the end users.

The athletes wear the device and are major influencers over the decision-making processes. Based on our discussions with college athletes, we concluded that if equipment impedes the athlete's performance or is uncomfortable, their coaches would not make them wear it. This was confirmed in our interviews with coaches and trainers who expressed they would only want equipment that is comfortable for athletes. Thus, athletes have a heavy influence over decision making.

Other influencers include university administrators, such as the Provost, Dean of Students, and President, have an incentive to avoid the bad press associated with an athletic medical emergency occurring to one of their athletes. They also can place significant pressure on the athletic programs and coaches giving them immense influence over the successful adoption of our product. Any positive press about how the program is focusing on player safety by incorporating new technologies into their workouts or play is also a huge win for the program.

In addition, the athletes' parents' priority is their children's safety and are capable of influencing coaches by personally entreating them to adopt our product to help protect their children.

These athletics programs largely learn about new products via word-of-mouth, trade shows, or reports from other programs. We spoke with several trainers and coaches who all had a top priority of keeping their athletes healthy and who invested time and money (sometimes using their own vacation time) attending trade shows and searching for new products. If a program begins using a new service and has a positive, successful experience, the technology is likely to spread among other universities. Athletic directors and team coordinators read news stories and attend athletic conventions. One of the biggest trade shows for us will be the Athletic Business Show. This is a major way that teams learn about newer technologies that are applicable to their programs.

While we do not intend on having any channel partners currently, we do see the potential to partner with a traditional sports equipment distributor should we decide in the future to open our product up to the general athletics and fitness market.

Existing Cardiac Monitoring Devices are Lacking

We observed that most cardiac monitoring and rhythm management products primarily offer either heart rate monitors or devices that provide home-based cardiac screening with advantages in brand identity, capital, and experience, but also feature little product differentiation. Most feature some combination of 1) wearability, 2) real-time, continuous monitoring, and 3) automated heart-health insights, but not all three. Therefore, our strategy involves filling this market gap by combining these features in one device, while focusing on auscultations instead of heart rate to differentiate ourselves from the industry entirely.

Chest strap heart monitors are the most significant direct competitor that we have in the athletics programs customer segment. These are devices that athletes wrap around their torso to measure and monitor their heart rate in real-time during workouts. Athletes prefer chest strap monitors because they are relatively unobtrusive and provide more accurate monitoring than wristband solutions during intense activity. Two commercial offerings in this category are the Polar H10 and the TICKR X Heart Rate 3 Monitor (Stables 2019). Both devices are integrated with their respective companies' fitness platforms which have mobile apps for viewing your fitness data. The fitness apps are ancillary services that the companies provide. The Polar H10 costs $89.95 ("H10 heart rate sensor") and the TICKR X costs $79.99 ("TICKR X", 2018). Both of these products are marketed towards athletes and can be purchased on the companies' websites or through online/brick-and-mortar retail distributors. Both companies also have very strong branding as the company to go to for chest strap heart monitors. While these are wearable devices that can provide real-time basic analytics, they only capture raw heart rate data, nothing more. Coaches and trainers expressed that they do not know how to interpret this data to derive performance insights or health insights, so these products are not good enough for them.

There are competing products that are within the medical space, not the athletics market. For example, there are traditional and electronic stethoscopes. Traditional stethoscopes must be operated by highly-trained individuals who know how to properly listen for structural cardiac defects. One of the most well-regarded stethoscope producers is Littmann ("The Choice of Professionals"). Littmann prides itself in making high quality, consistent stethoscopes that cardiologists can trust, a fact they highlight throughout on their website. Littmann is a subsidiary of the 3M conglomerate, which they can leverage by sharing their costs and management. The Littmann Master Cardiology Stethoscope retails for around $300. Traditional stethoscopes cannot be used for continuous monitoring.

The Eko CORE Stethoscope is a combination of an acoustic and electronic stethoscope. It is composed of an attachment and can be added onto existing stethoscopes (E, 2018). The CORE Stethoscope allows the user to listen to heart or lung sounds and transfer them to the Eko App, where the sounds can be visualized to better detect murmurs. The CORE is one of our indirect competitors because they also offer electronic analysis of heart and lung sounds. The entire stethoscope costs $299 and the CORE attachment by itself costs $199 (Goolkasian, 2019), which is more than the typical acoustic stethoscope, which can cost as low as $160 ("How Much Does a Stethoscope Cost?"). The product has been out for a few years and it is sold both on Eko's own website and on different online distributors. The stethoscope is marketed to doctors and there have been a positive few reviews online done by physicians. This device is not wearable and does not provide automated health insights, however.

There is also a category of competitive devices that monitors the electrical characteristics of the heart. Any ECG device that's currently on the market is also one of our direct competitors. Holter monitors are battery-operated devices that are either attached to the user's waist, taped to their body, or goes around a lanyard on the user's neck (Holter, 2018). The monitor is attached to multiple electrodes that constantly gathers data on the user's heart rhythm. There are many companies that produce Holters and they are currently the most popular form of portable ECG monitoring. It is less expensive than the other competitors described above. Current Holter monitors require users to wear it for a few days and then return the monitor to their physician for data analysis. The typical Holter is sold to hospitals and doctors who will then prescribe them to patients. However, Holter monitors are typically bulky and not wearable for athletics.

There are also wearable forms of ECGs, such as the iRhythm's Zio XT system. Zio is a small, pocket-watch sized, water-resistant patch that rests above the patient's heart ("Zio XT Patient Pamphlet"). The patch gathers data by performing an electrocardiogram (ECG) on the user's heart and users can also press a button on the patch when the user feels a symptom. After the user has worn the Zio monitor for the duration prescribed by their doctor, the user will remove the patch and mail it back to iRhythm where the iRhythm team will create a report on the data gathered by the monitor using an artificial intelligence algorithm (Tirrell, 2017). Then, the report will be shared with the user's doctor.

iRhythm's Zio is our major competitor because over one million patients have already used the Zio system for arrhythmia diagnosis. The system has already been cleared by the FDA to monitor many different kinds of irregular heart rhythms. Since their incorporation in 2006, iRhythm's main strategy is to do direct sales to get doctors to prescribe the Zio system for their patients. Their website is aimed at persuading physicians that Zio has a higher diagnostic yield over competitors and the images on their website also suggests that the product is directed at the elderly. The Zio system costs $360 for those on Medicare, which is higher than the $100-$150 pricetag of Holter monitors ("Holter Monitoring"). So, they target a subset who are willing to pay a premium for more portability and accuracy. iRhythm also does not seem to have any other products but they do a good job of promoting their Zio system. However, this device is not real-time, meaning it cannot analyze the heart health and alert coaches to any issues while the athlete is still on the field. Thus, this device does not meet the needs of coaching staff.

One final competitor is Biotricity, who makes Bioflux, a real-time, high-precision ECG monitoring system. Bioflux has some features that typical Holters do not such as mobile cardiac telemetry monitoring and automatic cloud upload and analysis of cardiac data in real-time ("Bioflux"). The user does not have to remove the device to upload data to their physician. Bioflux was launched in 2018 and it has been named the Best Remote Patient Monitoring Solution by MedTech Breakthrough (McCoin, 2019). Biotricity's strategy is to do direct sales across six states, selling directly to physicians who will then prescribe Bioflux to their patients. The company plans on expanding the current sales force and to expand their operations.

Biotricity is also planning on expanding their portfolio by filing for Biopatch, which is an alternative to the 3-lead system.

From our primary research -- interviews with coaches, athletes, and trainers, we found corroborating evidence that athletes only use heart rate monitors during training, as high-accuracy auscultation monitoring wearables do not exist on the market yet and customer awareness of other monitoring options is low. We also identified major problems with these heart rate monitors, and in accordance, we identified ways to incorporate better solutions into our own product. For example, EKG/ECG devices are bulky and are perceived to hinder athletes from performing at their best; moreover, few alert the user when something is actually concerning. Thus, we are designing HeartWare to be thin, compact, and comfortable enough such that athletes can wear it during intense workouts and even during games, and we are pairing it with a software system that performs the data analytics in real-time on the users' behalf. In these ways, HeartWare can generate demand in the market by providing a more comfortable, more wearable cardiac monitoring device that could also double as an alert system. Another important differentiating factor for our product which can create an entry barrier for other competitors is our proprietary technology (i.e. the machine learning algorithms that are comparable to professional cardiologist evaluations).

Sales and Marketing Plan

Not Just Another Heart Rate Monitor

We plan to target athletic programs as an entrypoint, offering more comprehensive monitoring than the heart rate monitors that athletes currently use. We will focus on addressing customer dissatisfaction with the current dominant heart monitoring brands (Polar heart rate monitors, primarily) as indicated by our primary research. We are positioning the product as a safety and educational tool, and prioritizing the three features that matter most to our decision-makers: comfortability, real-time analysis, and significant health insights. We are not just another heart rate monitor- we are a medical device.

However, to get an entrance point into our market, for the first few months, in addition to our safety features, we will also be putting more focus on our actionable performance insights. We will want to market our performance data analysis as better than the other heart monitors on the market.

Subscription Model

Our revenue model involves selling the device and first year subscription in a bundle for $160 and charging a $40 / quarter subscription fee thereafter. The monthly platform subscription has a higher margin than the initial device sale and will be a source of recurring revenue for us. Each subscription will get access to our real-time emergency alert system and to actionable performance insights.

Distribution

Our distribution strategy consists of having our sales team book orders for our products while making their visits to universities. We are going to use FedEx Ground for order fulfillment because one of interviewees who has a lot of experience in operations and order fulfillment told us positive things about the service.

Promotion via Social Media

HeartWare plans on advertising through social media campaigns. First, we will create accounts on social media and create a website homepage to accumulate followers. We will collaborate with social media accounts, who constantly share high-tech sports equipments. We will provide them with our product advertisement showing our product features emphasizing on comfortability and wearability. In order to further raise awareness and gain some prospects, we will run a social media campaign on the significance of sudden cardiac arrest (SCA) amongst athletes and how HeartWare helps save lives (Hyder, 2016). Similar to how the SickKids foundation ran a viral video promotion by sharing stories of patients to appeal at the public's emotions, we want to create videos about athletes that suffered SCA and then show how HeartWare detects heart problems (Krums, 2018). With this campaign, we want to target athletes and parents who may then show their coaches our product and request for it. We want to push the "We Listen" theme and emphasize that HeartWare speaks for those who may know that their hearts hurting.

Direct Sales to College Athletics Teams

HeartWare plan on starting sales around the Philadelphia area because this is where the founding team has the most connections. We will begin pitching the product to nearby schools as soon as we have a working MVP so we can start gathering interest. In the first few months after we finalize the product, the founders will personally reach out to customers we have direct personal connections with in order to create pre-sales. Within the first year, we will also hire 3 sales representatives.

Since our product will be new to its intended market, we will need to educate our customers on the importance of the product. Thus, we will need to teach our salesforce how our product works and how to effectively demonstrate it. The demonstration and pitch are essential for sales so we will need to take some time and go through a trial process to figure out what technique works best. To encourage our sales team to experiment with techniques, we will compensate them based on a salary based pay rather than a commission-based pay so they can focus on learning rather than focusing on achieving sales objectives. As our company grows and we move from the initiation phase, we will hire more sales reps and move into a commission-based pay so we can get more sales.

As we gain a customer base, our sales representatives will start cold-calling/contacting customers who we don't have a direct personal connection with. The representatives will also visit each new athletic program one or two times a year and offer training on the features of the products. As distribution expands, we will adjust to meet greater demand by increasing sales staff to make sure our customers are visited more frequently. This strategy will help us get first-person feedback on our product, which will help with product development and gives us data to show in our promotions and advertisements.

While we are starting in Philadelphia, in the first year, our sales team will expand to nearby large cities on the east coast such as Boston, Washington D.C., and New York City. Members of our founding team are from these locations so they will be able to leverage their personal connections to help make sales. By 2022, which is our third year of operations, we will expand our sales team to other regions in the country. The way that we're expanding our sales operations will be determined by the location of the collegiate athletic conferences. So, we're starting on the east coast with the Big 10, then moving to the ACC, then SEC, then XII, and finally to the Pac 12.

Operations

Overview

We will be keeping the software and hardware design in-house. In particular, we will develop the monitoring and machine learning software internally to ensure that we can rapidly modify it according to the feedback we receive. The software includes the alert software platform that provides red flags to coaches if athletes are experiencing signs of heart issues, as well as performance indicators to indicate whether athletes are in the healthy heart rate range. The engineering design of the physical product will also be kept proprietary for similar reasons, though the parts and materials will be outsourced. We will be subcontracting manufacturing; this will not be a special competency of ours. This way, we can obtain higher quality mechanical and electronic parts more quickly, more conveniently, and less expensively. As this will be a strictly buy and exchange relationship, we are not worried about IP risk.

An important factor for implementing our plan includes gathering engineering, sales, and legal talent. We will need technically skilled hardware and software engineers, sales reps with great communication skills and experience developing their own sales material for initial direct sales, and a legal team that specializes in IP to ensure we are designing the product without infringing on existing patents. We will leverage our existing personal networks to identify and recruit candidates for our team.

There are a few risks with our operations plan. First, there may be existing IP (patents, mostly) that our project, if not designed properly, could infringe on. To mitigate this, we will consult with IP attorneys and either get an FTO or redesign the product to minimize the risk. Second, the market for heart monitors is very competitive, and it may be difficult and expensive to make initial connections with athletic teams in our target market and let the decision-makers understand as to why HeartWare is a better product. We will mitigate this with concept, product, and market testing to ensure the demand for our product is strong and use that the opportunity warrants the time and costs. We will also go to sports equipment conferences, such as the Athletic Business Show, to get in touch with coaches and trainers (AB Show, 2019). Finally, there are always technological risks that can make it complicated, expensive, and time-consuming to engineer a highly accurate, external medical device based on internal heart sounds. However, we will prioritize developing the prototype before pouring costs and time into later endeavors and improvements.

Clear, 4-Phase Plan to Launch

Phase I: Product development. Within this stage, our engineers will be developing hardware to cleanly record and amplify the heartbeat audio signals, based on the market testing we have conducted to gauge customers' needs and purchasing intent. In addition, our software engineers will develop an initial version of the machine learning and data analysis algorithms that will process the recorded heartbeat and provide auscultation and heart rate insights. They will also create the web analytics platform and mobile alert app. The milestone for this phase is the completion of the product to be submitted for formal testing and FDA approval. Phase I will take 3 months and will be completed by the second quarter of 2020.

Phase II: FDA approval and clinical studies. Obtaining FDA approval as a Class II medical device is a crucial milestone for this phase. Concurrently, we will also be conducting clinical studies on the efficacy of product. While clinical studies are generally not required for Class II medical devices going through the 510(k) process, we learned from our primary research that trainers will not consider purchasing a new piece of athletic safety equipment without scientific evidence that it works as advertised. Phase II will be performed in 6 months and will be completed in the beginning of the fourth quarter of 2020.

Phase III: This phase will involve initial training of our direct sales force of 3-4 people. Phase III will be performed in 3 months beginning in the fourth quarter of 2020 and continuing until we make our first sale in the first quarter of 2021.

Phase IV: Marketing and direct sales to target market athletic teams. We will ramp up direct sales to target market teams. This involves sending our sales team to strategically chosen schools in order to acquire our initial customers. A major milestone for phase IV is making our first sale and collecting feedback about why specific teams are purchasing or not purchasing our product. By the end of this phase, we should have 8-10 customers and a repeatable sales model such that we can leverage coin-operated sales representatives to reach more customers. Phase IV is anticipated to take 6 months.

Scaling Plan

As the company grows, we will begin to produce these devices at a larger scale to provide them to larger athletic programs at major universities. This will require more manufacturing and support for a larger number of users with our analytics platform. In 2023, we are going to outsource our manufacturing to Taiwan. Doing so will allow us to minimize high fixed costs (such as for maintaining our own staffed manufacturing plant) and obtain dependable, industry-standard hardware less expensively. We will be better able to focus on our core competencies (design and software), and also re-allocate economic resources to the activities that provide tangible customer value, such product design. Product design, software development, and packaging will take place in-house, at one main office in San Francisco to be shared by executives and engineers. This is a location that provides a balance of affordable materials, existing talent, and access to ports (which at least gives us the option of purchasing parts overseas in the midst of scaling, should we decide to). Centralizing our core competencies will also aid in reducing transportation costs and IP risks.

In order to scale our sales strategy, we will hire a system of senior sales representatives (5 in total, 1 located in each of the main athletic conferences), who will each handle 2 outside sales agents local to their conference region. Though these agents will be conducting the personal visits to colleges, they will not be 'official' HeartWare employees in the sense of receiving a fixed compensation. Instead, they will be paid based on commission. This hired sales rep/unhired agent organizational sales system was suggested by an operations manager we interviewed, and was chosen to keep hiring expenses low while maintaining the amount of ''feet on the ground" as high as necessary during growth. Though a potential risk to this strategy may include a lack of motivation, or worse, a desire to sell at any cost, on the agents' end, we plan to mitigate those risks by paying them a baseline compensation for their efforts (the details of which are still amid discussion) and setting high ethical standards for our company culture.

Current Stage of Development

The HeartWare product needs 7 to 11 months before it is ready for launch. The HeartWare device consists of two main parts: 1) the hardware paired with a machine learning algorithm, and 2) a mobile platform for data recording and monitoring. Because there are existing algorithms capable of detecting heart murmurs better than professional cardiologists (Ren et. al., 2018; AI, 2018), building upon this software technology is feasible. Our key product development activity will be integration of this software with our flexible hardware design.

One team of software engineers will conduct algorithm development. At the same time, another team will build the mobile platform. These simultaneous processes should take 3-6 months. Once the machine learning algorithm is developed, we will create prototype devices to test the platform. We acknowledge that it may be wise to also begin acquiring customer data before algorithm completion, especially for various usability factors of the product (comfortability, mobility, resistance to damage, etc). To get feedback on partial prototypes, samples will be sent to medical schools and collegiate athletic programs for alpha/beta testing. With feedback from the various market segments, new improvements can be made on each aspect of the device based on what they seem to value most. The testing and improvement process will take 1-2 months.

The development of the machine learning algorithm, development of mobile platform, and feedback acquisition are the key engineering activities to be completed before launching. However, in addition to testing and development, we believe our device is a class II medical device as defined by the FDA which would require us to file a 510(k). Historically, approval for devices similar to ours takes about 6 months (How, 2017). We would also need to work out the issue of the patents. Depending on Motorola's willingness to cooperate, obtaining the necessary IP could take a couple of months up to a year. We believe a reasonable estimate for this is 6 months.

Therefore, under the assumption that we will be working on product development and IP concurrently, we predict that it will take one year before we are ready to hit the market.

Funding and Financials

Overview

HeartWare starts with five engineers who sought to save the lives of athletes dying from cardiac-related illnesses. In the beginning, we all pool together $25,000 in founders equity, at $5,000 each and collect another $25,000 in non-shareholder equity. We expect to spend around $220,000 the first year so we held also an initial seed/angel round to raise an additional $300,000 to cover costs. For the first three months, HeartWare's engineering team will work tirelessly to create an initial version of the product. The majority of the initial seed funding is spent on reimbursing the engineers for living expenses and on purchasing parts. During the next six months, the HeartWare team works on MVP development, refining the product, and obtaining FDA approval, and most of the costs go towards legal fees and clinical trials. During this product finalization process, the sales team is preparing for sales by reaching out to initial customers and forming a market penetration strategy.

Upon gaining FDA approval by the beginning of 2021 (our second year of operations), our sales team starts actively selling to customers. We gain our first initial customers and sell 10 units within the month, keeping this momentum up for the rest of the year. On average, each team buys 30 units and within our second year, we make 6,480 sales for $160.00 each. We use a subscription model where first time users need to purchase a full year of subscription for $160, which comes with a complementary device. For the second year onwards, users can choose to subscribe for $40/quarter. Our costs this year come mostly from reimbursing travel fees to our sales representatives and hiring 3 more representatives. We focus on training our sales representatives to find the best way to make sales. At the end of the year, we earn net revenues of $1 million and a net earnings of $(44,500), after also taking into account of the loss from the first year of operation.

In 2021, we will also raise a Series A round of $500,000 to cover the costs of years 3-5. This is because we will be spending $1 million in year 3 to move to a larger location and expand our sales to other regions in the United States. The $500,000 will keep us in the green so we can gain substantially more revenue in future years. We also don't want to raise much more cash than we need because we want to force ourselves to be more frugal and not overspend.

Our sales steadily increase in 2022. We hire two more sales agents who will work under our sales managers and we expand our outreach to neighboring cities on the east coast like New York City and Washington D.C. We sell 10,800 new units and 57,782 monthly subscription payments. Our engineering team continues improving the patch, and our customers provide feedback. Our past customers also continue purchasing our subscriptions (accounting for a churn rate) so past customers' sales accounted for 25% of our sales and new sales accounted for 75%. Over time, as we gain more loyal customers, more of our revenue will come from existing customers. At the end of the year, we earn $2.1 million in revenue and $440k in net profits, which is due to us moving to a larger office, hiring many more employees and focusing on sales expansion. This is our first year with positive net profits and this trend will continue in the future as we spend a smaller percentage of our revenues in operating expenses.

Our fourth year of operations (2023) sees another large growth in our team. We hire more sales associates and engineers, as well as a quality control manager. Our sales continue to grow as we expand across the east coast and we almost double our revenue to $4.3 million as our sales representatives become better trained. Our plan to make most of our revenue works because the percentage of revenue we get from old customers increases from 25% to 35%. While we do spend more on general and administrative expenses this year because this is when we start paying the majority of our founders, overall, we end our year with $1.2 million in net profits.

2024, which is our fifth year of sales, followed the general increase as the previous years and we release updates to our software. At this point, we have gained more than 1500 customers (teams) and we hope to expand our outreach by opening a small office on the west coast, where our founders have personal connections. This year, we have 10 sales representatives and we sold to 360 new schools for 21,600 new units. Our old customers makeup for 47% of our profits, which is a 10 percent increase from the following year. This year, our revenue is $6.6 million and our fixed costs have not changed substantially, such that profitability grows to $2.1 million. We see a good future ahead of us and expect that our growth will continue as we expand across the United States.

$6.5M in Revenue by 2024

We have a subscription based model where users can pay quarterly for the use of our analytics platform. Our revenue model involves charging a first year subscription fee of $160 for a 1 year subscription and complementary device bundle. We estimated units of devices sold based the number of schools that each sales agent can reach. Based on a direct sales model, each year, a sales agent can successfully sell to 72 teams with 30 athletes per team (2160 subscriptions/sales agent/year).

For the first two years of our five-year revenue projection, our focus is to penetrate the market and gain market shares. During our first year, as we will be working on refining our product for clinical trials, applying for FDA approval, and conducting market testing, so we will not generate revenue. We plan to begin using a sales team at the beginning of 2021, when we will focus on reaching initial customers. Our goal is to become profitable at the end of 2021 with a net revenue of $1 million. From 2022 to 2023, we plan to continue gaining market shares and engaging in marketing campaigns with highly influential sports associations, such as the NCAA, to grow our sales. We will become profitable at the beginning of 2023 with a net earnings of over $440k. From 2023 to 2024, we will switch our focus from gaining market shares to maintaining customer retention, because as we grow our revenue increase both from new subscriptions every year and accumulative subscriptions from previous users. To keep users on the subscription service and prevent switching, we will focus on gathering feedback through forums on our website, as well as customer relation channels by employing an HR team. At the end of 2024 we estimate a net revenue of $6.5 million with around half of the sales coming from cumulative subscriptions.

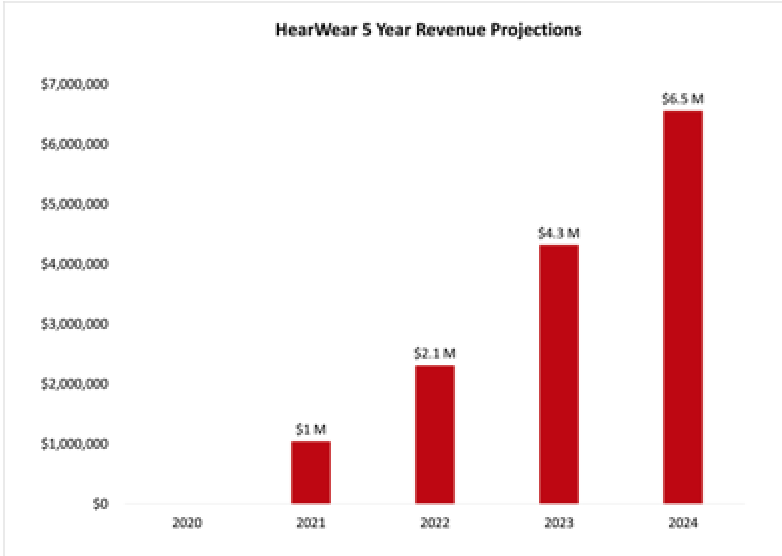

Figure : HeartWare Five-year Revenue Projections

High R&D Expenses in First Year

To elaborate on our product design focus in 2020, we will hire and pay one engineer for R&D of the initial prototypes. The initial engineer will be mainly compensated with equity, as he or she will be the earliest employee to join the company after the founders, as well as small salaries for living costs. Although his salary will be $70,000 which is less than the average engineering salary for Philadelphia, we anticipate they will be drawn in by the equity. During 2021-2022, we will add engineers who will receive a larger salary but less equity as compensation; as a result, the overall salary costs will increase. The CTO will get a raise in 2022 as he will also be incharge of management of three engineers. We will also hire technicians in 2022 to supplement the engineering team and handle lower level technical issues that any university teams may have with the devices. In terms of R&D-related expenses, the vast majority will go towards purchasing new potential sensor technology and cutting-edge components that the engineers believe will improve the product. Though medical device startups usually spend 6-12% of their revenues (Collins, 2018) on R&D, our costs will be low and fixed from 2020 to 2024 as the engineering effort will be geared towards integrating and building a complete product rather than experimenting with new technologies. Once a product is completed and sales have begun, there can be more resources allocated to R&D for newer, more accurate sensing technologies..

Sales/Marketing Expenses Increase Over Time

To elaborate on forming our sales team: after we have an initial prototype, we will hire 3 sales representatives to start surveying customers. These representatives will be compensated at a salary of around $50,000. In 2022, we will high 4 sales agents in total and one sales agent will be promoted to sales manager to gain further reach as we venture into more cities in Pennsylvania and along the east coast. The number of sales agents we hire will be heavily dependent on the sales we make during the previous years. In 2023, we will promote 5 sales managers and give them a raise with each of them in charge of one main athletic conference. As we make more sales, we will also hire a part-time customer service representative starting in 2021. Like the other four founders, the CMO will not start receiving a salary until 2023.

Most of the sales expenses will come from traveling, but we will lower those costs for the first few years by only traveling within a drivable distance of Philadelphia or traveling along the east coast to large cities that are easily accessible by bus or train. Since we will not be hiring additional marketing staff, we will hire freelance designers, photographers, and videographers as needed for our website design and social media promotions.

Other Expenses

General and Administrative Expenses

Under G&A expenses, we will record the costs for the following resources/activities: legal and accounting, office supplies and printing, payroll taxes, rent, officer salaries, non-sales travel, and utilities. (The critical activities of R&D and sales & marketing will not be considered under G&A, and have instead been designated separate expense documents.) Items such as rent and utilities will be constant over the five year period and dependent on our location in Philadelphia; in general, we want these expenses to be on the lower end (<$1500 per quarter) for 2020-2022. Two critical items on the sheet include legal and accounting fees, as well as officer salaries. Legal and accounting fees for startups tend to increase along with the number of founders involved, with future plans to raise third party capital, and with an increased public presence (Roberts, 2008). Since all of these scenarios apply -- in that we have 5 founders, the intention to raise venture capital within five years of launch, and designated plans to maintain a highly public, media-heavy presence in the collegiate athletic space, we will have to prepare to invest more into the legal activities throughout our first five years ($3,000 per quarter). These expenses will go toward obtaining trademarks and IP protection, and drafting liability and stock purchase agreements. In discussing ways to mitigate the high legal and logistical expenditures, we have planned to minimize non-sales travel and supplies expenditures, and to not take salaries for the company's first three years. Founder salaries will increase as sales warrant.

PP&E Expenses

Our PP&E consists only of expenses for office space and office equipment. The cost of plant and equipment is $0 for our company since we plan on outsourcing manufacturing. As we are still a young company and growing, we will keep manufacturing as a variable cost rather than a fixed cost. The two main sources of property expenses is rent for the company headquarters located in Philadelphia and desk spaces for the regional sales representatives. The headquarters will house the executive team, vice presidents of sales, branding, and accounting, and our engineers. We expect that we will need space for 20 employees in our headquarters six months from now and we will need to upsize to space for 40 employees by the end of our third year of operation. Using estimates of 75 sq ft / employee for open desk space, a reception area of 100 sq ft, a conference room for six people that is 200 sq ft, and 200 sq ft of hallways (10% of usable space), we have determined that we will need about 2000 sq ft of office space until we upsize (The Mehigan Company, 2019). With office space rent in Philadelphia at about $2 / sq ft / month, we will be spending $4000 / month in rent for our headquarters office. When we upsize we anticipate that our new office will be about 4000 sq ft, amounting to $8000 / month in rent (Offices.net, 2019). For our regional sales reps, we plan on renting individual desks in coworking spaces. A dedicated desk in a coworking space costs between $400 to $600 / month for depending on location (WeWork, 2019). With five regional reps, one per Power Five conference, the total expense would be around $2500 / month. In terms of office equipment, computers will be the most expensive purchase. We are budgeting $1000 / employee for computer equipment purchases under the assumption that computers last about three years before breaking or becoming outdated. If we purchase Macs, we would consider leveraging the financing options offered by Apple to small-businesses.

Bottom Line: HeartWare Saves Athletes' Lives.

Every day, athletes across the country are at risk of exercise-induced, heart-related emergencies. Fatigue and heat-related illnesses, which can be prevented by monitoring heart rate, are a threat coaches and trainers are becoming increasingly aware of. However, there are also perils that cannot be assessed by heart rate alone. Though arrhythmias and murmurs may not be fatal, they can often be precursors to more dangerous events such as Sudden Cardiac Arrest (SCA) -- the largest 'silent killer' in the United States. To track arrhythmias -- and to create a more complete picture of an individual's heart health -- heart sounds must also be monitored as well. Existing devices that attempt to address this problem, such as traditional heart rate or ECG monitors, are not wearable, real-time, easy to interpret, nor capable of identifying the clues in heart sounds. Because coaches and trainers may be not properly equipped to trained to examine raw heart rate info and detect heart issues, there is a gap in the industry between the customers' desire to protect their athletes and access to the adequate tools to do so.

Enter HeartWare. With five engineers seeking to protect athletes suffering from cardiac-related illnesses, and potentially even save their lives from sudden cardiac arrests, we are developing a wearable heart sound monitor that wirelessly transmits heart audio, analyzes it for signs of cardiac problems, sends alerts, and generates periodic performance insights based on heart audio and heart rate information.We want to create medical-grade device that delivers advanced, critical insights in real time, which we plan to deliver by building relationships with college sports programs.

HeartWare has strong revenue projections (estimated profitability in 2023). To launch and expand our reach, we are seeking a seed funding round of $300,000 in 2020, and a Series A investment of $500,000 in 2021 in return for common equity capital. This funding will allow us to deliver a high-quality, medical-grade device and analytics platform that keeps athletes safe. In 2024, we plan to exit via acquisition by an established medical device firm with a special interest in entering the cardiac monitoring industry. Seed investors can expect an ROI of 50%, while Series A investors can expect an ROI of 35%. We believe HeartWare will save athletes' lives and keep them healthy, and that our founding team is well equipped to deliver a product that solves this problem and addresses the problems that coaches, trainers, and athletes face.

XV. Appendices

Appendix A-Definition of Problem

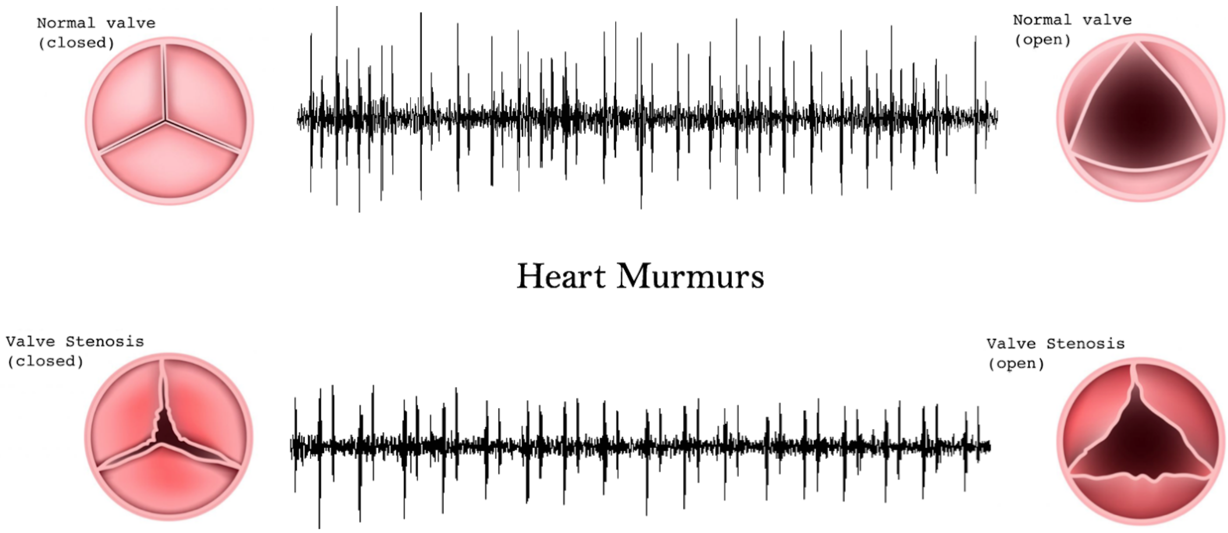

Figure A1: Comparison of Normal and Afflicted Heart PCG

Appendix B-Analog Hardware Design

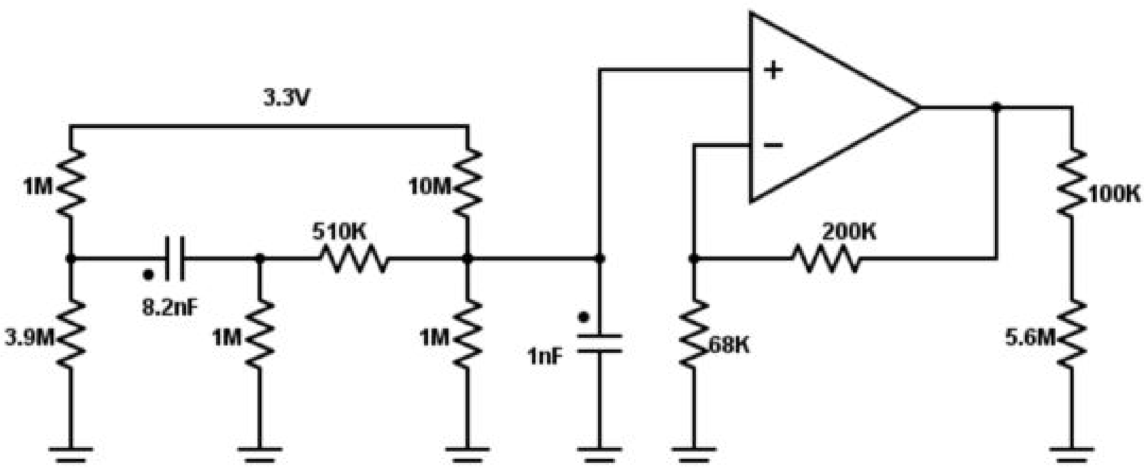

Figure B1: Designed Level Shift, Amplification, and Filter Analog Circuit

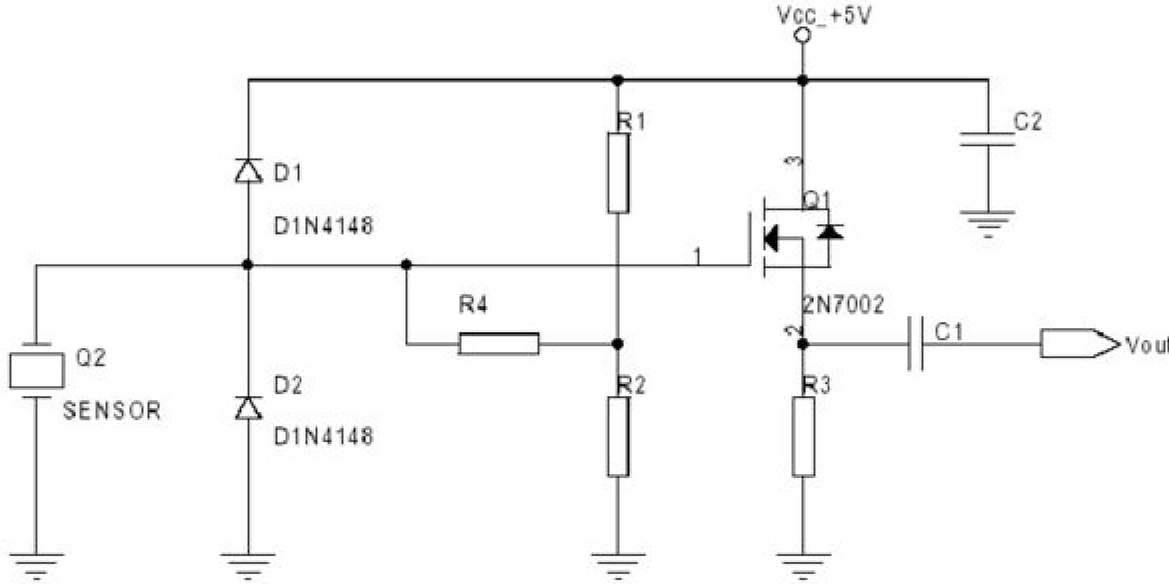

Figure B2: Internal Circuit of Sensor (Source: CM-01B Datasheet)

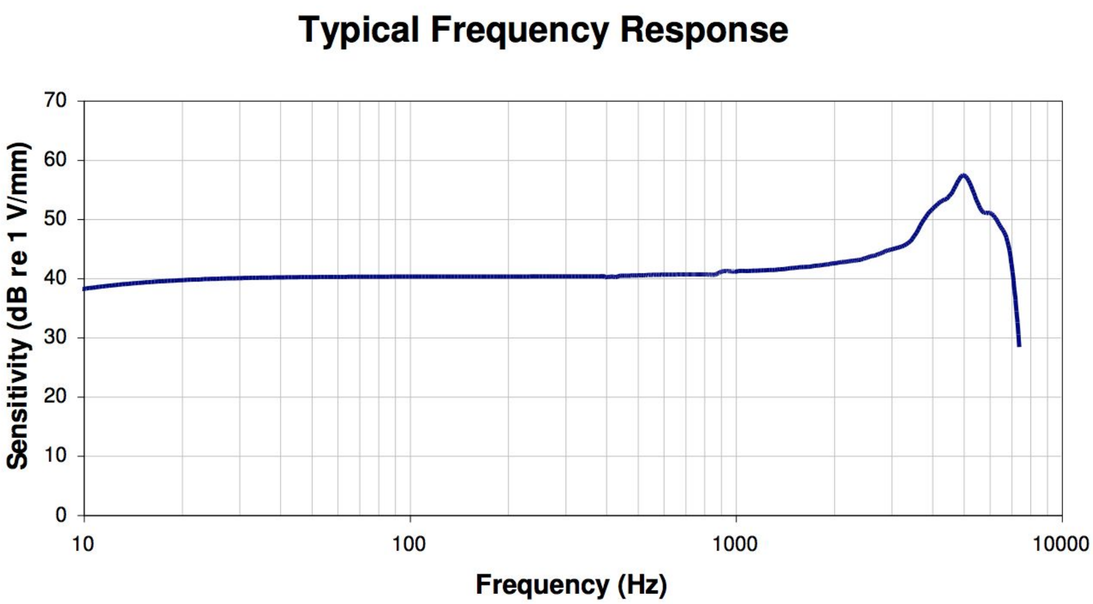

Figure B3: Sensor Frequency Response (Source: CM-01B Datasheet)
Note the excellent low frequency response

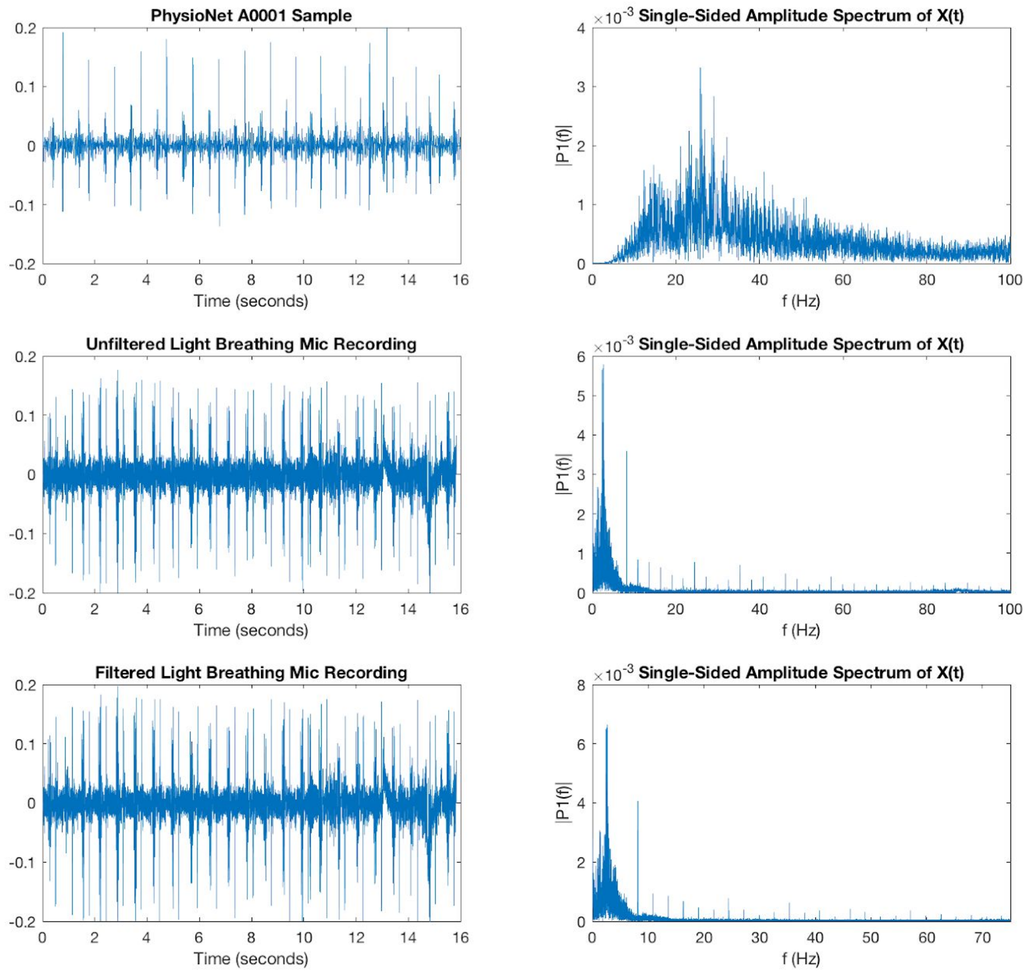

Figure B4: Example of Matlab Dashboard Used for Filter Testing

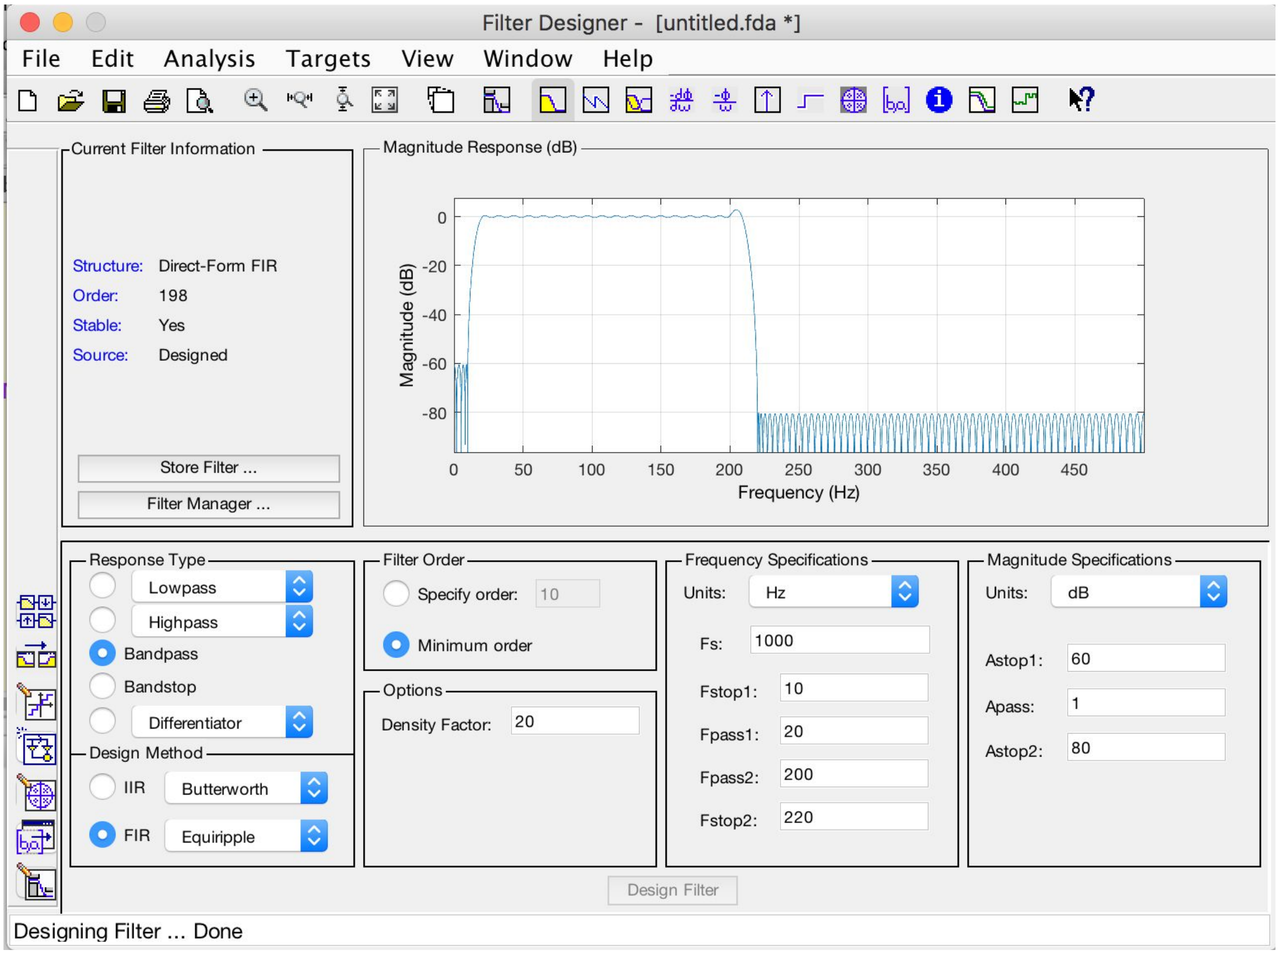

Figure B5: Example of Software Filter Designed Using Matlab's Filter Designer Tool

Appendix C-PCB and Digital Hardware Design

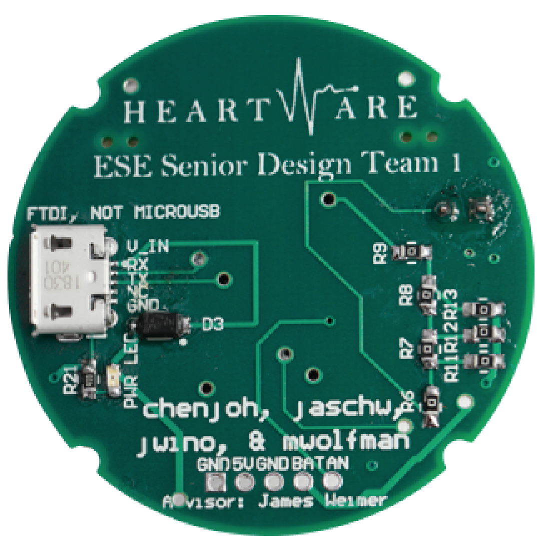

Figure C1: Photo of Manufactured and Assembled PCB

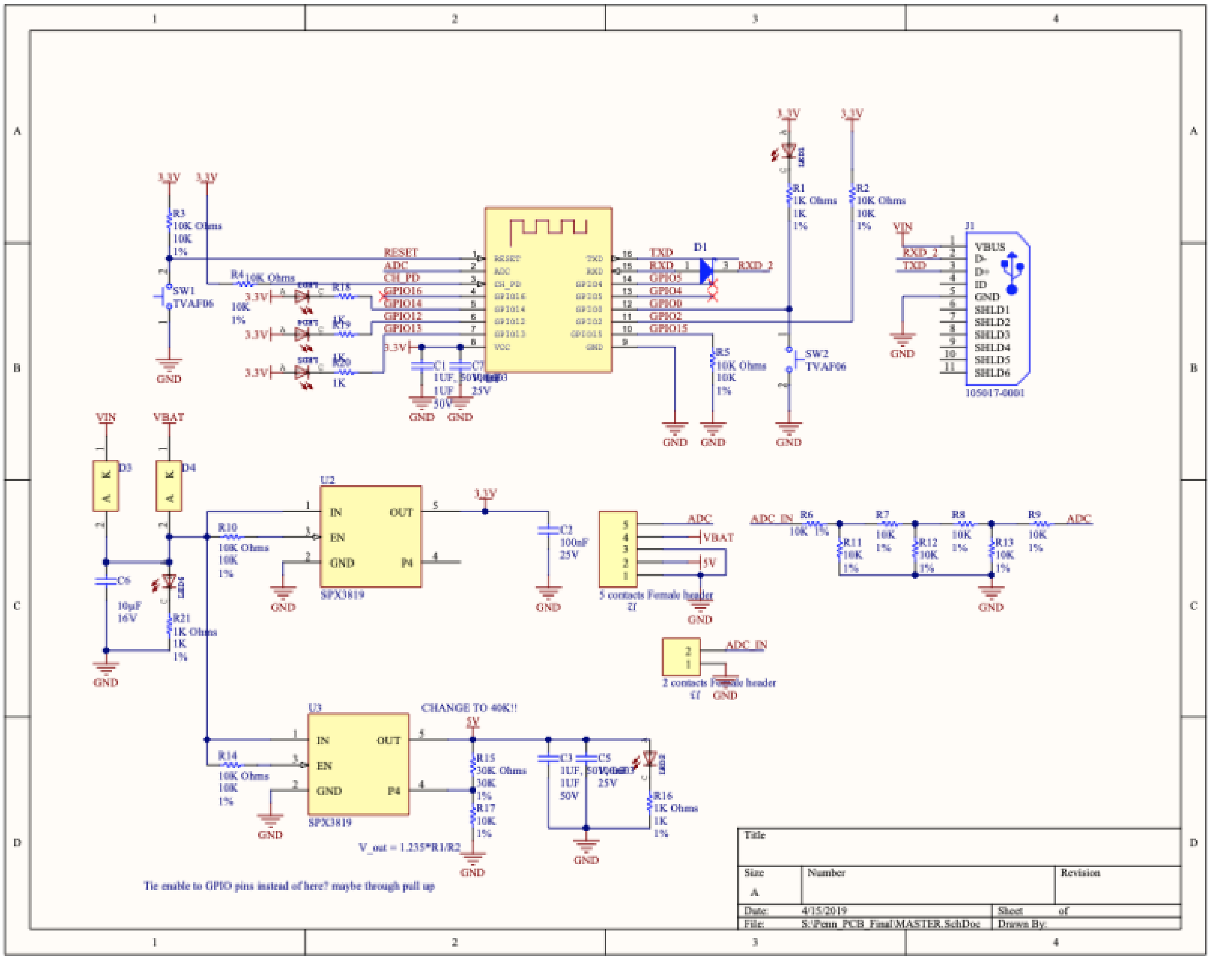

Figure C2: Schematic for PCB

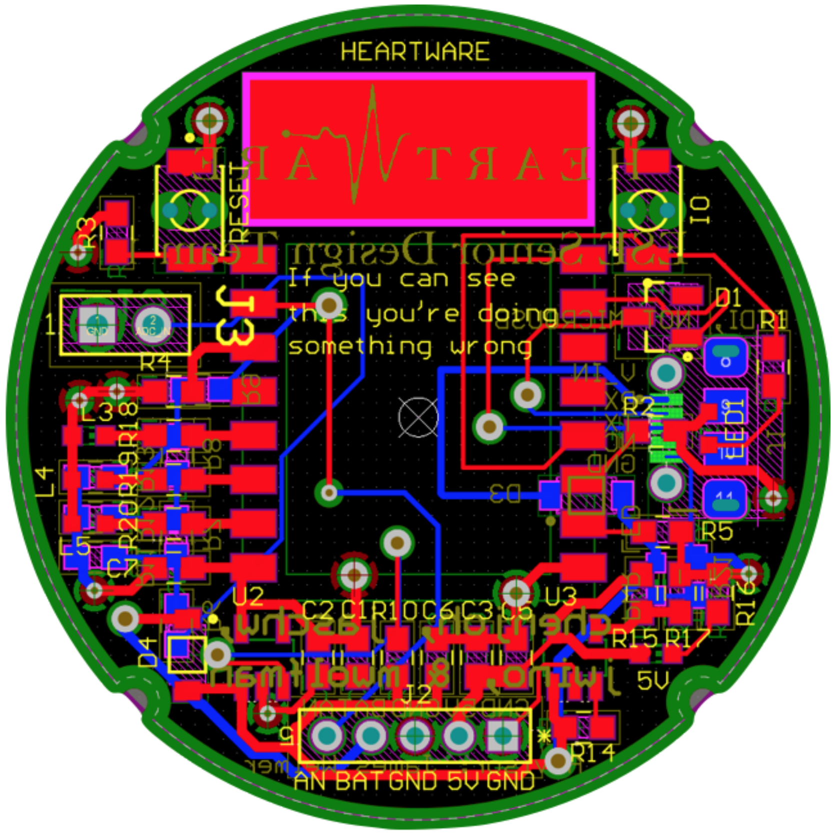

Figure C3: Altium PCB Layout

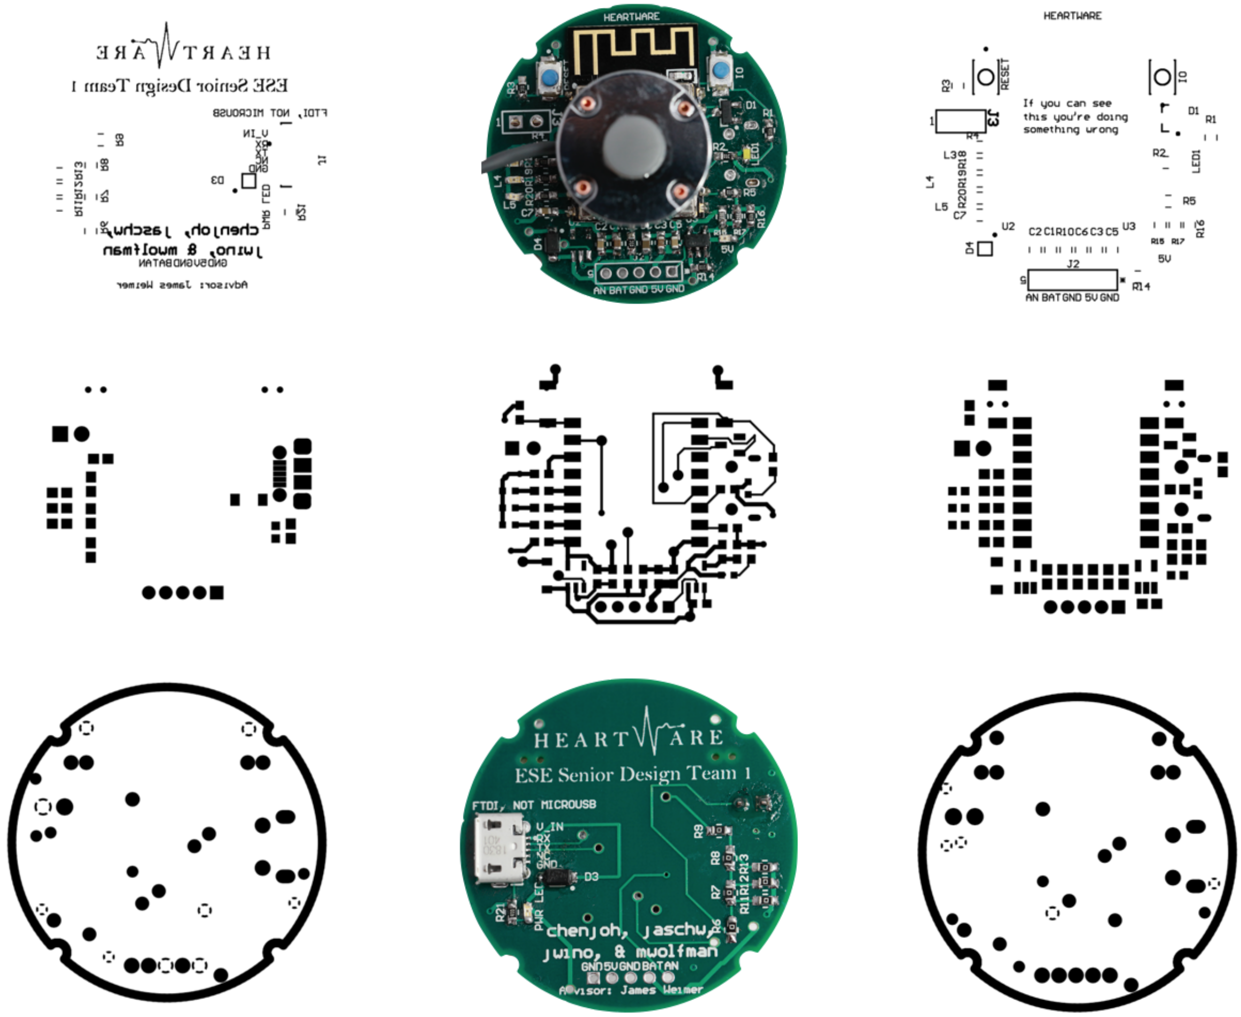

FIgure C4: Gerber Layers of PCB for Manufacturing

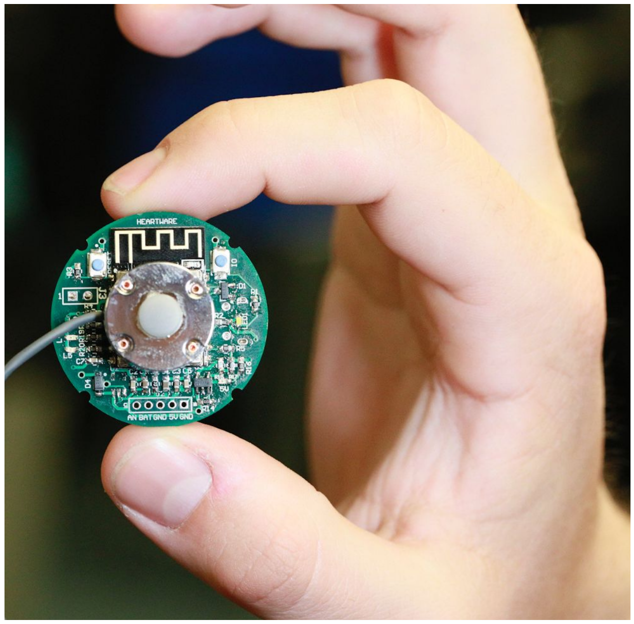

Figure C5: Assembled PCB and Sensor in the Theranos Position

Appendix D-Mechanical Mount and Complete Product

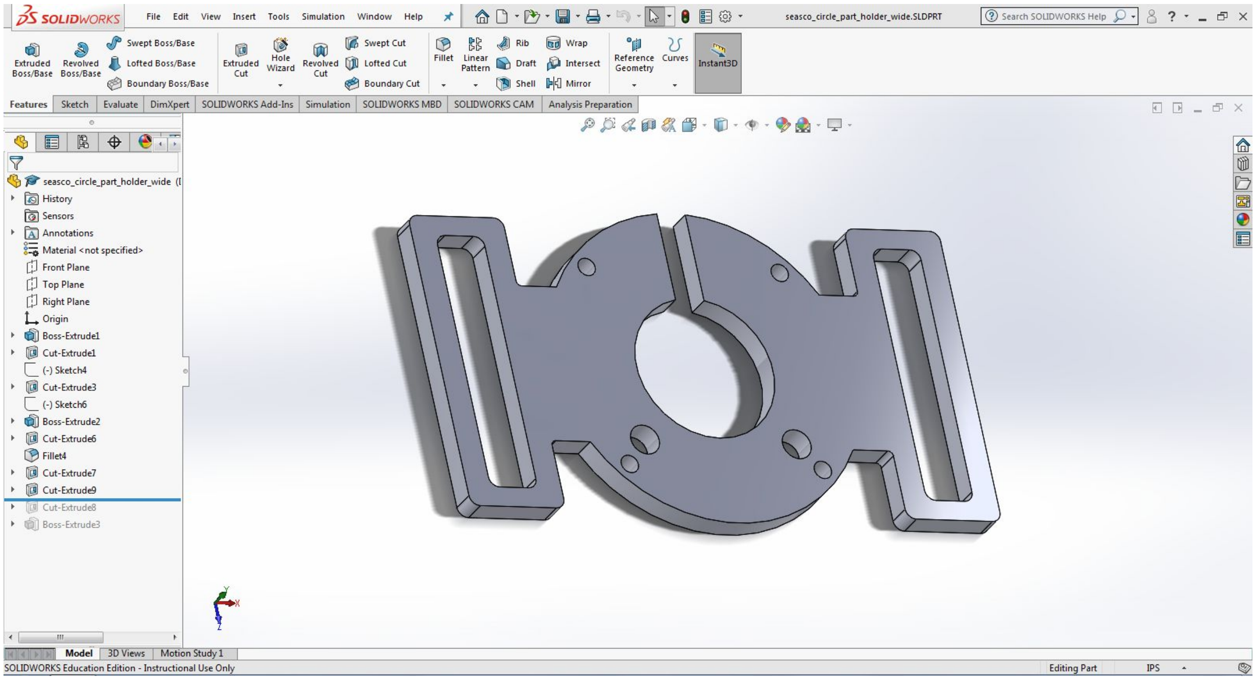

Figure D1: SolidWorks Model of Mechanical Mount

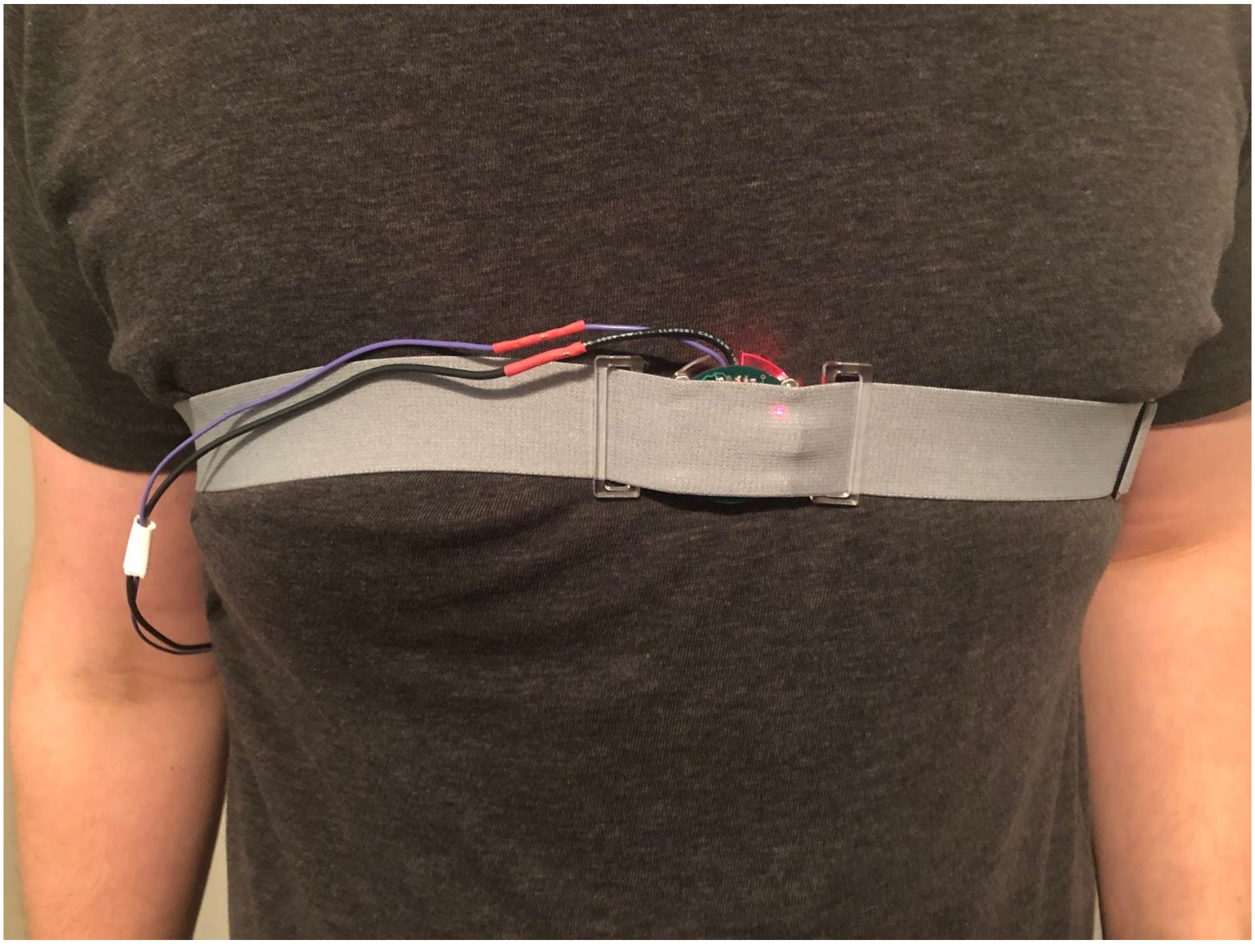

Figure D2: Device Worn on the Chest (Front)

Figure D3: Device Worn on the Chest (Back)

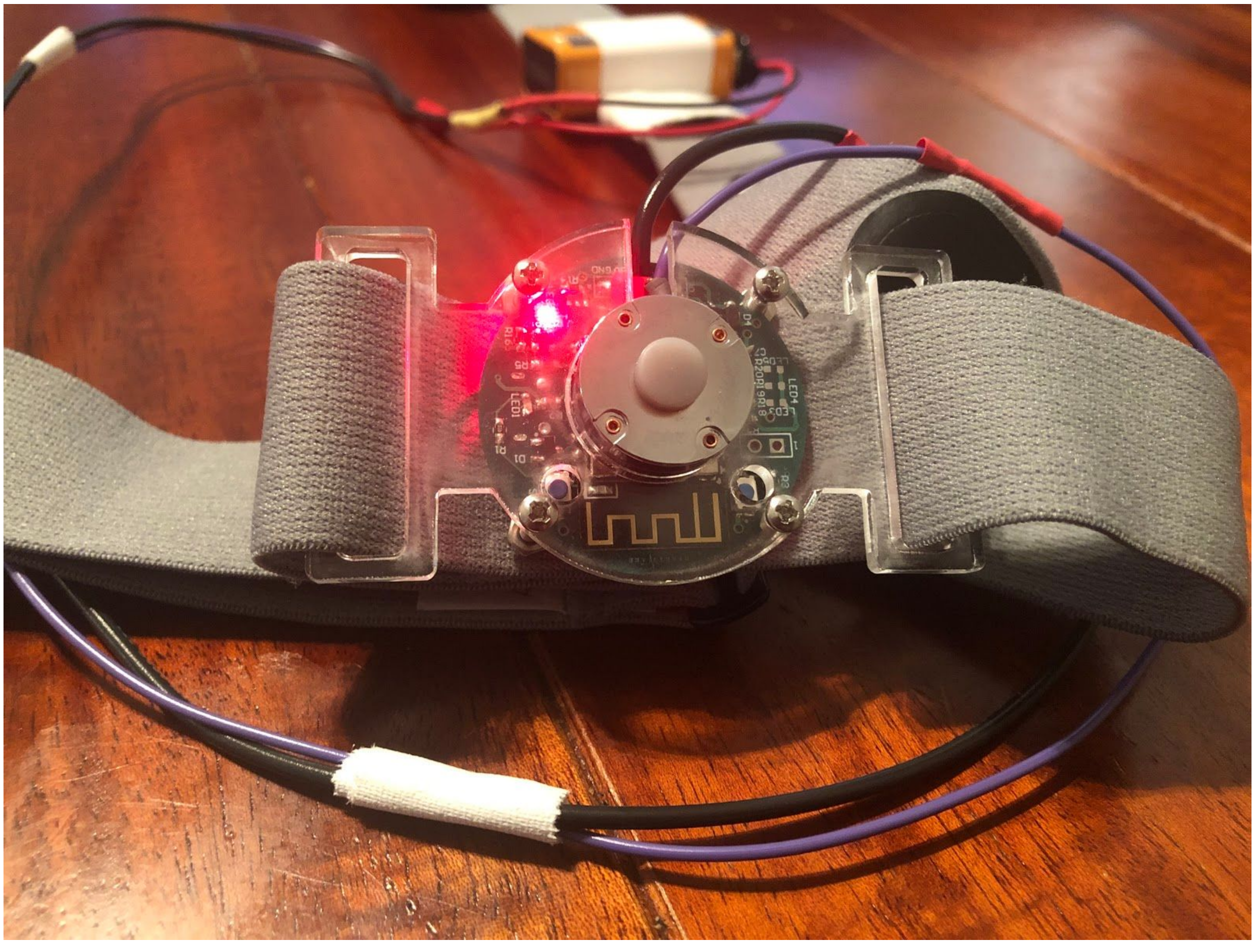

Figure D4: Mounted Device (front)

Figure D5: Back of Mounted Device

Appendix F-Financial Models

Unit Price	**$99**
Sensor	$40
Microcontroller	$7
Other electronics	$5
PCB	$5
Band and housing	$3
Total COGS	$60
Unit sale profit	**$39**

Figure F1: Cost of Goods Breakdown

College atheletic programs	770
Number of students per program	25
Student athletes	**19,250**
Cardiologists	23,000
Number patients per cardiologist	600
% to prescribe	5%
Cardioligist patients	**690,000**
Patients at home at risk	121,500,000
% to prescribe	5%
At home consumers	**6,075,000**
Unit Price	**$99**
Sensor	$40
Microcontroller	$7
Other electronics	$5
PCB	$5
Band and housing	$3
Total COGS	$60
Unit sale profit	**$39**
Yearly subscription price	**$99**

	Year 1	Year 2	Year 3	Year 4	Year 5
	Innovator 1	Innovator 2	Early Adopter 1	Early Adopter 2	Early Majority 1
Student athletes	1%	1.50%	5%	8.50%	15%
New users	192	288	962	1636	2887
Unit sales profit	$7,488	$11,232	$37,518	$63,804	$112,593
Subcriptions	$19,008	$47,520	$142,758	$304,722	$590,535
Segment profit	$26,496	$58,752	$180,276	$368,526	$703,128
		Innovator 1	Innovator 2	Early Adopter 1	Early Adopter 2
Cardiologists	0%	1%	1.50%	5%	8.50%
Users	*No entry*	6900	10350	34500	58650
Unit sales	-	$269,100	$403,650	$1,345,500	$2,287,350
Subcriptions	-	$683,100	$1,707,750	$5,123,250	$10,929,600
Segment profit	0	$952,200	$2,111,400	$6,468,750	$13,216,950
			Innovator 1	Innovator 2	Early Adopter 1
At home consumers	0%	0%	1%	1.50%	5%
Users	*No entry*	*No entry*	60750	91125	303750
Unit sales (half COGS)	-	-	$4,191,750	$6,287,625	$20,958,750
Subcriptions	-	-	$6,014,250	$15,035,625	$45,106,875
Segment profit	0	0	$10,206,000	$21,323,250	$66,065,625
Net income	**$26,496**	**$1,010,952**	**$12,497,676**	**$28,160,526**	**$79,985,703**

Figure F2: Revenue Model

Note: FIgures F1 and F2 were developed for the final round presentations at the M&T Summit

Appendix G- CNN Architecture

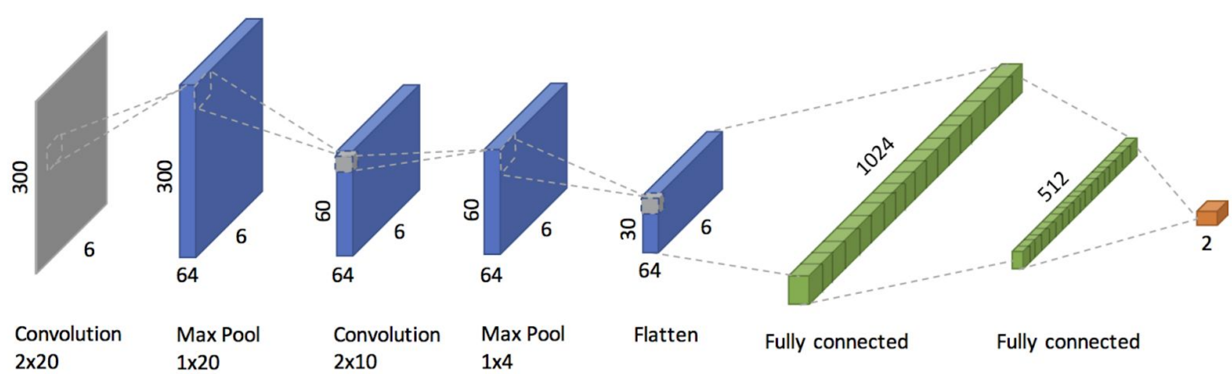

Appendix H - Preprocessing Techniques

Figure H1: Frequency Heatmap of Heartbeat Signal

Appendix I - Cloud Platform

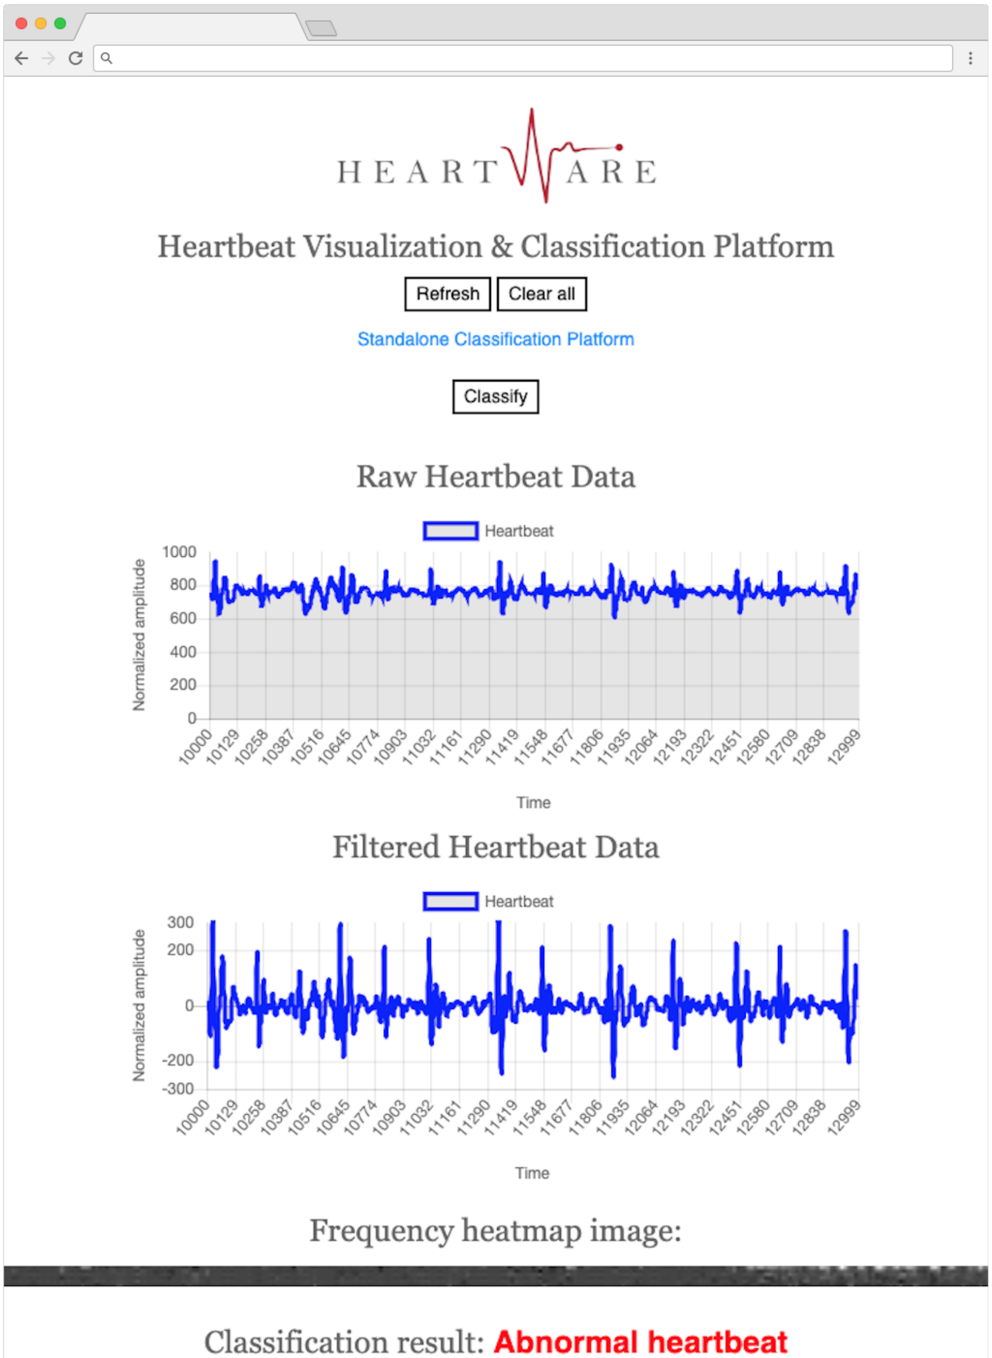

Figure I1: Screenshot of web application