

Detecting Heart Anomalies Using Mobile Phones and Machine Learning

Elhoussine Talab, Omar Mohamed, Labeeba Begum, Fadi Aloul, Assim Sagahyroun
Department of Computer Science & Engineering
American University of Sharjah
United Arab Emirates

Abstract— One out of four deaths is caused by heart related issues. Acting upon early signs of heart disease can, thus, drastically increase probability of saving lives. This paper discusses a cost-effective and reliable method of diagnosing heart abnormalities by using mobile phones that are nowadays typically available to an average user. A mobile application is developed to detect heart abnormal activities using either a digital stethoscope measurement as input, or a mobile recording of the heart beat using the mobile's microphone. To process the raw heart sound data, we first denoise the signal using wavelet transforms, and then apply machine learning techniques, namely, Convolutional Neural Networks for the classification of the stored heart sounds. A database consisting of recorded human heart sounds and their corresponding diagnosis is used to train the neural network. Moreover, neural network fine-tuning techniques such as ADAM Regularization is used to smoothen the prediction process. The proposed approach is tested on heart sound signals, that are 5 to 8 seconds long, and is shown to perform with an accuracy of 94.2% on the validation set.

Keywords—*Mobile Phones, Machine Learning, Convolutional Neural Networks, Digital Stethoscope, Cardiovascular Diseases*

I. INTRODUCTION

Cardiovascular disease (CVD) includes any disease that affects the heart or blood vessel's ability to function normally. There are various types of CVDs such as arrhythmias, prolapsed mitral valve, coronary artery disease, congenital heart disease and others. CVDs are one of the leading causes of morbidity and mortality globally with an estimated 17.9 million people deceased due to CVDs in 2016, representing 31% of all worldwide deaths. Lifestyle risk factors such as smoking, unhealthy diet, and physical inactivity account for ~80% of coronary artery and cerebrovascular disease [1]. Low-economic status is an independent risk factor and over three quarters of CVD deaths are found to occur in Low and Middle Income Countries (LMICs) [2]. Early diagnosis of CVD is crucial in order to decrease the risk factors resulting in deaths. However, people in LMICs find it difficult to obtain high quality CVD diagnosis as they have less access to effective and equitable health care services which meet their needs [2]. Heart sound auscultation, i.e. the act of listening to heart sounds through a stethoscope, is one of the basic techniques used for the diagnosis of CVDs. However, interpretation of heart sounds is not only highly subjective, but also a difficult skill that requires a professional to identify the sound. The medical expert uses a stethoscope for an initial screening and listens for irregularity of heart sounds from the patient's chest. However, human ears may tend to miss certain sounds as the higher frequency sounds could mask

the lower frequency sounds [3]. When auscultation was performed by medical students and primary care physicians, only a 20- 40% accuracy was achieved, and 80% accuracy by experienced cardiologists [4]. Driven by the above mentioned facts, a mobile application that listens to heart sound, cleans the signal and predicts whether the user requires medical attention or not, due to heart abnormalities, is proposed in this paper.

Mobile phones have shown the potential of being a powerful medical diagnostic tool and has been used in various medical applications such as the detection of obstructive sleep apnea [5], chronic obstructive pulmonary disease (COPD) [6], and Parkinson's disease [7]. In this paper, a mobile application is developed to detect heart abnormalities.

The following section gives an overview of related literature works. Section III discusses the proposed solution. Section IV provides testing and evaluation of the solution. The paper concludes in Section V.

II. LITERATURE REVIEW

In recent years, due to the crucial role of remote monitoring for Cardiovascular disease (CVD) patients, there has been significant efforts from research communities towards proposing and designing various CVD monitoring devices and mobile health applications. In Kown et al.'s [8] research, scholars utilized an iPhone's 3-axis accelerometer, which is sufficient and sensitive enough to perceive body movements caused by heart pumping, to measure the heart rate. The paper explains that the iPhone accelerometer readings were taken by mounting the phone on the middle of the chest in six different inactive or less-active postures. The six positions are easily posed in everyday life including sitting, standing, sleeping (supine, lateral, prone) and slow walking (1Km/h, 3Km/h) on a treadmill. Figure 1 shows the used postures. The paper discusses that the raw accelerometer data was filtered by a band pass algorithm and peak detection algorithm which is designed for real-time processing using a mobile phone. The authors specify that the accelerometer data was filtered with 5th order Butterworth high pass filter at cutoff frequency of 5Hz and then filtered with 5th order Butterworth low pass filter at the cut off frequency of 35Hz. The paper demonstrates that the iPhone can effectively function as a reliable heart rate extractor. However, the proposed approach doesn't diagnose the extracted heart signal.




Figure 1: Location of the mobile phone during each part in experimental protocol used in [8].

Landreani et al. [9] also presents a study that extracted heart rate from 9 healthy volunteers in different positions namely in supine and standing postures using mobile phones. While standing, the volunteers placed the mobile phone on their thorax, as shown in Figure 2 (POS1), and while supine, the mobile phone was placed on the thorax (POS1) and on the navel (POS2) to acquire seismocardiographic (SCG) signals. The signal was recorded for 3 minutes, the peak was detected using an algorithm, and compared to the RR intervals, which are the time intervals between successive heartbeats, from the ECG for evaluation. Again, this paper doesn't diagnose the extracted heart signal.




Figure 2: Schematization of the two positions and orientation of the mobile devices used during acquisition [9].

Low and Choo's [10] research explore using deep learning neural networks for heart sound classification. In their paper, the heart sounds are segmented per heartbeat and each segment is converted to form an intensity map that is classified using Softmax Regression (SMR) and a Convolutional Neural Network (CNN). The paper emphasizes that the approach followed removes the need to provide features to a supervised machine learning method and instead the features can be automatically determined through training a deep neural network model. The paper concludes that the CNN was able to provide better classification of heart sounds than SMR. It also highlights the simplicity in implementation and that this approach can be used for real-time classification of heart sounds from mobile phones. The paper doesn't use mobile phones for capturing the heart signal.

The dependence of accurate heart sound analysis on the quality of heart sound recordings is presented in a paper by Springer et. al [11]. The paper uses a Support Vector Machine (SVM) classifier in order to compute signal quality metrics from recordings. The paper found that one-third of all mobile phone recorded phonocardiograms (PCGs) were of high quality. The classifier used was able to determine the quality of PCGs for a range of stethoscopes accurately. This paper also concludes that a mobile phone is capable of recording high-quality PCG signals. From this paper, we conclude that mobile phone recordings can be classified with a high degree of accuracy using the SVM technique presented. Our proposed approach replaces SVMs with CNNs which are faster and doesn't require feature extraction.

Prasadh et al's paper [12] discusses the importance of having an efficient noise reduction algorithm for audio signals, which is required for the proposed mobile application [12]. The researchers state that using certain filters tweaked to specific parameters could help in decreasing the cost for noise removal, while still eliminating a large portion of the noise in audio signals. Their technique is different from the usual linear filtering methods. The algorithm adapts to each and every audio file and cancels out noise differently depending on the characteristics of the audio file and noise signal. In order to remove the white Gaussian noise which is the dominant noise in any audio signal they discuss a technique called discrete wavelet transform (DWT). A series of other filters are also used to achieve an even cleaner signal. Filters such as Kalman filters are used as they have very less response time which makes them extremely efficient. DWT was used in our proposed solution to initially clean the heart signal captured by the mobile phone.

III. PROPOSED SOLUTION

In this paper, we propose a mobile application that first records the heart sound using the internal built-in mobile phone's microphone, then cleans and processes the signal using signal processing techniques and eventually predicts whether the signal indicates a heart that functions normally, or not through the use of Convolutional Neural Networks (CNN)-based techniques.

The proposed solution is divided into the following three main stages, namely, a signal processing stage, followed by a machine learning stage, and finally a mobile application. These are described in the following sections.

A. Signal Processing

The user initially starts the mobile application which provides the user with two options to record their heart signal. The first option asks the user to place the mobile phone's microphone on the mitral valve location (under the 5th rib), as shown in Figure 3. A simple tutorial is shown on the mobile application to guide the user on how to place the microphone. The second option uses an external digital stethoscope to capture the heart signal. The external stethoscope is connected to the mobile phone through a headphone jack.




Figure 3: location of the microphone to properly acquire the heart signal.

After capturing the signal, the mobile application initiates the *signal processing stage* to clean the heart signal audio file to create meaningful phonocardiograms (PCG) readings which is then passed to the *machine learning stage*. This stage employs wavelet transforms in order to denoise the heart sound recordings.

In the process of gathering recordings of the heart using a mobile phone microphone we faced a challenge of denoising the signal coming in from the internal phone's microphone. Noise is defined as an unwanted signal that interferes with the communication or measurement of another signal. A noise itself is an information-bearing signal that conveys information regarding the sources of the noise and the environment in which it propagates.

Various types of noise exist; however, in our recordings, the main noise types were *white* noise which is a purely random noise that has an impulse autocorrelation function and a flat power spectrum. White noise theoretically contains all frequencies in equal power. We also had some *impulsive* noises which consist of short-duration pulses of random amplitude, time of occurrence and duration. Occasionally, we were also faced with *band-limited white* noise, which is very similar to white noise, however this noise comes with a flat power spectrum and a limited bandwidth that usually covers the limited spectrum of the device or the signal of interest. The autocorrelation of this noise is sine-shaped.

Our proposal approach passes the signal through a series of wavelet transforms to capture a cleaner heart signal. In our case, the low-frequency content is the most important part because it represents the desired signal and its identity; which is very important for our model's prediction. On the other hand, the high-frequency content imparts nuance. First, the signal passes through a one stage filter called *approximation and details* where we capture the high-scale low frequency components in the signal. We found that eliminating low-scale high frequency components does not distort the signal. Finally, *Coiflet 5* wavelet transform is utilized to denoise the signal. Figure 4 below illustrate the filter outputs.




Figure 4: The 1st filter stage [red: original feed. purple: processed feed]

B. Machine Learning

The Machine learning stage receives the processed data from the Signal Processing stage in the form of an audio file storing the heart signal of the patient. The system is trained using an on online open source dataset [13], which includes 176 data sets of healthy and unhealthy heart sounds that are acquired using a sample from the general public via the iStethoscope Pro iPhone application. The dataset is labelled into four categories: (1) normal, (2) murmur, (3) extra heart sound and (4) artifact. The *normal* category represents normal, healthy heart sounds. The *murmur* category represents heart murmurs which can be a symptom of many heart disorders. The *extra heart sound* category is to identify extra heart sounds which may not be a sign of an existing disease. For convenience, on suggestions from data scientists in [13], we have categorized normal and extra heart sounds as a single prediction to the user. The *artifact* category represents a wide range of different sounds such as feedback squeals, echoes, speech, music, and noise. Since the artifact category is considered as noise, we can instruct the user to try recording their heart sound again in case their prediction results indicate an artifact.

To design the Neural Network for predicting heart sound anomalies, we first explored existing architectures which was used in audio applications and improved on the architecture for our specific case. We decided to use the framework "KERAS" as it provides easier high-level API to Google's Tensorflow framework [14]. Before discussing the Convolutional Neural Network (CNN) architecture, we discuss the data pre-processing and features used for training.

Each example in the dataset is just an audio wav file. The application analyzes the file in terms of the value of the gain (amplitude) at a particular sampling rate. The sample rate we used was 44,100 Hz. Any file that does not have a sampling rate of 44,100 Hz gets converted to this sampling rate. Moreover, all audio files are forced to be of 9 seconds long. Shorter audio wav files are repeated until it is 9 seconds long. Longer audio wav files are cut to 9 seconds. This step is crucial, because it determines the features to be fed to the first layer of the CNN. The formula of equation 1 below illustrates the initial number of features given the sampling rate and the audio file length:

$$\text{sample rate (44,100 Hz)} \times \text{max audio length (9 seconds)} = \text{initial number of features (396900 feature)}. \quad (1)$$

Each feature to be fed to the neural network is just the amplitude value at a certain time in a particular audio wav file. However, we observed that a high number of the 396,900 features are irrelevant, which caused the model to have unreliable performance. Therefore, the audio file was down sampled [15] to reduce the number of features while keeping the overall shape of the audio file intact.

We down sampled three times with factors 8, 8 and 4 respectively based on trial and error, we ended up with the number of features per each example entering the first layer of the model given in equation 2.

$$\frac{396900 \text{ features}}{8 \times 8 \times 4} \approx 1551 \text{ (features)} \quad (2)$$

For the activation function, we used RELU function (shown in Figure 5) because it does not saturate at high values of x compared to other alternatives such as sigmoid function. Moreover, another advantage of RELU is sparsity, meaning that it provides output of 0 if x is less than 0. This will enable us to have less dense representations for false predictions. All together, these advantages add up to allow for faster learning and faster prediction.




Figure 5: RELU Function.

C. Mobile Application

The above stages were bundled in a hybrid user-friendly mobile application available in iOS and Android platforms using Ionic Framework. The users will be able to login, record their heart signal and get a prediction whether their heart signal is normal or in a murmur state. The latter indicates the necessity to seek medical attention from a specialized physician. Figure 6 shows screenshots from the developed mobile application.

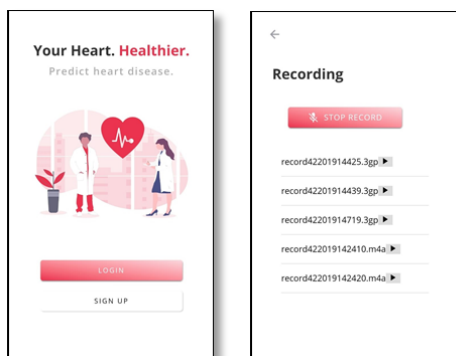


Figure 6: Sample mobile application screenshots.

The proposed system components, shown in Figure 7, consist of the patient's mobile phone application and cloud services. The mobile application fetches the signal and sends it through *RESTful API* to the cloud services. The cloud services apply all the processing required to the signal and replied to the *API* call with the prediction which is presented to the user via the user friendly mobile application.

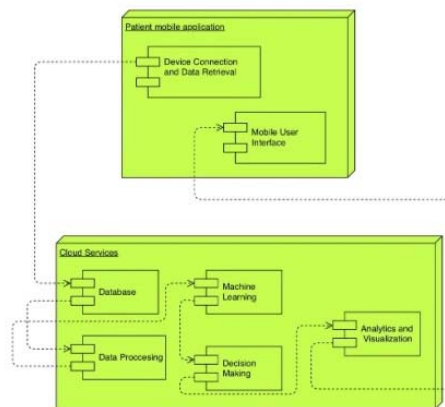


Figure 7: Components architecture of the proposed solution.

IV. TESTING & VALIDATION

When using the online open source dataset, which included 176 examples of healthy and unhealthy heart sounds, 30% of the dataset was reserved for validation. The remaining 70% was used for training the model. Using our developed model, the validation set scored an accuracy of 94.2%.

To further confirm the validity of the developed model, another dataset was used from MIT [16], known as "PhysioNet/CinC 2016", that consisted of audio heart signals labeled as normal or abnormal. 200 examples were used in testing the model which scored an accuracy of 90.15%.

Furthermore, we developed a *confusion matrix* to measure how well the model predicts each class. In our case, the possible predicted classes were:

1. Artifact: Non-heart sound
2. Normal/Extras: Normal heart sound
3. Murmur: Abnormal heart sound

To benchmark our model on a confusion matrix, we used the infamous *sklearn* library [17]; specifically, the *sklearn.metrics.confusion_matrix* where we passed the test set and the prediction of the test set after it was passed through the model. The *sklearn* function automatically determines the classes and constructs a confusion matrix. Table I shows the results scored in terms of confusion matrix.

		Prediction		
		Artifact	Normal/Extras	Murmur
Ground Truth	Artifact	10	0	0
	Normal/Extras	1	17	1
	Murmur	0	3	12

Table I: Confusion Matrix.

As the table illustrates, we observe the *true positive* of artifact class in our model as a score of 10, for Normal/Extras it scored 17 and it scored 12 for Murmur class. On the other hand, the false positive for Artifact was zero, for Normal/Extras it was 2 and for Murmur it was 3. Hence, we conclude that our model, when tested on test/validation set, it yielded a 100% Artifact prediction, 89% Normal/Extras prediction and 80% Murmur prediction accuracy.

V. CONCLUSIONS

In this paper, we have presented a detailed description of a mobile phone based stethoscope that can detect heart abnormalities. The solution is portable, low-cost and does not require a trained physician. The system utilizes mobile phone's internal microphone to extract raw heart data in an audio format. The data is then transmitted to the cloud where it is cleaned and processed using several wavelet transforms filters and then passed to Convolutional Neural Network (CNN) model to predict whether the user requires medical attention of not. This system will be especially useful for people in Low and Middle Income Countries (LMICs) who do not have access to basic healthcare services or expert cardiologists. Therefore, in non-hospital environments and rural areas, our system can be used as an initial tool for early diagnosis of CVDs.

REFERENCES

[1] L. Neubeck, N. Lowres, E. Benjamin, S. Freedman, G. Coorey, and J. Redfern, "The mobile revolution--using smartphone apps to prevent cardiovascular disease," in *Nature Reviews. Cardiology*; London, vol. 12, no. 6, pp. 350–360, June 2015.

[2] "Cardiovascular diseases (CVDs)." [Online]. Available: [http://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-\(cvds\)](http://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds)).

[3] S. Thiyagaraja, R. Dantu, P. Shrestha, A. Chitnis, M. Thompson, P. Anumandla, T. Sarma, and S. Dantu, "A novel heart-mobile interface for detection and classification of heart sounds," in *Biomedical Signal Processing and Control*, vol. 45, pp. 313–324, August 2018.

[4] S. Latif, M. Usman, R. Rana, and J. Qadir, "Phonocardiographic Sensing using Deep Learning for Abnormal Heartbeat Detection," in *IEEE Sensors Journal*, vol. 18, issue 22, pp. 9393-9400, September 2018.

[5] M. Al-Mardini, F. Aloul, A. Sagahyoon, and L. Al-Husseini, "Classifying Obstructive Sleep Apnea Using Smartphones," in *Elsevier Journal of Biomedical Informatics*, vol. 52, pp. 251–259, December 2014.

[6] F. Zubaydi, A. Sagahyoon, F. Aloul, and H. Mir, "MobSpiro: Mobile Based Spirometry for Detecting COPD," in *Proceedings of the 7th IEEE Annual Computing and Communication Workshop and Conference (CCWC)*, Las Vegas, USA, January 2017.

[7] A. Sahyoun, K. Chehab, O. Al-Madani, F. Aloul, and A. Sagahyoon, "ParkNosis: Diagnosing Parkinson's Disease Using Mobile Phones," in *Proceedings of the 18th IEEE International Conference on E-Health Networking, Application and Services (HealthCom)*, Munich, Germany, September 2016.

[8] S. Kwon, J. Lee, G. Chung, and K. Park, "Validation of heart rate extraction through an iPhone accelerometer," in *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 5260–5263, 2011.

[9] F. Landreani, A. Martin-Yebra, C. Casellato, C. Frigo, E. Pavan, P-F Migeotte, and E. Caiani, "Beat-to-beat heart rate detection by smartphone's accelerometers: Validation with ECG," in *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 525–528, 2016.

[10] J. Low and K. Choo, "Classification of Heart Sounds Using Softmax Regression and Convolutional Neural Network," in *Proceedings of the International Conference on Communication Engineering and Technology*, pp. 18–21, February 2018.

[11] D. Springer, T. Brennan, L. Zühlke, H. Abdelrahman, N. Ntusi, G. Clifford, B. Mayosi, and L. Tarassenko., "Signal quality classification of mobile phone-recorded phonocardiogram signals," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1335–1339, May 2014.

[12] S. Prasad, S. Natrajan, and S. Kalavani, "Efficiency analysis of noise reduction algorithms: Analysis of the best algorithm of noise reduction from a set of algorithms," in *Proceedings of the International Conference on Inventive Computing and Informatics*, pp. 1137–1140, November 2017.

[13] "Classifying Heart Sounds Challenge." [Online]. Available: <http://www.peterjbentley.com/heartchallenge/index.html>

[14] "Keras: The Python Deep Learning library." [Online]. Available: <https://keras.io/>

[15] "Python-based ecosystem of open-source software for mathematics, science, and engineering." [Online]. Available: <https://www.scipy.org/>

[16] "PhysioNet/CinC Challenge 2016: Training Sets." [Online]. Available: <https://www.physionet.org/physiobank/database/challenge/2016/>

[17] "Scikit-learn: Simple and efficient tools for data mining and data analysis" [Online]. Available: <https://scikit-learn.org/>