Proceedings of the 2007 IEEE Symposium on
Computational Intelligence and Data Mining (CIDM 2007)

Detection and Classification of Cardiac Murmurs using Segmentation Techniques and Artificial Neural Networks

S. L. Strunic, F. Rios-Gutierrez, R. Alba-Flores, G. Nordehn, S. Burns
University of Minnesota Duluth
1023 University Drive
Duluth, MN 55812
stru0121@d.umn.edu

Abstract- In this paper we present the implementation of a diagnostic system based on Artificial Neural Networks (ANN) that can be used in the detection and classification of heart murmurs. Segmentation and alignment algorithms serve as important pre-processing steps before heart sounds are applied to the ANN structure. The system enables users to create a classifier that can be trained to detect virtually any desired target set of heart sounds. The output of the system is the classification of the sound as either normal or a type of heart murmur. The ultimate goal of this research is to implement a heart sounds diagnostic system that can be used to help physicians in the auscultation of patients and to reduce the number of unnecessary echocardiograms-- those that are ordered for healthy patients. Testing has been conducted using both simulated and recorded patient heart sounds as input. Three sets of results for the tested system are included herein, corresponding to three different target sets of simulated heart sounds. The system is able to classify with up to 85$\pm$7.4% accuracy and 95$\pm$6.8% sensitivity. For each target set, the accuracy rate of the ANN system is compared to the accuracy rate of a group of 2nd year medical students who were asked to classify heart sounds from the same group of heart sounds classified by the ANN system. System test results are also explored using recorded patient heart sounds.

I. Introduction

Cost effective, accurate and early detection of cardiac illnesses is important to curb deaths caused by cardiovascular diseases. However, a substantial amount of health care money each year is spent on healthy people when it should be put toward assisting individuals who are ill [1]. One major cause of this misallocation of funds is the unnecessary request for echocardiograms.

When a patient visits the physician for auscultation, a heart murmur is the most common abnormal auscultatory finding. When a murmur is detected, the physician must decide whether it represents either a pathological or an innocent murmur [2]. An innocent heart murmur is one which does not represent current or future illness. Oftentimes, a physician who suspects that a patient is healthy will still order an echocardiogram for reassurance, even though the cost of an echocardiogram is high [1]. The result of this practice is a misallocation of healthcare funds, since echocardiograms cost between $750-1500 [3]. While it is clearly important to avoid type-I errors where healthy patients are sent for echocardiogram, it is also important to avoid type-II errors—the situation where a patient who has a pathological heart murmur is sent home without proper treatment.

The ability of primary care physicians to accurately diagnose a murmur is poor; as such, it is no wonder that studies show significant numbers of people who have innocent murmurs still have echocardiograms ordered [4]. This may occur if physicians are settling for a high rate of type-I errors because they assume their auscultation ability is not reliable enough to rule out type-II errors.

In this work, we present the implementation of a diagnostic system that can help to reduce the number of echocardiograms that are ordered for healthy patients. The diagnostic system is based on an easy-to-use graphical user interface that has been designed using MATLAB software and the ANN Toolbox. The system allows the user to interactively design and create a heart sound classification system that implements a set of ANNs as a means for classification. After a classifier has been designed and created, the user is able to select any audio files from the heart sound library on his or her computer and use them as input to the ANN system.

The ultimate goal of the diagnostic system is to provide physicians with an inexpensive classification tool to use along with auscultation. The classifier may provide helpful guidance in the event that a patient has a heart sound that is somewhat difficult for physicians to diagnose. We hope that in the near future, this type of guidance would reduce the tendency of the physician to refer the patient to echocardiogram when the physician is only somewhat uncertain of the diagnosis of the heart sound.

To show the validity of the system in the classification of heart sounds, testing has been performed using the ANN classifier to distinguish between different types of heart sounds in three different training target sets:

1. normal heart sounds, aortic regurgitation (a type of diastolic murmur denoted as AR), and aortic stenosis (a type of systolic murmur denoted as AS).
2. normal heart sounds, AR, AS, and mitral regurgitation (a type of systolic murmur, denoted as MR).
3. innocent (low-grade) AS and pathological (severe) AS.

Given an input heart sound of virtually any length or heart rate, a segmentation algorithm is applied to compute the representative single heart cycle for that sound.

The segmentation algorithm is used to identify the heart sound components S1 and S2. Once the positions of these components are located, individual heart cycles can be identified and the average of all the cycles within the sound is computed. The heart cycle obtained from the average will be the one that represents the particular sound under study. Then, the spectrogram of that heart cycle is computed. Using the spectrogram, the 195Hz frequency component (time vs. amplitude vector) of each sound in the design (or test) set are extracted and aligned. Through experimentation it was found that the 195Hz band of the heart sound contains the necessary information for the ANN to identify the type of murmurs we are considering in our study. The alignment algorithm then prepares the vectors that will be inputted into the ANN system for the design (or test) phase.

To design and test the ANN system, a heart sound library consisting of 96 simulated heart sounds is used. The simulated heart sounds were recorded from a device known as a Phono-Cardio-Simulator, which has been used by medical schools to teach auscultation to students. Also, recorded patient heart sounds are available for testing from the Murmur Study Library at the University of Minnesota Duluth. The sounds from the Murmur Study Library were recorded (with IRB approval) using a commercially available electronic stethoscope provided by Welch-Allyn. Echocardiogram reports from St. Luke's Hospital in Duluth, MN, have been used to define the correct condition of each recorded patient heart sound. Currently, the patient sounds used for testing in this research include 7 examples of normal, 4 examples of AS, 2 examples of MR and 2 examples of AR.

As a means for comparison of ANN system results to a real-world scenario, a group of medical students at the University of Minnesota Medical School Duluth were asked to classify the same set of simulated heart sounds that was used to design the system described in this paper. The classification accuracy for the students is provided and compared to ANN system performance.

II. Prior Work

Cardiology is a popular field for the different types of ANN applications [5-10, 19-21]. Many studies have worked toward designing practical murmur detector and classifier systems to improve the diagnostic accuracy of physicians in small practice settings. A report from 2004 described a murmur classification system that achieved accuracy over 73%, but the system was trained and tested using a small set of data—6 total sounds for training and an additional 38 sounds in the test set [11]. Furthermore, many of those sounds are taken from internet sources rather than more reliable recorded data. Other studies describe research using larger sets of data that were not collected from internet sources. As an example, one study used an ANN structure to classify pediatric heart sounds as either innocent or pathological [10]. Sounds were recorded at a hospital using a microphone, and the frequency spectrum was used as input to an ANN structure. The classifier achieved high accuracy, but the system required its users to select input vectors by hand based on the best observed cycles in terms of background noise and heart sound clarity. The need for a person to assess which signals are best to use for the ANN system makes the system non-ideal for a real-world scenario—a problem that is common to several reported classifier systems [5-6, 11, 21].

Various segmentation algorithms are used in an attempt to eliminate the need for manual selection of heart cycles from heart sounds [6, 8-10]. Segmentation algorithms extract individual heart cycles from heart sound recordings based on properties common to all cardiac signals. Still though, problems arise due to issues such as background noise while recording heart sounds. Also, some algorithms require the inclusion of the electrocardiogram signal to complete segmentation [5, 9-10].

The intent of this research is to first develop a working classifier system using ANNs that accepts heart sound recordings directly, processes the sounds, and provides a classification as the output based on how the ANN is trained. An important goal is to eliminate the need for human interaction in selecting the best part of the heart sound to use as input. The research begins by developing a system which works using simulated heart sounds, and then system testing is extended to include heart sounds recorded from patients using a commercially available electronic stethoscope and a common laptop running a Windows operating system.

III. Methods

The design and implementation of the classifier system will be described in the following order:

- Types of heart sounds that can be detected and classified by the ANN system.
- Heart sounds database used for design and testing.
- Type of data and pre-processing steps that are used to provide design and test vectors to the ANN classifier.
- ANN architecture used to create the classifier.
- Steps that taken to design the ANN system.
- Procedure used to test the ANN system.
- Procedure used to test 2nd year medical students.

A. Types of heart sounds that are detected by the ANN system

Normal Heart – A normal heart sound consists of only a distinct first and second sound, commonly referred to as S1 and S2, respectively. S1 marks the beginning of systole, during which the heart contracts and the ventricles pump blood through the aorta and pulmonary artery. S2 marks the beginning of diastole, which is relaxation of the heart while the ventricles fill with blood [12].

Aortic Stenosis (AS)– This heart murmur is associated with the abnormal narrowing of the aortic valve, which impedes normal blood flow during systole [12].

Aortic Regurgitation (AR)– Also known as aortic insufficiency, this heart murmur is the result of blood that leaks back from the aorta to the left ventricle during diastole due to the failure of the aortic valve to close properly [12].

Mitral Regurgitation (MR)– Also known as mitral insufficiency, this murmur is defined as the backward flow of blood into the atrium (from the left ventricle) due to failure of the mitral valve to completely close. Mitral regurgitation is defined as a systolic murmur [12].

All diastolic murmurs are considered pathological, while systolic murmurs can be either innocent or pathological depending on the degree of severity of the condition. The American College of Cardiology (ACC) suggests that if the auscultator finds that a murmur is grade 3 or higher, or if the murmur is end-systolic, holosystolic, or diastolic, then a referral to echocardiogram is recommended [12]. This grading system assigns grades of 1 to 6 to murmurs. It was developed in 1933 and is entirely subjective: grade 1 is soft; grade 2 is moderate, etc. [13].

B. Heart sounds database used for designing and testing

The heart sounds used to design the ANN system were recorded from the Phono-Cardio-Simulator. The heart sound simulator was provided by the University of Minnesota School of Medicine Duluth, and it is commonly used to aurally teach heart murmurs to medical students. The heart sound library consists of 24 normal sounds, 24 AR sounds, 24 AS sounds (12 innocent and 12 pathological), and 24 MR sounds (12 innocent and 12 pathological). The sounds within each group of 24 files varied in heart rate, amplitude and duration of murmur.

The Murmur Study Library from the University of Minnesota Duluth provides recorded patient sounds that can be used to test the system. These patient heart sounds were recorded from four locations on the chest using an electronic stethoscope provided by Welch Allyn and a laptop running Windows XP.

Currently, the heart sound databank includes a total of more than 110 sounds from over 28 patients, including more than four types of murmurs. One of the future endeavors for this research is more extensive testing using recorded patient sounds. At present, recorded patient sounds used for testing include 7 examples of normal, 4 examples of AS, 2 examples of MR and 2 examples of AR. These 15 patient heart sounds were selected as good representations of the heart murmurs which this research focuses on.

Since the set of 15 patient sounds is a very limited sample size, the average single cycle step was omitted for testing patient sounds. Instead, individual heart cycles were extracted from these 15 patient sounds, resulting in 22 normal cycles, 13 AS cycles, 4 MR cycles and 6 AR cycles. Each of these 45 cycles then represents the type of heart sound from which it was extracted. The application of the average cycle technique will be performed once a sufficient number of patient sounds is available for testing.

One limitation of the recorded patient sounds is excessive background noise which impairs correct operation of the segmentation and alignment algorithms necessary for optimal use with the ANN system. Once a filter module is designed and included in the complete system in order to sufficiently reduce the background noise, in-depth analysis of test results using more recorded patient sounds can take place. Another limitation of the recorded patient sounds is that the echocardiogram reports used to assign a particular type of murmur condition to each sound are often vague—several types of trace murmurs can accompany a severe murmur in a single patient sound, for example. Presently, the system has only been designed to detect one type of murmur present in a single sound. Future work could potentially explore multiple murmurs in single sounds.

C. Data Type and Pre-Processing Steps Used for ANN Vectors

The heart sound data that is provided to the complete system must be in .wav audio format. There are numerous commercially available electronic stethoscopes capable of recording heart sounds and transferring those sounds to a computer in .wav format. Heart sounds of virtually any heart rate or duration can be inputted to the system. Sounds with large file size, i.e. high sampling rates or long duration, will result in longer computation time for the system, but otherwise the system results are unchanged.

The first pre-processing step uses a segmentation algorithm to identify the heart sound components S1 and S2 within the original heart sound (Fig. 1). The segmentation method is based on the low-frequency components of the signal. The spectrogram of the full heart sound is computed and the 45Hz frequency band is extracted (Fig. 2). The high amplitude components of the 45Hz band of the heart sound are observed, and S1 and S2 are identified based on the timing between those high amplitude components. The fact that the time from S1 to S2 (systole) is always less than the time from S2 to S1 (diastole) is the basis for this algorithm.

The time locations of S1s and S2s are determined, and then single heart cycles within the whole heart sound can be identified and re-sampled to normalize their lengths. Only a single heart cycle is used for the analysis and classification of

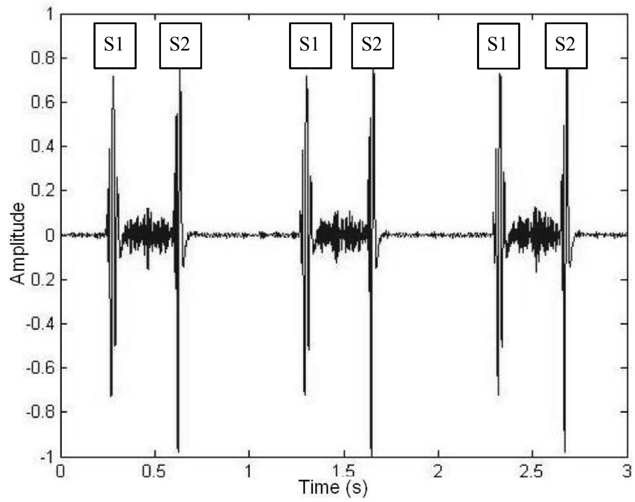

Fig 1. The time vs. amplitude plot of a heart sound with severe AS.

the sound. Instead of taking any of the individual cycles from a heart sound for the analysis, the average of the single cycles is computed and the resulting average heart cycle for that heart sound will be the one that is processed and inputted into the ANN for design (or testing) (Fig. 3). Compared to the use of a single extracted heart cycle from a heart sound, the average heart cycle technique provides a better representation of the complete heart sound. Also, the averaging technique helps to reduce the effect of anomalous heart cycles that may be present in a heart sound.

The next step in the pre-processing is to compute the spectrogram of the average heart cycle. Using the spectrogram, the 195Hz frequency band is extracted to serve as the input vector that will be fed into the ANN classifier. After studying and analyzing several frequency bands in the spectrogram for different heart sounds, we found that the 195Hz component

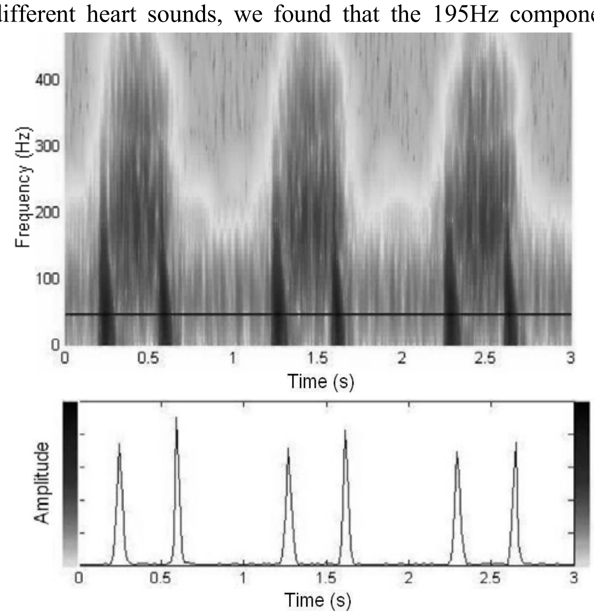

Fig 2. Illustration of the extraction of the 45Hz frequency component (bottom) of the heart sound using the spectrogram (top).

contains the necessary information for the ANN to identify the types of murmurs we are considering in our study. Fig. 4 illustrates the extraction of the 195Hz band from the spectrogram of the heart sound. The murmur is apparent in the 195Hz band—in the case of Fig. 4, the murmur occurs between S1 and S2. The group of 195Hz frequency bands from all average heart cycles in the design (or test) set makes up the set of vectors that will be inputted into the ANN system for design (or testing). First, though, an alignment algorithm is applied.

In order to account for the variability in S1-S2 intervals between heart sounds with different heart rates, an alignment algorithm is applied to the design (or test) vector set. The total number of data points to use for the input vectors to the ANN structure is first set—for example, 25 data points are used for every input vector to match the 25 input neurons of the ANN. The alignment is performed by sampling 60% of the desired number of total data points equally spaced from the end of S2 to the end of S1 and sampling the remaining 40% of total data points from the end of S1 to the end of S2. The result is that for different sounds, regardless of the time between S1 and S2, the same number of data points are taken from before (60% of total data points) and after (40% of total data points) the end of S1 (Fig. 5).

Sampling at a fixed number of data points across the single heart cycle maintains sufficient detail that was found in the original frequency slice vector while eliminating some redundancy of information. For example, using 50 data points the sample spacing results in a time resolution of 20 milliseconds per sample for a heart rate of 60 beats per minute (bpm). This resolution becomes finer as the heart rate increases, and thus sufficient detail is maintained within a heart cycle. The alignment algorithm ensures that the input data is always presented to the input neurons of the ANN in the same order: from the end of S2 to the end of the subsequent S2.

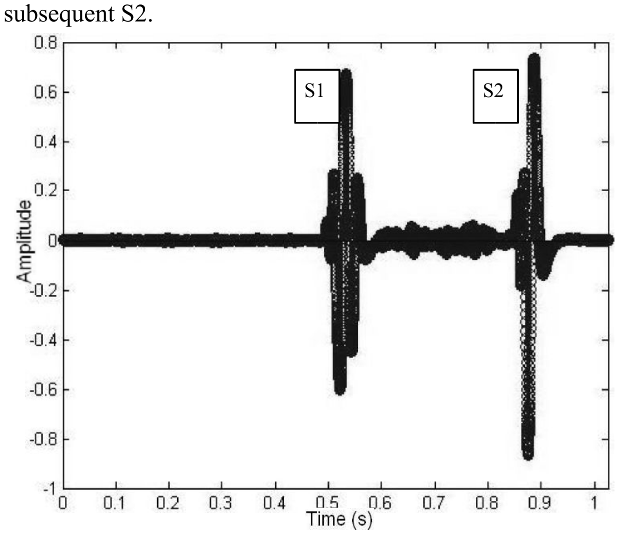

Fig 3. The average heart cycle—the result of averaging together the single cycles found by applying the segmentation algorithm to the heart sound in Fig. 1.

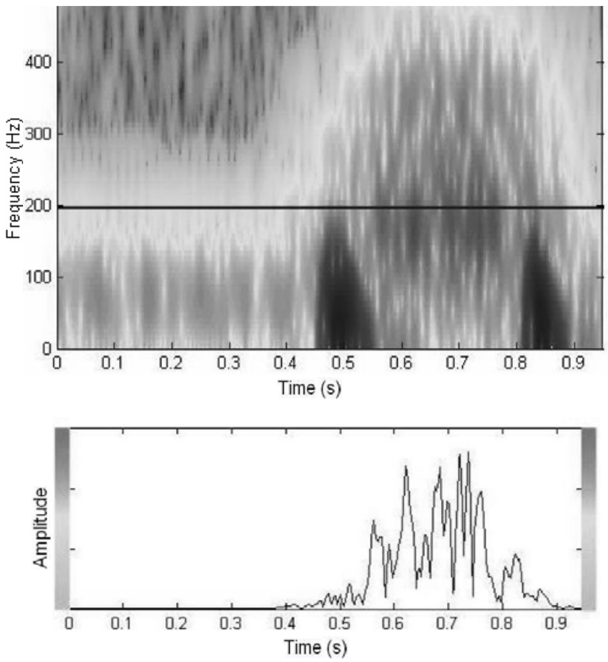

Fig 4. Illustration of the extraction of the 195Hz frequency component (bottom) of the heart sound using the spectrogram (top).

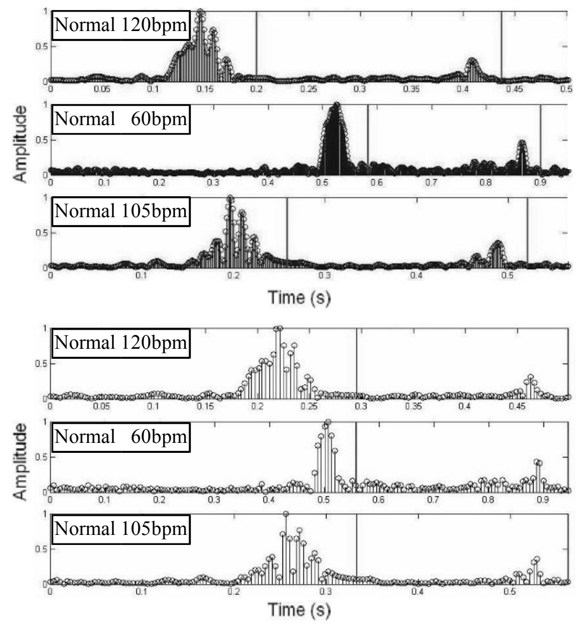

Fig 5. Before (top) and after (bottom) the alignment algorithm is applied to a set of 195Hz bands for 3 different normal heart sounds. The vertical lines mark the end of S1 (left) and S2 (right) in the plots.

D. ANN architecture used to create the classifier

Resilient back-propagation was used as the training algorithm for the feed-forward ANN classifier since it is recognized as a good algorithm for the purpose of pattern recognition [14]. The results shown in this report are from an ANN system using 3 hidden layers with 25 neurons each, 25 input neurons, one output neuron, a training mean squared error goal of 0.0005, and a learning rate of 0.0005. These parameters were optimized by repetition and comparison, with consideration from previous work [11].

E. Design of the ANN system

The process of designing the ANN system encompasses both the training and validation steps that are commonly referred to in ANN literature [15, 6-7]. In order to design the ANN classifier, a set of training vectors based on the desired training targets must first be provided. The user specifies separate directories on their computer for each desired target. Each of these directories contains only one type of heart sound—for example the group of 24 unique normal sounds or the group of 12 innocent AS sounds. When a directory is selected to represent a training target, the system randomly selects 50% of the files within that directory to use as the design set for that target, while the other 50% of the files are left to use as the test set. Once the user specifies that they are ready to train the neural network (i.e. all of the desired targets have been defined), the system is trained and all ANN characteristics are fixed.

As an example of a classifier design set, consider the target set which trains the ANN system to detect normal, AR and AS. Three training target directories are selected. Each of the three directories contains 24 unique heart sound files, thus each of the three training target sets consists of 12 unique heart sound files. The remaining files in the three directories that were not used for the design are used as the test set.

Once the ANN classifier has been designed, the test set(s) are used as input to the system.

F. Procedure used to test the ANN system

Once the ANN classifier has been designed and created, the user can either observe the output results using the test set or using any single heart sound file on their computer in order to use it as input. An output is then provided which indicates which training target is most similar to the current test input.

A comparison of the simulated heart sounds used in this research to heart sounds recorded from an electronic stethoscope shows that, after some signal processing, the sounds compare quite closely. Fig. 6 shows a comparison between two normal heart sounds: the first recorded from the Phono-Cardio-Simulator and the second recorded from an electronic stethoscope from Welch Allyn. Thus, designing the system with simulated sounds is expected to yield useful results for testing both simulated and patient-recorded heart sounds.

G. Procedure used to test 2nd year medical students

In order to put the ANN system accuracy into perspective, a group of 2nd year medical students at the University of Minnesota School of Medicine Duluth was asked to classify a set of simulated heart sounds taken from the same set used for ANN testing. To provide a side-by-side comparison, the accuracy of the 2nd year medical students is shown for each

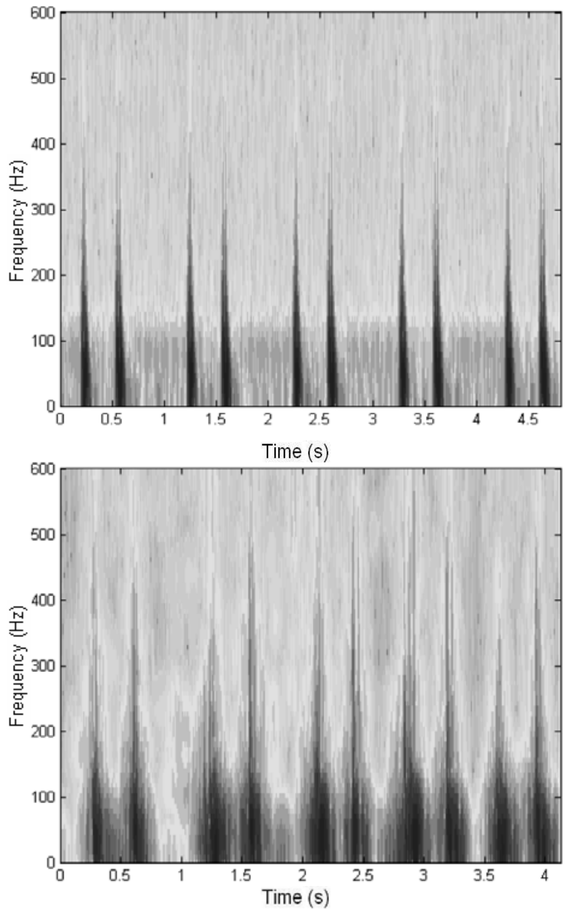

Fig. 6. Comparison of a heart sounds recorded from the Phono-Cardio-Simulator (top) and from the Welch Allyn stethoscope (bottom).

corresponding set of target sounds—for example, for the ANN accuracy using the normal vs. AR vs. AS target set, the medical student accuracy represents the accuracy for detecting only normal, AR and AS sounds.

IV. Results and analysis

The ANN architecture was determined through comparison of preliminary test results. To determine the optimal number of hidden layers and of neurons to use in the input and hidden layers, system complexity was reduced until performance began to degrade. The result of this procedure was an ANN structure using 3 hidden layers with 25 neurons in the input and hidden layers.

Note that for this research only one frequency component (195Hz) is used as input to the ANN system, while multiple frequency components could be used in one system. The GUI used to create the system allows users to select multiple frequency components, and a unique ANN system is created for each specific frequency. The single frequency component is described in this work in order to simplify the analysis of results, but the use of multiple frequency bands in a single ANN system remains to be explored further in future work.

Now that 3 hidden layers and 25 nodes per layer have been selected as the best configuration that has been analyzed, sensitivity and specificity can be explored. The sensitivity, given by (1), is a very important measure for this particular research since it is a measure of the percentage of patients with unhealthy hearts that are recognized as such. With a high sensitivity, the system has fewer Type II errors—the case when an unhealthy heart is classified as healthy.

$$Sensitivity = \frac{(true\ positives)}{(true\ positives + false\ negatives)} \qquad (1)$$

The specificity, given by (2), is the percentage of healthy cases that are classified as healthy. With a high specificity, the system has fewer Type I errors—the case when a healthy heart is classified as unhealthy.

$$Specificity = \frac{(true\ negatives)}{(true\ negatives + false\ positives)} \qquad (2)$$

For this system, high sensitivity is more important than high specificity. Higher sensitivity increases the number of patients with healthy hearts who are told they are unhealthy and sent to echocardiogram for further testing. More importantly, higher sensitivity reduces the number of patients with an unhealthy heart that are told they are healthy. On the other hand, higher specificity reduces the number of healthy patients who are referred to echocardiogram for further testing, while it increases the number of patients who have a heart murmur but are still told that they have a healthy heart—they are released from care with a potentially deadly heart condition. Therefore, high sensitivity is a much more desirable goal for this system.

The accuracy of the 2nd year medical students is compared to the accuracy of the ANN system in Table I for the three different target sets. The intervals listed are the 95% confidence intervals for the true mean values. The results are also displayed graphically in Fig. 7, 8, and 9.

The sensitivity for the medical students has not been reported in Table I because a detailed sensitivity analysis remains to be completed. However, the general trend is that, for the medical students, there are few false-negative diagnoses, i.e. the students rarely classified a heart sound as normal when a murmur was in fact present. The most common error is that the student recognizes that a murmur is present, but they are unable to identify the type of murmur.

Observing these results, there are several important features to note. First, the reported auscultative accuracy rate of primary care physicians is between 20-40% [1, 4, 16]. Although this range is the accuracy for physicians who are detecting heart murmurs using patient sounds, the accuracy rate for the 2nd year medical students observed in this study is similarly low when trying to classify simulated heart sounds. For each of the three target sets used in this research, the ANN system clearly has far better accuracy than the medical students.

Another important piece of information is that, in general, expert cardiologists have an auscultative accuracy rate of roughly 80% [1, 11]. The 95% confidence intervals given for the accuracy rate of the ANN system each include the 80% accuracy level, which indicates that if these results could be reproduced using patient sounds, this particular classifier system could prove very useful as an aide for primary care physicians before deciding whether or not to refer a patient for an echocardiogram.

System performance using patient sounds was less impressive than performance using simulated sounds. For the system trained to detect and classify simulated sounds representing normal, AS, and AR, the system achieved an average of 85% accuracy. However, when testing this same system with recorded patient sounds, the average accuracy was only 48.7 ± 12.7%.

There are several factors that could have led to the low accuracy when testing with patient sounds. One issue is that even after filtering the patient sounds, there is some background noise present. The system seems to have difficulty classifying sounds with excess background noise when it has been trained using simulated sounds with a very low level of background noise. Also, the alignment algorithm was not able to properly align several of the 45 patient heart cycles. In these cases the system would be less likely to correctly classify the sounds, since testing without properly applying the alignment algorithm means that S1 and S2, and subsequently systole and diastole, will not be applied to the same input neurons that were used during system training.

One potential solution to address the background noise in patient sounds is to train the system with noisy simulated sounds. Instead of using the simulated sounds directly to train the ANN, white Gaussian noise could be added to the simulated sounds, and then those noisy simulated sounds could be used to train the ANN system. Additionally, more advanced filtering methods could be applied to the patient sounds. Literature suggests that wavelet analysis is a powerful tool for extracting only the meaningful portions of heart sounds, and this may lead to successful noise suppression [6-9]. The noise reduction within the patient sounds would also help to reduce the failure rate of the alignment algorithm, since the main reason for failure of alignment is the presence of high amplitude anomalies at very low frequencies. It may also be possible to design a simple physical device to use with the stethoscope in order to reduce the likelihood of moving the chest-piece while recording patient sounds—a main cause of the low frequency background noise. Higher frequency noises, such as respiratory sounds, do not effect the alignment significantly since only low frequency bands are used to align sounds.

The current accuracy rate of the system for patient sounds is clearly not high enough to significantly improve a real-world situation. However, the suggested changes to the overall system design are expected to improve performance.

TABLE I
Comparison of Accuracy for Medical Students and ANN

Targets	Med Student Accuracy	ANN Accuracy	ANN Sensitivity
N, AS, AR	38.1 ± 11.7%	85 ± 7.4%	95 ± 6.8%
N, AS, AR, MR	27.9 ± 11.6%	76 ± 6.1%	89.7 ± 5.9%
AS-Inn., AS-Path.	19.0 ± 13.2%	70 ± 9.8%	67 ± 11.8%

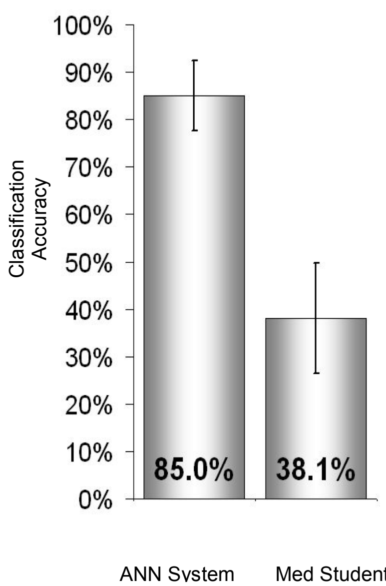

Fig. 7. Classification accuracy for the ANN system versus medical students for normal, AR, and AS.

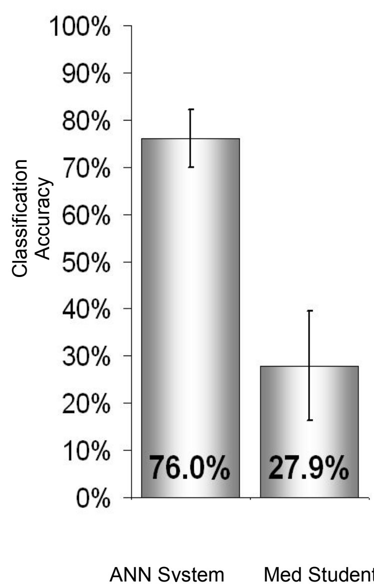

Fig. 8. Classification accuracy for the ANN system versus medical students for normal, AR, AS, and MR.

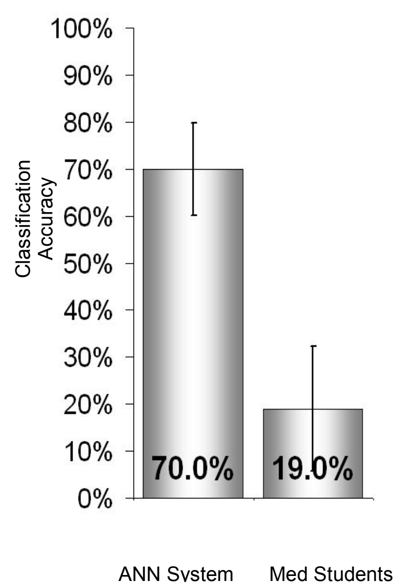

Fig. 9. Classification accuracy for the ANN system versus medical students for innocent AS and pathological AS.

V. CONCLUSION

The auscultative accuracy rate of the average physician is clearly low, and this fact leads to the referral of healthy patients for echocardiogram. Unnecessary referral to this costly procedure could be reduced if an inexpensive yet reliable diagnostic tool were available as an aide for physicians. The software system proposed in this work attempts to provide such a tool.

The software system provides the user with an output classification for an unknown heart sound, and this information could prove useful for a physician to consider when deciding whether or not to refer a patient for echocardiogram. It is expected that future research in noise reduction methods will lead to even better rates of classification.

Future work will emphasize pre-processing steps that will reduce the background noise levels present in the recorded patient sounds. As mentioned previously, wavelet analysis may be explored, and an attempt will be made to train the system using noisy simulated sounds. When the set of available patient sounds is large enough, an attempt will be made to train the system using patient sounds and observe performance. Additionally, research will include a slight change in the type of input vector used. Results are expected to improve if S1 and S2 are excluded from the input vector, since murmurs will occur over systole and diastole but they will be virtually undetectable during S1 and S2. In any case, it is important to show that the accuracy of the ANN system is high enough using patient sounds, not only simulated sounds, before it can be considered for a real-world application.

ACKNOWLEDGMENT

Thanks to the University of Minnesota Duluth Department of Electrical and Computer Engineering, the University of Minnesota School of Medicine Duluth and St. Luke's Hospital for motivation and providing the means to conduct this research. The following faculty members were incredibly helpful: Dr. Rocio Alba-Flores, Dr. Stanley Burns and Dr. Fernando Rios-Gutierrez from the UMD ECE Department; Dr. Glenn Nordehn from the University of Minnesota School of Medicine Duluth; Carolyn Otruba-Visina and the Diagnostics Lab staff from St. Luke's Hospital; Diane Anderson with the Whiteside Institute at SLH; and former UMD students Nick Andrisevic and Khaled Ejaz. Also, thanks to the Whiteside Institute for funding these research efforts so generously.

REFERENCES

[1] E. Etchells, C. Bell, and K. Robb, "Does this Patient Have an Abnormal Systolic Murmur," JAMA, vol. 277, pp. 564-571, 1997.

[2] AHA Scientific Statement: AHA Guidelines for Primary Prevention of Cardiovascular Disease and Stroke: 2002 Update, #71-0226 Circulation. 2002; vol. 106: pp. 388-391.

[3] American Society of Echocardiography, "Frequently Asked Questions about Cardiography," www.seemyheart.org/faq.php. Accessed October 1, 2006.

[4] S. Mangione and L.Z. Nieman, "Cardiac Auscultatory Skills of Internal Medicine and Family Practice Trainees: A Comparison of Diagnostic Proficiency," JAMA, vol. 278(9), pp. 717-722, Sept 1997.

[5] T. S. Leung, et al., "Classification of heart sounds using time-frequency methods and artificial neural networks," Proceedings of the 22"d Annual EMBS International Conference, Chicago IL, pp. 988-991, July 23-28, 2000.

[6] I. Hartimo and H. Liang, "A Heart Sound Feature Extraction Algorithm Based on Wavelet Decomposition and Reconstruction.," Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Vol. 20, No 3, 1998.

[7] I. Hartimo and H. Liang, "A Feature Extraction Algorithm Based on Wavelet Packet Decomposition for Heart Sound Signals," IEEE, pp. 93-96, 1998.

[8] C. N. Gupta, et al., "Segmentation and Classification of Heart Sounds," CCECE/CCGEI, Saskatoon, IEEE, pp. 1674-77, May 2005.

[9] Leung, T.S.; White, P.R.; Collis, W.B.; Brown, E.; Salmon, A.P.; "Classification of heart sounds using time-frequency method and artificial neural networks," Proceedings of the 22nd Annual EMBS International Conference, IEEE, vol. 2, pp 988-991, July 2000.

[10] S. R. Bhatikar, C. DeGroff, and R. L. Mahajan, "A classifier based on the artificial neural network approach for cardiologic auscultation in pediatrics," A.I. in Medicine, 2005; vol. 33, pp. 251-260.

[11] K. Ejaz, G. Nordehn, R. Alba-Flores, F. Rios-Gutierrez, S. Burns, and N. Andrisevic, "Heart Murmur Detection System Using Spectrograms and Artificial Neural Networks," IASTED Conference: Circuits, Signals and Systems (CSS 2004), Nov. 28, 2004 – December 1, 2004.

[12] R.O. Bonow, et al., "ACC/AHA Task Force Report," JACC, vol. 32, no. 5, pp. 1486-1588, November 1998.

[13] A.R. Freeman and S. A. Levine, "The Clinical Significance of the Systolic Murmur", Annals of Internal Medicine VI (1933): 1371-1385.

[14] M. T. Hagan, H. B. Demuth, and M. H. Beale, *Neural Network Design*, Martin Hagan, January 2002.

[15] A. Disher, C. Ornes, and J. Sklansky, "A Visual Neural Network that Learns Perceptual Relationships," IEEE, pp.873-875, 1998.

[16] Lam M. Z. C., et al., "Factors influencing cardiac auscultation proficiency in physician trainees," Singapore Medical Journal, 2005; vol. 25: pp. 11-14.

[17] AHA, Heart Disease and Stroke Statistics: 2006 Update. Dallas, TX, American Heart Association, 2003 [Online].

[18] MATLAB Help Documentation: Neural Networks Toolbox Help. MATLAB v6.5.0.1, June 18, 2002.

[19] Z. Sharif, et al., "An expert system approach for classification of heart sounds and murmurs," ISSPA, pp. 739-742, August, 2001.

[20] N. Shamsuddin, M. N. Mustafa, S. Husin, M. N. Taib, "Classification of heart sounds using a multilayer feed-forward neural network," IEEE, pp. 87-90, 2005.

[21] F. Rios-Gutierrez, R. Alba-Flores, K. Ejaz, G. Nordehn, N. Andrisevic, S. Burns, "Classification of Four Types of Common Murmurs using Wavelets and a Learning Vector Quantization Network," 2006 International Joint Conference on Neural Networks, pp. 2206-13, July 2006.