

Detection of murmurs from heart sound recordings with deep residual networks

Lei Hu, Wenjie Cai, Xinyue Li, Jia Li

University of Shanghai for Science and Technology, Shanghai, China

Abstract

Cardiac auscultation is an effective method to screen hemodynamic abnormalities. This paper aims to propose an automated algorithm to identify the presence of murmurs in heart sounds from multiple auscultation locations and to determine whether the heart sounds signal is normal. Two methods are explored. In method one, we perform a series of pre-processing such as denoising and segmentation of the heart sounds signal, extract Log Mel-Spectrogram as features, and use fastai's built-in xResNet 18 pre-trained model for classification. In method two, we extract Mel-frequency cepstral coefficients (MFCCs) as features without any pre-processing and build a customized model based on deep residual networks using one-dimensional convolutional neural layers. Our team, USST_Med, received a challenging score of weighted accuracy of 0.56 and cost of 11114.003 on the validation dataset. The proposed method has potential for clinical application.

1. Introduction

According to the World Health Organization, it is estimated that 17.9 million people died from cardiovascular diseases (CVDs). In 2019, accounting for 32% of all deaths worldwide. Most populations in low and middle-income countries do not have access to comprehensive primary health care systems, which may delay the diagnosis and treatment of CVDs, and lead to early death. Therefore, it becomes important to explore early prevention methods for heart diseases [1]. Auscultation of the heart plays an important role in the early detection for CVDs. However, traditional auscultation methods are highly dependent on physicians' experience, and large-scale auscultation screening of a population is a big challenge for doctors. Therefore, the automatic classification of heart sounds signal is of great practical importance for the screening and diagnosis of heart disease [2].

The fundamental heart sounds (FHS) signal usually includes the first (S1) and second (S2) heart sounds. S1 occurs at the beginning of ventricular systole and S2 occurs at the beginning of diastole. The diastole interval is usually

longer than the systole interval. Although the FHS is the most recognizable sound in the cardiac cycle, the mechanical activity of the heart may also cause other audible sounds such as the third heart sounds (S3), the fourth heart sounds (S4), systolic ejection clicks (EC), mid-systolic clicks (MC), diastolic sounds or open pops (OS), and cardiac murmurs caused by turbulent, and high-speed blood flow [3].

Heart sounds segmentation is a crucial step for feature extraction and classification of heart sounds. The primary purpose of segmentation is to divide the heart sounds into four parts: first heart sounds (S1), systole, second heart sounds (S2), and diastole [4]. Many heart sound segmentation methods have been proposed, including ECG signal-based segmentation methods, envelope-based segmentation methods, feature-based segmentation methods, machine learning-based segmentation methods, Hidden Markov Model (HMM)-based segmentation methods, and deep learning-based segmentation methods [5].

After the segmentation, feature extraction is used to convert the original high-dimensional heart sounds signal into low-dimensional features to facilitate the analysis of the heart sounds [4]. In general, the extracted features can be classified into three main types: time domain-based, frequency domain-based, and time-frequency domain-based features.

In the final step, the extracted features are fed into a classifier for feature learning and classification. Recently, with the rapid development of artificial intelligence, deep neural networks (DNN) have been explored for human heart sounds classification. The significant advantage of deep learning algorithms over traditional machine learning algorithms is the feature extraction function for complex heart sounds [6]. Most of the existing works use 1D/2D convolutional neural networks (CNN) [7], recurrent neural networks (RNN) [8], or deep convolutional and recurrent neural networks (CRNN) as their classifiers [6].

2. Methods

Two methods are applied in this challenge [9,10]. Method one uses the discrete wavelet transform (DWT) to denoise the heart sounds signal, and a proposed segmentation method to segment the signal. Then the Log

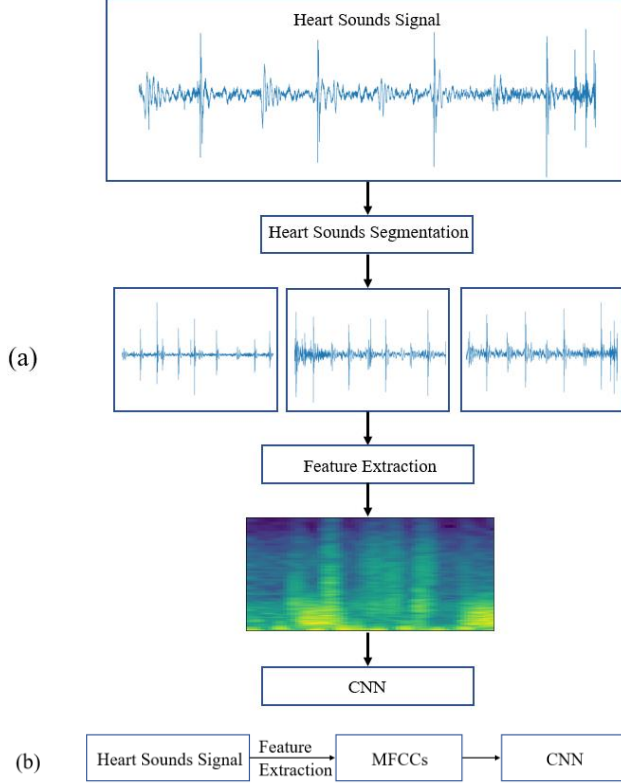


Figure 1. The system-level flow diagram of the proposed methods, (a) denotes method one, and (b) denotes method two.

Mel-Spectrogram is extracted as input features. Finally, a transfer learning method is used to classify the heart sounds signal. The deep learning model is implemented with fastai [11]. Method two requires feature extraction of Mel-frequency cepstral coefficients (MFCCs) and uses customized deep residual networks to classify the heart sounds. The system-level flow diagram of the proposed methods is shown in Figure 1.

2.1. Heart Sounds Segmentation

Heart sounds signal is interfered by various factors and noises, such as electromagnetic interference (EMI) from the surrounding environment, industrial frequency interference, electrical signal interference from the human body, and breath sound [12]. The presence of noise makes it difficult to localize S1 and S2. Thus, it is necessary to denoise the signal before segmentation.

In this paper, a fifth-level discrete wavelet decomposition using order six Daubechies filter of the original signal is performed to obtain the approximate coefficients, the third-level, fourth-level, and fifth-level detail coefficients. Then they are used to reconstruct the signal. The framework of the corresponding heart sounds signal wavelet denoising algorithm is shown in Figure 2. The denoising effect based on wavelet decomposition is

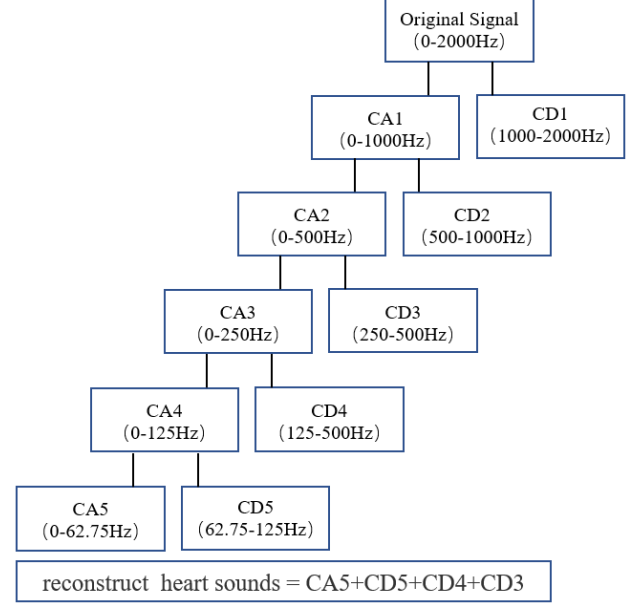


Figure 2. The framework of the corresponding heart sounds signal wavelet denoising algorithm.

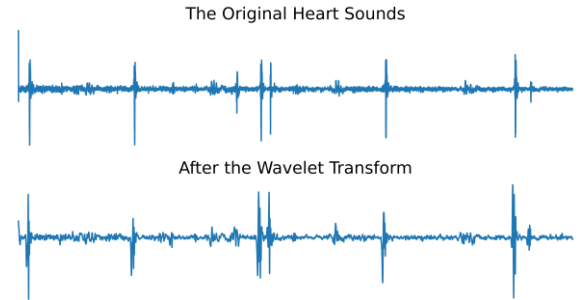


Figure 3. The original heart sounds signal and the signal after discrete wavelet transform.

shown in Figure 3.

After denoising, the heart sounds signal was normalized by equation 1 and equation 2:

$$X(t) = \frac{x(t) - \mu}{\sigma} \quad (1)$$

where $x(t)$ is the filtered heart sounds signal, μ is the mean of the filtered signal, and σ is the standard deviation of the filtered signal.

$$s(t) = \frac{X(t)}{\max(X(t))} \quad (2)$$

In this work, instead of choosing the envelope-based segmentation method, we directly localize the peak of the filtered signal.

First, the filtered signal $s(t)$ is differentiated to provide information about the slope by equation 3:

$$d[t] = s[t+1] - s[t] \quad (3)$$

Then, points with a positive slope are marked as 1, and others are marked as 0. Only when the current point is

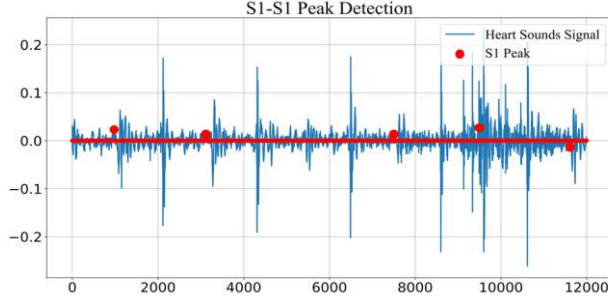


Figure 4. Heart sounds signal and S1 peaks.

marked as 1 and the next point is marked as 0, it is considered a potential S1 peak. We use the Python package PeakUtils to detect the final peak. The threshold is set as 0.3 and the minimum distance between each detected peak is set as 1800 after several experiments. At last, all S1 peaks are identified for segmentation. The S1 peak identification effect is shown in Figure 4.

After the S1 peak detection of the heart sounds signal, we segment the signal to every five cardiac cycles (S1-S1) until the last S1. Due to the inconsistent length of the segmented heart sounds signal, we fixed the signal length to 8800 and used zero padding for insufficient length. The distributions of two types of labels of the segments are shown in Table I and Table II.

Table I. The sample distribution of “Present” “Absent” and “Unknown”.

Labels	Number	Percentage
Present	1963	17.5
Absent	8962	80
Unknown	290	2.5
Total	11215	

Table II. The sample distribution of “Abnormal” and “Normal”.

Labels	Number	Percentage
Abnormal	5482	49
Normal	5733	51
Total	11215	

2.2. Feature extraction

The Log Mel-spectrogram and MFCC have been widely used in audio signal processing. In this stage, due to the excellent performance of CNN in image processing, we convert the audio signal processing problem into image processing by extracting frequency features. In method one, we extract the Log Mel-Spectrogram out of each segment heart sounds signal. Since the imbalance of our dataset, spectrogram augmentation techniques are applied [13]. A random time and frequency domain mask on the Log Mel-

spectrogram is used in this work.

In method two, we consider the MFCCs features as time-series data instead of images. For the classification of the presence, absence, or unclear cases of murmurs, we randomly segment MFCCs in a shape of (13, 2500). For the classification of the normal and abnormal clinical outcomes, we randomly segment MFCCs in a shape of (13, 1500). On both classification tasks, zero padding is used for insufficient length.

2.3. Classification

In method one, instead of building a model from scratch, we select the xResNet 18 pre-trained model built into fastai to classify the heart sounds signal. Fastai has a very easy-to-use workflow, which makes the process of debugging much simpler. We use the “lr_find” method to find the optimal learning rate, and in this work, we choose $3e-3$. Due to the “fit_one_cycle” method, the learning rate is uniformly varied over 20 training epochs, Callbacks and EarlyStopping are used to efficiently track training and validation losses at the specified learning rate.

The CNN network architecture in method two is shown in Figure 5. One-dimensional convolution layers are used to extract class-related features, and the main architecture is based on eight residual blocks [14]. For both tasks, the heart sounds signal from multiple auscultation locations is input to the network. All the outputs from the global max-pooling layer are concatenated and fed into fully connected layers with the Sigmoid function in the last layer to get the result of classification.

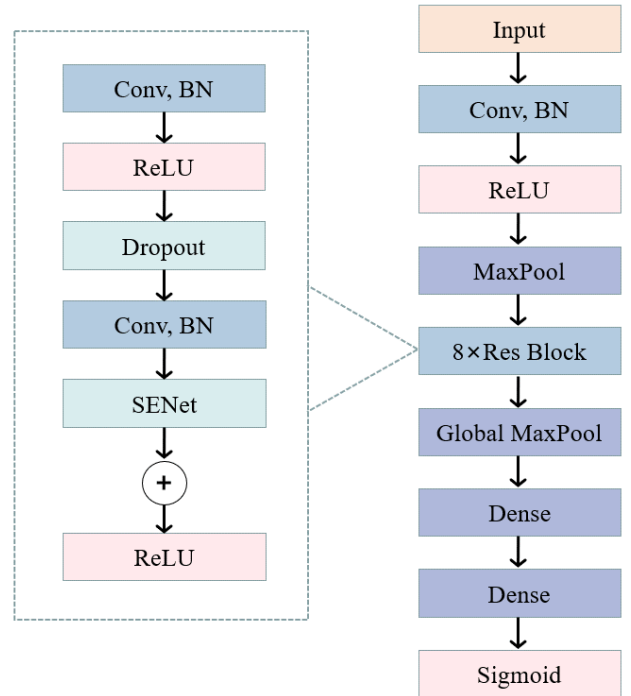


Figure 5. The CNN network architecture used in method two.

3. Results

After model training, we evaluated the performance of the classifier on the validation dataset. We got a score of weighted accuracy of 0.518 and cost of 11752.231 with method one. The second method achieved a score of weighted accuracy of 0.56 and cost of 11114.003. The official test score is based on method two. Its weighted accuracy is XXX and its cost is XXX.

4. Discussion and Conclusions

In this paper, we proposed two methods for heart sounds signal classification. Our results are not satisfactory. The pre-trained model used in method one is designed for natural images which are totally different from the extracted Log Mel-spectrogram. Due to relatively small amounts of data, fine-tuning does not achieve optimal results. The segmentation method is not robust enough to get an accurate result in the validation set, which limits the model's performance. In method two, we randomly segment fixed-length MFCCs features as model inputs. These inputs could be just some noisy signals, which may prevent the model from achieving good results.

In future work, we can try different denoising methods to obtain better quality heart sounds signal. We can conduct more in-depth studies for different hyperparameters (window size, hop length, number of mels, number of mfcc, FFT size, etc.) to obtain the optimal features when performing feature extraction. In addition, we can compare the classification effects of different types or sizes of pre-trained models. Finally, the CNN network architecture used in method two achieves a good result on electrocardiogram (ECG) classification, we can compare the differences between heart sounds and ECG to optimize the network architecture structure for better heart sounds signal classification.

In conclusion, although we did not achieve satisfactory scores, our approach at least demonstrates the good potential of deep learning for heart sounds classification.

References

- [1] J. Oliveira *et al.*, "The CirCor DigiScope Dataset: From Murmur Detection to Murmur Classification," *IEEE J. Biomed. Health Inform.*, vol. 26, no. 6, pp. 2524–2535, Jun. 2022.
- [2] J. F. Chen and X. Dang, "Heart Sound Analysis Based on Extended Features and Related Factors," in *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*, Xiamen, China, Dec. 2019, pp. 2189–2194.
- [3] C. Liu *et al.*, "An open access database for the evaluation of heart sound algorithms," *Physiol. Meas.*, vol. 37, no. 12, pp. 2181–2213, Dec. 2016.
- [4] W. Chen, Q. Sun, X. Chen, G. Xie, H. Wu, and C. Xu, "Deep Learning Methods for Heart Sounds Classification: A Systematic Review," *Entropy*, vol. 23, no. 6, p. 667, May 2021.
- [5] Y. He, W. Li, W. Zhang, S. Zhang, X. Pi, and H. Liu, "Research on Segmentation and Classification of Heart Sound Signals Based on Deep Learning," *Appl. Sci.*, vol. 11, no. 2, p. 651, Jan. 2021.
- [6] M. Deng, T. Meng, J. Cao, S. Wang, J. Zhang, and H. Fan, "Heart sound classification based on improved MFCC features and convolutional recurrent neural networks," *Neural Netw.*, vol. 130, pp. 22–32, Oct. 2020.
- [7] Z. Abduh, E. A. Nehary, M. Abdel Wahed, and Y. M. Kadah, "Classification of heart sounds using fractional fourier transform based mel-frequency spectral coefficients and traditional classifiers," *Biomed. Signal Process. Control*, vol. 57, p. 101788, Mar. 2020.
- [8] F. A. Khan, A. Abid, and M. S. Khan, "Automatic heart sound classification from segmented/unsegmented phonocardiogram signals using time and frequency features," *Physiol. Meas.*, vol. 41, no. 5, p. 055006, Jun. 2020.
- [9] M. A. Reyna *et al.*, "Heart Murmur Detection from Phonocardiogram Recordings: The George B. Moody PhysioNet Challenge 2022," *Health Informatics*, preprint, Aug. 2022.
- [10] A. L. Goldberger *et al.*, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals," *Circulation*, vol. 101, no. 23, Jun. 2000.
- [11] J. Howard and S. Gugger, "fastai: A Layered API for Deep Learning," *Information*, vol. 11, no. 2, p. 108, Feb. 2020.
- [12] M. N. Ali, E.-S. A. El-Dahshan, and A. H. Yahia, "Denoising of Heart Sound Signals Using Discrete Wavelet Transform," *Circuits Syst. Signal Process.*, vol. 36, no. 11, pp. 4482–4497, Nov. 2017.
- [13] D. S. Park *et al.*, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," in *Interspeech 2019*, Sep. 2019, pp. 2613–2617.
- [14] W. Cai, F. Liu, B. Xu, X. Wang, S. Hu, and M. Wang, "Classification of multi-lead ECG with deep residual convolutional neural networks," *Physiol. Meas.*, vol. 43, no. 7, p. 074003, Jul. 2022.

Address for correspondence:

Wenjie Cai
School of Health Science and Engineering,
University of Shanghai for Science and Technology,
516 Jungong Road, Yangpu Distric, Shanghai, China