Classification of Murmurs in PCG Using Combined Frequency Domain and Physician Inspired Features

Julia Ding[1], Jing-Jing Li[2], Max Xu[3]

[1] Emory University, Atlanta, GA, USA
[2] University of California at Berkeley, Berkeley, CA, USA
[3] Georgia Institute of Technology, Atlanta, GA, USA

Abstract

Physiological machine learning methods have a unique opportunity to augment deep-learning engineered features with additional features derived from prior pathological knowledge. We propose an phonocardiogram (PCG) classifier that combines raw spectrogram features with crafted, physician-inspired features with an end-to-end neural network architecture. Learned spectrogram features were obtained by training a convolutional neural network (CNN) directly on the raw mel-spectrogram representation of the PCG time-series. Crafted features were based on the four stages of the cardiac cycle (S1, systole, S2, and diastole) and the relative length intervals of each stage.

The spetrogram features have the advantage of introducing flexibility for the model to learn abstract, low-level information that captures a variety of different rhythmic abnormalities and the latter has the advantage of using segmentation to elucidate specific, high-level, human-interpretable information. Combined features are fed into a fully connected neural network which is able to learn the relationship between the two feature types.

1. Introduction

Auscultation of heart sounds for murmurs is vital in identifying cardiovascular disorders in children. In regions with lack of infrastructure or access to cardiology specialists, a non-invasive assessment of heart sounds via phonocardiogram (PCG) can provide life-saving information for pediatric patients with congenital or acquired heart disease. The analyzed data, provided by the Moody Challenge 2022, consisted of PCG recordings from 1568 pediatric patients collected in northeastern Brazil [1, 2].

Ensemble models that combine information across raw and curated feature types have found broad success, such as in the 2016 and 2020 PhysioNet Challenges [3–5]. We hypothesize that the neural network will be able to learn to identify heart murmur patterns within the time-frequency domain paired with our physician-inspired, human-interpretable information.

2. Methods

2.1. Preprocessing

For each patient, each of the PCG recordings at four classic auscultation locations are downsampled to 1000 Hz and preprocessed with a 25-400Hz 2nd order Butterworth band pass filter in order to remove noise. Spikes were removed by identifying deviations in the maximum absolute amplitude within 500ms windows [6].

2.2. CNN on Spectrogram Features

These time-series are then transformed into mel-spectrogram and concatenated to be used as inputs into a CNN composed of three sequential blocks of a convolution layer, leaky relu, max-pooling, and dropout. Following this, we utilize two fully-connected layers with a cross-entropy loss function. Batch normalization is applied after each block of layers.

2.3. Physician-inspired Features

Recordings were segmented into key components of the cardiac cycle (S1, systolic, S2, and diastolic intervals) using a pre-trained logistic regression, hidden semi-Markov model with a modified Viterbi decoding algorithm [7].

The mean, median, and standard deviation of these cardiac cycle intervals as well as between these intervals were calculated. Then, we used discrete Fourier transform to calculate frequency domain features post segmentation. For 10 Hz frequency bands between 30 and 800, the spectra of S1, systole, S2, and diastole were calculated [8].

These crafted features based on the cardiac cycle were concatenated with the spectrogram learned features and fed into a fully connected network to predict the label of the recording.

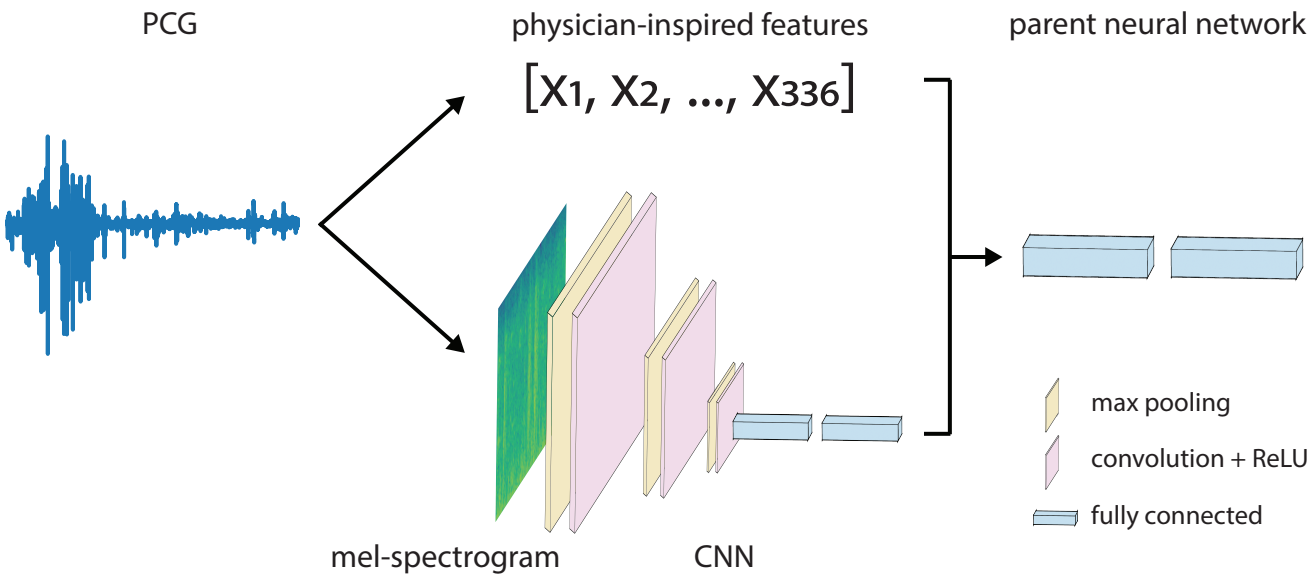

Figure 1. Network architecture. The PCG recording is used to extract both physician-inspired features and a mel-spectrogram. The spectrogram is input to a CNN containing three convolutional layers and two fully connected layers. Both the output of the CNN and the physician-inspired feature vector are input to a fully connected neural network, which predicts the label of the recording.

3. Results

3.1. Cross-validation

Here we report the 6-fold cross-validation performance of our method on the PhysioNet/CinC 2022 Challenge dataset (Fig. 2). Both loss and accuracy converged within 10 epochs, resulting in unweighted training accuracy of 73.8% and validation accuracy of 73.6%. Although accuracy plateaued after a few epochs, validation loss continued to decrease towards the end of training.

3.2. Challenge score

In the evaluation of the PhysioNet/CinC 2022 Challenge, our model achieved a weighted accuracy of 0.467. Our model was developed with Tensorflow-CPU version 2.4.1. With our entry's submission, our training took 7 hours 46 minutes, and model evaluation was completed in 1 hour and 20 minutes. Our model achieved a score of 13836 for outcome classification, as measured by screening cost.

4. Discussion

In conclusion, our model's end-to-end architecture allows it to flexibly combine physician-inspired and raw spectrogram features, and is able to achieve strong performance, well-above random chance. Additionally, we pro-

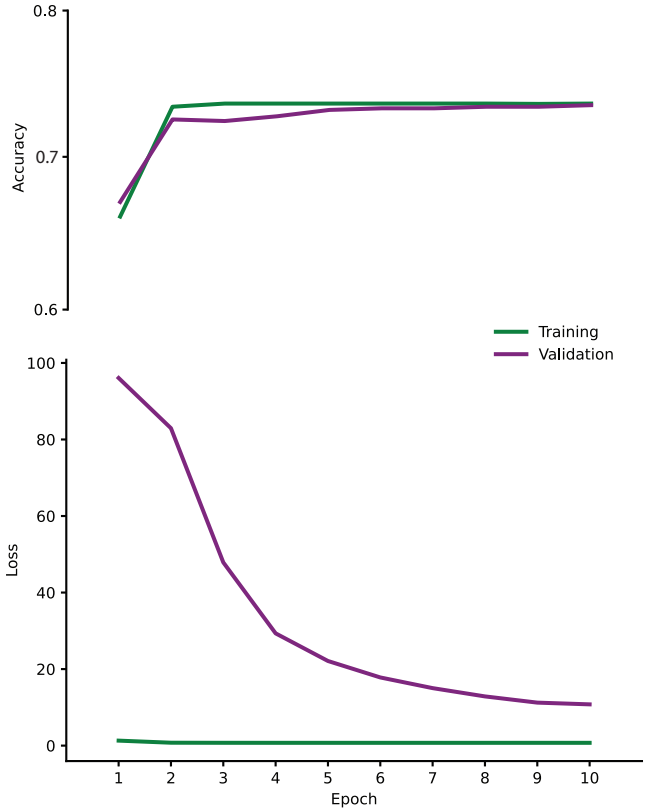

Figure 2. 6-fold cross-validation results.

vide the first public Python implementation of the widely-used PCG cardiac cycle segmentation algorithm [7] in our Github repository, which can be utilized by researchers in future works.

To further improve our algorithm, future research could optimize the hyper-parameters used in the model, which have not been carefully tuned. Extensive feature analysis could reveal the salience of each feature in classifying murmurs and highlight informative dimensions to assist diagnosis.

Most importantly, despite the deficiencies of our model, we hope our work could facilitate the development of murmur classification algorithms with human-level accuracy, efficiency, and reliability.

References

[1] Goldberger AL, Amaral LAN, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng C, Stanley HE. Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. Circulation June 2000;101(23):215–220.

[2] Oliveira J, Renna F, Costa PD, Nogueira M, Oliveira C, Ferreira C, Jorge A, Mattos S, Hatem T, Tavares T, Elola A, Rad AB, Sameni R, Clifford GD, Coimbra MT. The circor digiscope dataset: From murmur detection to murmur classification. IEEE Journal of Biomedical and Health Informatics 2022;26(6):2524–2535.

[3] Potes C, Parvaneh S, Rahman A, Conroy B. Ensemble of feature-based and deep learning-based classifiers for detection of abnormal heart sounds. In 2016 Computing in Cardiology Conference (CinC). 2016; 621–624.

[4] Oppelt MP, Riehl M, Kemeth FP, Steffan J. Combining scatter transform and deep neural networks for multilabel electrocardiogram signal classification. In 2020 Computing in Cardiology. 2020; 1–4.

[5] Alkhodari M, Hadjileontiadis LJ, Khandoker AH. Identification of cardiac arrhythmias from 12-lead ecg using beat-wise analysis and a combination of cnn and lstm. In 2020 Computing in Cardiology. 2020; 1–4.

[6] Schmidt SE, Holst-Hansen C, Graff C, Toft E, Struijk JJ. Segmentation of heart sound recordings by a duration-dependent hidden markov model. Physiol Meas Apr. 2010; 31(4):513–29.

[7] Springer DB, Tarassenko L, Clifford GD. Logistic regression-hsmm-based heart sound segmentation. IEEE Trans Biomed Eng Apr. 2016;63(4):822–832.

[8] Tang H, Dai Z, Jiang Y, Li T, Liu C. Pcg classification using multidomain features and svm classifier. BioMed Research International 2018;2018(Article ID 4205027).

Address for correspondence:

julia.ding@emory.edu