Phonocardiographic Murmur Detection by Scattering-Recurrent Networks

Philip A Warrick[1,2], Jonathan Afilalo[2]

[1]PeriGen Inc., Montreal, Canada, [2]McGill University, Montreal, Canada

Objectives: The phonocardiogram (PCG) is a recording of heart sound that provides a low-cost diagnostic tool to assess turbulent blood flow patterns ("murmurs") suggestive of heart valve dysfunction. Physicians' subjective interpretion of PCG sounds has historically low sensitivity and specificity for the diagnosis of pathological heart murmurs. To improve diagnostic performance, we developed a PCG heart murmur detector using the Scattering Transform (ST) and Long Short-Term Memory (LSTM) networks.

Methods: We used the PhysioNet/Computing in Cardiology Challenge 2022 training dataset that provides multiple PCG recordings for 942 patients. We presented the PCG signals to an ST representation layer, which reduced the sampling rate from f_{in}=4 kHz to f_{out}=250 Hz, and then 2 bidirectional LSTM layers, which captured feature trajectories over time. The final dense layer used cross-entropy loss during training to support three target classes: murmur present, absent or unknown. During training, we used the "murmur location" information to assign targets to corresponding recordings. We split recordings into 5 s segments for training batches. During evaluation, we classified each recording by selecting the highest average class probability over all segments for that recording. If any recording was classified as "murmur present" for a given patient, that patient was assigned a diagnosis of "murmur present". The system was developed with TensorFlow and the Kymatio ST package.

Results: Local 10-fold cross-validation test results achieved a Challenge score of 1064$\pm$278 (mean$\pm$standard deviation). Our entry ("team PCGPAW") successfully trained a model on the Challenge server and obtained a score of 1528 on the hidden test data consisting of 626 patients.

Conclusions: Our classifier architecture showed promising results in this first phase of development. In future work, we hope to improve the system by including more of the available training annotations, notably, the segmentation of diastole and systole timing.