Detection of Heart Murmurs in Phonocardiograms with Parallel Hidden Semi-Markov Models

Andrew McDonald*, Mark JF Gales, Anurag Agarwal

Department of Engineering, University of Cambridge
Cambridge, UK

Introduction: We describe a neural network and hidden semi Markov model (HSMM) approach to detect heart murmurs in phonocardiogram recordings. This model forms the 'CUED_Acoustics' entry to the 2022 George B. Moody PhysioNet challenge.

Methods: The algorithm first performs independent murmur detection on each patient's recordings. For numerous heart sound algorithms, including our third-placed 2016 PhysioNet challenge entry, the segmentation of the signal into its constituent sounds is a key feature extraction stage. However, many previous segmentation techniques assume the heart sound only contains healthy S1 and S2 sounds and so perform poorly in the presence of large-amplitude murmurs. Our approach directly models the presence of murmurs, simultaneously producing a murmur classification and signal segmentation. A frequency-normalised log-spectrogram is calculated from the recording and used as feature input for a recurrent neural network, which is trained to predict the heart sound state (silence, S1/S2, or murmur) of each feature time step. These predictions are used as observations for two parallel HSMMs. During systole, the first HSMM assumes a normal phonocardiogram with a 'silence' observation, whereas the second HSMM assumes a 'murmur' observation. The duration-dependent Viterbi algorithm is used to find the most likely state sequence for each model. A murmur is detected if the confidence in the murmur HSMM model fit is higher than in the normal HSMM. A gradient boosted decision tree then produces an overall patient classification using the prediction for each recording combined with patient biometrics such as age and weight.

Results: During the unofficial phase, the algorithm achieved a 5-fold cross-validation training score of 490 and a withheld validation score of 489. At phase end, we ranked 3rd out of 77 teams.

Conclusion: The model performs well and can potentially provide a method to screen for structural heart disease, improving early detection and reducing unnecessary referrals.