Deep Neural Networks for Heart Beat Detection

Fasil Cheema, Jabavu Adams, supervised by Yang Zhao, Dr. Peter Lian

Introduction

We used Deep Neural Networks (DNNs) to do real-time and robust heartbeat detection in potentially noisy and distorted electrocardiogram (ECG), and ballistocardiogram (BCG) signals. An ECG is a graph of the electrical activity in the heart. It is recorded by using electrodes placed on the skin. ECG is a common and useful diagnostic tool used to assess cardiac health and identify abnormalities.

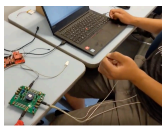

Collecting ECG

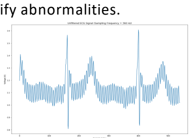

ECG Signal

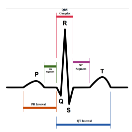

Ideal ECG Beat

Ballistocardiography measures the vibration of body tissues in response to blood pumped by each heartbeat. BCG sensors are cheaper, more rugged, and less invasive than ECG sensors. They can be integrated into seating, bedding, or wearable devices, because they do not require electrodes or electrical contact with the subject.

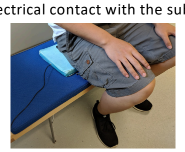

Collecting BCG

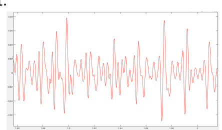

BCG Signal

Our ECG algorithms were trained on the MIT-BIH database[1][2] of annotated ECG signals. In this dataset, the peaks corresponding to heartbeats (R-peaks) have been annotated by cardiologists.

We are not aware of any datasets of BCG signals containing annotated heartbeat locations (J-peaks). Since the ECG measures electrical activity, and the BCG measures the motion resulting from blood pumping, there is a time-lag between a heartbeat in the ECG signal and in the BCG signal. We intend to use our ECG heartbeat detector and this time lag to generate heartbeat location annotations for BCG data, given simultaneously recorded ECG and BCG signals as input. We will then use machine learning (ML) to train a BCG heartbeat detector on this annotated BCG data.

Methodology

We use Deep Neural Networks (DNNs) as our machine learning model. DNNs are made of simple units (called neurons) that are interconnected in layers from input to output. The connections between units are weighted, and by "training", we mean modifying these numerical weights, so as to minimize the discrepancy between the computed output, and the desired output.

Our inputs are fixed-length subsequences of the ECG signal. We use 200 milliseconds of signal. Our output is a value representing the time at which the heartbeat (R-peak) occurs within this 200 millisecond window.

Before performing detection, we filter the signal to remove power-line noise and also high-frequency oscillations that are not relevant to heartbeats.

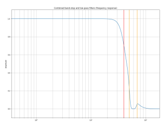

Filter frequency response

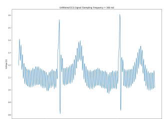

Unfiltered ECG

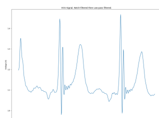

Filtered ECG

We created a data processing pipeline to chop long ECG recordings into short segments that were presented to our DNNs for training, after resampling, filtering and scaling.

Many choices for the number of neurons per layer, the number of layers, and other parameters are possible (the architecture). To find good choices of architecture, we developed software to queue up different choices of model parameters and learning settings. This allowed us to specify a sequence of training sessions to run to find good parameter choices.

We investigated two categories of models: a densely-connected network, and a set of two convolutional networks (CNNs). The dense network output a single value representing the peak time. The first-pass CNN detects whether any peak is present in the window. If it detects a peak, the second pass CNN is called and returns a discrete peak location range.

We also developed a signal viewer that can display a recorded signal or a live signal, while running a choice of detectors to find the peaks.

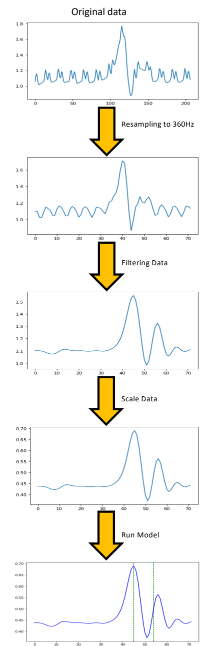

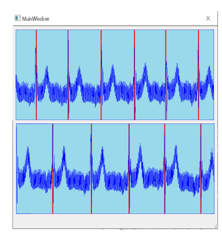

Results

Model	Dense DNN	CNN-1 (presence)	CNN-2 (range)
Accuracy	97.7%	97.3%	98.7%

We measured the accuracy of our models on test data from the MIT-BIH dataset that the model had not seen during training.

The best densely-connected model scored 97.7% accuracy (F1 score) where predicting the beat location to within 25 milliseconds was counted a hit, and outside that range was considered a miss, as in [6]

The first-pass CNN beat presence detector had an accuracy of 97.3% on the test set.

The second-pass CNN beat localizer had an accuracy of 98.7%

Future Work

We expect to be able to improve our models now that we have tools in place to rapidly try various parameter settings. We would also like to try Recursive Neural Networks (RNNs), as they show promise for time-series analysis. Finally, we were unable to examine the correlation between simultaneously-recorded ECG and BCG, in the time allotted. We hope to pursue this in the future.

References

1. Moody GB, Mark RG. The impact of the MIT-BIH Arrhythmia Database. IEEE Eng in Med and Biol 20(3):45-50 (May-June 2001). (PMID: 11446209)
2. Goldberger AL, Amaral LAN, Glass L, Hausdorff JM, Ivanov PCh, Mark RG, Mietus JE, Moody GB, Peng C-K, Stanley HE. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals (2003). Circulation. 101(23):e215-e220.
3. Yann LeCun, et al. (1999) Object Recognition with Gradient-Based Learning
4. Martin Abadi, et al. (2015) Tensorflow: Large-scale machine learning on heterogenous system, software available from tensorflow.org
5. Chollet, Francois et al (2015) Keras, software available from keras.io
6. Feifei Liu, Chengyu Liu, Xinge Jiang, et al., "Performance Analysis of Ten Common QRS Detectors on Different ECG Application Cases," Journal of Healthcare Engineering, vol. 2018, Article ID 9050812, 8 pages, 2018. https://doi.org/10.1155/2018/9050812.