Deep Learning for Improving Power-Accuracy of Heart Rate Monitors

Albert Gural (agural@stanford.edu)

Stanford University

Background

Continuous heart rate (HR) monitoring is becoming increasingly important for health tracking and diagnostics. Modern approaches use photoplethysmography (PPG) and typically consume a lot of power. This work explores techniques to reduce the power of PPG-based HR sensors to a level sufficient for continuous monitoring ($< 100\ \mu W$) while still maintaining robustness.

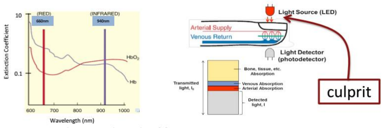

Basic principles of operation of PPGs, from [1].

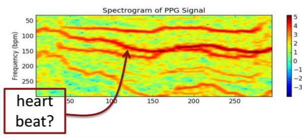

Prior Art

- [2]

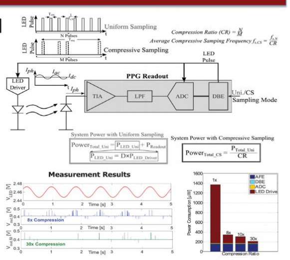

- [3]

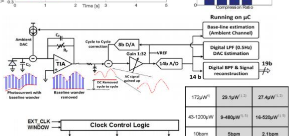

- [4]

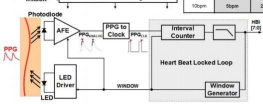

Experiment Design

Most work was conducted on data from the ICASSP signal processing cup challenge [5]. As in the challenge, we attempt to determine the heart rate for 8-second windows spaced 2-seconds apart. The target is to minimize average absolute error (avAE). We deviate from the challenge, however, in our goal to achieve low power by limiting sample rate to 12.5Hz (10x decimation of the original sample rate).

Training data comprises 12 samples of users undergoing a set exercise routine lasting 5 minutes long. The data includes PPG, accelerometer, and ECG data (ECG is for reference only).

We tested two approaches: a time-domain approach using sequential deep learning models and a frequency-domain approach using traditional deep learning. Because of the minimal amount of provided data, we augmented it with data synthesis.

Frequency-Domain Approach

Key idea: Feed FFTs of the PPG and accelerometer signals in each evaluation window to a neural network with categorical heart rate bins as outputs.

Details: Inspired by [6], we chose to feed two 64-point FFTs per signal to allow time-varying information to seep into the classification task. The network itself takes 4 FFTs (4x32) through three 512-node dense layers and a final softmax layer with 160 outputs (one for each frequency bin).

Results: On our most challenging synthetic dataset, the network achieves 58% accuracy. These results generalize weakly to the ICASSP data. ICASSP results below are the result of including the first six samples in the training set, so only the last six should be used to judge generalizability.

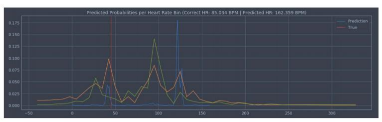

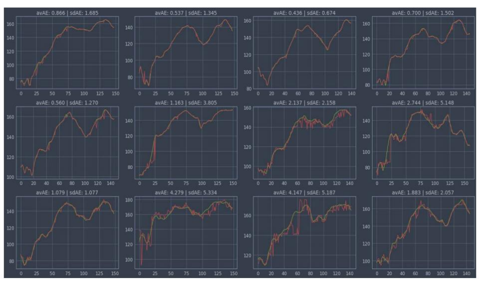

Time-Domain Approach

Key idea: Feed PPG and accelerometer signals for each 8-second evaluation window directly into an LSTM and output the heart rate as a category.

Details: The best model had PPG and accelerometer input (2x100) feed directly into a 512-unit LSTM, then to a 512-neuron dense layer, a dropout(0.5) layer, and finally a softmax layer with 160 outputs (one for each frequency bin 40…199 bpm).

Results: On a challenging synthetic dataset, the network achieves 70% accuracy. These results do not generalize well to the ICASSP data where there is a high standard deviation average error (sdAE).

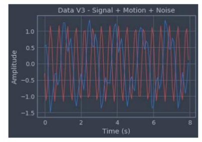

synthetic data improvement

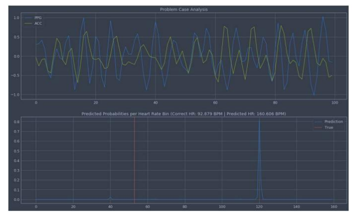

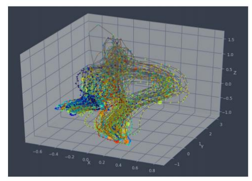

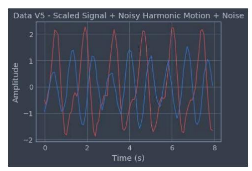

Results and Conclusion

The time-domain approach requires less up-front processing and is able to achieve very low error rates in most cases, but it is also very sensitive and fails often when run on the ICASSP data.

The frequency-domain approach requires FFT-preprocessing, but is a bit more robust especially on ICASSP, when compared to the time-domain approach. Because the evaluation windows are short, the FFTs have low frequency resolution and are thus less precise than the time-domain approach in optimal circumstances. It is possible that a combined TD and FD approach can achieve high precision and high reliability.

	Synthetic Data (1.5M Tr | 100k Dev)	ICASSP avAE | sdAE (884 Train | 884 Dev)
Time-Domain	72% | 70% (30 epochs)	5.9 BPM | 17.5 BPM
Freq-Domain	62% | 58% (30 epochs)	2.8 BPM | 8.3 BPM

Power: Independent (personal) testing on hardware suggests that the SNRs required for both of these algorithms implies an energy expenditure of ≈2µJ/sample. With these algorithms demonstrated at 12.5Hz, a system running at 25µW could achieve sub-3 BPM avAE, surpassing [4] in power, accuracy, and robustness.

However, this does not include the computational costs of performing multiple 512x512 matrix multiplies. When factoring the computational cost in, it is unlikely the system could achieve sub-100µW operation.

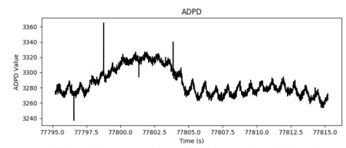

PPG signal from personally-owned hardware running at 2µJ/sample

References

1. Analog Devices. "An Introduction to PULSE OXYMETRY." 2016.
2. Rajesh, Pamula Venkata, et al. "22.4 A 172µW compressive sampling photoplethysmographic readout with embedded direct heart-rate and variability extraction from compressively sampled data." Solid-State Circuits Conference (ISSCC), 2016 IEEE International. IEEE, 2016.
3. Sharma, Ajit, et al. "A Sub-60-µA Multimodal Smart Biosensing SoC With> 80-dB SNR, 35-µA Photoplethysmography Signal Chain." IEEE Journal of Solid-State Circuits (2017).
4. Jang, Do-Hun, and SeongHwan Cho. "A 43.4 µW photoplethysmogram-based heart-rate sensor using heart-beat-locked loop." Solid-State Circuits Conference-(ISSCC), 2018 IEEE International. IEEE, 2018.
5. Z. Zhang, Z. Pi, B. Liu, TROIKA: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise, IEEE Transactions on Biomedical Engineering, vol. 62, no. 2, pp. 522-531, February 2015, DOI: 10.1109/TBME.2014.2359372
6. Temko, Andriy. "Accurate heart rate monitoring during physical exercises using ppg." IEEE Transactions on Biomedical Engineering 64.9 (2017): 2016-2024.