Arrhythmia Classification of Reduced-Lead Electrocardiograms by Scattering-Recurrent Networks

Philip A Warrick[1,2], Vincent Lostanlen[3], Michael Eickenberg[4], Masun Nabhan Homsi[5], Adrián Campoy Rodriguez[6], Joakim Andén[6]

[1]PeriGen Inc., Montreal, Canada, [2]McGill University, Montreal, Canada, [3]LS2N, CNRS, École Centrale de Nantes, France, [4]Flatiron Institute, New York, NY, [5]Helmholtz Centre for Environmental Research, Leipzig, Germany, [6]KTH Royal Institute of Technology, Stockholm, Sweden.

Objectives: Electrocardiogram (ECG) analysis is the standard of care for the diagnosis of cardiac abnormalities. Automated ECG analysis could aid cardiologists to make more accurate diagnoses, for timely treatment and reduced healthcare costs. Our team "BitScattered" developed an arrhythmia classifier of 12-lead and reduced-lead (6-, 3-, and 2-lead) ECG using the scattering transform (ST) and recurrent networks. We used PhysioNet/CinC Challenge 2021 data that includes normal sinus rhythms and 23 cardiac arrhythmias.

Methods: We analyzed ECG signals on a per-electrode basis with two orders of Morlet wavelet scattering, yielding 76 scattering "paths" at a sampling rate of $15.6\,\mathrm{Hz}$. A depthwise-separable convolution layer combined the scattering coefficients of different electrodes on a per-path basis and then combined the responses of different paths. Two bidirectional long short-term memory (BiLSTM) layers then captured feature trajectories over time. The final dense layer used binary cross-entropy loss during training to support multiple arrhythmia classes. Our decision rule chose any average class probability that exceeded the threshold $p = 0.5$; otherwise the maximum probability class was chosen. We developed the system with TensorFlow and the Kymatio package.

Results: We restricted our training procedure to datasets CPSC, CPSC2 and PTB-XL; indeed, extending to other datasets proved detrimental to the convergence of the validation loss. Ten-fold cross-validation results achieved challenge metrics 0.524 ± 0.0183, 0.486 ± 0.0225, 0.483 ± 0.0280 and 0.449 ± 0.0230 (mean$\pm$standard deviation, in decreasing order of lead count). Memory problems prevented us from generating similar models on the Challenge server. An entry trained only on datasets CPSC and CPSC2 succeeded. Our four-model hidden test scores were 0.218, 0.218, 0.202 and 0.155, respectively.

Conclusions: This architecture shows promising initial results. We will improve training pipeline scale and efficiency. Finally, we have observed promising reduced-lead models incorporating 12-lead information during training using canonical correlation analysis (CCA).