

Deep-learning based 12-lead ECG classification using convolutional neural networks with an attention layer and gated recurrent units

Martin Kropf^{1,2,3}, Aravind-Kumar Radhakrishnan^{1,2,3}, Narayanaraj Dharmalingam⁴

1. Department of Internal Medicine and Cardiology, Charité University Medicine, Campus Virchow-Klinikum
2. German Centre for Cardiovascular Research (DZHK), partner site Berlin, Berlin, Germany
3. Berlin Institute of Health (BIH), Berlin, Germany
4. KIS GmbH & Co. KG, Berlin, Germany

Introduction:

Standard 12 leads ECG is a noninvasive and inexpensive tool, which can identify patients at higher risk for cardiac disorders. Artificial intelligence has proved to be able to identify such disorders. However, for successful application of automated diagnosis algorithms, those algorithms must provide higher or equal reliability than state of the art ECG processing tools. We present a deep-learning based AI algorithm for automated classification of 12-lead ECG recordings, which had been developed in the course of the Computing in Cardiology Challenge 2020.

Methods:

To classify ECG recordings into the nine classes as defined by the challenge we used a convolutional neural network with an attention layer and bidirectional gated recurrent units. Preprocessing of the data included normalization of the signal, i.e. the mean of all signals was removed from each individual signal. Because the duration of the signals in the training set were not equal we filled shorter signals with zeros, so they would end up the same size as the longest signal in the dataset

We split all available ECGs from the training-set as provided for the Challenge to a ratio of 80% for training (=5501 records) and kept a hold-out set of 20% (=1376 records). The classifier only learned on the training-set and the best model (which was the one with the smallest binary-crossentropy-loss on the validation set) was chosen for evaluation on the hidden test set.

Result:

During the unofficial phase we achieved an F_{β} -score of 77.2% and an G_{β} -score of 58.2%, resulting in a geometric mean score of 67.0% on the hidden test dataset. Our best internal geometric mean score on our hold-out dataset was 72.0%.

Discussion:

Next steps will be to train our model on the full dataset using cross-validation, which should improve results compared to our current simple hold-out approach. Further we will employ model stacking, to build a meta-classifier based on several base-learners. Further we will apply weights based on the frequency of samples per class, as the distribution of classes is severely unbalanced.