

Identifying Abnormalities in a 12-Lead Electrocardiogram Image Stack with Convolutional Neural Networks

Jenny Venton*, Ashish Sundar, Ryan Moran, Spencer A Thomas, Peter M Harris, Nadia A S Smith, Merima Causevic, Alen Bosnjakovic, Vedran Karahodzic, Claudia Nagel, Axel Loewe, Nicolas Pilia, Olaf Dössel, Loïc Coquelin, Jane V Lyle, Philip J Aston

National Physical Laboratory
Teddington, UK

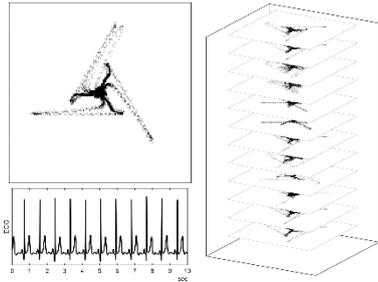
Aim: The aim of this study was to identify the clinical diagnoses present in a 12-lead electrocardiogram (ECG) by transforming the signals into images, then using a convolutional neural network to classify the image stack. We used Symmetric Projection Attractor Reconstruction (SPAR) to generate the 2D image of a single ECG signal.

Method: A 3D stack of twelve 2D SPAR images (one for each lead) was generated for each subject in the training set (see Figure). We used a convolutional neural network classifier with five convolutional layers to extract important features and four dense layers for the classification. To capture the relationships between the twelve ECG signals, the first convolutional layer utilised 3D convolutional filters coupled with 3D pooling filters. A custom loss function based on the F_β and G_β definitions was used to classify each subject into one or more diagnostic classes.

We aim to improve network performance by deriving additional features prior to the classification, including ECG intervals. Further image stacks comprising frequency-based scalogram and spectrogram images derived from the ECG signals will provide further features for classification. Finally, we will use techniques that make use of label dependence and label re-weighting to improve classification of poorer performing labels.

Results: Preliminary 3-fold cross validated training performance results using the SPAR image stack are $F_\beta = 0.47$, $G_\beta = 0.22$ ($\beta = 2$).

Conclusion: Visual representations of ECG signals make it possible to use the power of a convolutional neural network, that is typically used in image classification problems, to classify 12-lead ECG signals. Furthermore, using a 3D filter exploits the relationships between the signals for the different leads.



ECG signal and corresponding attractor (left) and a stack of images for all 12 leads (right).