

Electrogram-based Estimation of Myocardial Conduction Using Deep Neural Networks

Konstantinos Ntagiantas¹, Dimitrios Panagopoulos², Wing M. Poon², Jothi L. Mahendra Kumar²,
Danya Agha-Jaffar², Nicholas S. Peters², Chris D. Cantwell³, Anil A. Bharath^{1*}, Rasheda A.
Chowdhury^{2*}

¹ Department of Bioengineering, Imperial College London, London, United Kingdom

² National Heart and Lung Institute, Imperial College London, London, United Kingdom

³ Department of Aeronautics, Imperial College London, London, United Kingdom

Abstract

Contact electrograms (EGMs) can be used to guide catheter ablation in the treatment of atrial fibrillation. However, our understanding of the link between electrophysiology and the underlying myocardial substrate is limited. We use neural networks and EGMs to estimate the amount of collagen within the field of view of the recording electrodes. EGMs were recorded from rat ventricular slices ($n=15$). Samples were imaged using second harmonic generation (SHG) microscopy, allowing for quantification of collagen. A convolutional neural network (1D-ResNet) was trained to estimate collagen distribution from the recorded EGMs. Each electrogram, recorded for one cycle length, was paired with a collagen index for the corresponding electrode. The total number of samples was 91,239. We successfully estimated collagen index in the testing set, with an absolute error of 0.022 ± 0.024 , and a correlation coefficient of $R=0.81$ between the predicted and true collagen amount. The network identifies the variation in the QRS, as well as cases of fractionation, as useful features for quantifying collagen underneath the electrode. This work provides a framework and proof of concept that location of scar can be predicted from EGMS using neural networks.

1. Introduction

Contact EGMs have been used to help identify ablation targets, exploiting features of the electrogram that can be interpreted by clinicians. An example is fractionation, which is assumed to be indicative of slowly-conducting tissue (scar) [1]. However, there is no consensus as to how the morphology of EGMs correlates to the structural properties of the myocardium, and there have been studies that show that traditionally identified features are not robust markers of structural pathophysiology, and consequently

catheter ablation success rates remain low [2].

In recent studies, machine learning methods have been used to classify myocardial substrate from intracardiac electrograms using hybrid and/or in silico data [3]. However, analysis and characterisation of biological data alongside paired structural data, has been qualitative (classification) rather than quantitative (regression).

2. Methods

2.1. Myocardial slices preparations

All procedures were carried out according to the Animal (Scientific Procedures) Act 1986. Adult female Sprague-Dawley rats (350g-500g) underwent euthanasia by dislocation of the neck, after sedation, and confirmed by cessation of circulation. Hearts were explanted, and the ventricular tissue samples were prepared for the slicing procedure as previously described [4]. The tissue blocks were cut into a 6×6 mm surface area and attached to the specimen holder of the high-precision vibrating microtome (700 smz, Campden Instrument Ltd.UK). The thickness of each slice was set at $300 \mu\text{m}$ which was derived sequentially through the thickness of the ventricular wall.

2.2. EP data collection

The electrophysiological properties of the ventricular slices were measured through the multielectrode array (MEA) system (USB-MEA60-Inv, Multi-Channel Systems, Reutlingen, Germany). The MEA system used in this experiment consists of 60 gold electrodes arranged in an 8 matrix with inter-electrode distances of $700 \mu\text{m}$ and $100 \mu\text{m}$ electrode diameter. Stimulation was carried out using an STG stimulus generator programmed by MC Stimulus II software (version 3.4.4, Multi-Channel Systems). A biphasic stimulus of 2 ms in duration and be-

tween 1000 – 2500 mV in amplitude (120% above the threshold of activation) was applied for each 10 s recording with a pacing rate of 1 Hz, 1.33 Hz, 2 Hz, 3 Hz, or 4 Hz applied at all electrodes from one of the outer rows of the MEA grid. The slices were placed in the middle of the MEA, covering all electrodes and oxygenated Tyrode’s solution was superfused. The temperature of the solution was maintained at 37 °C. Slice positions on the MEA plate were imaged using a digital camera to record electrode locations.

2.3. EGM preprocessing

EGM recordings were imported into Python using Multi-Channel Systems MCDDataTool software, and processed with custom scripts. Each recording was partitioned into stimulation events which we refer to as activations. Activations of each experiment were z-score normalized, and the median signal of the 36 electrodes across each activation was subtracted. This ensured that the baseline and the stimulus artefact were removed.

2.4. Structural data collection

After electrophysiological recordings, slices were fixed with paraformaldehyde in preparation for imaging with a Leica SP5 inverted confocal microscope, and a multiphoton laser Newport Spectra-Physics Two Photon Laser – Mai Tai DeepSee was used for the second harmonic generation (SHG). SHG imaging was done at a wavelength of 845 nm [5]. To extract the collagen distribution of the slice and incorporate the different distributions from different slices, we processed the samples with an ImageJ pipeline, and labelled each pixel as collagen or not collagen based on a threshold.

After the processing of the microscopy images, the positions of the MEA electrodes were superimposed. We used a custom method based on landmarks to register the slice in the microscopy image and the image of electrode positions on top of the MEA. Each electrode was annotated with the proportion of pixels labelled as collagen within its field of view (FoV); we called this number the *collagen index*. The electrode field of view was set to be 350 μm , equal to half of the interelectrode distance of the MEA system.

2.5. Deep learning model

A one-dimensional ResNet architecture is used [6, 7]. First, the sample passes through a convolutional layer followed by batch normalization and ReLU activation. This is followed by four residual blocks, and finally three linear layers. All convolutional layers, including those of the residual blocks, use a kernel of size 17 and stride 2, except for the convolution on the identity branch of the residual

block, which uses a kernel of size 1 and stride 1, and the initial convolution of the network, which uses a kernel of size 17 and stride 1. Batch normalization and the ReLU activation function are used after every convolutional layer. The dimensions of the feature spaces after each layer are shown in Table 1.

Table 1. Output feature space sizes for the 1D ResNet architecture, given an input of 1×4096 .

Name	Output size
conv1	16×4096
resblock1	32×1024
resblock2	64×256
resblock3	128×64
resblock4	256×16
linear1	1×256
linear2	1×128
linear3	1×1

3. Results

3.1. Model training

We trained the model on the dataset created from 11 slices across 4 different rat hearts. One sample consisted of a single activation from one electrode and the corresponding collagen index for the electrode, as described in sections 2.2 and 2.4. To improve model robustness, we only selected samples whose EP mean wavelet coefficient was above 0.23. The threshold was chosen after looking at the mean of the wavelet transform of valid and invalid samples. Samples which were either corrupted with noise or lacked any signal, due to poor contact between the electrode and the slice, were deemed as invalid. We considered the first 165 ms of every sample, as the subsequent information was determined as baseline only and there was no useful information except for noise after that time. After this process, the final dataset consisted of 91,239 samples, from 396 locations. The distribution of collagen index in the dataset is shown in Figure 1.

We applied 10-fold cross-validation during training; for each fold 10% of the data was used as the testing set, and we trained on the remaining 90%. The train/test partitioning was done on the electrodes and not the individual samples, to ensure that the distribution of the testing data was unseen by the model. Training lasted for 20 epochs, with the Adam optimizer and the root mean square error (RMSE) as the loss function.

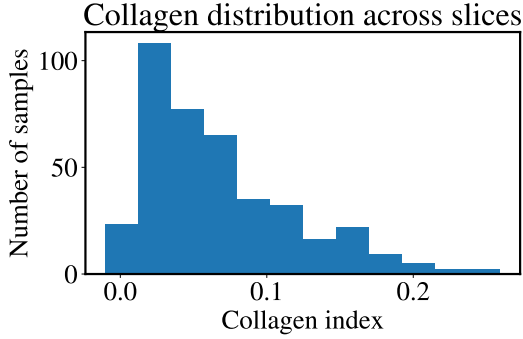


Figure 1. Distribution of the collagen index of the dataset. A value of 1 represents complete collagen coverage, while a value of 0 means that there is no collagen, in the field of view of the electrode.




Figure 2. Median predictions of the testing set across all folds of cross-validation. **Inset:** Distribution of interquartile range and standard deviation for the predictions of each electrode.

3.2. Performance

The model was evaluated on the testing set which consisted of 9,576 samples from 40 electrodes. All the metrics reported are averages across testing electrodes and across the ten cross validation folds. The RMSE in the testing set was 0.035. Comparatively, the distribution of collagen index had a standard deviation of 0.050, considerably larger than the RMSE of our predictions. The mean absolute error $|y_{true} - y_{predicted}| = 0.022 \pm 0.024$ also suggests that the model generalized well.

Due to the nature of the data, there are multiple activations for the same electrode, so two important metrics are:

1) the error in the median prediction across all samples for one particular electrode, and 2) the average standard deviation, sd, and the average interquartile range, IQR, of the predictions for one electrode. The median predictions per electrode are shown in Figure 2. The model captures the trend of the data, with a Pearson correlation coefficient $R_{y,\hat{y}} = 0.81$.

3.3. Interpretation

We investigated the response of the model to the morphology of the EGMs. After the first convolution of the trained network, only the parts of the EGM that are useful for the prediction are considered (highlighted) by the remainder of the network. We visualized which parts of EGMs are highlighted by the first convolution, and used in the downstream 16-dimensional feature space. The network highlights only the parts of the EGM during activation of the tissue, and the features are correlated with the slope of depolarization and repolarization, the peak-to-peak voltage and fractionation. An example of this is shown in Figure 3, for one of the 16 dimensions of the feature space.




Figure 3. EGMs from high-collagen (upper panel) and low-collagen (lower panel) locations, and the highlighted segments in one of the feature space dimensions of the convolutional network. The estimates of the collagen index are equal to the ground truth in these examples. In this dimension, the negative part of the downslope is the most important characteristic for the prediction.

4. Discussion

In this *ex vivo* slice model we accurately predicted the amount of collagen within a field of view of an electrode in regions of cardiac tissue previously unseen by the model using the contact electrogram. Physiologically this is analogous to estimating the degree of fibrosis, and the effective conductivity of the tissue. By using deep convolutional networks for this task, we utilized all of the available EGM signal without bias to regions of interest. An alternative option would be to use specific features, which are automatically either manually selected or generated by signal processing pipelines. By leveraging deep learning models and training on the raw recordings, morphological features were selected by the network rather than based on previous clinical assumptions. Furthermore, instead of limiting our work to the training and evaluation of the model, we sought to interpret the inner mechanics of the neural network by highlighting those morphological parts of the EGMs that activate the first layer of the network, thereby giving insight into the areas of importance. This practice only examines the very first part of the model, so it does not fully provide mechanistic explainability, but it offers useful insight and we confirmed that the features highlighted (such as voltage magnitude and depolarization slope, as shown in 3) have clinical relevance. The network also generalizes well to samples from unseen myocardium.

We observed that the predictions are less accurate as the collagen index increases (see Figure 2). This is expected since there is only a small number of samples with a low collagen index, as shown in Figure 1. The distribution of *sd* and *iqr* confirms that predictions are within a reasonable distance of the ground truth for the majority of the samples.

One of the most crucial steps of the work presented, is the localization of the electrodes on the tissue. The hypothesis that the EGM signal encodes information about the conductivity of the tissue requires that our estimated position is at least within the FoV of the true electrode position. However, there is no firm consensus on the size of the FoV of an electrode. We used a model-based approach and, given the dimensions of the electrodes, constructed the resulting transfer function that appropriately weights action potentials to produce EGMs [8]. As mentioned in section 2.4, we specified the radius of the field of view to be 350 μm . This is the radius where the weights of the action potentials are equal to half of the action potential in the location of the electrode.

5. Conclusion

We present a method of quantifying the amount of collagen within the field of view of an electrode on a MEA system, thus confirming our hypothesis that signals from

unipolar EGMs can be used to accurately predict structural information of the substrate. By using deep learning methods to solve this problem, the model was trained using all of the available EGM signal, and the important regions, which most strongly influence the prediction of the network, were highlighted, providing insight into the mechanics of the network and offering clinical explainability. This work can be extended with clinical data from high density mapping catheters, and may provide clinicians with a tool to accurately estimate the local degree of fibrosis in myocardial tissue.

Acknowledgments

This work was completed was supported by Wellcome Trust [Grant reference number: 2228, 45/Z/21/Z/Rosetrees Trust Grant M577 and BHF grant FS/4yPhD/F/21/34165.

References

- [1] Nademane K, Lockwood E, Oketani N et al. Catheter ablation of atrial fibrillation guided by complex fractionated atrial electrogram mapping of atrial fibrillation substrate. *Journal of Cardiology* 2010;55(1):1–12.
- [2] Verma A, Jiang Cy, Betts TR et al. Approaches to catheter ablation for persistent atrial fibrillation. *New England Journal of Medicine* 2015;372(19):1812–1822.
- [3] Sánchez J, Luongo G, Nothstein M et al. Using Machine Learning to Characterize Atrial Fibrotic Substrate From Intracardiac Signals With a Hybrid *in silico* and *in vivo* Dataset. *Frontiers in Physiology* 2021;12(July):1–15.
- [4] Chowdhury RA, Tzortzis KN, Dupont E et al. Concurrent micro- to macro-cardiac electrophysiology in myocyte cultures and human heart slices. *Scientific Reports* May 2018; 8(1):6947.
- [5] Tjin G, Xu P, Kable SH et al. Quantification of collagen I in airway tissues using second harmonic generation. *Journal of Biomedical Optics* mar 2014;19(3):036005.
- [6] He K, Zhang X, Ren S et al. Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 2016;2016-Decem:770–778. ISSN 10636919.
- [7] Ribeiro AH, Ribeiro MH, Paixão GM et al. Automatic diagnosis of the 12-lead ECG using a deep neural network. *Nature Communications* 2020;11(1):1–9.
- [8] Abdi B, van Schie MS, de Groot NM et al. Analyzing the effect of electrode size on electrogram and activation map properties. *Computers in Biology and Medicine* 2021; 134:104467.

Address for correspondence:

Anil A. Bharath

Royal School of Mines, Department of Bioengineering, Imperial College London, London SW7 2AZ, United Kingdom
a.bharath@imperial.ac.uk