Intracardiac electrical imaging using the 12-lead ECG: a machine learning approach using synthetic data

Mikel Landajuela[1], Rushil Anirudh[1], Joe Loscazo[2], Robert Blake[1]

[1]Lawrence Livermore National Laboratory, Livermore, California, EEUU
[2]Harvard Medical School, Boston, Massachusetts, EEUU

Abstract

Current state-of-the-art techniques for non-invasive imaging of cardiac electrical phenomena require voltage recordings from dozens of different torso locations and anatomical models built from expensive medical diagnostic imaging procedures. This study aimed to assess if recent machine learning advances could alternatively reconstruct electroanatomical maps at clinically relevant resolutions using only the standard 12-lead electrocardiogram (ECG) as input. To that end, a computational study was conducted to generate a dataset of over 16000 detailed cardiac simulations, which was then used to train neural network (NN) architectures designed to exploit both spatial and temporal correlations in the ECG signal. Analysis over a validation set showed average errors in activation map reconstruction below 1.7 msec over 75 intracardiac locations. Furthermore, phenotypical patterns of activation and the morphology of the activation potential were correctly reconstructed. The approach offers opportunities to stratify patients non-invasively, both retrospectively and prospectively, using metrics otherwise only available through invasive clinical procedures.

1. Introduction

The cardiac ECG plays an important role in diagnosing ventricular arrhythmias because it is non-invasive and cost-effective while still able of distinguishing a wide variety of diseases such as ventricular myocardial infarction and bundle branch blocks. The process of building a map of cardiac electrical activity from torso electrical recordings has been well studied in the literature as the cardiac inverse problem. Traditionally, this problem has been approached by introducing additional torso electrodes and building a detailed patient-specific geometrical model based on non-invasive diagnostic imaging. Numerically, the problem is notoriously ill posed and requires additional regularization to yield useful results. Recent methods using machine learning have bypassed the need for patient-specific geometries, but they still make use of dozens of torso electrodes [1] and/or are restricted to recordings on the epicardium [2].

This work presents an exploratory study on whether a machine learned algorithm can reconstruct ventricular transmural activation maps and full transmembrane voltage evolutions using only the 12-lead ECG as input. The approach is motived by results in the machine learning literature that show reconstruction of high-dimensional information from low-dimensional input streams, provided enough data is available. The core idea is to leverage decades of advances in first-principles modeling of cardiac electrophysiology by training NNs over thousands of detailed simulations. In this work, two imaging problems are considered: the reconstruction of the time activation map and the reconstruction of the full transmembrane voltage evolution. The resulting reconstruction systems are computationally efficient and can provide information of the electrical behavior inside the patient's cardiac ventricles, with a resolution equivalent to an AHA17 segment activation map. The proposed framework can increasingly improve its accuracy by assimilating new simulated and/or experimental data using transfer learning or physics-informed regularization techniques.

2. Methods

A comprehensive dataset of computational experiments was created in this work. Each experiment consists of intracardiac transmembrane voltage recordings and ECG signal pairs. The dataset was subsequently used to train machine learning systems on inverse reconstruction tasks.

2.1. Dataset of cardiac simulations

Cardiac simulations were carried out using Cardioid [3], a multiscale cardiac simulation code developed at Lawrence Livermore National Laboratory (LLNL). Simulations were performed over real bi-ventricular cardiac geometries obtained from the publicly available dataset [4] and resolved to a 200 μm resolution (see Figure 1). My-

ocardial fiber orientations were assigned based on a rule-based laplacian-driven algorithm for interpolating fiber geometries in the absence of DTMRI data. The high resolution simulations of the transmembrane voltages within the myocardium were used to compute the synthetic ECG signals via the pseudo-ECG method [5]. The locations of the pseudo-ECG electrodes were chosen based on locations derived from an existing torso mesh and then normalized to a 100 mm radius around the center of each mesh (see Figure 2b). Activation patterns were extracted from the literature and the cell models [6] included apex-to-base and transmural action potential duration (APD) heterogeneity.

For the recording of the transmembrane voltages within the myocardium, 30 points were selected by hand for each mesh — 17 endocardial points were selected in the left ventricle (LV), corresponding to standard AHA17 segment locations, and 13 points were selected in the right ventricle (RV). See Figure 2a for a schematic representation of the points over a Bull's-eye display of the heart. From these 30 points, 20 exterior wall points were programmatically identified based on minimum distances from the hand-selected endocardial points, and 25 mid-myocardial points were then found through interpolation. For each simulation performed, simulated transmembrane voltages were recorded for each of the 75 epicardial/midmyocardial/endocardial points. These transmembrane voltages were paired with the ECGs collected above for use in the machine learning classifier.

A wide range of physiological and pathophysiological parameters was considered, including variations in tissue conductivities, maximal conductance G_{Kr} of the rapid delayed rectifier current (0%, or blocked, and 50% with respect the original value in [6]), and basic cycle lengths of (600 ms and 1000 ms). All simulations were performed for 500 ms of simulation time with 200 μm resolution meshes and a time-step of 5 μs. The ECG and transmembrane voltages were recorded at a resolution of 1 ms.

In total, 16140 organ-level simulations were conducted in the course of this work. Simulations were performed at LLNL's *Lassen* supercomputer, concurrently utilizing 4 GPUs and 40 CPU cores. With the publication of this paper, the dataset is now made publicly available at `https://library.ucsd.edu/dc/object/bb29449106` (which should be cited as [7]). In that repository, the reader can find further details on the simulation settings, the mathematical models, the anatomical geometries and the parameter variations.

2.2. Deep intracardiac electroimaging

The simulation study described above produced pairs of 12-by-500×1ms ECG signals and 75-by-500×1ms transmembrane voltage signals. For the sake of notation, those signals are represented as matrices $X \in \mathbb{R}^{12\times 500}$ and

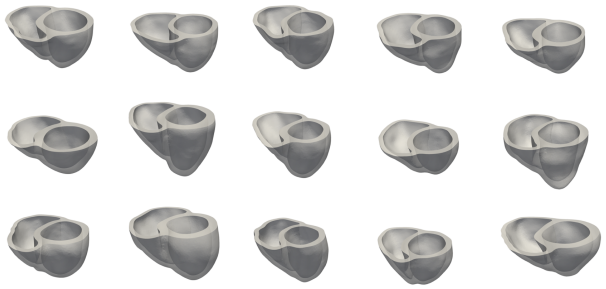

Figure 1: Bi-ventricular cardiac geometries.

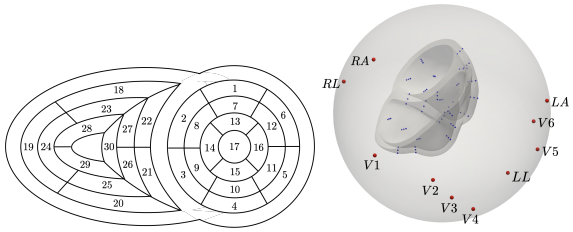

(a) (b)

Figure 2: Recording points. (2a) Location of 30 manually selected endocardial points. (2b) Location of the pECG electrodes (red) and transmembrane voltage recording points (blue).

$V \in \mathbb{R}^{75\times 500}$, respectively. The activation time vector $A \in \mathbb{R}^{75}$, corresponding to the initial activation time at each myocardial recording location, is defined as $A_i = \min_j V_{ij} > 0$. Two machine learning tasks were considered in this work: Task I (activation map reconstruction), which involved reconstructing $A \in \mathbb{R}^{75}$ from $X \in \mathbb{R}^{12\times 500}$, and Task II (transmembrane potential reconstruction), which involves reconstructing $V \in \mathbb{R}^{75\times 500}$ given $X \in \mathbb{R}^{12\times 500}$.

These tasks can be regarded as sequence-to-sequence prediction problems, where the goal is to transform a 500-length sequence of 12 dimensional vectors into a sequence of 75 dimensional vectors. For classification and compression of ECG signals using NNs, researchers have used a variety of architectures including 1D and 2D CNNs [8], and hybrid approaches combining CNNs and LSTM units [9]. For reconstruction tasks of heart surface potentials, the work in [2] uses a time-delayed NN to map the real recorded first lead of the ECG to the unipolar surface potential at the right ventricular apex. In this work, where reconstruction was considered over 75 intracardiac positions, the best results were achieved using 1D CNN architectures inspired by the SqueezeNet model [10]. Two different networks were considered:

- **Network I** (for Task I) : Network I was constructed using SqueezeNet (with 1 dimensional kernels) with a stride

of size 2 in the first convolutional layer and max pooling layers to progressively reduce the temporal dimension. Additional convolutional layers were added at the end to reduce the output dimension to $\mathbb{R}^{75}$. The total number of parameters in the network is 486,657.

• **Network II** (for Task II) : Network II was constructed using SqueezeNet (with 1 dimensional kernels). Additional convolutional layers were added at the end to produce outputs of dimension $\mathbb{R}^{75\times500}$. The total number of parameters in the network is 392,907.

The considered network architectures allow for both temporal and spatial information derived from the ECG signal to be combined and reorganized in a nonlinear way. For training the networks, each $X \in \mathbb{R}^{12\times500}$ tensor was normalized so that $\max_j(X_{ij}) - \min_j(X_{ij}) = 1, \forall i \in \{1,\ldots,12\}$. To train Network II, each $V \in \mathbb{R}^{75\times500}$ was normalized so that the value range was $[0,1]$. The dataset was randomly split into training and validation subsets containing 95% and 5% of the samples, respectively. Learning was performed over one GPU at LLNL's *Lassen* supercomputer, using the PyTorch library with Adam optimizer, mean squared error as loss, batch of size 32 and learning rate of magnitude 0.001.

3. Results

Figure 3 shows an example of a simulated ECG (Figure 3a) and activation map (Figure 3b) in the validation set. The reconstruction obtained by Network I is shown in Figure 3d. On average, the network incurs an error of 1.66 msec over all 75 recording points in the validation set with a mean standard deviation of 1.49 msec. These results show that Network I is able to reconstruct the activation map over the validation set of simulated data. In particular, the algorithm is able to capture and reproduce both septal and transmural activation times in cardiac tissue.

Figure 4 shows the reconstruction results obtained with Network II for the validation ECG in Figure 3a. Recall that this network produces as output the whole tensor $V \in \mathbb{R}^{75\times500}$. Figures 4a-4c show examples of reconstructed transmembrane voltages compared with reference simulated results at myocardial recording points 1, 17 and 67, which are shown in Figure 4e. The corresponding activation time vector $A \in \mathbb{R}^{75}$ is computed from $V \in \mathbb{R}^{75\times500}$ and plotted in Figure 4e.

The derived activation times are slightly slower compared to ground truth. Specifically, the mean error in the reconstruction of activation times using Network II for all points in the validation test is 6.5 msec with a mean standard deviation of 6.65 msec. The error is higher than in Network I, and is not surprising since in this case the reconstruction is not targeting the activation map itself but rather the whole temporal evolution of the transmembrane voltage within the myocardium.

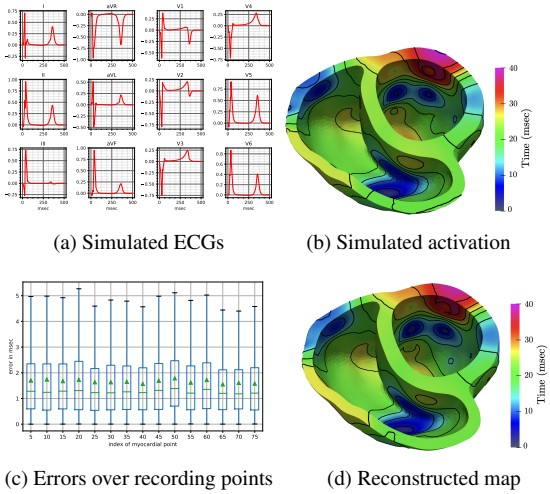

(a) Simulated ECGs (b) Simulated activation

(c) Errors over recording points (d) Reconstructed map

Figure 3: Activation time reconstruction using Network I. All times in ms, and all voltages in mV. Isochrones added every 10 ms. (3a) ECG signals used to reconstruct the activation times in (3d). (3b) Ground-truth activation times reconstructed from the 75 voltage samples. (3c) Box-and-whisker plot of the errors in activation times for a selection of 15 recording points over all the validation set. (3d) Predicted activation times from the ECG of this simulation. Result correspond to a validation case.

Figures 4a-4c show that the transmembrane voltage reconstruction obtained with Network II contains non-physical oscillations when compared to the simulated data. However, the overall amplitude and APD of the curves are well captured by the network (Pearson correlation coefficient is 0.9745 over all recording points in the validation set with a mean standard deviation of 0.0237.) The error in APD is 9.7892 msec over all recording points in the validation set with a mean standard deviation of 10.71 msec.

These results show that Network II it is able to capture the gross phenotypical patterns of activation, the morphology of the activation potential including the APD and the complete dynamical evolution of the depolarization and repolarization phases of the cardiac cycle.

4. Conclusions

In this work, an approach combining synthetic data and machine learning has been proposed for reconstruction problems in intracardiac electrical imaging from 12-lead ECGs. The proposed NN architectures exploit temporal correspondences in torso voltage recordings, thus mimicking the ability of clinicians to form diagnosis by looking at the time-context of the ECG and perform pattern matching on the shape of ECG waveforms. The approach has been shown to reproduce endocardial, epicardial, and mid-myocardial activation maps and transmembrane voltages

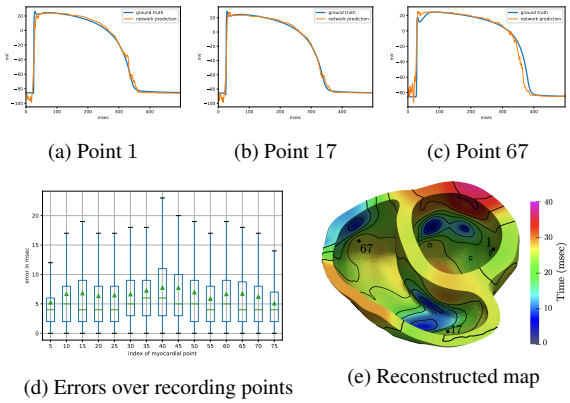

(a) Point 1 (b) Point 17 (c) Point 67

(d) Errors over recording points (e) Reconstructed map

Figure 4: Transmembrane voltage reconstruction using Network II. All times in ms, and all voltages in mV. Isochrones added every 10 ms. (4a-4c) Comparison of the ground-truth transmembrane voltage and the predicted transmembrane voltage. (4d) Box-and-whisker plot of the errors in activation times for a selection of 15 recording points over all the validation set. (4e) Reconstructed map with Network II.

when applied to a simulated dataset.

It remains to be studied the generalization to real clinical data. It is unknown what role abnormal patient geometries, myocardial infarction, or other clinical etiologies would play if this approach was applied in a clinical setting. Additional patient information beyond the bare 12-lead electrocardiogram (e.g., atrial data or additional synthetic ECGs) and physics-informed regularization might be required to build a truly robust clinical machine learning algorithm.

In conclusion, the proposed study is a necessary step towards an electrical cardiac imaging solution that could enable non-invasive stratification of patients based on metrics otherwise only available through invasive electroanatomical mapping studies. The resulting prediction tools do not require any special equipment, work with small number of torso electrodes and can be stored and deployed in devices with low memory and processing capabilities. The designed neural networks can also be readily used as a data-augmentation technique for downstream ECG-based machine learning algorithms.

Acknowledgments

We thank Livermore Computing and the Laboratory Directed Research and Development program at LLNL under the National Nuclear Security Administration (NNSA) for their support of this work through project number 18-LW-078. This work was performed under the auspices of the U.S. Department of Energy by LLNL under Contract DE-AC52-07NA27344 with release number LLNL-JRNL-799842.

References

[1] Karoui A, Bendahmane M, Zemzemi N. A spatial adaptation of the time delay neural network for solving ecgi inverse problem. In International Conference on Functional Imaging and Modeling of the Heart. Springer, 2019; 94–102.

[2] Malik A, Peng T, Trew ML. A machine learning approach to reconstruction of heart surface potentials from body surface potentials. In 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, 2018; 4828–4831.

[3] Richards DF, Glosli JN, Draeger EW, Mirin AA, Chan B, Fattebert JI, Krauss WD, Oppelstrup T, Butler CJ, Gunnels JA, et al. Towards real-time simulation of cardiac electrophysiology in a human heart at high resolution. Computer methods in biomechanics and biomedical engineering 2013;16(7):802–805.

[4] Duchateau N, Sermesant M, Delingette H, Ayache N. Model-based generation of large databases of cardiac images: synthesis of pathological cine mr sequences from real healthy cases. IEEE transactions on medical imaging 2017; 37(3):755–766.

[5] Plonsey R, Barr RC. Bioelectricity: a quantitative approach. Springer Science & Business Media, 2007.

[6] Ten Tusscher KH, Panfilov AV. Alternans and spiral breakup in a human ventricular tissue model. American Journal of Physiology Heart and Circulatory Physiology 2006;291(3):H1088–H1100.

[7] Landajuela M, Anirudh R, Blake R. Dataset of simulated intracardiac transmembrane voltage recordings and ecg signals. In Lawrence Livermore National Laboratory (LLNL) Open Data Initiative. UC San Diego Library Digital Collections, 2022; URL `https://library.ucsd.edu/dc/object/bb29449106`.

[8] Cantwell CD, Mohamied Y, Tzortzis KN, Garasto S, Houston C, Chowdhury RA, Ng FS, Bharath AA, Peters NS. Rethinking multiscale cardiac electrophysiology with machine learning and predictive modelling. Computers in biology and medicine 2019;104:339–351.

[9] Hong S, Wu M, Zhou Y, Wang Q, Shang J, Li H, Xie J. Encase: An ensemble classifier for ecg classification using expert features and deep neural networks. In 2017 Computing in Cardiology (CinC). IEEE, 2017; 1–4.

[10] Iandola FN, Han S, Moskewicz MW, Ashraf K, Dally WJ, Keutzer K. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and $<$ 0.5 mb model size. arXiv preprint arXiv160207360 2016;.

Address for correspondence:

Mikel Landajuela
7000 East Ave, Livermore, CA 94550
landajuelala1@llnl.gov