

Siamese Neural Networks for Small Dataset Classification of Electrograms

Bram Hunt, Eugene Kwan, Derek Dossdall, Rob S. MacLeod, Ravi Ranjan

University of Utah, Salt Lake City, Utah, United States

Abstract

Objective: In this work, we aimed to isolate endocardial electrograms with highly distinguishable activations in a small atrial fibrillation (AF) dataset by leveraging Siamese neural networks.

Methods: Unipolar endocardial electrograms were captured with a basket catheter from a paced canine model of persistent AF, and 1,006 unique electrograms were randomly sampled from the endocardial wall and we isolated those with highly distinguishable activations. We trained Siamese neural networks to compare pairs of samples and then to classify electrograms in the testing dataset.

Results: Using a reference from the validation dataset, the Siamese neural networks achieved a weighted accuracy of 90.3% and an F1 score of 0.94 on the classification testing dataset.

Conclusion: Even in an electrogram dataset with significant size constraints, we achieved high accuracy and improved over a conventional neural network classifier weighted accuracy of 71.4%.

1. Introduction

Atrial fibrillation (AF) is an arrhythmia originating from the upper chambers of the heart and presents as rapid and uncoordinated contractions of the atria. AF has been linked to the existence of localized disruptive phenomena implicated in its maintenance and progression.^[1-4] These electrical phenomena are known as drivers and have been used as a basis for treatment. Driving mechanisms are identified and confirmed through activation mapping, but the results of mapping are highly sensitive to measurement modality and choice of activation times.^[5-7] As such, the choices of activation timings and subsequent mechanisms behind AF are often ambiguous. Factors which obscure definitive mapping analysis include far-field effects, intricate AF wavefront propagation mechanisms, tissue-level effects like conduction heterogeneity from fibrosis or scar, catheter and electrode design decisions, and other noise generators.^[7-9] However, over the course of endocardial mapping, some signals during AF contain less ambiguous activation timings relative to other signals. Examples of electrograms with highly distinguishable and poorly distinguishable activations are shown in Figure 1.

An algorithm for isolating signals with highly distinguishable activations would allow creation of activation maps with less ambiguity in mechanism determination. In this study we investigated the feasibility of using a neural network-based method of identifying endocardial electrograms with highly distinguishable activations for the purpose of eliminating the ambiguity associated with activation mapping during AF.




Figure 1. Examples of 2-second unipolar electrograms with highly distinguishable activations (top) and with poorly distinguishable activations (bottom).

Conventional algorithms have difficulty classifying signals based on morphological features, a difficulty which is amplified by irregular signals such as those during AF. Deep neural networks have been shown to improve over conventional algorithms in image-based tasks and as such are promising for use in classifying electrograms based on morphology.^[10,11] Traditional classifications algorithms accept an input image and produce a classification vector. However, this typically comes at the expense of requiring extensive datasets and can be complicated by the use of datasets that have imbalanced class composition.

Siamese neural networks have been used for one-shot learning and small dataset classifications.^[12] These networks are trained to accept images from two classes and then produce a similarity metric assessing whether input images come from the same or different classes. For scoring, a score of 0 is dissimilar and a score of 1 is similar. After training, the Siamese neural networks can be tested by supplying a reference image for a target class and then iterating over the testing dataset. We sought to lever the

advantages of a Siamese neural network training schema to develop a classifier for isolating electrograms with highly distinguishable electrograms. For comparison, we also trained a convolutional neural network on a traditional electrogram-to-class training schema with an equivalent number of parameters.

2. Methods

We examined 1,006 non-contact endocardial unipolar electrograms of 2 seconds from a paced canine model of AF (n=6, female mongrel, 29±2 kg, 4 to 24 wks. pacing) during sustained AF with an Orion 64-electrode basket catheter and the Rhythmia mapping system (Boston Scientific). Electrograms were randomly sampled from both the left and right atria. We then manually identified 86 electrograms (8.5% of total) as having highly distinguishable activations. Classification criteria included low fractionation, the presence of a dominant single-potential morphology, and having few or no unclear activation timings.

For network training, we utilized a 75%/10%/15% training/validation/testing split. Splits were balanced such that classes were equally represented in each dataset. For the architecture of the Siamese networks, we adapted the Koch 2015 model for 1D inputs.^[12] Figure 2 shows the full model architecture. The network used for comparison was a single half of the Siamese networks, consisting of the four convolutional layers fed directly into the final linear layer. The model used a total of 48,145 parameters with each half of the twin networks having identical parameterization to the other.

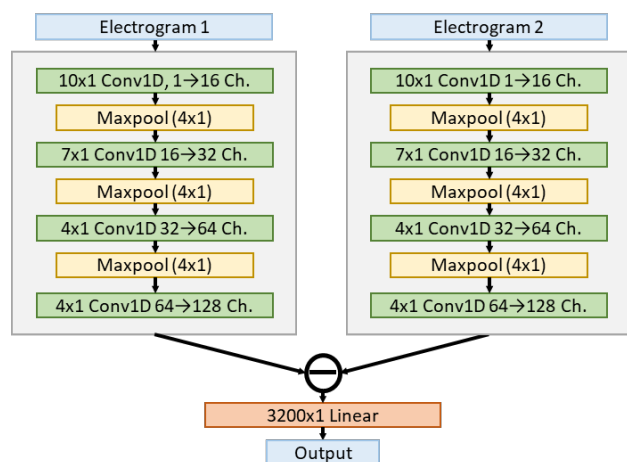


Figure 2. The architecture of the Siamese neural network model adapted for 1D signals from Koch (2015).^[12]

A comparison task and a direct classification task were both used in training. For the comparison task, two electrogram samples are randomly sampled from a dataset and networks are trained to determine if they originate

from the same class. For the direct classification task, networks are trained to directly determine the sample class. The Siamese neural networks model was trained on the comparison task and evaluated on the direct classification task, whereas the comparison network was trained and evaluated on the direct classification task only. We used binary cross entropy loss as an objective function for training and validation in both tasks. To adapt the Siamese networks for the direct classification task, we evaluated all electrograms in the validation dataset on the comparison task and selected the single sample with the highest accuracy when compared to other electrograms. That sample was fed into one of the two Siamese neural networks during testing, effectively serving as a reference electrogram. Testing electrograms were input to the other network for each example in the testing dataset.

The models and training algorithms were implemented in PyTorch, a Python 3 deep learning library.^[13] Training was performed on a computer with a single GeForce RTX 2060 graphics card. The Adam optimization algorithm was used with an initial learning rate of 1e-4 and batch size of 200.^[14] Loss weighting was used to address class imbalances. For the comparison task, we weighted the loss according to the inverse probability of having two electrograms come from different classes (weight of 6.67). For the direct classification task, the loss was weighted according to the inverse frequency of each class in the dataset (weight of 11.7 for non-highly distinguishable electrogram class loss).

3. Results

We randomly sampled pairs of electrograms from the training dataset, input the electrograms into the networks to determine similarity and loss, and trained until validation accuracy plateaued after 3 epochs. For the training, validation, and testing comparison tasks, we achieved weighted accuracies of 86.2%, 84.2%, and 85.4% respectively. Training and validation curves for loss and accuracy are shown in Figure 3.




Figure 3. Loss and accuracy for validation and training

during Siamese model training.

For the validation comparison task, the optimal reference electrogram is shown in Figure 4. This electrogram achieved a validation accuracy of 92.3% when evaluated in the comparison task against all validation electrograms.




Figure 4. Electrogram with highest validation weighted accuracy score. This electrogram was classified as having non-highly distinguishable activations due to the presence of activations with unclear timings.

We used the above reference electrogram on the classification task for the testing dataset with the Siamese neural networks and achieved a final weighted accuracy of 90.3%, an F1 score of 0.94, recall of 0.94, and precision of 0.95. This significantly improved over the testing accuracy of 71.4% of the comparison network that was only trained on the direct classification task.

4. Conclusions

We used Siamese neural networks to achieve clear improvement over a random prediction weighted accuracy of 50% and a conventional network accuracy of 71.4% in this electrogram classification task. This successfully demonstrated the feasibility of using this method of isolation on an AF dataset and lays a path to future work exploring its use in driver identification.

A limitation for this method of isolating electrograms with highly distinguishable activations is that electrograms with poorly distinguishable activations may contain important information for the identification or function of drivers. Activation mapping may need to be applied with other markers such as dominant frequency or peak-to-peak voltage to allow confident assessment of mechanisms. Additionally, it is unclear if the method would identify viable electrograms for all locations in the atria; there is potential that certain sites, such as those with low peak-to-peak voltage or highly fibrotic tissue, may not produce high quality signals over the course of a mapping session.

In addition to being potentially useful for driver identification, this methodology is promising for future electrogram morphology classification where target signals are rare, dataset size is small, or class imbalances are large. This technique can further be used in therapies which target regions based on electrogram features. For future work, non-AI algorithms may improve over this

method in terms of performance. Other methods of addressing small datasets in deep learning could be used on this task. For example, self-supervised learning on a related task, such as encoding electrograms into a latent space and decoding them from that space, could allow the use of a network with a greater number of parameters which could improve accuracy.^[15] Other transfer learning tasks such as image inpainting for the 1D signals could also be used to improve performance.^[16]

References

- [1] Konings KT, Kirchhof CJ, Smeets JR, Wellens HJ, Penn OC, Allessie MA. High-density mapping of electrically induced atrial fibrillation in humans. *Circulation*. 1994 Apr;89(4):1665-80.
- [2] Jalife J, Berenfeld O, Mansour M. Mother rotors and fibrillatory conduction: a mechanism of atrial fibrillation. *Cardiovascular research*. 2002 May 1;54(2):204-16.
- [3] Haïssaguerre M, Hocini M, Sanders P, Takahashi Y, Rotter M, Sacher F, Rostock T, Hsu LF, Jonsson A, O'Neill MD, Bordachar P. Localized sources maintaining atrial fibrillation organized by prior ablation. *Circulation*. 2006 Feb 7;113(5):616-25.
- [4] Narayan SM, Wright M, Derval N, Jadidi A, Forclaz A, Nault I, Miyazaki S, Sacher F, Bordachar P, Clémenty J, Jaïs P. Classifying fractionated electrograms in human atrial fibrillation using monophasic action potentials and activation mapping: evidence for localized drivers, rate acceleration, and nonlocal signal etiologies. *Heart Rhythm*. 2011 Feb 1;8(2):244-53.
- [5] Allessie MA, de Groot NM, Houben RP, Schotten U, Boersma E, Smeets JL, Crijns HJ. Electropathological substrate of long-standing persistent atrial fibrillation in patients with structural heart disease: longitudinal dissociation. *Circulation: Arrhythmia and Electrophysiology*. 2010 Dec;3(6):606-15.
- [6] Pandit SV, Jalife J. Rotors and the dynamics of cardiac fibrillation. *Circulation research*. 2013 Mar 1;112(5):849-62.
- [7] Costabal FS, Zaman JA, Kuhl E, Narayan SM. Interpreting activation mapping of atrial fibrillation: a hybrid computational/physiological study. *Annals of biomedical engineering*. 2018 Feb;46(2):257-69.
- [8] Rostock T, Rotter M, Sanders P, Takahashi Y, Jaïs P, Hocini M, Hsu LF, Sacher F, Clémenty J, Haïssaguerre M. High-density activation mapping of fractionated electrograms in the atria of patients with paroxysmal atrial fibrillation. *Heart Rhythm*. 2006 Jan 1;3(1):27-34.
- [9] Correa de Sa DD, Thompson N, Stinnett-Donnelly J, Znojkwicz P, Habel N, Müller JG, Bates JH, Buzas JS, Spector PS. Electrogram fractionation: the relationship between spatiotemporal variation of tissue excitation and electrode spatial resolution. *Circulation: Arrhythmia and Electrophysiology*. 2011 Dec;4(6):909-16.
- [10] Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*. 2017 May 24;60(6):84-90.
- [11] Hunt B, Kwan E, McMillan M, Dossall D, MacLeod R, Ranjan R. Deep Learning Based Prediction of Atrial Fibrillation Disease Progression with Endocardial Electrograms in a Canine Model. In 2020 Computing in

- Cardiology 2020 Sep 13 (pp. 1-3). IEEE.
- [12] Koch G, Zemel R, Salakhutdinov R. Siamese neural networks for one-shot image recognition. InICML deep learning workshop 2015 Jul 10 (Vol. 2).
 - [13] Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems. 2019;32:8026-37.
 - [14] Kingma DP, Ba J. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980. 2014 Dec 22.
 - [15] Zhuang F, Cheng X, Luo P, Pan SJ, He Q. Supervised representation learning: Transfer learning with deep autoencoders. Twenty-Fourth International Joint Conference on Artificial Intelligence 2015 Jun 27.
 - [16] Pathak D, Krahenbuhl P, Donahue J, Darrell T, Efros AA. Context encoders: Feature learning by inpainting. InProceedings of the IEEE conference on computer vision and pattern recognition 2016 (pp. 2536-2544).

Address for correspondence:

Bram Hunt
Nora Eccles Harrison CVRTI
95 S 2000 E, Salt Lake City, UT 84112
bram.hunt@utah.edu