Predicting Ventricular Arrhythmias using Upstream Electrograms from Intracardiac Devices

Zuzana Koscova[1], Faisal Merchant[2], Mikhael El-Chami [2], Albert Rogers [3], Gari D Clifford [1,4], Neal K Bhatia [2]

[1] Department of Biomedical Informatics, School of Medicine, Emory University, USA
[2] Division of Cardiovascular Medicine, School of Medicine, Emory University, USA
[3] Department of Cardiovascular Medicine, Stanford University School of Medicine, USA
[4] Department of Biomedical Engineering, Georgia Institute of Technology, USA

Abstract

Ventricular arrhythmias (VAs) can lead to sudden cardiac death if not promptly managed. Implantable cardioverter defibrillators (ICDs) typically deliver therapy after VA onset, leaving a limited window for intervention, and can lead to adverse effects such as inappropriate shocks. This study investigates predicting VA onset using intracardiac electrograms (EGMs) recorded by subcutaneous ICDs immediately preceding VA. The training set contained 10,913 EGM recordings, including 236 upstream EGMs recorded before VA onset, while the test set included 3,712 recordings with 51 upstream EGMs.

A deep learning model, pretrained on a large ECG dataset, achieved strong performance with a mean AUROC of 0.95 ± 0.01 and an AUPRC of 0.62 ± 0.06 in predicting VA on the test set. Temporal analysis (up to 37 seconds before VA onset) revealed that the proximity of upstream EGMs to VA onset did not affect AUROC, indicating that earlier detection could provide enough time for intervention before the arrhythmia onset.

Although the model demonstrated high sensitivity (0.89 ± 0.02) and specificity (0.93 ± 0.02), improvements in precision (0.15 ± 0.03) are necessary. These findings highlight the potential of predicting VA onset from intracardiac EGMs, allowing timely therapeutic intervention before VA fully develops.

1. Introduction

Ventricular arrhythmias (VAs), including ventricular tachycardia (VT) and ventricular fibrillation (VF), are life-threatening disturbances in cardiac rhythm that can lead to sudden cardiac death if not promptly treated. Implantable cardioverter defibrillators (ICDs) have become the cornerstone of secondary prevention in patients with a history of sustained VA, as well as in selected high-risk populations for primary prevention [1]. These devices continuously monitor the heart rhythm and deliver therapeutic interventions—such as anti-tachycardia pacing or defibrillation—to terminate arrhythmias and restore normal rhythm.

Despite their effectiveness, current ICD systems function in a reactive manner, delivering therapy only after the onset of a VA episode. This often leaves a narrow window for response and can be associated with adverse outcomes, including inappropriate shocks, patient discomfort, and psychological burden [2]. As such, there is a growing need for predictive models that can anticipate the onset of VAs with sufficient lead time to enable earlier, more personalized interventions—ideally before the arrhythmia fully manifests.

In this work, we explored the predictive potential of intracardiac electrograms (EGMs) recorded by subcutaneous ICDs, focusing on the period immediately preceding VA onset. Utilizing deep learning techniques, we aim to detect early indicators of impending arrhythmias. The ability to predict arrhythmia onset offers valuable insights into the mechanisms of VA initiation and provides an opportunity for timely therapeutic intervention. We hypothesized that deep learning can uncover novel features in intracardiac EGMs prior to an arrhythmic event, enabling the prediction of VA onset.

2. Dataset

Intracardiac electrograms (EGMs) from patients undergoing subcutaneous implantable cardioverter-defibrillator (ICD) implantation (Boston Scientific, Boston, MA) at Emory University Hospital and Stanford Hospital were captured. All metadata and EGMs were exported from a remote monitoring database. For patients who experienced a VA event, the corresponding portion of the EGM was segmented and annotated. Segments recorded prior to the VA event—typically during sinus rhythm or atrial

fibrillation—were labeled as ”upstream”. Segments during the VA event itself were labeled as ”downstream”, and recordings from patients without a VA event present at that time were labeled as ”presenting rhythm”. Examples of upstream, downstream EGMs are shown in Figure 1.

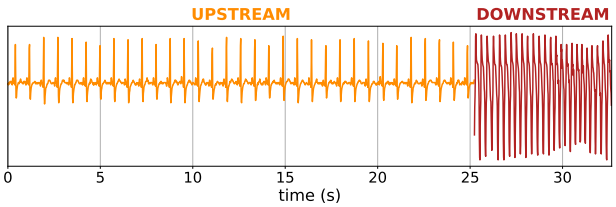

Figure 1. Example of a segmented electrogram (EGM) recorded using an implantable cardioverter-defibrillator. The **upstream** segment corresponds to the EGM recorded prior to the onset of a ventricular arrhythmia, while the **downstream** segment captures the episode of ventricular fibrillation.

Recordings from two independent centers were included in the study. For model development (training and validation), we used data from Emory University Hospital, comprising 10,913 EGM segments from 335 patients. Among these, 236 segments from 93 patients included upstream recordings preceding a VA episode. For independent evaluation, we used a test set from Stanford Hospital, which included 3,712 EGM recordings from 209 patients. Of these, 51 recordings from 24 patients included upstream EGMs. All recordings are sampled at 64 Hz.

The ICD device was configured to export a 12-second segment of the presenting rhythm when no VA event is detected. In contrast, when a VA event occurs, the duration of the upstream recording varies. In the Emory dataset, the upstream electrograms have a median duration of 35 seconds (interquartile range: 32–36 seconds), while in the Stanford dataset, the median duration is 36 seconds (IQR: 32–37 seconds).

3. Methods

Our objective was to predict the onset of VA from EGMs. To achieve this, we used both the upstream electrogram and the presenting rhythm, framing the task as a binary classification problem. The upstream electrogram was labeled as 1, indicating that it is followed by a VF/VT episode within the next few seconds. In contrast, the presenting rhythm was labeled as 0, as it is not followed by a VA episode. We performed 5-fold cross-validation on the Emory dataset, ensuring that the data were split at the patient level to avoid any patient overlap between the folds. The number of recordings in each validation fold, along with the prevalence of upstream electrograms, is summarized in Table 1. Finally, we tested our approach on an independent Stanford dataset, aggregating the predictions from all five models corresponding to five folds and computed the mean performance across the folds.

Table 1. Summary of 5-Fold Cross-Validation Splits.

Fold number	Validation size	Upstream EGMs (%)
1	2,160	54 (2.5)
2	2,544	51 (2.0)
3	1,759	51 (2.9)
4	2,151	47 (2.2)
5	1,811	49 (2.7)

3.1. Deep Learning Algorithm

We pretrained a deep learning model [3] on a large arrhythmia classification dataset, comprising over 86,000 recordings from the 2021 PhysioNet Challenge [4]. For VA/VT prediction from intracardiac EGMs, we reused the convolutional backbone from this pretrained model, which includes the initial convolutional and residual blocks. To adapt it for our specific task, we appended a custom output head with adaptive max pooling, fully connected layers, dropout, and batch normalization, followed by a sigmoid activation to estimate the probability of a VA/VT event. The model was then fine-tuned end-to-end on our EGM dataset. In total, we trained five models corresponding to the five folds for cross-validation.

All EGM recordings were preprocessed in the same manner as the data used for the original Challenge 2021 model. We resampled the data to a frequency of 500 Hz, applied a bandpass filter (1–47 Hz), zero-padded the recordings to a length of 8,192 samples, and performed z-score normalization. Since the original network was designed to handle a varying number of input ECG channels (1–12), the input matrix to the model is always 12x8,192. For our single-channel EGM data, we zero-pad the remaining leads.

Since the length of the presenting rhythm is fixed at 12 seconds, we truncated all upstream recordings to the same length. During training, we randomly selected segments from the upstream recordings, while for inference, we slid a 12-second window over the entire upstream recording with a one-second step. This approach enables us to investigate whether specific 12-second interval preceding a VA episode are more predictive than others. An example of how we iterate through the upstream recording is shown in Figure 2.

We trained the neural network for 30 epochs with an initial learning rate of 0.0001, using the Adam optimizer with a learning rate scheduler that reduces the rate by a factor of 0.1 every 6 epochs. As the loss function, we used a

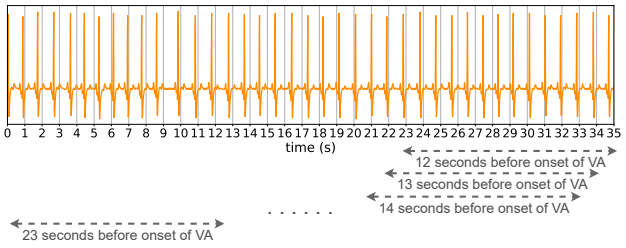

Figure 2. Segmentation of upstream EGMs into 12-second windows. Upstream EGMs were generally longer than the 12-second input required by the neural network. During inference, the upstream recording was divided into overlapping 12-second segments. This allowed us to evaluate the predictive power of each segment based on its temporal proximity to the onset of a ventricular arrhythmia (VA) episode.

Weighted Focal Loss, with weights assigned based on the prevalence of the minority class (upstream EGM), calculated as the total number of EGMs divided by the number of upstream EGMs.

To further investigate the regions on which the neural network is focused, we applied gradient-weighted class activation mapping (GradCam) to visualize the most important areas of upstream EGM for VA prediction.

4. Results

The results of five-fold cross-validation on the Emory dataset, evaluated using AUROC and AUPRC, are shown in Table 2. Performance was calculated as the mean AUROC and AUPRC across multiple 12-second segments extracted from the upstream portion of the recording, depending on their temporal proximity to the onset of VA. On the independent test set, we applied the five models trained during cross-validation and computed the median probability of VA; the results are reported in Table 3.

Table 2. Results of 5-fold cross-validation.

Fold	AUROC	AUPRC
1	0.94 (0.02)	0.56 (0.06)
2	0.95 (0.02)	0.57 (0.04)
3	0.90 (0.03)	0.61 (0.06)
4	0.89 (0.03)	0.51 (0.07)
5	0.95 (0.02)	0.75 (0.04)
Mean $\pm$ SD	0.92 $\pm$ 0.02	0.52 $\pm$ 0.06

On the test set, we also examined how the temporal proximity of upstream EGMs to the onset of VA influences prediction performance. Figure 3 shows the sensitivity and precision for different 12-second windows located between 0 and 25 seconds before the VA onset. Since the length of upstream EGMs varies across cases (ranging from 12 to 39 seconds), the figure also displays the prevalence of upstream segments available for each temporal window, relative to all classified segments (presenting and upstream combined).

Table 3. Results on an independent test set.

Metric	Mean $\pm$ SD
AUROC	0.95 (0.01)
AUPRC	0.62 (0.06)
Sensitivity	0.93 (0.03)
Specificity	0.94 (0.01)
Precision	0.15 (0.03)

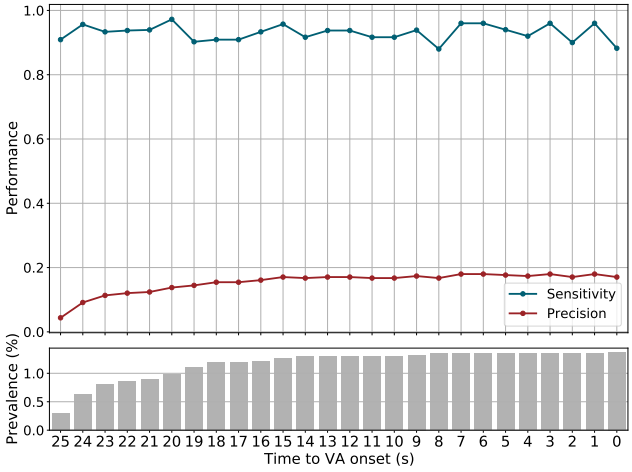

Figure 3. Precision and Sensitivity are reported for different 12-second segments of the upstream electrogram, based on their temporal proximity to the onset of ventricular arrhythmia (VA), ranging from 0-12 seconds (immediately preceding VA) to 25–37 seconds. The gray bar plot shows the prevalence of the upstream EGM class for each time window.

Analyzing the saliency map generated using the GradCam approach reveals that the ST segment and T-wave contribute more than other parts of the ECG to VA prediction, as shown in Figure 4, where the gradient is color-mapped.

5. Discussion and Conclusion

The results of the five-fold cross-validation demonstrate strong performance for the prediction of VAs based on upstream electrograms (EGMs), with a mean AUROC of 0.92 $\pm$ 0.02 and a mean AUPRC of 0.52 $\pm$ 0.06. These findings highlight the effectiveness of the model in distinguishing

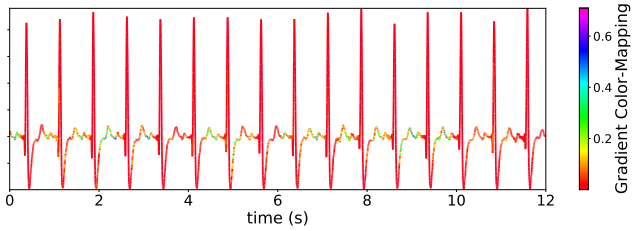

Figure 4. Saliency map generated using the GradCam approach, highlighting the regions of the upstream EGM, specifically the QT segment and T-wave, that mostly contribute to the prediction of ventricular arrhythmias.

between the upstream EGM preceding VA and the presenting rhythm (controls) across different temporal segments of the EGM data. Furthermore, performance on the independent test set showed slightly improved results, with an AUROC of 0.95 $\pm$ 0.01 and an AUPRC of 0.62 $\pm$ 0.06, indicating generalizability of the model to unseen data.

One of the key insights from this study was the evaluation of the temporal influence of upstream EGMs on prediction performance. As shown in Figure 3, the proximity of the upstream EGM to the onset of VA did not significantly impact the sensitivity. However, larger fluctuations were observed in the precision, which reflects the false positive rate. Specifically, precision was lower for segments farther from the VA onset, with performance improving as the window approached the onset. The observed decrease in precision for segments $>$18 seconds from VA onset may be also explained by the lower prevalence of upstream EGMs in these time intervals.

These results are consistent with previous studies that have shown the utility of machine learning models in predicting VAs from intracardiac EGM signals. For instance, a similar study by Cha et al. (2024) [5] reported an AUROC of 0.83 (95% CI: 0.79–0.89) for predicting VAs five seconds prior to event onset. In contrast, our study offers additional value by incorporating a more detailed temporal analysis of upstream EGMs, extending up to 25 seconds prior to VA onset. Our approach yielded a mean AUROC of 0.95 $\pm$ 0.01, highlighting both an improvement in predictive performance and the potential for earlier detection. The extended temporal window allows for the identification of segments in the EGM data well before the onset of VA, providing sufficient time for therapeutic intervention before the arrhythmia fully develops.

The neural network's emphasis on the QT segment and T-wave, identified through gradient mapping suggests that these features may be critical to the onset of VAs.

Although the model demonstrates promising performance, several limitations must be considered. Firstly, the dataset used in this study is relatively small, consisting of 335 subjects in the training set, with 28% exhibiting VA episodes, and 209 subjects in the test set, with 11% exhibiting VA episodes. While the challenge of limited data was partially addressed by pretraining the deep learning model on a larger ECG dataset, further validation with a more extensive dataset is necessary to confirm these findings. Additionally, while the model performs well in detecting VA events, the low precision (0.15 $\pm$ 0.03) suggests that further improvements may be needed for clinical application.

In conclusion, the results of this study highlight the potential of utilizing upstream EGMs for predicting ventricular arrhythmias through a robust deep learning approach, which achieved high sensitivity (0.93 $\pm$ 0.03) and specificity (0.94 $\pm$ 0.01). These findings could pave the way for improved clinical strategies for early prediction of VAs, offering the opportunity for timely therapeutic interventions that may ultimately contribute to the prevention of sudden cardiac death.

Acknowledgments

This work was funded by an unrestricted donation from AliveCor Inc.

References

[1] Goldenberg I, Huang DT, Nielsen JC. The role of implantable cardioverter-defibrillators and sudden cardiac death prevention: indications, device selection, and outcome. European Heart Journal June 2020;41(21):2003–2011.

[2] Pausch J, Mersmann J, Bhadra OD, Barten MJ, Tönnis T, Yildirim Y, Pecha S, Reichenspurner H, Bernhardt AM. Prognostic impact of implantable cardioverter defibrillators and associated adverse events in patients with continuous flow left ventricular assist devices. Frontiers in Cardiovascular Medicine 2023;10:1158248. URL https://doi.org/10.3389/fcvm.2023.1158248.

[3] Nejedly P, Ivora A, Smisek R, Viscor I, Koscova Z, Jurak P, Plesinger F. Classification of ECG using ensemble of residual CNNs with attention mechanism. In 2021 Computing in Cardiology (CinC), volume 48. 2021; 1–4.

[4] Reyna MA, Sadr N, Alday EAP, Gu A, Shah AJ, Robichaux C, Rad AB, Elola A, Seyedi S, Ansari S, Ghanbari H, Li Q, Sharma A, Clifford GD. Will two do? Varying dimensions in electrocardiography: The PhysioNet/Computing in Cardiology Challenge 2021. In 2021 Computing in Cardiology (CinC), volume 48. 2021; 1–4.

[5] Cha YM, Attia IZ, Metzger C, Lopez-Jimenez F, Tan NY, Cruz J, Upadhyay GA, Mullane S, Harrell C, Kinar Y, Sedelnikov I, Lerman A, Friedman PA, Asirvatham SJ. Machine learning for prediction of ventricular arrhythmia episodes from intracardiac electrograms of automatic implantable cardioverter-defibrillators. Heart Rhythm 2024;21(11):2295–2302. URL https://doi.org/10.1016/j.hrthm.2024.05.040.

Address for correspondence:

Zuzana Koscova; Dept of Biomedical Informatics, 101 Woodruff Circle, Atlanta, GA 30322; zuzana.koscova@dbmi.emory.edu