

Beyond the RVOT: Prognostic ECG Feature Discovery in Brugada Syndrome via Transfer Learning and Lead Ablation

Original

Beyond the RVOT: Prognostic ECG Feature Discovery in Brugada Syndrome via Transfer Learning and Lead Ablation / Casella, A., Randazzo, V., Gaita, F., Giustetto, C., Pasero, E.. - ELETTRONICO. - (2026), pp. 1-6. (2026 International Joint Conference on Neural Networks (IJCNN) Maastricht (Nld) 22-26 June 2026) [10.1109/ijcnn69854.2026.11729390].

Availability:

This version is available at: 11583/3016912 since: 2026-10-23T09:37:58Z

Publisher:

IEEE

Published

DOI:10.1109/ijcnn69854.2026.11729390

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Beyond the RVOT: Prognostic ECG Feature Discovery in Brugada Syndrome via Transfer Learning and Lead Ablation

Alessandro Casella
DET
Politecnico di Torino
Turin, Italy
alessandro.casella@polito.it

Vincenzo Randazzo
DET
Politecnico di Torino
Turin, Italy
vincenzo.randazzo@polito.it

Fiorenzo Gaita
Department of Medical Sciences
University of Turin
J Medical
Turin, Italy

Carla Giustetto
Department of Medical Sciences
University of Turin
Città della Salute e della Scienza
Turin, Italy

Eros Pasero
DET
Politecnico di Torino
Turin, Italy
eros.pasero@polito.it

Abstract—Arrhythmic risk stratification in Brugada Syndrome (BrS) patients is crucial for guiding clinical management, yet challenging due to the scarcity of high-risk events and the morphological variability of 12-lead electrocardiograms (ECGs). This study introduces a deep learning (DL) framework based on a Lightweight Residual Network (L-ResNet), which employs a compact convolutional neural network initially pre-trained on the Brugada pattern diagnostic task (with diagnostic accuracy of 96.5% on a three-class task) and subsequently fine-tuned using a retrospective Brugada cohort comprising 63 patients with documented events and 74 asymptomatic individuals. The method operates at beat level and aggregates predictions via both direct and hierarchical schemes. Two aggregation strategies were compared: a computationally efficient Beat-to-Patient pipeline and a multi-level Beat-to-ECG-to-Patient approach. Results demonstrate comparable performance (AUC = 0.757, Accuracy = 0.708). A post-hoc interpretability module based on lead ablation and integrated gradients confirm the dominant role of leads V_1 and V_2 , but unexpectedly highlights lateral leads such as aVL and V_5 as critical predictors of risk. The findings underscore the potential of interpretable DL to not only predict clinical outcomes but also to uncover previously underappreciated pathophysiological markers, supporting more refined risk stratification strategies in BrS.

Index Terms—Convolutional neural networks, risk stratification, ECG, transfer learning, arrhythmia prediction, explainable AI.

I. INTRODUCTION AND RELATED WORK

Despite decades of research, the identification of Brugada Syndrome (BrS) patients destined for life-threatening arrhythmias continues to rely on subjective and often inconclusive clinical criteria. BrS is an inherited condition linked to sudden cardiac death (SCD) in clinically stable individuals [1]. While diagnostic criteria for right precordial patterns are established [2], prognostic risk stratification remains uncertain, requiring a balance between preventing rare catastrophic events and avoiding invasive device complications [3]. This uncertainty

is amplified by the scarcity of arrhythmic events [4], heterogeneous acquisition [5], and confounding factors such as comorbid cardiac substrates and environmental triggers [6]. This diagnostic-prognostic disconnect underscores the need for interpretable tools to identify prognostic ECG signatures.

Deep Learning (DL) has shown high performance in large-scale ECG tasks [7], [8], but diagnostic accuracy does not inherently translate into prognostic utility, especially in rare diseases where outcomes reflect long-term events rather than instantaneous morphology. High-capacity architectures often overfit or lack interpretability when applied to small, unbalanced cohorts [9], [10]. In BrS, prior studies focused on patterns in leads associated with the right ventricular outflow tract (RVOT) [11], [12], while the prognostic role of non-canonical leads remains largely unexplored [13]–[17].

Post-hoc interpretability can improve transparency, but saliency-based explanations must be interpreted cautiously because attribution maps are sensitive to model and preprocessing choices and do not imply causal physiological mechanisms. Accordingly, in this study interpretability analyses are used as hypothesis-generating evidence rather than standalone mechanistic proof.

The work combines transfer learning for low-event BrS prognosis, comparison of patient-level aggregation strategies, and post-hoc interpretability analysis (lead ablation and integrated gradients).

II. METHODOLOGY

The proposed methodology for risk stratification relies on a two-stage transfer learning approach. First, a DL model is pre-trained on a large diagnostic dataset to learn robust ECG feature representations. Subsequently, this model is fine-tuned on a separate prognostic dataset to predict patient risk.

TABLE I
DETAILED COMPOSITION OF THE DIAGNOSTIC AND PROGNOSTIC DATASETS.

Dataset	Subgroup/Class	Patients (%)	Total ECGs (%)	Training / Val. (%)	Native Digital (%)	Digitized (%)
Diagnostic	Negative Controls	181 (59.3%)	244 (33.1%)	195 (79.9%) / 49 (20.1%)	200 (82.0%)	44 (18.0%)
	RBBB	77 (25.3%)	252 (34.3%)	201 (79.8%) / 51 (20.2%)	210 (83.3%)	42 (16.7%)
	Diagnostic/Suspicious BrS	47 (15.4%)	240 (32.6%)	194 (80.8%) / 46 (19.2%)	174 (72.5%)	66 (27.5%)
	<i>Total</i>	<i>305 (100%)</i>	<i>736 (100%)</i>	<i>590 (80.2%) / 146 (19.8%)</i>	<i>584 (79.3%)</i>	<i>152 (20.7%)</i>
Prognostic	High-Risk (Events)	63 (46.0%)	211 (52.0%)	Patient-wise	4 (1.9%)	207 (98.1%)
	Low-Risk (Asymptomatic)	74 (54.0%)	195 (48.0%)	5-fold CV	114 (58.5%)	81 (41.5%)
	<i>Total</i>	<i>137 (100%)</i>	<i>406 (100%)</i>	–	<i>118 (29.1%)</i>	<i>288 (70.9%)</i>

A. Study Population and Datasets

Two distinct datasets were employed for the diagnostic and prognostic tasks, with a stratified splitting to avoid overlapping in patients or recordings. Data collection adhered to the Declaration of Helsinki and received approval from the Institutional Review Board of Città della Salute e della Scienza Hospital, Turin, Italy (*Protocol code: 248/2017, approved 22 February 2017*). Informed consent was obtained from all subjects. The detailed composition of the datasets is summarized in Table I.

The *diagnostic* dataset comprises 12-lead ECG recordings from public repositories (including PTB-XL [18]) and private institutional archives. Organized for multi-class classification, it includes negative controls, right bundle branch block (RBBB), and BrS (encompassing definitive and suspected forms). Unlike the prognostic cohort, this dataset incorporates pharmacological test recordings obtained after sodium channel blocker administration.

The *prognostic* dataset consists of 137 patients diagnosed with BrS, stratified by long-term clinical outcomes. The high-risk cohort includes patients who experienced major arrhythmic events (SCD, aborted SCD, or appropriate implantable cardioverter defibrillator interventions), while the low-risk cohort consists of asymptomatic patients. A critical challenge in this dataset is the significant imbalance in signal acquisition formats: prospective native digital recordings are scarce among event-positive cases compared to the control group, necessitating the digitization of retrospective paper archives for high-risk patients.

B. Signal Preprocessing and Beat Segmentation

Both datasets underwent identical preprocessing pipelines to ensure morphological consistency. All recordings, whether native digital or digitized from paper, were resampled to a uniform frequency of 1000 Hz. Baseline wander was removed via linear detrending and high-pass filtering, while high-frequency noise was suppressed using a low-pass filter with cutoff at 40 Hz. Beat detection was performed using the continuous wavelet transform with a Morlet mother wavelet to robustly identify S-wave fiducial points. For each detected beat, a fixed temporal window of 800 ms, with the detected S-wave peak aligned at 300 ms, was extracted to capture complete depolarization and repolarization phases. Beat-wise processing was mandated by the lack of temporal synchronization in

digitized retrospective records, necessitating independent lead analysis. In order to reduce the exclusion of digitized useful signal, a technique was applied to create copies of the same heartbeat in a lead if such lead had far fewer examples than others of the same ECG.

C. Data Augmentation and Regularization Strategy

To enhance model robustness against the inherent variability of clinical ECG acquisition, a multi-faceted augmentation protocol was applied during training.

Temporal misalignment augmentation was designed to compensate for potential failures in the automated S-wave peak detection algorithm. In practice, the fiducial point detection may occasionally misalign across the twelve leads, particularly in recordings with low signal-to-noise ratio or atypical morphologies. To simulate these scenarios, a random inter-lead jitter was applied to a subset of leads during training. Each selected lead was shifted by an integer offset $\delta \sim \mathcal{U}[-15, 15]$ samples. Replicate padding was employed at signal boundaries to maintain waveform continuity without introducing artificial discontinuities.

Lead-wise regularization was implemented through a structured dropout mechanism. During each training iteration, individual leads were randomly suppressed with probability $p = 0.2$ by setting all samples to zero. This procedure forces the network to distribute feature extraction across the full set of available leads rather than relying on dominant individual channels.

D. Neural Network Architecture

The neural network backbone, detailed in Table II, is a compact residual architecture (L-ResNet) optimized for 12-lead ECG analysis on (12, 800) inputs. To address the scarcity of data typical of rare disease cohorts, the design prioritizes architectural regularization. The network initiates with a stem block featuring a large receptive field (kernel size 15) specifically designed to capture low-frequency morphological characteristics, such as QRS duration and P-wave morphology.

The subsequent feature extractor consists of three residual stages with progressive downsampling. Convolutional kernels are set to a temporal span of 5 samples to match the spectral properties of ECG signals. To further mitigate overfitting, structured dropout ($p = 0.2$) is embedded within every residual block between the activation and the second convolution.

TABLE II
L-RESNET BACKBONE SPECIFICATIONS

Layer	Output Dimensions	Kernel / Stride	Channels
Input Tensor	12×800	—	12
Stem Convolution	32×400	$15 \times 1 / 2$	32
Max Pooling	32×200	$3 \times 1 / 2$	32
Residual Stage 1 ($\times 2$ blocks)	32×200	$5 \times 1 / 1$	32
Residual Stage 2 Block 1	64×100	$5 \times 1 / 2$	64
Residual Stage 2 Block 2	64×100	$5 \times 1 / 1$	64
Residual Stage 3 Block 1	128×50	$5 \times 1 / 2$	128
Residual Stage 3 Block 2	128×50	$5 \times 1 / 1$	128
Global Average Pooling	128×1	—	128
Feature Vector $\mathbf{h}$	128	—	—

A final global average pooling layer collapses the temporal dimension to produce a 128-dimensional embedding vector.

E. Transfer Learning and Experimental Validation

The model training procedure is divided into two main stages: diagnostic pre-training to learn robust feature representations, and prognostic fine-tuning to adapt these features for risk stratification.

The *diagnostic* model was trained on the three class dataset described previously to perform classification among negative controls, RBBB, and Brugada patterns. The identical L-ResNet backbone was employed, with the final layer modified to output three logits. The network was initialized randomly and trained end to end from the first epoch, without progressive unfreezing. Optimization was performed using AdamW with a learning rate of 10^{-3} and weight decay of 10^{-4} . The loss function was standard cross entropy without class weighting, given the relatively balanced distribution of the three diagnostic categories. A scheduler tracked the validation loss and reduced the learning rate by a factor of 0.5 if no improvement was observed for 5 consecutive epochs. Training proceeded for a maximum of 50 epochs with early stopping based on validation accuracy (patience 15), since fold-wise monitoring did not show consistent late-phase validation gains after early convergence. The model achieving the highest validation accuracy was retained as the source for subsequent prognostic transfer.

The *prognostic* adaptation employs a distinct protocol tailored to the binary risk stratification task and the limited sample size. Unlike the diagnostic training, which optimizes all parameters from scratch, the prognostic phase initiates from the diagnostic weights and employs progressive layer unfreezing. A warmup phase of 5 epochs freezes the backbone and trains only the classification head (now with two outputs) using AdamW with learning rate 10^{-3} and weight decay 10^{-2} . This prevents catastrophic forgetting of morphological features while adapting the decision boundary. Subsequently, the entire network is unfrozen and fine tuned end to end with a reduced learning rate of 10^{-4} . Crucially, given the extreme class imbalance in the prognostic cohort, the cross entropy loss is weighted inversely to class frequency, assigning higher penalization to misclassification of the minority high risk group. A scheduler with factor 0.1 and patience 3 epochs

monitors validation loss for dynamic learning rate adaptation on plateau.

F. Baseline Models and Comparative Protocol

For comparative analysis, two baselines were evaluated under the same patient-wise folds and preprocessing used for the proposed method. The first baseline uses the same L-ResNet architecture trained from scratch on the prognostic task (no diagnostic pre-training), with the same optimizer family, stopping policy, and aggregation logic. The second baseline is a logistic regression classifier trained on patient-level descriptors including heart rate and RR variability, median QRS duration and QTc, and lead-specific morphology features (J-point amplitude, ST slope, S-wave depth, and QRS area) extracted from V_1 , V_2 , aVL, and V_5 and aggregated per patient by median across beats; regularization is selected within each training fold. Comparative results are reported as mean and standard deviation of AUC across folds, with paired fold-level analysis for statistical comparison.

G. Patient-Level Aggregation

From a clinical perspective, the neural network’s output cannot be reported at the granularity of individual heartbeats; its predictions must be translated into an overall interpretation of one or more ECG recordings to characterize the patient’s condition. Two aggregation logics are considered to map beat-level outputs to a patient-level score. Let $p_b \in [0, 1]$ denote the predicted probability for beat b belonging to the positive class.

The *Top-K* mean aggregation computes a patient score by selecting the beats with the highest predicted probabilities and averaging their probabilities. Let $\mathcal{B}_i$ be the set of beats associated with patient i . Let $\mathcal{T}_i \subseteq \mathcal{B}_i$ be the subset containing the K elements with the largest values of p_b . The patient score is defined as $s_i = \frac{1}{|\mathcal{T}_i|} \sum_{b \in \mathcal{T}_i} p_b$. The value of K is set according to the Top-K configuration.

The *Vote* aggregation converts beat probabilities into beat labels via a decision threshold τ , then assigns the patient label by majority voting across beats. Specifically, beat b is classified as positive if $p_b \geq \tau$. The patient label is obtained as the most frequent beat label within $\mathcal{B}_i$. For threshold independent evaluation, receiver operating characteristic (ROC) curves are obtained by varying τ .

Two architectural variants are also evaluated. *Direct* aggregation maps beats to a patient score using a single aggregation step on $\mathcal{B}_i$. *Hierarchical* aggregation introduces an intermediate ECG-level. Let $\mathcal{R}_i$ be the set of ECG recordings for patient i , and let $\mathcal{B}_{ir}$ be the set of beats belonging to recording $r \in \mathcal{R}_i$. A recording score s_{ir} is first computed by applying the same beat aggregation rule within $\mathcal{B}_{ir}$. The patient score is then computed as the mean of recording scores.

H. Interpretability and Lead Importance Analysis

To bridge the gap between DL predictions and clinical practice, a post-hoc interpretability framework was implemented to identify the features driving the prognostic output. This

analysis focuses on two scales: global lead importance and local morphological attribution.

The global contribution of each lead was quantified through a systematic lead ablation study. In this procedure, the importance of a specific lead is measured by evaluating the model performance on a modified version of the test set where the signal of that lead is entirely zeroed out. This occlusion sensitivity analysis allows for the mathematical identification of the leads that most significantly influence the area under the curve and is interpreted as a relative sensitivity test rather than a physiologically realistic counterfactual. A substantial decrease in performance upon the removal of a specific lead indicates that the network has prioritized the features within that channel to differentiate between risk groups.

Local interpretability was achieved using Integrated Gradients (IG) to generate high-resolution saliency maps at the beat level. This method assigns attribution scores to each time step of the ECG signal by integrating the gradients of the model prediction with respect to the input along a path from a zero-baseline to the actual heartbeat. Unlike simple gradient maps, this approach satisfies the axiom of completeness, providing a reliable visualization of which segments of the P-QRS-T complex, such as the J-point elevation or the ST-segment morphology, are most influential for the final risk score.

III. RESULTS AND DISCUSSION

A. Diagnostic Pre-training Performance

The initial experimental stage focused on the diagnostic pre-training task, designed to validate the L-ResNet model's ability to serve as a robust feature extractor. The model was trained on a dataset of 5,443 beat windows and evaluated on a distinct validation set containing 1,412 windows, with the objective of performing a three-class classification of BrS, RBBB, and normal ECG patterns.

Following the protocol described in Section II-E, the training process converged after 18 epochs, triggered by the early stopping criterion. The best performing model checkpoint, which was saved for the subsequent fine-tuning stage, achieved a peak validation accuracy of 96.53%.

This high level of performance confirms that the architecture effectively captured the complex and subtle morphological features inherent in 12-lead ECG signals.

B. Prognostic Training Dynamics

The convergence curves exhibit a consistent decrease in both training and validation loss across the five cross-validation folds, indicating a stable and well-behaved optimization process. Specifically, the model achieves rapid convergence within the first 20 epochs, settling at a validation loss of approximately 0.61. As training progresses, the dispersion narrows due to the application of early stopping, which halted individual folds once validation performance plateaued. The persistent alignment and the minimal gap between training and validation profiles suggest that the model generalizes effectively and does not suffer from overfitting across folds.




Fig. 1. Patient-level ROC curves for *Top-K* mean and *Vote* aggregation. The shaded regions indicate variability across cross validation folds.

C. Prediction at Medical Level

Post-processing must be applied to the beat-level prediction probabilities for a clinical outcome (see Section II-G).

The *Top-K* mean and *Vote* methods yield comparable discriminative performance. Both aggregation strategies achieve similar AUC values and accuracy metrics across the cross-validation folds (Fig. 1), with mean area under the curve (AUC) ranging approximately 0.76 and accuracy around 0.71. Although the average performance is similar, *Top-K* mean is preferable when the objective is to emphasize the contribution of the most informative beats.

Under *Top-K* mean, the mean AUC of the *Hierarchical* aggregation from beat-to-ECG-to-patient is slightly higher compared to the *Direct* aggregation from beat-to-patient, while the most evident effect is the reduction of variability across folds. *Direct* achieves $AUC = 0.757 \pm 0.081$ and $Accuracy = 0.708 \pm 0.080$, whereas *Hierarchical* achieves $AUC = 0.760 \pm 0.062$ and $Accuracy = 0.678 \pm 0.044$. The intermediate ECG-level acts as a smoothing stage that reduces variance, leading to more stable patient scores.

The difference in AUC is not statistically significant, with $p = 0.83$ under the paired fold level comparison. In practice, the added complexity of the hierarchical pipeline does not provide a measurable benefit in mean discriminative performance for the considered task; direct aggregation represents an adequate choice due to its computational simplicity and comparable AUC.

D. Baseline Comparison

Using the comparative protocol described in Section II-E, the proposed transfer-learning model achieved $AUC = 0.757 \pm 0.081$, compared with $AUC = 0.710 \pm 0.080$ for the from-scratch L-ResNet and $AUC = 0.650 \pm 0.070$ for logistic regression.

E. Global Feature Importance via Lead Ablation

The AUC drop observed when each ECG lead was independently masked in the validation set is shown in Fig. 2 to conduct the lead ablation study and its impact on performance.




Fig. 2. Performance degradation induced by single-lead ablation. Bars represent the AUC drop with respect to the full model (12-lead) on the validation set. Lateral leads such as aVL and V₅ show unexpectedly high importance, indicating a possible role in risk stratification beyond the classical diagnostic region.

The precordial leads V₁ and V₂ exhibit the highest importance, consistent with their established role in identifying the characteristic Brugada pattern localized to the RVOT. However, the most significant finding is the prominent contribution of lateral leads. Specifically, lead aVL demonstrates the second-highest predictive value, surpassing even V₁ in terms of AUC degradation. Leads V₅ and V₆ also show substantial importance, ranking among the top five contributors. Features in aVL, V₅, and V₆ may reflect high lateral wall conduction delays (with the exclusion of the fourth lateral lead I). In contrast, inferior leads (II, III, aVF) contribute minimally to the model’s decision-making process, indicating redundancy or limited discriminative power for this specific prognostic task.

More notably, the significant involvement of lateral leads such as aVL and V₅ suggests that the model is receptive to features extending beyond the classical diagnostic zone. These leads likely reflect subtle manifestations of conduction delay or abnormal depolarization patterns projected laterally, such as prominent or slurred S-waves and QRS fragmentation.

The importance attributed to these non-canonical leads implies that the predictive signature of arrhythmic risk may not be fully captured by right precordial abnormalities alone. Rather, prognostic risk appears to be modulated by broader depolarization dynamics, effectively projected onto lateral vectors of the 12-lead ECG.

F. Local Interpretability and Alignment Analysis

The IG saliency map for a representative true positive patient is illustrated in Fig. 3. The signal alignment indicates a successful application of the preprocessing pipeline, as the QRS complexes across most leads appear temporally consistent. A mild deviation is observed in lead III, where the S-wave occurs slightly later than 300 ms, suggesting a degree of inter-lead misalignment. This empirical observation supports the inclusion of inter-lead jitter augmentation during training (see Section II-C), which aimed to ensure robustness against sub-millisecond localization artifacts from automated beat segmentation.




Fig. 3. Saliency map on raw 12-lead ECG signal generated via integrated gradients for a representative true positive patient. Minor inter-lead timing variability (e.g. delayed S-wave in Lead III) highlights the relevance of jitter-based augmentation. Color overlay of attribution intensity: red indicates features contributing positively to risk classification (e.g. ST segment in V₁, deep S-waves in aVL, V₅, I), while blue denotes features aligned with the control class (e.g. S-waves in V₃, V₄ and T-wave in V₂).

Focusing on regions associated with positive attribution, the model assigns strong relevance to the ST segment and T-wave morphology in lead V₁, in agreement with known risk markers of repolarization alteration in BrS. In addition, significant positive attribution is observed on the S-wave downslopes in several lateral and inferior leads, notably aVL, V₅, I, II, and V₂. This distribution suggests sensitivity to delayed depolarization projected along alternative anatomical vectors, reflecting the potential extension of conduction abnormalities beyond the canonical right precordial region.

Regions with negative attribution, interpreted by the model as resembling non-risk phenotypes, appear primarily over the S-waves in V₃ and V₄, with additional negative attribution on the T-wave in V₂.

G. Limitations and Threats to Validity

Several limitations should be considered when interpreting these results. First, the prognostic cohort presents a marked acquisition-format imbalance (digitized vs. native digital ECGs) across classes, which may introduce residual confounding despite identical preprocessing. Second, although patient-wise cross-validation reduces leakage, the study does not include an external test cohort; therefore, generalization

to independent centers remains to be established. Finally, the beat replication strategy used to compensate lead-wise sparsity may introduce sampling bias. For these reasons, the present findings should be interpreted as hypothesis-generating and require prospective multi-center validation.

IV. CONCLUSION AND FUTURE DEVELOPMENTS

This work demonstrates the effectiveness of a transfer learning framework for arrhythmic risk stratification in BrS, transitioning from a large-scale diagnostic task to a specialized prognostic objective. It shows that DL is able to detect subtle differences in cardiac beat patterns between patients who experienced an event and those who never developed severe symptoms. The integration of global and local interpretability tools suggests that these models can support clinically meaningful hypothesis generation, beyond purely black-box prediction. Specifically, the lead ablation study reveals that while diagnosis is anchored to the right precordial leads, prognostic risk is significantly modulated by the lateral leads, with aVL and V₅ emerging as critical predictors. This finding suggests that the extension of the arrhythmogenic substrate beyond the RVOT may be an important determinant of life-threatening events.

Labeling patients based on observed events introduces a degree of uncertainty, as SCD or syncope may be influenced by undetected comorbidities, environmental triggers, or lifestyle factors that the ECG alone cannot capture. The operating threshold can be tuned to prioritize sensitivity or specificity according to the intended clinical use. From a clinical perspective, this conservative behavior reduces the risk of unnecessary and invasive cardioverter-defibrillator implantations, yet it requires careful calibration to avoid underestimating potentially fatal arrhythmic risks. Furthermore, local attribution analysis via IG (see Fig. 3) highlights current technical limitations in signal alignment. Despite the use of jitter augmentation, slight misalignments in certain leads can lead to morphological ambiguity, such as the potential misinterpretation of Q-waves as S-waves.

Future developments will explore the use of dedicated neural modules for precise lead-wise wave detection to refine beat centering. Additionally, while preliminary internal evaluations of multiple instance learning and attention mechanisms showed lower performance due to the limited dataset size, these end-to-end architectures remain a promising direction for future research as larger, multi-center cohorts become available.

REFERENCES

- [1] Y. G. Kim, S.-K. Oh, H. Y. Choi, and J.-I. Choi, "Inherited arrhythmia syndrome predisposing to sudden cardiac death," *The Korean Journal of Internal Medicine*, vol. 36, no. 3, p. 527, 2021.
- [2] A. B. de Luna, J. Brugada, A. Baranchuk, M. Borggrefe, G. Breithardt, D. Goldwasser, P. Lambiase, A. P. Riera, J. Garcia-Niebla, C. Pastore, *et al.*, "Current electrocardiographic criteria for diagnosis of brugada pattern: a consensus report," *Journal of electrocardiology*, vol. 45, no. 5, pp. 433–442, 2012.
- [3] P. Calvelli, N. Cerrato, C. Giustetto, A. Saglietto, M. Anselmino, and A. Curcio, "Which brugada patient deserves continuous ecg monitoring through implantable loop recorder? an evidence update," *Journal of Cardiovascular Medicine*, vol. 26, no. 2, pp. 64–71, 2025.
- [4] P. Bordachar, S. Reuter, S. Garrigue, X. Cai, M. Hocini, P. Jaïs, M. Haïssaguerre, and J. Clementy, "Incidence, clinical implications and prognosis of atrial arrhythmias in brugada syndrome," *European heart journal*, vol. 25, no. 10, pp. 879–884, 2004.
- [5] P. Kligfield, L. S. Gettes, J. J. Bailey, R. Childers, B. J. Deal, E. W. Hancock, G. Van Herpen, J. A. Kors, P. Macfarlane, D. M. Mirvis, *et al.*, "Recommendations for the standardization and interpretation of the electrocardiogram: part i: the electrocardiogram and its technology," *Journal of the American College of Cardiology*, vol. 49, no. 10, pp. 1109–1127, 2007.
- [6] D. F. d. A. F. Fernandes, "Adequation of the therapeutic profile, comorbidities and personal habits of brugada patients with icd and its link to dysrhythmic events," Master's thesis, Universidade de Coimbra (Portugal), 2017.
- [7] H. Zhu, C. Cheng, H. Yin, X. Li, P. Zuo, J. Ding, F. Lin, J. Wang, B. Zhou, Y. Li, *et al.*, "Automatic multilabel electrocardiogram diagnosis of heart rhythm or conduction abnormalities with deep learning: a cohort study," *The Lancet Digital Health*, vol. 2, no. 7, pp. e348–e357, 2020.
- [8] S. Gaurhar, A. K. Tiwari, and S. Deora, "A cnn-based classifier for detecting rhythm disorders, premature contractions, and conduction abnormalities from ecg," *Biomedical Signal Processing and Control*, vol. 113, p. 108834, 2026.
- [9] E. Pasero, F. Gaita, V. Randazzo, P. Meynet, S. Cannata, P. Maury, and C. Giustetto, "Artificial intelligence ecg analysis in patients with short qt syndrome to predict life-threatening arrhythmic events," *Sensors*, vol. 23, no. 21, p. 8900, 2023.
- [10] S. Caligari, V. Randazzo, C. Giustetto, F. Gaita, and E. Pasero, "Ecg images analysis for the prediction of fatal arrhythmic events in patients with brugada syndrome based on a vision transformer model," in *Advanced Neural Artificial Intelligence: Theories and Applications*, pp. 13–22, Springer, 2025.
- [11] E. Pasero, A. Casella, and V. Randazzo, "Towards explainable ai in cardiac diagnostic: Attention-based interpretation of brugada syndrome," in *2025 IEEE International Conference on Artificial Intelligence for Learning and Optimization (ICoAILO)*, pp. 1–6, IEEE, 2025.
- [12] V. Randazzo, A. Casella, S. Caligari, F. Gaita, C. Giustetto, and E. Pasero, "A three-class ai model for brugada syndrome detection to improve diagnostic accuracy in ecg analysis," in *2025 IEEE Medical Measurements & Applications (MeMeA)*, pp. 1–6, IEEE, 2025.
- [13] F. Gaita, N. Cerrato, C. Giustetto, L. Garberoglio, and L. Calò, "Brugada syndrome: identification of subjects at risk and therapy," *European Heart Journal Supplements*, vol. 26, no. Supplement_1, pp. i69–i73, 2024.
- [14] B. Asatryan, P. G. Postema, and A. A. Wilde, "Pushing prognostic boundaries in brugada syndrome: Trying to predict the unpredictable," *Heart Rhythm*, vol. 20, no. 10, pp. 1368–1369, 2023.
- [15] C. Giustetto, G. Nangeroni, N. Cerrato, B. Rudic, E. Tülümen, E. Gribaudo, D. F. Giachino, L. Barbonaglia, L. M. Biava, P. Carvalho, *et al.*, "Ventricular conduction delay as marker of risk in brugada syndrome. results from the analysis of clinical and electrocardiographic features of a large cohort of patients," *International Journal of Cardiology*, vol. 302, pp. 171–177, 2020.
- [16] F. Gaita, N. Cerrato, C. Giustetto, A. Martino, L. Bergamasco, M. Millesimo, L. Barbonaglia, P. Carvalho, D. Caponi, A. Saglietto, *et al.*, "Asymptomatic patients with brugada ecg pattern: long-term prognosis from a large prospective study," *Circulation*, vol. 148, no. 20, pp. 1543–1555, 2023.
- [17] C. T. Chung, G. Bazoukis, D. Radford, E. Coakley-Youngs, R. Rajan, P. T. Matusik, T. Liu, K. P. Letsas, S. Lee, and G. Tse, "Predictive risk models for forecasting arrhythmic outcomes in brugada syndrome: A focused review," *Journal of Electrocardiology*, vol. 72, pp. 28–34, 2022.
- [18] P. Wagner, N. Strodthoff, R.-D. Boussejot, D. Kreiseler, F. I. Lunze, W. Samek, and T. Schaeffter, "Ptb-xl, a large publicly available electrocardiography dataset," *Scientific data*, vol. 7, no. 1, pp. 1–15, 2020.