HEART - Hybrid ECG Analysis for Recognizing Chagas

Pinar Bisgin[1,2], Martin Sondermann[1], Harshini Eggoni[1], Ole Werger[3], Prabhudev Bengaluru Kumar[1], Florian Lübbe[1], Maximilian Fecke[1], Niklas Tschorn[1], Mostafa Kamal Mallick[1], Michael Pantförder[1], Hendrik Wöhrle[3,4]

[1] Fraunhofer Institute for Software and Systems Engineering ISST, Dortmund, Germany
[2] TU Dortmund University, Dortmund, Germany
[3]Fraunhofer Institute for Microelectronic Circuits and Systems, Duisburg, Germany
[4] University of Duisburg-Essen, Duisburg, Germany

Abstract

Chagas disease, caused by Trypanosoma cruzi, poses a major global public health challenge, particularly in endemic regions. With an estimated 6.5 million affected individuals and nearly 10,000 deaths annually, effective screening and diagnostic strategies are critical. This study leverages advances in machine learning and ECG data analysis to improve patient prioritization for confirmatory testing. The proposed HEART system employs a spectrogram-based convolutional neural network with a novel lead-aware attention mechanism, achieving promising results in identifying at-risk patients. Our model achieved a Challenge Score of 30.3 % *during the offical phase, demonstrating its effectiveness in ranking Chagas-positive patients among the top cases. These findings highlight the potential of algorithmic approaches to enhance diagnostic accuracy and reduce the morbidity associated with Chagas disease.*

1. Introduction

Chagas disease, a tropical parasitic infection caused by Trypanosoma cruzi, represents a significant public health concern, particularly in endemic regions. The parasite is primarily transmitted by triatomine insects, commonly known as "kissing bugs," and is estimated to infect about 6.5 million people worldwide, causing nearly $10,000$ deaths annually [1]. Importantly, the disease is not confined to endemic countries; it also affects migrants living in high-income nations [2]. Despite its severity, no vaccine is currently available, underscoring the need for improved screening and diagnostic measures.

The disease progresses through distinct phases, beginning with an acute stage—often occurring in childhood—that is usually mild or asymptomatic. At this stage, treatment with antiparasitic medication can effectively halt disease progression [3, 4]. However, untreated individuals may develop chronic disease, which can lead to severe complications such as Chagas cardiomyopathy. This condition is characterized by heart failure, arrhythmias, and an elevated risk of thromboembolism [5]. While diagnosis generally relies on serological testing, access to such resources remains limited in many regions. Notably, Chagas cardiomyopathy frequently manifests in electrocardiograms (ECGs), providing a valuable signal for identifying at-risk patients and guiding subsequent treatment decisions [6].

Against this backdrop, the George B. Moody PhysioNet Challenge 2025 offers a unique opportunity to advance screening methods for Chagas disease. The competition encourages teams to develop innovative algorithmic approaches that use ECG data to prioritize patients for confirmatory testing. By harnessing recent advances in artificial intelligence and machine learning, this initiative aims to improve diagnostic capabilities and ultimately reduce the morbidity and mortality associated with Chagas disease.

2. Methods

2.1. Data Introduction

The Challenge used the CODE-15% dataset [7], the SaMi-Trop dataset [8], the PTB-XL dataset [9], and multiple private datasets from Chagas-endemic areas. The CODE-15% dataset contains over $300,000$ 12-lead ECG recordings collected in Brazil between 2010 and 2016, with self-reported binary Chagas labels. The SaMi-Trop dataset includes $1,631$ validated Chagas-positive ECGs collected in Brazil between 2011 and 2012. The PTB-XL dataset provides 21,799 ECGs from non-Chagas patients in Europe between 1989 and 1996. All datasets differ in collection and validation procedures but together represent realistic clinical data. The training set combines CODE-15%

with weak labels and SaMi-Trop and PTB-XL with strong labels, while validation and test sets consist of strongly labeled data, reflecting the prevalence rates in endemic countries.

2.2. Data Preprocessing

Since the ECG recordings varied in sampling frequency and duration, we transform all signals to a fix samplerate of 400 Hz. The length is standardized to a length of 10 seconds using truncation or padding on both ends, resulting in $4,000$ samples per lead. Each of the 12 leads was normalized using per-lead Z-score normalization, with noise injection applied to flat signals to avoid numerical instability. From the preprocessed signals, spectrogram representations were extracted, primarily with the Short-Time Fourier Transform (STFT, window size of 256 samples, hop length of 48, Hann window), while Mel-frequency cepstral coefficients (MFCC) and Stockwell transforms were also evaluated as alternatives. This ensured a consistent and physiologically relevant time–frequency representation of the ECG data, as illustrated in Fig. 1.

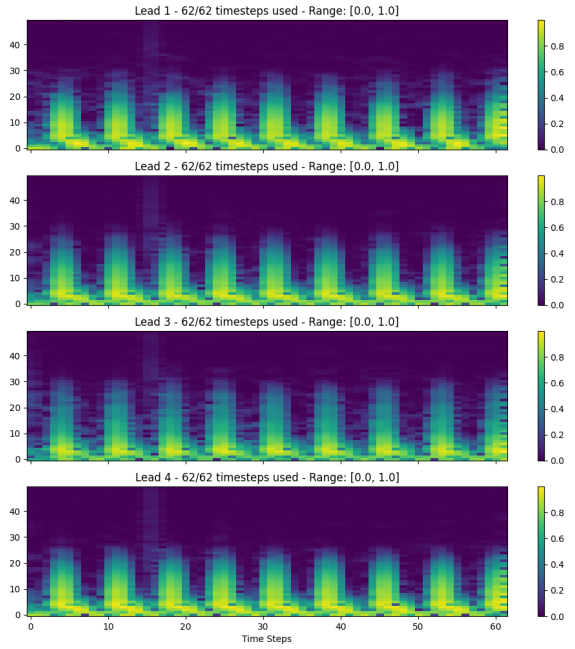

Figure 1. STFT features from four ECG leads, showing time (x-axis), coefficients (y-axis), and feature intensity (color), used for Chagas detection.

2.3. Data Augmentation

To mitigate overfitting and enhance model generalization, spectrogram-based augmentation techniques were applied. Specifically, frequency masking (up to 25 frequency bins with probability 0.4) and time masking (up to 30 frames with probability 0.4) were performed on the spectrograms during training. These augmentations simulated variability in signal acquisition and reduced dependency on narrow spectral features. In contrast, raw-signal augmentations such as time stretching, amplitude scaling, or Gaussian noise injection were disabled to preserve data consistency with the caching strategy.

2.4. Model Structure

The proposed architecture Fig. 2 is a spectrogram-based convolutional neural network enhanced with a lead-aware attention mechanism. Each ECG lead is first processed independently through a dedicated feature extractor consisting of three convolutional blocks. These blocks combine convolutional layers with batch normalization, SiLU activations, and max-pooling operations, followed by an adaptive average pooling layer. This design captures local time–frequency characteristics of the spectrogram while progressively reducing dimensionality.

The resulting per-lead feature maps are then passed to a lead attention module, which assigns an adaptive weight to each lead through a linear transformation and sigmoid activation. This mechanism emphasizes clinically relevant leads and down-weights noisy or less informative signals, improving both robustness and interpretability.

The weighted features from all leads are concatenated and fed into a fully connected classification head. This module consists of dense layers with batch normalization, SiLU activations, and dropout regularization, ensuring effective hierarchical feature learning and reducing overfitting. The final output layer provides a binary prediction indicating the presence or absence of Chagas disease.

Compared to conventional CNNs, the incorporation of lead-aware attention enables the model to handle missing or corrupted leads more effectively while highlighting leads strongly associated with Chagas-specific ECG abnormalities.

2.5. Training Details

The proposed model was trained using the *AdamW optimizer* [10] with an initial learning rate of 0.003 and no additional L1/L2 weight regularization. To stabilize convergence, a One-Cycle learning rate scheduler was employed, with a warm-up fraction of 0.3 and a minimum learning rate of 1×10^{-6}. The scheduler dynamically adjusted the learning rate to improve generalization and avoid premature convergence.

For loss optimization, we adopted a composite loss function combining Focal Loss ($\gamma = 2.5$) weighted at 0.8 with a Top-k True Positive Rate (TPR) objective. This de-

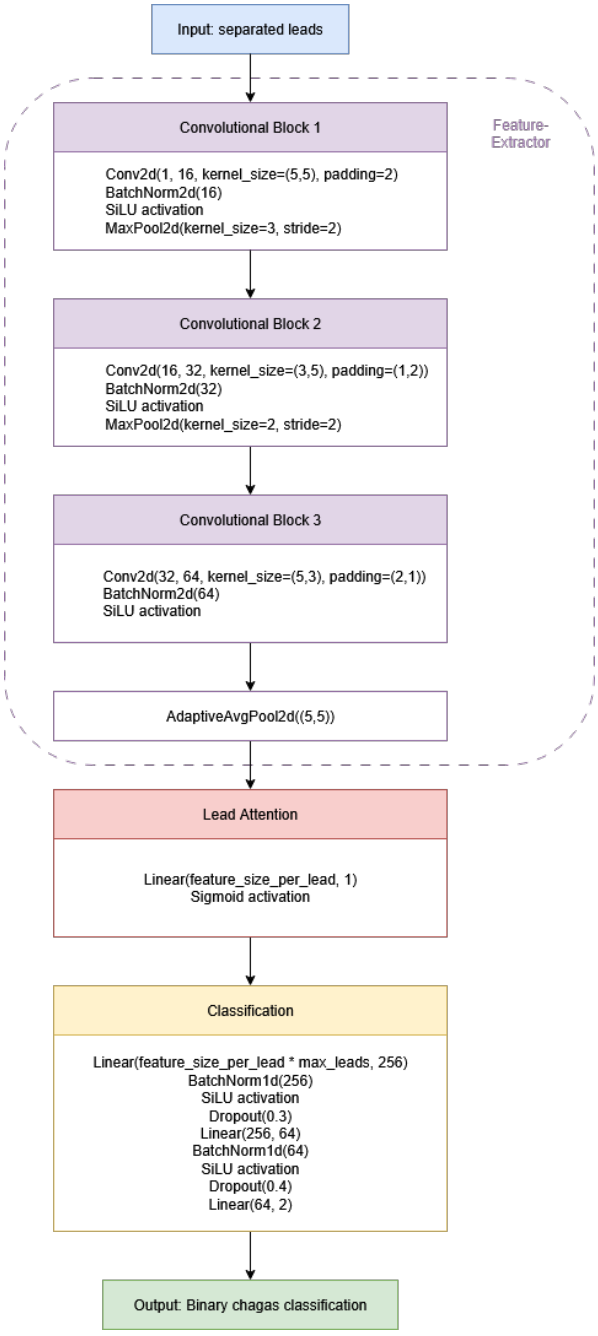

Figure 2. Architecture of the ECG-Spectrogram Classifier with convolutional feature extraction, lead attention, and dense layers for Chagas detection.

sign directly aligns the training process with the Challenge evaluation metric, while also addressing the strong class imbalance inherent in Chagas datasets. A decision threshold of 0.5 was applied for binary classification during inference.

The model was trained over 20 epochs with a batch size of 256, utilizing 80 % of the available data for training and 20 % for validation. All three datasets were used to develop a robust model. A stratified split strategy was employed to preserve the natural distribution of Chagas-positive cases.

To improve signal quality, ECG inputs were preprocessed with a Butterworth bandpass filter (0.5–45 Hz, order 2) and a notch filter at $50/60$ Hz to suppress powerline interference. Gradient clipping at 0.7 was applied to prevent instability from exploding gradients.

2.6. Evaluation Metric

The primary evaluation metric for the Chagas Challenge is the Challenge Score, which assesses algorithms' ability to prioritize patients for serological confirmation based on ECG data. In regions with limited testing capacity, this metric reflects the clinical utility of computational models by measuring their effectiveness in identifying high-risk individuals.

The Challenge Score is defined as the proportion of confirmed Chagas-positive patients ranked within the top 5 % of the patient cohort, emphasizing early and accurate detection to allocate limited tests to the most probable positive cases. Models are evaluated on a hidden test set, with the highest score determining the Challenge winner.

Each algorithm assigns a probability of Chagas disease to every ECG record, and patients are ranked accordingly. The score is computed as the fraction of true positives in the top-ranked subset relative to the total number of positives. Ties are resolved randomly, and the expected score value is reported. This evaluation scheme optimizes algorithms for both classification accuracy and prioritization efficiency, essential for real-world clinical deployment in resource-constrained settings.

3. Results and Discussion

In the PhysioNet Challenge for Chagas disease detection, the proposed model achieved a Challenge Score of 0.303. This score represents a central indicator of the algorithm's ability to identify patients at high risk of Chagas disease based on ECG analysis. Specifically, the model succeeded in ranking a substantial fraction of Chagas-positive patients within the top 5 % of the cohort, which is highly relevant in regions with limited testing capacity, as it enables more efficient allocation of scarce diagnostic resources.

The achieved score demonstrates that the model can contribute to earlier diagnosis and timely treatment, potentially reducing morbidity and mortality among affected populations. The incorporation of a lead-aware attention mechanism proved particularly valuable: by adaptively weighting ECG leads according to their diagnostic relevance, the model exhibited robustness against missing or corrupted leads, thereby enhancing both reliability and interpretability.

Despite these promising results, the obtained score of 0.303 also indicates room for improvement. Misclassifications and performance variability highlight the need for further refinements in data handling and model optimization. Future work should investigate advanced data augmentation strategies, alternative attention mechanisms, and ensemble modeling approaches to increase predictive performance. In addition, expanding the dataset with more diverse patient cohorts could improve generalization across different populations and clinical settings.

In summary, the results highlight the potential clinical utility of the proposed spectrogram-based CNN with lead-aware attention for prioritizing Chagas testing. While the current performance is encouraging, further methodological and data-driven improvements are required to enhance accuracy and support deployment in real-world healthcare environments.

4. Conclusion

This study presented HEART, a hybrid ECG analysis system leveraging spectrogram-based CNNs with a lead-aware attention mechanism for Chagas disease detection. Our approach achieved a Challenge Score of 0.303 while maintaining computational efficiency with only 5 million parameters.

Key contributions include: (1) A novel attention mechanism that adaptively weights ECG leads to ensure robustness against missing or corrupted data; (2) A comprehensive preprocessing pipeline combining STFT-based spectrograms with targeted augmentation strategies; (3) Successful integration of heterogeneous datasets with varying label quality through a unified training framework.

Our results demonstrate that automated ECG analysis can effectively prioritize patients for serological confirmation in resource-constrained settings. The model identified $30.3\ \%$ of Chagas-positive cases within the top $5\ \%$ of rankings, significantly exceeding random selection.

However, there are limitations: reliance on weak labels from the CODE-15 % dataset may have introduced noise, and the current architecture does not explicitly model temporal dependencies between heartbeats, potentially missing subtle rhythm abnormalities.

Future work will focus on: (1) Implementing semi-supervised learning techniques to better leverage weakly labeled data; (2) Integrating unsupervised pre-training approaches to enhance ECG representations; (3) Developing uncertainty quantification methods to identify cases requiring expert review. We also plan to explore multi-task learning frameworks that jointly predict Chagas status and related cardiac conditions to improve diagnostic accuracy and clinical interpretability.

References

[1] World Health Organization. Chagas disease (american trypanosomiasis). https://www.who.int/news-room/fact-sheets/detail/chagas-disease-(american-trypanosomiasis), 2021. Accessed: October 15, 2023.

[2] Centers for Disease Control and Prevention. Chagas disease. https://www.cdc.gov/parasites/chagas/index.html, 2021. Accessed: October 15, 2023.

[3] Rassi AJ, Rassi A, Marin-Neto JA. Chagas disease. The Lancet 2010;375(9723):1388–1402.

[4] Pérez-Molina JA, Molina I. Chagas disease. The Lancet 2018;391(10115):82–94.

[5] Moncayo A, Silveira AC. Current epidemiological trends for chagas disease in the americas. Current Opinion in Infectious Diseases 2009;22(5):420–425.

[6] Nunes MC, et al. Chagas cardiomyopathy: a review. Nature Reviews Cardiology 2018;15(10):683–695.

[7] Ribeiro AH, Paixao GM, Lima EM, Horta Ribeiro M, Pinto Filho MM, Gomes PR, Oliveira DM, Meira Jr W, Schon TB, Ribeiro ALP. Code-15%: a large scale annotated dataset of 12-lead ecgs, June 2021. URL https://doi.org/10.5281/zenodo.4916206.

[8] Ribeiro ALP, Ribeiro AH, Paixao GM, Lima EM, Horta Ribeiro M, Pinto Filho MM, Gomes PR, Oliveira DM, Meira Jr W, Schon TB, Sabino EC. Sami-trop: 12-lead ecg traces with age and mortality annotations, June 2021. URL https://doi.org/10.5281/zenodo.4905618.

[9] Wagner P, Strodthoff N, Bousseljot RD, Kreiseler D, Lunze F, Samek W, Schaeffter T. Ptb-xl, a large publicly available electrocardiography dataset, 2020. URL https://publica.fraunhofer.de/handle/publica/264664.

[10] Loshchilov I, Hutter F. Decoupled weight decay regularization, 2019. URL https://arxiv.org/abs/1711.05101.

Address for correspondence:

Pinar Bisgin
Team DataNinyas
Speicherstraße 6
44147 Dortmund, Germany
pinar.bisgin@isst.fraunhofer.de