

Graph Neural Networks Enhance CTRCD Detection from 12-Lead ECG by Modeling Inter-Lead Relationships: A Preliminary Study

Yifan Liang¹, Natsu Suyama¹, Yuki Ishizuka², Takio Kurita¹, Kazuko Tajiri², and Akira Furui¹

Abstract—Cancer therapy-related cardiac dysfunction (CTRCD) is a serious complication, posing significant risks to cancer survivors’ cardiovascular health. While early detection is crucial, current imaging-based methods are expensive and impractical for routine screening. This study proposes a deep learning framework that combines convolutional neural networks (CNNs) and graph convolutional networks (GCNs) to analyze 12-lead electrocardiography (ECG) data for CTRCD detection. By representing ECG leads as graph nodes and modeling their anatomical and physiological relationships through different adjacency matrices, our approach captures inter-lead dependencies overlooked by conventional methods. Experimental results demonstrate that the physiological function-based graph structure outperforms the conventional CNN approach, particularly in sensitivity and F1-score. Interpretability analysis reveals distinct lead-specific patterns, enhancing clinical understanding.

I. INTRODUCTION

Cancer therapy-related cardiac dysfunction (CTRCD) is a serious complication affecting cancer patients undergoing treatment, potentially leading to treatment interruption or persistent cardiac dysfunction even after cancer remission [1], [2]. Although echocardiography and magnetic resonance imaging (MRI) provide accurate detection, their high cost and specialized equipment requirements limit large-scale screening [3]. The 12-lead electrocardiography (ECG) offers a non-invasive, cost-effective alternative, but conventional ECG interpretation remains heavily dependent on clinical expertise and is limited by subjectivity [4], highlighting the need for automated analytical approaches.

Recent advances in deep learning have revolutionized automated ECG analysis. Convolutional neural networks (CNNs) have demonstrated remarkable performance in various cardiac disease classification tasks, including arrhythmia detection and myocardial infarction diagnosis [5], [6]. These models can automatically extract discriminative features from ECG signals, often achieving accuracy comparable to or exceeding that of experienced cardiologists. However, despite these successes in conventional cardiac conditions, the application of deep learning to CTRCD prediction remains largely unexplored, presenting a significant gap in current research.

Moreover, existing ECG deep learning models typically treat the 12 leads as independent channels, either processing

them separately or simply stacking them as multi-channel images. This approach overlooks the inherent spatial and physiological relationships among leads [7], [8]. In reality, ECG leads represent different projections of the heart’s three-dimensional electrical activity, and their correlations may contain crucial information for detecting subtle CTRCD-related changes. This limitation motivates the use of graph-based approaches that can explicitly model inter-lead relationships.

To address these challenges, we propose a graph convolutional networks (GCNs) framework for CTRCD detection that represents the 12-lead ECG data as a graph structure, where leads serve as nodes and their anatomical and physiological relationships define the edges. This graph-based representation enables our model to capture inter-lead dependencies that may reveal subtle CTRCD-related cardiac changes overlooked by conventional approaches. Furthermore, we incorporate Grad-CAM visualization to identify which ECG leads contribute most to CTRCD prediction, enhancing clinical interpretability.

The main contributions of this work are as follows:

- A GCN-based framework for CTRCD prediction leveraging 12-lead ECG inter-lead relationships;
- Systematic evaluation of different adjacency matrix designs encoding various physiological hypotheses;
- Interpretable visualization through Grad-CAM to reveal lead-specific contributions for clinical insights.

II. METHOD

A. Problem Definition

This study addresses CTRCD prediction as a binary classification problem using graph-based modeling of 12-lead ECG signals. Let $\mathbf{X} \in \mathbb{R}^{12 \times T}$ represent a 12-lead ECG sample with temporal length T . We model the 12 leads as nodes in a graph structure, where edges represent spatial and physiological relationships between leads. The objective is to predict whether a patient is at risk of CTRCD ($y = 1$) or not ($y = 0$) by leveraging both temporal patterns within individual leads and structural relationships across the lead graph.

B. Proposed Model Architecture

We propose a graph-based ECG analysis framework that combines convolutional neural networks (CNNs) and graph convolutional networks (GCNs) to enhance CTRCD prediction accuracy and interpretability. The architecture consists of two main components: a node feature encoder and a GCN module.

This work was supported by the National Cancer Center Research and Development Fund (2023-A-12).

¹Y. Liang, N. Suyama, T. Kurita, and A. Furui are with Graduate School of Advanced Science and Engineering, Hiroshima University, Higashi-Hiroshima, Japan (email: yifanliang@hiroshima-u.ac.jp)

²Y. Ishizuka and K. Tajiri are with Department of Cardiology, National Cancer Center Hospital East, Kashiwa, Japan.




Fig. 1. Overview of the proposed method. The 12-lead ECG signals are processed by a lead feature encoder f_{enc} to extract features $\mathbf{h}_i$, which serve as nodes in the GCN module. The adjacency matrix $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ defines inter-lead relationships, and the GCN output is fed to a CTRCD classifier.

1) *Node Feature Encoder*: The node feature encoder extracts temporal features from each lead independently using 1D-CNN. Let $\mathbf{x}_i \in \mathbb{R}^T$ denote the time-series signal of lead $i = 1, 2, \dots, 12$. Each lead is processed by a feature encoder $f_{\text{enc}} : \mathbb{R}^T \rightarrow \mathbb{R}^d$ to extract a latent feature vector $\mathbf{h}_i \in \mathbb{R}^d$ as follows:

$$\mathbf{h}_i = f_{\text{enc}}(\mathbf{x}_i) \quad (1)$$

The encoder f_{enc} employs 1D convolutional layers followed by pooling operations to capture temporal patterns and compress the signal into a fixed-length representation $\mathbf{h}_i \in \mathbb{R}^d$.

2) *GCN Module*: The GCN module models inter-lead relationships using a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where each of the 12 leads represents a node $v_i \in \mathcal{V}$, and edges $(v_i, v_j) \in \mathcal{E}$ encode spatial and physiological relationships between leads. Each node v_i is associated with a feature vector $\mathbf{h}_i$ obtained from the node feature encoder. The graph structure is defined by an adjacency matrix $\mathbf{A} \in \{0, 1\}^{12 \times 12}$, where $A_{ij} = 1$ indicates a connection between leads i and j .

The initial node feature matrix $\mathbf{H}^{(0)} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{12}]^T \in \mathbb{R}^{12 \times d}$ serves as input to the GCN. The graph convolution operation is computed as:

$$\mathbf{H}^{(l+1)} = \sigma \left(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^{(l)} \mathbf{W}^{(l)} \right), \quad (2)$$

where $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ is the adjacency matrix with self-loops, $\tilde{\mathbf{D}}$ is its degree matrix with $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$, $\mathbf{W}^{(l)}$ is the learnable weight matrix, and $\sigma(\cdot)$ is a nonlinear activation function. After graph convolution layers, the final node representations are aggregated and passed through a classifier for binary prediction.

C. Adjacency Matrix Design Strategies

To investigate how different structural assumptions among ECG leads affect GCN performance, we design four adjacency matrix strategies representing different hypotheses about inter-lead relationships. Fig. 2 illustrates these four graph structures.

1) *Self-loop graph*: This baseline assumes complete independence among leads, with adjacency matrix $\mathbf{A} = \mathbf{I}$. Only self-connections are retained, eliminating inter-lead message passing. This serves as an ablation study to isolate the contribution of graph structure from temporal feature learning.

2) *Fully-connected graph*: This model assumes all leads are potentially correlated, with $A_{ij} = 1$ for all i, j . It provides maximum flexibility for the GCN to learn arbitrary inter-lead dependencies without structural constraints. This serves as an upper bound for graph-based approaches while testing whether unconstrained connectivity improves performance.

3) *Physical proximity graph*: This strategy reflects physical electrode placement on the patient's body, under the hypothesis that spatially adjacent leads capture correlated cardiac activity. The limb leads (I, II, III, aVR, aVL, aVF) form a fully connected subgraph due to their overlapping spatial coverage of the heart's electrical axis. The precordial leads (V1–V6) are connected sequentially (V1–V2, V2–V3, ..., V5–V6) to model the continuous spatial transition across the chest wall from the right sternal border to the left midaxillary line.

4) *Physiological function graph*: This approach incorporates clinical domain knowledge by organizing leads according to their diagnostic roles in assessing specific myocardial regions. The limb leads form a fully connected subgraph due to their global cardiac coverage. The precordial leads are divided into three functional modules: V1–V2 (right ventricle/anterior septum), V3–V4 (mid-septum/anterior wall), and V5–V6 (lateral wall). Within each module, leads are fully connected to capture intra-regional dependencies. Inter-module connections link only boundary leads (V2–V3, V4–V5), while cross-connections between precordial and limb leads (e.g., V1, V2 $\leftrightarrow$ aVR, V5, V6 $\leftrightarrow$ aVL) reflect established clinical associations for comprehensive cardiac assessment [9].




Fig. 2. Four adjacency matrix designs for modeling ECG inter-lead relationships.

III. EXPERIMENT

A. Dataset

The dataset used in this study was provided by the National Cancer Center Hospital in Japan and consists of 12-lead ECG recordings from 616 cancer patients who had undergone anthracycline-based chemotherapy, including 87 males and 529 females. Although multiple ECGs were collected from each patient, only one ECG per individual was used in this study. Specifically, for positive cases, the ECG recorded on the day of CTRCD diagnosis was selected, while for negative cases, one ECG was randomly chosen from the available examinations. After quality control procedures, a total of 557 samples were included in the final dataset, comprising 43 positive and 514 negative cases. This study was approved by the review board of the National Cancer Center (research project no. 2022-152).

All ECG signals were sampled at 500 Hz with an original recording duration of approximately 11 ± 0.2 seconds. To ensure consistency across samples, each signal was truncated to a uniform length of 10 seconds (i.e., $T = 5000$) by removing excess data from the end of the recording [8].

B. Model Implementation

All models shared identical hyperparameter settings to ensure fair comparison. For the node feature encoder, we employed two 1D convolutional layers with kernel sizes of 3 and 5, respectively, followed by batch normalization, ReLU activation, and max pooling. Adaptive average pooling was applied to generate fixed-length feature vectors of dimensionality $d = 64$.

The GCN module consisted of five graph convolution layers with hidden dimensions of 64. Each layer was followed by ReLU activation and dropout with rate 0.3 to prevent overfitting. The final node representations were aggregated using global average pooling and passed through a fully connected layer with 32 hidden units for binary classification.

C. Model Comparison and Baseline

To comprehensively assess the effectiveness of the GCN-based architecture, we established a baseline model using a pure 1D-CNN approach. This baseline model employs the identical 1D-CNN encoder architecture to extract features from each of the 12 leads independently. The extracted feature vectors ($\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{12}$) are concatenated and passed

through a fully connected layer for binary classification. This approach does not consider any inter-lead relationships and serves as the primary comparison to evaluate the benefit of incorporating graph structure in ECG analysis.

To isolate the effect of different structural assumptions, we also compared GCN-based models with various adjacency matrices, including identity (self-loop only), fully connected, physical proximity, and physiological function graphs. This comparison allows us to examine the impact of different graph priors on the proposed GCN framework. All models used identical 1D-CNN encoders to ensure consistent feature extraction capacity, thereby isolating the effect of structural design on model performance.

D. Training Configuration

We adopted a five-fold stratified cross-validation strategy, which maintains the class distribution of CTRCD-positive and CTRCD-negative samples within each fold. Each fold was further divided into a training and a validation set. All models were trained and evaluated on identical data splits to ensure fair and consistent comparisons.

All models were trained using the Adam optimizer with the initial learning rate 0.0001, L2 weight decay 0.0001, and batch size 64. Training was conducted for 175 epochs. To address the class imbalance problem caused by the relatively low proportion of CTRCD-positive samples, we employed the focal loss function, which is particularly well-suited for imbalanced classification tasks. The focal loss parameters are configured as follows: $\alpha = [0.2, 0.8]$, $\gamma = 2$.

E. Evaluation Metrics

To comprehensively evaluate model performance in the CTRCD classification task, we reported standard classification metrics including accuracy, sensitivity, specificity, F1 score, and area under the receiver operating characteristic curve (ROC-AUC). All metrics are calculated on each validation fold and averaged across the five folds to ensure robust evaluation.

IV. RESULTS

A. Performance Evaluation

This study systematically evaluated the performance of GCN models with four different adjacency matrix designs: self-loop only, fully connected, physical proximity-based,

TABLE I
PERFORMANCE COMPARISON OF THE BASELINE CNN AND PROPOSED GCN MODELS WITH DIFFERENT ADJACENCY MATRICES.

Model	Accuracy	Sensitivity	Specificity	F1 score	ROC-AUC
Pure CNN	0.9192	0.5111	0.9533	0.4968	0.8785
CNN + GCN (ours)					
w/ Self-loop graph	0.9121	0.3249	0.9611	0.2824	0.8050
w/ Fully-connected graph	0.9174	0.5638	0.9475	0.5067	0.8686
w/ Physical proximity graph	0.9156	0.5138	0.9493	0.4440	0.8837
w/ Physiological function graph	0.9301	0.6722	0.9514	0.5956	0.9147




Fig. 3. Training loss curves for GCN models with four different adjacency matrix designs.




Fig. 4. Receiver operating characteristic (ROC) curves for the baseline CNN model and the proposed GCN model with physiological function adjacency matrix.

physiological function-based. All variants were compared against a traditional 1D-CNN baseline.

Fig. 3 shows the training loss curves across different adjacency matrix designs. The self-loop graph exhibited the highest loss throughout training, while the fully connected model achieved the lowest loss. The physiological and physical proximity-based graphs demonstrated intermediate loss values.

Table I summarizes the average results across five-fold cross-validation. The GCN model with the physiological function-based adjacency matrix achieved the highest sensitivity and F1 score among all models. The fully-connected and physical proximity-based GCNs also outperformed the self-loop and CNN baselines in most metrics. The self-loop model showed the lowest sensitivity and F1 score.

Compared to the pure 1D-CNN model, the GCN-based models demonstrated improved performance in sensitivity, F1 score, and ROC-AUC. These results indicate that modeling inter-lead relationships can enhance CTRCD prediction. Fig. 4 shows the ROC curves comparing the best-performing GCN model with the CNN baseline, with ROC-AUC improving from 0.8785 to 0.9147.

B. Interpretability Analysis

Grad-CAM was used to visualize the contribution of each ECG lead to CTRCD predictions [10]. For each node i , we computed importance scores by taking the mean of gradient-weighted activations from the final GCN layer across feature dimensions, yielding a single value representing each lead's contribution to the prediction. Fig. 5 show the visualization

results for correctly classified negative samples, with true positive cases in the top row and true negative cases in the bottom row.

The visualizations reveal different attention patterns between classes. For CTRCD-positive samples, leads V1, V2, and aVR exhibit higher attention weights. For CTRCD-negative samples, limb leads (I, II, III, aVL, aVF) along with aVR show greater contributions to the model's predictions.

V. DISCUSSION

This study systematically evaluated the impact of different adjacency matrix designs on GCN model performance for CTRCD prediction. The superior performance of the physiological function-based adjacency matrix demonstrates that incorporating clinical domain knowledge and functional correlations between leads enhances the model's ability to capture pathological ECG features. The substantial improvement in sensitivity (from 0.5111 to 0.6722) is particularly clinically relevant for CTRCD detection.

Grad-CAM visualizations reveal distinct attention patterns between classes. In CTRCD-positive cases, the model emphasizes leads V1, V2, and aVR, particularly V1 and V2, which may reflect detection of conduction abnormalities associated with cardiotoxicity. In contrast, limb leads (I, II, III, aVL, and aVF) contribute more in negative cases. However, interpretation requires caution: in binary classification, models may identify the absence of normal patterns rather than specific pathological features. The observed patterns could reflect disrupted electrical activity rather than CTRCD-specific abnormalities. Notably, leads V3–V6 showed low




Fig. 5. Grad-CAM visualizations showing lead-specific attention weights for correctly classified samples. (a) Three true positive CTRCD cases. (b) Three true negative cases. Node colors indicate the contribution level of each lead to the model's prediction, with red representing high attention and blue representing low attention.

importance despite their role in detecting left ventricular abnormalities typically associated with CTRCD. This finding suggests that CTRCD-related ECG changes may manifest more prominently through conduction abnormalities rather than direct left ventricular signals, highlighting the value of modeling inter-lead relationships. Further clinical investigation is needed to validate these patterns and enhance clinical interpretability.

VI. CONCLUSION

In this study, we proposed a graph-based deep learning framework integrating CNN and GCN for CTRCD detection from 12-lead ECG signals. By explicitly modeling inter-lead relationships through physiological function-based adjacency matrices, our approach significantly outperformed traditional CNN methods, achieving improvements. These results demonstrate that capturing functional correlations between ECG leads is crucial for detecting subtle CTRCD-related cardiac changes. Additionally, Grad-CAM visualizations enhanced clinical interpretability by revealing distinct attention patterns between positive and negative cases.

Future directions include exploring learnable adjacency matrices for data-driven graph discovery beyond predefined structures, potentially incorporating domain expertise to guide the learning process. We also plan to investigate graph attention networks to enhance interpretability through attention mechanisms. Clinical validation with medical professionals will verify whether model-identified relationships align with known CTRCD pathophysiology.

REFERENCES

- [1] M. W. Bloom *et al.*, "Cancer therapy-related cardiac dysfunction and heart failure: part 1: definitions, pathophysiology, risk factors, and imaging," *Circ. Heart Fail.*, vol. 9, no. 1, 2016, Art. no. e002661.
- [2] C. E. Hamo *et al.*, "Cancer therapy-related cardiac dysfunction and heart failure: part 2: prevention, treatment, guidelines, and future directions," *Circ. Heart Fail.*, vol. 9, no. 2, 2016, Art. no. e002843.
- [3] I. E. Perez, S. Taveras Alam, G. A. Hernandez, and R. Sancassani, "Cancer therapy-related cardiac dysfunction: an overview for the clinician," *Clin. Med. Insights: Cardiol.*, vol. 13, 2019, Art. no. 1179546819866445.
- [4] P. M. van Dam, M. Boonstra, E. T. Locati, and P. Loh, "The relation of 12 lead ECG to the cardiac anatomy: the normal CineECG," *J. Electrocardiol.*, vol. 69, pp. 67–74, 2021.
- [5] U. B. Baloglu, M. Talo, O. Yildirim, R. San Tan, and U. R. Acharya, "Classification of myocardial infarction with multi-lead ECG signals and deep CNN," *Pattern Recognit. Lett.*, vol. 122, pp. 23–30, 2019.
- [6] G. Petmezas, K. Haris, L. Stefanopoulos, V. Kilintzis, A. Tzavelis, J. A. Rogers, A. K. Katsaggelos, and N. Maglaveras, "Automated atrial fibrillation detection using a hybrid CNN-LSTM network on imbalanced ECG datasets," *Biomed. Signal Process. Control*, vol. 63, 2021, Art. no. 102194.
- [7] B. Andayeshgar, F. Abdali-Mohammadi, M. Sepahvand, A. Almasi, and N. Salari, "Arrhythmia detection by the graph convolution network and a proposed structure for communication between cardiac leads," *BMC Med. Res. Methodol.*, vol. 24, no. 1, 2024, Art. no. 96.
- [8] N. Suyama, A. Furui, T. Kurita, and K. Tajiri, "Multi-scale feature learning with CNN-RNN-attention framework for ECG-based cancer therapy-related cardiac dysfunction detection," in *Proc. 47th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2025.
- [9] G. S. Wagner and D. G. Strauss, *Marriott's Practical Electrocardiography*, 13th ed. Lippincott Williams & Wilkins (LWW), 2020.
- [10] P. E. Pope, S. Kolouri, M. Rostami, C. E. Martin, and H. Hoffmann, "Explainability methods for graph convolutional neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 10 772–10 781.