IMLE-Net: An Interpretable Multi-level Multi-channel Model for ECG Classification

Likith Reddy[1], Vivek Talwar[1], Shanmukh Alle[2], Raju. S. Bapi[3], U. Deva Priyakumar[2]

[1]International Institute of Information Technology, Healthcare and AI Center, Hyderabad
[2]International Institute of Information Technology, CCNSB Lab, Hyderabad
[3]International Institute of Information Technology, Cognitive Science Lab, Hyderabad
Email: likith.reddy@ihub-data.iiit.ac.in

Abstract— Early detection of cardiovascular diseases is crucial for effective treatment and an electrocardiogram (ECG) is pivotal for diagnosis. The accuracy of Deep Learning based methods for ECG signal classification has progressed in recent years to reach cardiologist-level performance. In clinical settings, a cardiologist makes a diagnosis based on the standard 12-channel ECG recording. Automatic analysis of ECG recordings from a multiple-channel perspective has not been given enough attention, so it is essential to analyze an ECG recording from a multiple-channel perspective. We propose a model that leverages the multiple-channel information available in the standard 12-channel ECG recordings and learns patterns at the beat, rhythm, and channel level. The experimental results show that our model achieved a macro-averaged ROC-AUC score of 0.9216, mean accuracy of 88.85% and a maximum F1 score of 0.8057 on the PTB-XL dataset. The attention visualization results from the interpretable model are compared against the cardiologist's guidelines to validate the correctness and usability.

I. INTRODUCTION

Heart diseases are the leading cause of death globally. It is estimated that 17.9 million people died of heart diseases in the year 2016 which represents about 31% of global human mortality according to the World Health Organization. Early detection of heart disease is crucial for effective treatment and reducing mortality. Electrocardiography is an inexpensive, simple and non-invasive procedure that helps us to understand the functioning of a heart which in turn helps us diagnose heart diseases. An Electrocardiogram (ECG) records information about the electrical activity of the heart based on which a cardiologist can identify abnormal functioning of the heart to diagnose various conditions. But this process of analyzing an ECG recording is time-consuming that requires a trained expert's time and attention. Moreover, it is prone to human error. Misdiagnosis of early signs of heart diseases by cardiologists from an ECG recording is still a major concern [1]. So work has been done to automate the ECG signal classification to aid in human diagnosis.

Most existing approaches use traditional machine learning-based methods for the detection of heart diseases from a patient's ECG recording [2, 3]. Such methods involve manually crafted feature extraction and the use of classifiers on the extracted features. Asgari et al. [4] use the stationary wavelet transform to extract features and a support vector machine (SVM) for the detection of atrial fibrillation. Poddar et al. [5] employed the heart rate variability (HRV) features computed in time-domain, frequency-domain and using non-linear methods, the features are then classified using k-nearest neighbor (KNN) and SVM classifiers. Unsupervised methods using time series clustering were proposed in Annam et al. [6].

The advent of deep learning-based methods for ECG signal classification has progressed in recent years to reach cardiologist-level performance [7]. Such methods analyze ECG data without manual feature extraction in contrast to traditional machine learning approaches. The work done by Rajpurkar et al. [7] uses a time series ECG signal as an input to a 34-layer eural network (CNN) by mapping a sequence of ECG samples to a sequence of rhythm classes. The model exceeded the average cardiologist performance in both precision and recall. Mousavi et al. [8] developed an automatic heartbeat classification method using deep convolutional neural networks and sequence-to-sequence models. Murugesan et al. [9] proposed a combination of CNNs and a long short-term memory (LSTM) based feature extractor that can be directly trained without any preprocessing. Recurrent neural network (RNN) along with an attention mechanism is used by Schwab et al. [10] for the classification of a single-channel ECG recording.

An interpretable model along with reliable performance is crucial to instil confidence in medical practitioners to use computer-assisted electrocardiography. A lot of work has been done to improve the interpretability of deep learning models in computer vision tasks but the progress made on interpreting ECG classification models has been limited and has been gaining interest recently [11, 12]. Vijayarangan et al. [13] use a gradient-weighted class activation map [14] for visualizing saliency on single-channel ECG signals using their proposed model. Attention models have been utilized by Hong et al. [15] to improve interpretability in the time and frequency domains.

The majority of the interpretable deep learning work focuses on models which deal with single-channel ECG recordings. A single-channel ECG recording is used for basic heart monitoring whereas a multi-channel ECG recording provides information about the 3-dimensional electrical activity of the heart. From a clinical point of view, a cardiologist makes a diagnosis based on the standard 12-channel ECG recording [16]. When compared to a single-channel recording, a 12-channel recording accurately reproduces various features of

an ECG such as QRS complex, ST-segment and T waveforms which sometimes are poorly represented in a single-channel recording [17]. For example, ST-segment elevation which is an important characteristic for Myocardial Infarction (MI) is best identified using a 12-channel recording and the localization of myocardial injury cannot be identified without a 12-channel recording. Analysis of ECG recordings from a multiple-channel perspective has not been given enough attention and it is essential to analyze an ECG recording from a multi-channel perspective. To address this gap, we propose a model that leverages the multi-channel information available in a multi-channel ECG recording. Also, we make use of the hierarchical structure of the ECG signal to learn patterns at the beat, rhythm and channel level. The model is then evaluated on the PTB-XL dataset [18] which comprises of a total 21837 clinical 12-channel ECG recordings of 10 seconds length from 18885 patients. The performance of the model is then compared against the existing deep learning methods. The interpretability of our model is evaluated by comparing the attention visualizations with the cardiologist's rules for identifying the subtypes of MI. We perform an ablation study by reducing the number of available channels in an input ECG recording to evaluate the robustness of the proposed model. The source code of the proposed model and all the experiments performed is made publicly available to ensure reproducibility for future research[1].

II. METHODOLOGY

A. Preprocessing

In preprocessing, a multi-channel ECG signal is segmented into beat segments of length W for every individual channel. A popular approach to segment an ECG signal into beat segments is based on first detecting the R-peaks in an ECG signal using the Pan-Tompkins algorithm [19] and then using the adaptive searching window for segmenting beats which is used in Mousavi et al. [20]. Other approaches for the detection of R-peaks include standard wavelet transform [21] and Hamilton [22]. But these approaches perform poorly in the presence of signal noise and HRV irregularities. Instead, we opted for a sliding window approach with no overlap which was used by Hong et al. [15] to segment beats of the ECG signal. To get the k^{th} beat segment, the beat is spanned from $(k-1) \times W$ to $k \times W$ over the ECG signal, where W is the length of the window.

B. Architecture

The proposed model can be divided into two parts, the first part processes each channel separately to generate a channel encoding for an input ECG recording. The architecture of the first part is shown in Figure 1. The segmented beats from a multi-channel ECG recording are passed through a beat level block, the outputs from the beat level block are passed through a rhythm level block to get channel-wise encodings. These channel-wise encodings are passed into the second part of the model which pools these encodings from each channel and generates predictions as shown in Figure 2.
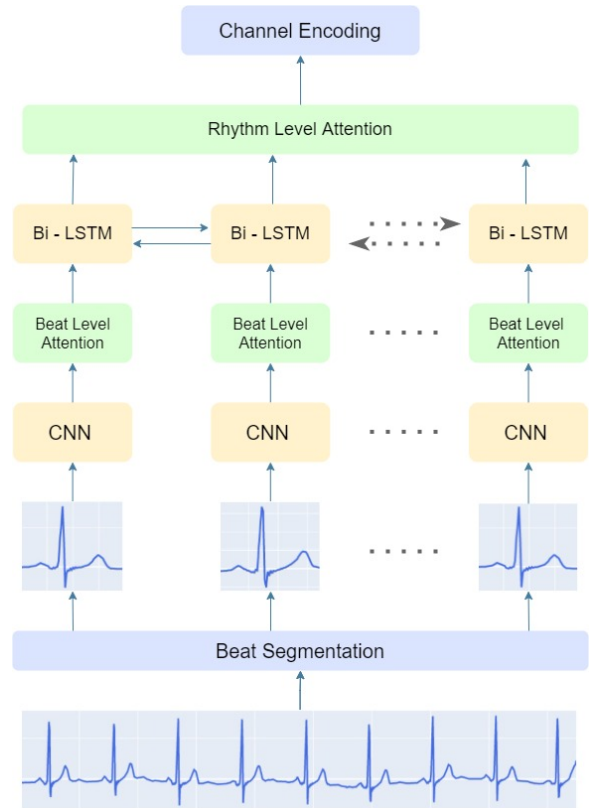
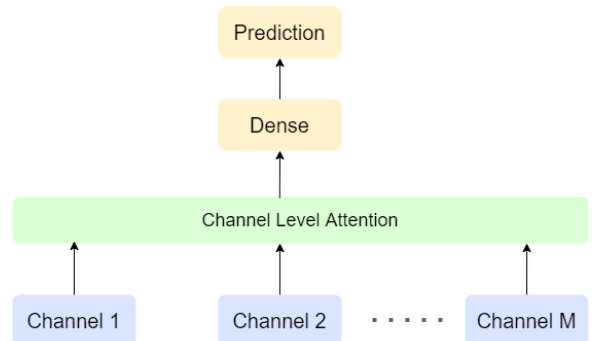Fig. 1. Architecture of the model at each individual channel

Fig. 2. Channel level architecture of the model

1) Beat Level Block: The beat level block involves the processing of input beat segments. The beats are first passed through a CNN and then an attention layer.

The architecture of the CNN used in the beat level block is shown in Figure 3. Inspired by the ResNet architecture [23], the CNN uses a total of 6 residual blocks. Skip connections in a residual block improve backpropagation in deep neural networks which optimize model training. In a single residual block, there are two convolutional layers and every convolutional layer is followed by a batch normalization layer and a ReLU activation function. Between two convolutional layers, we employ dropout to stabilise model training and to avoid overfitting. No downsampling is done in the first two residual blocks. In the remaining residual blocks, inputs

are downsampled by a factor of two after every alternate residual block. All the convolutional layers have a filter length of 8 with 32 starting filters. In every residual block where downsampling of inputs occurs, the number of filters is doubled. All the weights in a CNN are shared across all the beat segments and across multiple channels in the model.

Let T be the length of an ECG signal for the c^{th} channel and W be the beat length. The total number of beat segments are given by $N = \frac{T}{W}$. Let $(b_1^c, b_2^c, b_3^c, ... b_N^c)$ be the beat segments for the c^{th} channel of an ECG recording. The CNN takes in a time series of k^{th} beat segment b_k^c of length W. The output from the CNN with input as b_k^c beat segment be $\bar{b}_k^c$ with an intermediate sequence length of $\bar{W}$.

We know that all the areas within a heartbeat segment are not equally relevant in identifying an abnormal heartbeat. The attention mechanism helps in identifying more important areas within a beat segment by giving higher attention scores to them. It is a two-layer neural network that takes inputs from the CNN outputs and gets the attention scores by computing softmax over the outputs from the two-layer neural network as given in eqs. (1) and (2). The weights of the beat level attention layer are shared across all beat segments similar to CNN.

$$\alpha_k^c = softmax(V_b \tanh(W_b \bar{b}_k^c + b_b)) \tag{1}$$

$$B_k^c = \sum_{i=1}^{\bar{W}} \alpha_{k,i}^c \bar{b}_{k,i}^c \tag{2}$$

Where, V_b, W_b and b_b are beat level attention parameters that are learned during training. α_k^c is the beat level attention score for the k^{th} beat segment in an ECG signal for the c^{th} channel. B_k^c is the beat context vector for the k^{th} beat segment in an ECG signal for the c^{th} channel. The beat context vector B_k^c is a weighted sum of $\bar{b}_{k,i}^c$ with weights given by attention scores $\alpha_{k,i}^c$.

2) Rhythm Level Block: The rhythm level block deals with the input ECG rhythm signal which is constructed from several beat segments. It consists of two layers, a Bi-directional LSTM and a rhythm level attention layer. Bi-directional LSTM is used as the current state can process both forward and backward directions of the input sequence information. The input to the Bi-LSTM is a sequence of beat context vectors B_k^c for the c^{th} ECG channel with N beat sequences. In the Bi-LSTM, the forward and the backward direction hidden states for the current step k are concatenated to obtain output r_k^c.

The output from the Bi-LSTM r^c with a sequence length of N for the c^{th} channel is passed through a rhythm attention layer. The rhythm attention mechanism used here identifies the important beat segments within an ECG rhythm by giving higher attention scores to them and is given by eqs. (3) and (4).

$$\beta^c = softmax(V_r \tanh(W_r r^c + b_r)) \tag{3}$$

$$R^c = \sum_{i=1}^{N} \beta_i^c r_i^c \tag{4}$$
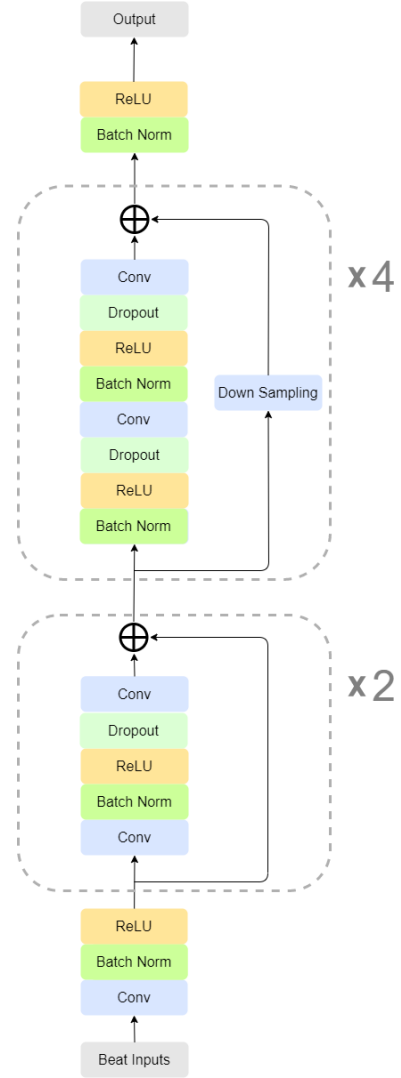

Fig. 3.   Architecture of the CNN used at the beat level.

Where, V_r, W_r and b_r are rhythm level attention parameters that are learned through training and β^c is the rhythm level attention score for the c^{th} channel. R^c is the rhythm context vector for the c^{th} channel which encodes entire ECG signal for a particular channel into a vector (Channel Encoding).

3) Channel Level Block: The channel level block processes the information present in several ECG channels. The information present in individual ECG channels differ. The characteristic information for diagnosing a particular condition might be present only in few channels. So we need a mechanism to not treat each channel equally always and give more importance to the channels that hold the characteristics for a particular heart disease. So a channel-level attention mechanism is used here which receives inputs $(R^1, R^2, ... R^c ... R^M)$, where M is the total number of channels available in an ECG recording. The channel level attention scores are calculated similar to beat and rhythm level attention given by eqs. (5) and (6).

$$\gamma = softmax(V_c \tanh(W_c R^c + b_c)) \qquad (5)$$

$$C = \sum_{i=1}^{M} \gamma_i R_i^c \qquad (6)$$

Where, V_c, W_c and b_c are the channel level parameters that are learned during training. γ is the channel level attention score and C is a context vector that encodes the entire ECG signal for all the available channels. The context vector C is passed through a linear neural network for final predictions.

III. EXPERIMENTS

A. Dataset

The dataset used is the PTB-XL dataset [18] which is the largest openly available dataset that provides clinical 12 channel ECG waveforms. It comprises 21837 ECG records from 18885 patients of 10 seconds length which follow the standard set of channels (I, II, III, aVL, aVR, aVF, V1–V6). The dataset is balanced concerning sex with 52% male and 48% female and covers age ranging from 0 to 95 years. The dataset covers a wide range of pathologies with many different co-occurring diseases. The ECG waveform records are annotated by two certified cardiologists. Each ECG record has labels assigned out of a set of 71 different statements conforming to the Standard communications protocol for computer assisted electrocardiography (SCP-ECG) standard. The ECG waveform was originally recorded at a sampling rate of 400 Hz and downsampled to 100 Hz. All the experiments in our work were performed using the 100 Hz sampling rate.

B. Evaluation

We select the samples from the PTB-XL dataset [18] in a way to cover a huge cohort of diseases and use it to evaluate the model. Each SCP statement annotated in the ECG recording belongs to one of the groups of form, rhythm, and diagnostic statement groups, with a few SCP statement annotations belonging to both rhythm and diagnostic statements. For evaluating the model in our work only the diagnostic statements are considered, a total of 21430 ECG recordings are present in the dataset with at least one diagnostic statement. A specific ECG recording can have multiple diagnostic statements annotated against it. For the diagnostic statement labels, hierarchical groups are provided in terms of 5 superclasses and 24 subclasses. In our experiments, we have used the 5 superclass labels in which the ECG recordings in a particular superclass belong to the same family of pathologies. The superclass disease groups are Normal ECG (NORM), Conduction Disturbance (CD), Myocardial Infarction (MI), Hypertrophy (HYP), and ST/T changes (STTC). Splitting of the dataset for training and testing is based on the stratified folds provided by the dataset. Each ECG recording is assigned to one of the ten folds, folds 1 to 8 are used for training the model, the 9th fold is used as the validation set and the 10th fold is used

as a test set. The train and test sets are standardized before training.

The proposed model is compared against Resnet 101 model [23], Mousavi et al. [8], ECGNet [9] and Rajpurkar et al. [7]. The performance of the model is evaluated using metrics such as the mean accuracy, macro averaged Area under Receiver Operating Characteristics (ROC-AUC) and maximum F1 score. Along with these metrics, class-wise accuracy and class-wise ROC-AUC scores are also reported.

C. Interpretability

We used the beat attention scores, rhythm attention scores and channel attention scores to gain insights into the model at beat, rhythm and channel level.

As the training samples prepared in evaluating the model are grouped according to the same family of pathologies, such training data is not ideal to validate the channel level interpretations as the grouped ECG recordings do not hold many variations in abnormal beat characteristics across different ECG channels. For validating channel level interpretations, we chose two specific sub-diagnostic diseases of MI, namely, Anteroseptal Myocardial Infarction (ASMI) and Inferior Myocardial Infarction (IMI) along with normal (NORM) ECG recordings and rejected the rest of the subtypes of MI mainly due to the paucity of the number of ECG recordings for these subtypes. A total of 1344 (ASMI), 916 (IMI), and 1000 (NORM) samples are obtained. Normal ECG samples are obtained by randomly selecting 100 samples from each fold. We have chosen the stratified folds from 1 to 8 as training data and folds 9 and 10 as test data.

D. Ablation Study

We experimented our model with different combinations of ECG recording channels (leads), by doing so we can estimate how well these combinations of channels contribute to the classification of superclass labeled data which was used in the evaluation of the model. The channel combinations selected are limb leads (I, II, III, aVL, aVR and aVF), precordial leads (V1-V6), leads I, II and II and all the available 12 leads in the ECG recording. The performance of our proposed model is compared against different combinations of ECG channels and class-wise ROC-AUC and class-wise accuracy scores are reported.

E. Implementation Details

The model is trained for a maximum of 60 epochs in mini-batches of 32 with an early stopping criterion. The Cross-Entropy loss function has been used for training the model with Adam optimizer [24] with the first momentum value set to 0.9 and the second momentum value set to 0.99. An adaptive learning rate is used for minimizing the loss with an initial learning rate set to 0.001. The learning rate is decreased by a factor of 10 after every 10 epochs if no improvement in performance is observed. Training is done with window size (W) of 50 time-points and a dropout rate of 0.5. An L2 regularisation loss with a coefficient of ($2 \times 1e$-5)

has been used to avoid overfitting of the model. The model is developed using Tensorflow and all the experiments are trained on an Nvidia RTX 2080Ti 11GB GPU workstation.

IV. RESULTS

A comparison of class-wise ROC-AUC and class-wise accuracy scores across the models is shown in Tables I and II, respectively. Our model achieves the best class-wise ROC-AUC score across all the classes except for the HYP class and the class-wise accuracy score is highest for all classes except for the CD class. Table III shows that our model achieved the best performance across the metrics macro ROC-AUC, mean accuracy and maximum F1 score as compared to the other models. This is due to the effective use of the multi-channel information used by our model compared to other models by incorporating a channel level attention. The processing of the hierarchical structure of an ECG signal is efficiently used by using beat and rhythm level attention. The model parameters are shared across multiple channels instead of having separate model parameters for each channel resulting in fewer final model parameters and better generalization.

TABLE I

CLASS-WISE ROC-AUC COMPARISON

	CD	HYP	MI	NORM	STTC
Resnet101 [23]	0.8928	0.8636	0.8921	0.9098	0.9178
Mousavi et al. [8]	0.8555	0.8093	0.8708	0.9096	0.8817
ECGNet [9]	0.9102	0.8842	0.9021	0.9336	0.9205
Rajpurkar et al. [7]	0.8910	**0.9185**	0.9066	0.9309	0.9306
IMLE-Net	**0.9196**	0.8834	**0.9244**	**0.9450**	**0.9354**

TABLE II

CLASS-WISE ACCURACY COMPARISON

	CD	HYP	MI	NORM	STTC
Resnet101 [23]	87.74	90.89	83.12	84.69	87.47
Mousavi et al. [8]	86.26	87.19	82.47	81.78	83.26
ECGNet [9]	88.71	90.66	84.60	85.89	86.87
Rajpurkar et al. [7]	**90.01**	88.76	85.94	85.94	88.90
IMLE-Net	88.53	**91.58**	**87.97**	**87.10**	**89.08**

TABLE III

OVERALL PERFORMANCE COMPARISON

	Macro ROC-AUC	Mean Accuracy	F1 score(Max)
Resnet101 [23]	0.8952	86.78	0.7558
Mousavi et al. [8]	0.8654	84.19	0.7315
ECGNet [9]	0.9101	87.35	0.7712
Rajpurkar et al. [7]	0.9155	87.91	0.7895
IMLE-Net	**0.9216**	**88.85**	**0.8057**

The performance metrics for the selected subtypes of MI are shown in Table IV. The metrics shown in this table are class-wise accuracy and class-wise ROC-AUC scores along with mean accuracy and macro ROC-AUC score. The best

TABLE IV

ACCURACY AND ROC-AUC METRICS FOR TWO SUBTYPES OF MYOCARDIAL INFARCTION

	ASMI	IMI	NORM	Overall
Accuracy	89.47	87.24	94.09	90.27
ROC-AUC	0.9594	0.9336	0.9906	0.9612

performance is observed for the NORM class followed by the ASMI and IMI classes.

The beat level attention scores, the rhythm level attention scores and the channel level attention scores are visualized for an ECG recording in Figure 4 with Figure 5 describing the importance of a specific channel responsible for the classification. The patient in the ECG recording shown in Figure 4 has an inferior myocardial infarction (IMI) present in leads II, III and aVF. Our proposed model was able to identify the channels responsible for this particular subtype of MI in Figure 5. These findings have been validated by an independent clinician. The channel level attention scores assess the relevance of specific channels in an ECG recording. This setting is crucial as the abnormal beat characteristics might be spread only over a few of the channels and may be completely absent in the rest of the channels. So along with the beat level, rhythm level interpretations, it is also essential to capture such channel level interpretations.

TABLE V

ACCURACY SCORE COMPARISON ON DIFFERENT ECG LEAD CONFIGURATIONS (ABLATION STUDY)

	CD	HYP	MI	NORM	STTC	Mean Accuracy
Limb Leads	87.97	89.59	83.63	84.48	85.85	86.38
Precordial Leads	87.47	90.98	84.46	84.97	88.48	87.27
Leads I,II,III	87.51	88.95	83.40	85.25	85.85	86.19
12 Leads	**88.53**	**91.58**	**87.97**	**87.10**	**89.08**	**88.85**

TABLE VI

ROC-AUC SCORE COMPARISON ON DIFFERENT ECG LEAD CONFIGURATIONS (ABLATION STUDY)

	CD	HYP	MI	NORM	STTC	Macro ROC-AUC
Limb Leads	0.8725	0.8407	0.8787	0.9292	0.9120	0.8866
Precordial Leads	0.8972	0.8627	0.8683	0.9290	0.9261	0.8966
Leads I,II,III	0.8755	0.8447	0.8609	0.9259	0.9106	0.8836
12 Leads	**0.9196**	**0.8834**	**0.9244**	**0.9450**	**0.9354**	**0.9216**

Tables V and VI report the class-wise accuracy and class-wise ROC-AUC scores for several different channel combinations as part of the ablation studies. It can be seen that the performance metrics when utilizing all the available 12

Fig. 4.   Visualization of normalized attention scores with red having a higher attention score and yellow having a lower attention score


Fig. 5.   Channel Importance scores

channels in the ECG recording are the highest for each class and overall when compared against different combinations of channels. The precordial leads gave the second-best overall performance when compared against limb leads and leads I, II and III. Further, disorder classes such as HYP and STTC were classified with better metrics. Limb leads and leads I, II and III gave similar performance metrics with marginally better performance on the limb leads. In short, the standard 12-lead ECG setting gives the best performance and the performance drops as the leads are removed due to loss of information in the channels, again emphasizing the importance of multi-channel information in diagnosing heart conditions.

V. CONCLUSION

In this study, we proposed an interpretable multi-level model for multi-channel ECG classification. The proposed model has been evaluated on the PTB-XL dataset and it outperforms several of the existing models. An ablation study was performed to understand the changes in the performance of the model with different combinations of ECG channels.

The attention scores generated were visualized and compared with guidelines that cardiologists follow to diagnose two subtypes of myocardial infarction. The interpretability aspect of the proposed model enables a better understanding of how it can be useful in clinical decision-making in a real-life scenario. In future, the model interpretability can be evaluated against a wide range of heart pathologies.

REFERENCES

[1]   McCabe JM et al. "Prevalence and Factors Associated With False-Positive ST-Segment Elevation Myocardial Infarction Diagnoses at Primary Percutaneous Coronary Intervention–Capable Centers". In: *Arch Intern Med* 172.11 (2012), pp. 864–871.

[2]   Mohamed Boussaa et al. "ECG signals classification using MFCC coefficients and ANN classifier". In: *2016 International Conference on Electrical and Information Technologies (ICEIT)*. 2016, pp. 480–484.

[3] Ataollah Ebrahim Zadeh et al. "Classification of the electrocardiogram signals using supervised classifiers and efficient features". In: *Computer Methods and Programs in Biomedicine* 99.2 (2010), pp. 179–194.

[4] Shadnaz Asgari, Alireza Mehrnia, and Maryam Moussavi. "Automatic detection of atrial fibrillation using stationary wavelet transform and support vector machine". In: *Computers in biology and medicine* 60 (2015), pp. 132–142.

[5] M. G. Poddar, Vinod Kumar, and Yash Paul Sharma. "Heart rate variability: Analysis and classification of healthy subjects for different age groups". In: *2015 2nd International Conference on Computing for Sustainable Global Development (INDIACom)*. 2015, pp. 1908–1913.

[6] Jagadeeswara Rao Annam and Bapi Raju Surampudi. "AAMI Based ECG Heart-Beat Time-Series Clustering Using Unsupervised ELM and Decision Rule". In: *ICIT*. IEEE, 2016, pp. 137–141.

[7] Masoumeh Haghpanahi Awni Y Hannun Pranav Rajpurkar. "Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network". In: *Nature Medicine* 25 (2019), pp. 65–69.

[8] Sajad Mousavi and Fatemeh Afghah. "Inter- and Intra-Patient ECG Heartbeat Classification for Arrhythmia Detection: A Sequence to Sequence Deep Learning Approach". In: *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2019, pp. 1308–1312.

[9] Balamurali Murugesan et al. "ECGNet: Deep Network for Arrhythmia Classification". In: *2018 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*. 2018, pp. 1–6.

[10] Patrick Schwab et al. "Beat by beat: Classifying cardiac arrhythmias with recurrent neural networks". In: *2017 Computing in Cardiology (CinC)*. 2017, pp. 1–4.

[11] Sajad Mousavi et al. "ECGNET: Learning Where to Attend for Detection of Atrial Fibrillation with Deep Visual Attention". In: *2019 IEEE EMBS International Conference on Biomedical Health Informatics (BHI)*. 2019, pp. 1–4.

[12] Yong-Yeon Jo et al. "Explainable artificial intelligence to detect atrial fibrillation using electrocardiogram". In: *International Journal of Cardiology* 328 (2021), pp. 104–110.

[13] Sricharan Vijayarangan et al. "Interpreting Deep Neural Networks for Single-Lead ECG Arrhythmia Classification". In: *2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*. 2020, pp. 300–303.

[14] Ramprasaath R. Selvaraju et al. "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization". In: *2017 IEEE International Conference on Computer Vision (ICCV)*. 2017, pp. 618–626.

[15] Shenda Hong et al. "MINA: Multilevel Knowledge-Guided Attention for Modeling Electrocardiography Signals". In: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, July 2019, pp. 5888–5894.

[16] Drew BJ and Finlay DD. "Standardization of reduced and optimal lead sets for continuous electrocardiogram monitoring: where do we stand?" In: *Journal of Electrocardiology* 41 (2008), pp. 458–465.

[17] Paul Kligfield. "Value and limitations of 12-lead ECG monitoring". In: *Datex–Ohmeda Clin Window Web J* 6 (2001), pp. 1–6.

[18] Patrick Wagner et al. "PTB-XL, a large publicly available electrocardiography dataset". In: *Scientific Data* 7.1 (2020), p. 154.

[19] Jiapu Pan and Willis J. Tompkins. "A Real-Time QRS Detection Algorithm". In: *IEEE Transactions on Biomedical Engineering* BME-32.3 (1985), pp. 230–236.

[20] Sajad Mousavi, Fatemeh Afghah, and U Rajendra Acharya. "HAN-ECG: An interpretable atrial fibrillation detection model using hierarchical attention networks". In: *Computers in Biology and Medicine* 127 (2020), p. 104057.

[21] T. Abdelmalik M. Merah and B. Larbi. "R-peaks detection based on stationary wavelet transform". In: *Computer methods and programs in biomedicine* 121.3 (2015), pp. 149–160.

[22] P. Hamilton. "Open source ECG analysis". In: *Computers in Cardiology*. 2002, pp. 101–104.

[23] Kaiming He et al. "Deep Residual Learning for Image Recognition". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2016.

[24] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: `1412 . 6980 [cs.LG]`.