

Large Language Model-Informed ECG Dual Attention Network for Heart Failure Risk Prediction

Chen Chen , Lei Li , Marcel Beetz , Abhirup Banerjee , *Member, IEEE*, Ramneek Gupta ,
and Vicente Grau 

Abstract—Heart failure (HF) poses a significant public health challenge, with a rising global mortality rate. Early detection and prevention of HF could significantly reduce its impact. We introduce a novel methodology for predicting HF risk using 12-lead electrocardiograms (ECGs). We present a novel, lightweight dual attention ECG network designed to capture complex ECG features essential for early HF risk prediction, despite the notable imbalance between low and high-risk groups. This network incorporates a cross-lead attention module and 12 lead-specific temporal attention modules, focusing on cross-lead interactions and each lead's local dynamics. To further alleviate model overfitting, we leverage a large language model (LLM) with a public ECG-Report dataset for pretraining on an ECG-Report alignment task. The network is then fine-tuned for HF risk prediction using two specific cohorts from the U.K. Biobank study, focusing on patients with hypertension (UKB-HYP) and those who have had a myocardial infarction (UKB-MI). The results reveal that LLM-informed pre-training substantially enhances HF risk prediction in these cohorts. The dual attention design not only improves interpretability but also predictive accuracy, outperforming existing competitive methods with C-index scores of 0.6349 for UKB-HYP and 0.5805 for UKB-MI. This demonstrates our method's potential in advancing HF risk assessment with clinical complex ECG data.

Index Terms—Large language model, multi-modal learning, heart failure, risk prediction, interpretable artificial intelligence, electrocardiogram.

I. INTRODUCTION

HEART failure (HF) is a complex cardiovascular syndrome, where the heart fails to pump sufficient blood to meet the body's demands. Common causes of HF are cardiac structural and/or functional abnormalities, including heart attack, cardiomyopathy, and high blood pressure. HF is a chronic and progressive disease. In England, admissions due to HF have notably escalated, witnessing an increment from 65,025 in the year 2013/14 to 86,474 in 2018/19, representing a 33% surge, as reported by the British Heart Foundation [1]. It has been found that around 50% of deaths in HF patients presented with a sudden and unexpected pattern [2], [3], which leaves a tremendous burden on patients with HF, their families, and healthcare systems worldwide.

Preventing HF early is crucial to reduce its health and economic impacts, yet HF diagnosis often occurs late when patients have already developed serious symptoms [4], [5]. A promising strategy for improving HF management is the development of risk prediction models for future HF events. These models generate a risk score for patients over a specific timeframe, taking into consideration of the patient's specific characteristics. With such a personalized risk assessment, more tailored HF management strategies and/or treatment recommendations can be provided. In this context, the low-cost 12-lead electrocardiogram (ECG), a medical test commonly used in clinical practice, serves as a valuable resource for evaluating a patient's cardiovascular health and uncovering the risk. Recent studies have already found that several markers detected from clinically acquired 12-lead ECG are associated with future HF events, such as prolonged QRS duration [6], [7], [8], [9], [10], conduction disorders (left and right bundle-branch blocks) [11], [12], etc. A significant limitation of these previous research lies in its reliance on a limited set of biomarkers (e.g., QRS duration), which are identified through predetermined rules based on ECG data. Moreover, much of this research has employed simple linear models for modeling the risk ratio associated with HF. While linear models offer a straightforward and interpretable framework for risk assessment, they may fail to capture the intricate, complex subtleties and nuances embedded within the ECG for early-stage HF risk prediction.

Received 30 November 2023; revised 22 March 2024; accepted 20 June 2024. Date of publication 30 January 2025; date of current version 15 May 2025. The work of Abhirup Banerjee was supported by Royal Society under Grant UR/R1221314. The work of Abhirup Banerjee and Vicente Grau was supported by the British Heart Foundation Project under Grant PG/20/21/35082, in part by the CompBioMed 2 Centre of Excellence in Computational Biomedicine, and in part by European Commission Horizon 2020 Research and Innovation Programme under Grant 823712. This work was supported in part by the Novo Nordisk collaborative research fund and in part by U.K. Biobank Resource under Grant 40161. Recommended for acceptance by Special Issue On Large Language Models, Cross-Modal Generative AI. (*Corresponding author: Chen Chen.*)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by PTB-XL: The Institutional Ethics Committee approved the publication of the anonymous data in an open-access database (PTB-2020-1). <https://physionet.org/content/ptb-xl/1.0.3/> UKB: The UK Biobank data has approval from the North West Research Ethics Committee (REC reference: 11/NW/0382).

Chen Chen is with the Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, OX1 2JD Oxford, U.K., also with Imperial College London, SW7 2AZ London, U.K., and also with the University of Sheffield, S10 2TN Sheffield, U.K. (e-mail: work.cherise@gmail.com).

Lei Li is with the Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, OX1 2JD Oxford, U.K., and also with the University of Southampton, Southampton SO17 1BJ, U.K.

Marcel Beetz, Abhirup Banerjee, and Vicente Grau are with the Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford OX1 2JD, U.K.

Ramneek Gupta is with Novo Nordisk Research Centre Oxford (NNRCO), Oxford OX3 7FZ, U.K.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TBDATA.2025.3536922>, provided by the authors.

Digital Object Identifier 10.1109/TBDATA.2025.3536922

In recent years, deep learning-based approaches with neural networks have shown great capacities to automatically extract features from raw data and utilize them to perform a wide range of tasks. In the field of ECG analysis, deep neural networks have shown competitive performance on a wide range of tasks, including disease classification, waveform prediction, rhythm detection, mortality prediction, and automated report generation, with higher accuracy compared to traditional approaches with handcrafted features [13], [14], [15]. Yet, two prevalent limitations associated with deep learning-based methods are their over-fitting issues as well as poor interpretability. Overfitting occurs when a model directly memorizes the training data instead of learning task relevant features, reducing its ability to perform well on new, unseen data. The large number of parameters and layers in deep neural networks can further exacerbate the risk of overfitting, especially when the available training data is limited. Additionally, the complex architecture of these networks makes it difficult to trace and understand the decision-making process behind their predictions, leading to challenges in interpretability.

In this work, we aim to develop a deep learning-based HF risk prediction model with improved feature learning, higher data efficiency, and explainability. To this end, we design a novel, lightweight ECG dual attention network. This network is capable of capturing intricate cross-lead interactions and local temporal dynamics within each lead. The dual attention mechanism also enables the visualization of lead-wise attention maps and temporal activation across each lead for improved explainability of neural network behaviors. To further alleviate the over-fitting with improved data efficiency and more importantly, to learn clinically useful representations from ECG data for higher precision, we adopt a two-stage training scheme. We first employ a large language model (LLM) to pre-train the network using a large public ECG-Report dataset covering a wide spectrum of diverse diseases and then finetune the network for the HF risk prediction task on two specific cohorts with specific risk factors collected from the U.K. Biobank study. Specifically, we employ ClinicalBERT [16] to extract text embeddings from ECG report and force the extracted ECG features to be aligned with corresponding text features at the pretraining stage. We hypothesize that such an ECG-text feature alignment learning paradigm can better facilitate the deep neural network to capture clinically useful patterns, which provide a more holistic picture of patients and potentially yield more accurate risk predictions. To the best of our knowledge, this is the first work that applies LLM-informed pre-training to benefit the training of the downstream ECG-based HF risk prediction model.

To summarize, our contribution is two-fold:

- 1) *A novel deep neural network architecture to enhance both the representation learning of ECG features and the model's interpretability:* The proposed network not only yields a quantitative risk score but also offers qualitative, interpretable insights into the neural network's reasoning through a dual attention mechanism. This unique feature acts as a transparent medium, enabling both clinicians and readers to observe and understand the intricate relationships between different ECG leads as well as the dynamic temporal patterns within each lead, highlighting

the particular leads and time segments that hold the greatest importance for reliable risk prediction.

- 2) *An effective model pretraining strategy with LLM:* We design an LLM-informed multi-modal pretraining task so that clinical knowledge can be transferred to the downstream risk prediction task. Training a deep risk prediction network is challenging due to a lack of sufficient event data. In this work, we advocate for the strategic use of large language models, coupled with *structured ECG reports with confidence scores*, to guide the pretraining process for the benefits of more data-efficient, accurate risk prediction models.

II. RELATED WORK

A. Risk Prediction

Risk prediction aims to estimate the chance $h(t)$ that a patient will have a certain event in an infinitesimal time interval $D \in [t, t + \Delta t)$: $h(t) = \lim_{\Delta t \rightarrow 0} \frac{p(D \in [t, t + \Delta t) | D \geq t)}{\Delta t}$, given that the event has not occurred before. A common approach is the Cox proportional hazards formulation (CoxPH) [17]. The CoxPH model is a statistical model for predicting the time-to-event outcomes based on the assumption that the hazard function $h(t)$ is proportional to a set of features or variables associated with the subject being studied. Mathematically, it can be defined as: $h(t|x) = h_0(t) \exp(g_\theta(x))$, where $h(t|x)$ models the hazard rate at which events occur at a time t taking the subject information x into account, $h_0(t)$ is the baseline hazard function shared by all observations, depending only on the time t , and $\exp(g_\theta(x))$ is the risk score for observation x . Conventional CoxPH models assume that the input observation is a set of covariates (e.g., sex, age, medical history, smoking) and constrain the function $g(\cdot)$ to the linear form: $g_\theta(x) = \theta^T x = \theta_1 \times x_1 + \theta_2 \times x_2 + \dots + \theta_p \times x_p$ with θ being the weights to the set of p covariates in $x : x_1, x_2, \dots, x_p$ [17]. In the existing body of literature pertaining to HF risk prediction, the predominant methodologies employ a predefined set of variables to construct a linear model, subsequently isolating variables that exhibit a high correlation for the purpose of risk prediction. This approach includes parameters such as prolonged QRS duration [6], [7], [8], [9], [10], various conduction disorders (specifically left and right bundle-branch blocks) [11], [12], elongated QT intervals [7], abnormalities in QRS/T angles and T wave patterns [7], [18], [19], [20], [21], and ST-segment depression in the V5 precordial lead [7], [22].

More recently, deep learning approaches such as DeepSurv [23] consider replacing the linear function $g(\cdot)$ with a neural network f , which is a deep architecture parameterized by the weights of the network θ . This relaxes the risk model to a non-linear form; also, the input can be high-dimensional without needing to be linearly independent, as the neural network is capable of extracting hierarchical features for risk score estimation in a non-linear fashion. Such an approach has been found to have superior performance against traditional approaches [23] on different risk prediction or survival prediction tasks. For example, researchers have successfully applied neural networks to automatically discover latent features from high-dimensional




Fig. 1. (a) Overview of the ECG dual attention encoder-based risk prediction network (ECG-DAN). A 12-lead ECG recording $\mathbf{x}$ is sent to the ECG dual attention encoder, which is capable of simultaneously extracting both cross-lead relationships, as well as temporal dynamic patterns within each lead, for better feature aggregation. Then, features from two routes are added and then sent to a max-average pooling layer, producing a flattened feature vector $\mathbf{z}_{ecg}$. Finally, we employ a multi-layer perceptron (MLP) module to map from a high-dimensional feature space into a risk score (scalar) r for heart failure. (b) Overview of the core attention module used in lead attention and temporal attention modules. See Section III-B for more details.

data, such as whole-slide pathology images [24], [25], 4D cardiac shape with motion information [26], etc., for risk/survival prediction. In this work, we adhere to a similar philosophy and concentrate on creating neural networks designed to autonomously extract features from intricate 12-lead ECG waves, bypassing the reliance on a predefined set of ECG parameters.

B. Large Language Model for Healthcare

Large language models (LLMs) are a class of artificial intelligence (AI) algorithms to understand human language. They can answer questions, provide summaries or translations, and create stories or poems [27]. Recent studies have found that LLMs can be effective in guiding representation learning on image data [28], enabling the knowledge transfer to several downstream vision tasks such as image classification, object detection, and segmentation. In the medical domain, such kind of multi-modal pretraining has been exploited for better understanding of imaging data, such as chest X-rays and magnetic resonance images, to benefit the downstream disease classification and medical image segmentation tasks [29], [30], [31], [32], [33]. Apart from image data, most recently, concurrent works have been made on exploring the connection between natural language and signals (ECG, EEG) for better disease classification [34], [35], [36]. To the best of our knowledge, there is no existing risk prediction work that explores the benefit of combining ECG reports with ECG waves for better representation learning. Compared to disease classification, the task of risk prediction is more challenging due to the lack of event records for training, which amplifies the overfitting issue associated with deep neural networks. Our work provides a promising transfer learning approach to alleviate the need for a large number

of event data, by utilizing LLMs and additional public large ECG-report datasets to conduct multi-modal pretraining.

III. METHODS

A. Overview

Assume we have a dataset of N triplets $\{\mathbf{x}_i, \delta_i, t_i\}_{i=1}^N$ to record the HF events in a population. Here, $\mathbf{x}_i \in \mathbb{R}^{12 \times T}$ is a 12-lead ECG signal (I, II, III, AVL, AVR, AVF, V1-6) with a recording length of T ; δ_i indicates whether there is known date of HF; t_i is the number of month to the censoring time if there is no reported HF event during the follow-up period ($\delta_i = 0$, right censored) or the number of months until the patient was diagnosed with HF during the follow-up ($\delta_i = 1$, uncensored). The objective is to have a risk prediction model f parameterized by θ so that it can predict a patient's risk of HF given the patient's ECG data. To this end, we design an ECG dual attention network (ECG-DAN), as shown in Fig. 1(a), where the input is a 12-lead ECG median wave recording $\mathbf{x}$, and the output of the network is a single node r , which estimates the risk score $r = f(\mathbf{x}; \theta)$.

B. ECG-DAN Network

A signature of ECG-DAN is a dual attention module, which is designed to extract morphological and spatial changes and relationships across different leads as well as the temporal dynamics inside each lead for a more comprehensive understanding of the heart's electrical activity. We first process each lead signal via a group of 1D convolution layers for noise filtering and feature extraction, which gives C_{in} -channel features for each lead at each time point. We then employ a set of K residual blocks with K $2 \times$ down-sampling layers along the temporal dimension to extract features at different scales. The output for each lead

consists of feature maps with C_{out} channels and a reduced time dimension of $\frac{T}{2^K}$: $\mathbf{h}_i \in \mathbb{R}^{C_{out} \times \frac{T}{2^K}}$. Those features from all 12 leads $\mathbf{h} : \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{12}\}$ are then chained and sent to the Lead Attention (LA) block, facilitating the learning of cross-lead interactions to enhance feature aggregation globally. Concurrently, a 12-lead Temporal Attention (TA) component is employed to capture crucial temporal patterns within each lead $\mathbf{h}_i$ along the time dimension. Here, for each lead $\mathbf{h}_i$, we apply an individual temporal attention module TA_i across the time domain separately, given the fact that different leads correspond to different directions of cardiac activation within three-dimensional (3D) space. The outcomes of the two attention modules are added and then summed with a concat-pooling module, producing concatenated flattened features from max-pooling and average pooling operations following common practice from previous ECG analysis work [13]:

$$\begin{aligned} \mathbf{h}' &= \mathbf{h}^{LA} + \mathbf{h}^{TA} \\ \mathbf{z}_{ecg} &= \text{Concat}(\text{Maxpooling}(\mathbf{h}'); \text{Avg-pooling}(\mathbf{h}')). \end{aligned} \quad (1)$$

With this latent feature $\mathbf{z}_{ecg}$, we then send it to a multi-layer perceptron network (MLP), which consists of three linear layers, reducing the feature by half and then projecting the feature into a 3D feature space, finally, regressing it to a risk score r .

Next, we will explain the two key modules for improved feature learning and model explainability: the lead attention module and the temporal attention, in more detail.

1) *Lead Attention*: We apply the multi-head attention-based encoder module introduced in transformers [37] to capture the cross-lead interactions. Specifically, we first reshape the feature matrix ($\mathbf{h} \in \mathbb{R}^{12 \times \frac{C_{out}T}{2^K}}$; 12 is number of leads) where each lead feature is treated as a token. As shown in Fig. 1(b), a sinusoid positional encoding PE [37] is added to the input feature $\mathbf{h} \leftarrow \text{LayerNorm}(\mathbf{h} + \text{PE}(\mathbf{h}))$ to encode the contextual information with layer normalization [38], followed by a multi-head self-attention module to capture cross-lead interactions. At a high level, the input to the attention are: key matrix ($\mathbf{K} \in \mathbb{R}^{12 \times \frac{C_{out}T}{2^K}}$), query matrix ($\mathbf{Q} \in \mathbb{R}^{12 \times \frac{C_{out}T}{2^K}}$), and value matrix ($\mathbf{V} \in \mathbb{R}^{12 \times \frac{C_{out}T}{2^K}}$), where key and query matrices are used to compute the weights to re-weight the value matrix. Mathematically, it can be expressed as:

$$\mathbf{h}^{LA} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{D_k}}\right)\mathbf{V} = \mathbf{A}^{LA}\mathbf{V} \quad (2)$$

where $\mathbf{A}^{LA} \in \mathbb{R}^{12 \times 12}$ is the attention weight matrix, and D_k is the feature dimension ($\frac{C_{out}T}{2^K}$). We use the same input $\mathbf{h}$ for computing $\mathbf{K}, \mathbf{Q}, \mathbf{V}$.¹

In addition to the self-attention, a fully connected feed-forward network (FFN) is applied to refine the re-weighted

features with residual connection:

$$\begin{aligned} \mathbf{h}^{LA} &= \text{LayerNorm}(\mathbf{h} + \mathbf{h}^{LA}) \\ \mathbf{h}^{LA} &= \text{FFN}(\mathbf{h}^{LA}) \\ &= \mathbf{h}^{LA} + \max(0, (\mathbf{h}^{LA}W_1 + b_1)W_2 + b_2) \end{aligned} \quad (3)$$

where W_1, W_2 and b_1, b_2 are weights and biases of two linear layers in FFN; LayerNorm is a layer normalization layer [38] for feature normalization. In this way, the output feature is adjusted with all the information from other leads into account.

2) *12-Lead Temporal Attention (TA)*: The structure of the temporal attention module is very similar to the lead attention. The only difference is that now we have 12 separate attention modules for different lead features so that each module is adapted to a specific lead to extract temporal dynamic features locally. Specifically, we split the feature $\mathbf{h}$ along the lead dimension, and for each lead feature $\mathbf{h}_i$, we first add sinusoid positional encoding PE information along the time domain and then process each lead with a separate temporal attention module where each time point is treated as a token. In other words, the input becomes a sequence of C_{out} -dimensional features with a sequence length of $l = \frac{T}{2^K}$ and the corresponding temporal attention matrix for lead i becomes: $\mathbf{A}_i^{TA} \in \mathbb{R}^{l \times l}$. Similar to the lead attention module process, we reweight and normalize the lead feature $\mathbf{h}_i$ with the generated temporal attention matrix and then employ an FFN module for feature refinement for each lead. The output of the temporal attention is a concatenation of output from the 12-lead temporal attention modules. The whole process can be defined as follows:

$$\begin{aligned} \mathbf{h}_i &\leftarrow \text{LayerNorm}(\mathbf{h}_i + \text{PE}(\mathbf{h}_i)) \\ \mathbf{h}_i^{TA} &= \text{LayerNorm}(\mathbf{h}_i + \mathbf{A}_i^{TA}\mathbf{h}_i^\top) \\ \mathbf{h}_i^{TA} &= \text{FFN}_i(\mathbf{h}_i^{TA}) \\ \mathbf{h}^{TA} &= \text{Concat}(\mathbf{h}_1^{TA}, \mathbf{h}_2^{TA}, \dots, \mathbf{h}_{12}^{TA})^\top. \end{aligned} \quad (4)$$

C. Training

Deep learning-based risk prediction often struggles with small datasets, particularly when events are rare, as in our case where HF events are below 5%. To overcome this and prevent overfitting, we adopt a two-stage training approach, see Fig. 2. Initially, we train our network on a large, diverse public ECG dataset. Of note, this dataset does not contain any HF records for risk prediction. After that, we initialize the model with the pretrained parameters and then conduct fine-tuning, focusing on risk prediction for particular populations with documented HF from the U.K. Biobank study [41]. Subsequently, we fine-tune the pretrained model on the U.K. Biobank data, specifically targeting HF risk prediction. The pretraining incorporates human-verified ECG reports to align features with clinical knowledge, aiming to improve the model's ability to discern pathological ECG patterns relevant to HF risk.

1) *Large Language Model Informed Model Pre-Training*: For pretraining, we employ the PTB-XL dataset [39], [40], a large, expert-verified collection of 21,799 clinical 12-lead ECGs with accompanying text reports. It features annotations by

¹In practice, for computational efficiency, following [37], we apply the multi-head attention trick, which first linearly projects the queries, keys, and values h times with different learnable linear projection matrices to lower dimensions (D_k/h) respectively to compute the attention matrix and re-weight the projected value matrix, then concatenate all the output h heads to recover the feature dimension.

(a) LLM-informed model pretraining

Example of an SCP code-based ECG report with confidence score:
 $\{IMI:15, LNGQT:100\}$


1. Search SCP code-statement knowledge base

SCP code	Category	SCP-ECG Description
IMI	Myocardial Infarction	inferior myocardial infarction
LNGQT	other ST-T descriptive statements	long QT-interval
...		

2. Construct structured reports

Myocardial Infarction:
inferior myocardial infarction

Other ST-T descriptive
statements: long QT-interval



(b) Finetuning for risk prediction




Fig. 2. Training overview. Our model is (a) first pretrained on the ECG-Report alignment task and the signal reconstruction task on a large-scale public dataset (PTB-XL [39], [40]), and then (b) finetuned on the heart failure risk prediction task with two specific cohorts from the U.K. Biobank where the future HF event data is available. Here, in PTB-XL dataset, each report has been abstracted to a set of SCP codes with SCP-ECG statement description and confidence score (annotated by human experts). We construct a structured report based on SCP-ECG protocol [39] and then send it to a frozen LLM to extract clinical knowledge for better representation learning guidance. As one ECG may have multiple SCP-code relevant statements, we extract text features separately and then use confidence-based reweighting to aggregate features for feature summation. See below texts for more details.

cardiologists according to the SCP-ECG standard and classifies waveforms into five categories: Normal (NORM), Myocardial Infarction (MI), ST/T Change (STTC), Conduction Disturbance (CD), and Hypertrophy, with possible overlap due to concurrent conditions. We follow the dataset creators' protocol, using folds 1-8 for training and folds 9-10 for validation and testing during pretraining [39].

Extracting latent text code z_{text} from reports using large language model: We employ an LLM to extract the knowledge embedded in ECG reports. Specifically, we use the medical domain language model: BioClinical BERT [42],² which has been trained on a large number of electronic health records from MIMIC III [43]. We consider two ways of extracting text embeddings z_{text} :

- *Latent text code from raw ECG report (raw):* For a piece of ECG report y (in English),³ we simply feed it to the LLM to get $z_{text} = \text{LLM}(y)$, following [32], [35].
- *Latent text code from structured ECG report weighted by confidence (structured with confidence):* The original PTB-XL dataset also provides ECG-SCP codes

generated from ECG reports. Specifically, each report has been abstracted to a set of SCP codes with SCP-ECG statement description and confidence score (annotated by human experts). In this case, we build a structured sentence with linked SCP statement category information and SCP description for each SCP code: $y(SCP) = \{\#Statement \ Category(SCP)\} \{\#SCP-ECG \ Statement \ Description(SCP)\}$.⁶ For example, as shown in Fig. 2(a), given an SCP code: LNGQT, the input becomes: 'other ST-T descriptive statements: long QT-interval'. We send this type of structured input to the LLM, which gives an embedding $z^{SCP} = \text{LLM}(y(SCP))$ for each SCP code. It is important to note that a single recording may encompass multiple SCP codes. In that case, multiple SCP embeddings are initially derived. These embeddings are then aggregated with weights based on corresponding confidence scores c to get the text embedding for the corresponding ECG: $z_{text} = \sum_j \frac{c_j z_j^{SCP}}{c^{sum}}$, where c^{sum} is the sum of all confidence scores for all corresponding statements for the input ECG signal.

2) *ECG-Report Alignment Loss:* To align the ECG to report, similar to [28], we first project the latent ECG code z_{ecg} and the latent text code z_{text} to e_{ecg} , e_{text} with two learnable projection functions p_x , p_y , respectively, so that the two embeddings $e_{ecg} =$

²[Online]. Available: https://huggingface.co/emilyalsentzer/Bio_ClinicalBERT

³Since the original reports are written in a mixture of German and English, we used the open-source machine translation tool: EasyNMT⁴ for batch translation following [34] and then used ChatGPT⁵ to further refine those failed cases with a prompt: *translate the ECG report into English:#text*.

⁴[Online]. Available: <https://github.com/UKPLab/EasyNMT>

⁵[Online]. Available: <https://chat.openai.com/>

⁶Database containing SCP code-statement mappings can be found at https://physionet.org/content/ptb-xl/1.0.1/scp_statements.csv and https://physionet.org/content/ptb-xl/1.0.1/ptbx1_database.csv.

$p_x(\mathbf{z}_{ecg}), \mathbf{e}_{text} = p_y(\mathbf{z}_{text})$ are of the same dimension, as shown in Fig. 2(a). We use a distance loss $\mathcal{D}$ to quantify the dissimilarity between the two:

$$\mathcal{L}_{align} = \frac{1}{n} \sum_{i=1}^n \mathcal{D}((\mathbf{e}_{ecg}, \mathbf{e}_{text})_i), \quad (5)$$

where $\mathcal{D}$ is a cosine embedding metric function: $\mathcal{D}(\mathbf{e}_{ecg}, \mathbf{e}_{text}) = 1 - \cos(\mathbf{e}_{ecg}, \mathbf{e}_{text})$, a metric that is commonly used for measuring the distance between two embeddings. The loss value is an average value over a batch of N paired embeddings.

3) *Pretraining Loss*: The total loss for pretraining is a combination of the ECG-report alignment loss and a signal reconstruction loss, defined as:

$$\mathcal{L}_{pretraining} = \mathcal{L}_{recon} + \mathcal{L}_{align}$$

$$\mathcal{L}_{recon} = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|^2 \quad (6)$$

where we compute the mean-squared-error between every input signal $\mathbf{x}_i$ and reconstructed signals $\hat{\mathbf{x}}_i$ in a batch and average them. Adding a signal reconstruction loss is necessary as it can help uncover latent generic features in the ECG signal and serve as a regularization term to avoid latent space collapse problems. During pre-training, we only update the parameters in the ECG encoder, ECG decoder, and two projectors, while keeping the parameters of the language model *frozen* for training stability and efficiency, as suggested by prior work [29]. Detailed network structures can be found in the Appendix.

4) *Finetuning Loss*: After pre-training, we copy the model weights to initialize the risk prediction model, see Fig. 2(b) and then finetune the risk prediction network. The finetuning loss is also a multi-task loss, including the self-supervised signal reconstruction loss, and a risk loss [23], which aims to minimize the average negative log partial likelihood of the set of uncensored patients ($\delta_i = 1$: developed into HF during the follow-up). The risk loss is defined as:

$$\mathcal{L}_{risk} = -\frac{1}{n_{\delta=1}} \sum_{i:\delta_i=1} \left[f_{\theta}(\mathbf{x}_i) - \log \sum_{j:t_j \geq t_i} \exp(f_{\theta}(\mathbf{x}_j)) \right] \quad (7)$$

where $n_{\delta=1}$ is the number of uncensored subjects in a batch, and $f_{\theta}(\mathbf{x})$ is a predicted risk score.

The total loss is then defined as:

$$\mathcal{L}_{finetuning} = \alpha \mathcal{L}_{recon} + (1 - \alpha) \mathcal{L}_{risk} \quad (8)$$

where α is a trade-off parameter to balance the contribution of two losses.

During model optimization, an issue with $\mathcal{L}_{risk}$ is that it can be very sensitive to the number of uncensored subjects $n_{\delta=1}$ (number of subjects who develop HF in the follow-up) in the training batch. This causes training instability problems if the number becomes zero or jumps from large to small and vice versa between batches. In real-world datasets, this issue is amplified due to a high-class imbalance between censored and uncensored subjects. To address this problem, we, therefore, perform a modified version of stochastic gradient descent for

TABLE I
STATISTICS OF STUDIED POPULATION(S)

characteristics		total	had HF during follow-up?	
			Yes	No
UKB-HYP	Age at examination [#]	66.31 (45-82)	69.57 (49-81)	66.27 (45-82)
	Sex (men/women)	6730/4851	109/53 (1.6%/1.1%)	6621/4798
UKB-MI	Age at examination [#]	68.21 (49-81)	67.28 (50-79)	68.25 (49-81)
	Sex (men/women)	650/150	27/5 (4.3%/3.4%)	623/145

[#]Values represent mean (minimum-maximum).

model optimization, where we select a batch of n observations with stratified random sampling to ensure every batch maintains a comparable ratio of censored to uncensored observations, mirroring the ratio found in the entire training population.

IV. EXPERIMENTS

A. Study Population

In this study, we focus on two populations: patients with hypertension (HYP) and MI, which are highly related to disease progression to HF. Subjects are selected from the U.K. Biobank (UKB) study, which is a large-scale biomedical database containing genetic, demographic, and disease information and is regularly updated with comprehensive follow-up studies, from approx. 500,000 subjects. The UKB dataset consists of a large portion of healthy subjects as well as those with a range of cardiovascular and other diseases. To assess the future risk of HF, our analysis is confined to individuals who have ECG together with imaging data available and have not been diagnosed with HF either prior to or at the time of the ECG evaluation.

1) *Ukb-Hyp*: HF-free subjects with prevalent HYP (had HYP before or during the ECG examination) at baseline time from the UKB dataset are studied [41]. We identified 11,581 HF-free HYP subjects. Follow-up time was defined as the time from the baseline ECG measurement until a diagnosis of HF or death or the end of follow-up (January 5, 2023). Most ECG recordings together with images were taken between 2014 and 2021. Records with less than two years of follow-up time were excluded. Among 11581 participants, 162 (1.2%) participants developed HF. The median follow-up time is 56 months (4.7 years), and the maximum follow-up time is 87 months (7.3 years).

2) *Ukb-Mi*: HF-free subjects at baseline with prevalent MI records are studied. Similar to the above selection procedure, we identified 800 subjects. Among them, 32 subjects (4%) developed HF during the follow-up period. The median follow-up time is 53 months (4.4 years), and the maximum follow-up time is 83 months (6.9 years). Code lists used for the retrieval of disease information can be found in the Appendix.

B. Implementation Details

For ECG signals, we use 12-lead median waveforms (covering a single beat) with a frequency of 500 Hz as input. Each lead is preprocessed with z -score normalization and then zero-padded to a length of 1024. We employ 5 residual blocks ($K = 5, C_{in} = 4, C_{out} = 16$) to obtain multi-scale features. The dimension of latent ECG feature $\mathbf{z}_{ecg}$ is 512, and the dimension of projected

TABLE II
COMPARISON OF RISK PREDICTION PERFORMANCE USING DIFFERENT
PRE-TRAINING TASKS

Pretraining Tasks	Study population	
	Hypertension (UKB-HYP)	Myocardial Infarction (UKB-MI)
-	0.6122 (0.0190)	0.5065 (0.0776)
SR	0.6327 (0.0165)	0.5069 (0.0770)
SR + Classification	0.6370 (0.0216)	0.5220 (0.0475)
SR + ECG-R Alignment (raw)	0.6088 (0.0189)	0.5796 (0.0570)
SR + ECG-R Alignment (structured with confidence)	0.6349 (0.0156)	0.5805 (0.0580)

SR: Signal Reconstruction; Classification: ECG Disease Classification; ECG-R: ECG-Report.

All experiments are performed using the same proposed network architecture. reported values are average C-index over 5 cross-validations using different random splits.

features $\mathbf{e}_{ecg}$, $\mathbf{e}_{text}$ is 128. The default convolutional kernel size is 5. Further network details are provided in the Appendix. We use a batch size of 128 ($n = 128$) for model updates at pre-training, with random lead masking for data augmentation [44]. AdamW optimizer [45] with stratified batches is used for training. For fine-tuning in post-MI risk prediction (UKB-MI), we maintain a batch size of 128. Given the low HF incidence ($< 2\%$) in UKB-HYP, we increase the batch size to 1024 to include enough uncensored subjects for calculating $\mathcal{L}_{risk}$. Batch-wise dropout is used to stabilize fine-tuning [24] as well as avoid overfitting. The loss function's weighting parameter (α) is set at 0.5. Pre-training and fine-tuning are conducted over 300 and 100 epochs, respectively, to ensure convergence.

C. Evaluation Metrics and Evaluation Method

We report the concordance index (C-index) [46] as the primary evaluation metric. This metric measures the accuracy of the ranking of predicted time based on the ratio of concordant pairs: $C\text{-index} = \frac{\# \text{ concordant pairs}}{\# \text{ concordant pairs} + \# \text{ discordant pairs}}$. A concordant pair refers to when the predicted time ranking of two subjects aligns with their actual ranking, while a discordant pair means the opposite. A pair (i, j) is concordant if $t_i < t_j$ and risk scores $r_i > r_j$, or vice versa. A value of 1 denotes perfect prediction, while a value of 0.5 indicates a prediction quality equivalent to random guessing.

Robust evaluation with multiple two-fold stratified cross-validation: Due to the scarcity of uncensored subjects (with HF events), using conventional deep learning data splits (e.g. five-fold cross validation or 7:1:2 for training, validation and testing) would result in too few HF events for reliable C-index evaluation. Hence, we opt for a 1:1 training-to-testing ratio with *stratified* cross-validation, aligned with patient HF status during follow-up. This approach constitute a two-fold *stratified* cross-validation, which also helps to maintain the same proportion of each class as in the original dataset in each split. To further avoid over-fitting, 20% of the training data is randomly allocated as a validation set for hyper-parameter search (e.g. choice of the learning rate) and model selection. Specifically, we split the dataset into two folds, one for training and validation, the

rest for testing. This process is then repeated once more, with the roles of the two parts reversed. To ensure that the result is not biased by the splitting, we perform the above procedure five times, each with a different split across two datasets (UKB-HYP and UKB-MI). The final model performance is reported as the average (with standard deviation) of these 10 trials.

V. RESULTS

A. LLM-Informed Pre-Training Improves the Accuracy of the Downstream Risk Prediction

We first compare our proposed pretraining strategy against various pretraining tasks:

- No pretraining (represented by a hyphen), serving as a baseline, with networks trained for 400 epochs (equivalent to the sum of pretraining and finetuning epochs);
- Pretraining on signal reconstruction only;
- Pretraining on signal reconstruction combined with multi-label disease classification (NORM, MI, STTC, CD, hypertrophy), using a cross-entropy loss for classification;
- Pretraining on signal reconstruction and ECG-Report alignment using raw text reports;
- Our proposed method, pretraining on signal reconstruction and ECG-Report alignment with structured report and confidence information, as detailed in Section III-C1.

Results are shown in Table II. It is clear that the models pretrained using the proposed strategy (signal reconstruction + ECG2Text alignment (structured w/ confidence)) consistently obtain the highest accuracy among the two tasks, obtaining the average C-index: 0.6349 (0.0156) on UKB-HYP and 0.5805 (0.0580) on UKB-MI.

B. Comparison Study on Risk Prediction Using Different Network Architectures

We further compare model performance using different encoder architectures, including 1) the encoder used in a variational autoencoder (VAE)-like network architecture, which has been found effective for ECG signal feature representation learning in previous works [47], [48], [49], [50]; 2) XResNet1D [13], which is the top performing network architecture on a wide range of different ECG analysis tasks such as ECG disease classification, age regression, and form/rhythm prediction on the public PTB-XL benchmark dataset [39] and ICBEB2018 dataset [51]. We report their performance trained w/o and w/ the proposed language model informed pre-training strategy with structured SCP report in Table III. The proposed ECG attention network has the fewest parameters. Yet, models with this type of network architecture and the proposed pre-training strategy obtain the highest C-index scores on both two tasks.

C. Comparison Study Between Traditional ECG Parameter-Based Risk Prediction vs ECG Dual Attention

We further compare our method to the traditional method with well-established ECG parameters. Specifically, we collect a set of ECG parameters from the ECG signals, which have been

TABLE III
COMPARISON OF RISK PREDICTION PERFORMANCE USING DIFFERENT TYPES OF DEEP NEURAL NETWORKS

Network architectures	# parameters	LLM-informed Pretraining	Study population	
			Hypertension (UKB-HYP)	Myocardial Infarction (UKB-MI)
CNN-VAE ([47], [48], [49], [50])	7.0M	✗	0.6215 (0.0237)	0.5675 (0.0353)
CNN-VAE ([47], [48], [49], [50])	7.0M	✓	0.6346 (0.0218)	0.5484 (0.0269)
XResNet1D ([13])	1.9M	✗	0.5601 (0.0273)	0.5357 (0.0485)
XResNet1D ([13])	1.9M	✓	0.6234 (0.0143)	0.5431 (0.0387)
ECG dual attention (The proposed)	1.4M	✗	0.6122 (0.0190)	0.5065 (0.0776)
ECG dual attention (The proposed)	1.4M	✓	0.6349 (0.0156)	0.5805 (0.0580)

By default, all network has been initialized with pre-trained on the proposed ECG-Report alignment and reconstruction tasks (300 epochs) and then finetuned on the risk prediction task (100 epochs). For fair comparison, models without any pre-training are trained for 400 epochs (300+100). To avoid overfitting, models with the highest C-index score on the validation set are selected as the final model. Reported values are average C-index over 5 cross-validations using different random splits (10 trials in total).

TABLE IV
COMPARISON OF RISK PREDICTION PERFORMANCE USING TRADITIONAL ECG PARAMETER-BASED AND OUR DEEP LEARNING MODEL (ECG DUAL ATTENTION)

Method	Input	Model Type	Study population	
			Hypertension (UKB-HYP)	Myocardial Infarction (UKB-MI)
Traditional	15 ECG parameters	Linear	0.6149 (0.0125)	0.5398 (0.0200)
ECG dual attention (the proposed)	12-lead Median waves	Nonlinear	0.6349 (0.0156)	0.5805 (0.0580)

Reported values are average C-index over five times of 2-fold cross-validations using different random splits (10 trials in total).

identified in previous studies for incident HF prediction and mortality estimation [6], [7], [8], [9], [11], [12], [18], [19], [20], [21], [22]. All of these parameters are automatically extracted by the ECG device (using the GE CardioSoft V6⁷) with supporting evidence from previous work: 1) Ventricular rate; 2) Left-axis deviation; 3) Right-axis deviation; 4) Prolonged P-wave duration (> 120 ms) [22]; 5) Prolonged PR interval (> 200 ms) [22]; 6) Prolonged QRS duration (> 100 ms) [22]; 7) Prolonged QT interval (≥ 460 (women)/450 (men) ms using the Framingham formula [52]; 8) Delayed intrinsicoid deflection (DID time) (the maximum value in leads V5 and V6 > 50 ms) [53]; 9) Abnormal P-wave axis (values outside the range of 0° and 75°) [54]; 10) Left ventricular hypertrophy [55]; 11) Abnormal QRS-T angle (> 77° (women)/88° (men)) [22]; 12) Low QRS voltage [56]; 13) ST-T abnormality [56]; 14) Right bundle-branch block [56]; 15) Left bundle-branch block [56].

Following [22], we fit these ECG parameters into the conventional CoxPH linear regression model [17] and use it as a strong competing model for HF risk prediction. Table IV shows the performances using our approach and the traditional ECG parameter-based approach, and Fig. 3 shows Kaplan-Meier plots which depict the survival probability estimates over time, stratified by risk groups defined by each model's predictions. We further plot the features learned in the last hidden layer ($\in \mathbb{R}^3$) of ECGs from the risk prediction branch in the proposed network and representative ECG waves with the lowest/highest risk score in Fig. 4.




Fig. 3. Kaplan-Meier risk curves for a) conventional model using a composite of 15 predefined ECG parameters/measurements, and b) the proposed ECG dual attention risk prediction model with the language-informed pretraining. For both models, patients were divided into low- and high-risk groups with a cutoff value referenced from the top 98th percentile (for UKB-HYP) or top 96th percentile (UKB-MI) risk scores predicted by the model, reflecting the statistics of the datasets in Table I.

TABLE V
ABLATION EXPERIMENTS

Method	Study population	
	Hypertension (UKB-HYP)	Myocardial Infarction (UKB-MI)
w/o time attention module	0.6038 (0.0243)	0.5855 (0.0353)
w/o lead attention module	0.6018 (0.0418)	0.4920 (0.0884)
w/o lead+time attention module	<u>0.6167 (0.0105)</u>	0.5195 (0.0539)
The proposed	0.6349 (0.0156)	<u>0.5805 (0.0580)</u>

Numbers in bold represent the highest values, while numbers with underlines denote the second highest.

D. Dual Attention Mechanism Improves Model Stability and Interpretability

1) *Quantitative Ablation Study:* We also evaluate the impact of the dual attention modules on whether they can help to enhance the accuracy of risk prediction. We conduct the ablation study experiments using the same training strategy with the same network but either the lead and/or the temporal attention module removed. Results are shown in Table V. It can be observed that the propose model containing both attention modules consistently produces the most stable performance, with the best performance on the UKB-HYP and the second best performance

⁷[Online]. Available: <https://biobank.ndph.ox.ac.uk/ukb/ukb/docs/CardiosoftFormatECG.pdf>




Fig. 4. 3D visualization of the last 3-dim hidden feature learned in the risk prediction subnetwork along with the visualization of input ECG waves with lowest predicted risk score (dark purple) and highest predicted risk score (light yellow) on the (a) UKB-HYP and (b) UKB-MI datasets.




Fig. 5. Visualization of lead attention patterns and differences between low-risk and high-risk groups across two, different populations: UKB-HYP and UKB-MI.

on the UKB-MI, yielding the best average performance across the two datasets.

2) *Lead Attention Matrix Visualization*: We further visualize the average lead attention matrix at the population level, averaging matrices in different risk groups. For both datasets, patients are divided into low- and high-risk groups with a cutoff value referenced from the top 98th percentile (for UKB-HYP) or top 96th percentile (UKB-MI) risk scores predicted by the model, where the threshold is chosen based on the statistics of the datasets. Of note, all risk scores and attention matrices are obtained using the same cross-validation strategy, where the predicted risk is computed by averaging predictions from models trained with data excluding that subject. Fig. 5 illustrates the contribution of each lead for feature re-weighting across all 12-lead features for different populations. We can find similar patterns though the underlying study cohorts are different.

3) *Temporal Attention Activation Map Visualization*: To visualize the temporal attention matrix A_i^{TA} in the original ECG

time length T , we adapt GradCAM for ECG leads to highlight the model's focus. We condense the attention matrix by summing column values for attention scores across time, and then map them to the ECG input features, creating Grad-CAM maps. These maps, weighted by their attention scores, reveal the network's focus areas on the ECG. More details of implementation can be found in the Appendix. The visualization is presented in Fig. 6.

VI. DISCUSSION

Importance of language informed pretraining: In the experiments, we first studied the impact of different pretraining tasks for downstream risk predictions and highlighted the value of LLM-informed pretraining for risk prediction in Table II. In general, pre-training tasks enhance the risk prediction performance compared to those without pretraining, especially on the smaller dataset (UKB-MI). It is interesting to see that the performance of the deep risk prediction model (average C-index: 0.5065) can be inferior to the traditional approach (average C-index 0.5398) if without proper pretraining, see Tables II and IV. This indicates that pre-training is important to alleviate model over-fitting. Our study further highlights the critical role of targeted model pretraining in identifying both generic and pathological features for downstream HF risk prediction. Pre-training on structured ECG reports outperforms methods using simple disease labels or unstructured text inputs. We credit this improved performance to the integration of detailed context from structured reports and supplementary confidence information. Fig. 7 shows the UMAP [57] visualization of latent text codes of structured reports z^{SCP} extracted by LLM, which suggests that LLM distinguishes between disease-specific embeddings (e.g., separated clusters) and captures the interrelations among various diseases. For instance, the embedding for the normal or healthy category (red dots) is nearer to ST/T wave change embeddings (blue dots) and notably distant from those for MI




Fig. 6. Visualization of cross-lead (a), (b) and 12-lead temporal attention maps (d), (e) obtained from a high HF risk ECG with HYP (a), (d) and a high HF risk ECGs with MI (b), (e). (c) is a schematic standard ECG for illustrative purpose. Source: Wikimedia Commons.




Fig. 7. U-map visualization of latent code embeddings $\mathbf{z}^{SCP}$ from the large language model using different structured SCP statements. Different colors represent the categorization of statements with disease labels.

(green dots), which is aligned with the fact that ST/T variations can be non-pathological, unlike the distinct disease pattern of MI.

ECG dual attention network exhibits high effectiveness of risk stratification: By comparing the results in Table III, it is clear that the suggested approach surpasses leading deep learning approaches in terms of both computational costs and accuracy. Interestingly, we observe that language-informed pretraining, while continuously boosting the performance on the ResNet-based structure and our attention model, does not enhance the CNN-VAE-based model's performance. This phenomenon probably attributed to the over-parameterization in the CNN-VAE networks (7.0 M). In that case, the prioritization of strong regularization to shape the latent representation to like a Gaussian distribution to ensure feature independence [58]

becomes crucial. Such regularization might conflict with the language-informed training process.

Results in Table IV also show that our approach surpasses the performance of a traditional methodology, which relies on a predefined set of ECG parameters with a simple linear model, as well as current leading deep learning-based models in terms of both computational costs and performance. The high overlapping region between high-risk and low-risk curves of MI populations (left bottom) when using the traditional approach in Fig. 3 reveals the challenge when the underlying population is limited. By contrast, our method can consistently stratify the patients into different risks with a clear gap. Fig. 4 demonstrates that our network effectively discriminates between low-risk (upper right, purple dots) and high-risk ECGs (bottom left, yellow dots) in the latent space, correlating high-risk ECGs with clinical markers like prolonged QRS duration and longer QT intervals, in alignment with established studies [6], [7], [8], [9], [10].

ECG dual attention network exhibits high interpret-ability:

Furthermore, the proposed two dual attention modules offer model interpretability, a feature highly valued in the clinical setting. The integration of a dual attention mechanism serves as a window into the network's thought process, allowing clinicians and readers alike to visualize and comprehend the complex interactions between different ECG leads, as shown in Fig. 5. First, we found the three augmented leads (aVR, aVL, aVF) contributed the least (see blue regions). We believe that this may be because aVR, aVL, and aVF can be derived from Lead I, II, and III using Goldberger's equations [59], thus containing redundant information for feature aggregation and risk prediction. Second, we found that the network typically pays more attention

to chest lead I and precordial leads (V2-V6) (see red regions). In clinical practice, comparing the morphological changes across the precordial leads (V1-V6) can help to identify ST/T wave abnormality, which could be an indicator of future HF [7], [18], [19], [20], [21] and sudden cardiac death [60]. On the other hand, upon comparing the attention patterns derived from high-risk groups to those from low-risk groups, as illustrated in the column titled ‘Difference’ in Fig. 5, we observed elevated activation values within the high-risk groups. This observation underscores the network’s heightened sensitivity in identifying abnormal features.

For better understanding, we visualize the cross-lead and temporal activation maps of two high-risk cases in Fig. 6. For temporal ones, we apply GradCAM [61] to map the low-resolution temporal attention maps back to the original input signal level (see Appendix for more details). It can be observed that cross-lead module provides an overview to identify the uniqueness or exploit synergy or cross-lead interactions among various leads, whereas temporal lead attention focuses on identifying local areas in each lead important to the prediction. For example, leads III and II stand out in the cross-lead attention maps (a) and (b) respectively, highlighting unique pathological pattern (prolonged negative QRS in lead III, see (d) or strong abnormal noises in lead II, see original signals in (e)), in contrast to other leads. Additionally, the orange diagonal clusters in panel (b) showcase good R wave progression from V2 to V6, see (e). On the other hand, the temporal attention maps in (d) and (e) reveal more intricate details, showing the neural network’s capability to focus on clinically significant features such as P, Q, S, T waves, and R peaks, even with strong noise present, and highlight pathological abnormalities like abnormal R peaks, T wave irregularities and prolonged QRS and QT intervals. This implies that the dual attention network can autonomously discover clinically relevant biomarkers from ECG data without explicitly being taught during its training phase. Future research will aim at collaborating with medical professionals to verify potential novel biomarkers using these visual tools.

Limitations: One of the limitations of the current work is that it only considers information from ECG signals. In the future, we will consider extending our approach to a broader spectrum of data. The input features can be blood test results, demography (age, sex, ethnicity) [47], smoking, chronic disease condition such as diabetes, imaging-derived features [49], [50], [62] as well as genetic information to create a more holistic and accurate characterization of the patient [8], [9], [63], [64]. Moreover, it is interesting to increase the inclusivity and diversity of our study by considering a broader population base.

Broader Impact: We believe the proposed LLM-informed pre-training scheme is not limited to the risk prediction task. It also holds the potential to inspire new approaches and applications for ECG signal analysis, such as disease diagnosis and ECG segmentation. The LLM model, utilizing textual reports, acts as an additional source of contextual supervision, which guides the ECG network to better understand and extract complex patterns. We would also like to highlight that the utilized public pre-training dataset which comprises a diverse and comprehensive

pathologies is crucial for model generalization. In the future, it is also interesting to further enhance the model interpretability through the incorporation of the generative capabilities of LLM, providing structured report for further explanation.

VII. CONCLUSION

This paper presents a study with a novel ECG dual attention network for ECG-based HF risk prediction. This network distinguishes itself through its unique blend of being both lightweight and efficient, outperforming existing models in the field. Its standout feature is the ability to generate not just a quantitative risk score but also to provide a qualitative interpretation of its internal processes. This is achieved through the generation of attention visualization maps, which span across and delve within individual leads, offering a granular view of the network’s focus and decision-making process. We hope to contribute to the development of more transparent and interpretable AI-assisted systems, fostering trust and facilitating broader adoption in clinical settings. Additionally, the study presents a multi-modal pretraining approach for risk prediction models, which leverages external public ECG reports and confidence scores from *diverse* populations, combined with advanced large language models, to address the challenges posed by limited future event data in risk prediction tasks.

ACKNOWLEDGMENT

The authors express no conflict of interest.

REFERENCES

- [1] British Heart Foundation, “Heart failure hospital admissions rise by a third in five years,” Nov. 2019. Accessed: Oct., 2023. [Online]. Available: <https://www.bhf.org.uk/what-we-do/news-from-the-bhf/news-archive/2019/november/heart-failure-hospital-admissions-rise-by-a-third-in-five-years>
- [2] G. F. Tomaselli and D. P. Zipes, “What causes sudden death in heart failure?,” *Circulation Res.*, vol. 95, no. 8, pp. 754–763, 2004.
- [3] R. E. Lane, M. R. Cowie, and A. W. Chow, “Prediction and prevention of sudden cardiac death in heart failure,” *Heart*, vol. 91, no. 5, pp. 674–680, 2005.
- [4] C. W. Yancy et al., “Clinical presentation, management, and in-hospital outcomes of patients admitted with acute decompensated heart failure with preserved systolic function: A report from the acute decompensated heart failure national registry (ADHERE) database,” *J. Amer. College Cardiol.*, vol. 47, no. 1, pp. 76–84, Jan. 2006.
- [5] B. S. Thompson and C. W. Yancy, “Immediate versus delayed diagnosis of heart failure: Is there a difference in outcomes? Results of a Harris interactive patient survey,” *J. Cardiac Failure*, vol. 10, no. 4, Aug. 2004, Art. no. S125.
- [6] R. M. F. L. da Silva, N. M. Kazzaz, R. M. Torres, and M. D. C. V. Moreira, “P-wave dispersion and left atrial volume index as predictors in heart failure,” *Arquivos Brasileiros de Cardiologia*, vol. 100, no. 1, pp. 67–74, Jan. 2013.
- [7] B. Triola et al., “Electrocardiographic predictors of cardiovascular outcome in women: The national heart, lung, and blood institute-sponsored women’s ischemia syndrome evaluation (wise) study,” *J. Amer. College Cardiol.*, vol. 46, pp. 51–56, 2005.
- [8] S. S. Khan et al., “10-year risk equations for incident heart failure in the general population,” *J. Amer. College Cardiol.*, vol. 73, no. 19, pp. 2388–2397, 2019. [Online]. Available: <https://www.jacc.org/doi/abs/10.1016/j.jacc.2019.02.057>
- [9] S. S. Khan et al., “Development and validation of a long-term incident heart failure risk model,” *Circulation Res.*, vol. 130, no. 2, pp. 200–209, Jan. 2022.

- [10] L. Ilkhanoff et al., "Association of QRS duration with left ventricular structure and function and risk of heart failure in middle-aged and older adults: The Multi-Ethnic study of atherosclerosis (MESA)," *Eur. J. Heart Failure*, vol. 14, no. 11, pp. 1285–1292, Nov. 2012.
- [11] Z.-M. Zhang, P. M. Rautaharju, R. J. Prineas, L. Loefer, W. Rosamond, and E. Z. Soliman, "Ventricular conduction defects and the risk of incident heart failure in the atherosclerosis risk in communities (ARIC) study," *J. Cardiac Failure*, vol. 21, no. 4, pp. 307–312, Apr. 2015.
- [12] Z.-M. Zhang et al., "Different patterns of bundle-branch blocks and the risk of incident heart failure in the Women's Health Initiative (WHI) study," *Circulation, Heart Failure*, vol. 6, pp. 655–661, 2013.
- [13] N. Strodthoff, P. Wagner, T. Schaeffter, and W. Samek, "Deep learning for ECG analysis: Benchmarks and insights from PTB-XL," *IEEE J. Biomed. Health Inform.*, vol. 25, no. 5, pp. 1519–1528, May 2021.
- [14] V. A. Ardeti, V. R. Kolluru, G. T. Varghese, and R. K. Patjoshi, "An overview on state-of-the-art electrocardiogram signal processing methods: Traditional to AI-based approaches," *Expert Syst. Appl.*, vol. 217, May 2023, Art. no. 119561.
- [15] J. W. Hughes et al., "A deep learning-based electrocardiogram risk score for long term cardiovascular death and disease," *NPJ Digit. Med.*, vol. 6, no. 1, Sep. 2023, Art. no. 169.
- [16] E. Alsentzer et al., "Publicly available clinical BERT embeddings," in *Proc. 2nd Clin. Natural Lang. Process. Workshop*, Minneapolis, MN, USA, 2019, pp. 72–78. [Online]. Available: <https://www.aclweb.org/anthology/W19-1909>
- [17] D. R. Cox, "Regression models and life-tables," *J. Roy. Statist. Soc. Ser. B*, vol. 34, no. 2, pp. 187–220, 1972.
- [18] P. M. Rautaharju et al., "Electrocardiographic predictors of incident heart failure in men and women free from manifest cardiovascular disease (from the atherosclerosis risk in communities [ARIC] study)," *Amer. J. Cardiol.*, vol. 112, pp. 843–849, 2013.
- [19] P. M. Rautaharju, C. Kooperberg, J. C. Larson, and A. LaCroix, "Electrocardiographic predictors of incident congestive heart failure and all-cause mortality in postmenopausal women: The women's health initiative," *Circulation*, vol. 113, pp. 481–489, 2006.
- [20] P. M. Rautaharju, R. J. Prineas, J. Wood, Z.-M. Zhang, R. Crow, and G. Heiss, "Electrocardiographic predictors of new-onset heart failure in men and in women free of coronary heart disease (from the atherosclerosis in communities [ARIC] study)," *Amer. J. Cardiol.*, vol. 100, pp. 1437–1441, 2007.
- [21] Z.-M. Zhang et al., "Comparison of the prognostic significance of the electrocardiographic QRS/T angles in predicting incident coronary heart disease and total mortality (from the atherosclerosis risk in communities study)," *Amer. J. Cardiol.*, vol. 100, no. 5, pp. 844–849, Sep. 2007.
- [22] W. T. O'Neal et al., "Electrocardiographic predictors of heart failure with reduced versus preserved ejection fraction: The Multi-Ethnic study of atherosclerosis," *J. Amer. Heart Assoc.*, vol. 6, no. 6, May 2017, Art. no. e006023.
- [23] J. Katzman, U. Shaham, J. Bates, A. Cloninger, T. Jiang, and Y. Kluger, "DeepSurv: Personalized treatment recommender system using a cox proportional hazards deep neural network," *BMC Med. Res. Methodol.*, vol. 18, pp. 1–12, 2018.
- [24] S. Jiang, A. A. Suriawinata, and S. Hassanpour, "MHAttnSurv: Multi-Head attention for survival prediction using Whole-Slide pathology images," *Comput. Biol. Med.*, vol. 158, May 2023, Art. no. 106883.
- [25] Z. Li, Y. Jiang, M. Lu, R. Li, and Y. Xia, "Survival prediction via hierarchical multimodal co-attention transformer: A computational histology-radiology solution," *IEEE Trans. Med. Imag.*, vol. 42, no. 9, pp. 2678–2689, Sep. 2023.
- [26] G. A. Bello et al., "Deep learning cardiac motion analysis for human survival prediction," *Nat. Mach. Intell.*, vol. 1, pp. 95–104, 2019.
- [27] K. Singhal et al., "Large language models encode clinical knowledge," *Nature*, vol. 620, pp. 172–180, Dec. 2022.
- [28] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
- [29] C. Liu et al., "M-FLAG: Medical Vision-Language pre-training with frozen language models and latent space geometry optimization," in *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Interv.*, 2023, pp. 637–647.
- [30] X. Zhang, C. Wu, Y. Zhang, W. Xie, and Y. Wang, "Knowledge-enhanced visual-language pre-training on chest radiology images," *Nat. Commun.*, vol. 14, no. 1, Jul. 2023, Art. no. 4542.
- [31] J. Liu et al., "Clip-driven universal model for organ segmentation and tumor detection," in *Proc. Int. Conf. Comput. Vis.*, 2023, pp. 21095–21107.
- [32] J. Qiu et al., "Multimodal representation learning of cardiovascular magnetic resonance imaging," Apr. 2023.
- [33] Ö. Turgut et al., "Unlocking the diagnostic potential of ECG through knowledge transfer from cardiac MRI," 2023. [Online]. Available: <https://arxiv.org/abs/2308.05764>
- [34] J. Qiu et al., "Transfer knowledge from natural language to electrocardiography: Can we detect cardiovascular disease through language models?," in *Proc. ACM Eur. Conf. Field Comput. Linguistics*, Dubrovnik, Croatia, 2023, pp. 442–453.
- [35] C. Liu, Z. Wan, S. Cheng, M. Zhang, and R. Arcucci, "ETP: Learning transferable ECG representations via ECG-text pre-training," *IEEE Int. Conf. Acoust., Speech Signal Process.*, Seoul, Korea, 2024, pp. 8230–8234, doi: [10.1109/ICASSP48485.2024.10446742](https://doi.org/10.1109/ICASSP48485.2024.10446742).
- [36] W. Han et al., "Can brain signals reveal inner alignment with human languages?," 2023. [Online]. Available: <https://aclanthology.org/2023.findings-emnlp.120.pdf>
- [37] A. Vaswani et al., "Attention is all you need," in *Proc. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [38] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *Neural Inf. Process. Syst. 2016 Workshop*, Jul. 2016. [Online]. Available: <https://arxiv.org/abs/1607.06450>
- [39] P. Wagner et al., "PTB-XL, a large publicly available electrocardiography dataset," *Sci. Data*, vol. 7, no. 1, May 2020, Art. no. 154.
- [40] N. Strodthoff et al., "PTB-XL+, a comprehensive electrocardiographic feature dataset," *Sci. Data*, vol. 10, no. 1, 2023, Art. no. 279.
- [41] C. Sudlow et al., "UK biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age," *PLoS Med.*, vol. 12, no. 3, Mar. 2015, Art. no. e1001779.
- [42] E. Alsentzer et al., "Publicly available clinical BERT embeddings," in *Proc. 2nd Clin. Natural Lang. Process. Workshop*, Minneapolis, MN, USA, 2019, pp. 72–78. [Online]. Available: <https://aclanthology.org/W19-1909>
- [43] A. E. W. Johnson et al., "MIMIC-III, a freely accessible critical care database," *Sci. Data*, vol. 3, May 2016, Art. no. 160035.
- [44] P. Nejedly et al., "Classification of ECG using ensemble of residual CNNs with attention mechanism," in *Proc. 2021 Conf. Comput. Cardiol.*, 2021, pp. 1–4.
- [45] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. Int. Conf. Learn. Representations*, 2019, Art. no. 26395.
- [46] F. E. Harrell Jr, R. M. Califf, D. B. Pryor, K. L. Lee, and R. A. Rosati, "Evaluating the yield of medical tests," *J. Amer. Med. Assoc.*, vol. 247, no. 18, pp. 2543–2546, May 1982.
- [47] Y. Sang, M. Beetz, and V. Grau, "Generation of 12-lead electrocardiogram with subject-specific, image-derived characteristics using a conditional variational autoencoder," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2022, pp. 1–5, doi: [10.1109/ISBI52829.2022.9761431](https://doi.org/10.1109/ISBI52829.2022.9761431).
- [48] M. Beetz, A. Banerjee, Y. Sang, and V. Grau, "Combined generation of electrocardiogram and cardiac anatomy models using multi-modal variational autoencoders," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2022, pp. 1–4, doi: [10.1109/ISBI52829.2022.9761590](https://doi.org/10.1109/ISBI52829.2022.9761590).
- [49] M. Beetz, A. Banerjee, and V. Grau, "Multi-domain variational autoencoders for combined modeling of MRI-based biventricular anatomy and ECG-based cardiac electrophysiology," *Front. Physiol.*, vol. 13, 2022, Art. no. 886723. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/fphys.2022.886723>
- [50] L. Li, J. Camps, A. Banerjee, M. Beetz, B. Rodríguez, and V. Grau, "Deep computational model for the inference of ventricular activation properties," in *Proc. Int. Workshop Statist. Atlases Comput. Models Heart*, Springer, 2022, pp. 369–380, doi: [10.1007/978-3-031-23443-9_34](https://doi.org/10.1007/978-3-031-23443-9_34).
- [51] F. Liu et al., "An open access database for evaluating the algorithms of electrocardiogram rhythm and morphology abnormality detection," *J. Med. Imag. Health Inform.*, vol. 8, no. 7, pp. 1368–1373, 2018.
- [52] A. Sagie, M. G. Larson, R. J. Goldberg, J. R. Bengtson, and D. Levy, "An improved method for adjusting the QT interval for heart rate (the Framingham Heart Study)," *Amer. J. Cardiol.*, vol. 70, no. 7, pp. 797–801, Sep. 1992.
- [53] W. T. O'Neal et al., "Electrocardiographic time to intrinsicoid deflection and heart failure: The multi-ethnic study of atherosclerosis," *Clin. Cardiol.*, vol. 39, no. 9, pp. 531–536, Sep. 2016.
- [54] Y. Li, A. J. Shah, and E. Z. Soliman, "Effect of electrocardiographic P-wave axis on mortality," *Amer. J. Cardiol.*, vol. 113, no. 2, pp. 372–376, Jan. 2014.
- [55] R. B. Devereux, P. N. Casale, R. R. Eisenberg, D. H. Miller, and P. Kligfield, "Electrocardiographic detection of left ventricular hypertrophy using echocardiographic determination of left ventricular mass as the reference standard. comparison of standard criteria, computer diagnosis and physician interpretation," *J. Amer. College Cardiol.*, vol. 3, no. 1, pp. 82–87, Jan. 1984.


- [56] R. J. Prineas, R. S. Crow, and Z.-M. Zhang, *The Minnesota Code Manual of Electrocardiographic Findings*. London UK: Springer, Nov. 2009.
- [57] L. McInnes, J. Healy, J. Saul, and L. Großberger, "UMAP: Uniform manifold approximation and projection," *J. Open Source Softw.*, vol. 3, no. 29, Sep. 2018, Art. no. 861.
- [58] D. P. Kingma and M. Welling, "Auto-Encoding variational bayes," in *Proc. Int. Conf. Learn. Representations*, 2014, pp. 1–14.
- [59] (2023) EKG/ECG leads, electrodes, limb leads, chest (precordial) leads. ECG Waves. Accessed: Jul. 01, 2023 [Online]. Available: <https://ecgwaves.com/topic/ekg-ecg-leads-electrodes-systems-limb-chest-precordial/>
- [60] J. T. Tikkanen, T. Kentt , K. Porthan, H. V. Huikuri, and M. J. Junttila, "Electrocardiographic T wave abnormalities and the risk of sudden cardiac death: The Finnish perspective," *Ann. Noninvasive Electrocardiol., Official J. Int. Soc. Holter Noninvasive Electrocardiol., Inc.*, vol. 20, no. 6, pp. 526–533, Nov. 2015.
- [61] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis., Venice, Italy, October 22–29, 2017*, 2017, pp. 618–626, doi: [10.1109/ICCV.2017.74](https://doi.org/10.1109/ICCV.2017.74).
- [62] C. Grafton-Clarke et al., "Cardiac magnetic resonance left ventricular filling pressure is linked to symptoms, signs and prognosis in heart failure," *ESC Heart Failure*, vol. 10, no. 5, pp. 3067–3076, Oct. 2023.
- [63] M. Cikes et al., "Machine learning-based phenogrouping in heart failure to identify responders to cardiac resynchronization therapy," *Eur. J. Heart Failure*, vol. 21, no. 1, pp. 74–85, Jan. 2019.
- [64] S. Biton et al., "Atrial fibrillation risk prediction from the 12-lead electrocardiogram using digital biomarkers and deep representation learning," *Eur. Heart J. Digit. Health*, vol. 2, no. 4, pp. 576–585, Dec. 2021.



Chen Chen received the MSc and PhD degrees in advanced computing from Imperial College London, in 2016 and 2022, respectively. From 2022, she worked as a research associate with Imperial College London. After that, she joined Oxford BioMedIA group, University of Oxford in 2023. Her research interests include medical image analysis and multi-modal AI.



Lei Li received the BS degree from the Department of Medical Information Engineering, Sichuan University, in 2016 and the PhD degree from the School of Biomedical Engineering, Shanghai Jiao Tong University, in 2021. She is working with the Department of Engineering Science, University of Oxford. Her research interests include cardiac image analysis and multi-modal AI.



Marcel Beetz is currently working towards the PhD degree with the Institute of Biomedical Engineering (IBME), Department of Engineering Science, University of Oxford. His research interests include the development and application of machine learning models for the analysis of multimodal medical data.



Analysis, etc.

Abhirup Banerjee (Member, IEEE) received the BSc (Hons) and master's degrees in statistics from the Indian Statistical Institute, Kolkata, India and the PhD degree in computer science from the Indian Statistical Institute, in 2017. He joined the University of Oxford as postdoctoral researcher in 2017. He received the prestigious Royal Society University Research Fellowship and started as the Faculty member with the Department of Engineering Science, University of Oxford, U.K., in 2022. His research interests include Multimodal Machine Learning, Biomedical Image



Ramneek Gupta received the doctor of Philosophy in bioinformatics from the Technical University of Denmark 1997–2001. He is the Digital director, Global Drug Discovery with Novo Nordisk Research Center Oxford. He brings experience from previous roles with Technical University of Denmark and Eli Lilly and Company. His research focuses on using in silico pipelines, novel algorithms and advanced machine learning activities to identify novel targets for precision medicine and other research projects.



Vicente Grau is a professor of biomedical image analysis with the Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford. His research focuses on the development of novel computational methods for the analysis of biomedical images, with a particular interest on the interaction between these and models of cardiovascular function.