

HeartLLM: Discretized ECG Tokenization for LLM-Based Diagnostic Reasoning

Jinning Yang^{1,2}, Wenjie Sun^{1,2}, Wen Shi^{3,4*}

¹Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, Shenzhen, Guangdong, China

²University of Chinese Academy of Sciences, Beijing, China

³Department of Neurology, Massachusetts General Hospital, Boston, MA, USA

⁴Department of Neurology, Harvard Medical School, Boston, MA, USA

jn.yang@siat.ac.cn, wshi3@mgh.harvard.edu

Abstract

Electrocardiography (ECG) plays a central role in cardiovascular diagnostics, yet existing automated approaches often struggle to generalize across clinical tasks and offer limited support for open-ended reasoning. We present *HeartLLM*, a novel framework that integrates time-series (TS) and language modeling by enabling large language models (LLMs) to process 12-lead ECG signals for clinical text generation tasks. Our approach discretizes continuous ECG embeddings into quantized codes using a lead-wise encoder and quantization module. These quantized codes are then mapped to an extended ECG vocabulary to form ECG tokens, enabling the model to process both ECG and natural language inputs within a unified framework. To bridge the modality gap, we pretrain the model on an autoregressive ECG token forecasting task, allowing the LLM to capture temporal dynamics through its inherent language modeling capability. Finally, we perform instruction tuning on both ECG question answering and diagnostic report generation. Without modifying the core model, *HeartLLM* achieves strong performance across tasks while maintaining generalization to out-of-distribution settings. Extensive experiments demonstrate the effectiveness of each component and highlight the potential of integrating discretized ECG tokens into LLMs for medical reasoning.

Code — <https://github.com/yangjinning/HeartLLM>

Introduction

Cardiovascular diseases (CVDs) are the leading cause of mortality and morbidity worldwide (Di Cesare et al. 2023), and electrocardiography (ECG) remains a fundamental tool for early diagnosis and monitoring. With the increasing volume and complexity of ECG data in clinical settings, computer-aided interpretation has become an essential supplement to clinical expertise. Recent advances in deep learning have significantly improved performance in clinical decision support and ECG classification tasks (Ge et al. 2024; Xue et al. 2023a; Guhdar, Mohammed, and Mstafa 2025; Na et al. 2024a; Nejedly et al. 2021; Oh et al. 2022; Xue et al. 2023b), surpassing traditional signal-processing-based pipelines (Chen et al. 2022a). However, most existing models are tailored to closed-set classification and rely heavily

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

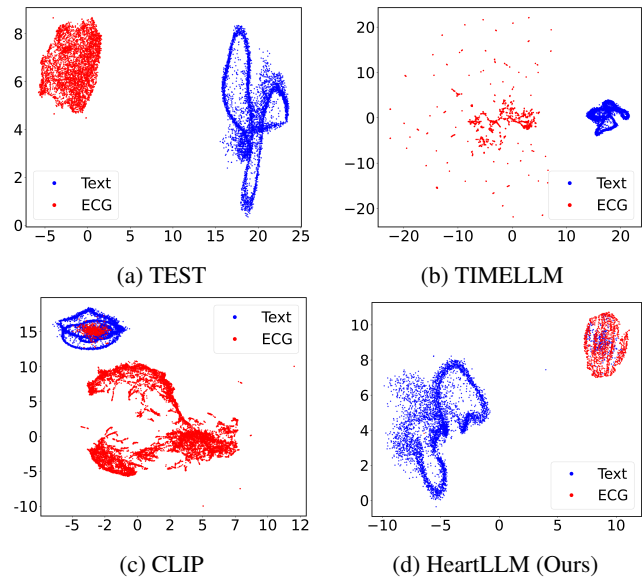


Figure 1: **UMAP visualizations of ECG and text embedding distributions across different models.** (a-b) TEST and TIMELLM exhibit clear modality separation. (c) CLIP shows partially aligned clusters, but still fragmented modality boundaries. (d) HeartLLM aligns all ECG representations with text in a shared semantic space, indicating effective modality unification without explicit contrastive pairing.

on supervised training. When confronted with new diagnostic categories or unseen clinical tasks, they require extensive retraining with labeled data, limiting their scalability and adaptability. Moreover, these models are typically designed for narrow classification objectives and cannot support more flexible reasoning tasks such as open-ended question answering or narrative report generation.

In parallel, large language models (LLMs) have demonstrated impressive generalization in vision-language domains (Alayrac et al. 2022; Huang et al. 2023), prompting efforts to extend such capabilities to physiological TS data like electrocardiograms (ECGs). However, this extension poses unique challenges: unlike symbolic and semantically structured text, ECG signals are continuous, noisy, and

lack explicit semantic boundaries. These fundamental differences complicate the integration of ECG and text into a unified representation space. Recent studies have explored multimodal LLMs for clinical applications, including ECG-based question answering (Liu et al. 2024a; Zhao et al. 2024) and report generation (Wan et al. 2024). Most adopt a two-stage paradigm where an ECG encoder produces continuous representations that are subsequently consumed by a language model. However, this design suffers from two major limitations. First, the modality gap between continuous ECG embeddings and discrete textual semantics remains underexplored. Prior alignment strategies in time series–language modeling (Jin et al. 2024; Sun et al. 2024) or contrastive learning in vision–language settings (Radford et al. 2021) often fail to achieve consistent cross-modal alignment. These limitations, when directly applied to ECG, result in either overly clear or fragmented modality boundaries, as illustrated in Figure 1. Second, ECG signals exhibit high redundancy due to waveform repetition and stable low-variance segments. Directly feeding such continuous features into LLMs can lead to overfitting and reduced generalization to rare or subtle signal variations. These challenges highlight the need for a more structured and symbolic interface between ECG signals and language models.

To address the challenges of bridging physiological signals and language models, we introduce *HeartLLM*, a unified framework that enables LLMs to reason over ECG signals in an open-ended, instruction-driven manner. Instead of relying on paired ECG–text supervision or complex cross-modal alignment losses, *HeartLLM* represents ECGs as discrete symbolic tokens, naturally compatible with the LLM’s vocabulary.

Our method is built on three core contributions: *First*, we design a **lead-wise encoder** that processes each ECG lead independently using an identical encoder built from temporal inception enhancement blocks. This architecture captures waveform-specific temporal patterns while avoiding inter-lead interference, resulting in precise and interpretable representations. *Second*, we introduce a **discretization-based ECG tokenizer** that converts continuous encoder outputs into symbolic tokens via a fixed-scale quantizer (FSQ) (Mentzer et al. 2023). These tokens are optimized through autoregressive pretraining and embedded in a shared space with language, enabling seamless integration with LLMs without contrastive learning or paired ECG-text data. *Third*, we apply **lightweight instruction tuning** for ECG report generation and question answering. By updating only low-rank modules and prompting with structured clinical inputs, our method achieves strong zero-shot generalization across tasks and datasets with minimal supervision.

HeartLLM is evaluated on two ECG understanding benchmarks: question answering and diagnostic report generation. It achieves state-of-the-art performance across multiple datasets and exhibits strong zero-shot generalization. Further analysis confirms the contribution of each design component, and visualizations illustrate interpretable attention over clinically meaningful waveform regions.

Related Work

Cross-modal ECG Text Generation Recent work has begun to explore the integration of LLMs into ECG interpretation tasks, particularly for open-ended question answering and diagnostic report generation. ECG-ReGen (Tang et al. 2025) adopts a retrieval-augmented framework that identifies similar historical cases to assist report generation in a zero-shot setting. PULSE (Liu et al. 2024a) constructs an instruction-tuning dataset based on ECG images, enabling the fine-tuning of image-text multimodal models. MEIT (Wan et al. 2024) introduces TS embeddings directly into the LLM’s attention layers to bridge modality differences, while (Tang et al. 2024) treats ECG features as soft prompts and applies meta-learning for rapid adaptation to new tasks. Despite these advances, most existing methods depend on paired ECG–text data to supervise modality alignment, which requires costly annotation and constrains applicability. In contrast, our approach aligns ECG and language representations through token-level pretraining, avoiding the need for paired supervision while maintaining flexibility across tasks.

LLMs for Time Series Tasks Large language models have been increasingly applied to TS tasks, typically via two main directions. One leverages LLMs to extract textual features to guide TS modeling (Wang et al. 2024a; Liu et al. 2025b,a). The other treats LLMs as core TS processors, mainly via two paradigms. The TS-as-text approach (Wang et al. 2025; Ansari et al. 2024; Wang et al. 2024b; Gruver et al. 2023) tokenizes numerical sequences into text for direct input into LLMs, but suffers from length and attention complexity limitations. The TS-embedding-based approach instead encodes signals into continuous embeddings and feeds them into frozen LLMs, often using prompt tuning (Zhou et al. 2023; Liu et al. 2024c). Some recent work aligns TS embeddings with language semantics, such as S2IP (Pan et al. 2024), TEST (Sun et al. 2024), and TIME-LLM (Jin et al. 2024), through contrastive or reprogramming strategies. Others like VITRO (Bellos, Nguyen, and Corso 2025) and UniTime (Liu et al. 2024b) propose task-specific vocabularies or reconstruction pretraining. However, most focus on forecasting or classification, with limited support for open-ended language tasks grounded in physiological signals. In contrast, our method discretizes ECG features into symbolic tokens, enabling autoregressive pretraining and unified ECG-text processing for clinical QA and report generation.

Problem Statement

Given a collection of 12-lead ECG signals $X = \{x_i\}_{i=1}^n$, where each $x_i \in \mathbb{R}^{L \times T}$ consists of $L = 12$ leads and $T = 5000$ time steps, we consider two tasks: ECG question answering (ECG-QA) and diagnostic report generation (ECG-Report).

ECG-QA. Each signal x_i is paired with a set of natural language questions $Q_i = \{q_i^{(k)}\}_{k=1}^{m_i}$ and corresponding answers $A_i = \{a_i^{(k)}\}_{k=1}^{m_i}$. We consider three question types:




Figure 2: Overview of the HEARTLLM framework. The model consists of three stages: (1) **ECG Tokenizer**, where 12-lead ECG signals are encoded by 12 lead-wise encoders and discretized with FSQ into symbolic ECG tokens; (2) **ECG Token Pretraining**, where an LLM is autoregressively pretrained on ECG tokens to jointly optimize the extended ECG vocabulary and the model using teacher-forcing next-token prediction; and (3) **QA/Report Tuning**, where the pretrained model is adapted to downstream tasks using lightweight LoRA tuning. Prompts include structured tabular features and textual instructions to guide generation or question answering.

- **Verify**: yes/no questions, $A_{\text{verify}} = \{\text{"Yes"}, \text{"No"}\}$. Example: “Does this ECG show any abnormal symptoms? Yes.”
- **Choose**: select one or more from predefined options, $A_{\text{choose}} = \{\text{"attr}_1", \text{"attr}_2", \dots, \text{"none"}\}$. Example: “Which noise does this ECG show, static noise or baseline drift? Baseline drift, static noise.”
- **Query**: open-ended, $A_{\text{query}} \subseteq S$, where S is the space of short free-text spans. Example: “What diagnostic symptoms does this ECG show? Incomplete right bundle branch block.”

Let $D_{\text{QA}} = \{(x_i, q_i^{(k)}, a_i^{(k)})\}_{i,k}$ denote the full set of QA triples. The goal is to learn a model $f_{\text{QA}} : (x_i, q_i^{(k)}) \mapsto \hat{a}_i^{(k)}$ that accurately predicts $\hat{a}_i^{(k)} \in A$.

ECG-Report. In this task, the model generates a diagnostic report r_i for each ECG signal x_i . The goal is to learn a mapping $f_{\text{report}} : x_i \mapsto \hat{r}_i$ that produces clinically coherent and diagnostically relevant reports.

These two tasks present complementary challenges:

ECG-QA requires precise grounding of language in signal features across diverse question formats, while ECG-Report demands the generation of coherent medical narratives from high-dimensional, often noisy physiological input. Successfully addressing both tasks requires the model to align temporal ECG patterns with linguistic representations in a flexible and generalizable manner.

Methods

Overview

We propose *HeartLLM*, a unified framework that enables large language models to perform open-ended clinical reasoning over ECG signals through symbolic representation and lightweight adaptation.

First, we design a lead-wise encoder that captures fine-grained temporal features independently for each ECG lead, enhancing representation fidelity without cross-lead interference. *Second*, we introduce a discretization mechanism that maps continuous features to symbolic tokens using an FSQ, forming an ECG-specific vocabulary compatible with

LLMs. These tokens are pretrained via autoregressive prediction, naturally aligning time series and text without paired supervision. *Finally*, HeartLLM is instruction-tuned on clinical tasks such as ECG report generation and question answering, using only LoRA modules for efficient adaptation. An overview of the full framework is shown in Figure 2.

ECG Tokenizer

Our ECG tokenizer converts a raw 12-leads ECG signal $\mathbf{X} \in \mathbb{R}^{L \times T}$ into a sequence of discrete tokens suitable for language modeling. This process includes three components: (1) 12 lead-wise encoders that extract multi-scale TS representations from each lead, which are then aggregated to capture the global cross-lead context; (2) a fixed-scale quantizer (FSQ) that first compresses TS representations into a low-dimensional latent space, discretizes this latent space to obtain quantized codes, and finally dequantizes the codes back into the original TS representation latent space; (3) a decoder that mirrors the encoder architecture and, with a self-supervised reconstruction loss, ensures that the learned tokens preserve physiological information. The quantized codes are first converted into scalar indices, which are subsequently mapped to the ECG vocabulary. This vocabulary extends the original vocabulary of the LLM, allowing for unified modeling of text and ECG signals.

Lead-wise Encoder. Each lead x_ℓ is first z-normalized and independently encoded using a stack of convolutional blocks. The encoder consists of two key components: an Inception Block to capture diverse temporal patterns and a Temporal Inception Enhancement (TIE) Block for deep contextual modeling.

Inception Block. The Inception Block applies multiple 1D convolutions with kernel sizes $\{1, 3, 5, 7\}$ in parallel to extract multi-scale temporal features.

Temporal Inception Enhancement (TIE) Block. The outputs from the Inception Block are summed and passed through GroupNorm and a LeakyReLU activation, followed by a residual connection to preserve information flow:

$$\mathbf{h}^{(i)} = \text{LeakyReLU} \left(\text{GN} \left(\sum_k \text{Conv1D}_k(\mathbf{h}^{(i-1)}) \right) + \mathbf{h}^{(i-1)} \right).$$

TIE blocks are applied sequentially to capture long-range signal dependencies, followed by convolutional downsampling layers that progressively compress the temporal resolution. This hierarchical structure allows the encoder to balance fine-grained waveform preservation with contextual abstraction. The encoded representations from all lead-wise encoders are concatenated along the lead dimension to form the final latent ECG representation $\mathbf{Z} \in \mathbb{R}^{T' \times D'}$.

Fixed-Scale Quantization (FSQ). To discretize continuous ECG representations, each latent ECG representation at timestep t $\mathbf{z}_t \in \mathbb{R}^{D'}$ is projected into a lower-dimensional space \mathbb{R}^D using a linear layer with parameters $\mathbf{W} \in \mathbb{R}^{D \times D'}$ and $\mathbf{b} \in \mathbb{R}^D$:

$$\mathbf{z}_t^{\text{cont}} = \sigma(\mathbf{W}\mathbf{z}_t + \mathbf{b}),$$

where $\sigma(\cdot)$ is the sigmoid function. The quantized code is obtained by uniformly quantizing each projected dimension into K discrete levels:

$$\mathbf{z}_t^{\text{disc}} = \frac{1}{K-1} \cdot \text{round}(\mathbf{z}_t^{\text{cont}} \cdot (K-1)),$$

with rounding applied element-wise. A straight-through estimator (Bengio 2013) is used to enable gradient flow through the quantization step. The quantized code $\mathbf{z}_t^{\text{disc}}$ is then projected back to the latent space of \mathbf{z}_t , yielding $\mathbf{z}'_t \in \mathbb{R}^{D'}$.

Token Indexing and Vocabulary Construction. To generate symbolic ECG tokens, the quantized code $\mathbf{z}_t^{\text{disc}}$ is mapped to a unique token index using positional base conversion:

$$\text{token_id}_t = \sum_{i=1}^D v_{t,i} \cdot K^{i-1},$$

where $v_{t,i} = \text{round}(z_{t,i}^{\text{disc}} \cdot (K-1)) \in \{0, 1, \dots, K-1\}$. This yields $\text{token_id}_t \in \{0, \dots, K^D - 1\}$. These indices are used to look up embedding vectors from an ECG-specific vocabulary:

$$\text{tok}_t^{\text{ECG}} = \text{ECG_Vocabulary}[\text{token_id}_t],$$

where $\text{ECG_Vocabulary} \in \mathbb{R}^{|\mathcal{V}_{\text{ECG}}| \times d_{\text{LLM}}}$ and $|\mathcal{V}_{\text{ECG}}| = K^D$.

Decoder and Autoencoder Objective. To ensure that discretization preserves physiologically meaningful information, we employ a symmetric decoder mirroring the encoder structure. It uses transposed 1D convolutions for temporal upsampling and stacked TIE blocks to reconstruct fine-grained waveform details. The reconstruction is trained using a mean squared error loss:

$$\hat{X} = f_{\text{dec}}(\mathbf{Z}'), \quad \mathcal{L}_{\text{recon}} = \frac{1}{LT} \sum_{\ell=1}^L \sum_{t=1}^T (x_{\ell,t} - \hat{x}_{\ell,t})^2.$$

ECG Tokenization for Language Modeling. The resulting ECG token sequence $\{\text{tok}_t^{\text{ECG}}\}_{t=1}^{T'}$ can be used as input to a language model. ECG tokens are initialized from the same distribution as the LLM’s text tokens. In the next stage, we leverage these ECG token sequences for autoregressive pretraining, using the same pretraining objective as for text tokens. Therefore, ECG token embeddings are drawn from the same embedding space as text tokens, the LLM can process ECG and text in a unified fashion.

Autoregressive Pretraining on ECG Tokens

To integrate ECG signals into the language modeling framework, we pretrain the model on an autoregressive forecasting task over discretized ECG tokens. This allows the LLM to build temporal dependencies within ECG sequences and learn to embed symbolic ECG tokens into its semantic space.

Given an input ECG signal \mathbf{X} , the tokenizer converts it into a discrete token sequence $\{\text{tok}_t^{\text{ECG}}\}_{t=1}^{T'}$, where each

token is embedded in the same space as textual input. During pretraining, a random slice of this sequence is selected and split into a historical context and a prediction target (typically in a 9:1 ratio). The LLM is trained to autoregressively predict the next ECG tokens from the context:

$$\mathcal{L}_{\text{AR}} = \frac{1}{\sum_{i=1}^N |\mathcal{M}_i|} \sum_{i=1}^N \sum_{t \in \mathcal{M}_i} \text{CE}(\text{LM}_{\theta}(\text{tok}_{i,\leq t}^{\text{ECG}}), \text{tok}_{i,t+1}^{\text{ECG}}),$$

where N is the batch size, \mathcal{M}_i is the set of prediction steps, $\text{LM}_{\theta}(\cdot)$ is the LLM output logits over the ECG vocabulary, and $\text{CE}(\cdot, \cdot)$ is the cross-entropy loss.

Training Strategy. The pretraining data is constructed from the MIMIC-IV-ECG database (Gow et al. 2023), yielding over 1.5 million time-series slices. To enable parameter-efficient adaptation, we fine-tune only a subset of the LLM parameters: the ECG embedding table and output classification head are fully updated, while the linear projections within the transformer blocks (e.g., QKV and feedforward layers) are adapted using low-rank LoRA modules. This strategy allows the model to efficiently acquire ECG-specific temporal structure while preserving the general language capabilities of the backbone model.

Instruction Tuning for Clinical Tasks

In the final stage, we fine-tune the model on downstream clinical tasks, including ECG diagnostic report generation and open-ended question answering. This stage enables the pretrained model from the previous stage to align symbolic ECG representations with textual outputs in task-specific contexts. Following common instruction-tuning practice, the embedding table and output classification head are frozen, and only the linear projection layers of the LLM (e.g., QKV and feed-forward layers) are updated with LoRA.

Each training sample is structured as a prompt that concatenates five components: (1) a brief dataset description (e.g., “The dataset contains electrocardiogram (ECG) time-series data.”), (2) a task-specific instruction (e.g., “Based on the given ECG signal embeddings, generate an ECG diagnostic report.”), (3) tabular features (including patient ID, age, sex, height, weight, recording date, nurse ID, device ID, and recording site), (4) an optional natural language question when the task is ECG-QA, and (5) ECG indicators (e.g., $\langle | \text{start_ecg} | \rangle$ and $\langle | \text{end_ecg} | \rangle$) that mark the position for inserting ECG token sequences. The model is trained to autoregressively generate either a diagnostic report or an answer span, depending on the task. The supervision signal comes from expert-written reports or curated answers, enabling the model to learn both structured reasoning and domain-specific terminology. This unified format allows the model to flexibly handle both classification-type questions and free-form report generation, without needing task-specific heads or architectural changes.

Experiments

Experimental Settings

Datasets and Tasks. We evaluate our method on two ECG-related language generation tasks: ECG-based ques-

tion answering (ECG-QA) and diagnostic report generation (ECG-Report). The ECG-QA task follows the setup of (Tang et al. 2025), where 10,598 questions are curated based on six types of ECG attributes—extra systole, heart axis, signal noise, numerical indicators, SCP codes, and infarction stages—using recordings from PTB-XL (Wagner et al. 2020) and MIMIC-IV-ECG (Gow et al. 2023). Only single-ECG questions are considered. The ECG-Report task uses the same datasets with official diagnostic reports. Each dataset is split into training, validation, and test sets in a 7:1:2 ratio by patient ID to prevent data leakage. For MIMIC-IV-ECG, only ECGs without missing values are retained, and one ECG is randomly selected per patient, resulting in 78,358 ECGs. PTB-XL contains 21,797 ECGs.

Baselines. We compare our method with a range of TS-Text multimodal models. These include two TS-as-text baselines: ChatTime (Wang et al. 2025) and LLMTIME (Gruver et al. 2023), which directly feed raw numerical sequences into language models. We further consider five TS-embedding-based methods: TIME-LLM (Jin et al. 2024), TEST (Sun et al. 2024), MEIT (Wan et al. 2024), and LLaVA (Liu et al. 2023) with either ST-MEM (Na et al. 2024b) or W2V-SMSC-RLM (Oh et al. 2022) as the ECG encoder. In addition, we include two strong multimodal baselines specific to ECG-QA: Fusion Transformer (Oh et al. 2023) and M3AE (Chen et al. 2022b). Note that Fusion Transformer and M3AE do not support text generation and are thus excluded from ECG-Report evaluation. Moreover, M3AE requires pretraining with classification labels, which are not available in MIMIC-IV-ECG.

Evaluation Metrics. We use exact match accuracy (EM ACC) to evaluate ECG-QA (Tang et al. 2025). For ECG-Report, standard text generation metrics are reported, including BLEU (Papineni et al. 2002), METEOR (Banerjee and Lavie 2005), and ROUGE-F1 (Lin 2004).

Implementation Details. All models based on large language models use LLaMA-3.2-3B as the default backbone, except ChatTime, which is built on LLaMA-2-7B. For methods involving LLM tuning, we apply 4-bit quantization with LoRA. LoRA rank is set to 64 with a scaling factor of 16. Models are trained for one epoch with a constant learning rate of $1e-4$ and a batch size of 10. For ECG discretization, the number of discrete levels is set to $K = 6$ and the dimension of the quantized code to $D = 4$, resulting in an ECG vocabulary of 1296 unique tokens. All experiments are conducted on one A100 GPU.

Main Results

Comparison with SOTA Table 1 summarizes the performance of HeartLLM and various baselines across ECG-QA and ECG-Report tasks on the MIMIC-IV-ECG and PTB-XL datasets. For ECG-QA, our model consistently achieves the highest exact match accuracy on both datasets, reaching **56.63** on MIMIC-IV-ECG and **54.98** on PTB-XL, with significant margins especially in the *Query* setting, which poses the greatest semantic challenge. This indicates that HeartLLM can more effectively capture and interpret the

Dataset	Task	Metric	Models									
			Ours	TIMELLM	TEST	MEIT	STMEM	W2V-CMSC-RLM	LLMTIME	ChatTime	Fusion Trans.	M3AE
MIMIC-IV-ECG	QA-Verify	EM ACC	78.12	66.00	67.86	67.69	69.73	71.57	69.97	68.78	<u>72.91</u>	–
	QA-Choose	EM ACC	69.73	28.86	41.21	50.18	49.39	<u>65.80</u>	55.77	52.37	64.38	–
	QA-Query	EM ACC	22.05	3.24	8.65	6.67	8.22	<u>15.24</u>	8.16	6.26	14.35	–
	QA-Average	EM ACC	56.63	32.70	39.24	41.51	42.45	<u>50.87</u>	44.63	42.47	50.55	–
PTB-XL	QA-Verify	EM ACC	72.66	67.46	66.43	65.34	68.70	<u>71.48</u>	67.96	67.65	64.27	67.73
	QA-Choose	EM ACC	59.74	29.42	47.25	48.06	46.61	<u>58.66</u>	45.89	43.19	50.95	31.15
	QA-Query	EM ACC	32.55	9.24	23.58	14.89	22.62	<u>31.57</u>	20.15	18.98	25.30	22.71
	QA-Average	EM ACC	54.98	35.37	45.75	42.76	45.98	<u>53.90</u>	44.67	43.27	46.84	40.53
MIMIC-IV-ECG	Report	BLEU-1	44.99	14.49	28.40	27.20	24.89	<u>35.89</u>	28.92	26.29	–	–
		BLEU-2	37.61	7.68	19.50	18.78	15.92	<u>28.00</u>	19.71	16.92	–	–
		BLEU-3	33.59	5.69	15.51	14.73	12.73	<u>24.14</u>	15.84	12.98	–	–
		BLEU-4	29.74	3.52	12.24	11.23	9.83	<u>20.60</u>	12.27	9.62	–	–
		METEOR	46.99	15.20	28.76	31.45	25.26	<u>37.33</u>	29.15	26.86	–	–
		ROUGE-1	58.67	26.57	41.85	39.42	36.43	<u>49.06</u>	41.81	37.96	–	–
		ROUGE-2	45.41	11.86	25.68	23.72	20.45	<u>35.02</u>	25.50	21.79	–	–
		ROUGE-L	58.63	26.39	41.82	39.41	36.40	<u>49.04</u>	41.77	37.92	–	–
PTB-XL	Report	BLEU-1	49.61	21.03	32.37	17.82	38.07	<u>47.65</u>	22.18	21.38	–	–
		BLEU-2	43.35	13.02	25.42	12.27	31.86	<u>40.98</u>	16.12	15.75	–	–
		BLEU-3	37.50	5.75	20.74	8.54	25.86	<u>34.87</u>	13.28	12.98	–	–
		BLEU-4	33.60	4.22	17.83	6.17	22.20	<u>30.92</u>	12.03	11.82	–	–
		METEOR	53.90	23.94	34.14	25.14	41.11	<u>53.33</u>	23.64	23.12	–	–
		ROUGE-1	59.53	38.74	47.27	34.94	49.39	<u>57.08</u>	38.35	37.58	–	–
		ROUGE-2	45.12	18.24	31.68	19.58	34.20	<u>42.31</u>	23.37	22.86	–	–
		ROUGE-L	59.24	38.69	47.01	34.81	49.30	<u>56.72</u>	38.29	37.56	–	–

Table 1: Evaluation results on ECG question answering and diagnostic report generation tasks. All metrics are reported as percentages (%), where higher values indicate better performance. For each dataset, the best-performing result in each task is highlighted in **bold**, while the second-best is underlined.

clinical semantics of ECG signals compared to both TS-as-text and embedding-based LLM methods. In the ECG-Report task, HeartLLM also obtains the best results across all evaluation metrics. On PTB-XL, it achieves a BLEU-4 score of **33.60**, METEOR of **53.90**, and ROUGE-L of **59.24**, outperforming competitive models like TIME-LLM and LLaVA+W2V-CMSC-RLM. Our method also achieves strong performance on MIMIC-IV-ECG in a zero-shot setting, indicating robust generalization across datasets. These results validate the advantage of our discretization-based representation and token-level alignment between TS and language modalities.

Ablation Study

We conduct ablations on PTB-XL to isolate the effects of three core components in HEARTLLM. Results are shown in Figure 3. Removing the discretization module (w/o DISC) causes consistent drops across all tasks, suggesting that symbolic ECG tokens are more compatible with LLMs than continuous embeddings. Without fine-tuning the LLM (w/o FT), performance notably declines—especially on QA-Choose—indicating that LoRA-based adaptation helps the model capture task-specific semantics. Excluding tabular features in the prompt (w/o TAB) slightly weakens generation and QA-Verify, reflecting their complementary role in providing patient context. Together, these results confirm that each component contributes meaningfully to downstream performance.

Analysis

Effectiveness of Our Encoder. We compare our lead-independent ResIncept encoder with four representative alternatives: MultiCNN (Oh et al. 2023), used in Fusion Transformer and MEIT; CausalCNN (Sun et al. 2024), the default encoder in TEST; Unet2D, a 2D convolution-based U-Net to extract channel-dependent features; and IndUnet1D, a structure that applies independent 1D U-Nets to each lead. As shown in Table 2, our encoder consistently outperforms all baselines across both datasets and all QA tasks. Notably, it yields the largest gains on QA-Query, which requires nuanced reasoning over subtle waveform features. Compared to MultiCNN and Unet2D, which either couple all leads early, our encoder better preserves fine-grained temporal features while maintaining lead-specific context. The improvement over CausalCNN suggests that fixed dilation patterns are less effective than our multi-scale design for ECGs.

Evaluating Latent Alignment and Signal Reconstruction. To assess whether the model captures clinically meaningful regions, we perform token-masking attribution: sequentially masking each ECG token and measuring the resulting drop in EM ACC to estimate its importance, which is then mapped back onto the original ECG waveform to visualize model attention. Figure 4(a) shows that when asked about *non-diagnostic T abnormalities*, the model emphasizes the T-wave segments. In contrast, for the *myocardial infarction* query in Figure 4(b), attention shifts toward P and QRS complexes—consistent with diagnostic criteria. This

Dataset	Task	Unet2D	Unet1D	MultiCNN	CausalCNN	Ours
MIMIC-IV-ECG	QA-Verify	75.09	76.58	73.09	75.20	78.12
	QA-Choose	67.29	68.01	63.66	68.07	69.73
	QA-Query	18.71	20.28	16.71	18.74	22.05
PTB-XL	QA-Verify	71.47	70.89	69.53	70.54	72.66
	QA-Choose	59.49	57.02	54.38	59.08	59.74
	QA-Query	29.49	31.14	29.50	29.45	32.55

Table 2: Performance comparison of different ECG encoders on QA tasks across MIMIC-IV-ECG and PTB-XL datasets. Each score indicates the EM ACC (%).

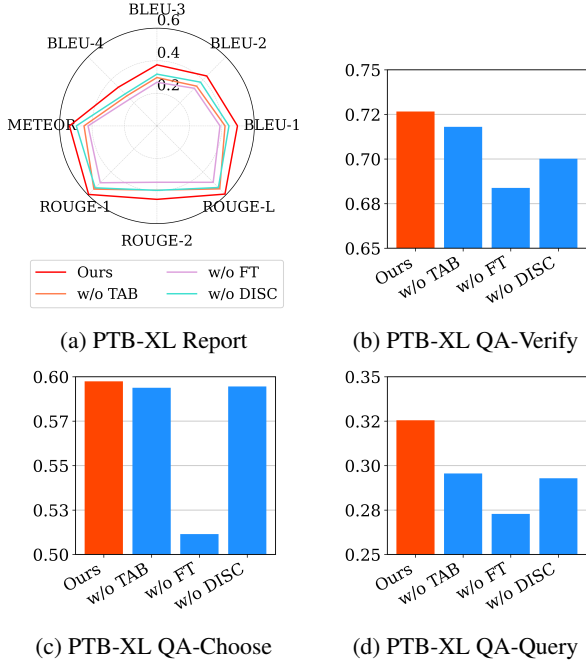


Figure 3: Ablation study on the PTB-XL. (a) Report generation evaluated by BLEU, METEOR, and ROUGE. (b–d) EM accuracy on QA-Verify, QA-Choose, and QA-Query tasks. We compare the full model (Ours) against three variants: w/o TAB, w/o FT, and w/o DISC.

context-dependent focus suggests that our alignment strategy successfully links ECG segments with query semantics. Additionally, Figure 4(c) demonstrates that discretization acts as an implicit regularizer, suppressing overfitting to noisy ECG segments while preserving clinically informative waveform morphologies through accurate reconstruction.

Cross-Dataset Generalization via Discretization. To examine the robustness of our learned representations, we evaluate zero-shot generalization by training on MIMIC-IV-ECG and directly testing on PTB-XL. As reported in Table 3, our method achieves the highest QA-Verify and QA-Query scores among all baselines, and the best average EM accuracy overall. Notably, although ST-MEM slightly outperforms our model on QA-Choose, its QA-Query performance remains weak. In contrast, our discretized token representation provides more balanced gains across all QA

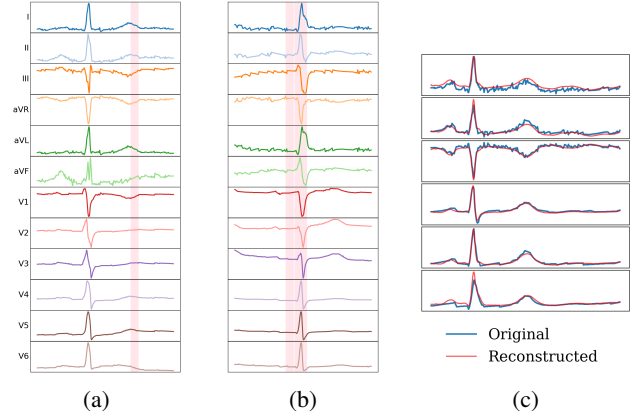


Figure 4: Visualization of latent alignment and signal reconstruction. (a-b) show attention maps for clinical queries, with shaded areas denoting high-impact segments. (c) shows reconstructed waveforms from symbolic tokens, demonstrating that discretization suppresses noise overfitting while preserving key information.

Methods	QA-Verify	QA-Choose	QA-Query	QA-Average
TIMELLM	66.81	26.62	11.62	35.02
TEST	31.96	26.45	1.40	19.94
MEIT	66.60	25.75	7.02	33.12
LLMTIME	66.50	37.88	7.22	37.20
ChatTime	67.66	39.95	6.59	38.07
ST-MEM	67.61	41.08	7.94	38.88
W2V-CMSC-RLM	61.82	36.46	4.85	34.38
w/o DISC	64.93	37.32	5.06	35.77
Ours	68.86	39.56	13.70	40.71

Table 3: Zero-shot performance on PTB-XL using models trained on MIMIC-IV-ECG.

subtasks. These results indicate that the proposed symbolic ECG representations, implicitly aligned with the LLM’s textual embedding space, enhance domain transfer and resist overfitting to dataset-specific artifacts.

Conclusion

We present *HeartLLM*, a unified framework that enables language-based clinical reasoning over ECG signals through symbolic tokenization and lightweight adaptation. By discretizing multi-lead ECG signals into a structured vocabulary, HeartLLM bridges the modality gap between continuous physiological data and discrete language representations. It avoids reliance on explicit ECG-text alignment or contrastive supervision, instead leveraging autoregressive pretraining and instruction tuning to support open-ended tasks such as question answering and report generation. Extensive experiments across multiple benchmarks demonstrate that HeartLLM achieves state-of-the-art performance and generalizes well to unseen ECG distributions. While effective, HeartLLM assumes offline processing and does not support real-time ECG analysis. Future work may extend our approach to streaming settings and incorporate medical knowledge to improve interpretability.

References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Ansari, A. F.; Stella, L.; Turkmen, C.; Zhang, X.; Mercado, P.; Shen, H.; Shchur, O.; Rangapuram, S. S.; Arango, S. P.; Kapoor, S.; et al. 2024. Chronos: Learning the language of time series.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 65–72.
- Bellos, F.; Nguyen, N. H.; and Corso, J. J. 2025. VITRO: Vocabulary Inversion for Time-series Representation Optimization. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Bengio, Y. 2013. Estimating or Propagating Gradients Through Stochastic Neurons. arXiv:1305.2982.
- Chen, S. W.; Wang, S. L.; Qi, X. Z.; Samuri, S. M.; and Yang, C. 2022a. Review of ECG detection and classification based on deep learning: Coherent taxonomy, motivation, open challenges and recommendations. *Biomedical Signal Processing and Control*, 74: 103493.
- Chen, Z.; Du, Y.; Hu, J.; Liu, Y.; Li, G.; Wan, X.; and Chang, T.-H. 2022b. Multi-modal masked autoencoders for medical vision-and-language pre-training. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 679–689. Springer.
- Di Cesare, M.; Bixby, H.; Gaziano, T.; Hadeed, L.; Kabudula, C.; McGhie, D. V.; Mwangi, J.; Pervan, B.; Perel, P.; and Piñeiro, D. 2023. World heart report 2023: Confronting the world’s number one killer. Technical report, World Heart Federation.
- Ge, Z.; Cheng, H.; Tong, Z.; He, Z.; Alhudhaif, A.; Polat, K.; and Xu, M. 2024. A knowledge-driven graph convolutional network for abnormal electrocardiogram diagnosis. *Knowledge-Based Systems*, 296: 111906.
- Gow, B.; Pollard, T.; Nathanson, L. A.; Johnson, A.; Moody, B.; Fernandes, C.; Greenbaum, N.; Waks, J. W.; Eslami, P.; Carbonati, T.; et al. 2023. MIMIC-IV-ECG: Diagnostic electrocardiogram matched subset. *Type: dataset*, 6: 13–14.
- Gruver, N.; Finzi, M.; Qiu, S.; and Wilson, A. G. 2023. Large language models are zero-shot time series forecasters. *Advances in Neural Information Processing Systems*, 36: 19622–19635.
- Guhdar, M.; Mohammed, A. O.; and Mstafa, R. J. 2025. Advanced deep learning framework for ECG arrhythmia classification using 1D-CNN with attention mechanism. *Knowledge-Based Systems*, 315: 113301.
- Huang, S.; Dong, L.; Wang, W.; Hao, Y.; Singhal, S.; Ma, S.; Lv, T.; Cui, L.; Mohammed, O. K.; Patra, B.; et al. 2023. Language is not all you need: Aligning perception with language models. *Advances in Neural Information Processing Systems*, 36: 72096–72109.
- Jin, M.; Wang, S.; Ma, L.; Chu, Z.; Zhang, J. Y.; Shi, X.; Chen, P.-Y.; Liang, Y.; Li, Y.-F.; Pan, S.; and Wen, Q. 2024. Time-LLM: Time Series Forecasting by Reprogramming Large Language Models. In *The Twelfth International Conference on Learning Representations*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, C.; Xu, Q.; Miao, H.; Yang, S.; Zhang, L.; Long, C.; Li, Z.; and Zhao, R. 2025a. Timecma: Towards llm-empowered multivariate time series forecasting via cross-modality alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 18780–18788.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. arXiv:2304.08485.
- Liu, P.; Guo, H.; Dai, T.; Li, N.; Bao, J.; Ren, X.; Jiang, Y.; and Xia, S.-T. 2025b. Calf: Aligning llms for time series forecasting via cross-modal fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 18915–18923.
- Liu, R.; Bai, Y.; Yue, X.; and Zhang, P. 2024a. Teach multi-modal llms to comprehend electrocardiographic images.
- Liu, X.; Hu, J.; Li, Y.; Diao, S.; Liang, Y.; Hooi, B.; and Zimmermann, R. 2024b. Unitime: A language-empowered unified model for cross-domain time series forecasting. In *Proceedings of the ACM Web Conference 2024*, 4095–4106.
- Liu, Y.; Qin, G.; Huang, X.; Wang, J.; and Long, M. 2024c. Autotimes: Autoregressive time series forecasters via large language models. *Advances in Neural Information Processing Systems*, 37: 122154–122184.
- Mentzer, F.; Minnen, D.; Agustsson, E.; and Tschannen, M. 2023. Finite scalar quantization: Vq-vae made simple.
- Na, Y.; Park, M.; Tae, Y.; and Joo, S. 2024a. Guiding masked representation learning to capture spatio-temporal relationship of electrocardiogram.
- Na, Y.; Park, M.; Tae, Y.; and Joo, S. 2024b. Guiding masked representation learning to capture spatio-temporal relationship of electrocardiogram.
- Nejedly, P.; Ivora, A.; Smisek, R.; Viscor, I.; Koscova, Z.; Jurak, P.; and Plesinger, F. 2021. Classification of ECG using ensemble of residual CNNs with attention mechanism. In *2021 computing in cardiology (CinC)*, volume 48, 1–4. IEEE.
- Oh, J.; Chung, H.; Kwon, J.-m.; Hong, D.-g.; and Choi, E. 2022. Lead-agnostic self-supervised learning for local and global representations of electrocardiogram. In *Conference on Health, Inference, and Learning*, 338–353. PMLR.
- Oh, J.; Lee, G.; Bae, S.; Kwon, J.-m.; and Choi, E. 2023. Ecg-qa: A comprehensive question answering dataset combined with electrocardiogram. *Advances in Neural Information Processing Systems*, 36: 66277–66288.
- Pan, Z.; Jiang, Y.; Garg, S.; Schneider, A.; Nevmyvaka, Y.; and Song, D. 2024. IP-LLM: Semantic space informed

- prompt learning with LLM for time series forecasting. In *Forty-first International Conference on Machine Learning*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Sun, C.; Li, H.; Li, Y.; and Hong, S. 2024. TEST: Text Prototype Aligned Embedding to Activate LLM’s Ability for Time Series. In *The Twelfth International Conference on Learning Representations*.
- Tang, J.; Xia, T.; Lu, Y.; Mascolo, C.; and Saeed, A. 2024. Electrocardiogram-Language Model for Few-Shot Question Answering with Meta Learning.
- Tang, J.; Xia, T.; Lu, Y.; Mascolo, C.; and Saeed, A. 2025. Electrocardiogram Report Generation and Question Answering via Retrieval-Augmented Self-Supervised Modeling. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Wagner, P.; Strodthoff, N.; Boussejot, R.-D.; Kreiseler, D.; Lunze, F. I.; Samek, W.; and Schaeffter, T. 2020. PTB-XL, a large publicly available electrocardiography dataset. *Scientific data*, 7(1): 1–15.
- Wan, Z.; Liu, C.; Wang, X.; Tao, C.; Shen, H.; Peng, Z.; Fu, J.; Arcucci, R.; Yao, H.; and Zhang, M. 2024. MEIT: Multi-modal electrocardiogram instruction tuning on large language models for report generation.
- Wang, C.; Qi, Q.; Wang, J.; Sun, H.; Zhuang, Z.; Wu, J.; Zhang, L.; and Liao, J. 2025. Chattime: A unified multi-modal time series foundation model bridging numerical and textual data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 12694–12702.
- Wang, F.; Wu, F.; Tang, Y.; and Yu, L. 2024a. CTPD: Cross-Modal Temporal Pattern Discovery for Enhanced Multi-modal Electronic Health Records Analysis.
- Wang, X.; Feng, M.; Qiu, J.; Gu, J.; and Zhao, J. 2024b. From news to forecast: Integrating event analysis in llm-based time series forecasting with reflection. *Advances in Neural Information Processing Systems*, 37: 58118–58153.
- Xue, B.; Said, A. S.; Xu, Z.; Liu, H.; Shah, N.; Yang, H.; Payne, P.; and Lu, C. 2023a. Assisting clinical decisions for scarcely available treatment via disentangled latent representation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5360–5371.
- Xue, B.; Shah, N.; Yang, H.; Kannampallil, T.; Payne, P. R. O.; Lu, C.; and Said, A. S. 2023b. Multi-horizon predictive models for guiding extracorporeal resource allocation in critically ill COVID-19 patients. *Journal of the American Medical Informatics Association*, 30(4): 656–667.
- Zhao, Y.; Kang, J.; Zhang, T.; Han, P.; and Chen, T. 2024. Ecg-chat: A large ecg-language model for cardiac disease diagnosis.
- Zhou, T.; Niu, P.; Sun, L.; Jin, R.; et al. 2023. One fits all: Power general time series analysis by pretrained lm. *Advances in neural information processing systems*, 36: 43322–43355.