CardioAgent: A Multi-Agent Framework for Integrated Cardiovascular Decision Support

Qian Cai
School of Health Science and Engineering
University of Shanghai for Science and Technology
Shanghai, China
233352446@st.usst.edu.cn

Xiangyong Kong*
School of Health Science and Engineering
University of Shanghai for Science and Technology
Shanghai, China
Kxy@usst.edu.cn

Minghui Chen
School of Health Science and Engineering
University of Shanghai for Science and Technology
Shanghai, China
cmhui.43@163.com

Enjun Zhou
School of Health Science and Engineering
University of Shanghai for Science and Technology
Shanghai, China
2335050631@st.usst.edu.cn

Yuhao Liu
Department of Biomedical Engineering, School of Life Science and Technology
Huazhong University of Science and Technology
Wuhan, Hubei, China
Hh15180393122@gmail.com

Xiankai Li
Department of Cardiology
Shanghai Tenth People's Hospital, School of Medicine, Tongji University
Shanghai, China
lixiankai@tongji.edu.cn

Abstract

The complexity of cardiovascular diseases (CVDs) poses significant challenges to precision diagnosis and treatment. Existing AI tools are often isolated "black-box" models, lacking interpretability and integration with clinical workflows. To address these limitations, this paper proposes CardioAgent, a multi-agent decision support framework for cardiovascular health management. The framework adopts a hierarchical architecture that coordinates agents for user intent understanding, context management, and task scheduling, while integrating three types of functional agents, patient services, clinical decision support, and research assistance, to provide role-customized services. The system incorporates a high-precision ECG analysis model based on Trans-Unet-ECG and Retrieval-Augmented Generation (RAG) technology, achieving a deep fusion of data-driven signal recognition and knowledge-driven evidence-based reasoning. Evaluation results show that the ECG model achieves over 99% accuracy on datasets such as MIT-BIH. Case studies demonstrate the effectiveness of the system in scenarios such as TAVR valve selection and patient consultation for palpitations. Through multi-agent collaboration and multimodal integration, CardioAgent establishes an interpretable and scalable intelligent ecosystem for cardiovascular health, offering a new paradigm for the deep integration of AI in clinical decision-making, patient management, and scientific research.

*Corresponding author.

ACM ISBN 979-8-4007-1511-2/2025/10
https://doi.org/10.1145/3777577.3777726

CCS Concepts

• **Computing methodologies** → Artificial intelligence; Distributed artificial intelligence; Multi-agent systems.

Keywords

CVDs, multi-agent, clinical decision-making

ACM Reference Format:
Qian Cai, Xiangyong Kong, Minghui Chen, Enjun Zhou, Yuhao Liu, and Xiankai Li. 2025. CardioAgent: A Multi-Agent Framework for Integrated Cardiovascular Decision Support. In *2025 6th International Symposium on Artificial Intelligence for Medical Sciences (ISAIMS 2025), October 24–26, 2025, Wuhan, China.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3777577.3777726

1 INTRODUCTION

Cardiovascular diseases (CVDs) encompass a broad spectrum of disorders involving the heart and blood vessels, including coronary artery disease, hypertension, heart failure, arrhythmias, and valvular calcification. These conditions are characterized by high prevalence and mortality worldwide, accounting for approximately 17.9 million deaths annually [1]. The prolonged disease course, diverse clinical manifestations, and multifactorial etiology of CVDs present significant challenges for accurate diagnosis and effective management. Recent advancements in artificial intelligence (AI) have demonstrated exponential growth and transformative potential across various domains, including healthcare[2]. The increasing availability of medical data, coupled with the evolving capabilities of AI, enables the analysis of complex datasets and the identification of critical patterns, offering new opportunities for the early detection and intervention of cardiovascular events. Despite notable progress in diagnostic methodologies[3], translating these capabilities into comprehensive, user-friendly, and actionable clinical support systems remains a formidable challenge. Current AI tools are often deployed as isolated "black-box" models, providing predictions without contextual interpretation or integration with established medical knowledge. Such fragmentation hinders clinical adoption and limits the potential impact of AI on improving patient outcomes.

The research proposes CardioAgent, a novel multi-agent architecture tailored for the cardiovascular domain, aiming to transcend the limitations of deploying isolated models. In this architecture, intelligent agents function as the cognitive core of the system, responsible for understanding user queries, managing contextual information, orchestrating task execution, and coordinating interactions among various tools. This approach enables the seamless integration of multiple functionalities, including automated signal analysis, personalized patient engagement, clinical decision support, and research assistance. Furthermore, CardioAgent incorporates a deep learning-based ECG analysis model, Trans-Unet-ECG, along with retrieval-augmented generation (RAG) for knowledge-enhanced reasoning. By integrating heterogeneous AI models with evidence-based knowledge sources, the framework facilitates a synergistic fusion of data-driven insights and evidence-based medicine. This results in a comprehensive, intelligent ecosystem for cardiovascular health management.

2 RELATED WORK

Previous research has explored various aspects of AI in cardiology, primarily focusing on task-specific applications. In the domain of ECG analysis, multiple deep learning architectures, such as convolutional neural networks (CNNs), Transformers, and U-Nets, have been developed for tasks including arrhythmia classification. Hasan et al. proposed a lightweight deep learning model combined with ensemble methods for efficient cardiovascular disease prediction from ECG images, achieving high accuracy and robustness in multi-class classification tasks [4]. Wilsgaard et al. introduced a novel biomarker, "ECG-derived age," by leveraging deep neural networks to infer biological age from digitized 12-lead ECGs, and demonstrated its independent predictive value for cardiovascular disease risk[5]. Yu et al. presented a hybrid framework that integrates a convolutional recurrent autoencoder with ResNet1d101 to capture spatiotemporal dynamics of ECG signals, achieving superior performance in multi-class arrhythmia classification [6]. Although these models have attained high precision, they are typically deployed as standalone systems. Other studies have investigated the use of large language models (LLMs) in medical information retrieval and summarization[7], as well as in clinical report generation[8]. However, integrating these fragmented capabilities into a coherent, user-centered system remains an open challenge. Multi-agent systems have been proposed in other domains for task orchestration and collaboration. For example, Ajmi et al. designed a computerized decision-support framework that combines multi-agent systems with dynamic optimization to coordinate and reschedule emergency department patient pathways in real time, thereby enhancing resource utilization and care quality[9].Nevertheless, the application of such approaches in the complex, multi-stakeholder environment of cardiovascular care remains at an early stage. Building upon these foundations, CardioAgent is designed to address the unique requirements of cardiovascular healthcare delivery through a specialized, agent-based framework.

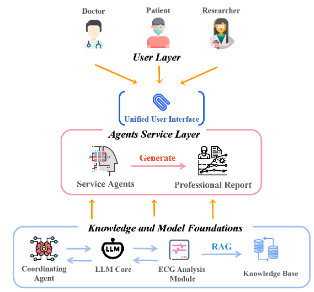

Figure 1: Overall Architecture of CardioAgent.

3 METHODOLOGY

3.1 Overall Architecture

The architecture of CardioAgent is organized into three primary layers, as illustrated in Figure 1.

The top layer is the user layer, which encompasses three major categories of end users: patients, clinicians, and medical researchers. These groups correspond to diverse needs, including health management, clinical decision support, and scientific exploration. The second layer is the interaction and coordination layer, which acts as a bridge between users and the system's intelligent functionalities. It provides unified access through multiple channels, including web platform, mobile applications, and API services, and integrates a central coordinating agent. This agent is responsible for parsing user requests, managing conversational context, orchestrating downstream agent tasks, and integrating multi-source responses to generate coherent and accurate outputs. The third layer is the agent service layer, which consists of multiple function-specific software agents, such as the patient service agent, clinical support agent, and research assistance agent. Each agent operates independently based on domain-specific knowledge and task logic, thereby enabling fine-grained functional specialization. At the bottom layer lies the knowledge and model foundation, which serves as the core infrastructure. It integrates structured medical databases, literature resources, and multimodal datasets, while incorporating LLMs, ECG analysis models, and RAG mechanisms. These components provide reliable reasoning capabilities and dynamic knowledge retrieval support for the upper-layer agents. This four-tiered architecture enables end-to-end collaboration from user interaction to intelligent decision-making, ensuring scalability, modularity, and clinical applicability of the system.

3.2 Core Components

3.2.1 Coordinating Agent. The coordinating agent serves as the central control unit of the CardioAgent system, responsible for intelligent parsing of user requests, task scheduling, and multi-agent collaborative decision-making. To achieve automated handling of complex user requests, it adopts a hybrid strategy that combines zero-shot task planning based on large language models with rule-based templates. Specifically, once a user request is parsed, the coordinating agent first invokes its core LLM (Qwen3: 30B in this

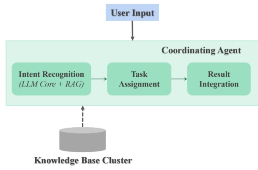

Figure 2: Workflow of the Coordinating Agent.

study) to perform task decomposition. This process is guided by carefully designed system prompts, which direct the LLM to break down complex queries into a series of logically executable atomic subtasks. The main workflow is illustrated in Figure 2.

Its core responsibilities include the following: First, through natural language understanding techniques, the coordinating agent performs intent recognition on user inputs, distinguishing different query types such as symptom descriptions and clinical decision-support requests, thereby providing precise guidance for subsequent processing. Second, it possesses robust context management capabilities, dynamically maintaining conversational states, tracking user interaction history and key information, and ensuring consistency and coherence across multi-turn dialogues. Building on this foundation, the coordinating agent executes task planning and decomposition, breaking down complex composite requests into a series of executable subtasks, such as literature retrieval, treatment comparison, and risk assessment. Based on task categories, it automatically assigns the most suitable service agent. Subsequently, through its task orchestration mechanism, the coordinating agent distributes subtasks to the corresponding specialized agents, such as the patient service agent, clinical support agent, or research assistance agent, while monitoring execution progress and outputs, thereby enabling coordinated information flow across agents. Finally, it performs result integration by semantically fusing and logically consolidating heterogeneous responses from multiple agents into a unified, accurate output that adheres to medical expression standards. To support such complex reasoning, the coordinating agent leverages an LLM as its core inference engine, augmented with RAG, enabling real-time access to the system's foundational knowledge and model layer, including professional medical databases, clinical guideline repositories, and the latest scientific literature. This design ensures that outputs are both generative and evidence-grounded, thereby enhancing the system's credibility and clinical applicability.

3.2.2 Service Agents. The system constructs three specialized service agents to deliver fine-grained, role-specific support tailored to different user groups. Their main workflow is illustrated in Figure 3.

The patient service agent focuses on providing end-to-end health management services for patients, encompassing health education, preliminary symptom assessment, personalized treatment recommendations, and rehabilitation tracking. Based on symptom descriptions and basic health data provided by the user, this agent generates accessible health interpretation reports and issues personalized alerts through risk prediction models, such as cardiovascular event risk warnings. In doing so, it enhances patient health literacy and self-management capacity. The clinical support agent is designed for healthcare professionals and aims to deliver evidence-based decision support. Its functions include, but are not limited to, recommendations on interventional device selection, comparative analysis of alternative treatment strategies, drug usage guideline queries, and dynamic updates of the latest clinical practice guidelines. By integrating authoritative medical knowledge bases with real-world clinical data, this agent ensures that its recommendations are both scientifically grounded and clinically actionable. The research assistant agent serves medical researchers, with the objective of accelerating scientific discovery. It supports automated literature retrieval, identification of similar studies, analysis of research gaps, and can even propose potential research hypotheses and experimental design schemes based on existing findings, thereby assisting researchers in efficiently completing early-stage tasks ranging from topic selection to methodological design.

Under the unified orchestration of the coordinating agent, these three service agents operate collaboratively while encapsulating their own knowledge bases and task logic. This design forms high-cohesion, low-coupling functional modules, ensuring both professional accuracy and domain specificity, while simultaneously enhancing the system's scalability and maintainability.

3.2.3 Knowledge and Model Foundation. The knowledge and model foundation layer constitutes the bottom support architecture of the system, providing reliable, updatable, and multimodal knowledge resources as well as domain-specific AI model capabilities for upper-layer agents. This layer is composed of two core components: a structured knowledge base cluster and a specialized AI model integration platform. For decision-support queries, a sentence-level RAG-based module is employed. The knowledge base integrates multiple authoritative medical data sources, including over 42,000 PubMed cardiovascular abstracts since 2020 and a cardiovascular drug and medical device manual (200,000 tokens). Texts are segmented using a 256-token sliding window (stride = 128), embedded via MedCPT-Retriever (768 dimensions), and indexed with FAISS-IVF1024. On a reserved QA set (n = 1000), single-stage retrieval achieves an accuracy of R@5 = 0.89, a level deemed sufficient to support downstream generation tasks.

For automated ECG analysis, the system employs Trans-Unet-ECG, an innovative multi-scale fusion network that integrates the strengths of CNNs, U-Net architectures, and Transformer mechanisms. Feature extraction begins with a CNN block utilizing convolution kernels of size (3,1) and stride 1 to capture local morphological characteristics of ECG signals. The number of kernels increases progressively from 64 to 512, with each layer followed by batch normalization and ReLU activation to stabilize training and introduce non-linearity. To effectively integrate multi-scale features and mitigate information loss during downsampling, the model incorporates a U-Net inspired structure consisting of an encoder, a bridge,

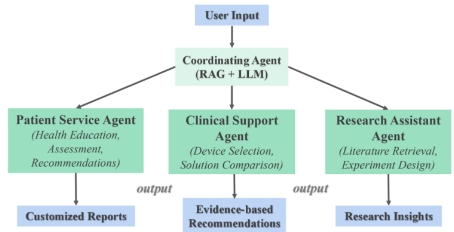

Figure 3: Workflow of the Service Agents.

and a decoder. The encoder abstracts features through downsampling, while the decoder restores spatial resolution via upsampling and employs skip connections to fuse shallow detail features from the encoder with deep semantic features in the decoder.

To capture long-range temporal dependencies and global contextual information, multiple multi-head attention (MHA) modules are embedded between the encoder and decoder. This mechanism enables the model to attend to different parts of the input sequence in parallel across multiple subspaces, thereby enhancing its ability to comprehend complex ECG patterns. The core self-attention mechanism is defined as:

$$Attention\,(Q, K, V) = SoftMax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{1}$$

where Q (Query), K (Key), and V (Value) are matrices derived from input feature transformations, and dkd_kdk is the dimensionality of the Key vector used for scaling to stabilize the SoftMax function. During training, the model employs cross-entropy loss and the Adam optimizer. Experiments were conducted on an NVIDIA GeForce RTX 4090 GPU for 30 epochs. The model's performance has been validated on multiple benchmark datasets, including the MIT-BIH Arrhythmia Database [10], the European ST-T Database , and the CU Ventricular Tachycardia Database [11]. Evaluation metrics include accuracy, precision , and F1-score:

$$accuracy = \frac{TP + TN}{TP + FP + TN + FN} \times 100\% \tag{2}$$

$$precision = \frac{TP}{TP + FP} \times 100\% \tag{3}$$

$$F1 = 2 \times \frac{pre \times rec}{pre + rec} \times 100\% \tag{4}$$

To ensure fair comparison and reproducibility, input formats and data augmentation processes were standardized. The input format was a single-lead, 10-second segment: 3600 samples (sampling rate 360 Hz) for MIT-BIH data, and 2500 samples (sampling rate 250 Hz) for ST-T and CU data. Normalization employed min–max scaling to map signal values into the range [−1, 1]. Data augmentation strategies included: (1) random cropping of ±5% of total length; (2) baseline wander simulation using 0–0.5 Hz sine waves with random amplitude σ =0.02 mV; (3) additive white Gaussian noise with SNR = 25 dB; and (4) random polarity inversion with probability p=0.5.

When user queries involve physiological signal interpretation, the coordinating agent passes the raw signal data into this model to obtain structured outputs such as rhythm type statistics, key waveform feature explanations, and report summaries. These outputs are seamlessly integrated into the final response, significantly improving accuracy and reliability in diagnostic support tasks. The foundation layer adopts a modular design, supporting continuous updates of knowledge bases and flexible integration of new models, thereby ensuring long-term maintainability and clinical adaptability of the system.

3.2.4 Data Flow and Integration. CardioAgent supports collaborative processing of multimodal inputs, enabling simultaneous reception of textual queries and physiological signal files, followed by semantic fusion and joint reasoning across modalities. When a user submits a request containing an ECG file, the coordinating agent first identifies the multimodal characteristics of the input and decomposes it into two parallel processing pathways. On one hand, the raw ECG signal is routed to a dedicated deep learning analysis module, which automatically performs waveform detection, rhythm classification, and abnormality identification, producing a structured analysis report that includes diagnostic results, key feature parameters, and confidence scores. On the other hand, the user's natural language query is parsed by the coordinating agent for intent recognition and contextual understanding. Subsequently, the system enters the multimodal fusion stage: the coordinating agent aligns the structured diagnostic results from the ECG model with the semantic information from the text query, and constructs a joint contextual representation by incorporating user role and historical interaction state. Building upon this representation, the system employs a retrieval-augmented generation mechanism to dynamically query the underlying knowledge base for content relevant to the current diagnosis and symptom description—such as the clinical significance of a specific arrhythmia, common triggers, complication risks, and therapeutic recommendations. Finally, the system integrates physiological data analysis results, medical knowledge evidence, and personalized context to generate coherent, accurate, and clinically interpretable responses. For instance,

Table 1: Performance of the Trans-Unet-ECG Algorithm Across Different Datasets.

Dataset	Accuracy	Precision	F1-score
MIT-BIH	0.9982	0.9667	0.9551
ST-T	0.9928	0.9615	0.9500
CU	0.9912	0.9600	0.9485

Table 2: Comparison results on accuracy of baseline models.

Method	MIT-BIH Accuracy	ST-T Accuracy	CU Accuracy
U-net	0.9981	0.9914	0.9881
CNN	0.9978	0.9910	0.9906
Transformer	0.9923	0.9922	0.9875
Trans-Unet-ECG	**0.9982**	**0.9928**	**0.9912**

when a patient uploads an ECG file while reporting palpitations, the system can not only identify the presence of paroxysmal atrial fibrillation, but also leverage the knowledge base to explain the association between this arrhythmia and palpitations, while providing lifestyle recommendations and medical consultation advice. This multimodal integration mechanism substantially enhances the system's understanding capacity and service depth in complex clinical scenarios, achieving a transition from "single-text QA" to a closed loop of perception–understanding–reasoning–explanation.

4 PRELIMINARY EVALUATION DISCUSSION

4.1 Framework Validation

The effectiveness of the CardioAgent framework depends on the performance of its constituent components. The deep learning model for ECG analysis has been rigorously validated on standard datasets (MIT-BIH, European ST-T, and CU Ventricular Tachyarrhythmia), achieving an accuracy exceeding 99%. The test results are summarized in Table 1.

Furthermore, compared with baseline models (U-Net, CNN, Transformer), our approach achieved higher accuracy, as shown in Table 2.

This establishes a solid foundation for the reliability of the system's core analytical capabilities. Integrating this model into the CardioAgent framework ensures that all provided decision support is grounded in accurate, data-driven signal interpretation.

4.2 Case Studies

We present two illustrative scenarios to demonstrate the operational logic of the framework.

Case Study 1: Clinical Decision Support

A cardiovascular specialist inquires: *"For a patient with a small annulus (annular area 390 mm^2) undergoing TAVR, how should one choose between Evolut and Sapien valves?"* The Coordinator Agent identified this query as a complex interventional decision-making scenario and accordingly initiated the clinical decision support process. Using RAG-based retrieval, the system integrated the latest evidence-based medicine, with particular emphasis on data from the SMART trial published in *JACC* in 2023. This was the first randomized controlled trial (n = 716) directly comparing Evolut PRO/PRO+ with SAPIEN 3/Ultra in patients with small annuli. The analysis revealed that the Evolut group demonstrated superior outcomes in the primary endpoint (1-year composite of death/disabling stroke: 9.4% vs. 10.6%, $P < 0.001$), along with hemodynamic advantages characterized by lower transvalvular gradients (7.7 mmHg vs. 12.6 mmHg) and fewer cases of valve dysfunction (9.4% vs. 41.6%). Furthermore, the system generated a structured comparative table (Table 3), quantifying the differences across three dimensions: hemodynamic performance, durability prediction, and anatomical compatibility. Based on these findings, personalized decision pathways were suggested: Evolut was recommended for younger patients, those with reduced baseline LVEF, or asymmetric annular calcification. Sapien was more suitable for patients with pre-existing right bundle branch block or prolonged PR interval, as well as cases requiring minimal femoral access sheath size or precise deployment control. In addition, the system provided supplementary post-procedural management recommendations, including 72-hour ECG monitoring, pre-discharge echocardiographic assessment, and 3-month CT follow-up. The final recommendation highlighted that, based on SMART trial evidence, Evolut demonstrates clear hemodynamic superiority in patients with a 390 mm^2 small annulus, and in the absence of specific conduction system risks, Evolut PRO/PRO+ should be considered the preferred option. This case study exemplifies the agent's ability to synthesize high-quality clinical trial evidence, perform multidimensional comparative analysis, and deliver individualized decision support, underscoring its advanced clinical utility in complex cardiovascular interventions.

Case Study 2: Patient Consultation

A patient uploaded an ECG recording and inquired: *"What does this mean? I have been experiencing palpitations."* The Coordinator Agent recognized this request as a patient consultation. It dispatched the ECG file to the ECG analysis module, which diagnosed ventricular premature contractions (VPCs). Subsequently, the Coordinator Agent employed RAG to retrieve relevant information regarding VPCs, including common etiologies, typical symptoms, and general management principles. If a patient profile was available, the system also checked for any known comorbidities. The Patient Service Agent then synthesized these findings into a clear,

Table 3: Structured Comparison of Valve Function Constructed by the System.

Comparison Dimension	Advantages of Evolut PRO/PRO+	Advantages of SAPIEN 3/Ultra
Hemodynamics	Lower transvalvular gradient, fewer valve dysfunction cases	Slightly lower trend for pacemaker implantation
Durability Prediction	More uniform stress distribution (supravalvular design)	Stronger radial force
Anatomical Fit	Self-expanding feature facilitates deployment in small annuli	Balloon expansion is more controllable

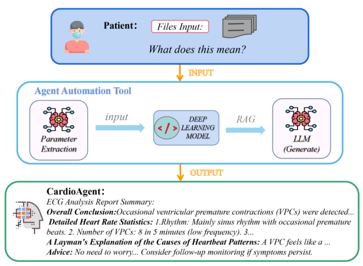

Figure 4: of a Patient Consultation Dialogue

non-technical report. The report explained the diagnostic results, reassured the patient that isolated VPCs are relatively common, and advised seeking medical consultation should the symptoms worsen or occur with higher frequency. This case study illustrates the integration of automated diagnostic interpretation with patient-friendly communication, demonstrating the framework's capacity to bridge clinical accuracy and accessibility for end-users. The dialogue process is illustrated in Figure 4.

4.3 Discussion

The design of CardioAgent embodies a modular, scalable, and semantically enriched intelligent system architecture tailored for cardiovascular disease support. Its core advantage lies in the agent-based hierarchical framework, which achieves functional decoupling and system flexibility, enabling each service agent to be independently developed, tested, and iteratively updated. This architecture demonstrates strong scalability, allowing seamless integration of new knowledge sources, AI models, or specialized agents in response to the continuous expansion of medical knowledge and the advancement of technologies, thereby adapting to the dynamic evolution of clinical practice. The adaptability of the system further benefits from the Coordinator Agent's robust capabilities in task planning and dynamic orchestration, enabling it to interpret diverse user intents, manage multi-turn conversational contexts, and provide role-specific, customized responses. Most importantly, the system realizes collaborative reasoning across data-driven and knowledge-driven paradigms. By combining deep learning models for automated analysis of physiological signals (e.g., ECG) with retrieval-augmented knowledge reasoning, CardioAgent can not only detect abnormal patterns but also generate explainable, evidence-based recommendations grounded in the latest clinical findings. This integration delivers more comprehensive and reliable support compared to single-paradigm approaches.

5 CONCLUSION

CardioAgent, as a multi-agent decision support framework for cardiovascular health management, successfully overcomes the limitations of traditional "black-box" AI models by achieving a deep integration of data-driven and knowledge-driven paradigms. Through the introduction of a Coordinator Agent for task orchestration and context management, the system can accurately interpret the multidimensional needs of patients, clinicians, and researchers, thereby delivering personalized and explainable intelligent services. The ECG analysis model, validated on multiple authoritative datasets with an accuracy exceeding 99%, provides a solid foundation for the clinical reliability of the framework. Furthermore, illustrative case studies—including TAVR prosthesis selection and patient consultation for palpitations—demonstrate the practicality and scalability of CardioAgent in complex clinical scenarios.

Looking forward, we plan to continuously enhance the Coordinator Agent's intent recognition and task planning capabilities, strengthen the mechanisms for dynamic knowledge base updates, and promote multi-center clinical trials to comprehensively assess the system's safety and usability. This framework offers a viable pathway for building a scalable and intelligent ecosystem for cardiovascular health management, advancing AI from isolated models toward collaborative and systematized clinical assistance.

References

[1] Elvas, L. B., Almeida, A., & Ferreira, J. C. (2025). The Role of AI in Cardiovascular Event Monitoring and Early Detection: Scoping Literature Review. *JMIR Med Inform*, *13*, e64349. 10.2196/64349.

[2] Shao, Z.-G., Lu, Y., Wang, X.-Y., Ma, J., Yuan, C.-D., Wang, F., Sang, Q., Xu, W.-T., & Yang, X. (2024). Trend analysis of stroke risk research based on three-dimensional metrics. *Biomedical Engineering Communications*, *3*, 12–16. https://doi.org/10.53388/BMEC2024014.

[3] Mohammed, R. K., Alrawi, A. T. H., & Dawood, A. J. (2025). A dual-approach to genomic predictions: leveraging convolutional networks and voting classifiers. *Biomedical Engineering Communications*, *4*, 1–6. https://doi.org/10.53388/BMEC2025002.

[4] Hasan, M. N., Hossain, M. A., & Rahman, M. A. (2025). An ensemble based lightweight deep learning model for the prediction of cardiovascular diseases from electrocardiogram images. *Engineering Applications of Artificial Intelligence*, *141*, 109782. https://doi.org/10.1016/j.engappai.2024.109782.

[5] Wilsgaard, T., Rosamond, W., Schirmer, H., Lindekleiv, H., Attia, Z. I., Lopez-Jimenez, F., Leon, D. A., & Iakunchykova, O. (2025). A New Biomarker of Aging Derived From Electrocardiograms Improves Risk Prediction of Incident Cardiovascular Disease. *JACC: Advances*, *4*(6, Part 1), 101764. https://doi.org/10.1016/j.jacadv.2025.101764.

[6] Yu, H., Lu, Y., & Zheng, S. (2024). Inferring spatial–temporal dynamics of ECG signals with deep neural networks for cardiovascular diseases diagnosis. *Biomedical Signal Processing and Control*, *97*, 106668. https://doi.org/10.1016/j.bspc.2024.

106668.
[7] Liu, Y., Li, X., Luo, Y., Du, J., Zhang, Y., Lv, T., Yin, H., Tang, X., & Liu, H. (2025). Toward a Large Language Model-Driven Medical Knowledge Retrieval and QA System: Framework Design and Evaluation. *Engineering*, *50*, 270-282. https://doi.org/10.1016/j.eng.2025.02.010.
[8] Liu, J., Koopman, B., Brown, N. J., Chu, K., & Nguyen, A. (2025). Generating synthetic clinical text with local large language models to identify misdiagnosed limb fractures in radiology reports. *Artificial Intelligence in Medicine*, *159*, 103027. https://doi.org/10.1016/j.artmed.2024.103027.
[9] Ajmi, F., Ajmi, F., Ben Othman, S., Zgaya-Biau, H., Dotoli, M., Renard, J.-M., & Hammadi, S. (2024). Real time patient scheduling orchestration for improving key performance indicators in a hospital emergency department. *Journal of Computational Science*, *82*, 102422. https://doi.org/10.1016/j.jocs.2024.102422.
[10] Moody, G. B., & Mark, R. G. (2001). The impact of the MIT-BIH Arrhythmia Database. *IEEE Engineering in Medicine and Biology Magazine*, *20*(3), 45-50. 10.1109/51.932724.
[11] Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C. K., & Stanley, H. E. (2000). PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation*, *101*(23), E215-220. 10.1161/01.cir.101.23.e215.