Toward Agentic AI in Cardiology: A Multimodal Multi-Agent Framework for Cardiovascular Risk Assessment

Hayat Bihri [1,[0000-0002-8783-4985]], Achraf El Mounafih[2], Aissam El Hasnaoui[2], Salma Azzouzi[1,[0000-0001-8122-4185]] . My El Hassan Charaf[1,2[0000-0002-0388-461X]]

[1]LaRI Laboratory, Faculty of Sciences, Ibn Tofail University, Kenitra, Morocco

[2]Master Informatique et Intelligence Artificielle, Faculty of Sciences, Ibn Tofail University, Kenitra, Morocco

Abstract— **Cardiovascular diseases (CVDs) remain the leading cause of mortality worldwide, highlighting the need for advanced and integrated decision-support systems. This paper proposes a conceptual multimodal multi-agent framework for cardiovascular risk assessment, designed to address the limitations of traditional centralized and unimodal approaches. The proposed architecture combines heterogeneous data sources, including structured clinical data, medical imaging, and unstructured textual information, within a distributed system of specialized agents. The framework consists of five main components: a Clinical Agent for risk estimation from structured data, an Image Agent for cardiomegaly detection from chest X-rays, a Narrative Agent for extracting insights from clinical text, a Fusion Agent for multi-modal reasoning, and a Decision Agent for generating final clinical recommendations. The system leverages state-of-the-art machine learning models and large language models to enable robust, interpretable, and context-aware analysis. This work is purely methodological and does not include implementation at this stage. It aims to provide a scalable, modular, and explainable foundation aligned with real-world clinical workflows. The proposed framework highlights the potential of agent-based and multimodal AI systems for next-generation cardiovascular decision support.**

I. Introduction

Cardiovascular diseases (CVDs) remain the leading cause of mortality worldwide, accounting for a substantial proportion of global deaths each year [1]. Therefore, early detection and continuous monitoring of cardiovascular risk factors are essential for improving patient outcomes and reducing healthcare burdens [2]. In recent years, artificial intelligence (AI), particularly machine learning and deep learning techniques, has demonstrated significant potential in enhancing cardiovascular prediction, diagnosis, and clinical decision-making [3],[4]. Recent advances in AI have led to the emergence of agentic AI systems characterized by autonomy, adaptive reasoning, and dynamic interaction with their environment [5]. The development of foundation models and large language models (LLMs) has further accelerated the adoption of intelligent agent-based architectures across domains such as healthcare, finance, and industry [6]. Agentic AI refers to systems capable of autonomously decomposing complex tasks, coordinating specialized agents, selecting appropriate tools, and executing multi-step reasoning processes with limited human intervention [7].

These capabilities are particularly relevant in cardiovascular healthcare, where medical data are inherently heterogeneous and multimodal, including electronic health records, laboratory tests, wearable sensors, medical imaging, and unstructured clinical narratives. Among these modalities, medical imaging plays a central role in cardiovascular diagnosis by providing detailed structural and functional information. However, integrating imaging data with clinical and textual information remains challenging for conventional centralized AI systems, especially in dynamic and distributed healthcare environments.

AI agents differ from traditional AI approaches through their ability to collaborate, adapt dynamically, and integrate heterogeneous models and external tools within distributed environments [8]. Recent studies in multi-agent AI have explored decentralized coordination, reinforcement learning-based collaboration, explainable neurosymbolic reasoning, and adaptive decision-making [9]. In particular, decentralized multi-agent reinforcement learning has been applied to supply chain optimization and energy management systems [10], [11], while neurosymbolic and brain-inspired multi-agent architectures have emphasized explainability, traceability, cooperation, and adaptive reasoning capabilities[12]. These developments highlight the growing importance of distributed intelligence and autonomous coordination in modern AI ecosystems. Despite these advances, most cardiovascular AI systems still rely on centralized architectures that face limitations regarding scalability, interoperability, adaptability, and multimodal integration. Furthermore, existing approaches often focus on isolated modalities or single-task learning, limiting their ability to support comprehensive and context-aware cardiovascular assessment.

To address these challenges, this paper proposes a conceptual multimodal multi-agent framework for cardiovascular risk assessment and clinical decision support. The proposed architecture relies on specialized autonomous agents dedicated to tasks such as clinical data analysis, medical image interpretation, narrative processing, multimodal fusion, and decision support. The framework integrates structured clinical data, medical imaging, and unstructured narratives within a distributed and collaborative ecosystem. Rather than presenting a fully implemented system, this work aims to establish a theoretical and methodological foundation for future research on agentic AI in cardiology. The main contributions of this work are:

- the design of a distributed multimodal multi-agent architecture for cardiovascular healthcare;
- the integration of heterogeneous clinical, imaging, and narrative data sources within a unified framework;
- the incorporation of specialized autonomous agents for distributed reasoning and decision support;
- the positioning of agentic AI as a promising paradigm for scalable, explainable, and context-aware cardiovascular risk assessment.

II. State of the Art

Artificial intelligence (AI) has emerged as a major driver of innovation in cardiovascular medicine, with applications spanning risk prediction, diagnosis, medical imaging analysis, clinical workflow optimization, and precision medicine. Recent studies highlight a significant transition from isolated predictive models toward integrated and multimodal AI systems capable of combining heterogeneous data sources, including medical imaging, electronic health records, physiological signals, and patient-generated data [13]. This evolution reflects the inherent complexity of clinical decision-making, which relies on the integration of diverse and complementary information. However, despite these advances, the integration of AI into routine clinical practice remains limited due to challenges related to generalizability, interoperability, safety, and regulatory constraints. In parallel, the emergence of large language models (LLMs) has broadened the scope of AI applications in cardiology. These models are being explored for patient interaction, automated interpretation, and decision support. Nevertheless, their adoption raises significant concerns regarding data privacy, diagnostic reliability, and their safe integration into clinical environments[14].

Medical imaging represents one of the most advanced and impactful domains of AI in cardiology. Deep learning techniques have demonstrated high performance across multiple imaging modalities, including cardiac computed tomography (CT), magnetic resonance imaging (MRI), echocardiography, and nuclear imaging, particularly for tasks such as segmentation, classification, phenotyping, and prognostic prediction [15]. These advances have enabled improved detection of cardiovascular abnormalities, reduced analysis time, and earlier diagnosis. A study published in 2024 in Nature Medicine reported an AI-based cardiac MRI system achieving an average AUC of 0.991 across multiple cardiovascular disease classes, illustrating near-expert performance while emphasizing the complementary role of AI in clinical practice [16]. However, most AI systems in medical imaging remain limited to a single modality or narrowly defined tasks. Their large-scale adoption is hindered by challenges such as lack of data standardization, insufficient external validation, limited interpretability, and difficulties in integration into clinical workflows. These limitations are particularly critical in cardiology, where imaging data must be combined with clinical and longitudinal patient information[17]. To address these challenges, multimodal AI approaches are gaining increasing attention. These approaches aim to integrate diverse data sources to better reflect real world clinical decision-making processes[18]. Although promising, multimodal AI in cardiology remains underdeveloped. Existing studies are limited in number, and real-world implementations are still scarce. Key challenges include data heterogeneity, fusion strategies, bias, security, and model generalization [17]. These issues highlight the need for architectures capable of efficiently handling distributed data while ensuring robust and interpretable results[13].

In this context, Multi-Agent Systems (MAS) represent a promising approach for modeling distributed and collaborative environments. A MAS relies on the coordination of multiple autonomous agents, each responsible for specific tasks, making it particularly suitable for complex healthcare systems. Prior work has demonstrated the usefulness of MAS in telemedicine, remote monitoring, and decision support [19]. In cardiology, MAS have been mainly applied to chronic disease management, particularly heart failure. These systems can assist clinicians and improve over time through learning mechanisms, while also depending on non-technical factors such as data privacy and healthcare organization [19]. Despite their potential, current MAS applications remain limited and do not fully leverage recent advances in deep learning, multimodal AI, and agent orchestration. Consequently, their integration into modern cardiovascular decision-support systems remains insufficient[20]. More recently, the concept of agentic AI has emerged as an evolution of intelligent agents. These systems are capable of reasoning, planning, tool usage, and coordinating multiple specialized modules. A recent review demonstrated that multi-agent systems outperform traditional LLMs in complex tasks, particularly when their architecture is aligned with task complexity[21]. In healthcare, these approaches offer promising perspectives for chronic disease management, real-time data integration, and advanced decision support. However, their adoption still faces significant challenges, including data quality, bias, explainability, privacy, and clinical integration. In medical imaging, agentic systems may improve reproducibility and interpretability but also raise ethical and regulatory concerns, especially in high-risk domains such as cardiology [22]

Finally, the literature highlights several critical gaps: the predominance of unimodal systems, the lack of real-world multimodal applications, the limited exploitation of modern MAS, and the emerging nature of agentic AI in cardiology, particularly for integrated architectures combining clinical, imaging, and patient-generated data [23], [24]. Therefore, the development of architectures combining multi-agent systems, multimodal AI, and advanced predictive models appears essential. Such distributed and modular approaches could better address the complexity of cardiovascular care pathways and represent a promising direction for next-generation intelligent healthcare systems [25].

III. Methodology

This work proposes a conceptual agentic artificial intelligence framework based on a multi-agent system (MAS) for cardiovascular risk assessment and clinical decision support.

Unlike traditional centralized AI systems, the proposed architecture adopts a distributed and modular paradigm in which specialized autonomous agents collaborate to process heterogeneous and multimodal healthcare data. The framework is particularly suited to cardiovascular medicine, where patient information is progressively available and inherently distributed across structured clinical records, medical imaging, and unstructured narratives.

A major challenge in this domain is the limited availability of large-scale multimodal datasets combining all modalities simultaneously. To address this limitation, the proposed framework relies on modality-specific datasets, allowing each agent to operate independently using specialized data sources. The architecture is composed of five main agents: (i) Clinical Agent, (ii) Image Agent, (iii) Narrative Agent, (iv) Fusion Agent, and (v) Decision Agent.

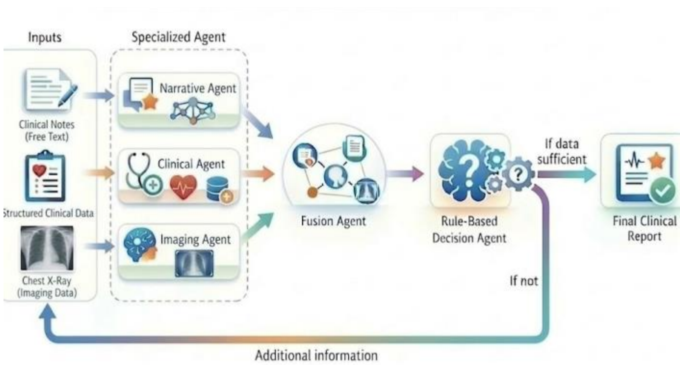

Fig. 1. A Multimodal Multi-Agent Cardiovascular Diagnostic Pipeline.

A. Clinical Agent

The Clinical Agent is designed to estimate cardiovascular risk using structured clinical data progressively collected throughout the patient care pathway. These data originate from multiple sources, including patient self-reports, measurements performed by healthcare professionals, and laboratory test results as they become available. This progressive acquisition reflects real-world clinical settings, where information is often incomplete and evolves over time. To model this dynamic process, the agent is implemented using a CatBoostClassifier. CatBoost [26] is particularly suitable for handling heterogeneous, noisy, and partially missing clinical data. It natively supports categorical variables without requiring complex preprocessing (e.g., one-hot encoding), demonstrates robustness to missing values, and mitigates overfitting through its ordered boosting mechanism, thereby improving generalization performance. These characteristics make it well-adapted to real-world medical datasets. The model is trained using both mandatory and optional clinical features. Mandatory variables, such as age, sex, chest pain type, and resting blood pressure, are assumed to be available at early stages. Additional features, including cholesterol, maximum heart rate, Oldpeak, and ST_Slope, are progressively incorporated as they become available. A missing data simulation strategy is applied during training to enhance robustness, enabling the model to generate initial predictions with limited inputs and refine them iteratively as new data are acquired. During inference, input data are structured in a standardized format, with appropriate handling of missing values. The Clinical Agent outputs a cardiovascular risk probability, a binary classification, an interpretable risk level, and a structured JSON summary, ensuring interpretability and seamless integration within the multi-agent system.

B. Image Agent

The Image Agent constitutes the visual analysis component of the system, designed to estimate the probability of cardiomegaly from chest X-ray images. By incorporating imaging data into the pipeline, it contributes to a more comprehensive and multimodal cardiovascular assessment. The agent is based on a DenseNet-121 convolutional neural network, selected for its dense connectivity pattern, which promotes feature reuse and mitigates the vanishing gradient problem. The architecture is adapted for binary classification by replacing the final layer with a single-node output. The model is initialized with pretrained weights obtained from cardiomegaly detection tasks.

During inference, input images undergo preprocessing, including RGB conversion, resizing to 224×224 pixels, and normalization. The model outputs a probability score via a sigmoid activation function. The Image Agent produces three primary outputs: cardiomegaly probability, a binary decision based on a threshold of 0.70, and a confidence score.

Additionally, a basic image quality assessment is performed using brightness analysis. To enhance interpretability, Grad-CAM visualizations are generated to highlight image regions contributing to the prediction, which is critical for clinical validation. All outputs are formatted in structured JSON to ensure traceability and interoperability within the multimodal framework.

C. Narrative Agent

The Narrative Agent processes unstructured clinical narratives and patient-generated textual information in order to extract clinically relevant cardiovascular knowledge. Since textual narratives may contain implicit symptoms, contextual information, and temporal medical descriptions, the agent relies conceptually on the Gemini 2.5 Flash large language model (LLM), selected for its advanced natural language reasoning capabilities.

To reduce hallucination risks and improve reliability, the agent operates through structured prompt engineering constrained to predefined medical extraction tasks. Standardized JSON outputs include extracted symptoms, cardiovascular risk factors, medical history elements, severity indicators, confidence scores, and uncertainty flags. Additional safeguards such as confidence-based filtering, rule-based post-validation, and semantic normalization are integrated to ensure consistency, explainability, and interoperability with the other agents.

C. Fusion Agent

The Fusion Agent represents the central coordination component of the framework. Its role is to aggregate, synchronize, and interpret outputs generated by the Clinical, Image, and Narrative Agents in order to produce a coherent multimodal cardiovascular assessment.

The orchestration process combines parallel and sequential coordination mechanisms. Specialized agents first operate independently on their respective modalities. Their outputs are then synchronized into a unified patient representation integrating predictions, confidence scores, uncertainty indicators, and contextual explanations.

To manage inconsistencies between modalities, the Fusion Agent incorporates a confidence-aware conflict-resolution mechanism. When contradictory or uncertain outputs are detected, the agent prioritizes clinically reliable modalities and generates uncertainty alerts requiring additional medical validation. The Fusion Agent is conceptually implemented using Gemini 2.5 Flash due to its reasoning and multimodal integration capabilities. It produces multimodal diagnostic hypotheses, severity assessments, confidence levels, missing-data alerts, and recommended clinical actions. Structured reasoning traces are also maintained to improve interpretability and support future clinician-oriented validation.

C. Decision Agent

The Decision Agent constitutes the final decision-support layer of the proposed architecture. Its objective is to transform multimodal outputs into clinically actionable recommendations while preserving transparency and clinician supervision. Rather than replacing medical expertise, the agent is designed as a clinician-support component.

The agent operates through a rule-based reasoning mechanism integrating severity level, confidence scores, multimodal hypotheses, uncertainty indicators, and recommended examinations. Based on these inputs, it determines urgency level, follow-up recommendations, need for specialist referral, and suggested clinical actions.

To reinforce transparency and accountability, the agent generates interpretable decision traces describing activated rules, involved modalities, and confidence levels associated with each recommendation. In situations involving uncertainty or conflicting multimodal evidence, the system triggers clinician-review alerts rather than producing fully autonomous decisions. This human-in-the-loop strategy contributes to clinical safety and safer management of uncertain AI outputs.

F. Explainability, Safety and System Limitations

Explainability and safety constitute essential requirements for agentic AI systems in healthcare. In the proposed framework, interpretability is addressed through structured outputs, confidence-aware reasoning, modality-specific explanations, and decision traces generated by the different agents. The Image Agent incorporates Grad-CAM visualizations, while the Narrative, Fusion, and Decision Agents generate structured reasoning outputs and contextual explanations in standardized JSON format.

Special attention is also given to LLM-related risks, including hallucinations, instability, and uncertainty. To mitigate these issues, the framework integrates conceptual safeguards such as restricted prompting, structured outputs, confidence filtering, post-processing validation, and rule-based consistency checks. In addition, uncertain or conflicting outputs are systematically flagged for clinician review through a human-in-the-loop strategy.

Despite its potential, the proposed framework remains conceptual and does not include a fully implemented prototype or experimental validation. Furthermore, multimodal integration may be affected by heterogeneity, incomplete information, temporal misalignment between modalities, and biases inherited from pretrained models and LLMs. Consequently, future work will focus on implementation, clinical validation, scalability assessment, and evaluation of explainability and safety mechanisms in real healthcare environments.

IV. Discussion

The proposed multi-agent framework introduces a conceptual paradigm for cardiovascular risk assessment by integrating heterogeneous and multimodal healthcare data within a distributed and collaborative architecture. Unlike traditional centralized systems, the proposed approach decomposes complex clinical tasks into specialized and interoperable agents dedicated to clinical analysis, imaging interpretation, narrative processing, multimodal fusion, and decision support. This modular organization is particularly suitable for healthcare environments, where patient information is progressively available, heterogeneous, and distributed across multiple sources.

One of the main strengths of the framework lies in its ability to combine structured clinical data, medical imaging, and unstructured textual information into a unified cardiovascular assessment. The Clinical Agent supports risk estimation under incomplete and evolving data conditions, while the Image Agent contributes visual biomarkers such as cardiomegaly detection.

The Narrative Agent enriches the analysis through contextual information extracted from clinical narratives. These multimodal outputs are synchronized and aggregated by the Fusion Agent to produce coherent diagnostic hypotheses and confidence-aware assessments, while the Decision Agent translates them into structured clinical recommendations. Such distributed reasoning may improve flexibility, interoperability, and explainability compared with conventional monolithic AI systems.

The framework also highlights the growing relevance of agentic AI and large language models (LLMs) in multimodal healthcare intelligence. The integration of LLM-based agents enables contextual reasoning, semantic interpretation, and structured knowledge extraction from unstructured narratives. However, despite these advantages, LLM-based systems also introduce important risks related to hallucinations, instability, uncertainty, prompt sensitivity, and potential biases inherited from pretrained data. In clinical settings, such limitations may affect reliability and patient safety if not properly controlled. To mitigate these risks, the proposed framework incorporates several conceptual safeguards, including restricted task-oriented prompting, structured JSON outputs, confidence-aware filtering, rule-based post-validation, uncertainty flags, and human-in-the-loop supervision. Rather than enabling fully autonomous decision-making, uncertain or conflicting outputs are systematically redirected for clinician review.

Another important challenge concerns multimodal patient-level data alignment in real-world deployment scenarios. Since the proposed agents rely on modality-specific datasets and independently processed information sources, effective synchronization between clinical records, medical imaging, and textual narratives remains a critical issue. In operational healthcare settings, this alignment would require standardized patient identifiers, temporal synchronization mechanisms, interoperable healthcare data standards, and robust data-governance strategies to ensure consistency across modalities. Although the proposed framework conceptually addresses multimodal coordination through the Fusion Agent, practical deployment would require additional validation regarding data integration reliability and cross-modal consistency.

The proposed architecture also raises important considerations regarding clinical accountability and responsibility. The Decision Agent is designed exclusively as a clinician-support component and not as an autonomous diagnostic authority. Final medical decisions remain under the responsibility of healthcare professionals, particularly in situations involving urgent referral, additional examinations, or high-risk cardiovascular predictions. To support transparency and accountability, the framework incorporates interpretable decision traces describing the reasoning pathway, activated rules, confidence levels, and modalities contributing to each recommendation. This explainability-oriented design may facilitate clinician validation and improve trust in AI-assisted decision-support systems.

Despite its conceptual contributions, several limitations must be acknowledged. First, the current work remains theoretical and does not include a fully implemented or clinically validated prototype. Second, multimodal integration may be affected by incomplete information, heterogeneous data quality, temporal misalignment, and biases inherited from pretrained models and LLMs. Third, the orchestration of multiple agents and large AI models may introduce computational complexity and latency challenges in real-world healthcare infrastructures.

V. Conclusion

This work proposes a novel conceptual framework for cardiovascular risk assessment based on a multi-agent and multimodal AI architecture. By integrating structured clinical data, medical imaging, and unstructured textual information through specialized agents, the proposed approach highlights the potential of distributed and collaborative systems to address the complexity of real-world healthcare environments.

The framework emphasizes modularity, scalability, and interpretability, while leveraging recent advances in machine learning and large language models to enable comprehensive and context-aware analysis. It also aligns with current trends toward multimodal and agent-based AI in medicine. However, this study remains theoretical and focuses on methodological design rather than implementation. Future work will involve the practical development, experimental validation, and evaluation of the proposed system in real clinical settings, as well as addressing challenges related to data integration, reliability, and ethical compliance. Overall, this framework provides a promising foundation for next-generation intelligent decision-support systems in cardiovascular medicine.

References

[1] L.-M. Desroche, A. Trimaille, T. Garban, and C. Bouleti, "Artificial intelligence in cardiovascular medicine: An exoskeleton for perception, reasoning and action," *Archives of Cardiovascular Diseases*, vol. 118, no. 11, pp. 613–620, Nov. 2025, doi: 10.1016/j.acvd.2025.07.012.

[2] Global Cardiovascular Risk Consortium *et al.*, "Global Effect of Cardiovascular Risk Factors on Lifetime Estimates," *N Engl J Med*, vol. 393, no. 2, pp. 125–138, Jul. 2025, doi: 10.1056/NEJMoa2415879.

[3] H. Bihri, L. A. Charaf, S. Azzouzi, and M. E. H. Charaf, "A Robust Stacking-Based Ensemble Model for Predicting Cardiovascular Diseases," *AI*, vol. 6, no. 7, p. 160, Jul. 2025, doi: 10.3390/ai6070160.

[4] A. Niazai, H. Jamil, M. Hameed, S. Sheikh, and M. R. Nisar, "Artificial intelligence in cardiovascular diagnostics: a systematic review and descriptive analysis of clinical applications and diagnostic performance," *BMC Cardiovasc Disord*, vol. 25, no. 1, p. 849, Nov. 2025, doi: 10.1186/s12872-025-05327-x.

[5] W. Zhao *et al.*, "An agentic system for rare disease diagnosis with traceable reasoning," *Nature*, pp. 1–10, Feb. 2026, doi: 10.1038/s41586-025-10097-9.

[6] Y. Kim *et al.*, "Towards a Science of Scaling Agent Systems," Dec. 17, 2025, *arXiv*: arXiv:2512.08296. doi: 10.48550/arXiv.2512.08296.

[7] A. Bandi, B. Kongari, R. Naguru, S. Pasnoor, and S. V. Vilipala, "The Rise of Agentic AI: A Review of Definitions, Frameworks,

Architectures, Applications, Evaluation Metrics, and Challenges," *Future Internet*, vol. 17, no. 9, p. 404, Sep. 2025, doi: 10.3390/fi17090404.

[8] M. A. Tawfeek, I. Alrashdi, M. Alruwaili, and H. Allahem, "Cardiovascular disease detection: A hybrid machine learning-AI framework for personalized diagnosis and risk assessment," *PLOS ONE*, vol. 20, no. 10, p. e0335421, Oct. 2025, doi: 10.1371/journal.pone.0335421.

[9] A. A. Melih, Y. Singh, K. L. Agarwal, P. Mukherjee, K. Pattnaik, and H. Bhatia, "Human Machine Social Hybrid Intelligence:A Collaborative Decision Making Framework for Large Model Agent Groups and Human Experts," Nov. 11, 2025, *arXiv*: arXiv:2510.24030. doi: 10.48550/arXiv.2510.24030.

[10] P. G. Giannopoulos, V. Malamas, and T. K. Dasaklis, "Decentralized pricing in supply chain management: a blockchain-enabled multi-agent Reinforcement Learning approach," in *2025 11th International Conference on Control, Decision and Information Technologies (CoDIT)*, Jul. 2025, pp. 1571–1576. doi: 10.1109/CoDIT66093.2025.11321740.

[11] T. K. Dasaklis, P. G. Giannopoulos, V. Malamas, G. Tantis, and C. Patsakis, "Multi-Agent Reinforcement Learning for EV energy management and trading using the Lightning Network," in *2025 11th International Conference on Control, Decision and Information Technologies (CoDIT)*, Jul. 2025, pp. 1–6. doi: 10.1109/CoDIT66093.2025.11321879.

[12] A. Mertzani, P. G. Giannopoulos, V. Malamas, P. Chountalas, and T. K. Dasaklis, "<p>A Neurosymbolic Multi-Agent System for Explainable and Traceable Decision-Making in Organizational Contexts</p>," Dec. 15, 2025, *Social Science Research Network, Rochester, NY*: 5946554. doi: 10.2139/ssrn.5946554.

[13] R. Khera *et al.*, "Transforming Cardiovascular Care With Artificial Intelligence: From Discovery to Practice: JACC State-of-the-Art Review," *J Am Coll Cardiol*, vol. 84, no. 1, pp. 97–114, Jul. 2024, doi: 10.1016/j.jacc.2024.05.003.

[14] G. Quer and E. J. Topol, "The potential for large language models to transform cardiovascular medicine," *Lancet Digit Health*, vol. 6, no. 10, pp. e767–e771, Oct. 2024, doi: 10.1016/S2589-7500(24)00151-1.

[15] M. E. Rivera Boadla *et al.*, "Multimodal Cardiac Imaging Revisited by Artificial Intelligence: An Innovative Way of Assessment or Just an Aid?," *Cureus*, vol. 16, no. 7, p. e64272, Jul. 2024, doi: 10.7759/cureus.64272.

[16] "Screening and diagnosis of cardiovascular disease using artificial intelligence-enabled cardiac magnetic resonance imaging | Nature Medicine." Accessed: Mar. 18, 2026. [Online]. Available: https://www.nature.com/articles/s41591-024-02971-2

[17] G. Azarfar *et al.*, "Responsible adoption of multimodal artificial intelligence in health care: promises and challenges," *Lancet Digit Health*, vol. 7, no. 12, p. 100917, Dec. 2025, doi: 10.1016/j.landig.2025.100917.

[18] M. Milosevic, Q. Jin, A. Singh, and S. Amal, "Applications of AI in multi-modal imaging for cardiovascular disease," *Front Radiol*, vol. 3, p. 1294068, 2023, doi: 10.3389/fradi.2023.1294068.

[19] D. Isern and A. Moreno, "A Systematic Literature Review of Agents Applied in Healthcare," *J Med Syst*, vol. 40, no. 2, p. 43, Feb. 2016, doi: 10.1007/s10916-015-0376-2.

[20] N. Mohammadzadeh, R. Safdari, and A. Rahimi, "Multi-agent system as a new approach to effective chronic heart failure management: key considerations," *Healthc Inform Res*, vol. 19, no. 3, pp. 162–166, Sep. 2013, doi: 10.4258/hir.2013.19.3.162.

[21] A. Gorenshtein, M. Omar, B. S. Glicksberg, G. N. Nadkarni, and E. Klang, "AI Agents in Clinical Medicine: A Systematic Review," *medRxiv*, p. 2025.08.22.25334232, Aug. 2025, doi: 10.1101/2025.08.22.25334232.

[22] E. Tzanis *et al.*, "Agentic systems in radiology: Principles, opportunities, privacy risks, regulation, and sustainability concerns," *Diagn Interv Imaging*, vol. 107, no. 1, pp. 7–16, Jan. 2026, doi: 10.1016/j.diii.2025.10.002.

[23] H. Bihri, S. Azzouzi, and M. E. H. Charaf, "A Multi-Modal Approach for Cardiovascular Disease Prediction by Combining Both Images and Non-Image Data," in *Progress in Intelligent Computing and Secure Communication Systems*, C. El Mokhi, H. Hachimi, N. Hmina, and A. Addaim, Eds., Cham: Springer Nature Switzerland, 2025, pp. 407–416. doi: 10.1007/978-3-031-99997-0_36.

[24] H. Bihri, S. Sraidi, H. Jamal, S. Azzouzi, and M. E. H. Charaf, "Real-Time Cardiovascular Health Monitoring Through a Multimodal Data Integration Framework," in *Proceedings of Tenth International Congress on Information and Communication Technology*, X.-S. Yang, S. Sherratt, N. Dey, and A. Joshi, Eds., Singapore: Springer Nature, 2025, pp. 255–267. doi: 10.1007/978-981-96-6435-1_20.

[25] A. A. Borkowski and A. Ben-Ari, "Multiagent AI Systems in Health Care: Envisioning Next-Generation Intelligence," *Fed Pract*, vol. 42, no. 5, pp. 188–194, May 2025, doi: 10.12788/fp.0589.

[26] N. H. Elmasry and M. K. Elshaarawy, "Hybrid metaheuristic optimized Catboost models for construction cost estimation of concrete solid slabs," *Sci Rep*, vol. 15, no. 1, p. 21612, Jul. 2025, doi: 10.1038/s41598-025-06380-4.