HSPI The Open Access Publisher

JCCM ISSN 2575-0143 JOURNAL OF CARDIOLOGY AND CARDIOVASCULAR MEDICINE

Comprehensive Review

# A Comprehensive Review on Automated Cardiac Report Generation Using Generative AI

**Sofia Munawar***

Institute of Data Science, University of Engineering and Technology, Lahore, Pakistan

**More Information**

***Corresponding author:** Sofia Munawar, Institute of Data Science, University of Engineering and Technology, Lahore, Pakistan,
Email: 2023msds13@student.uet.edu.pk

**Submitted:** June 20, 2026
**Accepted:** June 25, 2026
**Published:** June 26, 2026

**Citation:** Munawar S. A Comprehensive Review on Automated Cardiac Report Generation Using Generative AI. J Cardiol Cardiovasc Med. 2026; 11(3): 19-24. Available from: https://dx.doi.org/10.29328/journal.jccm.1001226



**Keywords:** Automated cardiac report generation; Generative AI; Large language models (LLM); Multimodal deep learning; Semantic consistency; Clinical decision support; Electrocardiogram (ECG); Echocardiography; AI in cardiology; Domain-specific language models

Check for updates

OPEN ACCESS

## Abstract

Cardiovascular diseases (CVDs) have been the major cause of death in all parts of the world, with a tremendous burden on healthcare systems. Accurate and timely clinical reporting is essential to diagnosis, treatment planning, and management of patients. Still, the manual generation of reports is time-consuming, prone to error, and not consistently reflective. The use of Generative Artificial Intelligence (AI) and especially transformer-based Large Language Models (LLMs) and multimodal deep learning models offers transformative potential to automate the process of cardiac report generation and improve it. This review focuses on the current state-of-the-art AI-based methods with a focus on semantic understanding, multimodal understanding, and clinical relevance. EchoGPT, MEIT, and ECG-GPT are developed as advanced models that prove it is possible to create human-like clinical narratives, but note the weaknesses of factual accuracy and practical implementation. We present a comparative study of traditional and modern machine learning techniques, comment on the technical, clinical, and ethical issues, and lay out the strategic priorities of further research to enhance reliability, interpretability, and implementation of automated cardiac reporting systems.

## Introduction

Cardiology is a discipline that consumes data in large amounts where numerous forms of information exist that include electronic health records (EHRs), electrocardiograms (ECGs), imaging like echocardiography, echocardiography, CT, MRI and more recently wearable sensor data. The creation of correct cardiology reports has always been a manual process and demands extensive knowledge of the domain, significant attention to the interpretation of multimodal inputs and strict quality control. It is time-consuming, prone to inconsistencies, and prone to human error, and this can cause delays in making vital clinical decisions.

The recent developments in Generative AI, particularly the transformer-based models, have transformed natural language processing in healthcare. These frameworks can learn both structured and unstructured medical information to generate consistent and context-sensitive clinical histories. Generative AI has the potential to change the focus from just assisting diagnostic tools to auto-documentation with the creation of semantically rich, clinically relevant reports. Although this progress has been made, there are still problems that are encountered in ensuring semantic accuracy, effective multimodal integration, and smooth alignment with clinical workflows. The review presents the synthesis of existing research, the existing gaps, and the directions of future development of automated cardiac report generation using generative AI.

### Importance of automated cardiac report generation

Cardiovascular diseases (CVDs) are the most common cause of death throughout the world; the precision and prompt reporting of this issue directly influence the diagnosis and patient outcomes.

Manual reporting is time-consuming, inaccurate, and inconsistent, thereby generating delays in the course of treatment decisions.

Generative AI creates a new paradigm of transformation: semantically rich report generation, fully automated, and clinically relevant (Figures 1,2).


**Figure 1:** Automated cardiac report generation process.

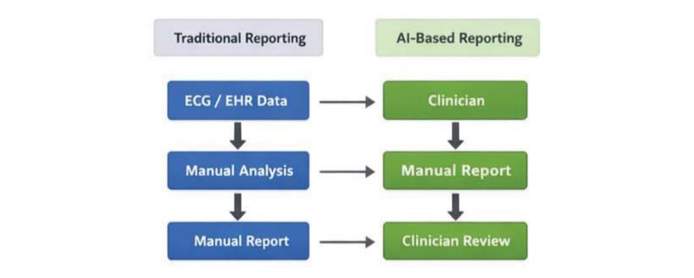


**Figure 2:** Workflow of traditional *vs.* AI-Based cardiac report generation.

# Literature review

Cardiovascular diseases (CVDs) are a major cause of death all over the world, and the emergence of patient-centered electronic health records (EHRs), wearables, and imaging technologies has created a need for intelligent, effective, and semantically correct clinical records. The latest developments in Generative Artificial Intelligence (AI) have demonstrated great potential in automating cardiology reports, providing the answer to the problem of efficiency, accuracy, and workload of clinicians.

## Artificial intelligence-cardiovascular diagnosis and risk prediction

Bhatt et al. [1] established the effectiveness of machine learning (ML) models such as the Random Forest, MLP, and XGBoost in predicting heart disease based on large-scale datasets of patients with AUCs of > 0.9. They increase the significance of feature engineering and clustering (K-Modes) to boost the model performance, but they pay more attention to the diagnosis, rather than to the semantic report generation.

Srinivasan and Sharma [2] provided a wider perspective of AI in CVD detection, including its use in echocardiography, CT, MRI, and ECG. Their review has indicated that the use of AI models could not only be used to detect cardiac abnormalities but also be used to predict risks (e.g., by automated coronary calcium scoring). Semantic interpretation and report generation are, however, underresearched.

The gap was strengthened by Kasartzian and Tsiampalis [3], who reviewed the transformative nature of AI in CVD risk prediction, demonstrating that multimodal input models are more accurate and personalized than traditional models. However, there is still the semantic layer of documentation that is not discussed properly.

## Generative AI in cardiology reporting

Chao, et al. [4] assessed echocardiography reporting with Large Language Models (LLMs). It was discovered that EchoGPT, a fine-tuned Llama-2, was useful in summarizing impressions of raw echoes. Although the research has supported that generative models can be used in generating cardiologist-aligned summaries, it also presented weaknesses in terms of numerical accuracy and clinical interpretability.

On the same note, Wan, et al. [5] proposed MEIT, a multimodal teaching-tuning framework that matches the ECG responses with clinical text to produce reports. It also demonstrated attention-based modal fusion, scoring much higher in report relevance and zero-shot generalization than traditional models, establishing a general pattern of full automation and semantically sensible ECG recording.

Khunte, et al. [6] suggested ECG-GPT, which is a vision encoder-decoder trained on ECGs. In contrast with the previous techniques that used raw signal data, ECG-GPT decoded ECGs in image formats and allowed the creation of diagnostic reports in resource-constrained environments. Its format-free and real-time deployment to the web points to its point-of-care application potential.

## Semantic consistency and structured reporting

Liu, et al. [7] dealt with semantic labeling on the diagnosis of coronary arteries based on the Multi-Graph Matching (MGM) framework. They merge an important aspect of proper AI-generated reports, which is semantic structure, although they deal with semantic annotation as opposed to narrative reports.

Similarly, BioASQ [8] presented multilingual semantic indexing and named entity recognition on cardiology texts, where the cross-lingual semantic extraction is considered significant. Their domain-specific generative model tasks are based on their clinical entity recognition activities and should be semantically and syntactically strong.

## Cardiology note generation with large language models

Jung, et al. [9] used the domain-specific LLM, Mistral-7B, to write discharge notes for cardiac patients. When compared to the record-writing by clinicians, their results state that LLMs may generate high-fidelity clinical notes. Nevertheless, the issues of clinical validation and workflow integration still exist.

https://doi.org/10.29328/journal.jccm.1001226

Chang Liu, et al. [10] suggested a contrastive learning-based bootstrapping technique with coarse-to-fine decoding of radiology reports. Though they worked in the area of radiology, the approach is mostly in line with the cardiology research requirements of domain-specific language adaptation and hierarchical report generation.

## Artificial intelligence policy and ethical application to cardiology reporting

Inam, et al. [11] examined the best journal guidelines on using generative AI in the writing of scientific papers in the cardiology field. As authors, their work has indicated the importance of disclosure, data integrity, and AI banning. These lessons are important in elucidating ethical scenarios of clinical documentation systems using generative AI.

## Closing gaps in generative AI in cardiology

There are several gaps in the literature. The majority of diagnostic systems prioritize accuracy rather than semantic completeness as their priority. Research papers by Rodriguez and Singh [12], Patel and Gupta [13], and Esteva, et al. [14] highlight that AI is not semantically fluent and has no contextual insights in report writing (Table 1).

There are also difficulties in incorporating multimodal information (structured EHR entries, imaging, and biosignals). Moreover, not many models have been made specifically to be trained for semantic consistency, narrative structure, or readability by clinicians in a cardiology setting [15,16].

## Developments in generative AI in cardiology

- Large Language Models (LLMs) that operate through the use of transformers and multimodal deep learning models allow the incorporation of ECG, imaging, as well as EHR data.
- Models such as EchoGPT, MEIT, and ECG-GPT have shown almost human-level quality clinical narratives.
- Multimodal architectures: Multimodal architectures improve the accuracy of semantics, zero-shot generalization, and the generation of reports in real time (Figure 3).

## Literature review conclusion

Although AI and ML have revolutionized cardiac diagnostics, the application of Semantic Cardiology Report

**Table 1:** Semantic cardiology reports literature review.

| Author(s) & Year | Title | Focus | Methodology | Key Findings |
|---|---|---|---|---|
| Bhatt et al. (2023) | Effective Heart Disease Prediction Using ML | Heart disease prediction using ML models and K-Modes clustering | ML models (DT, RF, MLP, XGBoost), feature engineering, 70k Kaggle dataset | MLP had highest AUC (0.95), 87.28% accuracy |
| Miller & Slomka (2024) | AI in Nuclear Cardiology | AI in myocardial perfusion imaging (MPI) | ML for image acquisition, segmentation, risk prediction | AI improves diagnostic accuracy, reduces radiation, streamlines MPI workflows |
| Srinivasan & Sharma (2025) | AI in Cardiovascular Disease Detection | Review of AI for CVD diagnosis | AI across imaging, ECG, wearable devices, predictive analytics | AI enhances diagnosis and risk prediction but faces clinical integration issues |
| Gibb et al. (2024) | Stent Therapy for Aortic Coarctation | Outcomes of stent use in small children | Retrospective study of 19 pediatric patients with CoA | 100% procedural success, 26% needed reintervention |
| Franke et al. (2024) | MYH7 Variant in Hypertrophic Cardiomyopathy | Genetic analysis of MYH7 in 5-generation family with HCM | Genetic testing + echocardiography | MYH7 mutation linked to HCM, supports family screening |
| Jung et al. (2024) | Discharge Note Generation Using LLM | LLMs for cardiac discharge note automation | Mistral-7B model tested on cardiology dataset | LLM-generated notes matched clinician quality, improved efficiency |
| Bernelli et al. (2024) | Radiation Exposure in Cardiology Labs | Health and gender effects of radiation in interventional cardiology | Survey of staff across institutions | Need for better PPE, radiation monitoring, policies for pregnant staff |
| Kasartzian & Tsiampalis (2025) | Transforming CVD Risk Prediction | AI/ML advancements in CVD risk prediction | Review of risk modeling, case study on calcium-omics | ML improves prediction but generalizability, trust remain concerns |
| Tiwari et al. (2025) | AI-based Risk Stratification in Autoimmune Women | AI for CVD/stroke risk in women with autoimmune disorders | ML models using imaging, clinical, and vitamin D data | AI improved prediction using imaging biomarkers |
| Chao et al. (2025) | LLMs in Echo Reporting | Evaluating open-source LLMs for echo summaries | EchoGPT, Mayo Clinic dataset | EchoGPT performed best, but factual accuracy concerns persist |
| Wan et al. (2024) | MEIT: ECG Instruction Tuning | Multimodal ECG report generation | Instruction tuning with attention-based fusion | High zero-shot accuracy, robust to signal perturbation |
| Inam et al. (2024) | AI Policy in Cardiology Journals | Journal guidelines on AI use in writing | Review of top 25 cardiology journal policies | AI allowed with disclosure, not as author |
| Khunte et al. (2024) | ECG-GPT for Image-based Reporting | AI-generated ECG reports from image format | Vision encoder-decoder on 2.6M ECGs | AUROC â‰¥ 0.81 across settings; good for low-resource clinics |
| Liu et al. (2024) | Zero-Shot ECG Classification (MERL) | Multimodal ECG classification with clinical knowledge prompts | Cross-modal alignment with clinical report grounding | Outperformed eSSL methods by 3.2% |
| Liu et al. (2024) | Radiology Report Generation via LLM | LLMs for chest X-ray report generation | Contrastive learning + Coarse-to-fine decoding | Outperformed SoTA on IU X-Ray & MIMIC-CXR |
| Liu et al. (2025) | ML for CVD Risk via EHRs | Systematic review/meta-analysis of ML for CVD risk | Meta-analysis of 32 ML vs 26 classical models | RF and DL had better AUC (0.865, 0.847) than traditional scores |
| Zhao et al. (2024) | Multi-Graph Matching for Coronary Artery Labeling | Semantic labeling of coronary arteries | Graph-based representation and matching of ICA | Labeling accuracy 94.7%, stenosis detection 91.5% |
| Nentidis et al. (2024) | BioASQ at CLEF2024 | Semantic QA & multilingual entity recognition in cardiology | BioASQ shared tasks: MultiCardioNER, QA. | Promotes NLP/QA tools for cardiology |

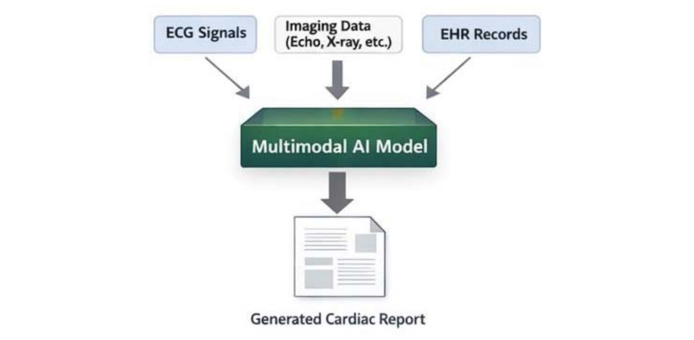


**Figure 3:** Multimodal generative AI framework for cardiac reporting.

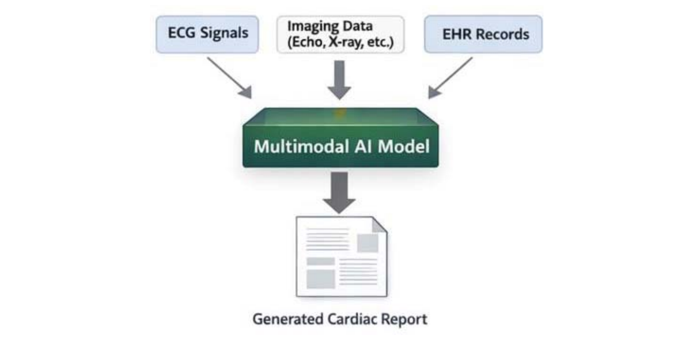


**Figure 4:** Comparative performance of traditional ML and generative AI.

Generation through Generative AI is a new field. LLMs such as EchoGPT, MEIT, and ECG-GPT have potential in creating structured and context-like reports. However, the issues of semantic consistency, clinical validation, data diversity, and ethical use need to be resolved. Future studies are advised to establish strong and multimodal generative models that can generate patient-specific, semantically enhanced, and scientifically accepted cardiology reports.

## Comparison of existing approaches

The classical machine learning (ML) models have been very much concerned with prediction accuracy based on structured clinical data, where the results are in the form of risk scores, disease classification, or probability estimates. Although models such as Random Forests, MLPs, and XGBoost are strong at diagnosing and assessing risks, they do not have the capability of producing natural language reports, which restricts clinical interpretability and usability.

Generative AI systems, p.s., are based on combining various data sets with ECG, imaging, and EHR records, and generate textual accounts with context-related predictive power. Vision-encoder-decoder systems, large language models, and multimodal instruction-tuning systems provide a few benefits over the conventional ones: improved semantic knowledge, comprehensible text, and compatibility with clinical procedures. It is interesting to note that these models exhibit better performance defined by the production of structured and coherent reports coupled with clinical relevance, and the gap between AI-driven diagnostics and actionable documentation.

## Comparison between traditional ML and generative AI

- **The conventional ML models**: Number prediction (score of risk, classification), weak semantic interpretability, and are not capable of generating human-readable reports (Figure 4).
- **Generative AI models:** Multimodal data: process data, generate semantically consistent and contextually aware text, and are more conducive to clinical adoption of workflow.
- Key advantages are that it will have improved readability of reports, semantic coherence, and clinically friendly outputs.

**The major findings of the comparison consist of:**

1. Generative AI models provide semantically sensible, contextualized texts.
2. Compared to single-source models, multimodal architectures are more accurate and relevant to clinical concerns.
3. Large Language Models improve report readability, which enables clinicians to adopt them for routine documentation.

# Discussion

The use of generative AI in cardiology is a paradigm shift, as it is a shift to a more diagnostic aid and not full-scale automated documentation. Systems like EchoGPT, MEIT, and ECG-GPT demonstrate that it is possible to generate clinical narratives that are similar to human-written ones. The key to the success of these models is quality access to domain-specific datasets, advanced training methods, and adaptation to the specifics of the clinical setting. Multimodal architectures that combine structured EHR data and imaging and physiological data provide a better quality of reports, but introduce technical challenges such as data fusion, model optimization, and deployment issues. In addition to technical performance, it is important to make sure that factual, semantic, and clinician trust is in place. To be adopted successfully, close cooperation between AI researchers, clinicians, and healthcare IT professionals is required to develop systems that are dependable and that can be seamlessly incorporated into the current workflows.

## Semantic accuracy and structured reporting

- Semantic consistency is of the utmost importance in precise cardiology reports (Figure 5).
- Multi-Graph Matching (MGM) and clinical entity recognition (BioASQ) offer the structures to be used to guarantee semantic labeling, cross-lingual indexing, and generated structured reports.

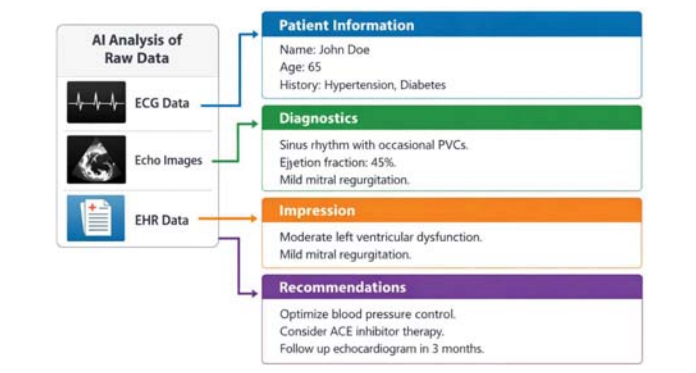


**Figure 5:** Semantic structure in AI-generated cardiac reports.


**Figure 6:** Challenges in automated cardiac report generation.

- System-specific LLMs in the field of cardiology enhance hierarchical report writing and context-awareness.

## Challenges and limitations

Generative AI-based automated generation of cardiac reports encounters several interconnected issues. Data-wise, there are still very limited high-quality labeled datasets, which are typically fragmented and subject to stringent privacy laws, which constrain the performance of models to generalize naturally and work well under different clinical settings. At the technical level, semantic consistency across reports, the ability to support multimodal data, including ECGs, imaging studies, and electronic health records, and avoiding hallucinations in outputs of Large Language Models are current challenges that need complex solutions. The clinical implementation of AI-generated reports into the current hospital workflow, demonstrating the accuracy of model performance in practice, and building trust among clinicians are critical obstacles to implementation. In addition to technical and operational considerations, ethical and legal concerns such as bias during training datasets, transparency during model decision-making, and compliance with regulatory standards are important factors that can be used in ensuring safe and responsible deployment of AI in healthcare.

## Technical and clinical problems

- **Data issues:** Data scarcity, heterogeneity, and privacy.
- **Technical issues:** Multimodal data fusion, semantic accuracy, hallucinations in the output of LLMs.
- **Workflow validation and clinician trust:** Clinical Challenges (Figure 6).
- **Ethical dilemmas:** Discrimination, legal adherence.

## Future research directions

To overcome these difficulties and improve the field, future research in the future has to be concentrated on several main areas. The creation of multimodal generative models that will successfully combine ECG data, imaging research, and EHR data will increase the semantic accuracy and clinical utility of AI-produced reports. Large-scale models: Domain-specific LM models that are trained or directed at cardiology should be better able to enhance report accuracy and context, whereas explainable AI (XAI) solutions can enhance understanding, leading to increased clinician trust and adoption. The attempts to introduce real-time clinical integration will allow generating patient-related reports directly in hospital processes, which will guarantee their practical use and instant clinical benefit. Lastly, standardized and publicly available datasets and evaluation criteria are needed in order to encourage reproducibility, transparency, and cross-generative AI model comparison to establish the basis of reliable and real-world application of automated cardiac reporting systems.

## Future work

- **Multimodal generative frameworks:** ECGs, imaging, and EHRs are combined to increase the accuracy of semantics.
- **Domain-specific LLMs:** Higher-fidelity cardiology language models.
- **Explainable AI (XAI):** Support assistibility and physician confidence.
- **Goal:** Integrate patient-specific AI-generated reports in hospitals in real-time.
- **Standardized datasets and benchmarks:** This will allow reproducibility, fair comparisons, and strong evaluation (Figure 7).


**Figure 7:** Roadmap for future research in AI-Based cardiac reporting.

## Conclusion

The use of generative AI to produce cardiac reports automatically is a breakthrough in healthcare documentation. Although traditional ML models enhance diagnostic accuracy, they do not generate reports that can be read by human beings and have a semantic nature. To overcome this drawback, generative AI systems, especially LLLMs and multimodal systems, can be used to generate context-oriented, clinically applicable narratives. Issues like lack of data, semantic issues, ethics, and incorporating into clinical processes still exist. Further investigations with a focus on multimodal systems, domain-specific language models, explainable AI, and publicly accessible datasets will play a crucial role in ensuring the reliability of real-world implementation, which, in the long run, will lead to the efficiency of clinical work, patient outcomes, and the quality of cardiology reports.

## References

1. Bhatt CM, Patel P, Ghetia T, Mazzeo PL. Effective heart disease prediction using machine learning techniques. Algorithms. 2023;16(2):88. Available from: https://doi.org/10.3390/a16020088?urlappend=%3Futm_source%3Dresearchgate.net%26utm_medium%3Darticle
2. Miller RJ, Slomka PJ. Artificial intelligence in nuclear cardiology: an update and future trends. Semin Nucl Med. 2024 Sep;54(5):648-57. Available from: https://doi.org/10.1053/j.semnuclmed.2024.02.005
3. Srinivasan SM, Sharma V. Applications of AI in cardiovascular disease detection: a review of the specific ways in which AI is being used to detect and diagnose cardiovascular diseases. In: AI Disease Detection: Advanced Applications. 2025. p. 123-46. Available from: https://doi.org/10.1002/9781394278695.ch6
4. Gibb JJ, Kim WC, Barlatay FG, Tometzki A, Pateman A, Caputo M, et al. Medium-term outcomes of stent therapy for aortic coarctation in children under 30 kg with new generation low-profile stents: a follow-up study of a single centre experience. Pediatr Cardiol. 2024;45(3):544-51. Available from: https://doi.org/10.1007/s00246-023-03402-8
5. Franke M, Książczyk TM, Dux M, Chmielewski P, Truszkowska G, Czapczak D, et al. A MYH7 variant in a five-generation family with hypertrophic cardiomyopathy. Front Genet. 2024;15:1306333. Available from: https://doi.org/10.3389/fgene.2024.1306333
6. Bernelli C, Di Fusco SA, Matteucci A, Zilio F, Nesti M, Barbero U, et al. Working in interventional cardiology laboratories: the perceived impact of radiation exposure as a health and gender hazard. A NEXT generation ANMCO initiative. Int J Cardiol. 2024;401:131682. Available from: https://doi.org/10.1016/j.ijcard.2023.131682
7. Kasartzian DI, Tsiampalis T. Transforming cardiovascular risk prediction: a review of machine learning and artificial intelligence innovations. Life (Basel). 2025;15(1):94. Available from: https://doi.org/10.3390/life15010094
8. Mathur P, Srivastava S, Xu X, Mehta JL. Artificial intelligence, machine learning, and cardiovascular disease. Clin Med Insights Cardiol. 2020;14:1179546820927404. Available from: https://doi.org/10.1177/1179546820927404
9. Chowdhury MA, Rizk R, Chiu C, Zhang JJ, Scholl JL, Bosch TJ, et al. The heart of transformation: exploring artificial intelligence in cardiovascular disease. Biomedicines. 2025;13(2):427. Available from: https://doi.org/10.3390/biomedicines13020427
10. Ali MT, Ali U, Ali S, Tanveer H. Transforming cardiac care: AI and machine learning innovations. Int J Multidiscip Res Growth Eval. 2024. Available from: https://www.allmultidisciplinaryjournal.com/article/2645/transforming-cardiac-care-ai-and-machine-learning-innovations
11. Tiwari E, Shrimankar D, Maindarkar M, Bhagawati M, Kaur J, Singh IM, et al. Artificial intelligence-based cardiovascular/stroke risk stratification in women affected by autoimmune disorders: a narrative survey. Rheumatol Int. 2025;45(1):14. Available from: https://doi.org/10.1007/s00296-024-05756-5
12. Chao CJ, Banerjee I, Arsanjani R, Ayoub C, Tseng A, Delbrouck JB, et al. Evaluating large language models in echocardiography reporting: opportunities and challenges. Eur Heart J Digit Health. 2025;6(3):326-39. Available from: https://doi.org/10.1093/ehjdh/ztae086
13. Wan Z, Liu C, Wang X, Tao C, Shen H, Peng Z, et al. MEIT: multi-modal electrocardiogram instruction tuning on large language models for report generation [Preprint]. arXiv. 2024. Available from: https://doi.org/10.48550/arXiv.2403.04945
14. Inam M, Sheikh S, Minhas AMK, Vaughan EM, Krittanawong C, Samad Z, et al. A review of top cardiology and cardiovascular medicine journal guidelines regarding the use of generative artificial intelligence tools in scientific writing. Curr Probl Cardiol. 2024;49(3):102387. Available from: https://doi.org/10.1016/j.cpcardiol.2024.102387
15. Khunte A, Sangha V, Oikonomou EK, Dhingra LS, Aminorroaya A, Coppi A, et al. Automated diagnostic reports from images of electrocardiograms at the point-of-care [Preprint]. medRxiv. 2024. Available from: https://doi.org/10.1101/2024.02.17.24302976
16. Liu C, Wan Z, Ouyang C, Shah A, Bai W, Arcucci R. Zero-shot ECG classification with multimodal learning and test-time clinical knowledge enhancement [Preprint]. arXiv. 2024. Available from: https://doi.org/10.48550/arXiv.2403.06659
17. Liu C, Tian Y, Chen W, Song Y, Zhang Y. Bootstrapping large language models for radiology report generation. Proc AAAI Conf Artif Intell. 2024;38(17):18635-43.
18. Liu T, Krentz A, Lu L, Curcin V. Machine learning-based prediction models for cardiovascular disease risk using electronic health records data: systematic review and meta-analysis. Eur Heart J Digit Health. 2025;6(1):7-22. Available from: https://doi.org/10.1093/ehjdh/ztae080
19. Zhao C, Xu Z, Baral P, Esposito M, Zhou W. Multi-graph graph matching for coronary artery semantic labeling [Preprint]. arXiv. 2024. Available from: https://doi.org/10.48550/arXiv.2402.15894
20. Nentidis A, Krithara A, Paliouras G, Krallinger M, Sanchez LG, Lima S, Tutubalina E. BioASQ at CLEF2024: the twelfth edition of the large-scale biomedical semantic indexing and question answering challenge. In: European Conference on Information Retrieval. Cham: Springer Nature Switzerland; 2024. p. 490-97.