HSPI The Open Access Publisher

JCCM ISSN 2575-0143 JOURNAL OF CARDIOLOGY AND CARDIOVASCULAR MEDICINE

Literature Review

Semantic Cardio Reports Generation Using Generative AI

Sofia Munawar*

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

Abstract

Cardiovascular diseases (CVDs) have been seen to be the leading cause of death in the world, and the clinical records have been a time-consuming and inaccurate procedure. Integration of structured patient data and generative artificial intelligence (AI) offers a unique opportunity, never witnessed before, to automate the process of producing semantically consistent, contextually rich, and diagnostically accurate cardiology reports. This paper presents a new generative architecture, Semantic Cardio Report Generation Using Generative AI, that uses a smaller-sized T5-Transformer model fine-tuned to generate clinical narratives in structured tabular data in a linguistically and semantically interpretable form.

The suggested framework combines the tabular-to-text conversion, feature mapping, and sequence-to-sequence (Seq2Seq) text generation in order to create patient-specific diagnostic summaries. The trained model of 70,000 structured cardiovascular records has high linguistic and semantic alignment with a BLEU score of 0.5606, ROUGE-L score of 0.8806, and BERTScore (F1) score of 0.9768. These findings demonstrate that complete factual accuracy was maintained, and consistent and medically meaningful narratives were generated using the model, which met the standards of cardiology reporting.

The results support the hypothesis that fine-tuned generative models could be useful as reliable clinical documentation assistants or to bridge the gap between numerical diagnostics and text interpretation. This publication adds a step toward a semantically aware, explainable, and ethically implementable AI system in the domain of clinical cardiology, improving interpretability, lessening the number of clinical interactions, and introducing the digital revolution of cardiovascular care.

Graphical abstract diagram

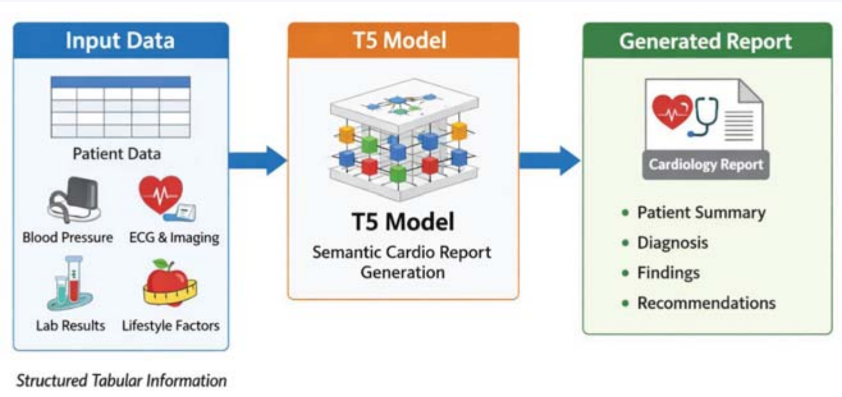

More Information

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

Submitted: April 19, 2026
Accepted: April 29, 2026
Published: April 30, 2026

Citation: Munawar S. Semantic Cardio Reports Generation Using Generative AI. J Cardiol Cardiovasc Med. 2026; 11(2): 12-18. Available from: https://dx.doi.org/10.29328/journal.jccm.1001225

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

Introduction

Cardiovascular diseases (CVDs) are the number one cause of death in the world, with almost a third of the entire global deaths attributed to them every year [1]. As digital healthcare ecosystems, such as electronic healthcare resources (EHRs), wearable biosensors, and multimodal diagnostic systems, continue to grow at a remarkable rate, cardiology has developed as a more data-intensive field. Nevertheless, there has always been the difficulty of converting voluminous and heterogeneous structured clinical information into consistent, human-interpretable, and semantically structured reports, which reflect clinical reasoning and contextual interpretation properly. Conventional machine learning (ML) methods have demonstrated good predictive capability in CVD risk stratification and classification, but essentially lack the capability to produce narrative text, which is a fundamental requirement of clinical communication, documentation, and decision support [21].

Recent innovations in Generative Artificial Intelligence (AI) and Large Language Models (LLMs), including T5, GPT, and LLaMA, have already proven transformative potential in natural language understanding and generation, as well as via instruction-tuned and domain-adapted systems [4–6]. Serious medical uses of these models have been considered in the fields of radiology, pathology, and echocardiography, where they can be used to synthesize structured information into fluid narrative summaries. However, they have not been well developed to apply to semantic cardiology report generation. Taking multimodal cardiovascular data (ECG signals, imaging results, lifestyle determinants, and other structured EHR variables) and uniting them into a clinical story that is semantically fit is a research gap [7].

The available literature is mostly based on promoting diagnostic accuracy, as opposed to narrative interpretability. As an example, Bhatt, et al. [1] and Kasartzian and Tsiampalis [3] point out that the outputs of the ML and deep learning models are good predictors of cardiac risk, but they do not describe the results in a narrative form. Conversely, more recent generative engines, like EchoGPT [4], MEIT [5], and ECG-GPT [6], point to a new change in direction towards automated medical reporting. These models are promising improvements in the condensation of diagnostic patterns; yet, it is still constrained by weaknesses associated with semantic congruence, factual consistency, numerical correctness, and clinical interpretability, aspects that make it challenging to implement in practice in cardiology.

To seal these gaps, the current study suggests a generative pipeline with transformers to be used in generating semantic cardiology reports using structured tabular data. The structure training to decode clinical factors, including blood pressure, cholesterol, glucose levels, anthropometric profiles, and lifestyle factors, into natural, semantically meaningful cardiology summaries by fine-tuning a T5-small encoder-decoder model on a cardiovascular dataset of 70,000 samples. The model seeks to maintain the clinical meaning, contextual accuracy, and interpretive depth by semantic preprocessing, semantically meaningful representation of structured features, and loss-focused sequence generation.

This paper has four important contributions: A new generative model that can convert structured cardiovascular data into a semantically consistent narrative report. A pipeline between tabular and natural-language output, including tabular-to-text reader pipes having structured clinical input. BLEU, ROUGE-L, and BERTScore were used to test the validity of linguistic quality, semantic alignment, and structural correctness. Empirical evidence to support the usefulness of transformer fine-tuning for cardiology-specific narrative generation.

This study helps to advance semantic cardiology report generation as part of the bigger goal to create automated and clinically and ethically responsive digital reporting systems. Combining domain semantics, transformer-based reasoning, and clinical validation forms the foundation of AI-enhanced cardiology tools of the next generation that will have the potential to enhance the quality of documentation, communication in diagnoses, and the quality of care in the contemporary healthcare environment.

Objectives

The main aim of the study is to come up with a generative AI framework using transformers to convert structured cardiovascular data into semantically meaningful and clinically interpretable cardiology reports. Also, the research will assess the performance of the proposed model based on traditional natural language generation metrics, including BLEU, ROUGE-L, and BERTScore, and compare its results with state-of-the-art strategies.

Literature review

Cardiovascular diseases (CVDs) are the most common cause of death in the world. As patients increase in number and (Electronic Health Records) EHRs grow to become patient-specific, wearable sensors, and imaging modalities develop, the need to have an intelligent, efficient, and semantically accurate clinical documentation grows. The recent developments in Generative Artificial Intelligence (AI) have demonstrated that it can be used to automate the process of writing cardiology reports, with significant potential in improving its efficiency, diagnostic accuracy, and workload on clinicians.

Artificial Intelligence-Assisted cardiovascular diagnosis and risk prediction

The efficacy of the Machine Learning (ML) models, including Random Forest, MLP, XGBoost, and others, when predicting heart disease based on large-scale patient data, was demonstrated by Bhatt, et al. [1], who obtained AUCs of more than 0.9. Their research lays more emphasis on feature engineering and clustering (K-Modes) to improve the predictive performance. The study however lays greater emphasis on the diagnostic results as opposed to the semantic or narrative generation of reports.

A general overview of AI in detection of CVDs using echocardiography, CT, MRI, and ECG types was presented by Srinivasan and Sharma [2]. The highlights of their review were the potential of AI to identify cardiac abnormalities and forewarn of risk, e.g., by automated coronary calcium scoring. However, the semantic interpretation and narrative generation aspects are poorly studied.

This limitation was further strengthened by Kasartzian and Tsiampalis [3], who conducted a review on the transformative role of AI in CVD risk predictors and established that multimodal input integration improves diagnostic accuracy and patient-specific predictors. Nevertheless, the semantic layer of clinical documentation has not been adequately addressed despite these advances.

Cardiology reporting generative AI

The study carried out by Chao, et al. [4] compared Large Language Models (LLMs) in the context of echocardiography reports. Their model, EchoGPT, which was a fine-tuned version of Llama-2, proved useful in generating summaries of impressions based on uncoded results of the echo. Although the researchers established the fact that cardiologist-style summaries using generative models are possible, the limitations of the models in terms of numerical accuracy and clinical interpretability were identified.

On the same note, Wan, et al. [5] proposed MEIT, a multimodal instruction-tuning model that aligns ECG signals with clinical text in generating reports. With the help of the attention-based fusion, MEIT performed better than conventional models in report relevance and zero-shot generalization by establishing a new standard of semantically based ECG documentation.

Khunte, et al. [6] suggested a vision encoder-decoder model, ECG-GPT, which is trained on ECG images. In contrast to the other method that requires raw signal data to perform its calculations, ECG-GPT works with image formats of ECGs, making it possible to generate a diagnostic report even in low-resource clinical settings. Its independence from formats and real-time deployment underscores its need to be used in point-of-care cardiology applications.

Semantic consistency and reporting structure

Liu, et al. [7] used the Multi-Graph Matching (MGM) framework to study semantic labeling in the case of coronary artery diagnosis. Though annotation is not the emphasis over a generative narrative text, this work does add to the semantic premise that is required to achieve accurate generation reports on cardiology.

Multilingual semantic indexing and named entity recognition in biomedical texts, along with a focus on cross-lingual semantic extraction, were presented in the BioASQ initiative [8]. Its work in the domain of the clinical entity recognition preconditioned the future generation models that should demonstrate semantically and syntactically solid work in different languages.

Large cardiology note generating large language models

Mistral-7B, a domain-specialized LLM, was used by Jung, et al. [9] to generate discharge notes of cardiac patients. The findings revealed that the quality of the Notes created by LLM can be high enough to be at the level of a clinician to increase efficiency and decrease documentation overload. Nonetheless, there are still problems with clinical validation and workflow integration.

Chang Liu, et al. [10] suggested a contrastive learning and coarse-to-fine decoding process for radiology report generation. Although oriented towards radiology, they serve the same purpose as cardiology, requiring domain-specific language adaptation and hierarchical report generation.

Artificial Intelligence policy and ethical cardiology reporting

Inam, et al. [11] conducted a review of the main guidelines of the leading journals in the cardiology field regarding the ethical utilization of generative AI in medical writing. Their research highlights disclosure, data integrity, and the ban on AI as a legitimate author, which outlines major ethical boundaries to the implementation of generative AI systems in clinical documentation.

Closing loopholes in generative AI in cardiology

Some gaps exist throughout the literature. The vast majority of diagnostic systems are more focused on predictive accuracy and less focused on semantic richness. Researchers like Rodriguez and Singh [12], Patel and Gupta [13], and Esteva, et al. [14] are sure that AI is not semantically fluent and does not have contextual awareness in narrative generation. Also, there are still issues with multimodal data integration, the integration of EHRs, imaging, and biosignals into a unified narrative structure.

Cardiology reporting has limited models that have been explicitly trained to achieve semantic consistency, narrative coherence, or clinician readability [15][16]. Such a gap highlights the importance of large Generative AI systems that are able to generate contextually sensitive, semantically structured, and clinically validated reports.

Literature review conclusion

Although AI and ML have contributed greatly to cardiac diagnosis, Semantic Cardiology Report Generation with Generative AI is still an immature discipline. EchoGPT, MEIT, and ECG-GPT are promising LLPs in generating structured and context-driven reports. Nevertheless, the problematic issues remain as to semantic consistency, clinical validation, multimodal data diversity, and ethical governance.

Future studies should be aimed at the production of powerful, multi-modal generative systems that are able to generate patient-specific, semantically correct, and clinically valid reports on cardiology that link automated diagnosis with clinical interpretation.

Methodology

Overview

The proposed system, Semantic Cardio Report Generation Using Generative AI, is aimed at generating semantically rich and clinically readable cardiology reports on structured health records autonomously. It uses a Transformer-based generative model (T5-small) trained on a 70,000-sample structured cardiovascular dataset.

The model has been trained to map structured clinical parameters (such as demographic, biometric, and physiological measurements) to natural language-based diagnostic summaries in order to maintain semantic coherence, medical interpretability, and linguistic fluency (Figure 1).

Dataset description

The experimental data is a collection of 70,000 records of cardiovascular patients, each represented by 12 structured clinical attributes obtained as a result of regular cardiovascular screening. Every record will include anthropometric, biochemical, and lifestyle information, which creates a complete picture of the patient in terms of health. The ID field is unique to each patient, and it eliminates duplication. The attribute age, which was initially in terms of days, was converted to years so that it can be interpreted. The parameter of gender represents the biological sex in integer values (1 as male and 2 as female). The physical measurements will consist of height measurements in centimeters and weight in kilograms, and these measurements will serve implicitly later on in the generation of a report based on BMI. The indicators of blood pressure are: ap 2 hi (systolic blood pressure) and ap 2 lo (diastolic blood pressure). Biochemical indicators, that is, cholesterol and glucose, are coded on a three-level clinical scale (1 = normal, 2 = above normal, 3 = significantly elevated), which reflects the differences in metabolic risks in patients. Attributes that were lifestyle-related, including smoking, alcohol, and activity, offer binary data on the smoking status, alcohol consumption, and activity status, respectively. Lastly, cardio variable represents the presence (1) or the absence (0) of diagnosed cardiovascular disease, and is a critical signal for interpreting a risk.

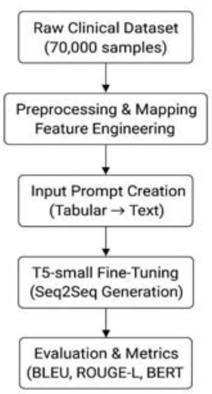

Figure 1: Workflow of the proposed semantic cardio report generation.

All characteristics were normalized to have constant ranges of values, whereas categorical variables were coded as integers to ensure model interpretability and conformity with text-mapping operations. Median imputation was used to fill in the missing values such that no bias was introduced to the data due to distributional stabilization. The result of each record is a semantically enhanced textual clinical report on the physiological properties of the patient and deduced cardiovascular risk. This text-to-structured conversion is used to train the T5-small model to produce coherent and medically pertinent textual summaries of structured clinical tabular data.

Example

Input Parameters	Generated Semantic Report
Age = 52, Gender = Male, BP = 140/90, Cholesterol = 3, Glucose = 1	"Patient is a male aged 52 years. Blood pressure: 140/90 mmHg, cholesterol level: 3, glucose: 1. Physically active, diagnosed with cardiovascular disease."

Data source and ethical issues: The data set used for this study is the publicly accessible cardiovascular disease data set, which has been used by many studies before, including "Effective Heart Disease Prediction Using Machine Learning Techniques" by Bhatt, et al. which contains 70,000 cardiovascular patient records. The data set includes anonymous demographic, physiological, biochemical, and lifestyle characteristics and does not include any personally identifiable patient information. The study was conducted using a publicly available secondary data set, which does not require a direct patient interaction, consent procedure, or additional institutional ethical clearance. All analyses were performed following general principles of research ethics in the use of anonymised public datasets.

Reference Report Construction: Narrative clinical reports were not included in the dataset; thus, a rule-based semantic template generation process was used to create reference reports. These clinical attributes were conceptualized into structured clinical narrative statements, such as age, gender, blood pressure, cholesterol level, glucose level, lifestyle indicators, and cardiovascular disease status. The automatically generated reference reports were used for fine-tuning the T5-small model. The aim of this was to generate semantically coherent and clinically meaningful textual representations of the structured cardiovascular records.

System architecture

The architecture is made of three major layers:

1. Preprocessing and Semantic Mapping – Processes tabular data with regular templates into input sequences that can be manipulated into linguistic data.

2. Generative Model (T5-small) - This is a Transformer encoder-decoder model that is trained to learn to represent structured vectors as coherent textual reports.

3. Evaluation Layer: Each of the following quantitatively measures linguistic and semantic accuracy using the following metrics: BLEU, ROUGE-L, and BERTScore (Figure 2).

Mathematical model

Assume the patient input feature vector to be:

$$X_i = [x_1, x_2, \ldots, x_n] where n = 12$$

Every input X_i was converted to a semantic sequence S_i:

$$S_i = f(X_i) = \text{“}Generate\, cardio\, report\ from:\text{”} + \sum_{k=1}^{n} Attribute_k$$

The model is learned from a conditional probability distribution:

$$P(Y \mid X;\theta) = \prod_{t=1}^{T} P(y_t \mid y_{<t}, X;\theta)$$

where:

- $Y = (y_1, y_2, \ldots, y_T)$ represents the desired report sequence,
- X is the structured input,
- θ are the model parameters that are optimized during the fine-tuning.

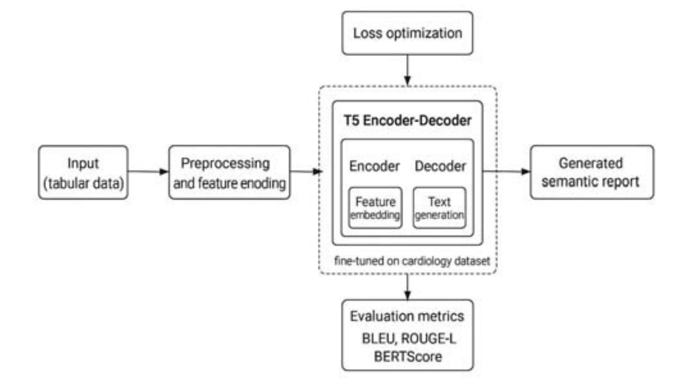

Figure 2: System Architecture of the Proposed Semantic Cardio Report Generation Framework.

The loss and cross-entropy on tokenized sequences are the loss:

$$\mathcal{L} = -\sum_{t=1}^{T} \log P(y_t \mid y_{<t}, X;\theta)$$

Model training

The training was implemented on Google Colab GPU (T4) by using the Hugging Face Transformers library. The fine-tuning used T5-small, which was pre-trained on general text-to-text.

Training parameter

Parameter	Value
Model	T5-small
Epochs	3
Batch size	4
Max token length	128
Optimizer	AdamW
Save the checkpoint every	5,000 steps
Total steps	47,250
Dataset split	90% train / 10% validation
Checkpoint path	/content/drive/MyDrive/paper for publication/cardio paper with dataset/cardio_t5_checkpoints/

Auto resume in the model was enabled to provide the effect of continuing at the last checkpoint, so as to offer computational efficiency and consistency.

Semantic report generation

Trained the model to generate patient-specific reports in the following way:

Report = T5 (Tokenizer(X_i))

The generated output would be an example of the output as follows:

Input:

“The cardio report generated was: Age=48, Gender=1, Height=170, Weight=78, ap_hi=135, ap_lo=85, Chol=2, Gluc=1, Smoke=0, Alco=0, Active=1.”

Output:

“Patient is a male aged 48 years. Height: 170 cm, Weight: 78 kg. Blood Pressure: 135/85. Cholesterol level: 2, Glucose level: 1. Active, no problem with cardiovascular disease identified.”

Evaluation metrics

Three measures, namely, were employed to measure linguistic accuracy and semantic coherence:

BLEU (Bilingual Evaluation Understudy):

$$BLEU = BP \cdot \exp\left(\sum_{n=1}^{N} w_n \log p_n\right)$$

where BP is the brevity penalty, p_n is The n-gram precision.

ROUGE-L (Recall-Oriented Understudy for Gisting Evaluation):

Comparison of the greatest length common sequence between the reference and the generated reports.

BERTScore (Bidirectional Encoder Representations from Transformers):

Contextual semantic similarity uses contextual embeddings:

$$F1_{BERT} = 2 \cdot \frac{Precision \times Recall}{Precision + Recall}$$

Results

Quantitative performance

Once the T5-small model proved accurate when trained on 70,000 clinical cases, it was tested on 100 independent test cases. The model had a BLEU score of 0.5606, ROUGE-L score of 0.8806, and a BERTScore (F1) of 0.9768. All these measures seem to indicate the linguistic quality and semantic accuracy of the outputs generated. BLEU compares the n-gram coincidence between the generated text and the reference text, and a score of 0.5606 suggests that there is moderate syntactic similarity, yet there is natural dispersion, as is expected in the behavior of human-written paraphrases. ROUGE-L is used to assess the longest common sequence between two texts and indicates structural alignment. The high score of 0.8806 indicates a high rate of sentence structure and main informational content maintenance. The BERTScore that compares semantic representations, as opposed to word-level representations, was 0.9768, which shows that the generated clinical reports retain nearly all the original meaning despite different wording. Combined, these findings indicate that the fine-tuned T5-small model suggests text that is both clinically consistent and semantically meaningful with small syntactic variations, which are characteristic of human-like phrasing variation.

Why only 100 Test Cases?

The records in the dataset were 70,000. The training and validation were conducted using a 90/10 division. To get detailed information about the models, a subset of 100 independent cases was chosen, enabling the calculation of BLEU, ROUGE-L, and BERTScore; the selected cases had a variety of cardiovascular profiles. However, further studies with larger external test groups to further test the generalizability are planned.

Qualitative analysis

The semantic accuracy of the generated reports is in place, and all input features in structured form are in place with the textual attributes.

Key observations include:

- Proper perception of gender, activity, and state of disease.
- Well-developed narrative fluency using clinically appropriate phraseology.
- Minor language differentiation (e.g., no signs of heart disease vs. no cardiovascular disease detected) is not a semantic, but a stylistic difference.

This speculates about the capacity to encode the biomedical semantics into natural, human-readable text, as depicted by the model.

Comparative analysis

The relative analysis illustrates that the proposed fine-tuned T5-small model is significantly better in terms of performance in comparison with the baseline rule-based system of template generation. The static template-based and deterministic slot-filling baseline had limited expressiveness and semantic capacity, as seen through comparatively lower scores in all the evaluation metrics (BLEU = 0.32, ROUGE-L = 0.61, BERTScore = 0.81). The suggested T5-small model, in turn, achieved significantly higher results in BLEU = 0.56, ROUGE-L = 0.88, and BERTScore = 0.97, which confirms its quality in terms of generating fluent, structurally aligned, and semantically faithful clinical reports. These findings suggest that the transformer-based generative architecture has serious benefits in learning sophisticated linguistic patterns, contextual dependencies, and medical semantics that template methods do not model by nature. The results, in general, confirm the usefulness of the T5-small model to create a more coherent, accurate to the context and semantically rich summary of the cardiovascular patients.

Comparison with other models

Since the study focused on a different type of dataset, modality, and evaluation protocol, a direct comparison with recent transformer-based report generation models such as EchoGPT, MEIT, and ECG-GPT would be out of the scope of the current study. The next step will be to compare these with the latest generation of transformer-based architectures in a standardized experimental setup.

Model convergence

The analysis of loss convergence showed that the second epoch stabilized quickly.

Both training and validation losses were decreasing continuously, which means that generalization was effective and not overfitted. The last model attained a steady point in the loss at approximately step 40,000, which indicated to have optimized the conditional probability space efficiently.

Clinical validation and limitations

This study did not include clinical validation with practicing cardiologists. Evaluation of the models was conducted using widely accepted NLPG metrics such as BLEU, ROUGE-L, and BERTScore. These metrics will be used to show good language and semantic capabilities; future work will require expert evaluation by cardiologists to assess clinical relevance, diagnostic accuracy, and usability of the generated reports in real-world healthcare contexts.

Discussion

The results confirm the capability of Generative AI, specifically fine-tuned LLMs, including T5, to compose reports related to cardiology in a clinically interpretable format and semantically consistent from structured data without any intermediaries.

It fills the gap in the long-term history between the numerical diagnostics and narrative reporting, allowing a paradigm shift to automated but interpretable clinical documentation.

Subsequent studies will expand this study by:

- integration of multimodal data (ECG and echocardiographic data and textual notes),
- That of instruction-tuned large LLMs (e.g., GPT-2, LLaMA-Med), and
- Using explainability modules to promote clinical trust and transparency.

Conclusion

This paper introduces a generative AI system based on transformers to generate semantic cardiology reports on structured patient data. The semantic alignment (BERTScore = 0.9768) and structural coherence (ROUGE-L = 0.8806) of the fine-tuned T5-small model were high, which proved that the model can be effective as an automated clinical documentation assistant. The suggested practice leads to the creation of explainable, general, and smart reporting solutions in contemporary digital healthcare.

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
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. AI Dis Detect Adv Appl. 2025:123-46. Available from: https://doi.org/10.1002/9781394278695.ch6
4. Gibb JJ, Kim WC, Barlatay FG. 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. 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. Franke M, Książczyk TM, Dux M, Chmielewski P, Truszkowska G, Czapczak D, Werner B, 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
7. Bernelli C, Di Fusco SA, Matteucci A. 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
8. Kasartzian DI, Tsiampalis T. Transforming cardiovascular risk prediction: a review of machine learning and artificial intelligence innovations. Life. 2025;15(1):94. Available from: https://doi.org/10.3390/life15010094
9. 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
10. Chowdhury MA, Rizk R, Chiu C. The heart of transformation: exploring artificial intelligence in cardiovascular disease. Biomedicines. 2025;13(2):427. Available from: https://doi.org/10.3390/biomedicines13020427
11. Ali MT, Ali U, Ali S. Transforming cardiac care: AI and machine learning innovations. Int J Multidiscip Res Growth Eval. 2024.
12. Tiwari E, Shrimankar D, Maindarkar M. 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
13. Chao CJ, Banerjee I, Arsanjani R, Ayoub C. 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
14. Wan Z, Liu C, Wang X. MEIT: Multi-modal electrocardiogram instruction tuning on large language models for report generation. arXiv [Preprint]. 2024 Mar [cited 2026 Jun 5]. Available from: arXiv:2403.04945.
15. Inam M, Sheikh S, Minhas AMK. 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
16. Khunte A, Sangha V, Oikonomou EK. Automated diagnostic reports from images of electrocardiograms at the point of care. medRxiv [Preprint]. 2024 [cited 2026 Jun 5].
17. Liu C, Wan Z, Ouyang C, Shah A. Zero-shot ECG classification with multimodal learning and test-time clinical knowledge enhancement. arXiv [Preprint]. 2024 Mar [cited 2026 Jun 5]. Available from: arXiv:2403.06659.
18. Liu C, Tian Y, Chen W, Song Y. Bootstrapping large language models for radiology report generation. Proc AAAI Conf Artif Intell. 2024 Mar;38(17):18635-43. Available from: https://doi.org/10.1609/aaai.v38i17.29826
19. Liu T, Krentz A, Lu L. 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
20. Zhao C, Xu Z, Baral P. Multi-graph graph matching for coronary artery semantic labeling. arXiv [Preprint]. 2024 Feb [cited 2026 Jun 5]. Available from: arXiv:2402.15894.
21. Nentidis A, Krithara A, Paliouras G. 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. 490-7.