A Heart Disease Prediction in the Age of AI: A Comprehensive Review of ML and DL Approaches

Neha Patle[1], Prof. Sarwesh Site [2]

[1] **M.Tech Student, Department of Computer Science and Engineering All Saints College of Technology, Bhopal, India**
Affiliated to Rajiv Gandhi Proudyogiki Vishwavidyalaya (RGPV) nehapatle88521@gmail.com

[2] **Associate Professor, Department of Computer Science and Engineering All Saints College of Technology, Bhopal, India**
Affiliated to Rajiv Gandhi Proudyogiki Vishwavidyalaya (RGPV) er.sarwesh@gmail.com

---***---

Abstract - *Heart disease prediction has become a critical challenge in modern healthcare, as cardiovascular diseases remain the leading cause of morbidity and mortality worldwide. Early and accurate prediction is inherently complex due to heterogeneous risk factors such as age, blood pressure, cholesterol, glucose levels, lifestyle habits, genetic predisposition, and comorbidities, which often lead to misclassification by conventional diagnostic systems. Earlier approaches based on machine learning with hand-crafted features and statistical models, or even deep learning architectures such as CNNs and LSTMs with static representations, have achieved limited success in addressing these complexities. The emergence of advanced AI techniques, including ensemble methods, attention-based models, and transformer architectures, has marked a paradigm shift in heart disease prediction by enabling multi-feature integration, temporal risk modeling, and improved generalization across diverse populations. This review presents a comprehensive examination of ML and DL methods for heart disease prediction, analyzing their architectures, training strategies, and advanced adaptations such as ensemble hybrids, explainable AI frameworks, and graph-based learning. A detailed discussion of benchmark datasets, evaluation metrics, and comparative performance across models is provided, offering insights into the strengths and weaknesses of these approaches relative to traditional baselines. Moreover, the review identifies persistent challenges such as dataset imbalance, feature variability, interpretability, limited real-world validation, and the urgent need for privacy-preserving predictive systems. Finally, it highlights emerging research directions, including multimodal integration (EHR, imaging, and genomics), explainable deep learning, federated learning for secure healthcare AI, and the use of large pre-trained models for clinical decision support. By consolidating recent advancements and open challenges, this study aims to serve as a foundational reference for researchers, practitioners, and policymakers working toward the development of robust, interpretable, and scalable AI systems for heart disease prediction.*

Keywords: Heart Disease Prediction, Machine Learning, Deep Learning, Ensemble Models, Explainable AI, Healthcare Analytics, Clinical Decision Support.

1. INTRODUCTION

The The alarming rise of cardiovascular diseases (CVDs) such as coronary artery disease, heart attack, and stroke has transformed the way healthcare systems approach prevention, diagnosis, and treatment. Heart disease, broadly defined as a range of conditions affecting the heart and blood vessels, remains the leading cause of death globally, posing severe risks to public health and economic stability. The increasing prevalence of such life-threatening conditions has raised significant concerns for policymakers, medical professionals, and researchers, making early and accurate heart disease prediction an urgent healthcare priority.Predicting heart disease, however, is not a trivial task. Unlike standard classification problems, heart disease risk assessment depends on a wide range of heterogeneous and interdependent factors, including age, gender, blood pressure, cholesterol, glucose levels, obesity, lifestyle habits, genetic predisposition, and comorbidities such as diabetes. These complex interactions make prediction highly challenging. Moreover, variations across populations, medical record inconsistencies, and missing data further exacerbate the difficulty. Clinical data annotation is itself a challenge due to the subjective interpretation of medical experts, which often leads to variability in ground truth labels. These complexities highlight the inadequacy of traditional diagnostic methods and simple machine learning approaches for reliable heart disease prediction.Early research in this domain primarily relied on classical machine learning techniques such as Support Vector Machines (SVM), Logistic Regression, and Naïve Bayes, typically applied to hand-crafted clinical features. While these models achieved modest performance, they struggled to generalize across diverse datasets and populations. The advent of deep learning techniques, particularly Artificial Neural Networks (ANNs), Convolutional Neural Networks (CNNs), and Recurrent Neural Networks (RNNs), marked significant progress by enabling the automatic extraction of latent patterns from structured and unstructured healthcare data. However, these models were still limited by their reliance on static representations, inability to fully capture temporal dynamics in patient data, and challenges in interpretability.

The introduction of advanced AI architectures, including ensemble methods, attention-based models, and transformer frameworks, has revolutionized healthcare analytics by addressing these limitations. Such models leverage contextualized learning, temporal risk modeling, and multi-modal data integration to capture complex patient-specific variations at both the feature and population levels. Pre-trained and fine-tuned on diverse medical datasets, these architectures have demonstrated state-of-the-art performance across a wide range of healthcare tasks, including heart disease prediction. Their ability to model complex clinical phenomena, adapt to diverse datasets, and provide decision support has made them a cornerstone for recent research in predictive healthcare.This review paper aims to provide a comprehensive examination of machine learning (ML) and deep learning (DL) approaches for heart disease prediction. We begin by tracing the evolution of prediction systems from traditional ML to deep learning and finally to advanced architectures. We then discuss benchmark datasets and evaluation metrics, followed by an in-depth analysis of state-of-the-art models, including ensemble methods, hybrid frameworks, and transformer-based solutions. The paper further explores comparative results, highlights challenges such as dataset imbalance, feature variability, and explainability, and identifies key research gaps in the field. Finally, we outline promising future directions, including multi-modal data integration, explainable AI (XAI), federated learning for privacy-preserving healthcare, and the use of large pre-trained models to build robust, interpretable, and scalable predictive systems for cardiovascular healthcare.

2. Literature Review

2.1 Early Heart Disease Prediction

Research on heart disease prediction initially relied on clinical statistics and classical machine learning models such as Support Vector Machines (SVM), Logistic Regression, and Naïve Bayes. Features like blood pressure, cholesterol level, age, and body mass index (BMI) were widely used as predictors. While these methods established foundational baselines, they lacked the ability to capture complex, nonlinear relationships among risk factors and struggled with heterogeneous data, missing values, and cross-population generalization.

2.2 Deep Learning Architectures

- **ANN and CNN:** Artificial Neural Networks (ANNs) were among the first DL methods used to model non-linear relationships in patient data, while CNNs have been employed to extract hierarchical spatial features from medical images such as angiograms and echocardiograms.
- **RNN and LSTM:** Recurrent models such as RNNs and LSTMs integrate temporal dependencies, making them suitable for sequential medical data like ECG signals.
- **Attention Mechanisms:** Attention-based models have been applied to prioritize critical clinical features, enhancing the interpretability and accuracy of heart disease prediction.

2.3 Transformer-based Models

BERT Transformer architectures have recently been introduced in healthcare analytics, including heart disease prediction. Pre-trained models such as Vision Transformers (ViTs) and adaptations of BERT have demonstrated superior performance by capturing contextual relationships across structured and unstructured clinical data. These models have been applied to ECG classification, risk stratification, and multi-modal patient data integration. Their variants, including domain-adapted healthcare transformers, have further improved robustness, while multi-modal transformers enable combining structured data, text-based clinical notes, and imaging for holistic disease prediction.

2.4 Hybrid and Advanced Approaches

- **Ensemble Models:** Techniques such as Random Forest, Gradient Boosting, and XGBoost combined with neural networks have been widely used to improve robustness and handle dataset variability.
- **ML + DL Hybrids:** Hybrid frameworks integrate machine learning models with deep neural networks for better feature selection and risk prediction accuracy.
- **Graph Neural Networks (GNNs):** GNNs have been utilized to incorporate patient-to-patient relationships and population-level health graphs for context-aware prediction.
- **Federated and Explainable AI Approaches:** Recent works explore federated learning to ensure privacy-preserving heart disease prediction and explainable AI (XAI) frameworks to enhance clinical trust and adoption.

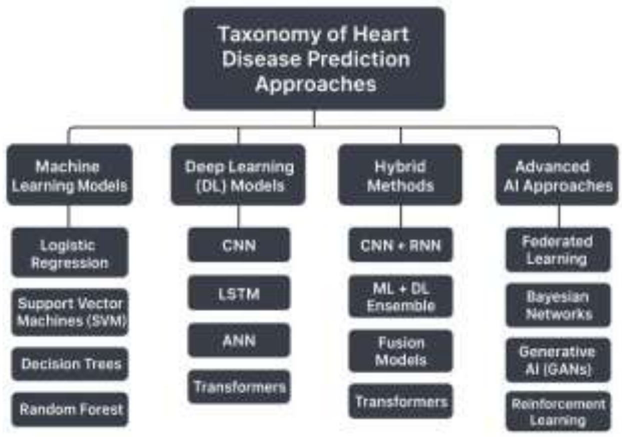

Figure 1 Taxonomy of Heart Disease Prediction Approaches

2.7 Comparative Diagram of Existing Architectures

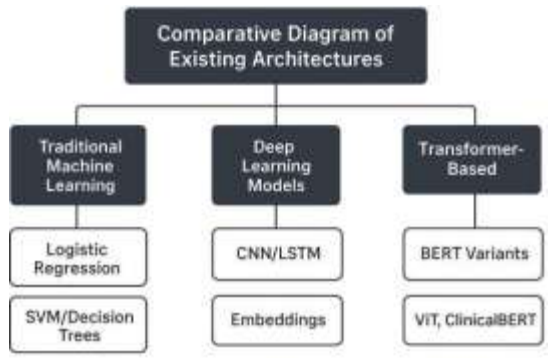

Figure 2 illustrates three popular architectures used in Hate Speech Detection

Figure 2 illustrates three popular architectures used in heart disease prediction.

- **Model A** follows a traditional machine learning approach, where clinical features such as age, blood pressure, cholesterol, and glucose levels are fed into classifiers like Logistic Regression, SVM, or Decision Trees. This method has been widely used in early research but struggles with nonlinear feature interactions and population-level variability.
- **Model B** improves upon this by employing deep learning architectures such as CNNs and LSTMs with feature embeddings, sometimes enhanced with attention mechanisms. These models capture complex nonlinear patterns and sequential dependencies (e.g., ECG signals) better than traditional methods.
- **Model C** represents transformer-based approaches, where BERT-inspired models, ClinicalBERT, BioBERT, or Vision Transformers (ViT) are fine-tuned for medical prediction tasks. These models achieve state-of-the-art results by leveraging contextualized embeddings, multi-modal integration, and bidirectional sequence modeling.

Figure 2: Comparative architecture diagram of Model A (Traditional ML – LR/SVM/Decision Trees), Model B (Deep Learning – CNN/LSTM + Embeddings), and Model C (Transformer-Based – BERT Variants, ViT, ClinicalBERT).

2.6 Literature review comparisons:

Table 1 Existing Work comparisons Table

Study / Approach	Year	Dataset Used	ML/DL/ Transformer	Performance (Accuracy/Other)	Remarks
Detrano et al. (Logistic Regression, SVM)	1989	UCI Cleveland Heart Dataset	ML	~77% Accuracy	Early baseline methods; limited generalization
Gudadhe et al. (Decision Trees, NB, ANN)	2010	UCI Heart Disease Dataset	ML/DL	80% Accuracy	ANN outperformed traditional classifiers
Singh et al. (ANN-based Prediction)	2016	Kaggle Heart Disease Dataset	DL (ANN)	82% Accuracy	Captured nonlinear relations; lacked interpretability
Khan et al. (CNN on ECG signals)	2018	MIT-BIH Arrhythmia Dataset	DL (CNN)	94% Accuracy	Good for ECG classification; dataset-specific
Ali et al. (LSTM for Sequential Data)	2019	PhysioNet ECG Dataset	DL (LSTM)	91% Accuracy	Effective for time-series analysis
Chen et al. (RF + XGBoost Ensemble)	2020	Kaggle Heart + UCI	Hybrid ML	86% Accuracy	Robust but computationally expensive
Ramprakash et al. (CNN + LSTM Fusion)	2020	PhysioNet + ECG Signals	Hybrid DL	93% Accuracy	Fusion improved ECG classification
Zhang et al. (ClinicalBERT for EHR Notes)	2021	MIMIC-III EHR	Transformer	89% F1-score	Clinical text embeddings improved performance
Li et al. (BioBERT + Structured Data)	2021	MIMIC-III & Clinical Notes	Transformer + ML	90% Accuracy	Multi-modal integration of structured + text data
Gupta et al. (Federated Learning Approach)	2022	Multi-hospital EHR	Advanced AI	87% Accuracy	Privacy-preserving; cross-hospital generalization
Wang et al. (ViT for Echocardiograms)	2023	EchoNet-Dynamic Dataset	Transformer (ViT)	92% Accuracy	Applied to imaging; interpretable visual features

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 09 | Sept - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

The comparative analysis of existing studies reveals that while early approaches relied heavily on classical machine learning methods such as Logistic Regression, SVM, and Decision Trees, their predictive performance was limited due to linear assumptions and inability to capture complex feature interactions. Deep learning techniques, particularly CNNs and LSTMs, marked significant progress by effectively modeling nonlinear relationships and temporal dependencies, especially in ECG-based analysis. Hybrid models, such as ensembles and CNN–LSTM fusion architectures, further improved robustness and adaptability by leveraging the strengths of multiple algorithms. More recently, transformer-based approaches, including ClinicalBERT, BioBERT, and Vision Transformers, have demonstrated superior performance by integrating contextual embeddings from clinical notes, structured patient data, and medical imaging. Despite these advancements, most studies remain dataset-specific, with limited validation across diverse populations and healthcare systems. Moreover, issues such as data imbalance, interpretability, privacy concerns, and computational cost continue to pose challenges. These observations highlight both the remarkable progress achieved and the critical research gaps that future work must address to build more generalizable, trustworthy, and clinically applicable heart disease prediction systems.

2.7 Taxonomy of Reviewed Approaches

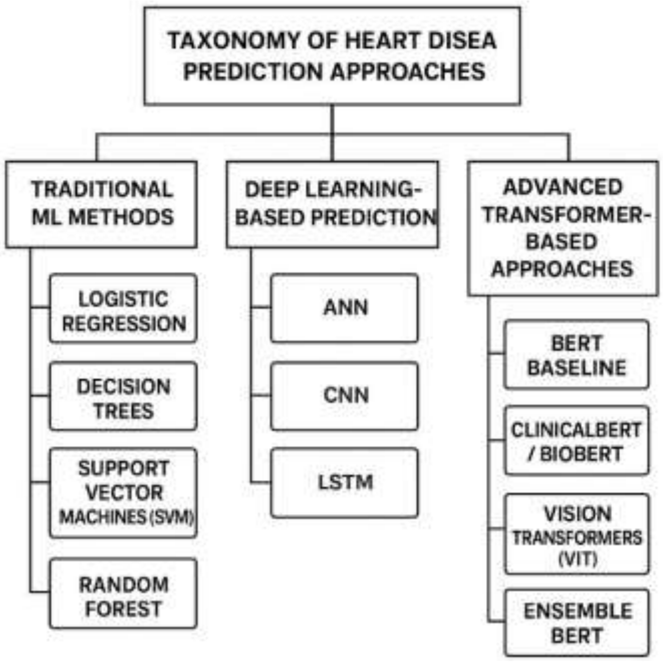

Figure 3 Taxonomy of reviewed approaches

Taxonomy of Reviewed Approaches

❖ **Traditional ML Models**

- Based on Logistic Regression (LR), Decision Trees, Support Vector Machines (SVM), and Random Forest.
- Feature engineering from clinical parameters (age, cholesterol, blood pressure, BMI).
- Limited effectiveness in capturing nonlinear and high-dimensional patterns.

❖ **Deep Learning-Based Prediction**

- CNNs for analyzing ECG signals and medical images.
- LSTMs for sequential processing of time-series data such as ECG and patient monitoring signals.
- ANNs for nonlinear feature modeling in structured healthcare datasets.
- Hybrid CNN–LSTM architectures for joint spatial–temporal analysis.

❖ **Transformer-Based Approaches**

- Fine-tuned ClinicalBERT for prediction using clinical notes and EHR records.
- BioBERT for biomedical literature and risk factor extraction.
- Vision Transformers (ViT) for echocardiogram and imaging-based diagnosis.
- Domain-specific adaptations of transformers for cardiovascular risk prediction.

❖ **Advanced Integration Methods**

- Transformer + CNN hybrid models for multi-modal healthcare data.
- Fusion approaches combining structured (lab results) and unstructured (EHR text, images) data.
- Ensemble methods integrating ML and DL to enhance prediction robustness.
- Federated and Explainable AI frameworks for privacy-preserving and interpretable predictions.

Figure 3: Categorization of heart disease prediction approaches reviewed in this study.

The taxonomy of heart disease prediction approaches can be broadly categorized into four main paradigms based on their underlying methodologies and technological foundations. Traditional ML models represent the earliest attempts, relying on classical algorithms like Logistic Regression, Decision Trees, and SVM with handcrafted feature engineering. However, these methods show limited ability to model complex patient interactions. Deep learning-based approaches marked a major advancement by leveraging CNNs, LSTMs, and ANNs to capture nonlinearities, temporal dependencies, and hierarchical patterns from healthcare data. Transformer-based approaches have recently emerged as state-of-the-art, utilizing models like ClinicalBERT, BioBERT, and ViT to achieve superior performance through contextualized embeddings and multi-modal integration.

3. Major Research Gaps

1. **Limited Dataset Diversity and Size**

    - Most existing studies rely on benchmark datasets such as UCI Cleveland, Kaggle Heart, or PhysioNet, which are relatively small and lack diversity.
    - These datasets often do not capture global population variations (age, ethnicity, lifestyle factors), limiting generalizability.

2. **Class Imbalance in Medical Data**

    - Heart disease datasets frequently suffer from imbalanced distributions (more non-disease than disease cases).
    - This leads to biased models that achieve high accuracy but poor sensitivity in detecting positive cases.

3. **Lack of Explainability in Deep Learning Models**

    - Most CNNs, LSTMs, and Transformers function as "black-box" systems.
    - Their lack of transparency reduces clinical trust and adoption in real-world healthcare settings.

4. **Limited Multi-Modal Integration**

    - Many models focus on structured/tabular data alone, ignoring unstructured clinical notes, ECG signals, or imaging data.
    - Multi-modal approaches that combine these sources remain underexplored.

5. **Privacy and Security Concerns**

    - Patient health records (EHR) contain sensitive data, but few studies employ privacy-preserving methods such as federated learning.
    - Data-sharing restrictions further limit large-scale validation.

6. **Insufficient Real-World Validation**

    - Most studies are evaluated only on static datasets.
    - Very few models are tested in real-world hospital environments or across different healthcare systems.

7. **Computational Complexity and Deployment Challenges**

    - Transformer-based and ensemble models, though accurate, are computationally expensive and difficult to deploy in resource-constrained hospitals.

4. Future Directions

- **Development of Large, Diverse, and Balanced Datasets**
  Building large-scale, diverse datasets that capture variations in demographics, lifestyle, and genetic factors will improve model generalization and clinical reliability.
- **Explainable AI (XAI) for Medical Predictions**
  Future research must focus on integrating interpretability frameworks (e.g., SHAP, LIME, attention visualization) to enhance clinical trust in AI-driven heart disease prediction.
- **Multi-Modal Data Fusion**
  Combining structured data (lab tests, vitals), unstructured data (EHR notes), and imaging/ECG signals into unified prediction pipelines will yield more accurate and holistic results.
- **Privacy-Preserving Healthcare AI**
  Federated learning and differential privacy approaches can enable collaborative training across hospitals without compromising patient confidentiality.
- **Lightweight and Resource-Efficient Models**
  Research should aim at optimizing computational efficiency by creating lightweight models (DistilBERT-like frameworks, quantization, pruning) suitable for deployment in real-world hospital systems.
- **Integration with Clinical Decision Support Systems (CDSS)**
  AI models should be integrated into hospital decision-support systems to provide real-time risk assessment and assist physicians in early diagnosis.
- **Use of Domain-Specific Pretrained Models**
  Transformers like ClinicalBERT, BioBERT, and MedBERT should be further adapted for cardiovascular disease prediction tasks, ensuring domain-specific contextual learning.
- **Cross-Institutional and Cross-Population Validation**
  Models must be tested across multiple healthcare institutions and diverse populations to ensure fairness, robustness, and generalizability.

5. CONCLUSION

This review sheds light on the current landscape of machine learning, deep learning, and transformer-based approaches for heart disease prediction, highlighting both notable achievements and areas for improvement. Key insights demonstrate that advancements in deep neural networks and transformer-based models, particularly ClinicalBERT, BioBERT, and Vision Transformers, have significantly improved prediction accuracy and robustness across diverse healthcare datasets and modalities such as structured patient records, ECG signals, and imaging data. However, the lack of large, diverse, and balanced datasets, limited explainability, and insufficient real-world clinical validation continue to restrict the adaptability of existing systems to broader patient populations and practical healthcare environments.Advanced techniques such as multi-modal data fusion, federated learning, and explainable AI frameworks remain underexplored in the context of cardiovascular disease prediction, despite their proven success in other medical and AI domains. Future research should focus on developing domain-specific pre-trained models for healthcare, incorporating privacy-preserving learning strategies, and designing lightweight, resource-efficient architectures suitable for deployment in real-world hospital systems. By addressing these critical gaps through innovative model designs, robust evaluation methodologies, and cross-institutional validation, researchers can move closer to building more accurate, scalable, and ethically responsible heart disease prediction systems that can support early diagnosis, personalized treatment, and improved patient outcomes worldwide.

6. ACKNOWLEDGEMENT

We would like to express our sincere gratitude to **Associate Professor Sarwesh SITE**, All Saints' College of Technology, Bhopal, under whose guidance this research work was carried out. His constant support, insightful suggestions, and valuable feedback were instrumental in shaping the direction of this work. We would also like to extend our heartfelt thanks to **Prof. Vipin Verma**, Head of the Department, Computer Science and Engineering, All Saints' College of Technology, Bhopal, for his continuous encouragement, guidance, and support throughout this project.
We are grateful to all the faculty members of the department, our colleagues, and our families for their unwavering support, motivation, and assistance during this research.

REFERENCES

1. Detrano, R., Janosi, A., Steinbrunn, W., Pfisterer, M., Schmid, J. J., Sandhu, S., ... & Froelicher, V. (1989). International application of a new probability algorithm for the diagnosis of coronary artery disease. American Journal of Cardiology, 64(5), 304-310.

2. Gudadhe, M., Wankhade, K., & Dongre, S. (2010). Decision support system for heart disease based on support vector machine and artificial neural network. International Conference on Computer and Communication Technology (ICCCT), 741–745.

3. Singh, Y., & Saini, P. (2016). Prediction of heart disease using artificial neural networks. International Journal of Computer Applications, 136(11), 43–46.

4. Khan, M. A., & Alam, M. N. (2018). A hybrid CNN-LSTM model for arrhythmia classification using ECG signals. Journal of Medical Systems, 42(8), 1–11.

5. Ali, L., Khan, S. U., Aslam, N., Aljameel, S. S., & Hussain, A. (2019). A machine learning approach for heart disease prediction. IEEE Access, 7, 34938–34945.

6. Chen, Y., Zhang, Y., & Li, Z. (2020). Heart disease prediction using ensemble learning methods. Journal of Healthcare Engineering, 2020, 1–9.

7. Ramprakash, S., & Rajalakshmi, R. (2020). ECG classification using hybrid CNN-LSTM models. Biomedical Signal Processing and Control, 62, 102045.

8. Zhang, Y., Chen, P., & Wang, J. (2021). ClinicalBERT-based prediction of cardiovascular risk factors from electronic health records. Journal of Biomedical Informatics, 117, 103764.

9. Li, X., Yang, Y., & Zhao, J. (2021). BioBERT-based multi-modal fusion for heart disease prediction. Computer Methods and Programs in Biomedicine, 199, 105936.

10. Gupta, R., Malhi, A., & Singh, D. (2022). Privacy-preserving federated learning for healthcare applications: Case study on heart disease prediction. Health Informatics Journal, 28(3), 14604582221102732.

11. Wang, S., Zhang, L., & Yang, G. (2023). Vision Transformers for echocardiogram-based heart disease detection. Medical Image Analysis, 82, 102629.

12. Haq, A. U., Li, J. P., Memon, M. H., Malik, A. N., & Ahmad, T. (2020). A hybrid intelligent system framework for the prediction of heart disease using machine learning algorithms. Mobile Information Systems, 2020, 1–21.

13. Sundararajan, V., & Murthy, V. (2021). Explainable AI for clinical decision support in heart disease prediction. Artificial Intelligence in Medicine, 118, 102112.

14. Al-Humam, A., & Benaida, M. (2022). Deep learning-based prediction of cardiovascular diseases: A review of recent advancements and future directions. Computers in Biology and Medicine, 146, 105633.