

A Structured Review of Machine Learning Methods for Early Heart Disease Prediction

Monika M M*, H N Prakash†

*Research Scholar, Dept. of CSE, Rajeev Institute of Technology, Hassan

Email: monikahassan46@gmail.com

†Professor, Rajeev Institute of Technology, Hassan

Email: prakash_hn@yahoo.com

Abstract—Cardiovascular disease (CVD) remains the foremost cause of mortality worldwide, underscoring the urgent need for intelligent and early-warning clinical decision support systems. Conventional single-modality diagnostic approaches are unable to holistically capture the multi-dimensional risk landscape of cardiac pathology. This survey reviews recent advances in machine learning and deep learning methods for early heart disease detection, with a focus on multimodal approaches that integrate electrocardiogram (ECG) signals, structured clinical parameters, demographic attributes, and lifestyle behavioural data. Key techniques examined include one-dimensional Convolutional Neural Networks cascaded with Gated Recurrent Units (1D-CNN-GRU) for spatiotemporal waveform feature extraction, Squeeze-and-Excitation Fully Connected Networks (SE-FCN) with residual skip connections for adaptive feature recalibration, and cross-attention fusion modules for modelling inter-modal dependencies. The survey also examines approaches to class imbalance remediation through SMOTE-ENN resampling, dimensionality reduction via Principal Component Analysis (PCA), and clinical transparency through SHapley Additive exPlanations (SHAP). Research gaps, open challenges, and future directions are discussed.

Index Terms—Cardiovascular Disease Prediction, Multimodal Deep Learning, CNN-GRU, Squeeze-and-Excitation Network, Cross-Attention Fusion, SMOTE-ENN, SHAP, ECG Signal Analysis, Clinical Decision Support

I. INTRODUCTION

Cardiovascular diseases (CVDs) account for approximately 17.9 million deaths each year, representing 31% of all global mortality [4]. The complexity of CVD onset—driven by a convergence of physiological abnormalities, lifestyle risk factors, and comorbidities—makes early and accurate prediction a formidable challenge. Conventional diagnostic pathways, including resting electrocardiography, echocardiography, and biochemical biomarker panels, tend to be episodic, resource-intensive, and incapable of integrating heterogeneous risk signals into a unified predictive framework.

The emergence of machine learning (ML) and deep learning (DL) has opened new avenues for automated cardiac risk stratification. Classical ensemble methods such as Random Forest and XGBoost have demonstrated competitive performance on structured clinical tabular data [7], while convolutional and recurrent neural network architectures have shown strong capability in learning discriminative features from raw electrocardiographic signals [11]. More recently, hybrid models that combine multiple network paradigms [13] and transformer-

based multimodal architectures [8] have substantially advanced the state of the art.

Despite these advances, significant limitations persist. First, most existing systems analyse a single data modality in isolation, thereby losing complementary diagnostic information present across physiological signal streams and structured clinical records. Second, medical datasets commonly exhibit pronounced class imbalance that biases model learning towards the majority class. Third, the black-box nature of many DL frameworks impedes clinical adoption, as practitioners require transparent and actionable explanations for algorithmic predictions.

This survey provides a structured review of relevant publications sourced from IEEE Xplore, PubMed, and Google Scholar, with particular emphasis on studies addressing multimodal cardiac risk prediction, deep learning for ECG analysis, explainable AI, and clinical decision support. The key themes examined are:

- 1) **Multimodal and dual-pathway deep learning architectures** that process ECG temporal signals and structured clinical-lifestyle data through dedicated, modality-specific encoders prior to fusion.
- 2) **Squeeze-and-Excitation** and residual network designs for adaptive channel-wise feature recalibration in tabular clinical data.
- 3) **Cross-attention fusion modules** that explicitly model inter-modal dependencies to capture synergistic cardiovascular risk indicators.
- 4) **SHAP-based post-hoc explainability** methods that provide quantitative, clinically meaningful feature importance rankings at both global and instance levels.

The remainder of this paper is structured as follows. Section II surveys related work across six research directions. Section III describes the common datasets and preprocessing strategies employed in the reviewed works. Section IV presents a taxonomy of architectural frameworks identified in the literature. Section V presents a synthesis of a representative multimodal pipeline. Section VI discusses key findings, research gaps, and future directions. Section VII concludes the survey.

A. Review Methodology

This review follows a structured literature search process. Relevant studies were identified through searches of IEEE Xplore, PubMed, and Google Scholar using combinations of keywords related to heart disease, cardiovascular disease, machine learning, deep learning, prediction, detection, ECG, multimodal learning, and explainable AI. The review primarily considered publications from 2017–2025.

Studies were included when they addressed machine learning or deep learning methods for cardiovascular risk prediction or detection and provided sufficient methodological information for analysis. Studies that were not directly related to cardiovascular applications or lacked sufficient methodological detail were excluded.

Foundational works were additionally included when they provided the theoretical basis for techniques used in the reviewed studies, such as SHAP, residual learning, and SMOTE (e.g. [14]–[16]). These works are explicitly identified as foundational references in Table I.

II. LITERATURE REVIEW

Recent progress in AI-driven cardiovascular risk assessment spans six principal research directions: traditional machine learning approaches, deep learning models applied to biosignals, multimodal or hybrid frameworks, explainable AI methods, ECG signal analysis, and clinical decision support systems. Table I summarises the reviewed works across these directions.

A. Machine Learning Approaches

Classical supervised algorithms including Logistic Regression, Naive Bayes, and Support Vector Machines (SVM) were among the first to be applied to benchmark cardiac datasets. Vasavi and Sujatha [7] conducted a comparative evaluation of five ML algorithms on the Framingham Heart Study dataset, concluding that stacked ensemble methods consistently outperform individual classifiers in this domain. Importantly, their SHAP analysis established a clear feature importance hierarchy, with cholesterol level, maximum heart rate, and age identified as the most predictive variables—findings that align closely with established clinical risk literature and provide a useful interpretability baseline for comparing subsequent deep learning approaches. Xue et al. [19] demonstrated that Random Forest combined with information fusion can achieve a high true positive rate for warfarin dosage prediction in cardiac surgery patients, illustrating that classical ensemble methods retain clinical utility even alongside more complex DL architectures, particularly in data-scarce or high-interpretability settings.

B. Deep Learning for Biosignal Analysis

Convolutional Neural Networks have been widely adopted for automatic feature learning from ECG time-series data, eliminating the need for hand-crafted feature engineering. Hybrid CNN-LSTM frameworks extend this capacity by simultaneously modelling local morphological patterns and long-range temporal dependencies. Dwivedi et al. [11] proposed

a Healthcare-IoT pipeline combining SVM for preliminary ECG screening with an Artificial Neural Network (ANN) for refined classification, demonstrating that a two-stage architecture can achieve high accuracy while remaining computationally practical for resource-constrained IoT deployments—an important practical consideration absent from most purely academic benchmarks. Panchumarthi et al. [12] introduced HRAE-LSTM, which integrates residual attention mechanisms with Bidirectional LSTM layers; their use of SMOTE-ENN resampling prior to training is notable as it directly addresses the class imbalance problem that plagues most cardiac datasets, and their strong results on the UCI Cleveland dataset demonstrate the combined benefit of architecture design and data-level remediation.

C. Hybrid and Multimodal Frameworks

Fusing data from multiple modalities has consistently demonstrated performance gains over single-source systems, though the specific fusion strategy varies considerably across works. Sherly et al. [8] proposed a transformer-based architecture combining a Vision Transformer for ECG spectrogram analysis with a Tab Transformer for structured clinical features; the use of attention mechanisms at both the intra-modal and inter-modal level represents a principled approach to capturing complex dependencies, though the computational cost is substantially higher than CNN-based alternatives. By contrast, Barla and Madam [5] combined Self-Attention Assisted Dense Capsule Networks with Chaotic Dwarf Mongoose Optimisation (CHDMO), a bio-inspired hyperparameter search strategy, demonstrating that architecture search can yield competitive results even without transformer-scale parameter counts. Macha [13] showed that a CNN-GRU hybrid with SHAP interpretability achieves strong performance on a large synthetic dataset, though the use of synthetic rather than real clinical data limits the generalisability of the findings. Karthikeyan and Sathiyamoorthy [9] uniquely extended multimodal fusion to brain-heart interaction analysis, combining EEG and ECG signals within a CNN-BiLSTM architecture, opening a promising but underexplored direction for patients with comorbid neurological and cardiac conditions.

D. Explainable AI in Cardiac Prediction

Explainability is increasingly recognised as a prerequisite for clinical deployment rather than an optional feature. Wu et al. [6] proposed ResSE-FNN-XGBoost, which incorporates squeeze-and-excitation blocks into a residual fully connected network; their SHAP analysis on the Framingham dataset identified blood pressure medications, systolic blood pressure, and education level as dominant predictors, with the education-level finding suggesting that socioeconomic factors may carry predictive signal that purely physiological models miss. Selvi et al. [10] combined PCA-based dimensionality reduction with Relief feature selection and a hybrid DenseNet-Residual model, demonstrating that coupling feature selection with deep learning can yield competitive results with lower model complexity. Lundberg and Lee [14] provided the theoretical foundation

for SHAP, whose adoption across multiple reviewed works underscores its status as the de facto standard for post-hoc interpretability in cardiac AI.

E. Deep Learning for ECG Signal Analysis

ECG signals present unique challenges for machine learning due to their high sampling rates, susceptibility to noise, and the need to capture both local waveform morphology and global rhythm patterns. CNNs address the former while recurrent architectures such as LSTM and GRU address the latter [9], [12]. Several reviewed works combine these complementary strengths within a single end-to-end architecture [11], [13]. A consistent finding across these works is that preprocessing quality—particularly wavelet-based denoising and spectrogram generation—significantly affects downstream model performance, suggesting that signal preprocessing deserves as much attention as architecture design [8], [10].

F. Clinical Decision Support Systems

Several studies have examined the integration of AI models into clinical decision support systems (CDSS), highlighting the gap between algorithmic performance and real-world deployment. Sutton et al. [20] identified clinician trust, workflow integration, and regulatory compliance as the primary barriers to CDSS adoption, noting that technical performance alone is insufficient for successful deployment. Mai et al. [18] addressed personalised therapy by developing a drug recommender system for hypertension management, illustrating how AI-driven CDSS can extend beyond diagnosis to treatment recommendation. Abul-Husn and Kenny [21] highlighted the additional value of integrating EHR data with genomic information for precision cardiac medicine. Wairimu [22] further demonstrated that lifestyle data integration improves both monitoring outcomes and intervention effectiveness, underscoring the importance of behavioural factors that purely physiological models tend to overlook.

III. DATASETS AND PREPROCESSING

The reviewed works employ a range of publicly available cardiovascular datasets and share several common preprocessing strategies. This section summarises the most frequently used datasets and the preprocessing techniques identified across the literature.

A. Commonly Used Datasets

1) *PTB-XL Electrocardiography Dataset*: The PTB-XL dataset is a large-scale, publicly available ECG dataset developed by PhysioNet and widely used for cardiovascular signal analysis and deep learning research [1]. It contains 21,837 clinical 12-lead ECG recordings collected from 18,885 patients, with each recording sampled at either 100 Hz or 500 Hz over a 10-second window. The dataset covers multiple diagnostic categories including myocardial infarction, conduction disturbances, hypertrophy, ST/T changes, and arrhythmias, making it suitable for multi-label cardiac classification tasks.

2) *Framingham Heart Study Dataset*: The Framingham Heart Study dataset is a longitudinal cardiovascular cohort dataset widely used for heart disease risk prediction [3]. It contains 4,240 patient records with 15 clinical and lifestyle features including demographic information, smoking history, BMI, glucose level, cholesterol, and blood pressure. The binary target variable represents the 10-year risk of coronary heart disease. It is a common benchmark for evaluating class-imbalance remediation strategies, as the at-risk category constitutes a minority of total samples.

3) *MIT-BIH Arrhythmia Database*: The MIT-BIH Arrhythmia Database is extensively used for arrhythmia analysis and automated cardiac diagnosis [2]. It consists of standard two-lead ECG recordings sampled at 360 Hz with annotated heartbeat abnormalities. It has been widely adopted in hybrid deep learning studies to evaluate CNN and recurrent architectures on temporal cardiac signal data.

4) *UCI Cleveland Heart Disease Dataset*: The UCI Cleveland Heart Disease Dataset [17] is one of the most widely cited benchmarks in cardiac ML research. It contains 303 patient records with 13 clinical features. Several reviewed works use this dataset for comparative evaluation due to its small but well-curated nature.

B. Preprocessing Strategies

1) *Missing Value Imputation*: Missing clinical feature values can be addressed using k -Nearest Neighbours (KNN) imputation, which estimates missing values from similar patient records [10].

2) *Class Imbalance Remediation*: Synthetic Minority Over-sampling (SMOTE) [16] is a widely used technique for generating synthetic minority-class samples to counteract class imbalance. Several reviewed works pair SMOTE with Edited Nearest Neighbours (ENN) filtering, a complementary technique that removes noisy or borderline samples near the decision boundary, forming the combined SMOTE-ENN pipeline that is adopted across the reviewed works to produce cleaner and more balanced training distributions [12].

3) *Feature Normalisation and Dimensionality Reduction*: Structured clinical features are typically scaled to a common range via Min-Max normalisation. Several studies additionally apply PCA for dimensionality reduction, retaining the variance most relevant to cardiac risk prediction [6], [10].

4) *ECG Signal Preprocessing*: Wavelet transform-based denoising is commonly used to remove baseline wander and high-frequency noise from ECG signals. Short-Time Fourier Transform (STFT) spectrograms are subsequently generated to produce time-frequency representations suitable for input to convolutional architectures.

IV. TAXONOMY OF ARCHITECTURAL FRAMEWORKS

Based on the reviewed literature, the methods surveyed can be organised into a four-level taxonomy, as illustrated in Figure 1. At the top level, approaches are divided into *unimodal* and *multimodal* systems. Unimodal systems are further subdivided into those operating on *tabular clinical data*

TABLE I
COMPARATIVE SUMMARY OF SELECTED STUDIES AND FOUNDATIONAL WORKS

No.	Author(s)	Methodology	Dataset(s)	Modality	XAI?	Key Limitation
1	Barla & Madam [5]	SDCN + CH-DMO + ResVGG	Not specified	Tabular	No	High architectural complexity; limited reproducibility
2	Wu et al. [6]	ResSE-FNN + XGBoost + SHAP	Framingham, Cleveland	Tabular	Yes (SHAP)	No ECG modality included
3	Vasavi & Sujatha [7]	RF + XGBoost (Stacked)	Framingham	Tabular	Yes (SHAP)	Single dataset; no temporal signal
4	Sherly et al. [8]	ViT + Tab-Transformer + Fusion	MIT-BIH, UCI	Multimodal	No	High compute cost; limited clinical validation
5	Karthikeyan & Sathiyamoorthy [9]	CNN-BiLSTM (EEG + ECG)	Proprietary	Multimodal	No	External validation absent
6	Selvi et al. [10]	PCA + DenseNet + Residual	Cleveland	Tabular	Partial	Small dataset; no ECG signal
7	Dwivedi et al. [11]	SVM + ANN + IoT	ECG benchmark	ECG	No	Two-stage pipeline; limited end-to-end learning
8	Panchumarthi et al. [12]	HRAE-LSTM + SMOTE-ENN	Cleveland	Tabular	No	No multimodal fusion
9	Macha [13]	CNN-GRU + SHAP	Synthetic	ECG	Yes (SHAP)	Synthetic data limits real-world applicability
10	Xue et al. [19]	Random Forest + Info. Fusion	Clinical longitudinal	Tabular	No	Domain-specific; not generalisable
11	Sutton et al. [20]	CDSS overview (survey)	N/A	N/A	N/A	No ML model proposed
12	Mai et al. [18]	Drug Recommender (CDSS)	Hypertension registry	Tabular	No	Narrow scope; single condition
13	Abul-Husn & Kenny [21]	EHR + Genomic Integration	EHR + Genomic	Multimodal	No	Privacy and consent challenges
14	Wairimu [22]	AI + Lifestyle Data	Public health data	Tabular	No	Lacks ECG or clinical validation
15	Lundberg & Lee [14]	SHAP (theoretical)	N/A	N/A	Yes (SHAP)	Foundational; not applied to cardiac data directly
16	He et al. [15]	Residual Learning (theoretical)	ImageNet	Image	No	Foundational; not cardiac-specific
17	Chawla et al. [16]	SMOTE (theoretical)	N/A	N/A	No	Foundational; does not address boundary noise

(using classical ML or FCN-based architectures) and those operating on *ECG temporal signals* (using CNN, LSTM, GRU, or hybrid designs). Multimodal systems are subdivided into *early fusion*, where raw inputs are concatenated prior to learning, and *late fusion* or *cross-attention fusion*, where modality-specific representations are first independently encoded and then merged. A cross-cutting dimension applicable to all four branches is the presence or absence of *explainability*, primarily implemented via SHAP across the reviewed works.




Fig. 1. Taxonomy of machine learning architectures reviewed for early heart disease detection, classified by data modality, fusion strategy, and architecture type. SHAP-based explainability (XAI) is applied across multiple branches of this taxonomy.

This taxonomy reveals a clear trend in the literature: the field has progressively shifted from unimodal tabular classifiers towards multimodal, attention-based architectures with integrated explainability. However, the majority of reviewed multimodal works still lack external clinical validation, representing the most significant open gap identified in this survey.

A. ECG Temporal Signal Encoders

The dominant design combines 1D-CNN layers for local morphological feature extraction with GRU or LSTM layers for long-range sequential dependency modelling [9], [13]. The Bidirectional GRU variant processes sequences in both temporal directions, enabling richer spatiotemporal representations and providing an alternative to unidirectional recurrent designs.

B. Structured Data Encoders

For tabular clinical data, FCNs augmented with Squeeze-and-Excitation (SE) blocks and residual skip connections are the most effective design identified in this survey [6]. SE blocks perform channel-wise feature recalibration, selectively emphasising the most informative clinical features. Residual connections [15] facilitate gradient flow and enable deeper networks without performance degradation. This SE-FCN design demonstrates competitive performance against plain FCN and classical ML baselines on the evaluated cardiovascular datasets.

C. Multimodal Fusion via Cross-Attention

Cross-attention fusion, as employed by Sherly et al. [8], represents the most principled approach to combining heterogeneous modalities identified in this review. Unlike simple concatenation or weighted averaging, cross-attention allows each modality stream to selectively attend to the most diagnostically relevant features of the other, capturing synergistic risk interactions that neither stream can model independently. This design has shown superior performance over single-modality baselines on combined ECG and clinical datasets.

D. SHAP-Based Explainability

SHAP [14] is one of the most frequently used post-hoc explainability methods across the reviewed works [6], [7], [13]. By decomposing predictions into additive per-feature contributions, SHAP enables both global feature importance rankings and local instance-level explanations. Across the reviewed works, SHAP consistently identifies age, cholesterol level, and blood pressure as dominant predictors, validating that model behaviour aligns with established clinical knowledge.

V. SYNTHESIS: A REPRESENTATIVE MULTIMODAL FRAMEWORK

To illustrate how the architectural components identified in Section IV are combined in practice, this section presents a system-level synthesis of a representative multimodal pipeline for early cardiovascular disease detection, drawn from the reviewed literature.

A. High-Level Pipeline

Figure 2 presents the complete prediction pipeline from raw data acquisition to final risk classification. The pipeline encompasses ECG signal acquisition and spectrogram conversion, structured clinical data ingestion, modality-specific feature encoding through dedicated network branches, cross-attention fusion of multimodal representations, and final binary

classification into cardiovascular risk categories. SHAP-based interpretability is applied after prediction to generate feature attribution maps for clinical transparency, addressing the trust barrier identified by Sutton et al. [20].




Fig. 2. High-Level Architecture of a Representative Multimodal Cardiovascular Risk Prediction Pipeline, synthesised from the reviewed literature.

B. Low-Level Component View

Figure 3 details the internal structure of the core processing modules. The ECG branch includes a wavelet transform denoiser, STFT spectrogram generator, 1D-CNN feature extractor, and Bidirectional GRU sequential encoder. The structured data branch encompasses KNN imputation, SMOTE-ENN resampling, PCA dimensionality reduction, and an SE-FCN with residual blocks. The fusion layer applies bidirectional cross-attention and combines both modality streams for the final classification decision.




Fig. 3. Low-Level Component View of the Representative Multimodal Pipeline.

VI. DISCUSSION

A. Key Findings and Cross-Paper Analysis

Three consistent findings emerge from the reviewed literature. First, multimodal architectures that fuse ECG temporal signals with structured clinical data generally demonstrate promising performance compared with unimodal approaches in the reviewed studies, confirming that complementary diagnostic information exists across modalities and that single-source systems are inherently limited in their predictive capacity [8], [12]. The performance advantage of multimodal systems is most pronounced on datasets with significant class imbalance, where the additional signal from a second modality helps the classifier learn more robust decision boundaries.

Second, the choice of fusion strategy matters significantly. Attention-based fusion mechanisms provide a more flexible approach than simple feature concatenation by allowing each modality to selectively attend to information from the other modality [8]. This suggests that the benefit of multimodal fusion is not merely additive but arises from modelling inter-modal dependencies that neither stream can capture independently.

Third, SHAP-based explainability is not merely a reporting tool but influences model behaviour: studies that incorporate explainability into their development pipeline tend to identify and correct feature leakage and spurious correlations that purely accuracy-focused pipelines miss [6], [7]. This finding has direct implications for clinical deployment, where regulatory bodies increasingly require that AI-assisted diagnostic tools provide auditable reasoning.

B. Research Gaps

Despite strong benchmark performance, several critical gaps persist. Most multimodal frameworks require complete, synchronised data from all modalities at inference time, which is not consistently achievable in real clinical settings where certain measurements may be missing or delayed [7], [13]. No reviewed work proposes a principled mechanism for handling partially missing modalities at test time, representing a significant barrier to deployment.

External validation across demographically diverse and geographically distinct patient cohorts is largely absent from the reviewed works. The majority of results are reported on a small set of well-curated public benchmarks—primarily the UCI Cleveland and Framingham datasets—which do not reflect the heterogeneity of real-world clinical populations. This limits confidence in the generalisability of the reported performance figures.

Ethical and regulatory considerations, including patient data consent, algorithmic bias across demographic subgroups, and compliance with medical device regulations, remain substantially underexplored. Only Sutton et al. [20] and Abul-Husn and Kenny [21] address these concerns directly, and neither does so in the context of a deployed deep learning system.

C. Implications and Future Directions

The findings of this survey point to four productive directions for future research. First, the development of *modality-robust* architectures that gracefully degrade when one or more input modalities are unavailable at inference time would substantially improve clinical applicability. Second, *federated learning* protocols would enable privacy-preserving multi-institutional model training, addressing both the data heterogeneity problem and the regulatory concerns identified above [9]. Third, incorporation of *imaging modalities* such as echocardiography and coronary CT angiography into the fusion pipeline would provide richer anatomical context than ECG signals alone. Fourth, integration of *genomic biomarkers* [21] and lifestyle factors [22] would enable personalised risk stratification that accounts for individual genetic predisposition alongside physiological and behavioural risk.

VII. CONCLUSION

This paper has presented a structured review of machine learning and deep learning methods for early heart disease detection, covering recent primary studies and foundational works across six research directions. The reviewed works indicate that multimodal deep learning approaches, which fuse ECG temporal signals with structured clinical and lifestyle data, can provide promising performance compared with single-modality and classical machine learning approaches on selected cardiovascular benchmarks. Key architectural patterns identified in the literature include 1D-CNN-GRU temporal encoders, Squeeze-and-Excitation FCN designs for structured tabular data, cross-attention fusion for inter-modal dependency modelling, and SMOTE-ENN resampling for class imbalance remediation.

SHAP-based explainability has emerged as an essential component for clinical adoption, consistently identifying established cardiovascular risk factors such as age, cholesterol level, and systolic blood pressure as dominant predictors, thereby validating the clinical coherence of the reviewed models. Despite strong benchmark performance, critical gaps remain in modality-robustness at inference time, external validation across diverse populations, and regulatory and ethical frameworks for real-world deployment.

The taxonomy and cross-paper synthesis presented in this survey provide a structured foundation for future work, and the identified gaps point to clear and impactful research directions in federated learning, imaging modality integration, and genomic biomarker incorporation for personalised cardiac risk stratification.

ACKNOWLEDGMENT

The authors thank the maintainers of the UCI Machine Learning Repository, the Framingham Heart Study public data repository, and the PhysioNet consortium (MIT-BIH Arrhythmia Database) for making their datasets freely accessible to the research community.

REFERENCES

- [1] P. Wagner, N. Strodthoff, R. Bousseljot, D. Kreiseler, F. I. Lunze, W. Samek, and T. Schaeffter, "PTB-XL, a Large Publicly Available Electrocardiography Dataset," *Scientific Data*, vol. 7, no. 154, 2020.
- [2] G. B. Moody and R. G. Mark, "The MIT-BIH Arrhythmia Database on CD-ROM and Software for Use with It," *Computers in Cardiology*, vol. 17, pp. 185–188, 1990.
- [3] T. R. Dawber, G. F. Meadors, and F. E. Moore, "Epidemiological Approaches to Heart Disease: The Framingham Study," *American Journal of Public Health*, vol. 41, no. 3, pp. 279–286, 1951.
- [4] World Health Organization, "Cardiovascular Diseases (CVDs) Fact Sheet," WHO, 2022. [Online]. Available: [https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-\(cvds\)](https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds))
- [5] M. B. Barla and A. K. Madam, "A Hybrid Deep Learning Framework for Early Cardiovascular Disease Prediction: SDCN-CHDMO-ResVGG Approach," in *Proc. 9th Int. Conf. Electronics, Communication and Aerospace Technology (ICECA)*, IEEE, 2025, pp. 1049–1053.
- [6] Y. Wu, M. Huang, Y. Shi, Q. Li, Y. Zhang, and B. Song, "An Interpretable Hybrid Deep Learning Method for Heart Disease Prediction," in *Proc. 5th Int. Conf. Digital Society and Intelligent Systems (DSInS)*, IEEE, 2025, pp. 248–252.
- [7] Ch. Vasavi and G. Sujatha, "Data Driven Heart Disease Prediction: A Comparative Study of ML Models," in *Proc. 4th Int. Conf. Innovative Mechanisms for Industry Applications (ICIMIA)*, IEEE, 2025, pp. 1711–1715.
- [8] I. Sherly S., T. A. Mohanaprakash, T. Sindhu, N. Subramanian, D. Bora, and K. Cinthuja, "Data Transformer – AI Driven Multimodal Fusion Network for Cardiovascular Risk Prediction Using Multimodal Data," in *Proc. 5th Int. Conf. Evolutionary Computing and Mobile Sustainable Networks (ICECMSN)*, IEEE, 2025, pp. 203–212.
- [9] L. Karthikeyan and E. Sathiyamoorthy, "Deep Learning for Predicting Cardiovascular Events in Patients with Neurological Disorders: A Multimodal Analysis of Brain-Heart Interactions," in *2025 Innovations in Power and Advanced Computing Technologies (i-PACT)*, IEEE, 2025, pp. 1–5.
- [10] C. Thirumarai Selvi, P. Seleena, R. S. Sankara Subramanian, A. K. Shruthika, P. Gandhi Ayyavu, and S. Srinidhi, "Heart Disease Prediction using Feature-Optimized ML and Hybrid Deep Learning Models," in *Proc. 3rd Int. Conf. Sustainable Computing and Data Communication Systems (ICSCDS)*, IEEE, 2025, pp. 1025–1028.
- [11] A. K. Dwivedi, A. Yadav, and R. S. Gamad, "Heart Disease Prediction using Hybrid Machine Learning with ECG Dataset for Healthcare-IoT Application," in *Proc. 4th OPJU Int. Technology Conference (OTCON)*, IEEE, 2025.
- [12] G. Panchumarthi, B. Vatram, A. S. Mothad, A. N. Shaik, S. Lalitha, N. Kalyani, and S. Moturi, "Hybrid Deep Learning Model Using Residual Attention and BiLSTM for Accurate Ischemic Heart Disease Classification," in *Proc. IEEE 6th Global Conf. for Advancement in Technology (GCAT)*, IEEE, 2025, pp. 1–6.
- [13] Y. Macha, "Predictive Analytics in Healthcare: Deep Learning Models for Early Cardiovascular Risk Assessment," in *Proc. IEEE Pune Section Int. Conference (PuneCon)*, IEEE, 2025.
- [14] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," *Advances in Neural Information Processing Systems*, vol. 30, pp. 4765–4774, 2017.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [16] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-Sampling Technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.
- [17] D. Dua and C. Graff, "UCI Machine Learning Repository," University of California, Irvine, School of Information and Computer Sciences, 2019. [Online]. Available: <https://archive.ics.uci.edu>
- [18] A. Mai, K. Voigt, J. Schübel, and F. Gräßer, "A Drug Recommender System for the Treatment of Hypertension," *BMC Medical Informatics and Decision Making*, vol. 23, no. 89, 2023.
- [19] L. Xue, S. He, R. K. Singla *et al.*, "Machine Learning Guided Prediction of Warfarin Blood Levels for Personalized Medicine Based on Clinical Longitudinal Data from Cardiac Surgery Patients," *International Journal of Surgery*, vol. 110, pp. 6528–6540, 2024.
- [20] R. T. Sutton, D. Pincock, D. C. Baumgart *et al.*, "An Overview of Clinical Decision Support Systems: Benefits, Risks, and Strategies for Success," *npj Digital Medicine*, vol. 3, no. 17, 2020.
- [21] N. S. Abul-Husn and E. E. Kenny, "Personalized Medicine and the Power of Electronic Health Records," *Cell*, vol. 177, no. 1, pp. 58–69, 2019.
- [22] K. Wairimu H., "Integrating Lifestyle Data into Personalized Health Solutions," *Research Output Journal of Public Health and Medicine*, vol. 5, no. 1, pp. 100–106, 2025.