Cureus Journal of Computer Science

Open Access Letter

Early Detection of Cardiac Abnormalities Based on Deep Learning Approaches

P Shoba [1] , Azha Periasamy [1]

1. Department of Electronics and Instrumentation, Bharathiar University, Coimbatore, IND

Corresponding authors: P Shoba, shobamurali2007@gmail.com, Azha Periasamy, azhaperiasamy65@buc.edu.in

Received 04/16/2025
Review began 04/30/2025
Review ended 05/16/2025
Published 05/21/2025

DOI: https://doi.org/10.7759/s44389-025-04619-9

Abstract

Cardiovascular disease is the leading cause of death worldwide and, as such, requires early and reliable screening methods. Traditional detection techniques are mostly time-consuming and resource-intensive. In this work, we provide an overview of some deep learning techniques for detecting heart abnormalities. We train and evaluate various models, including a multi-layer perceptron (MLP), recurrent neural network (RNN), convolutional neural network (CNN), deep Q-network (DQN), autoencoders, and bidirectional long short-term memory (Bi-LSTM). Our findings show that the MLP model achieved the highest accuracy of 86.88%, followed by the RNN and Bi-LSTM models, each with an accuracy of 85.24%. While the CNN and DQN models performed moderately well, the autoencoder model performed relatively poorly, with an accuracy of only 49.18%. These results suggest that deep learning approaches, specifically MLP, RNN, and Bi-LSTM, are effective tools for the early detection of cardiac abnormalities. Future work may focus on improving model performance through feature engineering, hybrid architectures, and the use of larger datasets.

Categories: Bioinformatics Algorithms, Biotechnology and Computational Biology, Deep Learning
Keywords: cardiac abnormalities, deep learning, multi-layer perceptron, recurrent neural network, convolutional neural network, autoencoder, bi-lstm, heart disease detection

Introduction

Cardiovascular diseases (CVDs) are the leading cause of death worldwide, responsible for nearly 8 million fatalities annually. Timely detection of cardiac abnormalities is essential for effective intervention, reducing mortality rates, and improving patient outcomes [1]. Traditional diagnostic techniques, such as electrocardiograms, echocardiography, and stress tests, require specialized medical expertise and may not always be accessible, particularly in remote areas. However, advancements in deep learning and artificial intelligence have led to the development of automated screening systems [2], which now serve as important tools for the early identification of cardiac disorders. Deep models, inspired by the neural structure of the human brain, have transformed several fields [3]. These models can detect hidden trends in medical data that human expertise may not recognize. Deep learning provides a non-invasive, cost-effective, and accurate diagnostic approach for heart disease [4]. It helps healthcare providers identify individuals at risk and take necessary preventive measures with the assistance of structured clinical data.

This study explores various deep learning algorithms for early diagnosis of heart problems using the Kaggle Heart Disease Dataset. The dataset contains 14 clinical variables, including age, gender, blood pressure, cholesterol levels, and others. Six deep learning architectures were used to test the effectiveness of various models: multi-layer perceptron (MLP), recurrent neural network (RNN), convolutional neural network (CNN), deep Q-network (DQN), autoencoders, and bidirectional long short-term memory (Bi-LSTM). Each model was trained and evaluated on the dataset to determine its prediction accuracy and classification efficiency in identifying individuals with or without heart disease.

Materials And Methods

This study uses a variety of deep models to predict cardiac abnormalities using structured clinical data. The methodology involved in this involves data preprocessing, model selection, training, and then evaluation. In the first step, the Heart Disease dataset was cleaned, normalized, and split into train-testing sets for the purpose of ensuring reliable performance assessment [5]. The six deep learning architectures used were: MLP, RNN, CNN, DQN, Autoencoder, and Bi-LSTM [6]. Each model was developed to capture a different aspect of feature relationships that will improve the accuracy of classification. Models were then chosen appropriately for their training with corresponding loss functions, optimizers, and activation functions to optimize their predictability. Finally, their efficacy in heart disease detection is gauged with the help of precision, accuracy, F1-score, and recall metrics. This methodology of approach provides a comparative study about deep learning techniques for medical diagnosis and also shows the potential to be used for automated healthcare systems [7]. Figure *1* demonstrates the proposed workflow of this study.

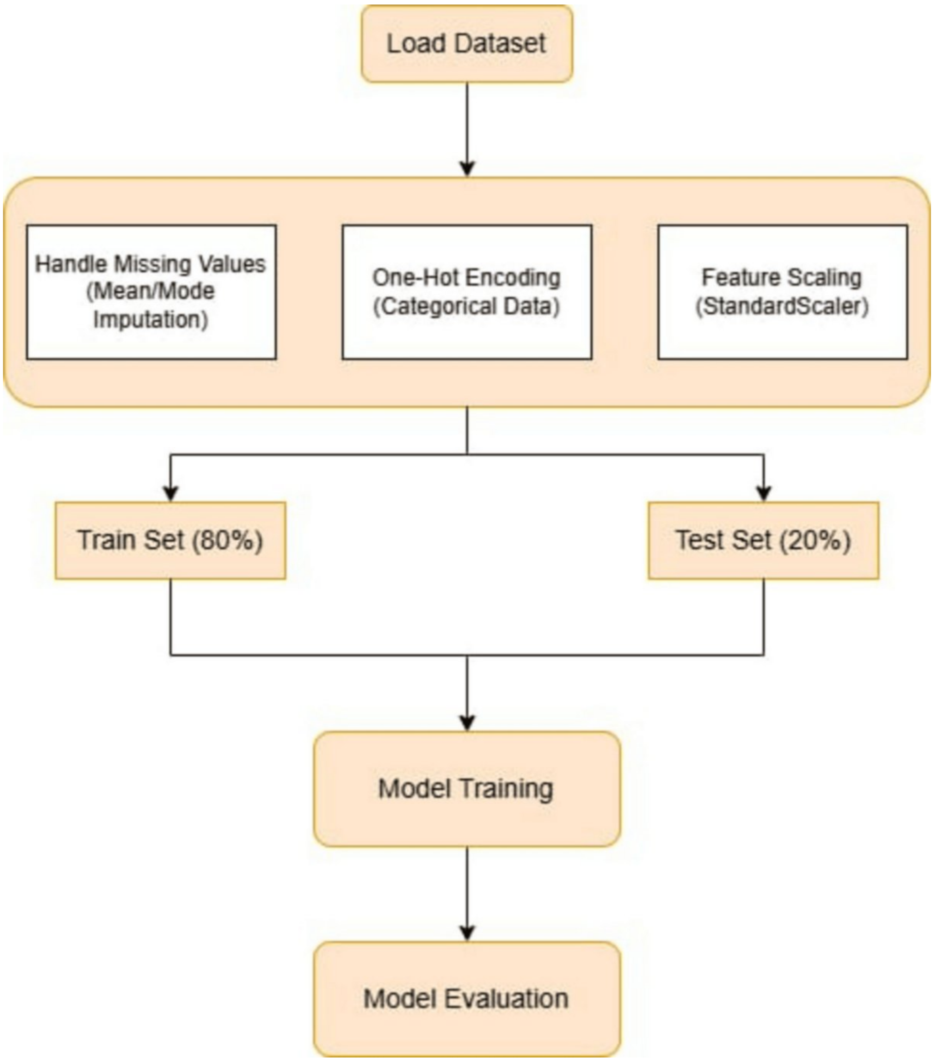

FIGURE 1: Work Flow

Dataset description and processing

For this heart disease research, a dataset from Kaggle was used [8]. It consists of 303 patient records with 14 clinical features and a binary target variable indicating the presence or absence of heart disease. The dataset includes attributes such as age, gender, chest pain type, resting blood pressure, cholesterol levels, fasting blood sugar, resting ECG outcomes, maximum heart rate, and exercise-induced angina, significant indicators for assessing cardiovascular health. Deep learning models were extensively preprocessed to ensure optimal performance. Missing values were identified and imputed using appropriate statistical techniques. Categorical variables, such as chest pain type and thalassemia, were encoded using one-hot encoding, while numerical variables were standardized using Min-Max Scaling to maintain uniform input distributions. To evaluate model performance, the dataset was split into an 80:20 ratio for training and testing purposes. Feature selection and exploratory data analysis were also conducted to identify the most influential factors in heart disease prediction. As a result, the deep learning models received clean and structured input, leading to improved accuracy and better generalization.

Preprocessing

In heart disease prediction using deep learning, preprocessing is essential to prepare the data for accurate modelling. Key steps include handling missing values, normalizing numerical features using standard scaling, and splitting the data into training and testing sets. These steps ensure the data is clean, consistent, and suitable for training deep learning models, leading to better performance and more reliable predictions.

Multi-layer perceptron

An MLP is a type of fully connected artificial neural network used for handling structured numerical data.

Its architecture consists of several hidden layers followed by a final output layer [9]. Each neuron applies an activation function to introduce non-linearity into the model. For this experiment, an MLP model was designed with three hidden layers using ReLU activation functions to enhance learning efficiency. The final output layer used a sigmoid activation function for binary classification. The Adam optimizer and binary cross-entropy loss function were employed to ensure stable and efficient model convergence during training. With a batch size of 32, the model was trained for 100 epochs, achieving an accuracy of 86.88%. The MLP model demonstrated strong performance by capturing complex feature relationships; however, it cannot process sequential information, which is why it was supplemented with RNNs and CNNs for comparative analysis.

Recurrent neural network

An RNN is designed to model sequential dependencies, making it useful for time-series and structured data analysis. In this study, an RNN-based architecture was implemented to capture interdependencies among cardiovascular health features. The RNN model consisted of a single recurrent layer followed by two fully connected layers. The tanh activation function was used to capture non-linearity, and a dropout regularization of 0.3 was applied to prevent overfitting. The model was trained using the Adam optimizer and binary cross-entropy loss, with a batch size of 32 for 100 epochs. The accuracy achieved by the RNN model was 85.24%, slightly lower than that of the MLP model. This difference may be attributed to the dataset's lack of sequential patterns, which are crucial for RNNs to perform optimally. Nonetheless, RNNs contributed an additional layer of interpretability by capturing feature dependencies across iterations, making them valuable for deeper analysis in heart disease prediction.

Convolutional neural network

A CNN, originally designed for image data, was used in this study to work with structured data for classification. CNNs were used to extract local features from the input data. This dimensionality-reduction approach included two dense layers and one convolutional layer with 64 filters combined with a max-pooling layer that flattened features, with activation from ReLU. A sigmoid function was used in the output layer as a binary activation function. Although the CNN model was quite innovative in its adaptation, it still only scored 75.40%, whereas MLP and RNN scored higher. This also indicates that CNNs may not be ideal for structured tabular data. However, the feature importance and hierarchical relationships brought insights from CNN models, which has a potential way of applying feature engineering in future studies.

Deep Q-network

Q learning with deep neural network that integrates with reinformed learning is defined as DQN. Even though DQNs were primarily designed for decision-making tasks, they were utilized in this experiment to learn optimal feature interactions to predict the heart disease prediction. This model employed a deep neural network consisting of three hidden layers and was trained through the Q-learning algorithm. DQN learned the policy that maximized the prediction accuracy through iterative updates based on reward-based learning. The DQN model reached an accuracy of 81.96%, which demonstrated the ability to identify feature dependencies in their learning process iteratively. However, DQNs require much more computational resources and require careful tuning of hyperparameters, making them not as practical compared to the supervised deep learning models like MLP and RNN.

Autoencoder

An unsupervised neural network autoencoder is used for both dimensionality reduction and anomaly detection. For this paper, an autoencoder was applied in order to reconstruct input data and then find the deviations, which relate to heart disease. The model consisted of the encoder part, which was essentially feature dimensionality reduction, and the decoder part. The bottleneck layer should have captured all the important feature representations, and hidden layers were using the ReLU activation function. While in the reconstruction error, there was an element of MSE used. Though it could extract features, the autoencoder was not performing well; it was showing only 49.18% accuracy. Thus, the autoencoder alone does not prove efficient for structured classification tasks. Yet, it also opens a pathway for further exploration in feature selection and anomaly detection.

Bidirectional long short-term memory

For the long dependencies between the features, a Bi-LSTM network was used. The input sequence was processed both forward and backward directions to make deeper relationships within structured data available for extraction. The architecture of the model is composed of one Bi-LSTM layer and two dense layers. A sigmoid function was used because the classification task is binary. To avoid the overfitting, a dropout rate of 0.3 was added. During training, Adam optimizer was used with a binary cross-entropy loss. It achieved an accuracy of up to 85.24%, not much different from the standard RNN. Providing improved feature dependencies apart, the lack of any kind of sequential patterns in the dataset limited its full potential. Nonetheless, Bi-LSTMs are incredibly valuable when dealing with time-series and longitudinal medical data.

Results

The efficiency of the deep learning model in predicting cardiac abnormalities was evaluated with the help of different evaluation metrics, such as precision, recall, F1-score, and accuracy (Table 1). These metrics provide a good understanding of how well the model is performing in prediction and how effective it is in classifying instances with and without heart disease. Accuracy, which is the proportion of correctly classified instances out of the total number of instances, is the simplest metric. It gives a general measure of how well the model is performing, but in imbalanced datasets, it may not be reliable. Thus, other metrics, such as precision and recall, were also studied for further insight. Precision, or positive predictive value, is the number of true cases of a condition (in this case, patients with heart disease) as a fraction of all predicted positive cases. A high precision score ensures that fewer false positives occur, which is critical in medical diagnoses to avoid unnecessary treatments or interventions. Recall, or sensitivity, is the proportion of actual positive cases that are correctly identified by the model. In heart disease detection, a high recall score is crucial as it minimizes the risk of false negatives, ensuring that patients receive timely treatment before their condition progresses to a more severe stage.

Model	Accuracy	Precision	Recall	F1-Score
Multi-Layer Perceptron (MLP)	86.88%	0.87	0.87	0.87
Recurrent Neural Network (RNN)	85.24%	0.85	0.85	0.85
Convolutional Neural Network (CNN)	75.40%	0.78	0.77	0.77
Deep Q-Network (DQN)	81.96%	0.85	0.85	0.85
Autoencoder	49.18%	0.42	0.46	0.39
Bidirectional Long Short-Term Memory (Bi-LSTM)	85.24%	0.85	0.85	0.85

TABLE 1: Performance Metrics of Different Deep Learning Models

For a balanced evaluation of precision and recall, the F1-score was used as an additional scoring measure. The F1-score is the average of precision and recall, providing a single score that represents the model's ability to maintain a balance between false positives and false negatives. A higher F1-score indicates better model performance, meaning the model can more effectively distinguish between diseased and non-diseased samples. Considering all these metrics, the study ensured that the evaluation of each deep learning model was robust and could be further compared to predict heart disease, taking into account their respective strengths and limitations.

Experimental setup

The proposed MLP model with three hidden layers and ReLU activation was used to enhance learning efficiency. For binary classification, the output layer utilized a sigmoid activation function. The model was trained using the Adam optimizer and binary cross-entropy loss to ensure stable and efficient convergence. Training was carried out using a batch size of 32 over 100 epochs. An RNN model featured a single recurrent layer, followed by two dense layers, with the Tanh activation function applied to capture the non-linear relationships in the data. To prevent overfitting, a dropout rate of 0.3 was included, and the Adam optimizer was chosen for its efficiency in training. The loss function used was binary cross-entropy, and the model was trained over 100 epochs using a batch size of 32. CNNs are effective at extracting local features from the input, employing a dimensionality-reduction strategy. The architecture consisted of a convolutional layer with 64 filters, followed by a max pooling layer, and two fully connected layers that flattened the extracted features using ReLU activation. The final output layer applied the sigmoid function for binary classification.

The DQN model employed a deep neural network with three hidden layers and was trained using the Q-learning algorithm. The bottleneck layer is designed to capture the essential feature representations, with the hidden layers employing the ReLU activation function. Additionally, the reconstruction error incorporates mean squared error as a key component. Model architecture includes one LSTM layer with the potential for bidirectional sequence processing and two dense layers. The output layer is activated by a sigmoid function, which is typical for models engaged in binary classification tasks. A dropout of 0.3 is applied to avoid overfitting. The optimizer is Adam, and the model trains under a binary cross-entropy loss.

Discussion

This study explored various deep learning methods for early findings of heart disease using structured medical data. The MLP model achieved the best performance with 86.88% accuracy, 0.87 precision, 0.87 recall, and 0.87 F1-score, showing its effectiveness for structured medical data. RNN and Bi-LSTM also performed well, highlighting the importance of feature relationships. In contrast, CNNs and autoencoders underperformed, suggesting they are less suited for this type of data.

Future improvements may include hybrid modelling, feature engineering, and the use of larger datasets to boost accuracy and real-world applicability. Overall, the results reinforce the value of deep learning in medical diagnostics, particularly for early detection of cardiovascular conditions. These enhancements can significantly improve the practical applicability of AI-driven tools in real-world clinical settings.

Conclusions

This work discusses a few deep learning approaches for early detection of cardiac abnormalities using the Heart Disease Dataset. The best model was the MLP model with an accuracy of 86.88%, 0.87 precision, 0.87 recall, and 0.87 F1-score, thus proving that fully connected neural networks are well suited for structured medical data. Models like RNN and Bi-LSTM also perform well, thereby suggesting that feature dependencies are highly important. CNNs and autoencoders performed the worst of all, indicating that image-based structures and unsupervised learning algorithms do not best suit structured medical data. Future opportunities into hybrid models, feature engineering, and larger datasets will improve the accuracy of predictions and real-world applicability. This study reaffirms the effectiveness of deep learning in medicine and, more importantly, highlights the need for data-driven early detection methods of cardiovascular diseases.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: P Shoba, Azha Periasamy

Acquisition, analysis, or interpretation of data: P Shoba, Azha Periasamy

Drafting of the manuscript: P Shoba, Azha Periasamy

Critical review of the manuscript for important intellectual content: P Shoba, Azha Periasamy

Supervision: P Shoba, Azha Periasamy

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

References

1. Annadate P, Bedekar M, Annadate M: Harnessing deep learning for early detection of cardiac abnormalities. International Journal of Computing and Digital Systems. 2025, 17:1-10. 10.12785/ijcds/1571031205
2. Wani MA, Bhat FA, Afzal S, Khan AI: Advances in Deep Learning. Singapore, Springer; 2020. 10.1007/978-981-13-6794-6
3. Dutta A, Batabyal T, Basu M, Acton ST: An efficient convolutional neural network for coronary heart disease prediction. Expert Systems with Applications. 2020, 159:113408. 10.1016/j.eswa.2020.113408
4. World Health Organization: Prevention of cardiovascular disease: pocket guidelines for assessment and management of cardiovascular risk: (WHO/ISH cardiovascular risk prediction charts for the African Region). World Health Organization, 2007.
5. Schmidhuber J: Deep learning in neural networks: An overview. Neural Networks. 2015, 61:85-117. 10.1016/j.neunet.2014.09.003
6. Zeng W, Shan L, Yuan C, Du S: Advancing cardiac diagnostics: Exceptional accuracy in abnormal ECG signal classification with cascading deep learning and explainability analysis. Applied Soft Computing. 2024, 165:112056. 10.1016/j.asoc.2024.112056

7. Menshawi A, Hassan MM, Allheeib N, Fortino G: A hybrid generic framework for heart problem diagnosis based on a machine learning paradigm. Sensors. 2023, 23:1392. 10.3390/s23031392
8. Heart Disease Dataset. https://www.kaggle.com/datasets/johnsmith88/heart-disease-dataset.
9. Goodfellow I, Bengio Y, Courville A: Deep Learning. MIT Press, Cambridge; 2016.