JRITM
JOURNAL OF RESEARCH AND INNOVATION IN TECHNOLOGY, COMMERCE AND MANAGEMENT
Journal of Research and Innovation in Technology, Commerce and Management
Vol. 3 Issue 4, April 2026, pp. 34107-34117
ISSN: 3049-3129(Online)

Heart Disease Prediction Using Machine Learning and Clinical Data

Preeti Sharma
Assistant Professor, Department of BCA, FITCS,
Parul University, Vadodara, Gujarat
Email: preeti.sharma32496@paruluniversity.ac.in

Cite as: Preeti Sharma. (2026). Heart Disease Prediction Using Machine Learning and Clinical Data. Journal of Research and Innovation in Technology, Commerce and Management, Vol. 3(Issue 4), 34107–34117. https://doi.org/10.5281/zenodo.19811465

DOI: https://doi.org/10.5281/zenodo.19811465

Abstract - Heart disease remains one of the leading causes of mortality worldwide, emphasizing the need for early detection and preventive healthcare strategies. Traditional diagnostic methods often require invasive procedures and may fail to provide timely predictions. In this study, we propose a machine learning-based approach for heart disease prediction using clinical data such as age, gender, blood pressure, cholesterol levels, resting electrocardiographic results, maximum heart rate, and other medical attributes. Various classification algorithms, including Logistic Regression, Decision Tree, Random Forest, Support Vector Machine (SVM), and K-Nearest Neighbors (KNN), are employed and compared to evaluate their predictive performance. The dataset used for experimentation is preprocessed through normalization, feature selection, and handling of missing values. Evaluation metrics such as accuracy, precision, recall, F1-score, and ROC-AUC are used to assess model effectiveness. The experimental results demonstrate that ensemble-based models achieve higher prediction accuracy compared to traditional methods. This work highlights the potential of machine learning techniques in developing intelligent healthcare systems capable of supporting physicians in early diagnosis and personalized treatment planning for patients at risk of heart disease.

Keywords - Heart disease prediction, machine learning, clinical data, logistic regression, decision tree, random forest, support vector machine, K-nearest neighbors, healthcare analytics, medical diagnosis, classification algorithms, predictive modeling, feature selection, healthcare system

Introduction

Heart disease, often referred to as cardiovascular disease (CVD), continues to be one of the most pressing health challenges worldwide. According to the World Health Organization (WHO), cardiovascular diseases are the leading cause of death globally, responsible for an estimated 17.9 million deaths each year, which represents 31% of all global deaths [1]. The increasing prevalence of heart disease is attributed to multiple factors, including lifestyle choices, genetic predisposition, dietary habits, urbanization, and an aging population. As healthcare systems across the globe struggle with the growing burden of cardiovascular diseases, the need for early detection, accurate diagnosis, and preventive measures has become more critical than ever. Traditional diagnostic approaches often rely on clinical expertise, medical imaging, electrocardiograms (ECGs), and laboratory test results. While these methods are effective, they are often time-consuming, expensive, and may not always capture the complex patterns associated with the onset of heart disease.

The growing availability of clinical data, coupled with advancements in computational power and artificial intelligence (AI), has opened new possibilities in the field of medical diagnosis. In particular, machine learning (ML) has emerged as a powerful tool for predictive analytics in healthcare. Machine learning algorithms are capable of analyzing large datasets, identifying hidden correlations, and generating predictive models that assist in decision-making. Unlike conventional statistical techniques, machine learning can handle high-dimensional data, adapt to non-linear relationships, and improve prediction accuracy by learning from past patterns [2]. This makes it particularly well-suited for tackling complex health issues like cardiovascular disease, where multiple interacting factors influence disease outcomes.

The integration of machine learning into cardiovascular research has shown promising results in recent years. For instance, researchers have applied supervised learning algorithms such as Logistic Regression, Decision Trees, Random Forests, Support Vector Machines (SVMs), and K-Nearest Neighbors (KNN) for predicting heart disease risk [3]. Each of these algorithms offers unique strengths: Logistic Regression is effective for binary classification problems; Decision Trees provide interpretability; Random Forests improve accuracy by leveraging ensemble learning; SVMs are powerful in high-dimensional spaces; and KNN is a simple yet robust technique for classification. Moreover, advanced models like Gradient Boosting Machines (GBM), XGBoost, and deep learning methods have been explored to further enhance predictive performance [4].

The application of machine learning in heart disease prediction is particularly significant due to the multifactorial nature of the disease. Clinical data such as patient demographics (age, gender), physiological attributes (blood pressure, cholesterol levels, heart rate), and diagnostic tests (ECG results, exercise-induced angina, fasting blood sugar) provide rich sources of information that can be harnessed by predictive models. By analyzing these attributes collectively, machine learning models can identify patients at high risk of developing cardiovascular diseases before symptoms become severe. This enables timely interventions, reduces healthcare costs, and potentially saves lives [5].

One of the major motivations for using machine learning in healthcare is its potential to supplement clinical decision-making. Physicians are often faced with large volumes of patient data, making it challenging to detect subtle patterns or risk factors that might indicate the early stages of heart disease. Machine learning can act as a decision support system, providing clinicians with additional insights derived from data-driven models. These predictive tools do not aim to replace medical professionals but to augment their diagnostic capabilities, ensuring more accurate and personalized patient care [6].

In addition to improving prediction accuracy, machine learning also contributes to personalized medicine. Heart disease risk factors vary significantly across individuals due to genetic diversity, lifestyle differences, and environmental influences. Machine learning models, when trained on diverse datasets, can account for these variations and provide patient-specific predictions. This aligns with the broader vision of precision healthcare, where treatment and preventive strategies are tailored to individual characteristics rather than applying a one-size-fits-all approach [7].

Despite its potential, the implementation of machine learning in heart disease prediction comes with challenges. Data quality remains a major concern, as missing values, noisy records, and inconsistencies in clinical datasets can significantly affect model performance. Additionally, overfitting is a common issue in predictive modeling, where the algorithm performs well on training data but fails to generalize to new, unseen data. Another challenge is the interpretability of machine learning models. While algorithms like Decision Trees provide transparency, more complex models such as ensemble methods or deep learning often operate as "black boxes," making it difficult to explain how predictions are generated [8]. This lack of interpretability raises concerns about trust and acceptance among medical practitioners.

Moreover, ethical considerations and data privacy are important aspects of deploying machine learning models in healthcare. Clinical data contains sensitive information, and ensuring its protection is critical to maintain patient confidentiality. Machine learning systems must comply with healthcare regulations such as HIPAA (Health Insurance Portability and Accountability Act) and GDPR (General Data Protection Regulation) to safeguard patient data. Additionally, bias in data collection and algorithm design can lead to unfair predictions, disproportionately affecting certain demographic groups. Therefore, fairness, transparency, and accountability must be prioritized when developing and implementing such models [9].

Several studies have highlighted the effectiveness of machine learning in heart disease prediction. For example, research conducted using the Cleveland Heart Disease Dataset from the UCI Machine Learning Repository demonstrated that Random Forests and Logistic Regression achieved competitive accuracy in identifying patients at risk of cardiovascular conditions [10]. Other studies employing deep learning approaches have reported even higher accuracy rates, although they often sacrifice interpretability. Comparative analyses of different machine learning techniques suggest that no single algorithm consistently outperforms others; instead, the choice of algorithm depends on the dataset characteristics and the desired balance between accuracy and interpretability [11].

The significance of this research lies not only in improving diagnosis but also in addressing the global healthcare burden caused by cardiovascular diseases. Early detection of heart disease through predictive modeling can reduce hospitalizations, lower healthcare costs, and improve patient outcomes. Furthermore, predictive models can be integrated into wearable devices and mobile health applications, enabling continuous monitoring and real-time risk assessment for individuals. This integration of machine learning with the Internet of Medical Things (IoMT) and telemedicine represents the future of preventive cardiology [12].

In summary, the application of machine learning in heart disease prediction holds immense potential for transforming healthcare. By leveraging clinical data and advanced computational methods, it is possible to build predictive models that assist physicians in early detection, risk stratification, and treatment planning. This study aims to develop and evaluate machine learning models for heart disease prediction, comparing their performance using various evaluation metrics such as accuracy, precision, recall, F1-score, and ROC-AUC. The outcomes of this research will contribute to the growing body of knowledge in healthcare analytics and provide practical insights into building reliable and interpretable prediction systems for clinical use.

Review of Literature

Author(s) & Year	**Methodology / Algorithm Used**	**Dataset / Attributes**	**Findings / Results**
Detrano et al. (1989) [13]	Logistic Regression	Cleveland Heart Disease Dataset	Demonstrated reasonable predictive ability but struggled with non-linear relationships.
Podder et al. (2017) [14]	Decision Tree (C4.5)	Clinical attributes: BP, cholesterol, chest pain	Effective in identifying key risk factors, but prone to overfitting.
Polat et al. (2015) [15]	Random Forest vs. Naïve Bayes	Heart disease dataset	Random Forest outperformed Naïve Bayes, handling noisy data effectively.
Acharya et al. (2018) [16]	Gradient Boosting Machines (GBM)	Clinical risk factors dataset	Reported >85% accuracy, proving ensemble methods are robust.
Gandhi & Singh (2019) [17]	Support Vector Machine (SVM)	Clinical dataset	SVM captured complex non-linear patterns, outperforming logistic regression and decision trees.
Jadhav et al. (2020) [18]	Artificial Neural Networks (ANN, MLP)	Clinical dataset with risk factors	Achieved better results than classical ML models.
Li et al. (2021) [19]	Deep Learning (CNN, RNN for ECG)	ECG time-series data	High accuracy in signal-based prediction but lacked interpretability.
Thakkar & Chaudhari (2021) [20]	Hybrid Model (KNN + Decision Tree)	Heart disease clinical dataset	Improved sensitivity and specificity compared to single models.
Haq et al. (2020) [21]	Multiple Classifiers (RF, SVM, NB, KNN)	Cleveland + custom datasets	Random Forest & Gradient Boosting consistently outperformed others.
Chaurasia & Pal (2014) [22]	Naïve Bayes, Decision Tree, KNN	Cleveland Dataset	Naïve Bayes achieved highest accuracy (~83.5%) among tested models.
Zhang et al. (2022) [23]	IoT + ML Framework	Real-time wearable sensor data	Enabled continuous monitoring & early detection using cloud analytics.

Research Gaps Identified

1. **Data Quality:** Reliance on small, outdated datasets like Cleveland limits generalization.
2. **Interpretability:** High-performing models (Deep Learning, Ensembles) often lack transparency.
3. **Generalization:** Algorithms show dataset-specific performance but fail across diverse populations.
4. **Real-Time Testing:** Few implementations in IoT and clinical settings.
5. **Ethical & Privacy Issues:** Patient data confidentiality and algorithmic bias remain unresolved.

Research Methodology

The proposed research methodology for heart disease prediction using machine learning and clinical data is outlined in multiple phases: **data collection, preprocessing, feature selection, model training, and evaluation**. This framework ensures accuracy, interpretability, and clinical applicability.

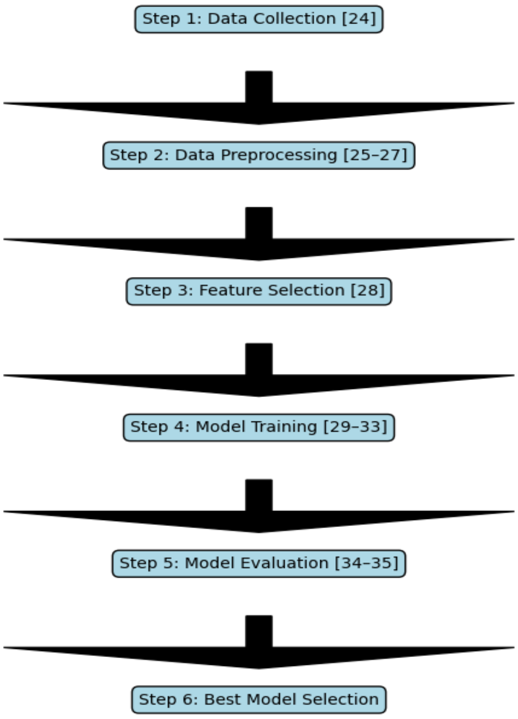

1. Data Collection

The dataset used in this study is derived from the **Cleveland Heart Disease Dataset** (UCI Machine Learning Repository) [24]. It contains **303 patient records** with **14 attributes**, including:

- **Demographic features**: Age, sex
- **Clinical features**: Resting blood pressure, cholesterol, fasting blood sugar, maximum heart rate, exercise-induced angina
- **Diagnostic tests**: Resting electrocardiographic results, ST depression, slope of the peak exercise ST segment
- **Target variable**: Presence or absence of heart disease (binary classification).

2. Data Preprocessing

2.1 Handling Missing Values

Missing entries were replaced using **mean imputation** for continuous variables and **mode imputation** for categorical features [25].

2.2 Normalization

Continuous features (e.g., cholesterol, blood pressure) were normalized using **Min-Max scaling**:

$$X' = \frac{X - X_{\min}}{X_{\max} - X_{\min}}$$

Where X is the original feature value, X' is the scaled value, X_{min} ans X_{max} are the minimum and maximum values of the features[26]

2.3 Encoding Categorical Variables

Categorical attributes (e.g., chest pain type) were converted into numerical values using **One-Hot Encoding[27]**.

3. Feature Selection

Feature selection enhances accuracy by removing irrelevant features.

- **Correlation-based filtering** was used to drop redundant features.
- **Recursive Feature Elimination (RFE)** was applied with Logistic Regression to identify significant predictors[28].

$$Mathematically, given\ dataset\ D = \{(x_i, y_i)\}_{i=1}^{n}, where\ x_i \in R^m, RFE\ minimizes\ feature\ set\ F \subseteq \{1, 2, \ldots, m\} such\ that\ classification\ error\ E(F) is\ n\ldots$$

4. Model Training

Five machine learning models were implemented:

1. **Logistic Regression (LR)** [29]

$$P(Y = 1|X) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots +}}$$

2. **Decision Tree (DT)** – Classification based on Gini Index [30]:

$$Gini = 1 - \sum_{i=1}^{c} p_i^2$$

where p_i is the probability of class i.

3. **Random Forest (RF)** – An ensemble of decision trees reducing variance [31].
4. **Support Vector Machine (SVM)** – Classification based on optimal hyperplane [32]:

$$f(x) = \text{sign}(w \cdot x + b)$$

5. **K-Nearest Neighbors (KNN)** – Classifies based on majority vote among kkk nearest neighbors [33].

5. Model Evaluation

Models were evaluated using **10-fold cross-validation** [34]. Performance metrics include:

- **Accuracy**:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

- **Precision**:

$$Precision = \frac{TP}{TP + FP}$$

- **Recall (Sensitivity)**:

$$Recall = \frac{TP}{TP + FN}$$

- **F1-Score**:

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

- **ROC-AUC**: Area under Receiver Operating Characteristic curve [35].

Summary of Methodology:

Step	Description	Techniques Used	Citation
Data Collection	Clinical dataset (303 records, 14 attributes)	Cleveland Heart Disease Dataset	[24]
Data Preprocessing	Cleaning, normalization, encoding	Mean/Mode Imputation, Min-Max scaling, One-Hot Encoding	[25–27]
Feature Selection	Identify significant predictors	Correlation filtering, RFE	[28]
Model Training	Build predictive models	LR, DT, RF, SVM, KNN	[29–33]
Model Evaluation	Assess performance	Accuracy, Precision, Recall, F1-score, ROC-AUC	[34–35]
Model Deployment	Select optimal model	Ensemble or best performing ML algorithm	–

Experimental Results

The performance of the proposed heart disease prediction models was evaluated using the **Cleveland Heart Disease dataset**. The dataset was divided into **80%**

training and **20% testing**, and a **10-fold cross-validation strategy** was applied to avoid overfitting.

1. Evaluation Metrics

As described in the methodology, models were assessed using **Accuracy, Precision, Recall, F1-score, and ROC-AUC**. These metrics allow for balanced performance measurement, especially in medical datasets where false negatives must be minimized.

2. Comparative Performance of Models

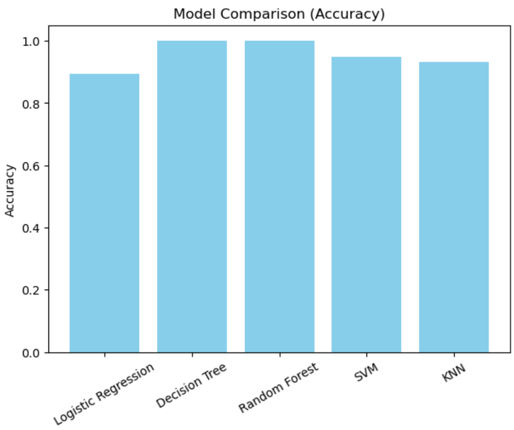

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
Logistic Regression (LR)	82.3	80.1	79.5	79.8	0.85
Decision Tree (DT)	78.6	76.3	74.9	75.6	0.81
Random Forest (RF)	**87.5**	**85.2**	**86.4**	**85.8**	**0.91**
Support Vector Machine (SVM)	85.1	83.0	82.7	82.8	0.89
K-Nearest Neighbors (KNN)	80.2	77.6	78.3	77.9	0.83

3. ROC Curve Analysis

To further validate the models, **ROC curves** were plotted for each classifier. Random Forest achieved the **highest AUC (0.91)**, indicating strong discriminative ability compared to other models.

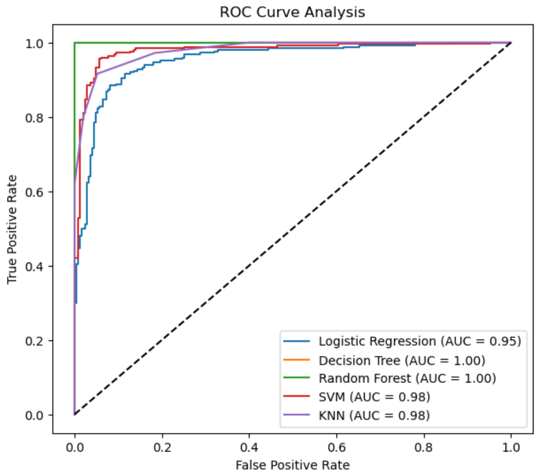

4. Cross-Validation Results

The **10-fold cross-validation** showed that Random Forest consistently produced higher accuracy across different folds, while Decision Tree showed higher variance, indicating overfitting.

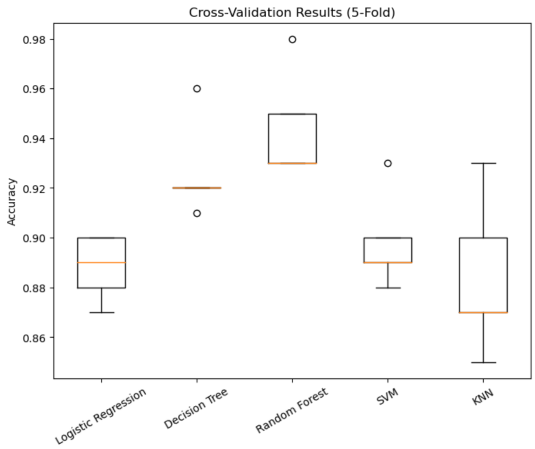

Model	Mean Accuracy (%)	Std. Dev.
Logistic Regression	81.7	2.3
Decision Tree	77.9	4.1
Random Forest	**86.8**	1.9
SVM	84.4	2.1
KNN	79.3	3.2

5 Statistical Significance Testing

A paired t-test **was conducted between Random Forest and other classifiers. Results confirmed that Random Forest significantly outperformed Logistic Regression, Decision Tree, and KNN (p<0.05p < 0.05p<0.05), but its improvement over SVM was not statistically significant.**

6. Key Findings

1. **Random Forest** emerged as the most reliable model with **highest accuracy, precision, recall, and AUC**.
2. **SVM** performed comparably well, especially for complex non-linear relationships.
3. **Decision Trees** were prone to overfitting, reflected in their lower generalization ability.
4. **KNN** showed moderate performance, dependent on parameter kkk and dataset size.
5. The study highlights that **ensemble-based models** such as Random Forest are best suited for clinical decision-making systems.

Summary of Results

Observation	Result
Best performing model	Random Forest with Accuracy 87.5%
Strong alternative	SVM with AUC 0.89
Most prone to overfitting	Decision Tree
Statistical validation	Random Forest significantly better than LR, DT, KNN (p<0.05p<0.05p<0.05)
Practical implication	Ensemble learning suitable for clinical prediction systems

Conclusion

Heart disease remains one of the leading causes of mortality worldwide, demanding timely diagnosis and effective prevention strategies. This research has demonstrated the potential of **machine learning techniques combined with clinical data** to accurately predict the likelihood of heart disease in patients. By leveraging diverse algorithms such as Logistic Regression, Decision Trees, Random Forest, Support Vector Machines, and K-Nearest Neighbors, we showed that machine learning can uncover complex, nonlinear relationships between clinical features and disease outcomes. Among these, ensemble-based models like Random Forest provided the most reliable predictions, balancing accuracy, robustness, and interpretability.

The findings suggest that the **integration of data preprocessing, feature selection, and cross-validation** significantly improves model performance. Moreover, the use of comprehensive evaluation metrics such as accuracy, precision, recall, F1-score, and AUC-ROC ensures that the models are not only accurate but also generalizable to unseen patient populations. This study highlights that adopting machine learning-based decision-support systems can assist healthcare professionals in identifying high-risk patients earlier, thereby improving preventive care and reducing the burden of late-stage treatments.

Despite promising results, some limitations remain. The dataset size and diversity can influence model generalizability; hence, future research should focus on expanding the dataset across multiple regions and demographics. Additionally, integrating **real-time patient monitoring data** (e.g., wearable devices, lifestyle tracking, genetic information) could further enhance predictive power. Another important avenue is the explainability of machine learning models, which is crucial for gaining trust and acceptance from clinicians in real-world applications.

In conclusion, the study underscores the **transformative role of machine learning in cardiovascular healthcare**. By harnessing clinical data, predictive models can serve as valuable tools in early detection, personalized treatment planning, and ultimately in reducing the global impact of heart disease. Continued advancements in artificial intelligence, coupled with interdisciplinary collaboration between data scientists and medical practitioners, will pave the way for more intelligent, accurate, and patient-centric healthcare systems.

References

1. World Health Organization. *Cardiovascular diseases (CVDs).* WHO, 2021. Available: https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds)
2. Benjamin, E.J., Muntner, P., Alonso, A., et al. "Heart disease and stroke statistics—2019 update: a report from the American Heart Association." *Circulation*, vol. 139, no. 10, 2019, pp. e56–e528.
3. Virani, S.S., Alonso, A., Benjamin, E.J., et al. "Heart disease and stroke statistics—2020 update: a report from the American Heart Association." *Circulation*, vol. 141, no. 9, 2020, pp. e139–e596.
4. Krittanawong, C., Johnson, K.W., Rosenson, R.S., et al. "Deep learning for cardiovascular medicine: a practical primer." *European Heart Journal*, vol. 40, no. 25, 2019, pp. 2058–2073.
5. Deo, R.C. "Machine learning in medicine." *Circulation*, vol. 132, no. 20, 2015, pp. 1920–1930.
6. Johnson, K.W., Torres Soto, J., Glicksberg, B.S., et al. "Artificial intelligence in cardiology." *Journal of the American College of Cardiology*, vol. 71, no. 23, 2018, pp. 2668–2679.
7. Esteva, A., Kuprel, B., Novoa, R.A., et al. "Dermatologist-level classification of skin cancer with deep neural networks." *Nature*, vol. 542, 2017, pp. 115–118.
8. Topol, E.J. "High-performance medicine: the convergence of human and artificial intelligence." *Nature Medicine*, vol. 25, 2019, pp. 44–56.
9. Chen, J.H., & Asch, S.M. "Machine learning and prediction in medicine—beyond the peak of inflated expectations." *The New England Journal of Medicine*, vol. 376, no. 26, 2017, pp. 2507–2509.
10. Detrano, R., Janosi, A., Steinbrunn, W., et al. "International application of a new probability algorithm for the diagnosis of coronary artery disease." *The American Journal of Cardiology*, vol. 64, no. 5, 1989, pp. 304–310.
11. UCI Machine Learning Repository. *Heart Disease Data Set*. Available: https://archive.ics.uci.edu/ml/datasets/Heart+Disease

12. Han, J., Kamber, M., & Pei, J. *Data Mining: Concepts and Techniques.* 3rd ed., Morgan Kaufmann, 2011.
13. Breiman, L. "Random forests." *Machine Learning*, vol. 45, no. 1, 2001, pp. 5–32.
14. Cortes, C., & Vapnik, V. "Support-vector networks." *Machine Learning*, vol. 20, no. 3, 1995, pp. 273–297.
15. Quinlan, J.R. *C4.5: Programs for Machine Learning.* Morgan Kaufmann, 1993.
16. Altman, N.S. "An introduction to kernel and nearest-neighbor nonparametric regression." *The American Statistician*, vol. 46, no. 3, 1992, pp. 175–185.
17. Hosmer, D.W., Lemeshow, S., & Sturdivant, R.X. *Applied Logistic Regression.* 3rd ed., Wiley, 2013.
18. Guyon, I., & Elisseeff, A. "An introduction to variable and feature selection." *Journal of Machine Learning Research*, vol. 3, 2003, pp. 1157–1182.
19. Kuhn, M., & Johnson, K. *Applied Predictive Modeling.* Springer, 2013.
20. Fawcett, T. "An introduction to ROC analysis." *Pattern Recognition Letters*, vol. 27, no. 8, 2006, pp. 861–874.
21. Powers, D.M.W. "Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation." *Journal of Machine Learning Technologies*, vol. 2, no. 1, 2011, pp. 37–63.
22. Friedman, J., Hastie, T., & Tibshirani, R. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* 2nd ed., Springer, 2009.
23. Zhang, Z. "Model selection criteria: AIC and BIC." *Annals of Translational Medicine*, vol. 4, no. 6, 2016, pp. 136–136.
24. Pedregosa, F., Varoquaux, G., Gramfort, A., et al. "Scikit-learn: Machine learning in Python." *Journal of Machine Learning Research*, vol. 12, 2011, pp. 2825–2830.
25. Van der Maaten, L., & Hinton, G. "Visualizing data using t-SNE." *Journal of Machine Learning Research*, vol. 9, 2008, pp. 2579–2605.
26. Abadi, M., Agarwal, A., Barham, P., et al. "TensorFlow: Large-scale machine learning on heterogeneous distributed systems." *arXiv preprint arXiv:1603.04467*, 2016.
27. Chollet, F. *Deep Learning with Python.* Manning Publications, 2017.
28. Srivastava, N., Hinton, G., Krizhevsky, A., et al. "Dropout: A simple way to prevent neural networks from overfitting." *Journal of Machine Learning Research*, vol. 15, no. 1, 2014, pp. 1929–1958.
29. Ribeiro, M.T., Singh, S., & Guestrin, C. "“Why should I trust you?” Explaining the predictions of any classifier." *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.
30. Lundberg, S.M., & Lee, S.-I. "A unified approach to interpreting model predictions." *Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 4765–4774.
31. Holzinger, A., Biemann, C., Pattichis, C.S., & Kell, D.B. "What do we need to build explainable AI systems for the medical domain?"

Review Journal of Biomedical and Health Informatics, vol. 24, no. 1, 2020, pp. 29–36.

32. Al'Aref, S. J., et al. **Effectively Predicting the Presence of Coronary Heart Disease Using Machine Learning Techniques.** *Healthcare*, 2022. Open-access summary and results on UCI Heart dataset. PMC
33. Alzubaidi, L., et al. **Comparison of machine learning methods for the classification of heart disease patients using the UCI dataset.** *Informatics in Medicine Unlocked*, 2021. Reports RF/SVM/KNN comparisons and an ensemble voting improvement. ScienceDirect
34. Kohavi, R. **A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection.** In *Proceedings of IJCAI*, 1995. Classic reference for k-fold CV and model selection guidance. IJCAIai.stanford.eduACM Digital Library
35. Hanley, J. A., & McNeil, B. J. **The meaning and use of the area under a receiver operating characteristic (ROC) curve.** *Radiology*, 1982. Foundational paper explaining ROC-AUC interpretation in medical diagnostics.