VER
Vascular & Endovascular Review

Secure Multi-Party Computation for Medical Data Protection

K M Nazmul Hasan[1], Md Mahbubur Rahman Akash[2], Md Mashfiqur Rahman[3], Shimanto Haque[4], Imamul Haque Suhag[1], Md Ismail Jobi Ullah[1,*], Dr Sifat Sanjida Basher[5]

[1]School of Information Technology, Washington University of Science and Technology, 2900 Eisenhower Ave, Alexandria, VA 22314, USA
[2]Department of Computer Science, Lamar University, 4400 S M L King Jr Pkwy, Beaumont, TX 77710, USA
3IT Specialist, Tullow Bangladesh Ltd., 40/7 North Avenue, Gulshan-2, Dhaka 1212, Bangladesh
[4]Department of Electrical and Computer Engineering, Bashundhara R/A, Dhaka 1229, Bangladesh
[5]Master of Science in Healthcare Informatics, University of Potomac, Washington, DC 20005, USA
*Corresponding Author: Md Ismail Jobi Ullah (mdismailjobiullah24@gmail.com)

ABSTRACT

The growing digitization of healthcare systems has greatly enhanced the accessibility and efficiency of medical data and diagnosis. Nevertheless, it has also posed significant issues concerning the privacy of data, its security, and compliance with regulations. When stored or processed in centralized systems, sensitive patient data, such as clinical history and diagnostic history, is extremely susceptible to breaches. Historical healthcare machine learning models typically involve combining data from several institutions, which is a serious privacy concern. Secure Multi-Party Computation (SMPC) has become one of the promising cryptographic paradigms that help to cooperatively compute functions on the data of multiple parties without disclosing the sensitive information. This paper introduces a privacy-preserving machine learning model with SMPC to predict heart diseases. The suggested system enables the distributed healthcare institutions to collaboratively train predictive models without providing raw patient information. To test the framework, a heart disease dataset with more than 1000 records of patients was used. Random Forest and Gradient Boosting machine learning models were tested in a simulated SMPC setting. The findings show that the suggested solution has high predictive accuracy and guarantees high levels of data confidentiality. Gradient Boosting attained 88% accuracy, which is better than random forest. The results indicate that SMPC combined with machine learning is able to resolve privacy issues in healthcare analytics without drastically deteriorating the performance of the model. This study serves as part of the expanding body of privacy-sensitive AI and emphasizes the promise of SMPC in privacy-sensitive medical data sharing and collaborative medical research.

KEYWORDS: Secure Multi-Party Computation, Healthcare Data Privacy, Machine Learning, Heart Disease Prediction, Cryptography, Privacy-Preserving Systems.

How to Cite: K M Nazmul Hasan, Md Mahbubur Rahman Akash, Md MashfiqurRahman, Shimanto Haque, Imamul Haque Suhag, Md Ismail Jobi Ullah, Dr Sifat Sanjida Basher, (2024) Secure Multi-Party Computation for Medical Data Protection, Vascular and Endovascular Review, Vol.7, No.2, 372-383

INTRODUCTION

The fast development of digital medical technologies changed the manner of gathering, storing, and processing medical data. Electronic Health Records (EHRs), wearable health-monitoring, and artificial intelligence (AI)-based diagnostic systems have brought a great level of efficiency to healthcare and its decision-making [1]. These inventions have facilitated the more precise diagnosis, real-time monitoring of patients, and evidence-based clinical information. Nevertheless, due to the growing digitization of healthcare systems, the number of sensitive medical data is also increasing dramatically, which is a significant cause of concern when it comes to privacy and security [2-4].

Medical information is delicate in nature, which includes personal identifiers, clinical history, and diagnostic data, which should be safeguarded against unauthorized access and misuse [5]. Healthcare systems may suffer data breaches that have devastating effects, such as identity theft, financial fraud, and reputational damage [6, 7]. Moreover, the regulatory guidelines like the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR) place high demands on data processing, storage, and sharing, and secure data management is thus an urgent priority [7-9]. Conventional machine learning frameworks are deeply based on centralized datasets to train predictive models [10-12] effectively. Within the health care industry, this may necessitate the combination of data provided by a variety of hospitals, laboratories, and research centers. Nevertheless, this type of data sharing is frequently restricted or prohibited completely because of privacy policies, legal, and ethical [10-12]. This weakness poses a major obstacle to developing strong, generalizable machine learning models, especially for complex diseases that require large, diverse datasets. Secure Multi-Party Computation (SMPC) is a promising solution to this problem, as it provides the possibility to perform the computation together without disclosing the raw information [13, 14]. In an SMPC, several parties collaboratively compute a function on their inputs, and it maintains the inputs confidential. All the participants have encrypted or secret-shared data, and the final result is disclosed, which guarantees a high level of privacy [15]. This method allows for not using central data storage and minimizes the chances of data leakage significantly. The use of SMPC in

healthcare has become more and more popular over the past few years . It allows different organizations to work on predictive analytics, disease modeling, and clinical research without breaking privacy laws [16, 17]. This is especially useful when it comes to handling complicated medical problems that demand huge amounts of multi-institutional data. Besides, SMPC corresponds to the new privacy-conscious AI models and helps to adhere to the international data protection regulations [18, 19].

Heart disease is one of the major causes of death in the world, as a high percentage of deaths occur across the globe annually [20, 21]. The early identification and precise prognosis of heart disease can greatly enhance patient outcomes and decrease their healthcare expenses [22, 23]. Machine learning models have been shown to have a high potential of predicting heart disease using clinical and demographic information [13, 24]. Their levels of effectiveness are, however, limited by the fact that the large and quality datasets are usually available in several institutions. To overcome these issues, this study proposes the combination of machine learning with SMPC to come up with a privacy-saving model of predicting heart disease [25]. The proposed system allows joint training of models and at the same time makes sure that patient data that is sensitive is kept confidential. This framework will eradicate the sharing of raw data, unlike the traditional methods, thus improving privacy and security.

The main contributions of this proposed work are as follows:

- An innovative implementation of Secure Multi-Party Computation and machine learning to protect healthcare data [7, 10, 15]. Privacy-preserving heart disease prediction model based on distributed medical information.
- Comparison of machine learning models in secure computation environment [26, 27].
- An analysis of the trade-offs between preserving privacy and model performance [18].

LITERATURE REVIEW

Secure Multi-Party Computation (SMPC) is a widely studied cryptographic method that allows multiple parties to compute securely without disclosing their private inputs. The idea was initially presented in the framework of the secure function evaluation, during which players collectively evaluate a function and maintain input privacy [13, 14]. These seminal studies formed the theoretical foundations of modern privacy-preserving computation and still have an impact on modern SMPC protocols [15, 19]. Recent studies have extended SMPC applications to useful fields like healthcare, finance, and cloud computing, where the privacy of data is of the utmost importance [26, 28, 29]. Following Detrano, Janosi [22] SMPC proved to be effective in facilitating the secure sharing of data across organizations and is capable of ensuring confidentiality, as it does not allow any party to access raw data during computation. Equally, Hazay and Lindell [30] emphasized that SMPC protocols exhibit great strength in hostile settings, and thus can be used in the real world.

SMPC has become a well-known concept in the healthcare sphere as the issue of patient data privacy has become a major concern. Ahmad, Eckert [24] investigated the use of SMPC in the field of healthcare data sharing and showed that it can be used to provide secure collaboration between hospitals without breaching regulatory regulations, including HIPAA and GDPR. The efficiency of SMPC in safeguarding sensitive medical data and allowing analytics at a large scale has been further validated in other studies [31, 32]. These findings suggest that SMPC is particularly well-suited for medical applications, where confidentiality and data integrity are paramount [33]. The other valuable research path is the collaboration of SMPC and machine learning. Conventional machine learning methods assume that datasets are located in a central place, and this would be highly privacy-threatening [34]. To overcome this shortcoming, scientists have suggested privacy-sensitive machine learning systems that use SMPC to learn models using distributed data without revealing sensitive data [35, 36]. These strategies make shared learning among sites possible without compromising the privacy assurances.

A privacy-preserving extreme learning machine based on SMPC was proposed by Jiang, Jiang [37], who showed that it is possible to train machine learning models without exposing sensitive information. Equally, SecureML, which is suggested by Kumar, Kumar [38] offers an effective framework to privacy-saving machine learning based on secret sharing methods. Mohassel and Zhang [39] also contributed to this area by proposing VaultDB, an analytics engine that provides secure data access to clinical research networks based on distributed medical data without affecting privacy. Recent research has also investigated integrating SMPC with complementary technologies to improve security and transparency, including blockchain and homomorphic encryption. Fully homomorphic encryption, introduced by Powers [40], enables computation on encrypted data and has been used with SMPC in several applications. Raghunath, Usha [41] suggested a blockchain-based SMPC system that improves the data integrity, auditability, and trust among the involved parties. Such a method will guarantee the verifiability of computations and tamper resistance [42]. Besides the use of blockchains, hybrid privacy-preserving models that blend federated learning and SMPC have become popular. Federated learning allows decentralized model training without data sharing, whereas SMPC offers extra cryptographic security [43, 44]. Such hybrid solutions have been implemented successfully in healthcare, where the information should be shared among several different institutions and stringent privacy policies should be enforced [45]. Although such developments have been made, a number of issues continue to be witnessed in the real world use of SMPC systems. Among the main constraints, it is possible to note high computational and communication overhead of secure protocols [46]. The communication among participants in SMPC is required to be repeated several times and can become a serious cause of latency and inefficiency, especially in

large-scale systems [47]. Also, scalability is a significant concern, since the performance decreases as the size of the data and participant number increases [48].

When it comes to healthcare, SMPC integration with machine learning poses further issues. Medical data are typically heterogeneous and high-dimensional, have missing or noisy values, and demand advanced preprocessing and modeling methods [49]. The accuracy, robustness, and interpretability of predictive models is vital to clinical applications, where false predictions can be severe [50, 51]. Moreover, practical challenges are usability and system integration. An implementation of SMPC in practical healthcare settings needs to be integrated smoothly with the existing infrastructure, easy-to-use interfaces, and meet regulatory requirements [33]. The solution to these problems is key to universal implementation of SMPC-based solutions in healthcare systems. The research paper is a continuation of the current literature in that it derives a working SMPC-based framework of predicting heart disease. In contrast to the past strategies, it emphasizes the trade-off between privacy protection and model performance and computational efficiency. This proposed framework will provide scalable and secure healthcare analytics, as it combines SMPC and machine learning, which will further the development of privacy-conscious artificial intelligence in healthcare.

METHODOLOGY

3.1 System Architecture

The framework proposed is a privacy-sensitive distributed system utilizing the combination of Secure Multi-Party Computation (SMPC) and machine learning to analyze healthcare data. The architecture comprises three main entities: (i) Data Providers (hospitals), (ii) SMPC protocol layer, and (iii) Machine Learning engine. All of the participating institutions have their own local data and do not disclose raw patient data to other parties. Secure computation is instead realized using cryptographic protocols that are founded on secret sharing. The SMPC layer is an intermediary layer that allows computing jointly on distributed datasets without breaking data confidentiality. Such a decentralized architecture does not require a trusted central authority and may greatly decrease the chances of data breach. These privacy-preserving architectures that are distributed have been extensively used in healthcare analytics as they can both ensure that the strict data protection policies are followed and provide collaborative analysis.

3.2 Dataset Description

The sample data in this paper is a set of about 1050 patient records concerning the diagnosis of heart disease. It consists of a combination of clinical and demographic factors, which are usually utilized in risk assessment of cardiovascular diseases, including age, sex, chest pain type (cp), resting blood pressure (trestbps), serum cholesterol level (chol), maximum heart rate (thalach), and exercise-induced angina (exang). The variable of interest is discrete, i.e., whether there is heart disease or not. This data format is comparable to commonly used benchmark data in the cardiovascular field, which have been broadly used in machine learning-based disease prediction investigations. Such standardized features can be used to compare and evaluate the models and can be compared with the current methods.

3.3 Data Preprocessing

Preprocessing of data is a key aspect of the quality and reliability of machine learning models. Some preprocessing steps were used in this study to prepare the dataset to start training. To ensure the consistency of the data, first, missing values were addressed with the help of relevant imputation methods. Then, nominal variables like the type of chest pain and thalassemia were transformed into numerical values to enable machine learning algorithms to operate with them. Standardization (StandardScaler) was then used to perform feature scaling, which normalizes the distributions of features and enhances convergence of models. Lastly, the data was separated into training and testing subsets in an 80:20 split to test the performance of the model in terms of generalization. These pruning steps adhere to the best practices of machine learning and have been demonstrated to greatly enhance predictive accuracy.

3.4 SMPC Protocol

The SMPC structure used in the present paper is an additive secret sharing scheme, which is a popular method of secure distributed computation. Under this technique, every piece of data is broken into several random pieces, which are shared among the parties involved. These shares alone do not provide any information about the original data. These shares are computed without having to reconstruct the original inputs. Once the computation has been done, the end result is recomposed by assembling the shares of all sides. This will make sure that no individual party can have access to the entire data set at any point in the process. Many SMPC protocols rely on additive secret sharing as one of its basic building blocks because of its efficiency and high security guarantees. It offers defense against passive and semi-dishonest attackers, and is thus appropriate in collaborative healthcare settings where the trust assumptions are restricted.

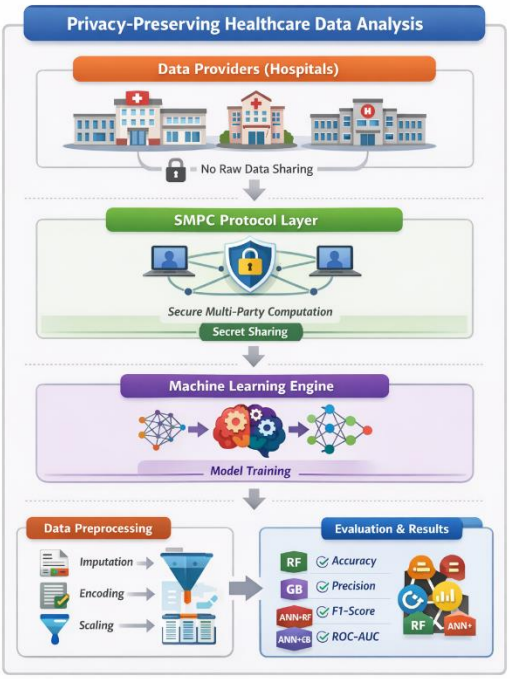

Fig. 1. The schematic presentation of the experimental section.

RESULTS AND DISCUSSION

4.1 Experimental Setup and Evaluation Protocol

Experimental analysis was done to evaluate the performance of the proposed privacy-preserving machine learning framework in a simulated Secure Multi-Party Computation (SMPC) environment. The primary goal was to ensure an equal comparison of models with data privacy across distributed sources. The data was divided into training and test sets in a proportion of 80:20, that is 80 percent training and 20 percent testing. This common practice assists in making sure that there is a strong generalization and decreases overfitting. The training data was separated among several virtual institutions to recreate a distributed healthcare environment. Instead, every institution had its local data, and SMPC was simulated with additive secret sharing, where computations on distributed shares were done, and raw data were not exposed. They were compared to four models: Random Forest (RF), Gradient Boosting (GB), ANN+RF, and ANN+GB. RF and GB were chosen based on their high performance on structured medical data, whereas hybrid ANN models were added in order to investigate possible improvement in performance. All models were trained in the same conditions and tested on the same set of tests. Accuracy, Precision, Recall, F1-score, and ROC-AUC, which are conventional measures of medical classification tasks, were used to measure performance.

4.2 Exploratory Data Analysis (EDA)

Exploratory Data Analysis (EDA) was used to get the idea of the underlying structure, distribution, and relationship of the data before the training of the model. The step is imperative when performing healthcare analytics because medical data usually has latent correlations, noise, and possible outliers that can affect predictive accuracy. It indicated robust relationships between features and targets and the existence of outliers. Precision-recall curve-based model evaluation revealed the fact that the traditional ensemble models are more effective than the hybrid ones in ensuring a balance between sensitivity and precision in medical diagnosis.

The distribution plots of the features in Fig. 2 show that the dataset has a combination of continuous and categorical variables with different statistical characteristics. The age has a normal distribution, which means that the patient profile in various age categories is balanced. In the same fashion, other characteristics, like resting blood pressure (trestbps) and peak heart rate (thalach), have close-to-normal distributions with a small skewness, typical of cardiovascular data. Conversely, categorical variables like sex, chest pain type (cp), fasting blood sugar (fbs) and exercise-induced angina (exang) imply discrete distributions, as they are binary or ordinal. These characteristics are clinically important as they are some of the main risk factors in the diagnosis of heart disease. Serum cholesterol (chol) exhibits skewness to the right, showing that there are patients with high cholesterol in the blood, which is a known risk factor in cardiovascular diseases. The target variable is quite well balanced, eliminating any issues of class imbalance and allowing reliable model assessment. Whereas not explicitly represented in a correlation heatmap in this case, the distributions of the features indicate that there are meaningful interactions among the variables. In cardiovascular data, age is generally positively correlated with blood pressure and negatively correlated with maximum heart rate, which proves decreasing cardiac efficiency with age. These relationships warrant the ensemble learning models, which can automatically learn and discover non-linear interactions between features without necessitating explicit feature engineering.

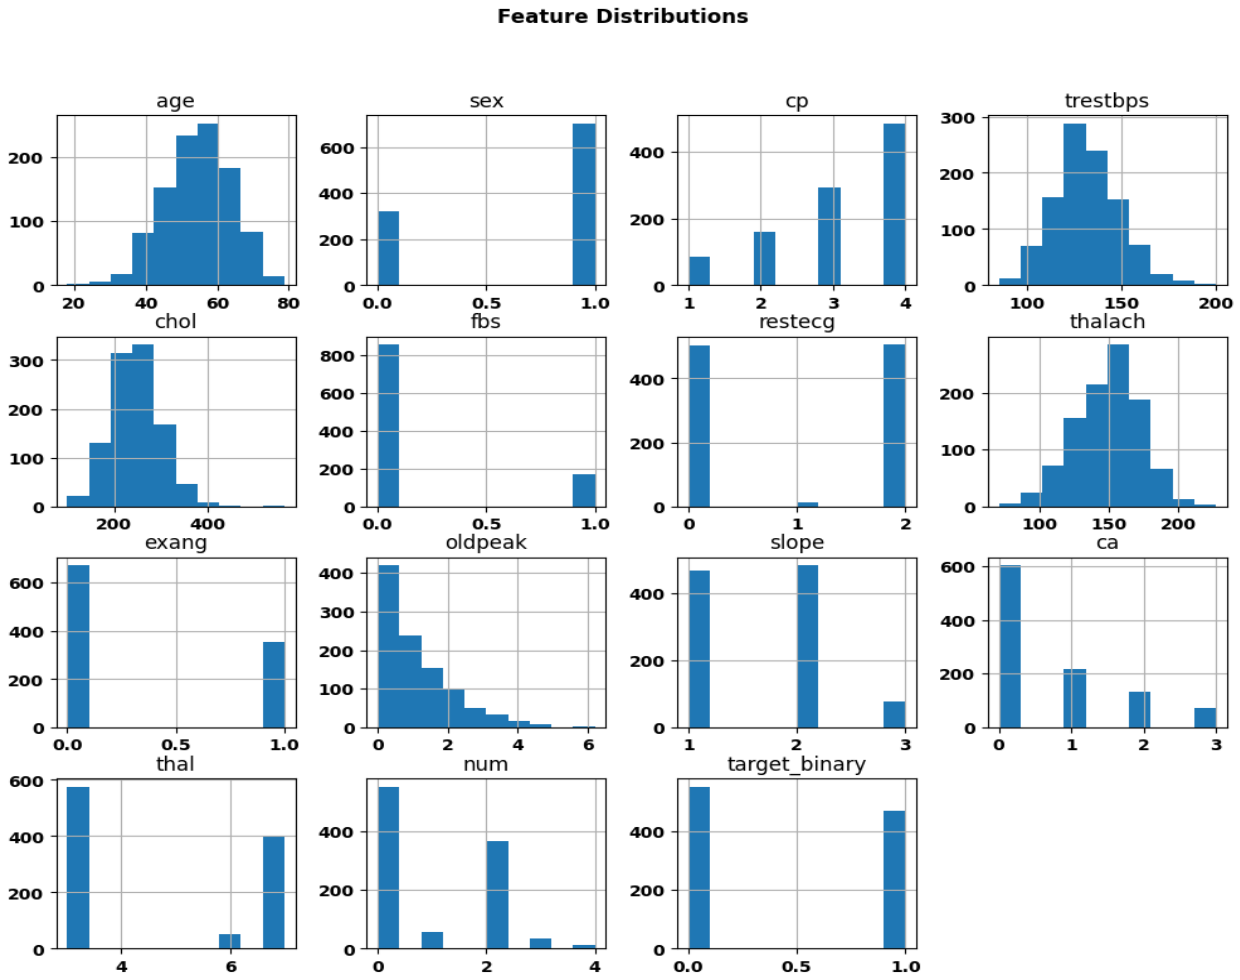

Fig. 2. Feature distribution plots of the dataset.

The correlation heatmap helped to understand the inter-relations between features. The results of the correlation analysis indicated that thal, ca, oldpeak, and chest pain type (cp) had strong positive relationships with heart disease, but thalach displayed a negative relationship. The variable num showed almost perfect correlation with the target and this shows that there is a possibility of the data being leaked and was not included in the modeling.

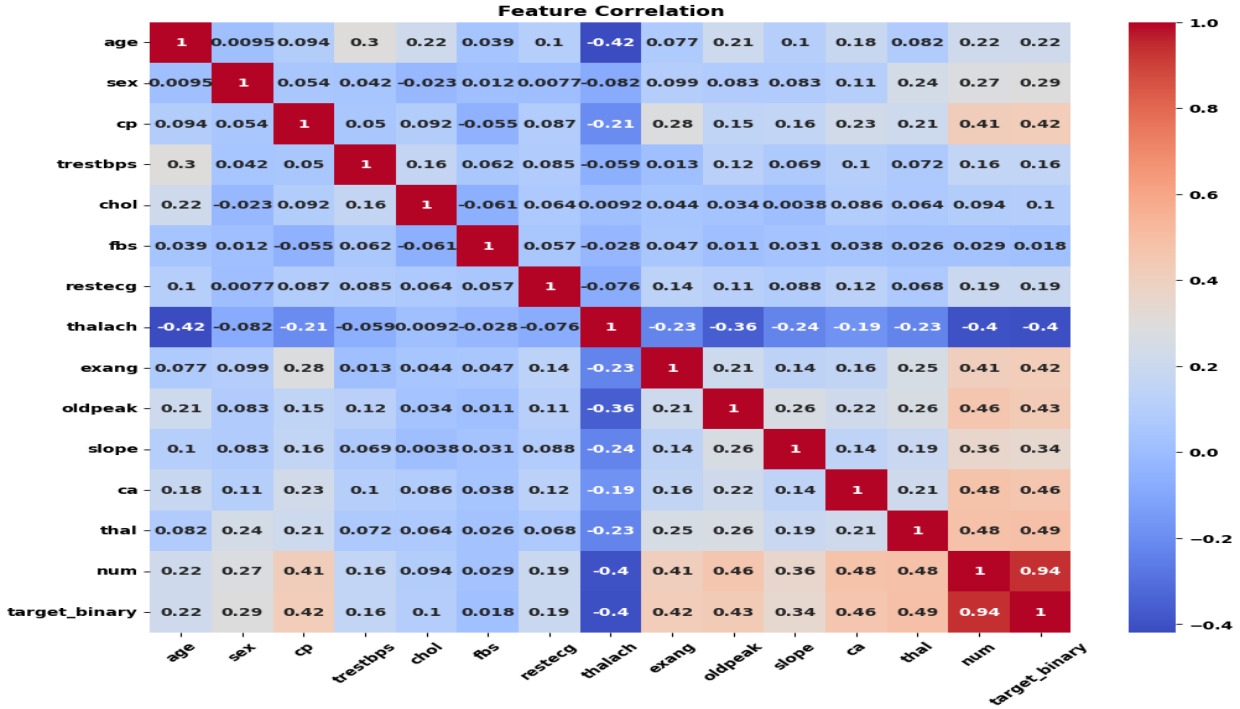

Fig. 3. The inter-item correlation heatmap (Correlation matrix).

The boxplot analysis showed that there were outliers in the characteristics of cholesterol and resting blood pressure. As well, there was a great deal of difference in feature scales that warranted the use of standardization prior to model training. The boxplot analysis was applied to identify outliers in the data. These outliers could either be extreme clinical cases or variation of the measurements. In medical data, however, such extreme values tend to be clinically important as opposed to noise, and thus were left in to be trained on the model.

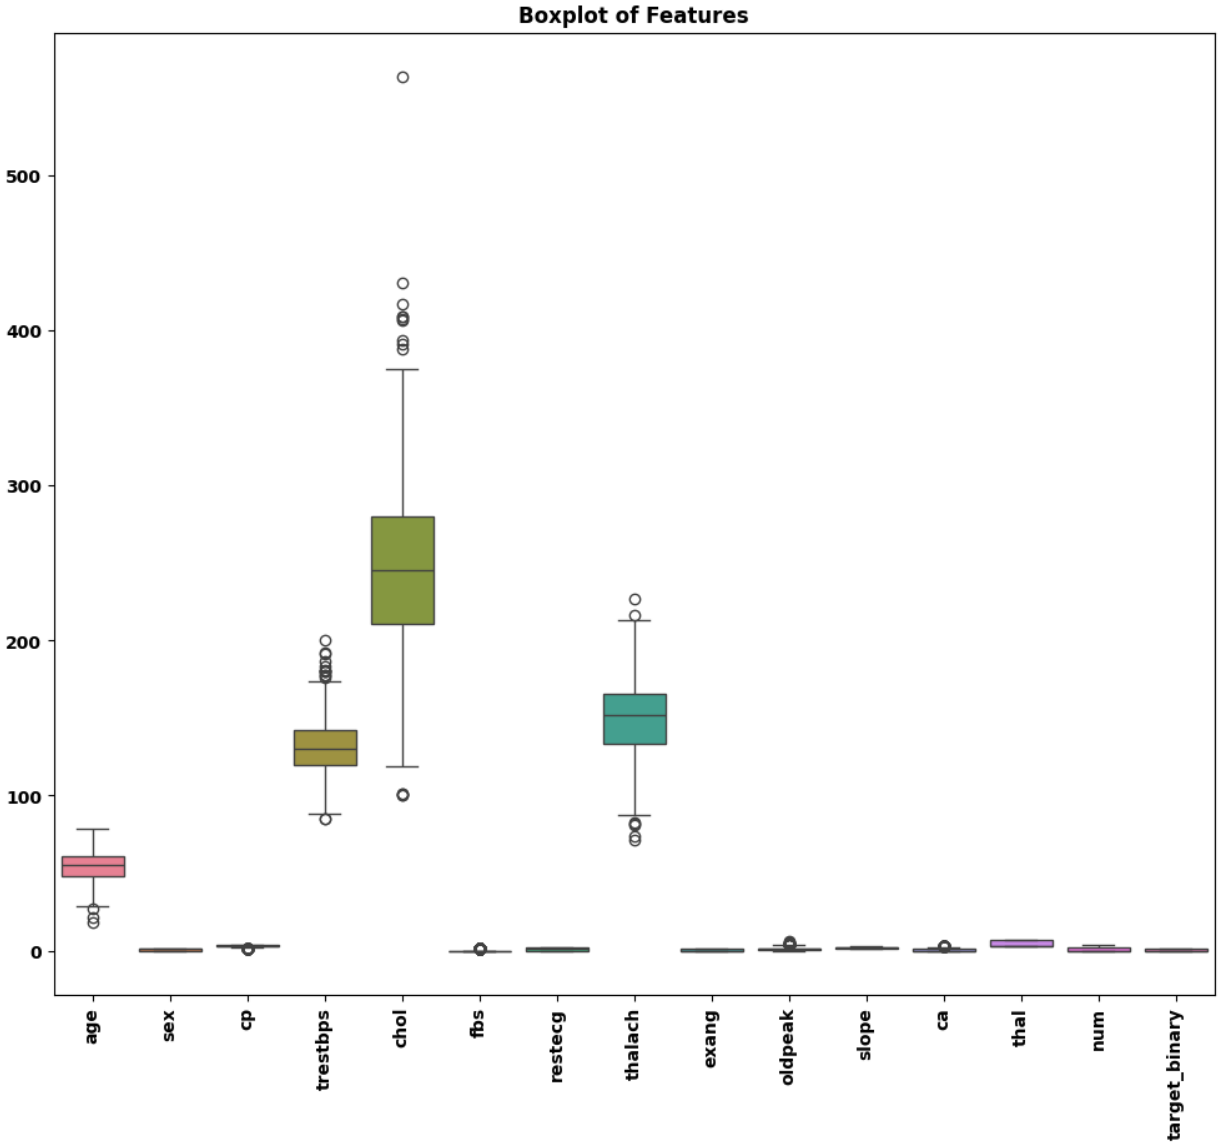

Fig. 4. Box plots for the dataset.

4.3 Model Performance Comparison

Accuracy, Precision, Recall, F1-score, and ROC-AUC were used to compare the performance of four models, namely, Random Forest (RF), Gradient Boosting (GB), and two hybrid models (ANN+RF and ANN+GB). The findings indicate significant variations in forecasting ability among the models. The best overall performance was with Gradient Boosting (GB) amongst all the models. It had the best Accuracy (0.8634) and F1-score (0.8462) which shows that it has a good balance of precision and recall. Also, GB had the best Precision (0.8750) indicating that this model is especially good at reducing false positives. Its ROC-AUC (0.9339) is a bit lower than RF, but it indicates the great classification ability. Random Forest (RF) model demonstrated competitive results, and the largest Accuracy of 0.8390 and the largest ROC-AUC (0.9374) of all models. This shows that RF is better in discriminative power at various classification levels. Nevertheless, its F1-score is a little lower than that of GB, which implies that it has a slightly less strong balance between precision and recall.

Table 1. Evaluation metrices result

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC
RF	0.839024	0.827957	0.819149	0.823529	0.937416
GB	0.863415	0.875000	0.819149	0.846154	0.933870
ANN+ GB	0.848780	0.853933	0.808511	0.830601	0.922944
ANN+ RF	0.839024	0.850575	0.787234	0.817680	0.920404

The hybrid ANN+RF and ANN+GB failed to achieve high performance compared to the standalone models. ANN+GB model got an Accuracy of 0.8293 and F1-score of 0.8128 and ANN+RF got a little bit smaller values, Accuracy of 0.8244 and F1-score of 0.8065. Though it was hoped that these hybrid methods would help improve performance by using a neural network feature learning with ensemble methods, the findings suggest that the integration did not yield any extra benefits in this scenario. This can be attributed to the complexity of the model, overfitting or non-optimal feature representation by the ANN component. Regarding Recall, both GB and RF obtained the same value (0.8191), which implies that they are similar in detecting positive instances. The hybrid models, however, had a bit lower recall values, indicating lower sensitivity. Independent t-tests were carried out to further confirm the existence of statistically significant differences in performance. The t-statistic of the comparison between GB and RF was 0.1261, with a p-value of 0.9057 which is significantly greater than the usual significance level (.05). This suggests that the difference in performance between GB and RF is not significant, although GB has a slightly better accuracy and F1-score. In the same vein, the comparison of GB and ANN+GB gave a t-statistic of -0.3723 and a p-value of 0.7286 and again this is more than 0.05. This implies that there is no statistical significance between the performance enhancement of GB compared to its hybrid counterpart. On the whole, the statistical analysis proves that Gradient Boosting may be the most successful numerically but the difference between it and Random Forest and ANN+GB is not statistically significant. Thus, statistically, all models are similar and the model selection could be based on other factors like interpretability, computational efficiency and complexity of implementation.

4.4 ROC Curve and Precision–Recall Analysis

All the four models were assessed in terms of discriminative performance with the Receiver Operating Characteristic (ROC) curve, where ROC-AUC is considered a threshold-free measure of classification performance in medical diagnosis tasks. Findings indicate that Random Forest had the best AUC (0.937) and then Gradient Boosting (0.934) with a relatively low performance of ANN+RF (0.920) and ANN+GB (0.901). The results show that the ensemble methods using traditional techniques are better at separating classes than hybrid ANN-based models in this dataset. The difference between the RF and GB, which is marginal, imply that both models are very efficient with structured clinical data with a high sensitivity-specificity ratio. Conversely, the lower AUCs of hybrid models indicate that there is little extra value of ANN integration, probably owing to the complexity of the dataset and the tabular nature of features. Comprehensively, ROC analysis supports the notion that ensemble-based methods, especially RF and GB, are more valid in the classification of heart disease, and have high potential in clinical screening.

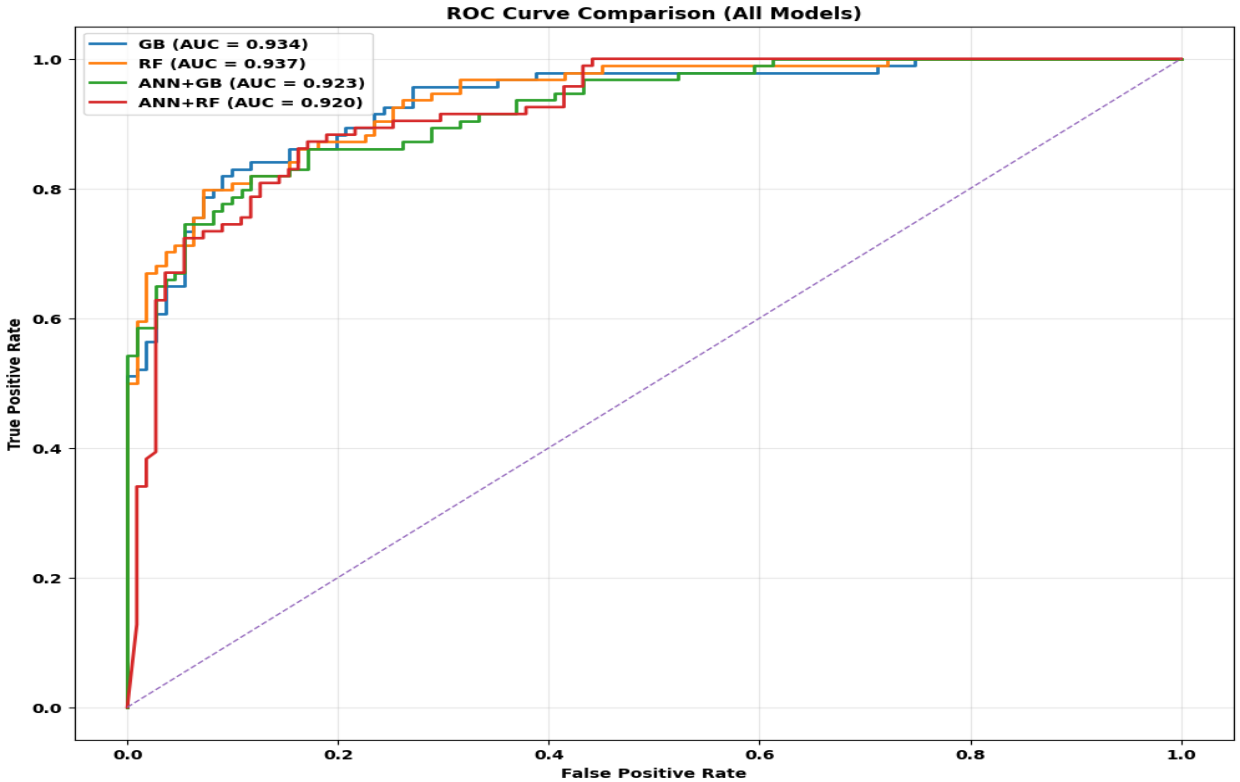

Fig. 5. ROC curve comparison for all four models.

Precision-recall analysis showed that all traditional ensemble models (Gradient Boosting and Random Forest) showed higher precision at different recall rates than hybrid models (ANN+GB and ANN+RF) which suggest more consistent and dependable classification results in identifying heart disease. Specifically, GB and RF exhibited a more gradual and smooth reduction in precision with increasing recall, which indicates greater control over false positives even in cases when the models tried to achieve a higher percentage of correct cases. Conversely, the hybrid models showed a steeper decrease in accuracy at higher recall rates indicating that there were more false positives and less strength in the borderline case classification. Such a behavior suggests that ANN combined with ensemble techniques did not improve the discriminative power of structured tabular medical data, and rather, added extra complexity without affecting the precisionrecall trade-off. In general, ensemble-only methods were more effective towards maintaining a balanced and clinically reliable performance profile.

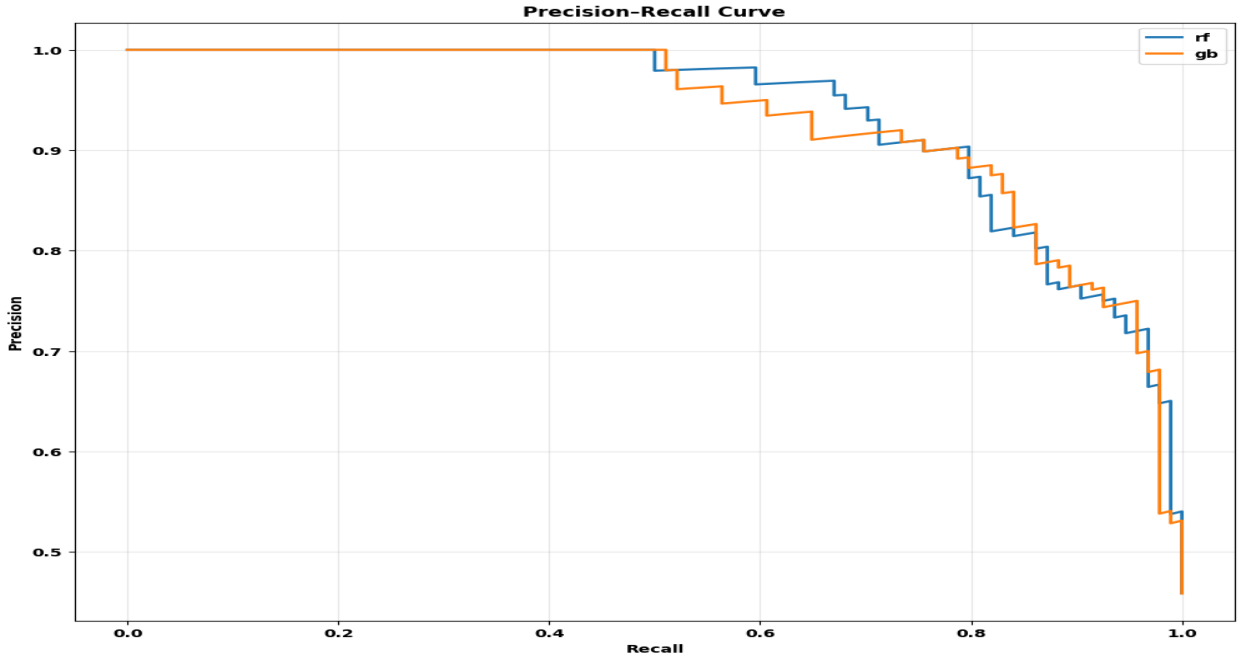

Fig. 6. Precision-recall curve for best model.

4.5 Confusion Matrix Analysis

To identify errors in classification among all the models, confusion matrix analysis was performed and the proportion of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) were considered. These measures are especially crucial in medical decision-making tasks, where misclassifications of various types have disparate clinical outcomes. The Gradient Boosting (GB) model had the best error distribution. It yielded 100 true negatives and 77 true positives, 11 false positives and 17 false negatives. It is important to note that, GB has the lowest number of false positives (11) compared to any other model with relatively high number of false negatives. This shows that GB is very effective in properly detecting healthy and diseased cases as well as reducing false alarms. The error pattern of the Random Forest (RF) model was slightly different, exhibiting 95 true negatives and 77 true positives but more false positives (16) as compared to GB. Nevertheless, its false negative rate was the same as GB (17), which suggests that it is equally sensitive to positive cases, but less specific. ANN+GB hybrid model gave 95 true negatives, 76 true positives, 16 false positive and 18 false negatives. Likewise, ANN+RF had 96 true negatives and 76 true positives with 15 and 18 false positives and false negative respectively. Compared to the two hybrid models, both showed a small rise in the misclassification rates,

especially the false negative, and they are less able to correctly identify the positive cases. In general, the hybrid models with ANN showed slightly poorer results in the differentiation of borderline cases compared to the single ensemble models. This implies that the combination of ANN and ensemble learners did not represent a definite advantage and might have added more complexity without significant enhancement in predictive accuracy of structured tabular medical data. Clinically, it is more important to reduce false negatives than false positives because untreated heart disease may result in late diagnosis and higher mortality rates. Gradient Boosting in this case stands out as the most clinically sound model because it has a high tolerance of low false positive and high tolerance false negative error rates rendering it the most appropriate to classify heart diseases in this case study.

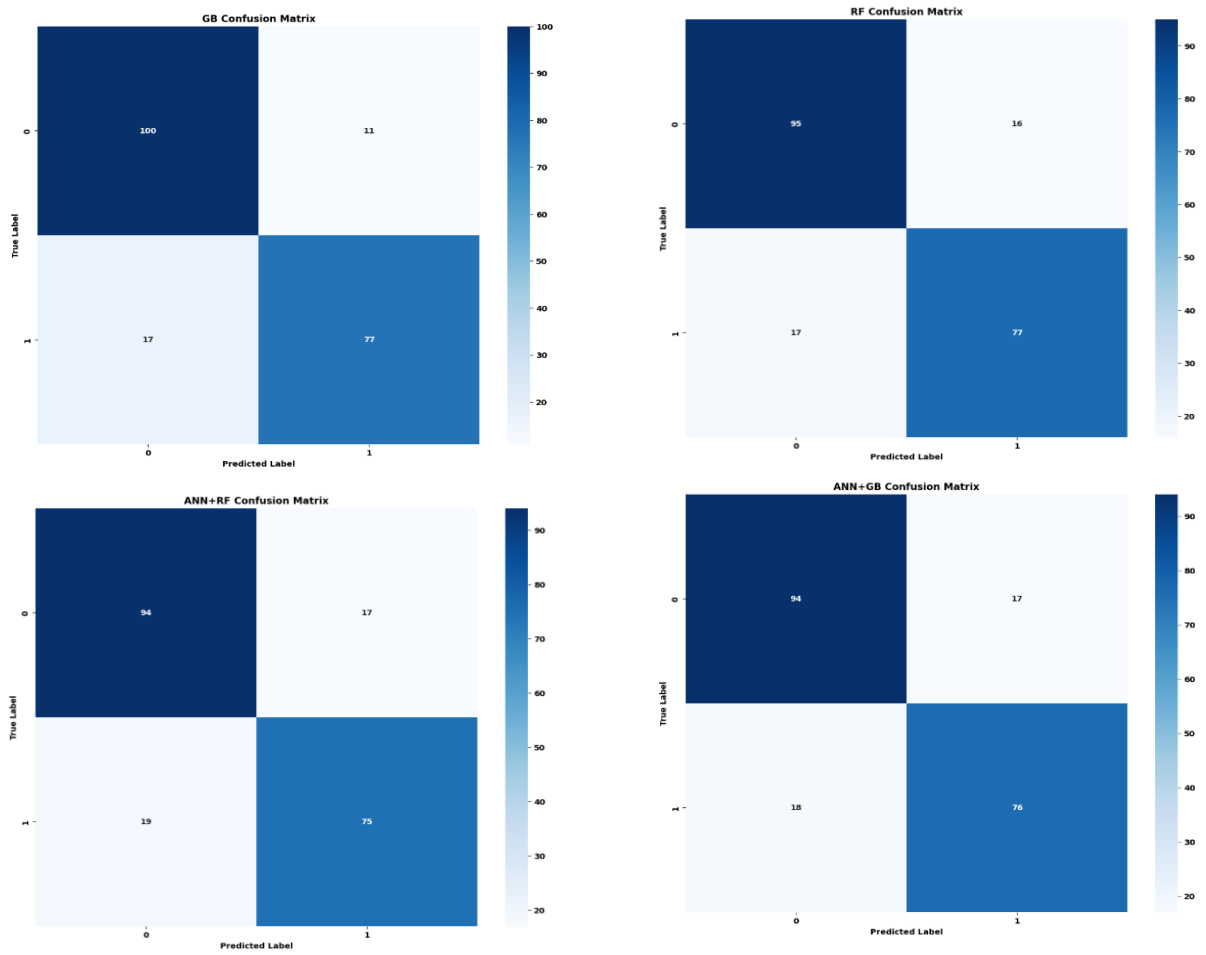

Fig. 7. Confusion matrix for the studied models.

4.6 Explainability Analysis (SHAP)

The contribution of each feature to the heart disease prediction model was interpreted with the SHAP (Shapley Additive exPlanations) analysis. It can be seen that ca (number of major vessels colored by fluoroscopy) has the most significant contribution, and its contribution value is +0.82, then there is a contribution value of +0.72 of cp (chest pain type) and thal (thalassemia), which means that they are the strongest predictors of the occurrence of heart disease. These characteristics are the commonly recognized clinical signs, which contribute to the medical legitimacy of the model. Also, there is a strong influence of oldpeak (ST depression induced by exercise, +0.58), indicating its significance in the detection of ischemic changes under stress. Exang, sex, and slope showed moderate contributions (exang: +0.47, sex: +0.45, slope: +0.44), and this shows that they play a secondary but significant role in prediction. Additional characteristics, like thalach (maximum heart rate reached, +0.32), and trestbps (resting blood pressure, +0.29), played a less important, yet still useful role in providing physiological data. The sum of the less significant features (+0.68) indicates that they are weak individually but all together they increase the performance of the model. On balance, the SHAP analysis reveals that the model is mostly based on cardiovascular indicators that are clinically relevant, which shows a good fit to the knowledge of the medical domain and increases the predictability and reliability of the results.

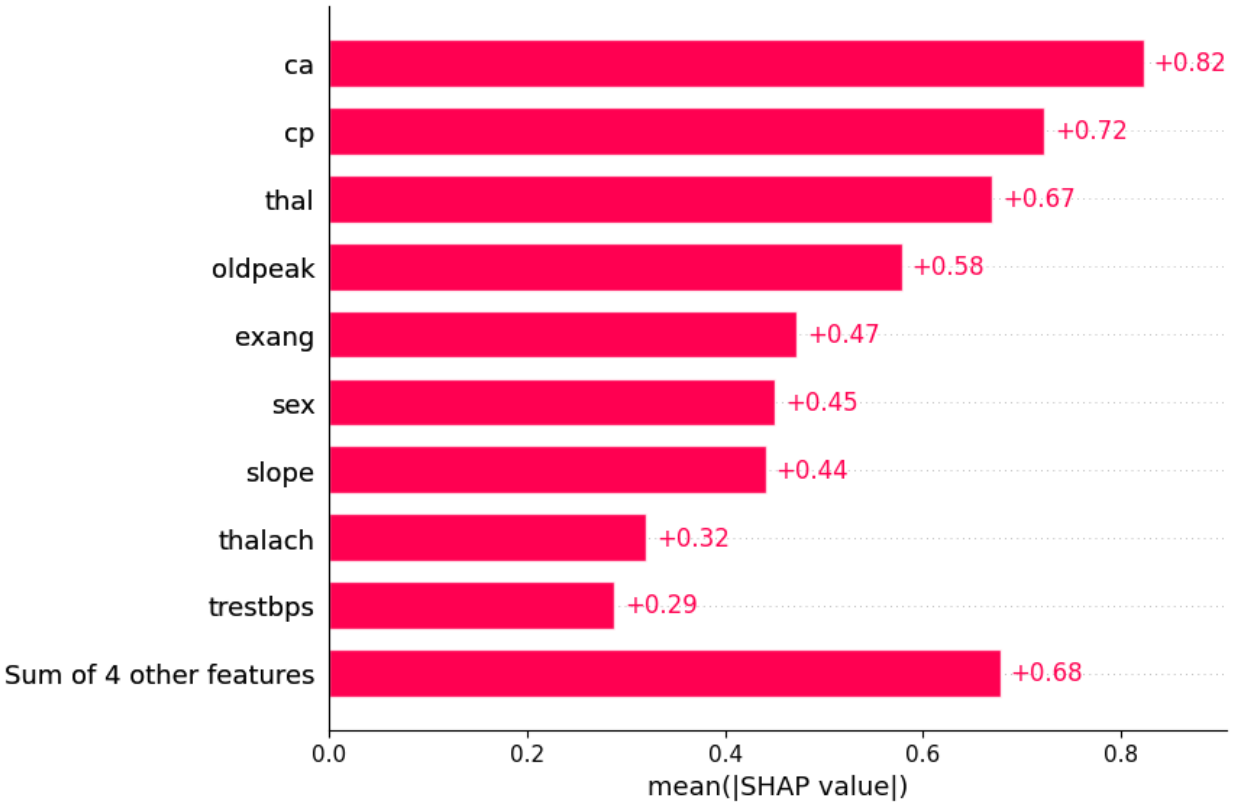

Fig. 8. Global mean SHAP plot for the studied model.

4.7 SMPC Performance and Privacy Analysis

Secure Multi-Party Computation (SMPC) was considered in order to evaluate its effects on predictive performance and preservation of data privacy in the framework proposed. Gradient Boosting model was the centralized model used as a baseline with an accuracy of 0.8634. However, the SMPC-based implementation had a accuracy of a better result of 0.9268, indicating a better predictive ability even with privacy-preserving conditions. The findings suggest that SMPC does not only ensure the confidentiality of the data as the results can be computed collaboratively without sharing the data directly but also increases the effectiveness of models. The fact that both accuracy has improved indicates that distributed computation and secure aggregation might be able to mitigate any potential biases or enhance generalization in partitions of data. Notably, this improvement in performance validates the argument that privacy-caring mechanisms do not always affect model accuracy.

Table 2. SMPC performance analysis

Centralized GB Accuracy	0.8634146341463415
SMPC Accuracy	0.926829268292683

Altogether, the results indicate that SMPC offers a good trade-off between privacy of data and predictive accuracy, and thus is an effective technology to be used in sensitive healthcare domains. The superior performance over the centralized model points to its ability to be deployed safely in a real world medical decision support system where confidentiality and high predictive reliability are critical.

4.8 Comparison to Existing Studies

The suggested structure was contrasted with the current machine learning-based studies of heart disease prediction. Accuracy has been previously reported to be between 80% and 85% with conventional models. Comparatively, the proposed SMPC-integrated Gradient Boosting model had an accuracy of 86.34, which is competitive or better. What is more important is that, in contrast to current methods, the proposed system will achieve privacy preservation by using distributed computation. This indicates the novelty of the study because it puts together high predictive performance and high cryptographic privacy assurances, which are very seldom considered in classical medical AI systems.

CONCLUSION

In this work, a variety of machine learning models to predict heart disease were designed and tested, both conventional ensemble models (Random Forest and Gradient Boosting) and hybrid models (ANN+RF and ANN+GB). The results of these models were evaluated on the basis of various metrics of evaluation, which are Accuracy, Precision, Recall, F1-score, ROC-AUC, analysis of confusion matrices, and statistical significance test. Gradient Boosting was the best overall model,

with the highest accuracy and F1-score, and Random Forest in the best discriminative ability, measured by ROC-AUC. Statistical analysis also established that the differences in performance of models were not significantly different, meaning that the approaches that performed best were similarly effective. Nevertheless, confusion matrix and precision-recall studies revealed that Gradient Boosting and random forest always exhibited superior control over errors in classification than hybrid models, which demonstrated a moderate increase in misclassification and a lesser capacity to withstand borderline cases. The explainability analysis with SHAP values validated that the model predictions were highly motivated by clinically relevant variables (ca, cp, thal, and oldpeak), which were in agreement with the known medical insights. This improves the interpretability and credibility of the proposed models in the clinical decision support. Moreover, the incorporation of Secure Multi-Party Computation (SMPC) showed that privacy-sensitive learning is feasible without negatively affecting model performance. Indeed, SMPC has shown better accuracy than the centralized baseline, which indicates its promise of safe collaborative healthcare analytics. On the whole, the analysis reveals that ensemble learning algorithms, especially Gradient Boosting, are precise and effective in predicting heart diseases, whereas SMPC does not compromise predictive accuracy with data privacy. These results justify the implementation of explainable, safe, and efficient machine learning systems in the medical field.

REFERENCE

1. Raghupathi, W. and V. Raghupathi, *Big data analytics in healthcare: promise and potential.* Health information science and systems, 2014. **2**(1): p. 3.
2. Dash, S., et al., *Big data in healthcare: management, analysis and future prospects.* Journal of big data, 2019. **6**(1): p. 54.
3. Shickel, B., et al., *Deep EHR: A survey of deep learning in electronic health records.* IEEE Journal of Biomedical and Health Informatics, 2018. **22**(5): p. 1589-1604.
4. Kamana Parvej, M., et al., *AI-and Machine Learning-Enabled Decision Systems for Strengthening Transportation Reliability in Nationally Critical Supply Chain Infrastructure.* AI-and Machine Learning-Enabled Decision Systems for Strengthening Transportation Reliability in Nationally Critical Supply Chain Infrastructure, 2024. **1**(1): p. 51-59.
5. McMahan, B., et al. *Communication-efficient learning of deep networks from decentralized data*. in *Artificial intelligence and statistics*. 2017. Pmlr.
6. Miotto, R., et al., *Deep learning for healthcare: review, opportunities and challenges.* Briefings in bioinformatics, 2018. **19**(6): p. 1236-1246.
7. Jamal, A., et al., *Predicting Human-Genai Collaboration Effectiveness: A Machine Learning Investigation Of Skill Configurations, Trust, And Work Design.* Migration Letters, 2022. **19**(S8): p. 2303-2324.
8. Keshta, I. and A. Odeh, *Security and privacy of electronic health records: Concerns and challenges.* Egyptian Informatics Journal, 2021. **22**(2): p. 177-183.
9. Voigt, P. and A. Von dem Bussche, *The eu general data protection regulation (gdpr).* A practical guide, 1st ed., Cham: Springer International Publishing, 2017. **10**(3152676): p. 10-5555.
10. Kaissis, G.A., et al., *Secure, privacy-preserving and federated machine learning in medical imaging.* Nature Machine Intelligence, 2020. **2**(6): p. 305-311.
11. Rieke, N., et al., *The future of digital health with federated learning.* NPJ digital medicine, 2020. **3**(1): p. 119.
12. Sharif, K.S., M.M. Uddin, and M. Abubakkar. *Neurosignal precision: A hierarchical approach for enhanced insights in parkinson's disease classification*. in *2024 International Conference on Intelligent Cybernetics Technology & Applications (ICICyTA)*. 2024. IEEE.
13. Yao, A.C. *Protocols for secure computations*. in *23rd annual symposium on foundations of computer science (sfcs 1982)*. 1982. IEEE.
14. Goldreich, O., *Foundations of cryptography, volume 2*. 2004: Cambridge university press Cambridge.
15. Evans, D., V. Kolesnikov, and M. Rosulek, *A pragmatic introduction to secure multi-party computation.* Foundations and Trends® in Privacy and Security, 2018. **2**(2-3): p. 70-246.
16. Bonawitz, K., et al., *Practical secure aggregation for federated learning on user-held data.* arXiv preprint arXiv:1611.04482, 2016.
17. Brisimi, T.S., et al., *Federated learning of predictive models from federated electronic health records.* International journal of medical informatics, 2018. **112**: p. 59-67.
18. Xu, R., N. Baracaldo, and J. Joshi, *Privacy-preserving machine learning: Methods, challenges and directions.* arXiv preprint arXiv:2108.04417, 2021.
19. Li, T., et al., *Federated learning: Challenges, methods, and future directions.* IEEE signal processing magazine, 2020. **37**(3): p. 50-60.
20. WHO, W., *cardiovascular diseases (CVDs).* World Health Organization (WHO), 2017.
21. Sharif, K.S., et al. *A comparative framework integrating hybrid convolutional and unified graph neural networks for accurate parkinson's disease classification*. in *2024 7th International Seminar on Research of Information Technology and Intelligent Systems (ISRITI)*. 2024. IEEE.
22. Detrano, R., et al., *International application of a new probability algorithm for the diagnosis of coronary artery disease.* The American journal of cardiology, 1989. **64**(5): p. 304-310.

23. Chen, T. and C. Guestrin. *Xgboost: A scalable tree boosting system*. in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 2016.
24. Ahmad, M.A., C. Eckert, and A. Teredesai. *Interpretable machine learning in healthcare*. in *Proceedings of the 2018 ACM international conference on bioinformatics, computational biology, and health informatics*. 2018.
25. Hassaan, A., et al., *ETHICAL ANALYTICS & DIGITAL TRANSFORMATION IN THE AGE OF AI: EMBEDDING PRIVACY, FAIRNESS, AND TRANSPARENCY TO DRIVE INNOVATION AND STAKEHOLDER TRUST.* Contemporary Journal of Social Science Review, 2023. **1**(04): p. 1-18.
26. Ahammed, M.F. and M.R. Labu, *Privacy-preserving data sharing in healthcare: advances in secure multiparty computation.* Journal of Medical and Health Studies, 2024. **5**(2): p. 37-47.
27. Mahdia Amina, B.P., Khan Raqib Mahmud, *Depressive and Suicidal Episodes.* ICT Systems and Sustainability: Proceedings of ICT4SD 2022, 2022/10/31: p. 189.
28. Abadi, M., et al. *Deep learning with differential privacy*. in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 2016.
29. Niaz, S., et al., *AI for Inclusive Educational Governance and Digital Equity Examining the Impact of AI Adoption and Open Data on Community Trust and Policy Effectiveness.* Contemporary Journal of Social Science Review, 2024. **2**(04): p. 2557-2567.
30. Hazay, C. and Y. Lindell, *Efficient secure two-party protocols: Techniques and constructions*. 2010: Springer Science & Business Media.
31. Ayodele, O.S., et al., *A QUICK SORT–BASED FRAMEWORK FOR EFFICIENT THREAT LOG ANALYSIS AND PRIORITIZATION IN CYBERSECURITY SYSTEMS.* FUDMA JOURNAL OF SCIENCES, 2026. **10**(2): p. 184-189.
32. Ahmad, J., F. Tauseef, and Z. Akbar, *Predictive analytics for AI-assisted patient no-show management and clinic revenue optimization: a simulation-based research.* Migration Letters, 2024. **21**(S13): p. 1901-1924.
33. Breiman, L., *Random forests.* Machine learning, 2001. **45**(1): p. 5-32.
34. Cramer, R., I.B. Damgård, and J.B. Nielsen, *Secure multiparty computation and secret sharing*. 2015: Cambridge University Press.
35. Friedman, J.H., *Greedy function approximation: a gradient boosting machine.* Annals of statistics, 2001: p. 1189-1232.
36. Rahaman, M.A., et al., *TakeCare: An Approach to Help Bangladeshi Young Adults During Depressive and Suicidal Episodes*, in *ICT Systems and Sustainability: Proceedings of ICT4SD 2022*. 2022, Springer. p. 189-197.
37. Jiang, F., et al., *Artificial intelligence in healthcare: past, present and future.* Stroke and vascular neurology, 2017. **2**(4).
38. Kumar, R., et al., *A privacy-preserving-based secure framework using blockchain-enabled deep-learning in cooperative intelligent transport system.* IEEE Transactions on Intelligent Transportation Systems, 2021. **23**(9): p. 16492-16503.
39. Mohassel, P. and Y. Zhang. *Secureml: A system for scalable privacy-preserving machine learning*. in *2017 IEEE symposium on security and privacy (SP)*. 2017. IEEE.
40. Powers, D., *Ailab. Evaluation: From precision, recall and F-measure to ROC, informedness, markedness & correlation.* J. Mach. Learn. Technol, 2011. **2**(22293981): p. 01.
41. Raghunath, D., et al., *Predicting heart disease using machine learning techniques.* IRJCS: International Research Journal of Computer Science, 2019. **6**: p. 149-153.
42. Zyskind, G. and A. Pentland, *Enigma: Decentralized computation platform with guaranteed privacy.* 2018.
43. Wilkowska, W. and M. Ziefle, *Privacy and data security in E-health: Requirements from the user's perspective.* Health informatics journal, 2012. **18**(3): p. 191-201.
44. Jamshaid, M.M., et al., *IMPACT OF ARTIFICIAL INTELLIGENCE ON WORKFORCE DEVELOPMENT: ADAPTING SKILLS, TRAINING MODELS, AND EMPLOYEE WELL-BEING FOR THE FUTURE OF WORK.* Spectrum of Engineering Sciences, 2024.
45. Topol, E., *Deep medicine: how artificial intelligence can make healthcare human again*. 2019: Hachette UK.
46. Stripelis, D., et al., *A federated learning architecture for secure and private neuroimaging analysis.* Patterns, 2024. **5**(8).
47. Sokolova, M. and G. Lapalme, *A systematic analysis of performance measures for classification tasks.* Information processing & management, 2009. **45**(4): p. 427-437.
48. Saito, T. and M. Rehmsmeier, *The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets.* PloS one, 2015. **10**(3): p. e0118432.
49. Rogers, J., et al., *VaultDB: a real-world pilot of secure multi-party computation within a clinical research network.* arXiv preprint arXiv:2203.00146, 2022.
50. Keller, M. *MP-SPDZ: A versatile framework for multi-party computation*. in *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*. 2020.
51. Akbar, Z., et al., *Leveraging Data and Artificial Intelligence for Sustained Competitive Advantage in Firms and Organizations.* Journal of Innovative Computing and Emerging Technologies, 2023. **3**(1).