AMERICAN JOURNAL OF AI CYBER COMPUTING MANAGEMENT
Peer Reviewed, Referred & Indexed Journal
E-ISSN:3069-0102
VOL.6, NO. 2(2026)

Edge-Enabled Diagnostic Engine with Homomorphic Encryption for Real-Time Medical Decision Support

B. Rajesh Reddy[1], Marri Vignesh Kumar[2], Manchala Shashi Pradeep[2], Dolka Snehith[2]

[1]Assistant Professor, [2]UG Student, [1,2]Department of Computer Science and Engineering

[1,2]Kommuri Pratap Reddy Institute of Technology, Ghanpur, Ghatkesar, 501301, Telangana, India.

*Correspondence: B. Rajesh Reddy (rajeshreddy.reddy54@gmail.com)

ABSTRACT

The rapid expansion of digital healthcare systems has intensified the demand for secure and privacy-preserving data analysis, particularly in sensitive domains such as medical diagnosis. Conventional approaches often process raw patient data directly, increasing exposure to security risks, unauthorized access, and privacy breaches, while also lacking robust authentication mechanisms and failing to ensure data confidentiality during both training and prediction phases. This creates a critical need for integrated frameworks that combine accurate prediction with strong privacy protection and secure system access. To address these challenges, this study proposes a privacy-preserving healthcare prediction system that integrates Machine Learning (ML) with encryption and secure authentication mechanisms. The system is developed using the Flask framework and incorporates Simple Mail Transfer Protocol (SMTP)-based One-Time Password (OTP) verification to ensure secure user login. Data privacy is maintained through Lightweight Privacy-Preserving Machine Learning Encryption (LPME), where dataset features are encrypted using polynomial-based cryptographic operations prior to model training, preventing data leakage. The system supports datasets such as heart disease and hypothyroid conditions, applying preprocessing techniques including normalization, label encoding, and Synthetic Minority Oversampling Technique (SMOTE) to handle class imbalance effectively. For prediction, Gaussian Naive Bayes (GNB) and Extreme Gradient Boosting (XGBoost) models are employed, with XGBoost serving as the primary model due to its ability to capture complex feature interactions and deliver higher predictive performance, while GNB is used for comparative analysis. The system's performance is evaluated using accuracy, precision, recall, and F1-score.

Keywords: Privacy-Preserving Machine Learning, Healthcare Data Security, One-Time Password (OTP) Authentication, Lightweight Encryption, Medical Diagnosis Prediction, Data Confidentiality,

1. INTRODUCTION

The healthcare system is a fundamental pillar of a nation's development, as it directly influences population health, workforce productivity, and economic growth. In recent years, the rapid spread of contagious and chronic diseases has placed immense pressure on healthcare systems worldwide, significantly impacting global economies and exposing limitations in traditional healthcare delivery models [1]. To address these challenges, the integration of Internet of Things (IoT)-based remote healthcare devices, commonly known as smart healthcare systems, has emerged as a transformative solution. These systems enable continuous, real-time monitoring of patients across different age groups without requiring frequent hospital visits, thereby improving accessibility, reducing healthcare costs, and supporting early detection and preventive care [2].

Despite these advantages, the increasing use of body-worn IoT devices introduces serious concerns related to data confidentiality, integrity, and secure transmission, as sensitive patient data is continuously collected, transmitted, and stored across distributed networks, making it highly vulnerable to cyberattacks, unauthorized access, and data breaches [3]. In addition, traditional cloud-

based analytics, although powerful in handling large-scale data processing, suffer from inherent latency issues due to the physical distance between the data source and centralized servers, which can lead to delayed responses in critical care scenarios where immediate decision-making is essential [4]. To overcome these dual challenges of privacy and latency, this work proposes an Edge-Enabled Diagnostic Engine integrated with Homomorphic Encryption (HE), which provides a secure and efficient framework for healthcare analytics. By shifting computational tasks from centralized cloud infrastructure to edge nodes located closer to patients, the system significantly reduces communication delays, enhances real-time processing capabilities, and ensures faster medical decision support [5]. Simultaneously, HE enables computations to be performed directly on encrypted data without requiring decryption, thereby preserving data confidentiality throughout the entire processing lifecycle and minimizing the risk of data exposure. The encryption and processing mechanisms are implemented within the edge middleware, ensuring that patients are not burdened with complex system configurations while maintaining seamless usability. The proposed system focuses on enabling real-time prediction of critical health conditions using data collected from remote patient monitoring sensors, thereby facilitating timely intervention and improving patient outcomes. Furthermore, this study addresses a crucial research challenge: how to maintain strict data confidentiality in compliance with Health Insurance Portability and Accountability Act (HIPAA) regulations while simultaneously achieving the low-latency requirements necessary for real-time healthcare diagnostics and decision-making [6].

2. RELATED WORK

2.1 Security, Authentication, and Access Control

Security and privacy are critical concerns in IoT-based healthcare systems due to the sensitive nature of medical data. Jia et al. [7] proposed an identity-based cross-domain authentication scheme that enables secure communication across heterogeneous IoT environments without relying on traditional certificate-based mechanisms. This approach improves interoperability while protecting against attacks such as impersonation and replay attacks. Similarly, Sun et al. [8] introduced a searchable personal health records framework with fine-grained access control in cloud-fog computing environments. Their system utilizes attribute-based access control along with searchable encryption to ensure secure and efficient retrieval of encrypted medical data. In addition, Patil et al. [10] developed a chain-of-custody framework using Attribute-Based Signcryption (ABSC), which guarantees data confidentiality, integrity, and authenticity while enabling traceability of data access and modifications. Collectively, these approaches enhance secure data sharing and controlled access in distributed healthcare systems.

2.2 Edge and Fog Computing for Healthcare

To overcome latency and scalability challenges in healthcare systems, edge and fog computing paradigms have been widely explored. Rahmani et al. [9] proposed a fog computing-based architecture that introduces an intermediate processing layer between IoT devices and cloud infrastructure. This approach enables real-time data processing closer to the source, thereby reducing latency, bandwidth consumption, and energy usage. Similarly, Raj [11] presented an optimized Mobile Edge Computing (MEC) framework for IoT-based medical sensor networks, focusing on improving cooperation among edge nodes. Their framework demonstrated enhanced task execution efficiency and better resource utilization using real-world wearable sensor datasets. These studies highlight the importance of decentralized computing in supporting time-sensitive healthcare applications.

2.3 Intelligent Healthcare Analytics and Anomaly Detection

Machine learning and deep learning techniques have been increasingly applied to improve healthcare monitoring systems. Kim et al. [12] proposed a convolutional variational autoencoder-based model for unsupervised anomaly detection in edge-based IoT systems. Their approach enables the detection of abnormal patterns in sensor data without requiring labeled datasets, making it highly suitable for real-time healthcare monitoring. Such intelligent models play a crucial role in identifying critical health conditions at an early stage and improving the reliability of healthcare systems.

2.4 IoT-Based Healthcare System Architectures

Several studies have focused on designing comprehensive IoT-based healthcare frameworks. Islam et al. [13] presented a layered architecture consisting of sensing devices, data aggregation, storage, and analytics components. Their work highlights the potential of IoT in enabling continuous patient monitoring and improving healthcare service delivery. At the same time, it identifies key challenges such as data security, interoperability, and scalability, which must be addressed to achieve efficient and robust healthcare systems.

2.5 Research Gap

Although existing studies address key aspects such as security, access control, edge/fog computing, and intelligent analytics, most solutions focus on these components independently. There is a lack of integrated frameworks that simultaneously ensure data confidentiality, low-latency processing, and real-time analytics. In particular, performing secure computations on encrypted healthcare data while maintaining edge-level efficiency remains underexplored. This motivates the need for a unified approach that combines privacy-preserving mechanisms with real-time healthcare intelligence.

3. PROPOSED SYSTEM

This research focuses on building a secure and intelligent healthcare analytics framework that enables accurate disease prediction while safeguarding sensitive patient information. With the increasing digitization of medical records, large volumes of health data are being generated and analyzed to support clinical decision-making. However, ensuring data privacy while extracting meaningful insights remains a major challenge.

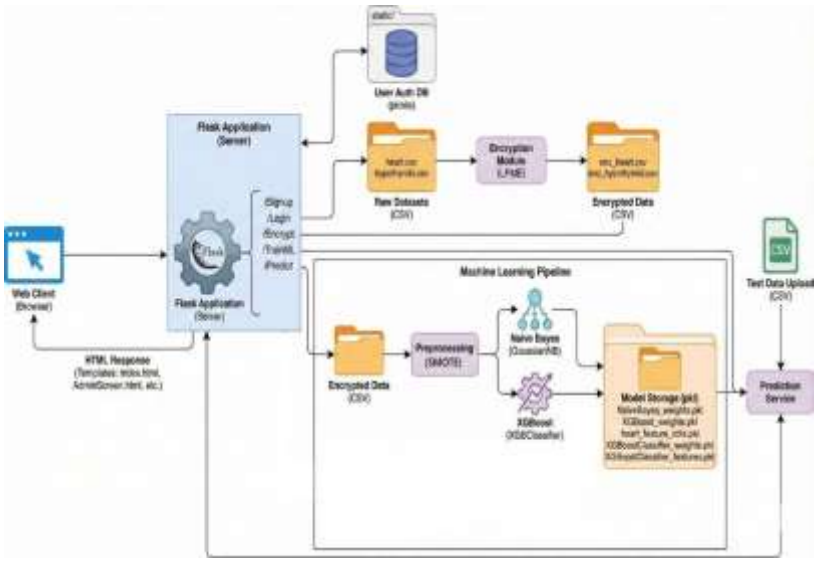

Fig. 1: Proposed System Architecture

The study addresses this issue by combining privacy-preserving techniques with machine learning methods to support reliable analysis of medical data related to heart and thyroid conditions. The study begins by transforming medical datasets using privacy-aware encryption mechanisms to protect confidential attributes before any analytical process is performed. This ensures that patient information remains secure during storage, training, and prediction stages. After encryption, the data undergoes preprocessing steps such as handling missing values, normalization, and feature alignment to maintain consistency and improve learning efficiency as shown in Fig. 1. These steps play a crucial role in enhancing the stability and accuracy of predictive models. The architecture also supports secure inference, allowing users to upload new test records that are automatically aligned, normalized, and processed before model prediction. Since the system never handles plaintext sensitive data after encryption, it provides strong protection against privacy breaches, unauthorized access, and insider threats. Through its combination of encryption, automation, and machine learning, the proposed system delivers a practical and secure solution for medical analytics, enabling healthcare organizations to leverage AI-based decision support while ensuring compliance with privacy requirements.

3.1 XGBoost Classifier

XGBoost is an advanced ensemble machine learning algorithm based on the gradient boosting framework, which builds multiple decision trees sequentially to improve prediction accuracy. It works by minimizing a loss function using gradient descent, where each new tree is trained to correct the errors of the previous trees. XGBoost introduces regularization techniques to prevent overfitting and supports parallel processing for faster computation. It efficiently handles missing values, large-scale datasets, and complex feature interactions, making it highly effective for classification and regression tasks. The algorithm combines predictions from multiple weak learners to produce a strong predictive model with high accuracy and robustness. Due to its performance and scalability. Fig. 2 illustrates the internal working process of the XGBoost algorithm.

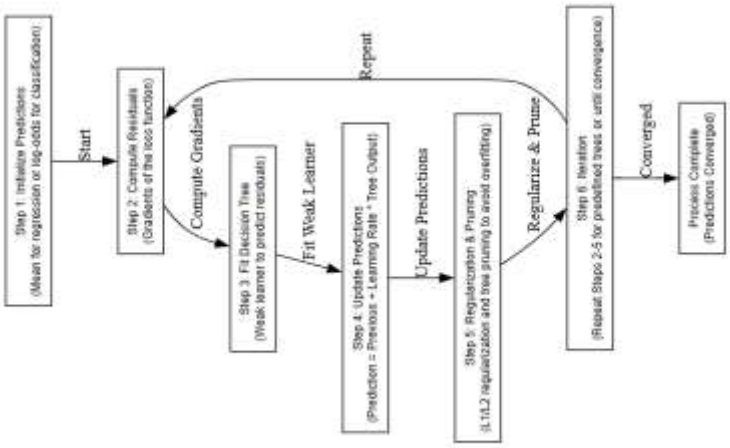

Fig. 2: workflow of XGBoost Classifier

Step 1 Initialize Predictions: Start with a simple prediction (e.g., mean of target for regression or log-odds for classification).

Step 2 Compute Residuals (Gradients): For each iteration, compute the gradient of the loss function with respect to the current prediction. This identifies how much each sample's prediction is wrong.

Step 3 Fit Decision Tree: Train a small decision tree (weak learner) to predict these residuals. The tree splits data based on features that reduce the loss most effectively.

Step 4 Update Predictions: Update the overall prediction by adding a fraction of the new tree's output (learning rate):

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + \eta f_t(x_i)$$

Step 5 Regularization & Pruning: Apply L1/L2 regularization on leaf weights and prune trees to prevent overfitting.

Step 6 Iteration: Repeat steps 2–5 for a predefined number of trees or until performance converges.

3.2 LPME

The LPME technique is a custom-designed encryption approach intended to safeguard sensitive medical information while still allowing meaningful machine learning operations on transformed data. Traditional encryption methods protect confidentiality but make it impossible to directly use encrypted data for analytics or model training. LPME addresses this challenge by converting raw numerical values into structured encrypted polynomial coefficients that preserve essential mathematical relationships required for downstream machine learning. The workflow of LPME is illustrated in Fig. 3

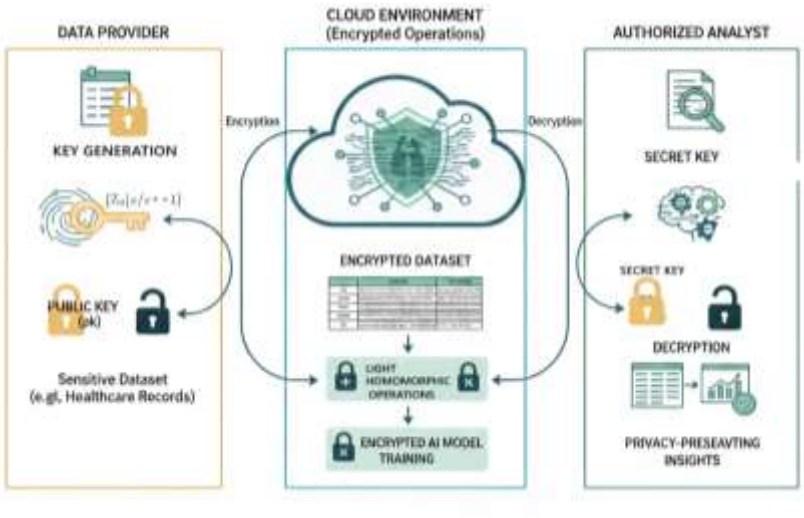

Fig. 3: LPME workflow

This enables the system to protect patient confidentiality without fully sacrificing the utility of the dataset. LPME operates on the principle of polynomial-based public key encryption, inspired by lattice and Learning-With-Errors (LWE) cryptographic constructions. It embeds each plaintext feature value into a polynomial representation and introduces controlled randomness and noise using secret and public keys. This ensures that even if an attacker gains access to the encrypted dataset, the original values cannot be recovered without the secret key. The encryption process produces ciphertext values that appear random but still maintain certain algebraic properties, allowing limited plaintext operations such as scaled addition or multiplication when necessary. In the context of this research, LPME is applied to all numerical features in medical datasets such as heart disease and hypothyroid data. Each feature value is transformed independently, ensuring strong protection for every record.

1. Dataset Loading: The workflow begins by loading the original medical dataset from a CSV file. The system reads each row and extracts all numerical features that require encryption. This ensures the raw data is properly structured before any transformation. Missing values in the dataset are replaced to maintain consistency during encryption. The entire dataset is then converted to a NumPy array for efficient processing.

2. Normalization of Feature Values: Each feature value is passed through the normalize value() function, which converts all inputs into integer form. Categorical symbols like 'T', 'F', '?', or empty values are mapped to numeric representations. This guarantees the encryption process receives clean, valid integers. Normalization also resolves inconsistent data formatting commonly found in medical datasets.

3. Key Generation: Before encryption begins, the system generates a public key and a secret key using polynomial-based key generation. The public key consists of two polynomials derived from uniform and noise distributions, while the secret key is a binary polynomial. These keys form the mathematical foundation for secure polynomial encryption. The keys ensure that encrypted values cannot be reversed without the corresponding secret key.

4. Plaintext Polynomial Encoding: Each normalized integer value is encoded into a polynomial format, where the first coefficient represents the actual message and the remaining coefficients are zeros. This conversion embeds the plaintext into a structure compatible with polynomial cryptography. Encoding as a polynomial also enables homomorphic-style operations. This step prepares the plaintext for integration with randomness and noise.

5. Message Scaling Using Modulus: The encoded polynomial message is scaled by computing delta = q/t, which maps the plaintext space into the ciphertext modulus space. Scaling ensures message components remain distinguishable after encryption noise is added. This safeguards the message from overflow or loss during polynomial operations. It also standardizes message representation across all features.

6. Randomness and Noise Generation: Three types of noise components are generated: a binary vector u and two error polynomials e1 and e2. These random elements strengthen encryption security by adding unpredictability to ciphertexts. Even identical plaintext values produce different encrypted outputs due to this randomness. Noise also prevents adversaries from performing statistical attacks.

7. Polynomial Ciphertext Computation: The encryption function computes ciphertexts (ct0, ct1) using polynomial multiplication and addition with the public key. ct0 combines the message, noise, and part of the public key, while ct1 retains the second half of the key structure. The result is a secure encrypted representation of the original value. These ciphertexts follow the principles of lattice/LWE-style encryption.

8. Encrypted Value Extraction and Storage: From each ciphertext, the system extracts the first coefficient from ct0, which becomes the encrypted value stored in the dataset. This coefficient is a large integer that hides the plaintext while preserving its analytic structure. The encrypted value is then written into a new CSV file. The process is repeated for all features in each dataset row.

9. Final Encrypted Dataset Output: After all rows and features are processed, the system exports a fully encrypted dataset ready for privacy-preserving machine learning. The output file contains only encrypted numerical values, protecting sensitive patient information. This ensures downstream ML models operate on data that never exposes raw medical attributes. The encrypted dataset maintains compatibility with normalization and training workflows.

4. Results description

Fig. 4 depicts the comparative visualization of the plain heart disease dataset and its encrypted counterpart. Part (A) illustrates the original medical dataset containing clinical attributes used for diagnostic analysis, while Part (B) shows the transformed encrypted representation of the same data. This comparison highlights how sensitive medical information is secured without losing structural integrity required for analysis. The figure emphasizes the effectiveness of encryption in protecting data confidentiality. It also demonstrates the system's capability to operate on encrypted medical datasets.

Fig. 5 depicts the output visualization of the hypothyroid dataset in both plain and encrypted formats. Part (A) represents the original hypothyroid medical records containing diagnostic features, while Part (B) illustrates the encrypted dataset generated for secure processing. This figure highlights the system's ability to handle multiple medical datasets using a unified encryption and analytics framework. The transformation preserves analytical usability while ensuring data security. It reinforces the applicability of the proposed system across diverse medical diagnostic scenarios.

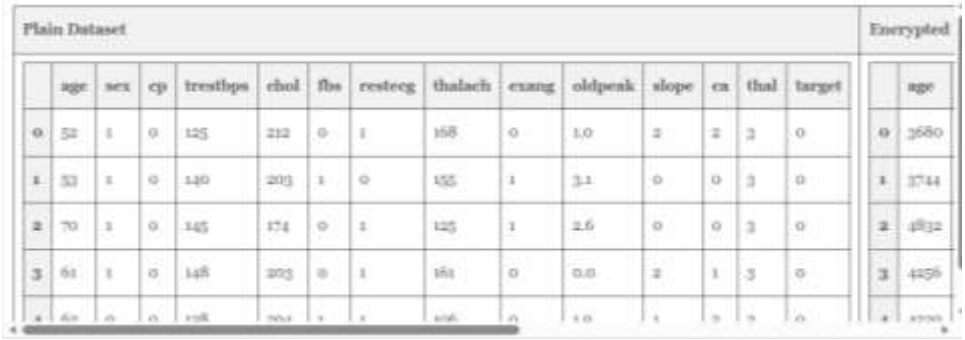

(A)

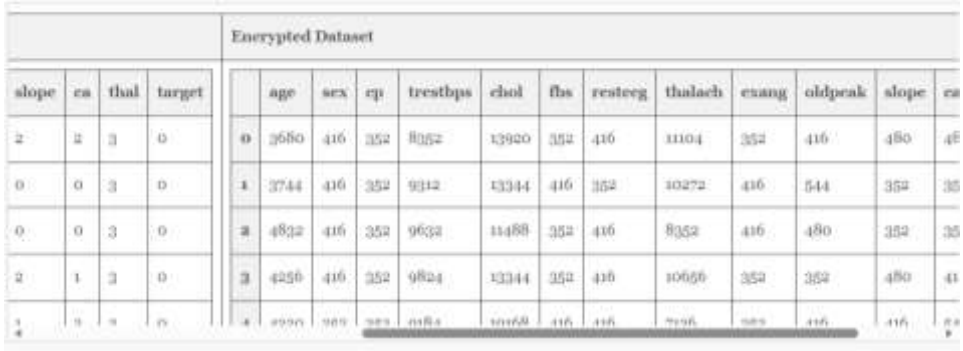

(B)

Fig. 4: (A) Plain and (B) Encrypted Heart Dataset Display Screen

(A)

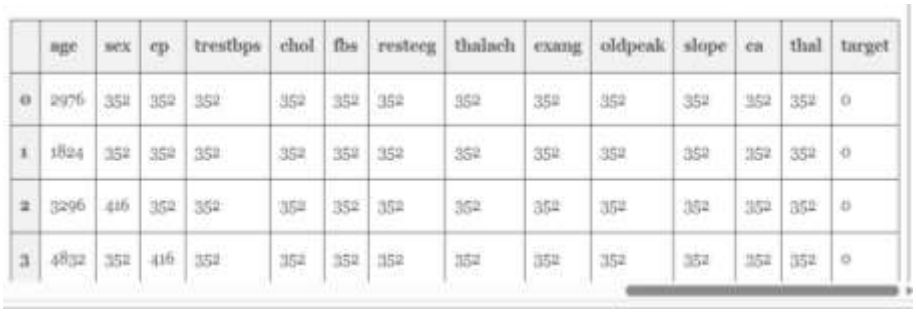

	age	sex	cp	trestbps	chol	fbs	resteeg	thalach	exang	oldpeak	slope	ca	thal	target
0	2976	352	352	352	352	352	352	352	352	352	352	352	352	0
1	1824	352	352	352	352	352	352	352	352	352	352	352	352	0
2	3296	416	352	352	352	352	352	352	352	352	352	352	352	0
3	4832	352	416	352	352	352	352	352	352	352	352	352	352	0

(B)

Fig. 5: (A) Plain and (B) Encrypted Hypothyroid Dataset Output Screen

Fig. 6 depicts the prediction result display screen generated after processing the uploaded test data. The figure shows the extracted feature values used by the trained machine learning model for medical condition assessment. It represents the system's ability to interpret patient-specific input attributes and produce diagnostic outcomes. This screen demonstrates real-time inference capability within the edge-enabled architecture. It confirms the integration of secure data handling and intelligent prediction in the proposed system.

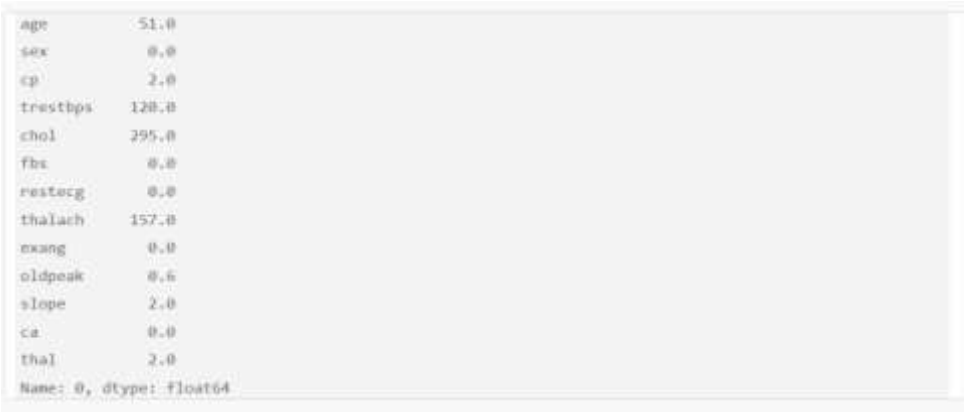

```
age         51.0
sex          0.0
cp           2.0
trestbps   120.0
chol       295.0
fbs          0.0
restecg      0.0
thalach    157.0
exang        0.0
oldpeak      0.6
slope        2.0
ca           0.0
thal         2.0
Name: 0, dtype: float64
```

Fig. 6: Prediction Result Display Screen

The comparison table 1 presents the performance of two machine learning algorithms Naive Bayes and XG Boost evaluated on the encrypted heart disease dataset. The metrics considered include Accuracy, Recall, and Specificity, which collectively measure how effectively each model distinguishes between heart disease and non–heart disease cases. Naive Bayes performs reasonably well with an accuracy of 81.95%, demonstrating its capability to classify most instances correctly. However, the recall value of 81.77% indicates that some positive cases are still misclassified. XG Boost, on the other hand, significantly outperforms Naive Bayes by achieving a perfect score of 100% across all three metrics. This indicates that XG Boost correctly identifies all positive and negative cases without any misclassifications. The results clearly show that XG Boost is the most robust and reliable model for this dataset, making it the preferred choice for accurate heart disease prediction within the proposed system.

Table 1: Comparison of Algorithm Performance

Algorithm Name	Accuracy (%)	Recall (%)	Specificity (%)
Naive Bayes	81.95	81.77	85.98
XG Boost	100.00	100.00	100.00

5. CONCLUSION

The study concludes that combining privacy-preserving techniques with machine learning provides an effective and secure approach for healthcare data analysis. By protecting sensitive medical information through encryption while enabling accurate disease prediction, the research addresses critical concerns related to data security and patient confidentiality. The use of advanced learning algorithms, along with proper data preprocessing and imbalance handling, improves prediction reliability for heart and thyroid conditions. Overall, the study demonstrates that secure, intelligent analytics can support early diagnosis and informed medical decision-making, contributing to the development of trustworthy and efficient healthcare systems. Comprehensive preprocessing steps, including normalization, encoding, and SMOTE balancing, significantly improve model accuracy and robustness. The evaluation results confirm that the system is capable of delivering reliable disease classification while maintaining strict privacy standards. This research demonstrates a practical and secure end-to-end solution for medical data analysis. It proves that advanced encryption techniques can be combined with machine learning to build systems that are both privacy-aware and performance-driven. The platform is well-suited for real-world healthcare applications where data confidentiality and predictive accuracy are critical.

REFERENCES

[1] Rancea, A.; Anghel, I.; Cioara, T. Edge Computing in Healthcare: Innovations, Opportunities, and Challenges. Future Internet 2024, 16, 329. https://doi.org/10.3390/fi16090329

[2] Rauniyar, A.; Hagos, D.H.; Jha, D.; Håkegård, J.E.; Bagci, U.; Rawat, D.B.; Vlassov, V. Federated Learning for Medical Applications: A Taxonomy, Current Trends, Challenges, and Future Research Directions. IEEE Internet Things J. 2024, 11, 7374–7398

[3] Rahman MA, Shahrior MF, Iqbal K, Abushaiba AA. Enabling Intelligent Industrial Automation: A Review of Machine Learning Applications with Digital Twin and Edge AI Integration. Automation. 2025; 6(3):37. https://doi.org/10.3390/automation6030037

[4] Mehrabi, A.; Yari, K.; Van Driel, W.D.; Poelma, R.H. AI-Driven Digital Twin for Health Monitoring of Wide Band Gap Power Semiconductors. In Proceedings of the 2024 IEEE 10th Electronics System-Integration Technology Conference (ESTC), Berlin, Germany, 11–13 September 2024; IEEE: Piscataway, NJ, USA, 2024; pp. 1–8.

[5] Ma, Q.; Niu, J.; Ouyang, Z.; Li, M.; Ren, T.; Li, Q. Edge Computing-Based 3D Pose Estimation and Calibration for Robot Arms. In Proceedings of the 2020 7th IEEE International Conference on Cyber Security and Cloud Computing (CSCloud)/6th IEEE International Conference on Edge Computing and Scalable Cloud (EdgeCom), New York, NY, USA, 22–24 August 2020; pp. 246–251.

[6] Hang, L.; Kim, D.-H., Design and implementation of an integrated iot blockchain platform for sensing data integrity. Sensors 2019, 19, 2228.

[7] Jia, X.; Hu, N.; Su, S.; Yin, S.; Zhao, Y.; Cheng, X.; Zhang, C., IRBA: An identity-based cross-domain authentication scheme for the internet of things. J. Electron. 2020, 9, 634.

[8] Sun, J.; Wang, X.; Wang, S.; Ren, L.; Mehmood, R. A searchable personal health records framework with fine-grained access control in cloud-fog computing. PLoS ONE 2018, 13, e0207543.

[9] Rahmani, A.M.; Gia, T.N.; Negash, B.; Anzanpour, A.; Azimi, I.; Jiang, M.; Liljeberg, P., Exploiting smart e-Health gateways at the edge of healthcare Internet-of-Things: A fog computing approach. Future Gener. Comput. Syst. 2018, 78, 641–658.

[10] Patil, P.; Mukane, S.; Nagpure, S.; Patil, R. Maintaining Chain of Custody using Attribute Based Signcryption (ABSC). In Proceedings of the 2024 International Conference on Intelligent Computing and Emerging Communication Technologies (ICEC), Guntur, India, 23–25 November 2024.

[11] Raj, J.S. Optimized Mobile Edge Computing Framework for IoT based Medical Sensor Network Nodes. J. Ubiquitous Comput. Commun. Technol. 2021, 3, 33–42.

[12] Kim, D.; Yang, H.; Chung, M.; Cho, S.; Kim, H.; Kim, M.; Kim, K.; Kim, E., Squeezed Convolutional Variational AutoEncoder for Unsupervised Anomaly Detection in Edge Device Industrial Internet of Things. In Proceedings of the 2018 International Conference in Information and Communication Technologies (ICICT), DeKalb, IL, USA, 23–25 March 2018; pp. 67–71.

[13] Islam, M.S.; Humaira, F.; Nur, F.N., Healthcare Applications in IoT. Global. J. Med Res. B Pharma Drug Discov. Toxicol. Med. 2020, 20, 1–3.

[14] Patyrykin, K., & Vasyukova, L. (2025). Environmental Accountability or Symbolic Compliance? A Critical Review of ESG Ratings, Greenwashing, and Indirect Emissions in the Global Insurance Sector. International Journal of Energy Economics and Policy, 15(6), 917–925. https://doi.org/10.32479/ijeep.22770