RESEARCH PAPER

Optimal Privacy Preservation of Structured and Unstructured Datasets in Machine Learning: A Homomorphic Encryption-Based Framework for Secure Healthcare Data Analytics

Snehal Chaudhary[1*], Sunita Dhotre[2], Trupti Patil[4]

[1,2]Bharati Vidyapeeth (Deemed to be University) College of Engineering, Pune, India.
[3]K J Somaiya Institute of Technology, Mumbai, India.

[1]sdchaudhary@bvucoep.edu.in

Abstract - Privacy of patients is of utmost importance during the training of Machine Learning (ML) and Deep Learning (DL) systems. The problem is particularly acute when these systems are applied in the healthcare sector which deals with very sensitive data. The authors of this paper present a fully privacy-preserving setting that relies upon homomorphic encryption (HE) schemes namely CKKS, BFV, and RSA which are used as both ML and DL pipelines for structured and unstructured medical data. These schemes are employed in an experimental and theoretical study of a chronic kidney disease dataset as well as a CT-based kidney image dataset divided into four classes. Performance is assessed for the original, AES-encrypted, and HE-encrypted situations. The test results illustrate that the Decision Tree and CNN techniques perform equally well without loss in accuracy (up to 100%) after encryption, while CKKS turns out to be the most effective HE schemes in terms of encryption time and key size resource saving. The study validates that privacy-preserving training and inference among others are possible with a relatively small decrease in the performance. This architecture facilitates the transition of cryptographic algorithms from theory to practice in AI applications and provides a privacy-compliant way to handle healthcare data.

Keywords - Deep learning in healthcare, homomorphic encryption, privacy-preserving machine learning, encrypted model inference, CKD prediction, and secure image classification

How to cite this article: Chaudhary S, Dhotre S, Patil T. Optimal Privacy Preservation of Structured and Unstructured Datasets in Machine Learning: A Homomorphic Encryption-Based Framework for Secure Healthcare Data Analytics. Int J Drug Deliv Technol. 2026;16(32s):1108-1118. DOI: 10.25258/ijddt.16.32s.123

Introduction

The growth of computer-based reasoning in the healthcare sector has tremendously increased exponentially that it has become a power capable of unknown potentials in the areas of disease forecasting, image recognition, and instantaneous diagnostic assistance. ML and DL procedures have succeeded in significant improvements in diagnostic accuracy and efficiency by using structured clinical data and unstructured image datasets. However, given that healthcare applications deal with extremely sensitive patient information, ensuring data privacy has become a major concern in the deployment of AI systems.

Conventional encryption techniques like AES (Advanced Encryption Standard) ensure the data safety during the storage and the transmission of data. However, traditional encryption methods are not enough to ensure secure computation. This limitation results in a privacy-utility dilemma, particularly in the cases when it is necessary to reveal some information about the data during the training and the inference of algorithms. HE via the fact that mathematical operations on the ciphertext result precipitate the same operations on the corresponding plaintext, is the technology that solves this problem. HE thus emerges as a promising solution to facilitate secure machine learning and deep learning in confidentiality-sensitive areas such as healthcare.

At the same time, the limitations of HE, namely, two-fold overhead in computation, compatibility with previous models, and potential accidental performance degradation based on the performance of prediction are the reasons, that HE based ML/DL have remained limited to the research field only and have not been actually implemented in the workflows thereof. What is more, there is a gap in research in the area of extensive practical studies measuring the efficiency of various HE methods used for the different models and data types.

1.1 Research Contributions and Novelty

In order to fill these gaps, this paper suggests and assesses a privacy-preserving AI framework for both structured and unstructured medical datasets that makes use of homomorphic encryption (HE). Below is a summary of our work's main contributions:

- Multi-level Evaluation Framework: We create a single experimental setup to assess how AES and three HE schemes (CKKS, BFV, and RSA) affect different models for DL and ML in two domains: image-based kidney classification and structured clinical datasets (chronic kidney disease).
- Cross-Domain Validation: Our study demonstrates consistent privacy-preserving model performance in both structured data analysis and image classification under encryption constraints, which is a relatively unexplored angle in the current literature.

*Author for Correspondence: *sdchaudhary@bvucoep.edu.in*

- Sturdy Performance Preservation: Our models, especially Decision Tree, CNN, and VGG19, maintain high accuracy (up to 100%) in spite of encryption, demonstrating that HE can protect data without compromising predictive effectiveness.
- Computational Efficiency with Optimised Key Management: We examine encryption and decryption times and show that key sizes can be significantly decreased (CKKS key, for example, from 689.08 KB to 199.32 KB) without affecting model results.
- Model Behaviour Analysis Post Encryption: We confirm that the majority of models maintain stable behaviour, with InceptionV3 demonstrating reasonable trade-offs, by offering classification reports and confusion matrices for all scenarios (original, AES-encrypted, and HE-encrypted).
- Practical Deployment Insights: Our research offers useful information about incorporating HE into ML/DL pipelines, recommending workable compromises for cloud-based or real-time healthcare applications.

In contrast with previous research that either concentrates on a single encryption scheme or restricts its analysis to structured data, our study provides a comprehensive viewpoint on the incorporation of privacy-enhancing methods in various AI models. This work provides a safe, scalable framework for healthcare analytics that protects data privacy without compromising AI functionality.

Literature Review

Because medical data is sensitive and regulatory requirements are growing, Integration has garnered a lot of interest privacy-preserving technologies into healthcare AI systems. To improve the confidentiality, integrity, and usefulness of AI-driven disease prediction and diagnostic systems, researchers have investigated a number of strategies, including secure multi-party computation, fuzzy logic, federated learning, and HE. The notable contributions covered in this review of the literature include deep learning adaptations for encrypted data environments, cloud-IoT integrated solutions, HE-based secure frameworks, and hybrid reasoning models.

A Privacy-Aware Disease Prediction Support System (PDPSS) based on hybrid reasoning was first proposed by Malathi D. et al. [1]. The synergistic advantages of integrating case-centered reasoning, k-nearest neighbour, and fuzzy set theory led to the development of better calculation results. In order to protect patients' sensitive information from unauthorised user access, the Disease Prediction Support System (DPSS) was enlarged to create the PDPSS. The experimental findings showed that PDPSS performed better than the other models in terms of prediction security and accuracy. The prediction system model was assessed using statistical evaluation metrics. The system produced positive outcomes. However, the system's processing and communication costs were high.

EPDP, a useful and private technology for e-healthcare that forecasts the likelihood of illness occurrence, was presented by Xue Yang et al. [2]. The two phases of sickness risk prediction that the EPDP successfully finished were disease design training and disease prediction with assured privacy protection. During the training stages of disease design, super-augmenting sequences were combined with homomorphic cryptography techniques to reliably obtain the set of symptoms for each disease. The forecast results were computed using the bloom filter technique during the illness risk prediction stage. Comprehensive performance assessments have also demonstrated that the system has achieved remarkable gains in processing and communication cost efficiency. The system had trouble approving some doctors, which made it hard to gain authority within the system.

Using IoT and cloud technology, Priyan Malarvizhi Kumar et al. [3] presented a healthcare plan that emphasises illness diagnosis and prediction. The method makes use of a fuzzy neural classifier. In this instance, diabetes was managed methodically. The related medical data generated to predict who will suffer from severe diabetes focused on medical sensors and the UCI Repositories dataset. In order to distinguish between the states of illness and serenity, the system also employed a fuzzy rule-based neural classification system. The studies make use of a standard dataset from the UCI repository in addition to the entire collection of medical records gathered from various institutions. The trial's outcomes showed that the technology performed better in predicting illness than the current method. Nevertheless, the cloud database's medical data was not adequately protected.

PPDP, an effective and privacy-preserving disease calculation system, was suggested by C. Zhang et al. [4]. PPDP protects patient privacy by encrypting and sending patients past medical records to a cloud server, where they can be used to create Single-Layer Perceptron learning prediction models.

Addressing data security issues requires efficient administration of electronic medical histories (HER). In the study, K. Munjal compared the HE and FHE systems and administered a survey. Initially, the work focused on the core ideas of HE, namely Somewhat Homomorphic Encryption (SHE) and Partial Homomorphic Encryption (PHE). Fully Homomorphic Encryption (FHE) necessitates many ideas. To determine the primary strategies within each of the four major categories into which the FHE processes were separated, a comparative analysis was carried out. A range of methods were collected and compared in the domains of bioinformatics and healthcare in order to accurately identify cardiovascular issues, typical heart rate, cancer, LQTC and the creation of a trustworthy healthcare query generation system sector. In conclusion, HE will be very beneficial to the health care industry due to the enhanced functionality and performance of FHE [5].

Ahmed EL-YAHYAOUI's study assessed the state-of-the-art in FHE systems. According to the research, there are several frameworks available for creating an FHE scheme. There is still much to be done to move fully homomorphic encryption (FHE) from theory to practice, despite the existence of several frameworks and the efforts of cryptologists to create a workable scheme. Boosted cryptanalysis of FHE algorithms is one way to

improve the design of FHE schemes and increase their runtime [6].

Several terms related to HE have been simplified in the research paper written by Majedah Alkharji et al. The current state of the art has been rigorously evaluated and presented, and the role of HE in the applications that are currently in use has been carefully investigated. Although HE techniques improve cloud computing's benefits, like customer satisfaction and data security, it's crucial to address its speed and capacity to manage large amounts of data. Therefore, more research is needed to improve these tactics and increase HE's efficacy. Practically speaking, the focus should be on developing more effective strategies [7].

A practical approach was presented by R. Shokri et al. that enables multiple participants to work together to jointly obtain an accurate neural-network model for a particular objective without revealing their input datasets. The author exploited the possibility that optimisation methods used in modern deep learning, namely stochastic gradient descent-based methods, could be parallelised and run asynchronously. During training, the suggested approach enables participants to selectively swap small parts of their models' critical parameters and train independently using their own datasets. This offers a compelling feature in the trade-off between privacy and usefulness: people can benefit from other participants' models while still protecting the privacy of their own data. Their learning accuracy is consequently improved beyond what they could accomplish using only their own inputs. Using benchmark datasets, the author verified the accuracy of our privacy-preserving deep learning method [8].

DPP, a cloud computing method that uses HE to protect data privacy, was introduced by Jing Wang and associates. Accuracy, compatibility, and security are guaranteed by this scheme. To preserve data privacy, the DPP framework uses HE. After building DPP using the Paillier HE technique, the author assessed its performance to ascertain its security. The experimental results demonstrate that the critical phases of the DPP scheme have realistic and reasonable durations [20].

A. Wood et al. provided a summary of the application of FHE in bioinformatics and medicine. Along with detailed information on current open-source implementations, the basic ideas and background of FHE were covered. A detailed analysis of FHE was conducted in order to safeguard privacy in bioinformatics and machine learning. The review also included explanations of how these strategies are actually applied in the encrypted domain [9].

First introduced by T.N. Van et al., HE preserves privacy while allowing deep learning models to process the encrypted data. This approach is applicable to almost all digital healthcare services, where data providers want to ensure that their data is not used without authorisation. The proposed encryption model's accuracy was only 0.01% different from the non-encryption models. Furthermore, the suggested model proved to be effective in practice [10].

A mobile application that moves user data to the cloud was developed by S. Carpov et al., along with a HE method that securely processes the data without disclosing any information to the cloud provider. The practical tests' outcomes demonstrate that the HE algorithm could review users' data in a respectable period of time, demonstrating the viability of this novel method for evaluating private data [11].

Dan Zhu et al. [12] introduced CREDO, a multi-level medical pre-diagnosis system that uses the ML-kNN (multiple-label k-nearest-neighbors) algorithm to accomplish both effectiveness and privacy protection. The service provider (SP) initially employed k-means clustering to lower the quantity of healthcare cases that required computation. The SP then used ML-kNN classification to provide services to healthcare consumers. The query vector was handled directly in the SP and encrypted before being delivered. Only the pre-diagnosis result was available to the medical user at the same time. The comprehensive analysis revealed that CREDO had substantially less computational complexity than the other system and could withstand a number of well-known security flaws.

N.N. Thilakarathne et al. [26] provided a comprehensive solution in their review [13] to address the challenge of learning about Medical IoT (MIoT) without transferring sensitive and private information to a central cloud. For this reason, they suggested federated learning (FL) as a viable option.

A suggested approach in [14] proposes to predict diseases by combining cloud computing with an Internet of Things (IoT) database. This approach predicts the occurrence of diseases by using patient data obtained from biosensors. Two approaches were put forth for the prediction: a classifier called GFIbALO, which is based on generalized fuzzy intelligence, and a regression technique.

Researchers N.D. Kathamuthu et al. [15] created the DQ-NNPP neural network architecture, which employs deep Q-learning to protect confidential patient medical data transmitted from medical IoT devices from external attacks. All of these approaches fall short in terms of data security and confidentiality. Moreover, all of the previously mentioned methods frequently have poor prediction accuracy and efficacy.

Common HE techniques can now work with real-valued numbers of any size and precision thanks to an encoding technique developed by A.B. Popescu et al. EEG data from two real-world scenarios idetifying seizures and predicting alcoholism susceptibility were used to evaluate the methodology. A direct (non-iterative) fitting method that requires a predefined and deterministic number of steps is used for training in a supervised machine learning methodology. Artificial data sets of various sizes and complexity levels were tested to determine how they affected execution time and error accumulation. While the inference time stays within the range of milliseconds, the models' training time increases but stays within reasonable bounds [16].

An inventive encryption technique developed by Kaushik Sinha and his associates guarantees anonymity when processing and storing sensitive medical data. By using a completely HE scheme that was safe, the proposed

method made it possible to perform computations using encrypted data that don't need decryption. Current public health monitoring systems are unable to query or compute encrypted data without the need for decryption. Using the COVID-19 pandemic as a case study, the author developed an algorithm for contact tracing and presented a novel computer model. The ElGamal encryption algorithm was used to simulate the proposed method in order to assess its accuracy and efficacy. The results show that the proposed method effectively met the computing requirements for contact-tracing while guaranteeing adequate security [17].

The Crypto tree framework, developed by Daniel Huynh[18], enables the application of Random Forests (RF) to HE. Especially when contrasted with linear regression, Random Forests are a very powerful learning method. The author achieved this by first converting a standard RF into a neural RF and then modifying it to conform to the HE schemes CKKS, which allow HE operations on real values. Using SIMD operations, the author's quick inference and prediction results beat the original RF on encrypted data.

S. Chaudhary et. al [19] offer an extensive overview of fully homomorphic encryption, covering its theoretical basis, recent advancements, existing challenges, practical use cases, and the tools currently available. It further highlights how integrating FHE with machine learning can support effective predictive modeling while ensuring strong protection of patient data privacy.

It is clear from the reviewed literature that although many methods show promise in terms of prediction accuracy and privacy protection, there are still a number of obstacles to overcome, especially when it comes to managing encryption overhead, attaining real-time performance, and making sure that deep learning frameworks are compatible. Though it still has a lot of promise, HE especially fully HE, or FHE has real-world drawbacks like scalability and computational cost. These shortcomings underline the necessity of privacy-preserving frameworks that are portable, effective, and flexible enough to be used in both structured and unstructured medical data scenarios, as this study suggests.

Methodology

The thorough approach used to examine the incorporation of privacy-preserving encryption methods with DL and ML models in healthcare applications is presented in this section. The process includes preparing the dataset, choosing a model, putting encryption schemes into practice, and assessing computational trade-offs and performance. Figures 1 and 2 show the system architecture of the implementation using DL and ML models on csv and image datasets.

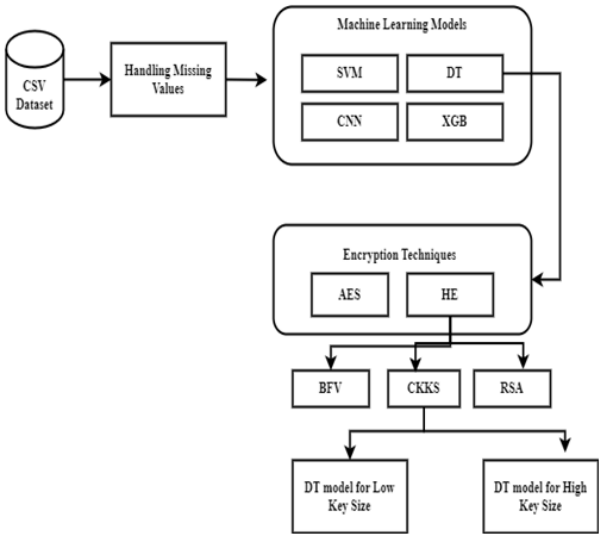

Fig. 1 System Architecture for CSV Based Dataset

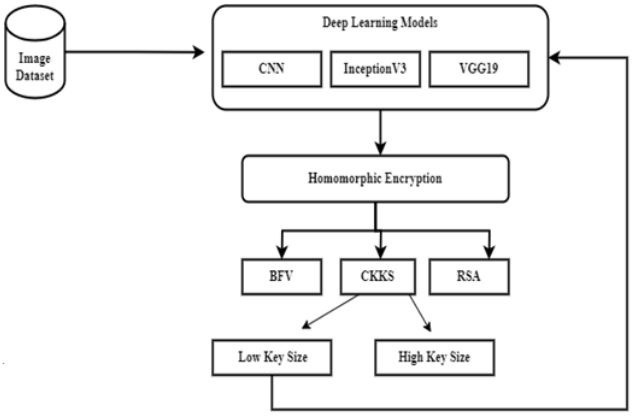

Fig. 2 System Architecture for Image Based Dataset

1.1 Dataset Description

Two different dataset types an image-based dataset and a structured tabular dataset in CSV frmat were used to evaluate the effect of encryption on predictive performance. Here is a description of the dataset.

1.1.1 CSV Based Dataset Description

For the structured data experiments, we used the chronic kidney disease (CKD) dataset, which is publically available on Kaggle and was originally sourced from the UCI ML Repository. Data preprocessing involved turning missing entries denoted by "?" into null values, followed by imputation using median values for numerical features and designated categories for categorical ones. This dataset includes 24 clinical attributes, ranging from demographic and biochemical measures to comorbidity indicators (e.g., haemoglobin, serum creatinine, hypertension). The dataset consists of 400 patient records that were gathered over a two-month period in India. In order to guarantee a representative distribution of CKD and non-CKD cases, after cleaning, the dataset was split into 120 testing samples and 277 training samples at random.

1.1.2 Image Based Dataset Description

The CT Kidney Dataset: Normal, Cyst, Tumour, and Stone, which is openly accessible on Kaggle, was used in the image-based kidney classification experiments. There are 12,446 CT scan images in the dataset, which are divided into four classes: Cyst (3,709), Tumour (2,283), Normal (5,077), and Stone (1,377). For model training, the DICOM-formatted images were resized to 224 x 224 pixels and converted to JPEG.

We ensured a balanced class distribution in our study by using training uses 80% of the dataset, whereas testing uses 20%. To improve generalisation, common preprocessing methods were used, such as normalisation and augmentation (rotation, zooming, brightness adjustment). Because every image was verified and annotated by medical experts, the dataset is trustworthy for classification tasks.

1.2 Machine and Deep Learning Models

Both traditional ML classifiers and sophisticated DL architectures were used to create baseline predictive performance and examine the effect of encryption. In the structured data area, three well-known machine learning models SVM, Decision Tree, and XGBoost were used. Because of its resilience in feature rooms with great dimensions, the SVM model with a linear kernel was employed. As a rule-based, non-parametric classifier, Decision Tree provided tabular data with interpretable classification. The ensemble-based gradient boosting method XGBoost was chosen because it can handle noise and non-linearity and perform well in structured learning tasks.

Three convolutional architectures a basic Convolutional Neural Network (CNN), VGG19, and InceptionV3 were employed for image-based classification. In order to learn low- to mid-level visual patterns, the CNN model was created using max-pooling, two convolutional layers, and fully linked layers. The kidney image dataset was used to refine VGG19, a deep convolutional network pre-trained on ImageNet, in order to take advantage of its deep hierarchical feature extraction capabilities. Known for its factorised convolutions and inception modules, InceptionV3 was also optimised to evaluate performance on challenging multi-class image classification tasks.

1.3 Privacy-Preserving Encryption Schemes

This study's main goal was to assess how well ML/DL models performed and how resilient they were when trained on encrypted data. In order to achieve this, two different encryption schemes were included: HE, which preserves privacy asymmetrically, and Advanced Encryption Standard (AES), which provides symmetric encryption.

The CKD dataset was first secured using AES before model training. Since decryption was necessary before feeding data to the model, AES's limitations became apparent during training, despite the fact that it effectively provided data confidentiality during storage and transmission. As a result, the AES-based configuration was mainly used as a baseline privacy-preserving method.

We combined CKKS, BFV, and RSA, three HE schemes, to allow calculation on material that has been encrypted without decryption. Because it enables approximation arithmetic operations on real numbers, the Cheon-Kim-Kim-Song (CKKS) technique is suitable for DL applications where floating-point processing is prevalent.

Tree-based models were trained on encrypted structured data using Brakerski/Fan-Vercauteren (BFV), which supports exact integer arithmetic. The overhead and comparative performance of RSA, a traditional public-key cryptosystem, were also assessed. A Python-compatible HE library that facilitates encryption and evaluation was used to implement each encryption scheme. During the experiments, encryption and decryption times were recorded, and key sizes were set to strike a balance between security and performance.

1.4 Evaluation and Metrics

A set of established classification metrics was used to thoroughly assess each ML and DL performance of the model with and without encryption. These were calculated on both training and testing sets and included accuracy, F1-score, recall, and precision. Confusion matrices created then analysed in order to examine model robustness and class-wise prediction in more detail. To understand the overhead caused by encryption, computational metrics like encryption time, decryption time, and total training time were measured in addition to predictive performance. The viability of HE schemes in resource-constrained environments was also assessed by analysing key size differences before and after optimisation.

1.5 Implementation Setup

Every experiment was carried out in a Jupyter Notebook environment using Python. The scikit-learn and XgBoost libraries were used to implement the ML models, and Kera with a TensorFlow backend was used to build the DL models. Python's native cryptographic toolkit was used to integrate RSA encryption, and HE libraries compatible with CKKS and BFV schemes were used to incorporate HE functionalities. The experimental system consisted of an NVIDIA GPU to expedite deep learning model training, an Intel Core i5 processor, and 16 GB of RAM. This setup ensured consistency in experimental procedures by conducting data encryption, model training, and performance evaluation end-to-end.

Results And Discussion

A thorough examination of the experimental findings from testing different ML and DL representations on both original then encrypted datasets is provided in this section. The assessment is divided into two sections: (i) structured dataset analysis with data on chronic kidney disease, and (ii) image dataset analysis with kidney classification tasks. We evaluate performance of the model for every dataset in three different scenarios: unencrypted, after AES encryption, and after using HE schemes (CKKS, BFV, RSA).

1.6 Performance on Structured Dataset (CKD CSV)

1.6.1 Baseline Model Performance (Unencrypted Data)

On the unencrypted CKD dataset, the Decision Tree and XgBoost classifiers' training and testing sets both attained perfect accuracy (100%) in both cases. Although SVM trained substantially quicker (0.02 seconds), its test accuracy was rather low at 69%. This suggests that the dataset's non-linear patterns are better captured by tree-based models. The machine learning model's train/test accuracy comparison on the unencrypted CKD dataset is displayed in Table 1.

Table 1. An Assessment of Machine Learning Models' Effectiveness

Model	Train Size	Test Size	Train Accuracy	Test Accuracy	Time
SVM	277	120	68	69	0.02
Decision Tree	277	120	100	100	0.01
XG Boost	277	120	100	100	0.74
CNN			98	95	3.14

For both the CKD and non-CKD classes, the Decision Tree model demonstrated optimal classification metrics with precision, recall, and F1-score values of 1.0, resulting in an overall test accuracy of 100%. Table 2 displays the Decision Tree model's classification report. Figure 3 illustrates how XgBoost replicated this performance with perfect classification reports and confusion matrices.

Table 2. Classification Report of DT Model

	Precision	Recall	F1-Score	Support
ckd	1	1	1	75
notckd	1	1	1	45
accuracy			1	120
macro avg	1	1	1	120
weighted	1	1	1	120

avg				

SVM's limitations in this medical dataset with mixed-type features and class imbalance, however, are reflected in its moderate classification scores (precision = 0.76 for CKD and 0.58 for not-CKD).

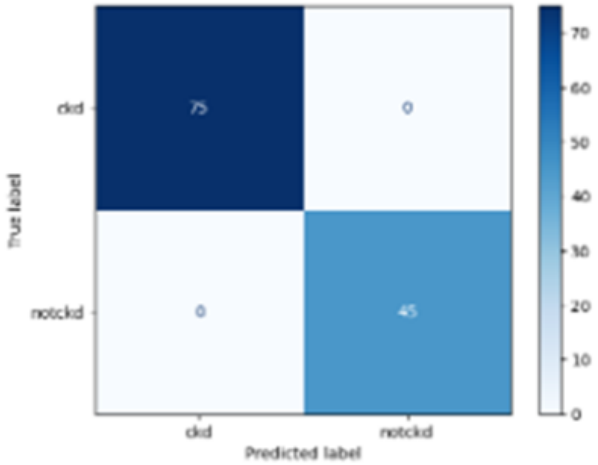

Fig. 3 Confusion Matrix of DT

1.6.2 AES-Encrypted Data

The model's performance slightly declined when AES encryption was used. The Decision Tree classifier obtained a test accuracy of 96% while maintaining a high training accuracy of 100%. Figure 4 displays the confusion matrix, and Table 3 highlights the Decision Tree Model's classification report on encrypted data. Precision and recall scores showed slight declines, especially for the non-CKD class (precision = 0.94, recall = 0.97). These findings demonstrate that it is possible to train models on symmetrically encrypted data, even though decryption before training still poses a privacy risk.

Table 3. DT Classification Report on AES-Encrypted Data

	Precision	**Recall**	**F1-Score**	**Support**
ckd	0.98	0.96	0.97	50
notckd	0.94	0.97	0.95	30
accuracy			0.96	80
macro avg	0.96	0.96	0.96	80
weighted avg	0.96	0.96	0.96	80

1.6.3 HE Results

With a test accuracy of 98%, the Decision Tree classifier maintained strong performance when HE specifically, the CKKS scheme was integrated. The classification report's outstanding class-wise metrics (CKD: F1-score = 0.99, not-CKD: F1-score = 0.98) verified minimal performance loss. Table 4 shows the HE Scheme's DT classification report, and Figure 5 shows the HE Scheme's DT confusion matrix.

Table 4. Decision Tree Classification Report on the HE Scheme

	Precision	**Recall**	**F1-Score**	**Support**
ckd	0.98	1	0.99	50
notckd	1	0.97	0.98	30
accuracy			0.99	80
macro avg	0.99	0.98	0.99	80
weighted avg	0.99	0.99	0.99	80

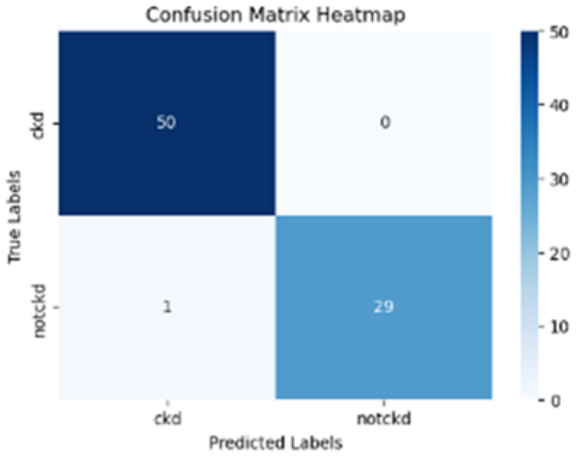

Fig. 4 Confusion Matrix of DT on AES-Encrypted Data

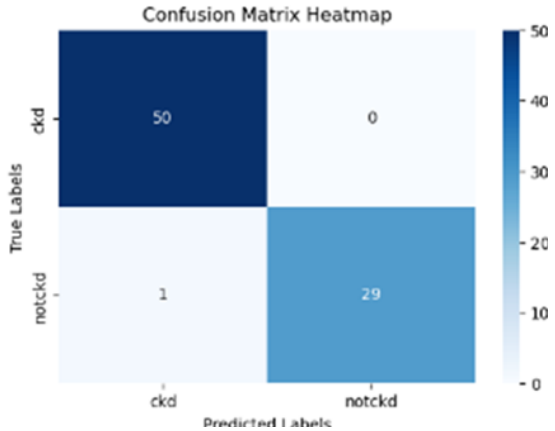

Fig. 5 Confusion Matrix of DT on HE Scheme

Table 5. Encryption Time Comparison for HE Schemes During Training and Testing Phases

Schemes	Train Enc Time	Test Enc Time
CKKS	5.84	1.09
BFV	44.64	9.19
RSA	49.85	13.07

Table 6. Key Size Optimization in CKKS HE Scheme

Key Size Reduction in CKKS	
Before	689.08 kb
After	199.32 kb

Table VI displays the key size optimisation in CKKS HE schemes, while Table V compares the encryption time for, HE schemes during the training and testing stages of the model. CKKS showed the optimal ratio between computational efficiency and model compatibility among the evaluated HE schemes. The CKKS scheme was significantly more efficient than the other techniques, requiring 5.84 seconds for training encryption and 1.09 seconds for test-time encryption. RSA turned out to be

the most computationally costly, requiring 49.85 seconds for training encryption and 13.07 seconds for testing encryption, while BFV required 44.64 seconds for training encryption and 9.19 seconds for testing encryption. These findings confirm that CKKS is the most sensible option for situations that call for frequent encryption and real-time processing. Additionally, CKKS's key size optimisation decreased the encryption key's size from 689.08 KB to 199.32 KB, greatly reducing memory overhead without sacrificing performance. This provides more evidence that CKKS is appropriate for implementing privacy-preserving machine learning systems in the real world.

1.7 Performance on Image Dataset (Kidney Image

Classification)

1.7.1 Baseline Model Performance (Unencrypted Data)

Both CNN and VGG19 models performed flawlessly on the unencrypted image dataset, achieving 100% accuracy on both training and testing samples. According to their classification reports, complete diagonal confusion matrices supported F1-scores of 1.0 for each of the four classes (tumour, normal, stone, and cyst).

Although the InceptionV3 model achieved 100% accuracy on the test set, it performed poorly during the training phase, achieving only 87.5% accuracy during training. Overfitting or inadequate layer fine-tuning on the small dataset is probably the cause of this discrepancy.

1.7.2 HE Results on Image Data

Both CNN and VGG19 maintained their 100% accuracy across training and testing sets when the HE schemes were applied to the image dataset (with an emphasis on CKKS). The fact that encryption had no effect on their precision, recall, or F1-score values shows how well CKKS works with DL architectures that use floating-point processes. Table 8 shows the classification report for the CNN model, while Table 7 shows the performance analysis of the DL model on the picture dataset. In Figure 6, the confusion matrix is shown.

Table 7. Performance Analysis of DL Model on Image Dataset

Model	Train Count	Test Count	Train Accuracy	Test Accuracy
CNN	80	20	100	100
VGG19	80	20	100	100
InceptionV3	80	20	87.5	100

Table 8. Classification Report of CNN Model

CNN	precision	recall	f1-score	support
cyst	1	1	1	5
normal	1	1	1	5
stone	1	1	1	5
tumour	1	1	1	5
accuracy			1	20
macro avg	1	1	1	20
weighted avg	1	1	1	20

On the other hand, InceptionV3's performance declined noticeably when encrypted. It demonstrated a test accuracy of 85% and a reduced training accuracy of 83.75%. All four classes had lower scores, with F1-scores averaging about 0.25, according to detailed metrics. Table 9's findings imply that intricate designs such as InceptionV3 are more susceptible to approximation errors introduced during encrypted computation.

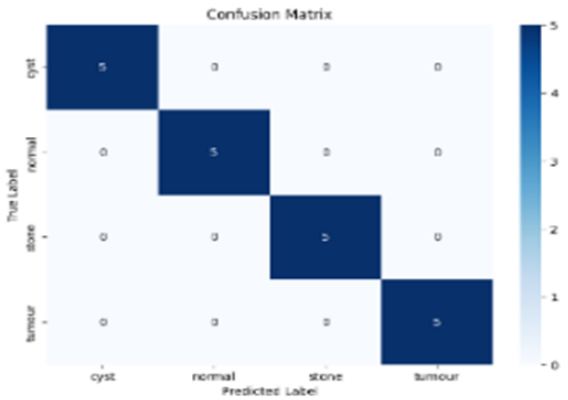

Fig. 6. Confusion Matrix of CNN Model

Table 9. Performance Analysis of DL Model on HE Scheme

Model	Train Count	Test Count	Train Accuracy	Test Accuracy
CNN	80	20	100	100
VGG19	80	20	100	100
InceptionV3	80	20	83.75	85

Additionally, an analysis of encryption times revealed that RSA once again had the highest encryption times (training: 49.85 seconds), while CKKS was the fastest of the HE schemes (training encryption: 5.84 seconds, testing: 1.09 seconds). With no change in classification accuracy, the initial key size of 689.08 KB for CKKS was reduced to 199.32 KB after optimisation, indicating a significant reduction in key size.

1.8 Comparative Insights and Practical Implications

The experimental findings provide compelling evidence for the applicability of HE, particularly CKKS, in privacy-preserving artificial intelligence applications. CNNs and decision trees performed consistently on both original and encrypted datasets, making them the most resilient models under encryption. AES offers a rudimentary privacy layer, but HE ensures end-to-end confidentiality by enabling model training and inference directly on encrypted data.

CKKS is unique from a deployment perspective because it strikes a balance between low latency, model compatibility, and a smaller storage footprint. RSA is less appropriate for high-throughput applications and has considerable computational delays, despite its security.

These results highlight the importance of carefully choosing an encryption scheme according to the data type, model architecture, and intended use. The suggested framework demonstrates that strong privacy can be attained without compromising precision or effectiveness, which is essential in delicate fields like cloud-based medical diagnostics.

CONCLUSION AND FUTURE WORK

In order to incorporate HE into structured and image-based healthcare classification tasks, this paper proposes a privacy-preserving machine learning framework. We show that secure model training and inference can be accomplished with low accuracy loss by implementing CKKS, BFV, and RSA schemes across a variety of ML and DL models. Even with encrypted data, CNN and Decision Tree models continuously performed well. Because it strikes a balance between accuracy preservation and computational efficiency, CKKS emerged as the most viable HE technique for real-time applications. All things considered, the suggested framework guarantees robust privacy without sacrificing prediction accuracy.

To further improve data decentralisation and privacy, future developments of this work might incorporate HE with federated learning. The scalability of the framework can be confirmed by more tests using bigger and more varied medical datasets, like MRI and EHR. Additionally, integrating explainable AI (XAI) techniques and investigating lightweight DL architectures (such as MobileNet and EfficientNet) may enhance interpretability in cloud or mobile healthcare settings while preserving safe, real-time performance.

Conflicts of Interest

The author(s) declare(s) that there is no conflict of interest regarding the publication of this paper.

REFERENCES Bottom of Form

[1] D. Malathi, R. Logesh, V. Subramaniyaswamy, V. Vijayakumar, and A. K. Sangaiah, "Hybrid Reasoning-based Privacy-Aware Disease Prediction Support System," *Computers & Electrical Engineering*, vol. 73, pp. 114–127, Jan. 2019, doi: 10.1016/J.COMPELECENG.2018.11.009.

[2] X. Yang, R. Lu, J. Shao, X. Tang, and H. Yang, "An efficient and privacy-preserving disease risk prediction scheme for E-healthcare," *IEEE Internet Things J*, vol. 6, no. 2, pp. 3284–3297, Apr. 2019, doi: 10.1109/JIOT.2018.2882224.

[3] P. M. Kumar, S. Lokesh, R. Varatharajan, G. Chandra Babu, and P. Parthasarathy, "Cloud and IoT based disease prediction and diagnosis system for healthcare using Fuzzy neural classifier," *Future Generation Computer Systems*, vol. 86, pp. 527–534, Sep. 2018, doi: 10.1016/J.FUTURE.2018.04.036.

[4] C. Zhang, L. Zhu, C. Xu, and R. Lu, "PPDP: An efficient and privacy-preserving disease prediction scheme in cloud-based e-Healthcare system," *Future Generation Computer Systems*, vol. 79, pp. 16–25, Feb. 2018, doi: 10.1016/J.FUTURE.2017.09.002.

[5] K. Munjal and R. Bhatia, "A systematic review of homomorphic encryption and its contributions in healthcare industry," *Complex & Intelligent Systems*, vol. 9, no. 4, p. 1, Aug. 2023, doi: 10.1007/S40747-022-00756-Z.

[6] A. El-Yahyaoui and M. Dafir, "Fully Homomorphic Encryption: State of Art and Comparison New cryptographic methods for big data and cloud computing View project," no. January 2016, 2016, doi: 10.6084/M9.FIGSHARE.3362338.

[7] M. Alkharji and H. Liu, "Homomorphic Encryption Algorithms and Schemes for Secure Computations in the Cloud," no. March, p. 19, 2016.

[8] R. Shokri, A. Houmansadr, and M. Nasr, "Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning

Machine Learning Privacy View project Privacy-Preserving Data Synthesis View project Comprehensive Privacy Analysis", doi: 10.1109/SP.2019.00065.

[9] A. Wood, K. Najarian, and D. Kahrobaei, "Homomorphic Encryption for Machine Learning in Medicine and Bioinformatics," *ACM Computing Surveys (CSUR)*, vol. 53, no. 4, Aug. 2020, doi: 10.1145/3394658.

[10] T. Nguyen-Van et al., "A Homomorphic Encryption Approach for Privacy-Preserving Deep Learning in Digital Health Care Service," Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 13758 LNAI, pp. 520–533, 2022, doi: 10.1007/978-3-031-21967-2_42.

[11] S. Carpov, T. H. Nguyen, R. Sirdey, G. Constantino, and F. Martinelli, "Practical Privacy-Preserving Medical Diagnosis Using Homomorphic Encryption," pp. 593–599, Jan. 2017, doi: 10.1109/CLOUD.2016.0084.

[12] D. Zhu et al., "CREDO: Efficient and privacy-preserving multi-level medical pre-diagnosis based on ML-kNN," Inf Sci (N Y), vol. 514, pp. 244–262, Apr. 2020, doi: 10.1016/J.INS.2019.11.041.

[13] N. N. Thilakarathne et al., "Federated Learning for Privacy-Preserved Medical Internet of Things," 2022. doi: 10.32604/iasc.2022.023763.

[14] A. Verma, G. Agarwal, and A. K. Gupta, "A novel generalized fuzzy intelligence-based ant lion optimization for internet of things-based disease prediction and diagnosis," Cluster Comput, vol. 25, no. 5, pp. 3283–3298, Oct. 2022, doi: 10.1007/S10586-022-03565-8/FIGURES/14.

[15] N. D. Kathamuthu, A. Chinnamuthu, N. Iruthayanathan, M. Ramachandran, and A. H. Gandomi, "Deep Q-Learning-Based Neural Network with Privacy Preservation Method for Secure Data Transmission in Internet of Things (IoT) Healthcare Application," Electronics 2022, Vol. 11, Page 157, vol. 11, no. 1, p. 157, Jan. 2022, doi: 10.3390/ELECTRONICS11010157.

[16] A. B. Popescu et al., "Privacy Preserving Classification of EEG Data Using Machine Learning and Homomorphic Encryption," Applied Sciences 2021, Vol. 11, Page 7360, vol. 11, no. 16, p. 7360, Aug. 2021, doi: 10.3390/APP11167360.

[17] K. Sinha, P. Majumder, and S. K. Ghosh, "Fully Homomorphic Encryption based Privacy-Preserving Data Acquisition and Computation for Contact Tracing," International Symposium on Advanced Networks and Telecommunication Systems, ANTS, vol. 2020-December, Dec. 2020, doi: 10.1109/ANTS50601.2020.9342834.

[18] D. Huynh, "Cryptotree: fast and accurate predictions on encrypted structured data," Jun. 2020.

[19] S. Chaudhary, S. Dhotre, and T. Patil, "Chronic Kidney Disease Prediction: A Study of Encrypted Datasets", Signal Processing, Telecommunication & Embedded Systems: AI and ML Applications, ICMEET 2024, Lecture Notes in Electrical Engineering, vol. 1430, Singapore: Springer, 2025.