Enhancing Privacy in Medical Data Classification with Homomorphic Encryption

PURDUE UNIVERSITY FORT WAYNE

Fairuz Haq

Advisor: Dr. Chao Chen

Department of Electrical and Computer Engineering

Introduction

Machine learning (ML) holds great potential for disease prediction and medical data classification. However, its reliance on sensitive patient data raises privacy concerns, further complicated by ethical guidelines and regulations.

Homomorphic Encryption (HE) offers a solution by allowing computations on encrypted data without the need for decryption, thus ensuring privacy throughout the processing. Unlike traditional encryption methods, which expose data to security risks during computation, HE enables secure analysis while maintaining confidentiality and regulatory compliance [1].

This research study examines the impact of the Cheon-Kim-Kim-Song (CKKS) HE scheme [2] on ML models, specifically Logistic Regression (LR) and Support Vector Machines (SVM), using the TenSEAL library [3]. By comparing the performance of learning on encrypted and unencrypted data, we evaluate trade-offs in accuracy, computational efficiency, and privacy, assessing the feasibility of HE for secure medical data analysis.

Methodology

We simulate the following scenario: The hospital retrieves a medical dataset from a database containing non-sensitive and publishable data. The selected dataset is then provided as plaintext to the ML service provider for model training. Meanwhile, privacy-sensitive current patient data is encrypted using CKKS HE scheme and supplied as a test dataset to the ML service provider, as depicted in Figure 1.

The ML service provider first trains the model on plaintext training data to achieve optimal learning. Since CKKS HE supports only addition and multiplication, the model is subsequently adjusted to process the encrypted test data. By using the encrypted test data without decryption, patient privacy is preserved. The ML service provider then returns encrypted predictions to the hospital, which decrypts them to obtain the final disease prediction results.

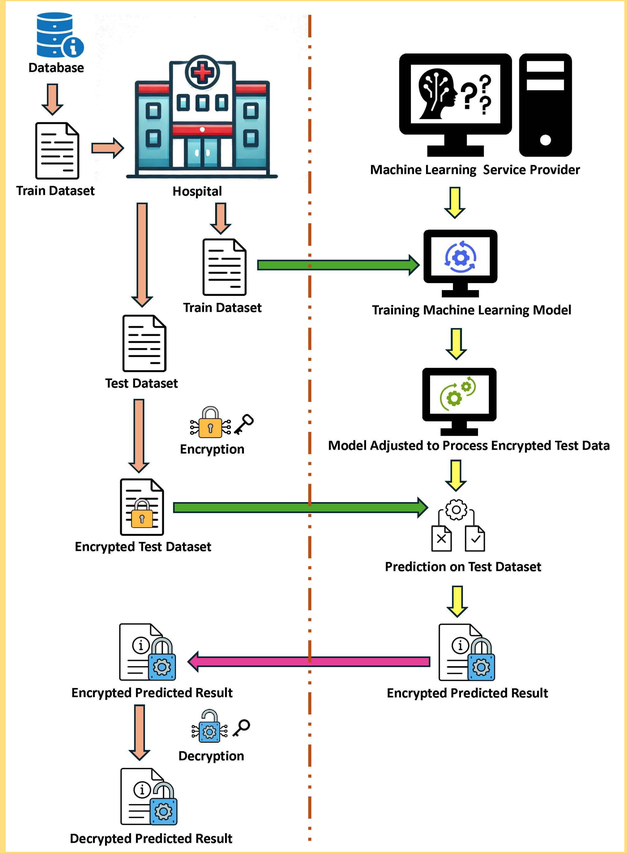

Figure 1. Methodology Overview

Implementation

The Wisconsin Diagnostic Breast Cancer (WDBC) dataset [4], comprising 569 samples, was used for binary classification to predict tumors as benign (0) or malignant (1) based on 30 features. The dataset was split into 90% for training (513 samples) and 10% for testing (56 samples). Logistic Regression and SVM machine learning models were employed, evaluated with and without CKKS. Training was conducted on unencrypted data, with 10-fold cross-validation to ensure robustness. The CKKS scheme was chosen to handle floating-point data, and experiments were run on Google Colab+ with an A100 GPU. Testing involved plaintext and encrypted evaluations. The trained model was adjusted to process encrypted inputs by prioritizing linear functions and extracting model weights and biases to ensure HE compatibility. Performance comparison included classification metrics (accuracy, precision, recall, F1-score), computational time differences, and decryption consistency to verify that decrypted predictions matched plaintext prediction results.

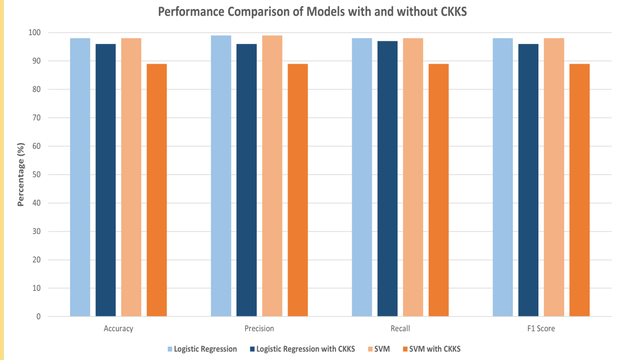

Figure 2. Performance Comparison of Classification Models with and without CKKS

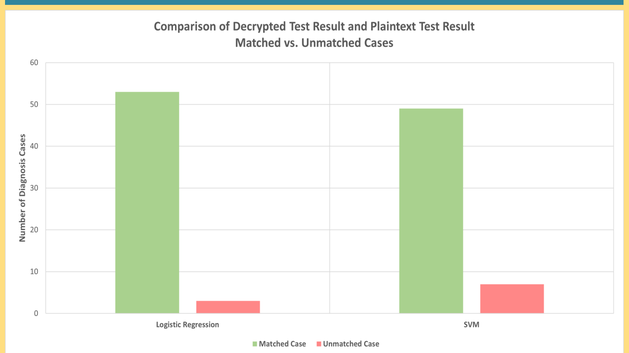

Figure 3. Comparison of Decrypted Test Result and Plaintext Test Result Matched vs. Unmatched Cases

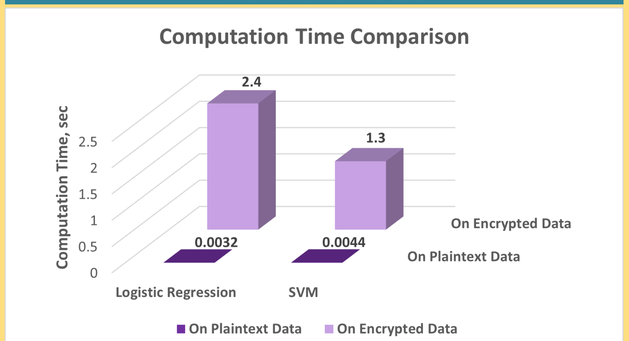

Figure 4. Computation Time Comparison of Classification Models with and without CKKS

Results

The impact of CKKS HE on model performance is illustrated in Figure 2. Both Logistic Regression and SVM models performed well with unencrypted test data. However, when the test data was encrypted with CKKS HE, both models experienced a slight decline in performance, indicating sensitivity to encryption overhead. Nevertheless, the results remained close to the performance with unencrypted test data.

Moreover, Figure 3 demonstrates that for both models, the majority of encrypted diagnostic predictions matched the unencrypted test results after decryption. This suggests that even with CKKS HE for privacy preservation, disease prediction closely aligns with the results obtained without encryption.

However, as shown in Figure 4, the computational time is significantly higher when CKKS encryption is introduced in the test data compared to using unencrypted test data for testing.

Conclusion

This research study demonstrates the effectiveness of HE in enabling privacy-preserving medical data classification without significantly compromising predictive accuracy. By evaluating two machine learning models under both encrypted and unencrypted conditions, we provide a comprehensive analysis of how encryption impacts model performance, including accuracy and computational efficiency. These findings highlight the potential of HE as a practical and scalable solution for secure, data-driven healthcare advancements.

In the future, this research study can be further extended to focus on reducing computational overhead, evaluating HE in more complex machine learning models, and exploring hybrid encryption techniques that combine HE with other privacy-preserving methods.

References & Acknowledgment

[1] Wood, A., Najarian, K., & Kahrobaei, D. (2020). Homomorphic encryption for machine learning in medicine and bioinformatics. ACM Computing Surveys (CSUR), 53(4), 1-35.

[2] Cheon, J. H., Kim, A., Kim, M., & Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. In Advances in cryptology–ASIACRYPT 2017: 23rd international conference on the theory and applications of cryptology and information security, Hong kong, China, December 3-7, 2017, proceedings, part i 23 (pp. 409-437). Springer International Publishing.

[3] Benaissa, A., Retiat, B., Cebere, B., & Belfedhal, A. E. (2021). Tenseal: A library for encrypted tensor operations using homomorphic encryption. arXiv preprint arXiv:2104.03152.

[4] Wolberg, W., Mangasarian, O., Street, N., & Street, W. (1993). Breast Cancer Wisconsin (Diagnostic) [Dataset]. UCI Machine Learning Repository. https://doi.org/10.24432/C5DW2B.

I would like to express my sincere gratitude to my advisor, Dr. Chao Chen, Professor in the Department of Electrical and Computer Engineering, for invaluable guidance, insightful feedback, and continuous support throughout this research.