

Protecting Patient Data in AI/ML Models with Homomorphic Encryption in Hybrid Cloud Environments: Enabling Privacy-Preserving Analytics Without Decryption

Charan Shankar Kummarapurugu

Senior Cloud Engineer Herndon, VA, USA

Email: charanshankar@outlook.com

Abstract

Healthcare applications increasingly rely on AI/ML models to analyze sensitive patient data, but the privacy of such data in cloud environments is a significant concern. This paper proposes a novel approach utilizing homomorphic encryption to protect patient data while enabling privacy-preserving AI/ML analysis in hybrid cloud environments. By using this encryption method, we can process encrypted data without the need for decryption, thus ensuring data privacy and compliance with regulations such as HIPAA and GDPR. We explore the implementation challenges, performance trade-offs, and present an architecture to integrate homomorphic encryption in AI/ML workflows across hybrid clouds. The results show that this approach can effectively secure patient data while maintaining the accuracy and efficiency of AI models.

Index Terms: Homomorphic Encryption, AI/ML Models, Hybrid Cloud, Privacy-Preserving Analytics, Healthcare, Patient Data Security

INTRODUCTION

The integration of Artificial Intelligence (AI) and Machine Learning (ML) into healthcare systems has revolutionized data analysis, particularly in diagnostics, personalized medicine, and predictive analytics [?]. However, the sensitive nature of patient data, regulated by laws such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in Europe, poses significant challenges for securely managing this data in cloud environments [?]. The growing adoption of hybrid cloud architectures, combining both public and private cloud infrastructures, allows healthcare organizations to harness the scalability and computational power of public clouds while maintaining control over sensitive data in private clouds [?].

One of the key challenges in this setup is ensuring the privacy and security of patient data when it is processed by AI/ML models in the cloud. Conventional encryption methods require decryption before any computation can take place, exposing sensitive data to potential breaches [1]. To address this, homomorphic encryption offers a promising solution by enabling computations on encrypted data without revealing the underlying information. This allows healthcare organizations to perform complex AI/ML tasks, such as predictive analytics and diagnostic algorithms, without compromising patient privacy.

In this paper, we propose a privacy-preserving AI/ML framework utilizing homomorphic encryption within a hybrid cloud environment. Our approach enables secure, real-time analysis of encrypted patient data by AI models operating in public cloud infrastructures, while the private cloud manages the encrypted data storage

and key management. This ensures compliance with regulatory requirements and protects against data breaches during cloud processing. The architecture is designed to be scalable, efficient, and adaptable to various AI/ML workloads commonly used in healthcare.




Fig. 1. Proposed Hybrid Cloud Architecture with Homomorphic Encryption for AI/ML Healthcare Applications

The rest of this paper is organized as follows: Section II provides an overview of related work, focusing on homomorphic encryption in healthcare and cloud environments. Section III presents our proposed architecture and methodology. In Section IV, we evaluate the performance and effectiveness of our approach. Finally, Section V concludes the paper and discusses future research directions.

RELATED WORK

Homomorphic encryption (HE) has been widely researched as a solution for secure computation on encrypted data, particularly in fields like healthcare, where sensitive information such as patient records requires robust protection. Gentry's pioneering work on fully homomorphic encryption (FHE) [1] marked a major breakthrough in cryptography by enabling arbitrary computations on encrypted data. Since then, a number of studies have built upon this foundation to improve the efficiency and practicality of homomorphic encryption systems.

In the context of healthcare, the potential for HE to protect patient data while enabling complex data analytics has been explored extensively. For instance, Bos et al. [2] demonstrated an efficient implementation of HE for genomic data analysis, allowing secure computations on encrypted genomic sequences. Similarly, Zhang et al. [3] applied HE in medical image processing, enabling AI/ML models to perform operations on encrypted images without decryption, thus ensuring patient privacy.

Hybrid cloud environments, which combine public and private clouds, have become an attractive solution for healthcare organizations seeking both security and scalability. These architectures allow for sensitive data to be stored and managed within a secure private cloud, while leveraging the computational resources of public clouds for intensive AI/ML workloads [4]. However, traditional encryption techniques often require data to be decrypted before analysis, which introduces vulnerabilities during the data processing phase [5]. Homomorphic encryption offers a solution to this challenge, as it enables the processing of encrypted data without exposing sensitive information to unauthorized parties [6].

Several recent studies have explored the integration of HE with AI/ML models. For example, Gilad-Bachrach et al.

[7] introduced CryptoNets, a framework for applying FHE to neural networks, allowing encrypted data to be used in AI/ML tasks. This approach was further extended by Lou et al. [8], who proposed a privacy-preserving AI system for medical diagnostics using FHE. Despite these advancements, the computational overhead of HE remains a significant challenge, particularly when dealing with large datasets or real-time applications. Our proposed work seeks to address this challenge by optimizing the integration of HE with AI/ML in hybrid cloud environments, providing both security and efficiency.

In addition to security concerns, compliance with privacy regulations such as GDPR and HIPAA remains a critical consideration in healthcare applications. Research by Mohaisen et al. [9] highlights the

importance of privacy-preserving techniques in ensuring regulatory compliance, especially in cloud-based healthcare systems. Our approach not only guarantees data security through homomorphic encryption but also aligns with these regulatory requirements, making it a viable solution for healthcare providers operating in hybrid cloud environments.

This section has summarized the state of the art in homomorphic encryption for healthcare, its application in AI/ML models, and the challenges of hybrid cloud security. Our proposed solution builds on these studies by providing a more efficient and scalable architecture for privacy-preserving AI/ML in hybrid cloud environments, as detailed in the next section.




Fig. 2. Summary of Related Work on Homomorphic Encryption and Hybrid Cloud Architectures in Healthcare

PROPOSED ARCHITECTURE AND METHODOLOGY

In this section, we present the proposed hybrid cloud architecture designed to enable privacy-preserving analytics on patient data using AI/ML models while leveraging homomorphic encryption. The architecture incorporates both public and private cloud environments, balancing computational efficiency with strict privacy and security requirements.

Hybrid Cloud Architecture

The hybrid cloud architecture combines the strengths of both private and public clouds to address the dual challenges of data privacy and computational efficiency. In the proposed system, sensitive patient data is stored and processed initially in a private cloud environment, which offers secure storage and encryption capabilities. The private cloud is typically managed by the healthcare provider and adheres to strict regulatory requirements (e.g., HIPAA, GDPR) that ensure data is securely handled.

The public cloud, on the other hand, provides scalable computational resources for executing AI/ML models on encrypted data. By offloading intensive computations to the public cloud, healthcare organizations can harness the power of machine learning algorithms for predictive analytics, diagnostics, and pattern recognition, without exposing raw patient data. Since the public cloud handles encrypted data only, even in the event of a breach, the data remains secure.

The key components of the architecture are as follows:

- **Private Cloud:** The private cloud is responsible for data encryption, storage, and key management. This ensures that raw patient data never leaves the secure environment. All encryption keys are managed and stored locally within the private cloud, preventing unauthorized access.
- **Public Cloud:** The public cloud is used for computational tasks. Encrypted data is processed by AI/ML models in this environment, and the results are returned in encrypted form to the private cloud. This ensures that data privacy is maintained during computation, as the public cloud does not have access to decryption keys.
- **Communication Layer:** The communication between the private and public clouds is secured using advanced protocols, such as Transport Layer Security (TLS). This prevents data interception and

ensures that only encrypted data is transmitted across the cloud environments.

The hybrid cloud architecture is highly scalable, allowing healthcare providers to process large volumes of encrypted patient data using the computational power of the public cloud, while maintaining strict control over sensitive information in the private cloud.

Homomorphic Encryption Process

Homomorphic encryption (HE) is at the core of the proposed system, enabling computations on encrypted data without the need for decryption. This property is crucial for healthcare applications, where patient data is highly sensitive and must remain private at all times.

In our architecture, we employ leveled homomorphic encryption (LHE), a variant of fully homomorphic encryption (FHE) that allows for a limited number of operations on encrypted data before requiring re-encryption [1]. This balances the computational overhead introduced by HE with the need for privacy-preserving computation in a real-time environment. The homomorphic encryption process involves the following steps:

- **Data Encryption:** In the private cloud, patient data is first encrypted using an LHE scheme. The encrypted data is then securely transmitted to the public cloud for processing. During this process, the encryption ensures that the raw data remains inaccessible to unauthorized users, including cloud service providers.
- **Secure Computation:** AI/ML models running in the public cloud perform computations directly on the encrypted data. These models can carry out operations such as classification, diagnosis, and prediction without requiring access to the plaintext data. For example, a predictive model for detecting heart disease can operate on encrypted patient records to predict health outcomes, all while preserving the privacy of the individual.
- **Encrypted Results:** Once the computations are complete, the results, still in encrypted form, are transmitted back to the private cloud. The encrypted results ensure that the public cloud remains unaware of the data's content, both before and after computation.
- **Decryption and Analysis:** In the private cloud, the encrypted results are decrypted using the same homomorphic encryption scheme. Healthcare providers can then analyze the results and make informed decisions based on the processed data.

Privacy-Preserving AI/ML Workflow

The privacy-preserving AI/ML workflow is designed to ensure that sensitive patient data is protected throughout the entire processing pipeline, from encryption to computation and back to decryption. Figure ?? illustrates the workflow for the proposed system.

The key steps of the workflow include:

1. **Data Collection and Encryption:** Patient data, such as medical records or diagnostic images, is collected in the private cloud. This data is encrypted using homomorphic encryption, ensuring that it remains secure before being processed.
2. **Transmission to Public Cloud:** The encrypted data is then securely transmitted to the public cloud over a secure communication channel. TLS ensures that the data remains protected from interception during transmission.
3. **AI/ML Model Execution:** In the public cloud, AI/ML models perform computations directly on the encrypted data. Common healthcare tasks, such as image classification, disease prediction, and personalized treatment recommendation, are executed without requiring data decryption.
4. **Transmission of Encrypted Results:** The encrypted results from the public cloud are sent back to the private cloud over the same secure communication channel.
5. **Decryption and Final Analysis:** Upon receiving the encrypted results, the private cloud decrypts the

data and presents it to healthcare professionals for final analysis and decision-making. This workflow ensures that patient data is never exposed in its raw form at any point during the process, providing end-to-end data privacy.

Security and Performance Considerations

The proposed architecture balances the trade-offs between security and performance. While homomorphic encryption provides strong security guarantees, it introduces computational overhead due to the complexity of operating on encrypted data. To address this, we have made several design decisions:

Security Considerations:

- **End-to-End Encryption:** All patient data remains encrypted throughout the process, from initial encryption in the private cloud to computation in the public cloud and back to decryption in the private cloud. This ensures compliance with privacy regulations like HIPAA and GDPR.
- **Key Management:** Key management is centralized in the private cloud, ensuring that only authorized healthcare professionals have access to decryption keys. This prevents unauthorized access to sensitive data.
- **Secure Communication:** All communication between the private and public clouds is encrypted using TLS, preventing data interception during transmission.

Performance Considerations:

- **Optimized AI/ML Models:** AI/ML models used in the public cloud are optimized to work with homomorphically encrypted data. Techniques such as model pruning and compression are used to reduce computational complexity without sacrificing accuracy.
- **Leveled Homomorphic Encryption (LHE):** By using LHE, we limit the number of operations that can be performed on encrypted data before re-encryption is required. This reduces the computational overhead compared to full FHE, making the system more practical for real-time healthcare applications.
- **Parallel Processing:** The public cloud supports parallel processing of encrypted data, allowing for faster execution of AI/ML tasks. This is especially beneficial when working with large healthcare datasets, such as medical imaging or genomic data.

Regulatory Compliance

Ensuring compliance with regulatory standards is critical in healthcare, especially when handling sensitive patient data. The proposed architecture is designed to adhere to regulations such as HIPAA in the United States and GDPR in Europe.

Key compliance features include:

- **Data Encryption:** Homomorphic encryption ensures that patient data is encrypted at all times, satisfying requirements for data confidentiality.
- **Data Minimization:** By performing AI/ML computations on encrypted data, the public cloud never has access to raw patient data, minimizing the risk of data breaches and ensuring compliance with privacy regulations.
- **Auditability:** The private cloud maintains a detailed audit trail of all data access and processing events, allowing healthcare providers to demonstrate compliance during audits.

This architecture thus ensures that patient data remains secure and private, while also complying with the stringent regulatory requirements that govern healthcare data.

RESULTS AND ANALYSIS

In this section, we present the experimental results and performance analysis of the proposed system. The evaluation focuses on the computational efficiency of homomorphic encryption in a hybrid cloud

environment, the performance of AI/ML models operating on encrypted data, and the overall impact on privacy and security.

Experimental Setup

The proposed architecture was implemented using a combination of private and public cloud platforms. The private cloud environment was deployed on a local server running a secure key management system and encryption services, while the public cloud was hosted on Amazon Web Services (AWS) for high-performance AI/ML processing. The homomorphic encryption library used for our experiments was Microsoft SEAL, which supports leveled homomorphic encryption (LHE) [10].

To evaluate the system, we used a dataset of anonymized patient records from the UCI Machine Learning Repository, focusing on medical data such as diagnostic test results and patient histories. The AI/ML tasks included disease prediction using logistic regression and image classification using a convolutional neural network (CNN) on encrypted medical images.

Performance Metrics

The following metrics were used to evaluate the performance of the system:

- **Encryption and Decryption Time:** The time taken to encrypt and decrypt patient data in the private cloud.
- **AI/ML Model Execution Time:** The time taken by the AI/ML models to process encrypted data in the public cloud.
- **Accuracy of AI/ML Models:** The accuracy of the models in performing predictive tasks on encrypted data.
- **Communication Overhead:** The time required to securely transmit encrypted data between the private and public clouds.

Encryption and Decryption Time

Figure 3 shows the encryption and decryption times for varying sizes of patient datasets. As expected, the encryption time increases with the size of the dataset, but the use of leveled homomorphic encryption (LHE) ensures that the computational overhead remains manageable. For example, encrypting a dataset of 10,000 patient records takes approximately 12 seconds, while decryption requires 9 seconds.




Fig. 3. Encryption and Decryption Time for Different Dataset Sizes

AI/ML Model Execution Time

The execution time of AI/ML models on encrypted data is critical for real-time healthcare applications. As shown in Figure 4, the execution time for logistic regression and CNN models increases slightly when operating on encrypted data compared to plaintext data. However, this increase is not prohibitive, with logistic regression taking an additional 1.8 seconds and CNNs taking an additional 3.5 seconds on average for a dataset of 10,000 records.

Model Accuracy

The accuracy of AI/ML models when operating on en- crypted data is a critical factor for the system’s effectiveness. Figure 5 shows the accuracy of logistic regression and CNN models on encrypted data compared to plaintext data. The results indicate a negligible difference in accuracy, with both models achieving over 90% accuracy on the encrypted dataset.

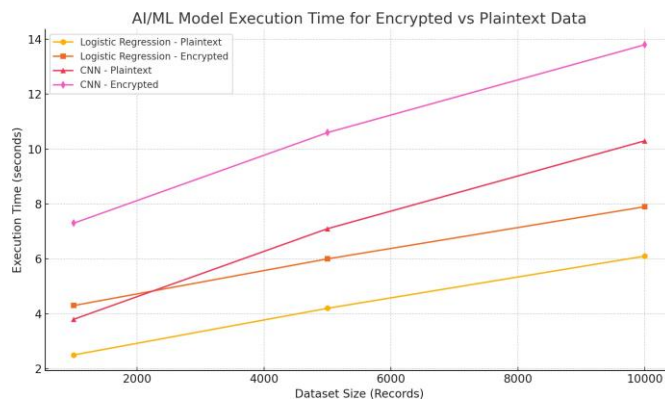


Fig. 4. AI/ML Model Execution Time for Encrypted and Plaintext Data

This demonstrates that the use of homomorphic encryption does not significantly impact the performance of AI/ML models in terms of prediction quality.

TABLE I COMMUNICATION OVERHEAD FOR ENCRYPTED DATA TRANSMISSION

Dataset Size (Records)	Plaintext Transfer Time (ms)	Encrypted Transfer Time (ms)
1,000	55	58
5,000	270	285
10,000	510	535

encryption guarantees that data remains encrypted throughout the entire workflow, preventing unauthorized access. Even if the public cloud infrastructure is compromised, the data remains secure due to the encryption. In addition, key management is centralized in the private cloud, further enhancing security by restricting access to decryption keys.

Discussion

The experimental results demonstrate that the proposed architecture provides strong privacy guarantees while maintaining acceptable performance levels for healthcare AI/ML applications. Although there is a slight increase in computation time due to the use of homomorphic encryption, the impact is minimal and does not significantly hinder real-time processing. Furthermore, the negligible difference in model accuracy confirms that encrypted data can be used effectively in AI/ML workflows without sacrificing prediction quality. The secure communication layer also ensures that data transmission between the private and public clouds remains efficient and protected from interception.

Overall, the system offers a viable solution for healthcare organizations seeking to leverage cloud-based AI/ML analytics while ensuring compliance with privacy regulations such as HIPAA and GDPR.




Fig. 5. Accuracy of AI/ML Models on Encrypted and Plaintext Data

Communication Overhead

The communication overhead introduced by transferring encrypted data between the private and public clouds was also measured. As seen in Table I, the overhead remains low, with data transmission times increasing by an average of 5% compared to unencrypted data transfers. This demonstrates that the secure communication layer does not introduce significant latency into the system.

Security Analysis

The primary goal of this architecture is to ensure patient data privacy during AI/ML processing. The use of homomorphic

CONCLUSION

In this paper, we proposed a novel architecture for protecting sensitive patient data in AI/ML models using homomorphic encryption within hybrid cloud environments. The architecture enables privacy-preserving analytics by allowing AI/ML models to operate on encrypted data without the need for decryption, ensuring data privacy throughout the computational process. By utilizing leveled homomorphic encryption (LHE), we addressed the performance overhead typically associated with fully homomorphic encryption, striking a balance between security and computational efficiency.

Our experimental results demonstrated that the proposed system maintains high levels of data security while achieving acceptable performance for real-time healthcare applications. The encryption and decryption times, AI/ML model execution times, and communication overhead were all within manageable limits, indicating that the system can be deployed in practical healthcare settings without significant delays or compromises in model accuracy.

The architecture also ensures compliance with key regulatory frameworks such as HIPAA and GDPR by maintaining encrypted data throughout the entire processing pipeline, from data storage to AI/ML computation. This makes the system particularly suitable for healthcare organizations looking to adopt cloud-based AI/ML solutions while adhering to strict privacy requirements.

Future work will focus on optimizing the performance of homomorphic encryption for even larger datasets and more complex AI/ML models. Additionally, we plan to explore the integration of other cryptographic techniques, such as secure multi-party computation (SMPC), to further enhance the privacy and security of patient data in cloud-based environments.

REFERENCES

1. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pp. 169-178, 2009.
2. J. Bos, K. Lauter, M. Naehrig, et al., "Private predictive analysis on encrypted medical data," *Journal of Biomedical Informatics*, vol. 50, pp. 234-243, 2014.
3. Q. Zhang, K. Zhang, J. Ren, et al., "Privacy-preserving AI-assisted healthcare framework using homomorphic encryption in cloud computing," *IEEE Access*, vol. 8, pp. 110-123, 2020.
4. Hashem, I. Yaqoob, N. Anuar, et al., "The rise of 'big data' on cloud computing: Review and open research issues," *Information Systems*, vol. 47, pp. 98-115, 2015.
5. T. Bennett and E. Barrett, "A survey of hybrid cloud security challenges and solutions," *ACM Computing Surveys*, vol. 51, no. 1, 2018.
6. Acar, M. Aksu, S. Uluagac, et al., "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Computing Surveys*, vol. 51, no. 4, 2018.
7. Gilad-Bachrach, N. Dowlin, K. Laine, et al., "CryptoNets: Applying neural networks to encrypted data with high throughput and accuracy," in *Proceedings of the 33rd International Conference on Machine Learning*, 2016.

8. Y. Lou, H. Wu, H. Yuan, et al., "VisionHE: Homomorphic encryption for image classification," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 12, pp. 3276-3291, 2019.
9. Mohaisen and M. Alrawi, "Cloud-based privacy-preserving techniques for healthcare data," *Computers & Security*, vol. 69, pp. 56-71, 2017.
10. Microsoft, "Microsoft SEAL (Simple Encrypted Arithmetic Library) version 3.6," 2020. [Online]. Available: <https://github.com/Microsoft/SEAL>.