Lightweight Privacy-Preserving AI Using Hybrid Homomorphic Encryption: Design and Evaluation of Guard AI

Dr. Raghad Tariq Al_Hassani
eng_raghadtariq@yahoo.com

Abstract: Artificial Intelligence (AI) has quickly revolutionized many industries, yet the growing susceptibility of AI models to adversarial attacks and privacy breaches is a significant obstacle to its proliferation. Homomorphic Encryption (HE) and other Privacy-Preserving Artificial Intelligence (PPAI) algorithms, which can compute on encrypted data, facilitate the implementation of privacy-preserving algorithms, but their implementation and application are typically limited by high computational and scaling costs, particularly on resource-limited systems. To overcome these drawbacks, this paper proposes a lightweight privacy-preserving AI model, Guard AI, on a Hybrid Homomorphic Encryption (HHE) scheme, which integrates both symmetric and homomorphic operations. The proposed design will minimize the computational complexity but ensure high data confidentiality when inferring the model. Guard AI is specially designed to support edge and low-resource devices and can classify data safely on encrypted data with no exposure of sensitive inputs or model parameters. To test the proposed framework, the given framework is applied to a practical healthcare application of health care use case on heart disease classification based on electrocardiogram (ECG) signals, which is very sensitive and highly vulnerable to privacy breaches. Through experiments, it is proven that the proposed HHE-based solution is characterized by a good trade-off between security, efficiency, and accuracy, where communication and computation overhead are low in comparison with other, HE schemes, whilst being competitive in comparison with inference without encryption. On the whole, this paper offers a scaled and effective framework of implementing privacy-sensitive AI systems in resource-constrained settings, indicating the promise of hybrid encryption methods in facilitating secure and lightweight intelligent systems.

Keywords: Privacy-Preserving Artificial Intelligence (PPAI), Hybrid Homomorphic Encryption (HHE), Lightweight AI Framework, Secure Edge Computing, Encrypted Data Classification

---

Dr. Raghad Tariq Al_Hassani, Ministry of Higher Education and Scientific, Minister's Office ,Iraq, Baghdad.

1. Introduction

Artificial Intelligence (AI) has become one of the most disruptive technologies in all industries, which substantially improves automation, decision-making, and analytics based on data. With these improvements, the introduction of AI into applications that handle sensitive and personal information has brought forth serious concerns over the privacy and security. Recent AI systems, especially those based on deep learning in distributed or resource-constrained settings, are becoming susceptible to adversarial attacks and information leakage. The risks do not only jeopardize the integrity of the systems, but also destroy the trust of the populace, thus restricting the large-scale implementation of AI in the most critical areas, including healthcare, finance, and government services.

In order to overcome these issues, Privacy-Preserving Artificial Intelligence (PPAI) has become an important research direction with the promise of providing secure data processing without affecting the performance of the model. Out of the numerous techniques proposed, Homomorphic Encryption (HE) has received significant attention due to its unique capability to perform computations on encrypted data, without the need to comprehend the underlying plaintext. There are operations such as addition and multiplication that can be done in a secure manner on the ciphertexts using this property. A range of optimized schemes have been developed since the launch of Fully Homomorphic Encryption (FHE) [1] like CKKS [2], TFHE [3], and BFV [4, 5], have been used in the context of Machine Learning as a Service (MLaaS) [13].

Nevertheless, even with such great security assurances, traditional HE methods have serious drawbacks, such as high computational cost, large ciphertexts, and inability to scale. These constraints do not allow them to be used in real-time systems and resource-limited environments, such as edge devices and Internet of Things (IoT) applications. As a result, an increased demand to have more efficient and lightweight encryption mechanisms that can maintain privacy without affecting performance is on the rise.

To overcome these weaknesses, some recent research has thought of Hybrid Homomorphic Encryption (HHE) in which the integration of the symmetric encryption and the homomorphic encryption is undertaken to create a balance between security and efficiency [14, 15]. In HHE, the information is encrypted by a symmetric key algorithm to make it computationally efficient then homomorphic encryption is applied to secure the symmetric key. This composite scheme lowers the size of ciphertext, reduces the computational burden, and communication overhead than conventional HE schemes. However, even more recent advances in HE-friendly symmetric ciphers, such as HERA, Rubato, have made HHE even more practical and efficient.

To expand on these advances, the current paper will suggest a lightweight privacy-sensitive AI framework, Guard AI, which is tailored towards resource-constrained settings. The suggested framework takes advantage of HHE to allow the secure and efficient inference over encrypted data and maintain the data confidentiality and model integrity. Guard AI offers a flexible, scalable, and feasible way of implementing privacy-preserving intelligence in the real world by combining scalable encryption methods with AI-enabled classification.

2. Related Work

Privacy-Preserving Artificial Intelligence (PPAI) is an area of research that has garnered considerable interest because of the increased demand of safe and effective AI systems that can manage sensitive information. Among the methods suggested, one of the basic methods is Homomorphic Encryption (HE) because it allows computing operations to be done on encrypted data without loss of confidentiality in the processing pipeline.

HE was initially proposed with Fully Homomorphic Encryption (FHE), allowing to do arbitrary calculations with ciphertexts. Nevertheless, the initial FHE protocols were too complex to be practically used. To overcome these shortcomings a number of optimized HE schemes have been suggested. The BFV scheme facilitates arithmetic on integer ciphertexts, and removes costly bootstrapping in some cases, and may be appropriate to moderate-depth computation [16]. Conversely, CKKS scheme allows the approximate arithmetic operation on floating-point data, which is very appropriate in the real-world AI applications that demand the accuracy of numbers [17]. Likewise, TFHE pays attention to enhancing the effectiveness of bootstrapping and encourages quick binary operations, which is why it can be used with logic-based and real-time AI processing [18].

Although these improvements have been made, the current methods of HE has significant limitations when used in a large-scale or resource-constrained setting. A number of studies have also shown incorporation of HE in privacy-preserving machine learning (PPML) systems. As an example, non-linear activation functions (executed using TFHE-based implementations) have been performed using lookup table (LUT) operations, with BFV and CKKS adopting a similar approach to permit this functionality [19]. Moreover, frameworks, like TAPAS and FHE-DiNN, demonstrated that HE can reach reasonable accuracy in encrypted inference settings [20]. Nevertheless, the approaches can be computationally expensive, larger in ciphertext, and have high latency, limiting their use in real-time and edge-based systems [21].

A recent research Hybrid Homomorphic Encryption (HHE) has been suggested as a potential solution to these problems. HHE uses a hybrid approach of symmetric

encryption and HE to minimize computation and communication costs. Early such applications employed conventional symmetric ciphers such as AES but the high multiplicative depth of these ciphers limited their performance in homomorphic environments [22]. New symmetric ciphers specifically to be HE compatible have thus been added.

A number of enhanced HHE schemes have been suggested to enhance efficiency and scalability. As an example, HERA can perform floating-point arithmetic in the CKKS framework, and employs the use of Weighted Modular Arithmetic (WMA) to improve performance. ELISabeth, which was developed to operate with TFHE, focuses on optimizing binary operations, and PASTA, which works with BFV, focuses on how best to run integer operations in large-scale artificial intelligence systems [23]. These strategies indicate how HHE can be flexible to accommodate various types of data and computation.

Although HHE significantly improves the performance of the HE-based systems, it is not yet applied in the real world in the context of AI applications. The current literature mainly concentrates on cryptographic optimization, as opposed to system-level interaction with AI systems, especially in resource-constrained systems. Additionally, there are no lightweight end-to-end systems to fuse HHE and AI inference without compromising security, computing power and model accuracy [24].

In order to fill these research gaps, this paper introduces a lightweight privacy-preserving AI, Guard AI, that uses HHE to achieve secure and efficient encrypted inference on resource-constrained devices. The proposed framework, in contrast to the current methods, is aimed at minimizing the computational cost without compromising data privacy or model integrity, and thus can be used in the real-world context of sensitive data, like the healthcare field. To facilitate a clear comparison of existing HE and HHE approaches, Table 1 summarizes their key features, advantages, and limitations, particularly in terms of efficiency and suitability for resource-constrained AI systems.

Table 1: Comparative Analysis of HE and HHE Schemes in Privacy-Preserving AI

Scheme	Encryption Type	Supported Data Type	Key Advantages	Limitations	Suitability for Lightweight AI
BFV	Homomorphic Encryption	Integer	Accurate arithmetic operations, no	Limited scalability, high computation	Low

			bootstrapping in some cases	for deep circuits	
CKKS	Homomorphic Encryption	Floating-point (approximate)	Efficient for real-valued AI computations	Ciphertext expansion, precision loss	Medium
TFHE	Homomorphic Encryption	Binary	Fast bootstrapping, supports real-time logical operations	High computational overhead	Low
HERA	Hybrid Homomorphic Encryption	Floating-point	Optimized for HE, improved performance using WMA	Requires advanced parameter tuning	Medium
Elisabeth	Hybrid Homomorphic Encryption	Binary	Efficient binary operations in TFHE	High computational complexity in large-scale tasks	Medium
PASTA	Hybrid Homomorphic Encryption	Integer	Scalable integer operations, efficient integration with BFV	Integration complexity	Medium
Proposed Guard AI	Hybrid Homomorphic Encryption (Lightweight)	Multi-type (Integer / Floating / Signals)	Reduced computational overhead, smaller ciphertext, privacy-preserving inference, edge compatibility	—	High

3. Methodology

This section presents the design and implementation of the proposed Guard AI framework, a lightweight privacy-preserving system based on Hybrid Homomorphic Encryption (HHE). The methodology is structured to describe the system architecture, protocol workflow, mathematical modeling, and security analysis.

3.1. System Architecture

The suggested Guard AI structure takes a distributed Privacy-Preserving Machine Learning (PPML) framework that consists of three key actors, including users, a Cloud Service Provider (CSP), and an analyst. Before uploading the sensitive information to the CSP, users encrypt their information with a mixture of HHE and symmetric encryption. CSP stores and processes encrypted data without access to any plaintext information, whereas the analyst processes encrypted inference results with the help of machine learning models. Every user uses his/her own cryptographic key in order to provide high data isolation and data confidentiality.

This architecture ensures that sensitive data is secured during the computation lifecycle, it cannot be accessed by unauthorized parties and it is possible to efficiently infer data encrypted by this architecture. In addition, the system is made to be deployed in resource-constrained environments, like edge and IoT devices.

As a means to facilitate a clear picture of the suggested Guard AI framework, the general system architecture is depicted, emphasizing the secure communication within the system among users, the Cloud Service Provider (CSP) and the analyst. The architecture focuses on the combination of symmetric encryption and Hybrid Homomorphic Encryption (HHE) to allow secure transmission of data and encrypted computation. The system guarantees privacy of end-to-end data by running machine learning inferences on ciphertexts without unveiling sensitive information. In addition, the architecture is built towards the lightweight deployment in the resource-constrained environments. Figure 1 shows the entire system model.

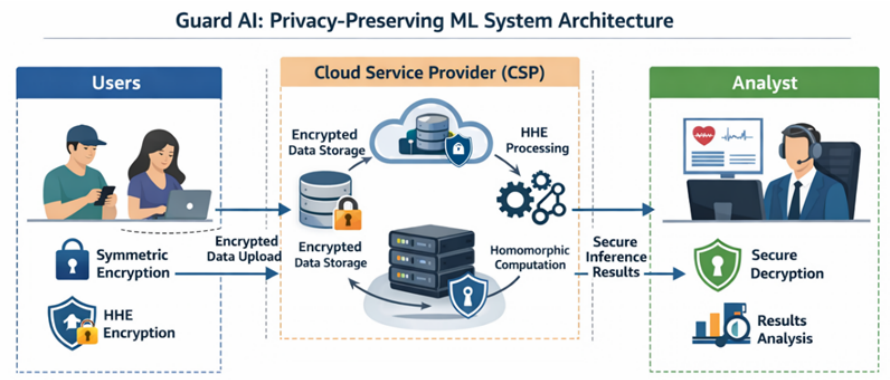

Figure 1: System Architecture

3.2. Guard AI Framework Design

The proposed Guard AI framework integrates lightweight encryption mechanisms with AI-based inference to achieve a balance between security and computational efficiency. The framework is designed based on the following key principles:

a) Minimizing computational overhead for edge deployment
b) Preserving data confidentiality during transmission and processing
c) Enabling secure inference over encrypted data
d) Supporting scalable AI applications in distributed environments

Unlike traditional systems based on HE, Guard AI employs HHE to reduce the scale of ciphertext and computation expenses with high privacy guarantees. To further elucidate on the structural structure of the proposed Guard AI framework, a conceptual representation is given to show the interaction between the main building blocks of the framework. The framework is organized into several levels which are data encryption, secure communication, cloud-based processing and inferential generation. All layers are built to guarantee the privacy of data and at the same time, be computationally efficient. The overall structure of the Guard AI is illustrated in Figure 2.

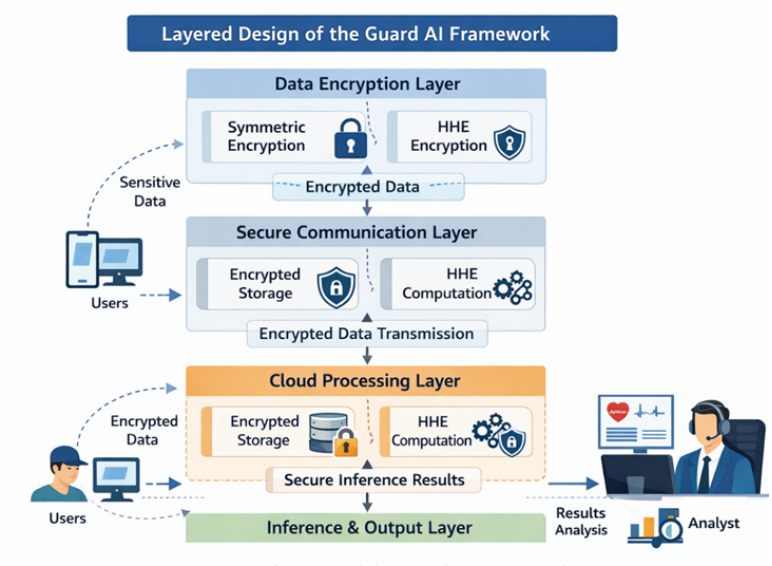

Figure 2: Layered design of the Guard AI framework for lightweight privacy-preserving AI using Hybrid Homomorphic Encryption (HHE).

In order to supplement the architectural design offered in Figure 2, a breakdown of the main elements of the Guard AI framework is also offered. Table 2 provides a summary of the key functions of each of the components, and indicates the role of each in the realization of lightweight encryption, secure computation and efficient data processing in the system.

Table 2: Core Components of the Guard AI Framework and Their Functions

Item	Component	Function
1	SKE (Secure Symmetric Encryption)	Provides efficient and lightweight encryption of user data before transmission
2	HHE (Hybrid Homomorphic Encryption)	Enables secure computation over encrypted data without exposing plaintext
3	CSP (Cloud Service Provider)	Performs storage and processing of encrypted data using HHE mechanisms
4	Analyst	Interprets decrypted results and performs final analysis

All the elements introduced in Table 2 allow the suggested Guard AI framework to create a balance between security and computation efficiency. Specifically, a combination of SKE and HHE is crucial to minimizing computing expenses and maintaining data confidentiality, which renders the framework applicable to the implementation in resource-limited settings.

3.3. 2GML Protocol Workflow

The 2GML protocol is the fundamental working mechanism of the proposed Guard AI framework, which allows the secure and effective interaction between users and the Cloud Service Provider (CSP). The protocol is particularly structured to enable privacy-friendly machine learning whereby encrypted data can be manipulated without revealing sensitive data. The 2GML protocol is organized into four successive phases of its workflow:

a)Phase of setup: This phase entails the production of cryptographic keys and safe initiation between the user and the CSP. Public and evaluation keys are shared whereas the private keys are confidential.

b)Upload Phase: The input data is encrypted with Secure Symmetric Encryption (SKE) and the resulting symmetric key is again encrypted with Hybrid Homomorphic Encryption (HHE). The encrypted data and keys are then transmitted to the CSP.

c)Evaluation Phase: The CSP executes machine learning inference on encrypted data with HHE operations, without accessing the encrypted data or sensitive information.

d)Classification Phase: The encrypted outputs are sent back to the user who decrypts them locally to get the final output of the prediction.

To give a better idea of the working sequence of the proposed protocol, a flowchart representation has been provided to show the stepwise execution of the 2GML protocol in Figure 3 illustrates the end-to-end execution of the 2GML protocol, where all the calculations are made on encrypted data without compromising the confidentiality of the data or the efficiency of the system. This workflow demonstrates the feasibility of the suggested methodology to be implemented in the real-world setting that is limited in resources.

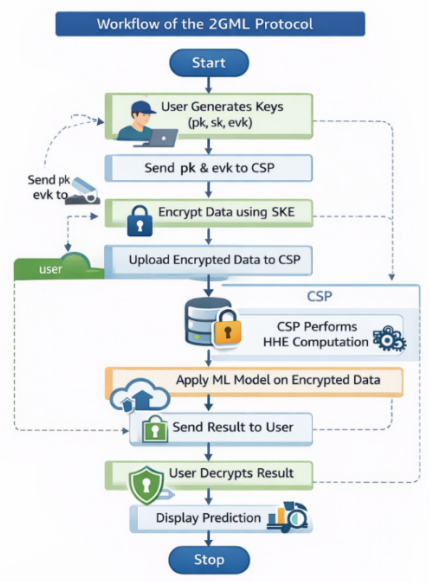

Figure 3: Workflow of the 2GML protocol within the Guard AI framework, demonstrating secure data encryption, HHE-based computation, and privacy-preserving inference across all processing stages.

3.4. Mathematical Formulation

This section illustrates the mathematical description of the proposed Guard AI framework, which combines Secure Symmetric Encryption (SKE) with Hybrid Homomorphic Encryption (HHE) to provide privacy-preserving machine learning with encrypted data.

Let:

- xi denote the input data sample of user iii
- Ki denote the symmetric encryption key
- ci denote the encrypted data
- cKi denote the encrypted symmetric key
- pku,sku denote the public and private keys of the user
- $evku$ denote the evaluation key used in homomorphic operations
- $f(\cdot)$ denote the machine learning model
- $cres$ denote the encrypted result
- res denote the final decrypted prediction

In the first step, the user encrypts the input data using symmetric encryption to ensure computational efficiency:

$$\mathrm{ci} = \mathrm{SKE.Enc}(\mathrm{xi},\mathrm{Ki})$$

To preserve the confidentiality of the symmetric key, it is further encrypted using HHE:

$$\mathrm{cKi} = \mathrm{HHE.Enc}(\mathrm{Ki},\mathrm{pku})$$

Thus, the transmitted data consists of:

$$\mathrm{Ci} = (\mathrm{ci},\mathrm{cKi})$$

At the Cloud Service Provider (CSP), computations are performed directly on encrypted data using HHE. The machine learning model $f(\cdot)$ is applied without decrypting the inputs:

$$\mathrm{cres} = \mathrm{HHE.Eval}(\mathrm{f},\mathrm{ci},\mathrm{cKi},\mathrm{evku})$$

This ensures that all intermediate computations remain encrypted, preserving data confidentiality.

After computation, the encrypted result is returned to the user, who decrypts it using their private key:

$$\mathrm{res} = \mathrm{HHE.Dec}(\mathrm{cres},\mathrm{sku})$$

The final output res represents the predicted class label or inference result.

To highlight the lightweight nature of the proposed framework, the total computational cost can be approximated as:

$$\mathrm{Ttotal} = \mathrm{TSKE} + \mathrm{THHE}_{\mathrm{Enc}} + \mathrm{THHE}_{\mathrm{Eval}} + \mathrm{THHE}_{\mathrm{Dec}}$$

Thus, the hybrid approach significantly reduces the overall computational overhead compared to pure HE systems.

3.5. Security Model and Threat Analysis

This paper provides a detailed security discussion of the suggested Guard AI system, tested in a semi-honest (honest-but-curious) adversarial paradigm. The participating entities in this model adhere to the protocol, but might seek to derive sensitive information based on available data when executing the protocol.

.3.5.1 Threat Model

The opponent A is supposed to have the following capabilities:

a. Hacking into the communication between users and the Cloud Service Provider (CSP).
b. Modification or replacement of relayed ciphertexts.
c. Attack on encrypted data or machine learning models.
d. There is risk of the CSP being compromised or collusion with bad users.

Even with such capabilities, the adversary cannot gain access to the private keys and cannot directly decrypt encrypted data.

.3.5.2 Security Guarantees

The suggested Guard AI architecture guarantees a number of core security guarantees that cumulatively guarantee strong protection to privacy-preserving machine learning processes.

a. Data Confidentiality: Data of users are encrypted when storing, transmitting, and computing data, which means that no sensitive data is revealed at any point of its processing.
b. Model Privacy: The machine learning model is safely stored in the Cloud Service Provider (CSP) and cannot be accessed by users or other parties.
c. Integrity and Authenticity: Digital signature schemes used will protect the integrity of all messages sent, and ensure integrity of data and reliable communication.
d. Computation Privacy: Inference is done with encrypted data (Hybrid Homomorphic Encryption (HHE)) so that the intermediate results of computation are not disclosed.

In order to assess the strength of the proposed framework further, the attack scenarios and associated mitigation strategies are presented in Table 3. This comparative study underlines the benefits of incorporating HHE and cryptographic primitives to improve the resilience of systems to protocol-level and sophisticated cryptographic attacks.

Table 3: Security Analysis of the Guard AI Framework Against Potential Attacks

Attack Type	Description	Potential Impact	Security Mechanism	Security Assurance
Ciphertext Substitution Attack	Adversary replaces valid ciphertexts during upload or evaluation phases	Incorrect computation results or data manipulation	Digital Signature (EUF-CMA secure)	Detects tampering with negligible forgery probability
Unauthorized Model Access Attack	Adversary attempts to access or extract the ML model from CSP	Loss of model confidentiality and intellectual property	HHE-based encrypted computation + restricted access	Model remains hidden and inaccessible

Data Leakage Attack	Adversary tries to infer sensitive information from encrypted outputs	Exposure of user data patterns or private information	End-to-end encryption using HHE	No plaintext exposure during computation
Cryptographic Attacks	Includes linearization, Gröbner basis, and differential attacks	Compromise of encryption scheme	Secure HE schemes + hybrid encryption design	Strong resistance against advanced cryptanalysis

As the analysis provided in Table 3 shows, the suggested Guard AI framework offers a high level of protection against a broad spectrum of attack vectors. With a combination of hybrid encryption methods and secure communication mechanisms, the framework succeeds in guaranteeing data confidentiality, model privacy, and adversarial manipulation resistance. These features render the suggested method very adaptable to implementation in resource-starved and critical security settings.

4. Results and Discussion

In order to determine the efficiency of the suggested Guard AI framework, the large experiments were performed with the help of ECG-based datasets under plaintext and encrypted circumstances. The analysis is conducted in three areas, namely, accuracy, computational performance and scalability.

To evaluate the effect of encryption on the performance of the models, the classification accuracy was measured in three conditions: plaintext (floating-point), plaintext (integer) and encrypted inference with HHE. The findings in Table 4 show that the suggested Guard AI framework is highly accurate with various input sizes. It is important to note that the error introduced by encryption is small and therefore the HHE-based method does not compromise model performance and provides data privacy. This validates the appropriateness of the framework in the real world privacy-sensitive applications.

Table 4: Accuracy Comparison under Plaintext and Encrypted Inference

Data Input	Plaintext (Float)	Plaintext (Integer)	Encrypted
1	100%	100%	100%
10	90%	90%	90%
20	90%	95%	90%
50	88%	92%	90%
100	86%	91%	90%
500	87%	87.2%	86.8%
1000	87.9%	87.3%	87.4%

2000	88.2%	87.4%	87.55%

A comparative analysis between the accuracy of classification before and after data scaling was conducted on various machine learning algorithms to further examine how preprocessing of data influences model performance. The findings point to the diversity in model sensitivity to scaling and show that preprocessing can have a great effect on predictive accuracy. The detailed comparison is presented in Figure4.

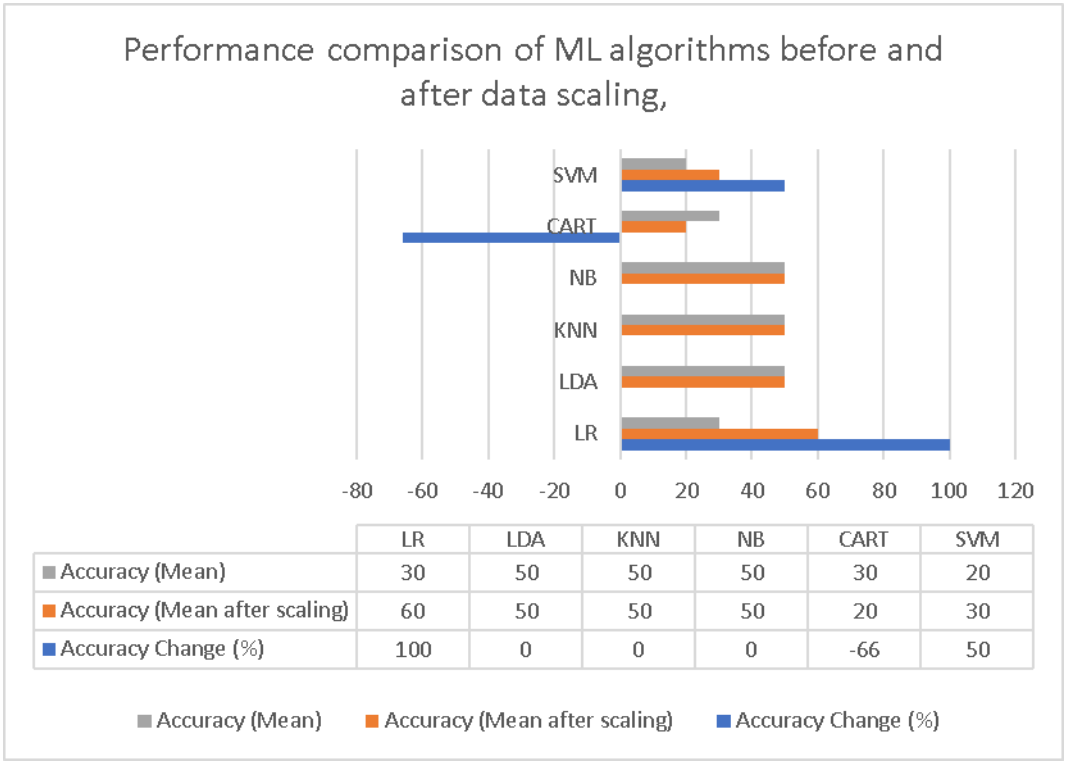

Figure 4: Comparison Plaintext versus encrypted inference with varying input sizes, which shows that there is little performance deterioration in the proposed Guard AI framework.

In order to determine the computational efficiency, the time it took to execute each step in the 2GML protocol was measured. Table 5 indicates that the evaluation stage is the largest part of the total computational cost as homomorphic operations are

complicated. Nevertheless, setup, upload, and classification have fairly low latency, proving the effectiveness of the lightweight design. This computational cost sharing is anticipated in systems based on HHE, and it substantiates the feasibility of offloading heavy computation to the CSP.

Table 5: Execution Time of 2GML Protocol Phases

Phase	User Time	Server Time	Total Time
Setup	243 ms	0	243 ms
Upload	607 ms	0	607 ms
Evaluation	0	3597.7 s	3597.7 s
Classification	900 ms	0	900 ms

An experimental study was conducted by changing the degree of the polynomial modulus to explore the nature of the security vs. computational efficiency trade-off inherent to homomorphic encryption solutions. The paper explores the impact of the rise in the level of security on the size of the ciphertext and the latency in the operations. This trade-off is critical to comprehend when developing lightweight privacy-preserving systems like Guard AI. Figure 5 shows the results.

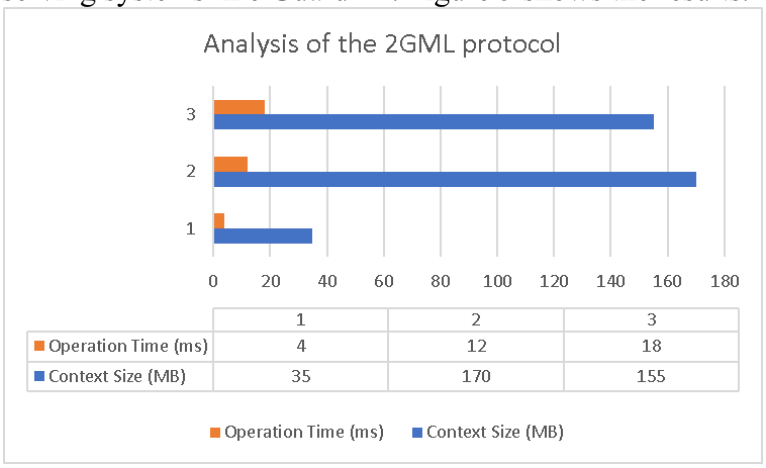

Figure 5: Execution time analysis of the 2GML protocol, illustrating the computational cost distribution across different phases, with emphasis on the evaluation stage as the most resource-intensive component.

Overall, the experimental results confirm that the proposed Guard AI framework achieves an effective balance between security, accuracy, and computational efficiency. Although the evaluation phase also adds computational overhead, the

lightweight design ensures user-side processing is reduced, which makes the framework extremely applicable to edge and resource-constrained settings. Moreover, the accuracy drop is very small, which proves the stability of the HHE-based method in maintaining the model performance with encryption.

5. Conclusion

This paper introduced Guard AI, a small privacy-preserving artificial intelligence system using Hybrid Homomorphic Encryption (HHE) to facilitate privacy-safe and efficient machine learning in resource-constrained settings. The suggested framework helps to overcome the most significant issues with conventional Homomorphic Encryption (HE), in particular, high computational costs, large ciphertext, and a lack of scalability. The Guard AI system is a combination of Secure Symmetric Encryption (SKE) along with HHE to enable the balance between efficient calculations and high data confidentiality. By introducing the 2GML protocol, the system is able to process encrypted data and make sure that the Cloud Service Provider (CSP) and any other third party do not have access to sensitive user data or model parameters. The design facilitates safe outsourced computation and is thus very applicable in the real-world applications like healthcare, finance, and edge computation systems. The experimental findings indicated that the suggested method has high classification accuracy at different input sizes with a slight degradation as compared to plaintext inference. Moreover, the performance analysis revealed that the evaluation step is the most expensive computationally because it entails homomorphic operations, whereas the rest of the steps such as setup, upload, and classification are efficient, which further adds to the lightness of the framework. Moreover, the security analysis showed that the Guard AI is highly resistant to various vectors of attacks, such as ciphertext substitution, data leakage, and unauthorized access to the model. The framework ensures that there is high level of privacy and that the system reliability is operational by incorporating both advanced cryptographic tools and realistic system design. To sum up, Guard AI is a scalable and efficient platform of implementing privacy-preserving AI-based systems in sensitive and resource-constrained environments. Future directions could be to optimize the homomorphic evaluation performance, add more sophisticated deep learning models and extend the framework to support real-time and large-scale distributed applications.

References

[1] C. Gentry, "Fully homomorphic encryption using ideal lattices," in Proc. 41st Annu. ACM Symp. Theory of Computing, May 2009, pp. 169–178.

[2] B. Li and D. Micciancio, "On the security of homomorphic encryption on approximate numbers," in Proc. Annu. Int. Conf. Theory and Applications of Cryptographic Techniques, Jun. 2021, pp. 648–677.

[3] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, "TFHE: Fast fully homomorphic encryption over the torus," Journal of Cryptology, vol. 33, no. 1, pp. 34–91, 2020.

[4] A. Kim, Y. Polyakov, and V. Zucca, "Revisiting homomorphic encryption schemes for finite fields," in Proc. Int. Conf. Theory and Application of Cryptology and Information Security, Dec. 2021, pp. 608–639.

[5] A. Kim et al., "General bootstrapping approach for RLWE-based homomorphic encryption," IEEE Trans. Computers, vol. 73, no. 1, pp. 86–96, 2023.

[6] R. Xu, N. Baracaldo, and J. Joshi, "Privacy-preserving machine learning: Methods, challenges and directions," arXiv preprint arXiv:2108.04417, 2021.

[7] Y. Li et al., "Toward secure and privacy-preserving distributed deep learning in fog-cloud computing," IEEE Internet of Things Journal, vol. 7, no. 12, pp. 11460–11472, 2020.

[8] O. Olasehinde, B. K. Alese, and O. Eqwuche, "Privacy-preserving artificial intelligence: Principles, methods, applications, and

challenges," Journal of Applied Artificial Intelligence, vol. 6, no. 2, pp. 60–70, 2025.
[9] R. Xu, N. Baracaldo, and J. Joshi, "Privacy-preserving machine learning: Methods, challenges and directions," arXiv preprint arXiv:2108.04417, 2021.
[10] I. Ahmad et al., "Machine learning meets communication networks: Current trends and future challenges," IEEE Access, vol. 8, pp. 223418–223460, 2020.
[11] H. C. Tanuwidjaja et al., "Privacy-preserving deep learning on machine learning as a service—a comprehensive survey," IEEE Access, vol. 8, pp. 167425–167447, 2020.
[12] S. M. Zobaed and M. Amini Salehi, "Confidential computing across edge-to-cloud for machine learning: A survey study," Software: Practice and Experience, vol. 55, no. 5, pp. 896–924, 2025.
[13] F. Thabit et al., "A novel effective lightweight homomorphic cryptographic algorithm for data security in cloud computing," International Journal of Intelligent Networks, vol. 3, pp. 16–30, 2022.
[14] A. Bakas, E. Frimpong, and A. Michalas, "Symmetrical disguise: Realizing homomorphic encryption services from symmetric primitives," in Proc. Int. Conf. Security and Privacy in Communication Systems, Oct. 2022, pp. 353–370.
[15] A. Shah and S. Sivakumar, "Encrypted intelligence: A comparative analysis of homomorphic encryption frameworks for privacy-preserving AI," Journal of Economy and Technology, 2025.
[16] J. Cho et al., "Transciphering framework for approximate homomorphic encryption," in Proc. Int. Conf. Theory and Application of Cryptology and Information Security, Dec. 2021, pp. 640–669.
[17] M. J. A. Dias, "Secure healthcare data analysis with hybrid homomorphic encryption," 2025.
[18] S. Belaïd et al., "Further improvements in AES execution over TFHE: Towards breaking the 1 sec barrier," Cryptology ePrint Archive, 2025.
[19] C. Dobraunig et al., "Pasta: A case for hybrid homomorphic encryption," IACR Trans. Cryptographic Hardware and Embedded Systems, no. 3, pp. 30–73, 2023.
[20] P. Gogoi and J. Arul Valan, "Homomorphic encryption for secure and scalable predictive healthcare analytics: A review and case study," Life Cycle Reliability and Safety Engineering, pp. 1–18, 2026.
[21] A. Stoian et al., "Deep neural networks for encrypted inference with TFHE," in Proc. Int. Symp. Cyber Security, Cryptology, and Machine Learning, Jun. 2023, pp. 493–500.

[22] Z. He et al., "Securebadger: A homomorphic encryption-based framework for secure medical inference," Digital Communications and Networks, 2025.
[23] Z. A. Abbood, D. Ç. Atilla, and Ç. Aydin, "Intrusion detection system through deep learning in routing MANET networks," Intelligent Automation & Soft Computing, vol. 37, no. 1, 2023.
[24] A. Vizitiu et al., "Framework for privacy-preserving wearable health data analysis: Proof-of-concept study for atrial fibrillation detection," Applied Sciences, vol. 11, no. 19, p. 9049, 2021.

نظام ذكاء اصطناعي خفيف الوزن مع الحفاظ على الخصوصية باستخدام التشفير المتماثل الهجين: التصميم والتقييم لنظام Guard A

د . رغد طارق الحسني
eng_raghadtariq@yahoo.com

المستخلص: أحدث الذكاء الاصطناعي ثورةً سريعةً في العديد من الصناعات، إلا أن تزايد قابلية نماذج الذكاء الاصطناعي للهجمات الخبيثة وانتهاكات الخصوصية يُشكل عائقًا كبيرًا أمام انتشاره. تُسهّل خوارزميات التشفير المتماثل (HE) وغيرها من خوارزميات الذكاء الاصطناعي الحافظة للخصوصية (PPAI)، القادرة على معالجة البيانات المُشفّرة، تطبيق خوارزميات الحفاظ على الخصوصية، ولكن عادةً ما يكون تطبيقها محدودًا بسبب ارتفاع تكاليف الحوسبة والتوسع، لا سيما على الأنظمة ذات الموارد المحدودة. للتغلب على هذه العيوب، تقترح هذه الورقة نموذجًا خفيف الوزن للذكاء الاصطناعي الحافظ للخصوصية، يُسمى Guard AI، يعتمد على نظام تشفير متماثل هجين (HHE)، يدمج العمليات المتناظرة والمتماثلة.
سيُقلل التصميم المقترح من التعقيد الحسابي مع ضمان سرية عالية للبيانات عند استنتاج النموذج. صُمم Guard AI خصيصًا لدعم الأجهزة الطرفية والأجهزة ذات الموارد المحدودة، ويمكنه تصنيف البيانات بأمان على البيانات المُشفّرة دون الكشف عن المدخلات الحساسة أو معلمات النموذج. لاختبار الإطار المقترح، طُبِّقَ على تطبيق عملي في مجال الرعاية الصحية، وتحديدًا على تصنيف أمراض القلب بناءً على إشارات تخطيط كهربية القلب (ECG)، وهو مجال شديد الحساسية وعرضة لانتهاكات الخصوصية. أثبتت التجارب أن الحل المقترح القائم على التشفير الهجين المتجانس (HHE) يتميز بتوازن جيد بين الأمان والكفاءة والدقة، حيث تكون تكلفة الاتصال والحساب منخفضة مقارنةً بأنظمة التشفير الهجين الأخرى، مع الحفاظ على تنافسيته مع الاستدلال بدون تشفير. إجمالًا، تقدم هذه الورقة إطارًا فعالًا وقابلًا للتطوير لتنفيذ أنظمة الذكاء الاصطناعي الحساسة للخصوصية في بيئات محدودة الموارد، مما يُشير إلى إمكانات أساليب التشفير الهجين في تسهيل إنشاء أنظمة ذكية آمنة وخفيفة الوزن.

الكلمات المفتاحية: الذكاء الاصطناعي الحافظ للخصوصية (PPAI)، لتشفير المتماثل الهجين (HHE)، إطار عمل ذكاء اصطناعي خفيف الوزن، الحوسبة الطرفية الآمنة، تصنيف البيانات المشفرة.

د. رغد طارق الحسني ، وزارة التعليم العالي والبحث العلمي ، مكتب الوزيـــر – بغداد – العراق.