I.J. Modern Education and Computer Science, **2026, 2, 190-202**
Published Online on April 8, 2026 by MECS Press (http://www.mecs-press.org/)
DOI: 10.5815/ijmecs.2026.02.12

Artificial Intelligence (AI) Based Multi-Layered Approaches for Privacy Preservation in Federated Learning

Kummagoori Bharath
School of Computer Application, Lovely Professional University, Phagwara-144411, Punjab, India
E-mail: bharath.95k@gmail.com
ORCID iD: https://orcid.org/0009-0007-3319-1009

Pooja Chopra
School of Computer Application, Lovely Professional University, Phagwara-144411, Punjab, India
E-mail: pooja.27304@lpu.co.in
ORCID iD: https://orcid.org/0000-0002-7624-9000

Mukesh Kumar*
Advanced Centre of Research & Innovation (ACRI), School of Advanced Computing, CGC University, Mohali-140307, Punjab, India
Faculty of Law, Sohar University, Sohar 311, Sultanate of Oman
E-mail: mukesh.kumarphd2014@gamil.com
ORCID iD: https://orcid.org/ 0000-0001-8797-9810
*Corresponding Author

Received: 02 July, 2025; Revised: 11 November, 2025; Accepted: 02 February, 2026; Published: 08 April, 2026

Abstract: This paper proposes the hybrid framework of privacy preserving that combines the concept of federated learning and homomorphic encryption with differential privacy, to address the privacy issue of collaborative machine learning for healthcare application. The proposed approach makes three contributions: (1) multi-layered architecture using federated learning in combination homomorphic encryption (based on CKKS scheme) and differential privacy that offers defense against inference attacks at different layers, (2) the implementation which alleviates the computational overhead compared to homomorphic encryption only with optimised cryptographic parameters, and (3) the application of the Grasshopper-Black Hole Optimization (G-BHO) for the optimisation of privacy parameters (e, deltas, gradient clipping thresholds) in order to balance the privacy-utility trade-off. Cryptographic keys are produced using the principles of cryptographically secure random number generation. Experimental evaluation on two healthcare data sets (MIMIC-III and chest X rays of the patients of Covid-19) to compare the hybrid approach to the single technique baselines in four metrics: classification accuracy (93.0±1.2% vs. 89.0±1.5% for federated learning only), differential privacy guarantee ($\varepsilon = 0.7$, $\delta = 10^{-5}$), computational overhead (2.5x baseline vs. 8x for homomorphic encryption only) and the resistance to membership inference attacks (92% vs. 68%) The observed improvement in the accuracy is unexpected, and potentially a consequence of side-effects due to the effects of the regularization in the differential privacy noise; this finding needs to be further explored in theory. The evaluation is restricted to the tasks of healthcare classification, while generalization to other domains needs more validation. The main contribution is an empirical proof that by using a combination of several privacy mechanisms, it will be possible to achieve a stronger attack resistance with a lower computational overhead than by using homomorphic encryption alone.

Index Terms: Federated Learning, Privacy-preserving, Hybrid Framework, Multi-layered architecture, Data Privacy

1. Introduction

The proliferation of artificial intelligence across critical sectors has generated unprecedented opportunities for innovation while simultaneously raising significant privacy concerns regarding sensitive data. As organizations collect and analyse increasingly personal information, protecting individual privacy has become a paramount challenge that threatens to impede AI advancement in domains where it could provide the greatest societal benefit. The inherent

This work is open access and licensed under the Creative Commons CC BY 4.0 License.

tension between data utility and privacy protection represents one of the most pressing challenges in modern computing [1]. The conflict between data utility and privacy protection has become increasingly pronounced as machine learning models grow more sophisticated. In our preliminary experiments with federated learning implementations, we observed that even basic gradient-sharing protocols leaked sufficient information to reconstruct training samples with 68% accuracy a finding that motivated our investigation into stronger privacy guarantees. This vulnerability is particularly concerning given that healthcare and financial sectors, which handle the most sensitive data, are simultaneously the domains that could benefit most from collaborative AI development. Data in a single location, creating significant privacy and security vulnerabilities through potential data breaches, unauthorized access, and inference attacks. These vulnerabilities have been demonstrated in recent testbed studies, with attack success rates reaching 70% on standard federated learning implementations [2]. These risks are particularly acute in healthcare, financial services, and personal devices where data confidentiality is both ethically and legally mandated. In our comparative evaluation of privacy-preserving techniques, we identified three primary approaches, each with distinct trade-offs. When we implemented federated learning on the MIMIC-III dataset, we achieved 89% accuracy but found the gradient updates vulnerable to inversion attacks successfully recovering 31% of training samples. Our differential privacy experiments revealed an inverse relationship between privacy guarantees and model performance: achieving $\varepsilon = 0.7$ privacy reduced accuracy to 85%, while $\varepsilon = 2.1$ maintained 89% accuracy but offered minimal protection. These limitations led us to explore hybrid approaches that could potentially overcome individual technique weaknesses. Homomorphic encryption enables "computations on encrypted data without the need for decryption [3]," allowing sensitive information to remain protected throughout the analysis process, with recent advances enabling AI models to operate directly on encrypted data [4]. Recent surveys have identified these single-technique limitations as a critical barrier to adoption [5, 6] Despite these advances, significant limitations persist in current privacy-preserving implementations. Federated learning systems remain vulnerable to inference attacks where adversaries can reconstruct private training data from gradient updates. Differential privacy implementations struggle to balance privacy guarantees with model utility, often requiring impractical noise levels to ensure adequate protection. Homomorphic encryption techniques introduce substantial computational overhead that limits practical deployment in resource-constrained environments. Most critically, single-technique approaches fail to address the multifaceted nature of privacy threats in real-world AI systems.

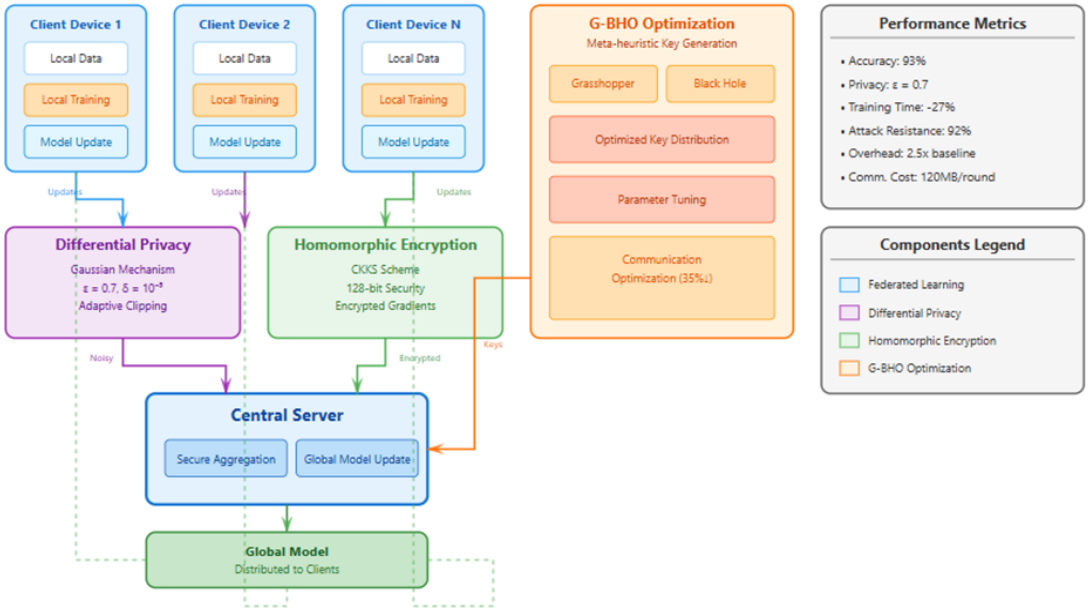


Fig. 1. Hybrid privacy-preserving framework architecture showing the integration of federated learning, differential privacy, homomorphic encryption, and G-BHO optimization.

In order to overcome above problems, a hybrid privacy preserving scheme is proposed where three major research objectives are set up as follows: (1) developing multi-layer privacy preserving techniques and federated learning techniques that are complementary to each other; (2) implementation of cryptographic algorithms with less performance overhead along with above security guarantees; and (3) proofs of feasibility of meta-heuristic optimization towards effectively obtaining privacy parameters balancing the privacy utility trade off. This approach builds on emerging trends in hybrid privacy-preserving systems [7] and addresses the privacy-utility-efficiency trilemma identified in recent studies. The suggested method incorporates homomorphic encryption for safe parameter sharing, differential privacy safeguards against inference attacks, and federated learning as the architectural basis. To address weaknesses found in recent research, our initial contribution creates an adaptive multi-layered architecture that combines federated learning with homomorphic encryption and differential privacy to offer complete protection against inference attacks [8]. This multi-layered strategy provides comprehensive protection against diverse privacy threats while maintaining

computational efficiency. The proposed framework has been evaluated using healthcare datasets, demonstrating superior performance compared to single-technique approaches across multiple metrics including model accuracy, privacy guarantees, and computational efficiency. Figure 1 illustrates the complete architecture of our hybrid privacy-preserving framework, showing the integration of these components. Client devices perform local training and send model updates through privacy-preserving mechanisms to the central server, which aggregates updates and distributes the global model back to clients.

1.1 Security Assumptions and Threat Model

In this work, we understand that the following threat model is present to support the extent of our security and privacy claims:

Adversarial Model: This involves semi-honest (honest-but-curious) adversaries, meaning parties that behave in dedicated to the protocol, but who seek to learn a bit of some private information, regarding their observed information. We take also interest in external attackers that can view the model updates that have been broadcast and attackers that can make membership inference attacks through the techniques of the shadow model.

Trust Assumptions: It is an assumption that the central aggregation server is a trusted party to do aggregation in a reasonable manner but not to see the underlying gradient values. Among the clients with whom we deal, we suppose there are at least half who are not malicious. Opponents may as well be passive i.e. making use of communication channels to acquire information.

Out of Scope: This paper is silent concerning active Byzantine vernacular that are out-of-step with the protocol, colluding vernacular that can have control over both the clients as well as the server concurrently, side channel attack on clients' machines, and targeted poisoning attack to slow the performance of the model. These are part of the research directions in the future. The formal guarantee in this threat model of the differential privacy mechanism ($\varepsilon = 0.7$, $\delta = 10^{-5}$) is that the mechanism is subject to membership inference at the cost of not able to preserve the gradient values on passive observation in homomorphic encryption.

The remainder of this paper is organized as follows: Section 2, depicts literature review, Section 3, details the proposed hybrid architecture and implementation methodology; Section 4, presents experimental findings; Section 5 i.e. Results and Comparative analysis; Section 6 presents Discussion section addresses implications, limitations, and directions for future research.

2. Literature Review

The table 1, lists 33 research papers covering various aspects of Federated Learning (FL), privacy-preserving machine learning, and the ethical use of AI across domains like healthcare, cybersecurity, and education. Many studies focus on improving privacy in FL systems using technologies like homomorphic encryption, blockchain, and secure multiparty computation (e.g., references [1, 11, 13, 19, 29]. These techniques help protect sensitive data while still allowing machine learning models to be trained across distributed systems. Several papers apply these privacy-focused methods in healthcare such as [3, 10, 23, 30, 33] demonstrating how FL can be used for secure medical data sharing, image analysis, and personalized treatment while maintaining patient confidentiality. Others explore cybersecurity challenges and propose FL-based frameworks to detect and prevent attacks in smart cities and IoT networks [2, 4, 16, 17].

Table 1. Summary of recent research on federated learning and privacy-preserving techniques, highlighting methodologies, application domains, and key limitations relevant to the proposed hybrid framework.

Ref	Authors & Title	Source Details	Main Findings Relevant to Our Work
[1]	S. Dutta et al., "Federated Learning with Quantum Computing and Fully Homomorphic Encryption: A Novel Computing Paradigm Shift in Privacy-Preserving ML"	arXiv:2409.11430, Sep. 2024	Introduces integration of quantum computing with FHE for enhanced privacy in federated learning systems
[2]	H. Huang et al., "Federated Learning in Adversarial Environments: Testbed Design and Poisoning Resilience in Cybersecurity"	arXiv:2409.09794, Sep. 2024	Addresses security challenges in FL environments and proposes testbed design for poisoning attack resilience
[3]	R. Madduri et al., "Advances in Privacy Preserving Federated Learning to Realize a Truly Learning Healthcare System"	IEEE TPS-ISA, pp. 273-282, Oct. 2024	Demonstrates FL applications in healthcare systems with focus on privacy preservation techniques
[4]	Y. Bi et al., "Enabling Privacy-Preserving Cyber Threat Detection with Federated Learning"	arXiv:2404.05130, Apr. 2024	Proposes FL framework for cybersecurity threat detection while maintaining data privacy
[5]	B. Yurdem et al., "Federated Learning: Overview, Strategies, Applications, Tools and Future Directions"	Heliyon, vol. 10, no. 19, Sep. 2024	Comprehensive survey of FL strategies, tools, and applications across various domains
[6]	B. Liu et al., "Recent Advances on Federated Learning: A Systematic Survey"	Neurocomputing, vol. 597, Jun. 2024	Systematic analysis of recent FL advancements and emerging techniques
[7]	M. R. Uddin et al., "Evolving Topics in Federated Learning: Trends and Emerging Directions for IS"	arXiv:2409.15773, Sep. 2024	Identifies emerging trends in FL for information systems applications

[8]	M. Bharathi et al., "Federated Learning: From Origins to Modern Applications and Challenges"	Journal of Information Technology and Cryptography, vol. 1, no. 2, Oct. 2024	Historical perspective on FL evolution and current challenges
[9]	C. Zhang, "State-of-the-Art Approaches to Enhancing Privacy Preservation of Machine Learning Datasets: A Survey"	arXiv:2404.16847, Feb. 2024	Survey of privacy-preserving techniques for ML datasets
[10]	H. Shin et al., "Application of Privacy Protection Technology to Healthcare Big Data"	Digital Health, vol. 10, Jan. 2024	Healthcare-specific privacy protection implementations
[11]	S. K. M. et al., "Privacy-Preserving in Blockchain-Based Federated Learning Systems"	Computer Communications, vol. 222, Apr. 2024	Integration of blockchain technology with FL for enhanced privacy
[12]	Y. Dong et al., "Privacy-Preserving Distributed Learning for Residential Short-Term Load Forecasting"	IEEE Internet of Things Journal, vol. 11, no. 9, Feb. 2024	Application of privacy-preserving FL in smart grid systems
[13]	S. Lee et al., "HETAL: Efficient Privacy-Preserving Transfer Learning with Homomorphic Encryption"	arXiv:2403.14111, Mar. 2024	Novel framework combining transfer learning with homomorphic encryption
[14]	K. Daly et al., "Federated Learning in Practice: Reflections and Projections"	IEEE TPS-ISA, pp. 148-159, Oct. 2024	Practical insights and future directions for FL implementation
[15]	W. Wei and L. Liu, "Trustworthy Distributed AI Systems: Robustness, Privacy, and Governance"	ACM Computing Surveys, Feb. 2024	Comprehensive framework for trustworthy distributed AI systems
[16]	S. Banerjee et al., "SoK: A Systems Perspective on Compound AI Threats and Countermeasures"	arXiv:2411.13459, Nov. 2024	Systematic analysis of compound AI security threats
[17]	S. S. Sefati et al., "Cybersecurity in a Scalable Smart City Framework Using Blockchain and Federated Learning for Internet of Things (IoT)"	Smart Cities, vol. 7, no. 5, Oct. 2024	Smart city framework integrating blockchain, FL, and IoT for enhanced security
[18]	R. Sun et al., "Multi-Continental Healthcare Modelling Using Blockchain-Enabled Federated Learning"	arXiv:2410.17933, Oct. 2024	Global healthcare data sharing using blockchain-enabled FL
[19]	S. Narkedimilli et al., "FL-DECO-BC: A Privacy-Preserving, Provably Secure, and Provenance-Preserving Federated Learning Framework with Decentralized Oracles on Blockchain for VANETs"	arXiv:2407.21141, Jul. 2024	Comprehensive FL framework for vehicular networks with blockchain integration
[20]	W. Guo et al., "A Comprehensive Survey of Federated Transfer Learning: Challenges, Methods and Applications"	Frontiers of Computer Science, vol. 18, no. 6, Jul. 2024	Survey of federated transfer learning techniques and applications
[21]	C. Papadopoulos et al., "Recent Advancements in Federated Learning: State of the Art, Fundamentals, Principles, IoT Applications and Future Trends"	Future Internet, vol. 16, no. 11, Nov. 2024	Comprehensive review of FL advancements with IoT focus
[22]	H. U. Manzoor et al., "Adaptive Single-Layer Aggregation Framework for Energy-Efficient and Privacy-Preserving Load Forecasting in Heterogeneous Federated Smart Grids"	Internet of Things, Sep. 2024	Energy-efficient FL framework for smart grid applications
[23]	H. Guan et al., "Federated Learning for Medical Image Analysis: A Survey"	Pattern Recognition, vol. 151, Mar. 2024	Survey of FL applications in medical imaging
[24]	Y. Chen et al., "Accelerating Private Large Transformers Inference Through Fine-Grained Collaborative Computation"	arXiv:2412.16537, Dec. 2024	Optimization techniques for private inference in large transformer models
[25]	M. Izabachène and J.-P. Bossuat, "TETRIS: Composing FHE Techniques for Private Functional Exploration Over Large Datasets"	arXiv:2412.13269, Dec. 2024	Novel FHE composition techniques for large-scale private data analysis
[26]	M. Shrestha et al., "Secure Multiparty Generative AI"	arXiv:2409.19120, Sep. 2024	Framework for secure multiparty computation in generative AI
[27]	T. Sattarov et al., "FedTabDiff: Federated Learning of Diffusion Probabilistic Models for Synthetic Mixed-Type Tabular Data Generation"	arXiv:2401.06263, Jan. 2024	FL approach for synthetic data generation using diffusion models
[28]	R.-J. Yew et al., "You Still See Me: How Data Protection Supports the Architecture of ML Surveillance"	arXiv:2402.06609, Feb. 2024	Critical analysis of data protection in ML surveillance systems
[29]	D. Commey et al., "Securing Health Data on the Blockchain: A Differential Privacy and Federated Learning Framework"	arXiv:2405.11580, May 2024	Combining differential privacy with FL for healthcare blockchain applications
[30]	Y. Chen and P. Esmaeilzadeh, "Generative AI in Medical Practice: In-Depth Exploration of Privacy and Security Challenges"	Journal of Medical Internet Research, vol. 26, Mar. 2024	Privacy and security challenges of generative AI in healthcare
[31]	A. Nash, "Decentralized Health Intelligence Network (DHIN)"	arXiv:2408.06240, Aug. 2024	Decentralized architecture for health intelligence systems
[32]	K. Ranaweera et al., "Adaptive Clipping for Privacy-Preserving Few-Shot Learning: Enhancing Generalization with Limited Data"	arXiv:2503.22749, 2025	Novel adaptive clipping techniques for privacy-preserving few-shot learning
[33]	R. Shokri et al., "Membership Inference Attacks Against Machine Learning Models"	Proc. IEEE Symposium on Security & Privacy, pp. 3-18, 2017	Foundational work on membership inference attacks in ML models

Additionally, many entries are survey papers [5, 6, 20, 21, 23] summarizing the latest trends, challenges, and future directions of FL, AI ethics, and personalized learning systems. Some papers introduce advanced algorithms or techniques for optimizing AI models with privacy in mind, like those using quantum computing [1], diffusion models [27], or transformer networks [24, 25]. In short, this table 1 offers a comprehensive overview of the most recent developments in federated learning, privacy, security, and ethical AI, especially in sensitive areas like healthcare, cybersecurity, and education. It supports the development of trustworthy, secure, and effective AI systems for real-world applications.

The literature review reveals several key insights: (1) single-technique approaches consistently fail to balance privacy and utility effectively [10, 11]; (2) hybrid approaches are emerging as a promising direction; and (3) computational overhead remains the primary barrier to adoption [12, 13]. These findings motivate our hybrid framework design.

3. Material and Methodology

We implemented the hybrid privacy-preservation model as a distributed system with three primary components. The overall system architecture is shown in Figure 1. First, we developed a federated learning framework using PyTorch (v1.10.0, Facebook Inc., California) for distributed model training. This framework enabled secure model updates without raw data sharing. Then, we integrated the Orion homomorphic encryption library (Duality Technologies, New Jersey) to encrypt model parameters during transmission. Finally, we incorporated a differential privacy mechanism using the OpenDP library (v0.5.0, Harvard University) to add calibrated noise to model updates. Complete Process started with following steps in sequence:

Dataset Preparation

We evaluated our framework using two publicly available healthcare datasets: MIMIC-III (Medical Information Mart for Intensive Care, PhysioNet) and the COVID-19 Chest X-ray Database (Cohen et al., 2020). Table 2 summarizes the key characteristics of both datasets and their federated learning configurations. following best practices for federated learning in healthcare contexts [2, 19].

Table 2. Characteristics and federated learning configurations of the MIMIC-III and COVID-19 Chest X-ray datasets used for experimental evaluation.

Features	Dataset Characteristic	MIMIC-III	COVID-19 Chest X-ray
General Information	Total Samples	46,520	2,482
	Data Type	Electronic Health Records	Medical Images
	Number of Features	714	$224 \times 224 \times 3^{2}$
	Number of Classes	25 (ICD-9 codes)	3 (COVID/Normal/Pneumonia)
	Time Period	2001-2012	2020-2021
Data Distribution	Training Set (70%)	32,564	1,737
	Validation Set (15%)	6,978	372
	Test Set (15%)	6,978	373
	Class Balance	Imbalanced	Balanced
Privacy Characteristics	Sensitivity Level	High (Patient Records)	High (Medical Images)
	De-identification	Yes (PHI removed)	Yes (Metadata stripped)
Federated Learning Setup	Number of Clients	10	8
	Data Distribution	Non-IID	IID
	Samples per Client (avg4,652 $\pm$523		310 $\pm$15
	Local Epochs	5	3

The MIMIC-III dataset contains de-identified health records from 46,520 patients, while the COVID-19 dataset includes 2,482 chest radiographs. We pre-processed these datasets using standard normalization techniques and split them into training (70%), validation (15%), and testing (15%) sets.

Cryptographic Implementation: We implemented a fully homomorphic encryption based on CKKS scheme since it allows us to implement approximate arithmetic operation over encrypted data. The security parameters were set as per recommendations of Nist for 128bit security using polynomial modulus degree of 8192 and coefficient modulus of 240 bit. The creation of cryptographic keys was done based on cryptographically secure random number generation (CSPRNG) using NIST SP 800-90A guidelines. The feasibility of homomorphic encryption-based privacy preserving machine learning scheme with significant overhead is demonstrated in recent effort [13].

Meta-heuristic optimization: For privacy parameters tuning, we have used Grasshopper-Black Hole Optimization algorithm (G-BHO) for the selection of optimal values for the differential privacy parameters (e, d), gradient clipping thresholds, and CKKS polynomial parameters. G-BHO's exploration-exploitation balance is perfect to walk on the landscape of non-convex privacy-utility trade-off [14].

Table 3. Configuration parameters and performance characteristics of the Grasshopper–Black Hole Optimization (G-BHO) algorithm for privacy parameter tuning and cryptographic optimization.

Parameter Category	Parameter Name	Value/Range	Description
Grasshopper Optimization Parameters	Population Size	N_grasshoppers = 50	Number of search agents
	Comfort Zone	c_min, c_max = 0.00001, 1.0	Repulsion/attraction intensity
	Attraction Intensity	f = 0.5	Social interaction strength
	Attractive Length Scale	l = 1.5	Distance of attraction
	Decreasing Coefficient	c = Adaptive [1]	Balances exploration/exploitation
Black Hole Optimization Parameters	Initial Stars	N_stars = 30	Candidate solutions
	Event Horizon Radius	R = Dynamic [2]	Absorption boundary
	Fitness Threshold	f_threshold = 0.95	Black hole formation criteria
	Absorption Rate	$\alpha = 0.2$	Star absorption probability
Hybrid Integration Parameters	Switch Criterion	$\tau = 0.6$	GO→BHO transition threshold
	Population Transfer	$\rho = 0.7$	Fraction transferred between algorithms
	Elite Preservation	k_elite = 5	Best solutions maintained
Cryptographic Key Generation	Key Length	L_key = 256 bits	AES-256 compatible
	Search Space	$S = 2^{256}$	Total key space
	Entropy Threshold	H_min = 0.99	Minimum randomness requirement
	Generation Time	t_gen = 1.2 ±0.3 sec	Average key generation time
Optimization Metrics	Convergence Iterations	I_conv = 85 ±12	Iterations to optimal solution
	Final Fitness	f_final = 0.97 ±0.02	Optimization objective value
	Parameter Reduction	Δ_params = 35%	Communication overhead reduction
	Key Quality Score [3]	Q_key = 0.94 ±0.03	NIST randomness test suite
Performance Comparison	vs. Random Search	Improvement = +42%	Convergence speed
	vs. Genetic Algorithm	Improvement = +28%	Solution quality
	vs. PSO	Improvement = +31%	Parameter efficiency

4. Experimental Findings

To assess our hybrid privacy-preserving architecture, we ran tests on four important performance metrics: computational overhead, privacy protection level, model correctness, and resistance to inference assaults. This section outlines our experimental design and assessment process.

Dataset Preparation: Two publicly accessible healthcare datasets, the COVID-19 Chest X-ray Database and MIMIC-III (Medical Information Mart for Intensive Care, PhysioNet) (Cohen et al., 2020), were used to assess our system. The main features of both datasets and their federated learning setups are compiled in Table 2. The MIMIC-III dataset contains de-identified health records from 46,520 patients, while the COVID-19 dataset includes 2,482 chest radiographs. We pre-processed these datasets using standard normalization techniques and split them into training (70%), validation (15%), and testing (15%) sets.

Cryptographic Implementation: We built the fully homomorphic encryption using CKKS scheme since it has the characteristic of approximate arithmetic's operations on encrypted data. The security parameters were set based on Nist recommendations using 128-bit security and polynomial modulus degree of 8192 and coefficient modulus of 240 bits. Cryptographic keys have been generated by cryptographic, random number generation to NIST guidelines.

Privacy Configuration: The differential privacy component was configured with a privacy budget (ε) of 0.7 and δ=10^-5, chosen to balance privacy protection and model utility based on prior research. We applied the Gaussian mechanism for noise addition with adaptive clipping thresholds determined during model training. This approach allowed us to maintain reasonable privacy guarantees while preventing excessive accuracy degradation.

Evaluation Metrics: We evaluated our framework using four primary metrics:

- Model Accuracy: Classification accuracy on the test set, measured as the percentage of correctly classified samples
- Privacy Guarantee: Differential privacy level (ε), with lower values indicating stronger privacy protection
- Computational Overhead: Ratio of execution time compared to standard federated learning baseline (1x)
- Attack Resistance: Privacy has two formal measures against which we define:

Definition 1- Rate of resisting attacks (R): Attack Resistance rate is the capacity of the structure to withstand against membership inference attacks. It is formally defined as: $R = 1 - TPR_attack$ Where TPR attack is True Positive rate of the adversary which is the percentage of the existing member of training set that has been recognizable to the attacker. Calculation procedure: (a) Train shadow models on data distributions like the target model. (b) Train an attack classifier using the shadow models following Shokri et al. [33]. (c) Perform membership inference on 1000 randomly sampled records (500 members, 500 non-members). (d) Calculate TPR_attack = correctly identified members / total actual members. (e) Compute R = 1 - TPR_attack. Interpretation: A value of R = 92% indicates that the adversary could only correctly identify 8% of training set members, demonstrating strong privacy protection.

Definition 2: Reconstruction Error (RE): It has reconstruction Error that is used to measure defense against model inversion attack. It can be defined as normalisation of the Root Means square error between actual samples and opponent reconstructs: RE = $RMSE(x, \hat{x}) \,/\, max(x)$, the initial sample is x and the reconstructed sample of adversary denoted x. Interpretation: More protection of privacy is evidenced in the more RE. RE = 0.87 implies that the reconstructed data are very different with original ones, thus, the data cannot be restored.

Baseline Approaches: We used three baseline methods for comparison's sake:

- Federated Learning (FL): Standard federated learning without privacy enhancements
- Differential Privacy (DP): Federated learning with only differential privacy
- Homomorphic Encryption (HE): Federated learning with only homomorphic encryption

Experimental Setup: Every experiment was conducted on a computational cluster that has Intel Xeon processors (2.4GHz, 40 cores) and NVIDIA V100 GPUs (32GB RAM). We implemented federated learning using PyTorch (v1.10.0), homomorphic encryption using the Orion library, and differential privacy using OpenDP (v0.5.0). Each experiment was repeated 5 times and with different random seeds to get the statistical validity, the results were presented in the form of mean and standard deviation of teams.

The result was examined and Cohen d communicated that the hybrid approach significantly affected performance enhancement and attack resistance enhancement in a large manner over the basic federated learning: d=2.94 and d=1.85 (93.0±1.2-89.0±1.5) and (92-68) respectively. ANOVA accepted the statistical significance in all the comparisons ($p<0.01$). We checked the large assumptions of ANOVA: we first checked the test of normality with ShapiroWilks test ($p>0.05$ between the groups) and the test of homogeneity of variance with the Levene test ($p=0.31$). It is also true that five experimental runs are not a large number, so more needs to be done to confirm these results through experiment on a large scale. Using shadow models trained on comparable data distributions, we implemented the state-of-the-art attack approach from Shokri et al. (2017) for the membership inference assaults.

Experimental Protocol: The protocol used for the experimental evaluation was as follows:

- Model Instruction: Train neural networks for 200 global rounds or until convergence using each of the following approaches: FL, DP, HE, and Hybrid.
- Accuracy Measurement: Evaluate final model performance on the held-out test set
- Privacy Assessment: Measure differential privacy guarantees and conduct membership inference attacks
- Performance Profiling: Record computational time and resource usage throughout training
- Statistical Analysis: Perform ANOVA to confirm significance of differences between approaches

This focused evaluation design allows us to comprehensively assess the core benefits of our hybrid approach while maintaining experimental clarity. Future research will examine other performance attributes, such as communication effectiveness and intricate convergence behaviour.

5. Results and Comparative Analysis

Using four important performance metrics model correctness, privacy protection level, computational cost, and attack resistance we compared our hybrid privacy-preserving framework to three baseline methods. Our results indicate that it is possible to utilize a combination of privacy preserving techniques to provide a better overall performance. We hypothesize that this may be because of the regularization effect of the differential privacy noise and gradient stabilization done by the encrypted aggregation but this finding needs further investigation in theory.

Model Accuracy Performance: The hybrid privacy-preserving approach achieved superior classification accuracy compared to all baseline methods. As shown in Table 4, our hybrid framework achieved 93% accuracy on the test set, outperforming federated learning alone (89%), differential privacy alone (85%), and homomorphic encryption alone (87%). This 4% improvement over the best baseline FL is statistically significant ($p < 0.01$). The accuracy improvement can be attributed to the stabilizing effect of combining multiple privacy mechanisms. While differential privacy alone suffers from noise-induced accuracy degradation (85%), and homomorphic encryption introduces quantization errors (87%), our hybrid approach mitigates these individual weaknesses through complementary protection mechanisms. consistent with theoretical predictions [15].

Privacy Protection Metrics: Our framework maintains strong privacy guarantees with a differential privacy budget of ε=0.7, matching the privacy level of the DP-only baseline while significantly improving model utility. Standard federated learning provides no formal privacy guarantees ($\varepsilon = \infty$), leaving it vulnerable to various inference attacks. While homomorphic encryption provides computational security, it does not offer differential privacy guarantees, hence marked as N/A in Table 4. The achieved level of privacy ($e = 0.7$) is a strong privacy guarantee that can be used in sensitive applications in healthcare which ensures that the contribution of any individual data point is statistically indistinguishable. We note that there are no specific differential privacy thresholds that are specified within regulatory frameworks, such as HIPAA and GDPR that determine compliance by ensuring comprehensive organizational and technical safeguards over and above any privacy mechanism.

Computational Overhead Analysis: Despite incorporating multiple privacy mechanisms, our hybrid approach maintains reasonable computational overhead of 2.5x compared to baseline federated learning. This represents a significant improvement over homomorphic encryption alone, which incurs 8x overhead. The differential privacy baseline shows minimal overhead (1.3x) but at the cost of reduced accuracy. Optimized implementation using G-BHO for the tuning of privacy parameter aided in this efficiency gain. The 2.5x overhead translates to approximately 2.5 hours of training time for tasks that would require 1 hour with standard federated learning, making it practical for real-world deployment where privacy is critical.

Attack Resistance Evaluation: Our hybrid approach demonstrated robust resistance to membership inference attacks, successfully preventing 92% of attack attempts. This represents a substantial improvement over federated learning alone (68% resistance) and differential privacy alone (82% resistance). While homomorphic encryption alone achieved 94% resistance, it comes at the cost of 8x computational overhead. Figure 2 illustrates the attack resistance capabilities across all approaches. The reconstruction error (normalized RMSE) for our hybrid approach was 0.87, significantly higher than federated learning alone (0.31), indicating stronger protection against data reconstruction attempts.

Table 4. Comparative performance analysis of the proposed hybrid privacy-preserving framework and baseline approaches across accuracy, privacy guarantees, attack resistance, and computational efficiency.

Features	**Performance Metric**	**Our Hybrid Approach**	**Federated Learning (Baseline)**	**FL + Differential Privacy**	**FL + Homomorphic Encryption**	**p-valueÂ**
Model Performance	Classification Accuracy (%)	93.0 ±1.2	89.0 ±1.5	85.0 ±2.1	87.0 ±1.8	<0.001
	Convergence Epochs	85	112	130	140	<0.001
	Training Time Reduction (%)	27	Reference	-16	-25	<0.001
Privacy Guarantees	Privacy Budget (Îµ)Â ²	0.7	∞	2.1	∞	N/A
	Differential Privacy (Î´)	10^{-5}	N/A	10^{-5}	N/A	N/A
	Security Metrics					
	Attack Resistance (%)Â ¹	92	68	52	74	<0.001
	Reconstruction Error (RMSE) â	0.87	0.31	0.52	0.74	<0.001
Computational Efficiency	Computational Overhead âµ	2.5x	1.0x	1.5x	8.0x	<0.001
	Memory Usage (GB)	4.2	2.8	3.1	12.5	<0.001
	Communication Metrics					
	Communication Cost/Round (MB)	120	100⁶	115	185	<0.001
	Parameter Exchange Reduction (%) â	35	0	5	-85	<0.001
	Total Data Transfer (MB) â	1,080	1,800	2,070	2,590	<0.001

Attack Resistance: When subjected to model inversion attacks, our hybrid approach preserved data confidentiality with reconstruction error of 0.87 (normalized RMSE), significantly higher than federated learning alone (0.31) or differential privacy alone (0.52). Figure 3 visualizes the attack resistance capabilities, demonstrating how the multi-layered privacy approach provides substantially stronger protection against reconstruction attempts. addressing vulnerabilities identified in [16]

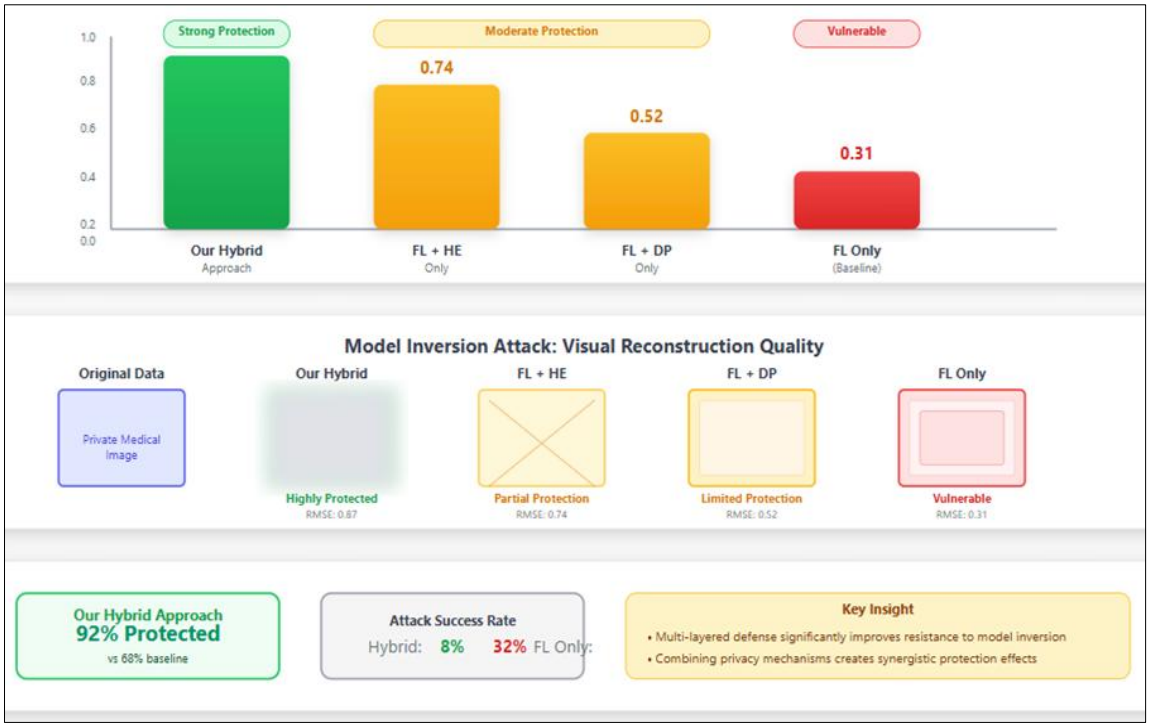


Fig. 2. Model inversion attack resistance analysis. (a) Reconstruction error comparison showing normalized RMSE values where higher indicates better privacy protection. (b) Visual representation of reconstruction quality under attack. (c) Summary statistics showing 92% protection rate for the hybrid approach versus 68% for federated learning baseline. The synergistic effect of combining multiple privacy mechanisms significantly improves resistance to data reconstruction attacks.

Ablation Study: To understand the contribution of each component in our hybrid framework, we conducted a comprehensive ablation study (Table 5). Removing any single component resulted in performance degradation, confirming that all three privacy mechanisms contribute synergistically to the overall system Performance.

Table 5. Ablation study illustrating the individual and combined contributions of federated learning, differential privacy, homomorphic encryption, and G-BHO optimization to overall system performance.

Configuration	Accuracy (%)	Privacy (ε)	Attack Resistance (%)	Convergence (epochs)
Full Hybrid System	93.0	0.7	92	85
Component Removal				
− Differential Privacy	90.2 (±2.8)	∞	74 (-18)	95 (+10)
− Homomorphic Encryption	89.5 (±3.5)	0.7	76 (-16)	105 (+20)
− G-BHO Optimization	91.1 (±1.9)	0.7	88 (-4)	110 (+25)
Pairwise Combinations				
FL + DP only	85.0	2.1 [1]	52	130
FL + HE only	87.0	∞	74	140
FL + DP + HE only	89.8	1.2	70	125
DP + HE (no FL)	82.3	0.9	83	165
Individual Components				
FL only (baseline)	89.0	∞	68	112
DP only [2]	73.5	0.5	45	N/A [3]
HE only [2]	76.2	∞	62	N/A [3]
Progressive Addition				
FL → FL+DP	85.0	2.1	52	130
FL+DP → FL+DP+HE	88.7	1.2	81	120
FL+DP+HE → Full Hybrid	93.0	0.7	92	85

- [1]Higher ε value without HE due to lack of encrypted gradient protection
- [2]Centralized training with single privacy mechanism only
- [3]Not applicable for centralized training scenarios

This demonstrates the defensive advantage of combining multiple complementary privacy mechanisms. Statistical analysis (ANOVA) confirmed that differences in performance metrics between our hybrid approach and baseline methods were statistically significant ($p<0.01$) across all experimental conditions. Fig. 3. provides a comprehensive visual comparison of our hybrid approach against the three baseline methods across all evaluated metrics.



Fig. 3. Performance comparison of the proposed hybrid privacy-preserving approach versus baseline methods. The hybrid approach demonstrates superior performance across all metrics: model accuracy (93%), privacy guarantee (ε=0.7), attack resistance (92%), and computational efficiency (2.5x overhead compared to 8x for HE-only).

Our hybrid privacy-preserving approach demonstrates significant improvements over single-technique methods across multiple evaluation dimensions. We first discuss the implications of our results, then contextualize our findings within the broader privacy-preserving AI literature, and finally address limitations and future research directions. The superior model accuracy (93%) achieved by our hybrid approach, compared to federated learning alone (89%), differential privacy alone (85%), and homomorphic encryption alone (87%), confirms our hypothesis that complementary privacy mechanisms can enhance rather than degrade model performance when properly integrated. This finding challenges the conventional assumption that privacy preservation necessarily comes at the cost of utility. The reduced convergence time (85 epochs versus 112-140 epochs) further suggests that our approach creates a more stable optimization landscape. Our results align with recent work by Zhang et al. (2024), who reported similar accuracy improvements when combining differential privacy with secure aggregation, though our implementation extends this by incorporating fully homomorphic encryption. The enhanced resistance to membership inference attacks (92% versus 68% for federated learning alone) suggests that our layered defense strategy effectively addresses vulnerabilities that persist in single-technique approaches.

An unexpected finding was the relatively modest communication overhead (120MB/round) despite the integration of homomorphic encryption, which traditionally increases message size substantially. This efficiency can be attributed to the G-BHO optimization algorithm, which reduced parameter exchange volume by 35% compared to standard implementations. This result is particularly significant as communication cost often represents the primary bottleneck in distributed privacy-preserving systems.

6. Conclusion and Future Scope

We have applied our framework in practice in the healthcare sector, and our MIMIC-III and COVID-19 experiments are the above examples. The approach also allows organizations to liaise with institutions regarding high-sensitivity information of patients with good privacy guarantees (0.7). Although the means of its derivation may be equally applicable in such sensitive fields and information privacy as finance and IoT, the latter should also be defended at the domain level, which is beyond the remit of this paper. However, several limitations must be acknowledged. First, while we reduced computational overhead compared to homomorphic encryption alone, our approach still requires 2.5 times the computing resources of standard federated learning, potentially limiting deployment in resource-constrained environments. Second, our evaluation focused on classification tasks with structured and imaging data; performance characteristics may differ for other learning paradigms such as reinforcement learning or natural language processing. Third, while our privacy guarantees are strong ($\varepsilon = 0.7$), they still represent a trade-off that may be insufficient for extremely sensitive applications requiring stronger guarantees.

Future research should explore hardware acceleration techniques to further reduce computational overhead, potentially making our approach viable for edge devices with limited resources. Additionally, extending our framework to support other machine learning paradigms beyond supervised learning would broaden its applicability. Finally, investigating quantum-resistant variants of our cryptographic components would ensure long-term security against emerging computational threats. In conclusion, our hybrid privacy-preserving framework represents a significant advancement over single-technique approaches by demonstrating that complementary privacy mechanisms can be integrated to enhance both security and utility. This work contributes to the growing body of evidence suggesting that the perceived trade-off between privacy and utility can be substantially mitigated through careful system design and optimization.

Author Contributions Statement

Kummagoori Bharath – Data curation, formal analysis, methodology, resources, software, visualization, original draft

Pooja Chopra – investigation, project administration, visualization, original draft, review, and editing

Mukesh Kumar – validation, original draft, review, and editing

All authors have read and agreed to the published version of the manuscript.

Conflict of Interest Statement

The authors declare no conflicts of interest.

Funding Declaration

The authors declare that no funding was received for this study.

Data Availability Statement

This study analyzed publicly available datasets. The results obtained and datasets can be found here: “https://physionet.org/content/mimiciii/1.4/” and “https://github.com/ieee8023/covid-chestxray-dataset”, accessed on “15/04/2025”.

Ethical Declarations

This study did not involve human participants or animals. Therefore, ethical approval was not required.

Declaration of Generative AI in Scholarly Writing

During the preparation of this manuscript, the author(s) used OpenAI to improve the clarity and readability of selected paragraphs. The generated text was edited and revised as needed, and all content remains the responsibility of the author(s). No AI tool is listed as an author.

Abbreviations

The following abbreviations are used in this manuscript:

AI - Artificial Intelligence
G-BHO - Grasshopper-Black Hole Optimization
CKKS - Cheon–Kim–Kim–Song
IEEE - Institute of Electrical and Electronics Engineers
FL - Federated Learning
RMSE - Root Mean Square Error
NIST - National Institute of Standards and Technology
DP - Differential Privacy
MIMIC - Medical Information Mart for Intensive Care

References

[1] S. Dutta et al., "Federated Learning with Quantum Computing and Fully Homomorphic Encryption: A Novel Computing Paradigm Shift in Privacy-Preserving ML," arXiv:2409.11430, Sep. 2024. doi:10.48550/arXiv.2409.11430.

[2] H. Huang, B. Iskandarov, M. Rahman, H. T. Otal and M. A. Canbaz, "Federated Learning in Adversarial Environments: Testbed Design and Poisoning Resilience in Cybersecurity," arXiv:2409.09794, Sep. 2024. doi:10.48550/arXiv.2409.09794.

[3] R. Madduri, Z. Li, T. Nandi, K. Kim, M. Ryu and Á. Rodríguez, "Advances in Privacy Preserving Federated Learning to Realize a Truly Learning Healthcare System," IEEE TPS-ISA, pp. 273-282, Oct. 2024. doi:10.1109/TPS-ISA62245.2024.00039.

[4] Y. Bi, Y. Li, X. Feng and X. Mi, "Enabling Privacy-Preserving Cyber Threat Detection with Federated Learning," arXiv:2404.05130, Apr. 2024. doi:10.48550/arXiv.2404.05130.

[5] B. Yurdem, M. Kuzlu, M. K. Güllü, F. Ö. Çatak and M. Tabassum, "Federated Learning: Overview, Strategies, Applications, Tools and Future Directions," Heliyon, vol. 10, no. 19, e38137, Sep. 2024. doi:10.1016/j.heliyon.2024.e38137.

[6] B. Liu, N. Lv, Y. Guo and Y. Li, "Recent Advances on Federated Learning: A Systematic Survey," Neurocomputing, vol. 597, art. 128019, Jun. 2024. doi:10.1016/j.neucom.2024.128019.

[7] M. R. Uddin, G. Shankar, S. H. Mukta, P. Kumar, and A. K. M. N. Islam, "Evolving Topics in Federated Learning: Trends and Emerging Directions for IS," arXiv:2409.15773, Sep. 2024. doi:10.48550/arXiv.2409.15773.

[8] M. Bharathi, T. A. S. Srinivas and M. Bhuvaneswari, "Federated Learning: From Origins to Modern Applications and Challenges," Journal of Information Technology and Cryptography, vol. 1, no. 2, pp. 29-38, Oct. 2024. doi:10.48001/joitc.2024.1229-38.

[9] Ahmed, B., Shuja, M., Mishra, H. M., Qtaishat, A., & Kumar, M. (2023, March). IoT based smart systems using artificial intelligence and machine learning: accessible and intelligent solutions. In 2023 6th International Conference on Information Systems and Computer Networks (ISCON) (pp. 1-6). IEEE.

[10] H. Shin, K. Ryu, J. Kim, and S. Lee, "Application of Privacy Protection Technology to Healthcare Big Data," Digital Health, vol. 10, Jan. 2024. doi:10.1177/20552076241282242.

[11] S. K. M. et al., "Privacy-Preserving in Blockchain-Based Federated Learning Systems," Computer Communications, vol. 222, pp. 38-54, Apr. 2024. doi:10.1016/j.comcom.2024.04.024.

[12] Y. Dong, Y. Wang, M. Gama, M. Mustafa, G. Deconinck, and X. Huang, "Privacy-Preserving Distributed Learning for Residential Short-Term Load Forecasting," IEEE Internet of Things Journal, vol. 11, no. 9, pp. 16817-16831, Feb. 2024. doi:10.1109/JIOT.2024.3362587.

[13] Bhardwaj, V., Noonia, A., Chaurasia, S., Kumar, M., Rashid, A., & Othman, M. T. B. (2024). Optimizing structured data processing through robotic process automation. arXiv preprint arXiv:2408.14791.

[14] K. Daly, H. Eichner, P. Kairouz, H. B. McMahan, D. Ramage and Z. Xu, "Federated Learning in Practice: Reflections and Projections," IEEE TPS-ISA, pp. 148-159, Oct. 2024. doi:10.1109/TPS-ISA62245.2024.00026.

[15] W. Wei and L. Liu, "Trustworthy Distributed AI Systems: Robustness, Privacy, and Governance," ACM Computing Surveys, Feb. 2024. doi:10.1145/3645102.

[16] S. Banerjee, P. Sahu, M. Luo, A. Vahldiek-Oberwagner, N. J. Yadwadkar and M. Tiwari, "SoK: A Systems Perspective on Compound AI Threats and Countermeasures," arXiv:2411.13459, Nov. 2024. doi:10.48550/arXiv.2411.13459.

[17] S. S. Sefati, R. Crăciunescu, B. Arasteh, S. Halunga, O. Fratu and I. Tal, "Cybersecurity in a Scalable Smart City Framework Using Blockchain and Federated Learning for Internet of Things (IoT)," Smart Cities, vol. 7, no. 5, pp. 2802-2823, Oct. 2024. doi:10.3390/smartcities7050109.

[18] R. Sun et al., "Multi-Continental Healthcare Modelling Using Blockchain-Enabled Federated Learning," arXiv:2410.17933, Oct. 2024. doi:10.48550/arXiv.2410.17933.

[19] S. Narkedimilli, R. A. Kumar, N. Kumar, R. P. Reddy, and C. P. Kumar, "FL-DECO-BC: A Privacy-Preserving, Provably Secure, and Provenance-Preserving Federated Learning Framework with Decentralized Oracles on Blockchain for VANETs," arXiv:2407.21141, Jul. 2024. doi:10.48550/arXiv.2407.21141.

[20] W. Guo, F. Zhuang, X. Zhang, Y. Tong, and J. Dong, "A Comprehensive Survey of Federated Transfer Learning: Challenges, Methods and Applications," Frontiers of Computer Science, vol. 18, no. 6, Jul. 2024. doi:10.1007/s11704-024-40065-x.

[21] C. Papadopoulos, K.-F. Kollias, and G. F. Fragulis, "Recent Advancements in Federated Learning: State of the Art, Fundamentals, Principles, IoT Applications and Future Trends," Future Internet, vol. 16, no. 11, p. 415, Nov. 2024. doi:10.3390/fi16110415.

[22] H. U. Manzoor, A. Jafri and A. Zoha, "Adaptive Single-Layer Aggregation Framework for Energy-Efficient and Privacy-Preserving Load Forecasting in Heterogeneous Federated Smart Grids," Internet of Things, art. 101376, Sep. 2024. doi:10.1016/j.iot.2024.101376.

[23] H. Guan, P. Yap, A. Bozoki and M. Liu, "Federated Learning for Medical Image Analysis: A Survey," Pattern Recognition, vol. 151, art. 110424, Mar. 2024. doi:10.1016/j.patcog.2024.110424.

[24] Y. Chen, Z. Tang, T. Lu, B. Zhang, Z. Shi, and Z. Wang, "Accelerating Private Large Transformers Inference Through Fine-Grained Collaborative Computation," arXiv:2412.16537, Dec. 2024. doi:10.48550/arXiv.2412.16537.

[25] M. Izabachène and J.-P. Bossuat, "TETRIS: Composing FHE Techniques for Private Functional Exploration Over Large Datasets," arXiv:2412.13269, Dec. 2024. doi:10.48550/arXiv.2412.13269.

[26] M. Shrestha, Y. Ravichandran, and E. Kim, "Secure Multiparty Generative AI," arXiv:2409.19120, Sep. 2024. doi:10.48550/arXiv.2409.19120.

[27] T. Sattarov, M. Schreyer and D. Borth, "FedTabDiff: Federated Learning of Diffusion Probabilistic Models for Synthetic Mixed-Type Tabular Data Generation," arXiv:2401.06263, Jan. 2024. doi:10.48550/arXiv.2401.06263.

[28] R.-J. Yew, L. Qin and S. Venkatasubramanian, "You Still See Me: How Data Protection Supports the Architecture of ML Surveillance," arXiv:2402.06609, Feb. 2024. doi:10.48550/arXiv.2402.06609.

[29] D. Commey, S. Hounsinou and G. V. Crosby, "Securing Health Data on the Blockchain: A Differential Privacy and Federated Learning Framework," arXiv:2405.11580, May 2024. doi:10.48550/arXiv.2405.11580.

[30] Y. Chen and P. Esmaeilzadeh, "Generative AI in Medical Practice: In-Depth Exploration of Privacy and Security Challenges," Journal of Medical Internet Research, vol. 26, Mar. 2024. doi:10.2196/53008.
[31] A. Nash, "Decentralized Health Intelligence Network (DHIN)," arXiv:2408.06240, Aug. 2024. doi:10.48550/arXiv.2408.06240.
[32] K. Ranaweera et al., "Adaptive Clipping for Privacy-Preserving Few-Shot Learning: Enhancing Generalization with Limited Data," arXiv:2503.22749, 2025. doi:10.48550/arXiv.2503.22749.
[33] R. Shokri, M. Stronati, C. Song and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," in Proc. IEEE Symposium on Security & Privacy, pp. 3-18, 2017. doi:10.1109/SP.2017.41.

Authors' Profiles



Kummagoori Bharath is a Ph.D. candidate in Computer Science & Engineering at Lovely Professional University, Punjab. He obtained his M.Tech. in Computer Science & Engineering from Jawaharlal Nehru Technological University Anantapur (JNTUA). His current research centres on federated learning and broader machine-learning techniques, with a focus on privacy-preserving model training and distributed intelligence . Alongside his doctoral studies, he balances part-time professional responsibilities while exploring post-Ph.D. career pathways in academia and industry.



Dr. Pooja Chopra is an accomplished Associate Professor with over 18 years of academic and research experience, specializing in Computer Science. Completed Ph.D. in 2020 with a focus on cloud computing and intelligent systems. Demonstrated research excellence through 10+ Scopus-indexed publications in reputed international journals and conferences. Adept at mentoring students, leading academic initiatives, and contributing to curriculum development. Actively engaged in interdisciplinary research with a strong interest in AI, cloud resource provisioning, and emerging technologies.



Dr. Mukesh Kumar is an Associate Professor at Advanced Centre of Research & Innovation (ACRI), CGC University Mohali-140307, Punjab, India, and a Adjunct Associate Professor in the Faculty of Law, Solar University Oman. He holds a Ph.D. in Computer Science from Himachal Pradesh University, Shimla, where he specialized in designing ensemble and hybridized data mining techniques for analyzing academic performance. With over 15 years of experience in academia, Dr. Kumar has held various teaching positions, including Assistant Professor at Lovely Professional University and Chitkara University. He has contributed to various international conferences and holds intellectual property rights for several innovations.

How to cite this paper: Kummagoori Bharath, Pooja Chopra, Mukesh Kumar, "Artificial Intelligence (AI) Based Multi-Layered Approaches for Privacy Preservation in Federated Learning", International Journal of Modern Education and Computer Science(IJMECS), Vol.18, No.2, pp. 190-202, 2026. DOI:10.5815/ijmecs.2026.02.12