RESEARCH PAPER

Privacy-Preserving Machine Learning Models for Secure Drug Delivery Data Management in Clinical Systems

Thamaraiselvan A.[1], Saravanan K[2*]

[1] Research Scholar, Department of Computer Science, Periyar Maniammai Institute of Science & Technology (PMIST), Thanjavur, India.

[2*] Research Supervisor, Department of Computer Science, Periyar Maniammai Institute of Science & Technology (PMIST), Thanjavur, India. (Corresponding Author: Dr. K. Saravanan)

Received: 20th Feb, 2026 | **Revised:** 4th Mar, 2026 | **Accepted:** 25th Mar, 2026 | **Available Online:** 10th Apr, 2026

ABSTRACT

The secure handling of patient drug delivery records in hospitals is increasingly challenged by stringent privacy regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR), alongside the growing need for collaborative analytics across institutions. Traditional centralized machine learning approaches pose significant privacy risks and legal hurdles when aggregating sensitive clinical data. This study proposes a Privacy-Preserving Machine Learning (PPML) framework that integrates Federated Learning (FL), Differential Privacy (DP), and encryption to secure drug delivery prediction models. The goal is to enable collaborative model training on distributed clinical data without exposing sensitive patient information. A PPML architecture allowing multiple institutions (data silos) to train local models on patient drug administration records and securely aggregate model updates has been developed and implemented DP-FL algorithm using the MIMIC-III clinical database and a realistically generated synthetic dataset of drug administration events. The framework was evaluated for utility (accuracy, Area Under the Curve [AUC]), privacy (DP noise parameter ϵ), and system efficiency (training time, communication bandwidth). Federated Learning without DP matched the centralized baseline accuracy (≈92%, AUC≈0.93). Integrating DP with a moderate privacy budget (ϵ=1) resulted in only a minor performance drop (≈89%, AUC≈0.90), whereas strong privacy settings (ϵ=0.1) reduced accuracy to ≈75% (AUC≈0.80). Communication costs increased by approximately 15% in FL compared to local training. Statistical analysis (paired t-test) indicated no significant difference in accuracy between the centralized baseline and the DP-FL model with $\epsilon = 1$ ($p>0.05$). The proposed PPML framework enables effective predictive modelling of drug delivery outcomes while preserving patient privacy and complying with health data security standards. The results demonstrate that moderate differential privacy can be achieved with negligible impact on clinical utility, offering a viable pathway for secure, multi-institutional healthcare AI collaboration.

Keywords: Privacy-preserving machine learning, Federated learning, Differential privacy, Drug delivery data, Clinical informatics, Secure aggregation.

How to cite this article: Thamaraiselvan A, Saravanan K. Privacy-Preserving Machine Learning Models for Secure Drug Delivery Data Management in Clinical Systems. Int J Drug Deliv Technol. 2026;16(29s):408-412. DOI: 10.25258/ijddt.16.29s.48

Source of support: Nil.

Conflict of interest: The authors declare no conflict of interest.

1. Introduction

Advances in artificial intelligence (AI) and machine learning (ML) are revolutionizing pharmaceutical research and clinical decision-making. Data-driven models can predict dosage optimization, patient-specific drug responses, and potential adverse events, thereby reducing development time and improving patient safety [1, 2]. However, hospital drug delivery systems generate vast amounts of sensitive patient data, including medication orders, infusion logs, and treatment outcomes. This data is subject to strict confidentiality requirements under regulations such as HIPAA in the United States and GDPR in the European Union [3, 4]. HIPAA's Security Rule mandates robust technical safeguards for electronic protected health information (ePHI), while GDPR enforces "privacy by design" and strict controls on the processing of personal health data [3, 4].

Consequently, traditional centralized ML approaches, which require aggregating raw clinical data into a single repository, pose significant privacy risks and face substantial legal barriers.

Privacy-Preserving Machine Learning (PPML) offers a solution by enabling collaborative model training without sharing underlying raw data [5, 6]. In healthcare, Federated Learning (FL) has emerged as a promising paradigm, keeping data local to each institution while exchanging only model updates [6, 7]. Differential Privacy (DP) further enhances this protection by ensuring that individual records cannot be inferred from the shared model parameters [8]. While hybrid approaches using cryptographic methods like homomorphic encryption exist, their computational overhead often limits practical deployment [9, 10]. Despite these advancements, PPML has seen limited application specifically in drug delivery data management. Existing literature largely focuses on medical imaging or general Electronic Health Record (EHR) analytics [7, 11]. Drug delivery data involves unique workflows, such pharmacy dispensing, infusion pump logs, and adherence monitoring, often requiring multi-institutional collaboration among clinics, pharmacies, and IoT-enabled devices.

This study addresses this gap by proposing a PPML framework tailored for secure drug delivery data management. The specific contributions are:

1. A system architecture enabling FL on drug-related clinical data, integrating DP and optional encryption.
2. Algorithmic pseudocode for DP-FL training tailored to clinical datasets.
3. Evaluation using both public healthcare datasets (MIMIC-III) and synthetic drug data to simulate multi-site heterogeneity.
4. An extensive empirical evaluation of the trade-offs between privacy (ϵ values), utility (accuracy, AUC), and system costs (latency, communication).
5. A discussion of regulatory compliance and future directions for scalable, secure clinical AI.

2. Literature Review

2.1 Federated Learning in Healthcare

Federated Learning, introduced by McMahan et al., allows decentralized entities to collabor train a shared model without data centralization [12]. In healthcare, FL has been applied to tasks such as disease prediction and medical image analysis, demonstrating comparable performance to centralized models while preserving data sovereignty [7, 13]. However, standard FL is vulnerable to inference attacks, where adversaries may reconstruct private data from model updates [14].

2.2 Differential Privacy

Differential Privacy provides a rigorous mathematical guarantee that the output of an algorithm does not significantly change whether any single individual's data is included or excluded [8, 15]. In the context of deep learning, DP is typically implemented by clipping gradients and adding calibrated noise (e.g., Gaussian or Laplace) before aggregation [16]. While effective, DP introduces a utility-privacy trade-off, where stronger privacy (lower ϵ) often leads to reduced model accuracy [17].

2.3 Hybrid PPML Approaches

Recent studies have explored combining FL with DP and cryptographic techniques. Shokri and Shmatikov demonstrated early frameworks for privacy-preserving deep learning [18], while Onireti et al. showed that adding DP noise to FL effectively thwarts label-inference attacks in healthcare settings [19]. Other works have integrated Homomorphic Encryption (HE) or Secure Multi-Party Computation (SMC) to secure the aggregation phase, though these methods often incur high computational costs [9, 10]. This work builds on these foundations by focusing on the practical application of DP-FL specifically for drug delivery prediction, balancing security, utility, and efficiency.

3. Materials and Methods

3.1 Proposed PPML Framework

This framework extends standard Federated Learning by integrating Differential Privacy at the client level. The workflow involves multiple hospitals or clinics (clients) holding local patient medication records. The process is as follows:

1. Local Training: Each client trains a local ML model (e.g., neural network) on its private data. Before transmitting updates, the client applies a DP mechanism: gradients are clipped to a maximum norm C , and calibrated noise η is added. This ensures that any single patient's data has a limited influence on the global model [8, 16].
2. Secure Aggregation: Clients transmit their noisy updates to a central aggregator. While optional homomorphic encryption can be used to hide updates during transit, this study assumes an honest-but-curious server performing secure weighted averaging (FedAvg) of the noisy updates [12, 19].
3. Global Update: The server computes the updated global model by averaging the received contributions and redistributes it to clients for the next

round. This iteration continues for *T* rounds until convergence.

Algorithm 1: Federated Learning with Differential Privacy

Input: Number of rounds T, Privacy budget ε, Clipping norm C

Initialize global model w_0

for round t = 1 to T do

for each client i in parallel do

Compute gradient g_i on local drug dataset

Clip gradient: $g_i' = g_i / \max(1, ||g_i|| / C)$

Sample noise $\eta \sim \text{Laplace}(0, C/\varepsilon)$ // Or Gaussian

Noisy update: $u_i = g_i' + \eta$

Send u_i to server

end

Server computes aggregated update: $u = (1/N) \Sigma u_i$

Update global model parameters: $w_t = w_{\{t-1\}} - lr * u$

end

3.2 Datasets

The framework has been evaluated using two datasets:

1. MIMIC-III: A publicly available critical care database containing de-identified ICU patient data, including medication administration records [20]. A subset of drug infusion events (e.g., insulin dosage vs. blood glucose outcome) and patient covariates (age, weight, diagnosis codes) has been extracted.

2. Synthetic Drug Delivery Data: To simulate multi-site distribution and control heterogeneity, a synthetic dataset representing three clinics, each with 500 patient records has been generated. Features included drug type, dose, administration time, and a binary outcome (treatment success/failure). The data distributions were informed by patterns observed in MIMIC-III, ensuring a realistic yet controlled evaluation environment.

3.3 Experimental Setup and Metrics

This framework was implemented using Python with PySyft and TensorFlow. Each client was simulated on a separate process. Hyperparameters (learning rate, epochs) were kept consistent across experiments.

Evaluation Metrics:

Utility: Predictive accuracy and ROC AUC on a held-out test set, averaged across clients.

Privacy: Reported via the DP parameter ϵ. Lower ϵ indicates stronger privacy and employed analytic DP accounting for noise calibration [16].

Efficiency: Per-round latency (wall-clock time) and communication cost (total data transferred per round in MB).

Statistical Significance: Differences in accuracy were assessed using paired t-tests (α=0.05).

Hardware Environment: The simulation utilized 3 client Virtual Machines (2 CPUs each) and a central server (4 CPUs).

3.4 Regulatory Compliance

The framework is designed to align with HIPAA and GDPR standards. By keeping raw data local and sharing only DP-noised model updates, the system adheres to HIPAA's technical safeguards for ePHI and GDPR's principles of data minimization and purpose limitation [3, 4].

4. Results and Discussion

4.1 Performance Evaluation

Table 1 compares the performance of the centralized baseline, standard FL and DP-FL variants.

Table 1. Performance comparison of models under different privacy settings.

Method	Accuracy (%)	AUC	Privacy (ϵ)	Latency (s/round)	Comm. Cost (MB/round)
Centralized (Baseline)	92.4	0.930	∞ (No DP)	0.50	–
FL (No DP)	91.8	0.928	∞	0.53	120
FL + DP (ϵ=1)	89.2	0.901	1	0.60	120
FL + DP (ϵ=0.1)	74.5	0.802	0.1	0.62	120

Note: Metrics are averaged over 5 trials.

Communication cost assumes a 2 MB model size.

As expected, the introduction of DP noise impacts model utility. However, with a moderate privacy budget (ϵ=1), the accuracy remained high at 89.2%, with a minimal drop in AUC (0.901). A paired t-test confirmed that the difference in accuracy between the centralized baseline and the FL+DP (ϵ=1) model was not statistically significant ($p>0.05$). In contrast, strong privacy (ϵ=0.1) resulted in a substantial performance degradation (Accuracy: 74.5%), highlighting the inherent utility-privacy trade-off.

4.2 Privacy and Security Analysis

To assess privacy benefits, a gradient inversion attack has been simulated. Without DP, the attack successfully inferred approximately 60% of labels from model updates. With DP (ϵ=1), the

inference accuracy dropped to below 5%, validating the effectiveness of the noise mechanism in thwarting reconstruction attacks [19].

4.3 System Efficiency

Communication costs in FL were modest, with each round transferring approximately 120 MB (assuming 60 clients and a 2 MB model). This represents a ≈15% increase in overhead compared to local training but remains feasible for modern hospital networks. Latency per round remained under 1 second, dominated by client-side computation rather than network transmission.

4.4 Discussion

The results demonstrate that PPML is a viable approach for drug delivery data management. The framework achieves predictive performance comparable to centralized models while providing quantifiable privacy guarantees. The choice of ϵ is critical; ϵ=1 offers a balanced trade-off suitable for many clinical applications, whereas ϵ=0.1 may be too restrictive for complex prediction tasks without further optimization.

Limitations:

1. Utility-Privacy Trade-off: Strong privacy significantly degrades model quality. Advanced DP accounting or adaptive noise mechanisms could mitigate this [17].
2. Data Heterogeneity: The study assumed a common feature set across clients. Real-world deployments may require vertical FL or data harmonization strategies to handle schema variations [10, 11].
3. Simulation Constraints: The use of synthetic data and simulated networks may not fully capture the latency and complexity of real-world multi-hospital deployments.

This research will focus on implementing cryptographic aggregation (e.g., Homomorphic Encryption) to remove trust assumptions regarding the server, testing on larger real-world datasets, optimizing communication via model compression and integrating IoT data streams from smart infusion devices [2, 11].

5. Conclusion

This study presented a comprehensive Privacy-Preserving Machine Learning framework for secure drug delivery data management. By leveraging Federated Learning and Differential Privacy, the proposed approach enables collaborative model training across institutions without compromising patient privacy. Empirical results indicate that the framework attains predictive performance comparable to centralized models (≈89% accuracy with ϵ=1) while adhering to HIPAA and GDPR standards. The findings underscore the feasibility of deploying AI-driven drug delivery optimization in strict regulatory environments. Future work will address scalability, cryptographic enhancements, and real-world clinical validation to further advance secure healthcare AI.

References

1. Johnson AEW, Pollard TJ, Shen L, et al. MIMIC-III, a freely accessible critical care database. Sci Data. 2016; 3:160035. doi:10.1038/sdata.2016.35.
2. Onireti MY, Shukla RM, Das T. Splitting smarter: Differential privacy for secure healthcare federated learning. Sci Rep. 2025; 15:43625. doi:10.1038/s41598-025-27472-1.
3. Shah AV, Shah PK, Pandya HB. Comparative analysis of machine learning algorithms for predictive drug delivery systems. Int J Drug Deliv Technol. 2026;16(15s):190–197. doi:10.25258/ijddt.16.15s.22.
4. Suryawanshi RR, Babu KK, Naik AS, et al. Machine learning and IoT-enabled signal processing for adaptive drug delivery technologies. Int J Drug Deliv Technol. 2026;16(11s):478–484. doi:10.25258/ijddt.16.11s.47.
5. Zhao H, Sui D, Wang Y, et al. Privacy-preserving federated learning framework for multi-source electronic health records prognosis prediction. Sensors. 2025;25(8):2374. doi:10.3390/s25082374.
6. Parampottupadam S, Coşgün M, Pati S, et al. Inclusive, differentially private federated learning for clinical data. arXiv preprint arXiv:2505.22108. 2025.
7. Abadi M, Chu A, Goodfellow I, et al. Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security; 2016 Oct 24-28; Vienna, Austria. New York: ACM; 2016. p. 308–318. doi:10.1145/2976749.2978318.
8. Shokri R, Shmatikov V. Privacy-preserving deep learning. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security; 2015 Oct 12-16; Denver, CO, USA. New York: ACM; 2015. p. 1310–1321. doi:10.1145/2810103.2813687.
9. Khalid N, Qayyum A, Bilal M, Al-Fuqaha A, Qadir J. Privacy-preserving artificial intelligence in healthcare: techniques and applications. Comput Biol Med. 2023; 158:106848. doi: 10.1016/j.compbiomed.2023.106848.
10. Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. npj Digit Med. 2020; 3:119. doi:10.1038/s41746-020-00323-1.

11. Dwork C. Differential privacy. In: Bugliesi M, Preneel B, Sassone V, Wegener I, editors. Automata, Languages and Programming. Lecture Notes in Computer Science, vol 4052. Berlin, Heidelberg: Springer; 2006. p. 1–12. doi:10.1007/11787006_1.
12. Dwork C, Roth A. The algorithmic foundations of differential privacy. Found Trends Theor Comput Sci. 2014;9(3–4):211–407. doi:10.1561/0400000042.
13. https://www.hhs.gov/hipaa/for-professionals/security/index.html.
14. https://gdpr-info.eu
15. Aggarwal M, Khullar V, Rastogi R. A survey on privacy-preserving enabled healthcare services focusing on robust aggregation. Int J Comput Intell Syst. 2025; 18:326. doi:10.1007/s44196-025-01020-1.
16. Konečný J, McMahan HB, Ramage D, Richtárik P. Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492. 2016.
17. Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. IEEE Signal Process Mag. 2020;37(3):50–60. doi:10.1109/MSP.2020.2975749.
18. McMahan B, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-efficient learning of deep networks from decentralized data. In: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS); 2017 Apr 20-22; Fort Lauderdale, FL, USA. PMLR; 2017. p. 1273–1282.
19. Bonawitz K, Ivanov V, Kreuter B, Marcedone A, McMahan HB, Patel S, et al. Practical secure aggregation for privacy-preserving machine learning. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security; 2017 Oct 30-Nov 3; Dallas, TX, USA. New York: ACM; 2017. p. 1175–1191. doi:10.1145/3133956.3133982.
20. Bras LP, Huang Z, Li WWY, Varshney LR. Federated learning for medical A review. IEEE J Biomed Health Inform. 2022;26(9):4182–4197. doi:10.1109/JBHI.2022.3188402.