Quantum-Aware Privacy Risk Evaluation for Differentially Private Federated Healthcare Analytics

Neha Sharma[1], Prasenjit Chatterjee[2]

[1,2] Lincoln University College, 47301, Petaling Jaya, Selangor Darul Ehsan, Malaysia

[1]pdf.nehasharma@lincoln.edu.my, [2]dr.prasenjitchatterjee6@gmail.com

Abstract: Federated learning enables collaborative healthcare analysis without requiring centralized patient data; however, the updates to the machine learning model may contain private information. In this work, a privacy risk scoring framework considering quantum attacks is proposed for differentially private federated healthcare analytics by applying UCI Diabetes 130-US Hospitals dataset to predicting 30-day readmission. The evaluation approach utilizes differential privacy, secure aggregation, quantum communication simulation, membership inference attack, and Quantum-Aware Privacy Risk Score. From the results, conventional Federated ML had a High QPRS of 97.5000, while DP-FL, DP-FL with secure aggregation, and QADP-PQSA-based evaluation method obtained Low QPRS with equivalent F1-score. With fixed non-IID evaluation, when DP noise multiplier is set to 1.0, the metrics of accuracy, F1-score, MIA accuracy, and QPRS were 0.7101, 0.2676, 0.5022, and 3.0937, respectively.

Keywords: Federated Learning; Differential Privacy; Secure Aggregation; Post-Quantum Communication; Membership Inference Attack; Quantum-Aware Privacy Risk Score; Healthcare Analytics; Privacy-Preserving Machine Learning; QADP-PQSA

1. Introduction

Machine learning plays an essential role in healthcare analytics in the area of predicting health issues, risk stratification, readmissions analysis, and decision-making. Sensitive healthcare data are scattered among institutions which cannot easily be centralized due to privacy regulations and other reasons (governance, ethics, and re-identification). Federated learning allows several institutions to train a common model without transmitting data to the central server but does not entirely ensure privacy preservation: gradient updates and model predictions still expose patient-level information. Membership inference attack can even detect patients' participation in training. Therefore, privacy-preserving federated healthcare requires additional measures to ensure differential privacy (by using clipping and random noise), secure aggregation (which hides each institution's update while aggregating gradients from different hospitals). Moreover, in the future, there will be the threat of communication insecurity since federated learning assumes repeated communication between hospitals and a server. Classical cryptographic tools can fail against quantum adversaries, and post-quantum communication will be required to ensure long-term security. However, post-quantum security is meant to protect channels but not model behavior, so privacy evaluation should include utility, leakage, communication protection, and quantum risk. The work proposes quantum-aware privacy-risk evaluation for differentially private federated healthcare analytics. Differential privacy, secure aggregation, post-quantum communication security overhead simulation, membership inference attacks detection, and Quantum-Aware Privacy Risk Score computation are integrated into a single pipeline. The goal of the research is to provide risk-aware evaluation framework beyond accuracy and F1-score. The main contributions are the compact design of quantum-aware privacy-

risk evaluation for federated healthcare learning, unified approach that combines differential privacy, secure aggregation, post-quantum communication overhead simulation, and membership inference attack auditing, application to 30-day hospital readmission prediction using a real-world dataset, and demonstration of how Quantum-Aware Privacy Risk Score helps to compare privacy-preserving federated learning approaches

2. Related work

FL is popular in healthcare analytics due to the opportunity to train a common model in hospitals without transferring private data. The related works concern medical image processing, smart healthcare, clinical prediction, and collaboration across different institutions [1]–[3]. They have decreased requirements for data centralization, but problems remain – including non-IID data, client distribution imbalance, communication overhead, and distrust in the aggregator. Different approaches of differential privacy (DP) are included to restrict individual patients' leakage with gradient clipping and adding noise [4]–[6]. FL with DP allows obtaining privacy-preserving analytics in healthcare, but there are reports only about privacy utility trade-off and not leakage through membership inference attack or post-quantum vulnerability. Secure aggregation and verifiable FL make updates hidden from the server, but they cannot prevent leakage by membership inference [7]. Even membership inference and gradient leakage may disclose sensitive information during training without revealing raw data. Post-quantum FL research deals with quantum-resistant communication using lattice cryptography [8].

3. Proposed Methodology

3.1 Dataset

Prediction of hospital readmission after 30 days was done based on the UCI Diabetes 130-US Hospitals dataset. It comprises 101,766 entries and 50 attributes from various hospitals. Binary classes: readmission within 30 days (positive), and everything else (negative). Data imbalance: 11.16% positive cases. Data processing involved replacing missing value placeholders, removing identifiers, deleting attributes that have too many missing values, filling in numerical and categorical values, merging infrequent categories, one-hot encoding categoricals, and normalizing numerical attributes. Attributes that had more than 40% missing values (such as weight, medical_specialty, max_glu_serum, A1Cresult) were deleted. After preprocessing and one-hot encoding, 524 features remained. The dataset was split in 75:25 proportion to get the training and testing datasets with an equal distribution of classes. Training set includes 76,324 observations; testing set has 25,442. In order to use federated learning, the training dataset was split into five fictional hospital clients, with IID and Non-IID distribution considered.

3.2 Proposed Model

The proposed approach provides a framework for quantifying the privacy risk of a federated learning system in the presence of quantum computers. The pipeline starts with preprocessing the dataset through dealing with missing values, removal of high-missing-value features, categorical feature encoding, numerical feature scaling, and definition of a 30-day readmission objective. Experiment 1 involves conducting standard Federated ML where each client (hospital) trains locally and sends model parameters to the server. In Differentially Private Federated Learning (DP-FL), local models are trained under gradient clipping and adding Gaussian noise to decrease the individual record's effect. A secure aggregation is simulated to prevent the server from obtaining any sensitive information on individual records. To

mitigate future security risks, the post-quantum communication overhead of quantum-resistant exchange is estimated. The research cannot be viewed as an implementation of a post-quantum setting yet since its goal is to measure the overhead needed for the additional security layer. Privacy leakage is determined with membership inference attacks that involve inferring whether a record was part of the training set with the help of model-output indicators (confidence, entropy, probability distribution). Calculation of the privacy-risk score (Quantum-Aware Privacy Risk Score (QADP-PQSA)) follows to compare privacy risks in different scenarios. The comparison between standard Federated ML, DP-FL, DP-FL with secure aggregation, and the QADP-PQSA-based approach is provided in terms of predictive utility, membership inference attack success rate, post-quantum overhead and privacy-risk score. The workflow of the experiment is shown in Figure 1..

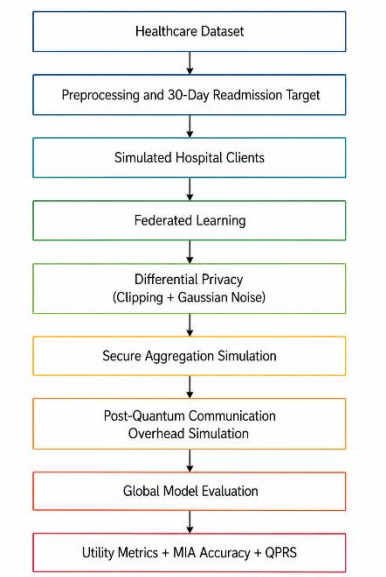

Figure 1. Quantum-aware privacy-risk evaluation pipeline for differentially private federated healthcare analytics.

4. Results and Discussion

The performance was analyzed in terms of predictive utility, MIA accuracy and Quantum Aware Privacy Risk Score (QPRS). The performance comparison between standard Federated ML, DP-FL, DP-FL with secure aggregation, and QADP-PQSA based system in IID FL is illustrated in Figure 2. While standard Federated ML exhibits the best performance in terms of F1-score (0.2704), it has the worst QPRS (97.5000), which means that even in case when only the locality is considered, the privacy risk remains high. With DP-FL, the QPRS decreases to 4.5703 and the F1-score slightly drops to 0.2671. Further, using secure aggregation reduces QPRS to 3.8672 while the F1 score is 0.2673. In case of QADP-PQSA based evaluation, the minimum QPRS of 2.5312 is achieved while F1 0.2673 and MIA accuracy of 0.5065 (near random) is maintained. The performance indicates that privacy-preserving pipeline can preserve privacy risks without affecting utility.

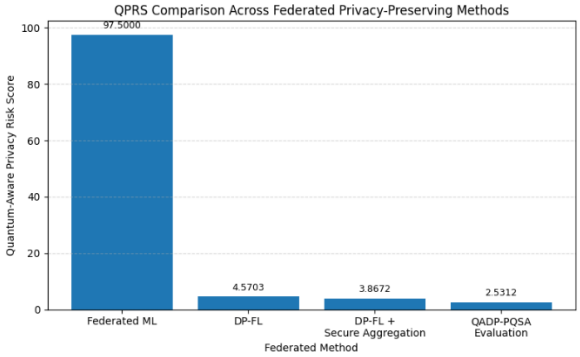

Figure 2. QPRS Comparison Across Federated Privacy-Preserving Methods

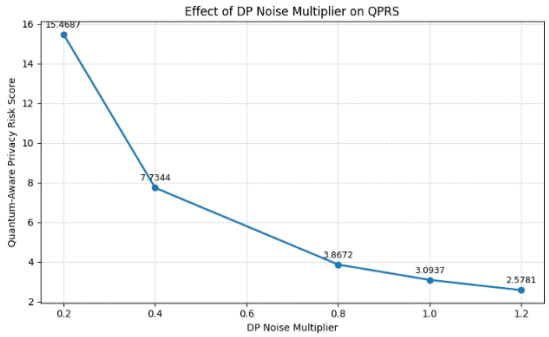

Figure 3. Effect of DP Noise on QPRS

As shown in Figure 3, QPRS decreased consistently as the DP noise multiplier increased. The score reduced from 15.4687 at a noise multiplier of 0.2 to 2.5781 at a noise multiplier of 1.2. A noise multiplier of 1.0 provided a balanced setting, achieving Low privacy risk with QPRS of 3.0937 while maintaining an F1-score of 0.2676 and MIA accuracy close to random guessing.

5. Conclusions

This work provides quantum-aware privacy-risk evaluation of differentially private federated healthcare analytics. The strengths of federated learning with the application of differential privacy, secure aggregation simulation, post-quantum communication overhead, membership inference auditing, and Quantum-Aware Privacy Risk Score (QPRS) are analyzed. Using the UCI Diabetes 130-US Hospitals dataset to predict 30-day readmission, it is demonstrated that federated learning produces comparable utility but high QPRS (97.5000), suggesting data locality alone does not guarantee privacy in federated healthcare analytics. Applying DP-FL, DP-FL with secure aggregation, and QADP-PQSA evaluation brings QPRS down to low-risk status while maintaining the same level of F1-score. The analysis of DP noise sensitivity revealed the optimal value of multiplier to be 1.0 with F1 0.2676, MIA accuracy 0.5022, and QPRS 3.0937. Privacy evaluation of federated healthcare analytics should focus on factors other than accuracy and F1-score because quantum-aware privacy risk score and auditing offer a more realistic privacy evaluation approach. Further research includes formal differential privacy accounting, complete implementation of post-quantum cryptography, privacy attack enhancement, use of multiple random seeds, and evaluation on additional datasets.

References

1. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of healthcare informatics research*, *5*(1), 1-19.
2. Li, J., Meng, Y., Ma, L., Du, S., Zhu, H., Pei, Q., & Shen, X. (2021). A federated learning based privacy-preserving smart healthcare system. *IEEE Transactions on Industrial Informatics*, *18*(3).
3. Adnan, M., Kalra, S., Cresswell, J. C., Taylor, G. W., & Tizhoosh, H. R. (2022). Federated learning and differential privacy for medical image analysis. *Scientific reports*, *12*(1), 1953.
4. Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., ... & Poor, H. V. (2020). Federated learning with differential privacy: Algorithms and performance analysis. IEEE transactions on information forensics and security, 15, 3454-3469.
5. Barnawi, A., Chhikara, P., Tekchandani, R., Kumar, N., & Alzahrani, B. (2024). A differentially privacy assisted federated learning scheme to preserve data privacy for IoMT applications. IEEE Transactions on Network and Service Management, 21(4), 4686-4700.
6. Xu, G., Li, H., Liu, S., Yang, K., & Lin, X. (2019). VerifyNet: Secure and verifiable federated learning. IEEE Transactions on Information Forensics and Security, 15, 911-926.
7. Hahn, C., Kim, H., Kim, M., & Hur, J. (2021). Versa: Verifiable secure aggregation for cross-device federated learning. IEEE Transactions on Dependable and Secure Computing, 20(1), 36-52.
8. Xu, P., Hu, M., Chen, T., Wang, W., & Jin, H. (2022). LaF: Lattice-based and communication-efficient federated learning. IEEE Transactions on Information Forensics and Security, 17, 2483-2496.