Adaptive RDP-FL: Enhancing Privacy-Preserving Federated Learning with Robust Differential Privacy Mechanisms

Ibtissem BEN OUHIBA[1] Zahra KODIA[2] and Nadia BEN AZZOUNA[3]

***Abstract*— Artificial Intelligence (AI) is revolutionizing information security, influencing both attack and defense strategies. Attackers leverage AI to automate cyberattacks and exploit vulnerabilities, while defenders utilize it for anomaly detection, predictive threat modeling, and automated responses. Federated Learning (FL), a privacy-preserving training method, remains vulnerable to inference attacks. To address this, we propose the Rényi Differential Privacy (RDP) based federated learning (RDP-FL) framework, which incorporates moment accounted noise scaling to dynamically regulate the privacy budget, achieving an optimal balance between privacy and utility. This method minimizes unnecessary noise addition while maintaining strong privacy guarantees, thereby preserving data integrity and enhancing model performance. Experimental validation on the Medical-MNIST and CIFAR-10 datasets demonstrates the effectiveness of RDP-FL, showing its ability to safeguard data privacy while ensuring high classification accuracy. This work advances the ongoing efforts to enhance cybersecurity in an AI-driven landscape.**

I. INTRODUCTION

Artificial Intelligence (AI) has become a pivotal force in modern information security, shaping both attack and defense strategies. While AI-driven methods enable attackers to automate cyber threats and exploit system vulnerabilities, they also strengthen cybersecurity [1] through automated anomaly detection, predictive threat modeling, and real-time threat response [2]. One of the most promising AI-driven privacy-preserving techniques is Federated Learning (FL) [3], which allows collaborative model training without exposing raw data. However, despite this enhanced data privacy, FL remains vulnerable to inference attacks, where adversaries attempt to extract sensitive information from model updates [4]. To address this vulnerability, FL is often combined with privacy-enhancing mechanisms such as Differential Privacy (DP). DP improves security by adding random noise to data or model updates, making it more difficult to identify individual contributions. However, excessive noise can negatively impact model accuracy, so finding the right balance between privacy and performance is essential.

In our approach, we integrate Rényi Differential Privacy (RDP) into the FL process to achieve a better trade-off between privacy and utility. Training starts locally on client devices, where each user computes model updates. These updates are then shuffled to break direct associations between users and their data. On the server side, we apply moment-accounted noise scaling before aggregation, making it harder for attackers to extract sensitive information. Finally, the server aggregates the processed updates to construct a secure global model. Although various privacy-preserving techniques like DP have been explored to enhance FL security [5], traditional DP methods often introduce excessive noise, which can significantly degrade model performance and utility [6]. Unlike traditional privacy-preserving approaches that apply fixed noise injection, our proposed framework leverages moment accounted noise scaling, a dynamic technique that adjusts the privacy budget adaptively based on model update sensitivity. This strategy optimally balances privacy protection and model utility, ensuring robust differential privacy guarantees while minimizing unnecessary noise perturbations. To validate our approach, we conduct comprehensive experiments on the Medical-MNIST and CIFAR-10 datasets, demonstrating that RDP-FL consistently outperforms existing privacy-preserving technique in both privacy protection and model accuracy. The results confirm the effectiveness of moment accounted noise scaling in mitigating inference risks while preserving high predictive performance. By reinforcing privacy in Federated Learning, our work contributes to the advancement of privacy-preserving AI systems, highlighting the need for robust differential privacy mechanisms in the evolving digital landscape.

Our contributions are summarized as follows:

- We propose RDP-FL, a novel Federated Learning framework that combines Rényi Differential Privacy with moment-accounted noise scaling for adaptive privacy budgeting.
- We enforce a gradient-sensitive, variance-aware noise mechanism to optimize the privacy-utility trade-off.
- We introduce a novel model based on moment-accounted noise scaling specifically for enhancing privacy in Federated Learning.
- We perform comprehensive experiments on benchmark datasets, showing that RDP-FL outperforms existing techniques such as SRR-FL in terms of both accuracy and privacy.

The rest of the paper is organized as follows. Section 2 provides a literature review. Section 3 outlined a detailed description of the proposed approach. Section 4 describes the implementation details and results. Section 5 concludes

*This work was not supported by any organization

[1]Ibtissem BEN OUHIBA is with University of Tunis, ISG Tunis, SMART-LAB, Cité Bouchoucha, 2000 Bardo, Tunis, Tunisia benouhiba.ibtissem@gmail.com

[2]Zahra KODIA is with University of Tunis, ISG Tunis, SMART-LAB, Cité Bouchoucha, 2000 Bardo, Tunis, Tunisia zahra.kodia@isg.rnu.tn

[1]Nadia BEN AZZOUNA is with University of Tunis, ESSECT, SMART-LAB, Cité Bouchoucha, 2000 Bardo, Tunis, Tunisia nadia.benazzouna@ensi.rnu.tn

the paper and presents future perspectives.

II. LITERATURE REVIEW

Federated Learning (FL) is a decentralized machine learning paradigm that enables multiple clients to collaboratively train models without exposing raw data [4]. While FL enhances data privacy by keeping information localized, it remains vulnerable to security threats such as inference attacks, gradient leakage, and data poisoning. To mitigate these risks, researchers have integrated Differential Privacy (DP) [5] into FL, injecting noise into model updates before aggregation [9] to enhance privacy protection while preserving model utility. However, adversaries can still infer sensitive information by analyzing model updates, as demonstrated by membership inference [7] and model inversion attacks [8], underscoring the need for stronger privacy mechanisms in FL. Traditional DP mechanisms ensure that adversaries cannot determine whether a particular data point was included in the training process, thereby mitigating privacy risks. Research in [10] introduced the Gaussian Mechanism-based DP-FL, where Gaussian noise is added to local model gradients before aggregation. While this technique provides strong privacy guarantees, it often degrades model accuracy due to excessive noise, making it difficult to balance privacy and utility effectively. To overcome this limitation, Staircase Randomized Response (SRR)-FL [4] was proposed, incorporating a staircase noise distribution to mitigate privacy leakage. While SRR-FL improves privacy protection, it suffers from high computational overhead, which hinders scalability in large-scale FL deployments. Similarly, Local Differential Privacy (LDP) in FL [11] prevents even the central server from extracting private information. However, LDP-based methods introduce significant noise at the client level, negatively impacting model convergence and overall performance. To enhance privacy guarantees while preserving model accuracy, Chamikara et al. [12] propose a LDP protocol designed for industrial settings in which data is locally perturbed on each device before training, ensuring privacy without the need for a trusted central server. Their approach offers stronger privacy guarantees than existing techniques and demonstrates high model performance even with small privacy budgets (e.g., $\epsilon = 0.5$). The study also addresses key challenges related to LDP implementation in FL, particularly in untrusted environments, with empirical evaluations validating its effectiveness. Recent advancements in differential privacy mechanisms for FL have further refined these approaches. Feng et al. [13] proposed a universally harmonized differential privacy mechanism, which optimizes noise distribution across different DP methods to improve both model accuracy and convergence speed. Their research underscores the importance of adaptive privacy mechanisms, demonstrating that dynamically adjusting privacy budgets can significantly enhance FL performance while ensuring strong privacy guarantees. Building on these advancements, we introduce Rényi Differential Privacy-based Federated Learning (RDP-FL), a novel framework that dynamically regulates privacy budgets using moment accounted noise scaling to achieve an optimal privacy-utility tradeoff. Unlike conventional approaches that apply fixed noise levels, RDP-FL intelligently adapts noise levels based on model update sensitivity, ensuring that privacy is preserved while minimizing unnecessary noise perturbations. This method provides rigorous differential privacy guarantees while improving model performance. By addressing these critical challenges, RDP-FL offers a scalable, efficient, and privacy-preserving solution for Federated Learning. Our framework not only strengthens protection against inference attacks but also enhances model performance, positioning it as a promising approach for secure and privacy-aware AI applications in distributed learning environments.

III. PROPOSED APPROACH: RÉNYI DIFFERENTIAL PRIVACY-BASED FEDERATED LEARNING (RDP-FL)

To address the trade-off between privacy and model performance in FL, we propose Rényi Differential Privacy-based Federated Learning (RDP-FL). Our approach introduces moment-accounted noise scaling, which sets it apart from traditional methods that adjust noise based on fixed schedules or iterations. Instead, RDP-FL dynamically adapts noise based on gradient variance, enabling a more responsive, data-aware privacy mechanism that improves privacy protection while maintaining model accuracy.

A. Rényi Differential Privacy (RDP)

Rényi Differential Privacy (RDP) [14] is a relaxation of Differential Privacy (DP) that provides a more flexible and analytically precise framework for quantifying privacy loss. Unlike traditional ϵ-DP, which enforces a strict upper bound on privacy leakage for a single interaction, RDP leverages Rényi divergence to track cumulative privacy loss more effectively, making it particularly useful for scenarios involving multiple iterative computations, such as machine learning, federated learning, and repeated queries on sensitive data.

RDP defines privacy loss using the α-Rényi divergence, which measures the distance between two probability distributions P and Q corresponding to model outputs when a single data point is included or removed. Mathematically, it is expressed as:

$$D_\alpha(P||Q) = \frac{1}{\alpha - 1} \log \mathbb{E}_{x \sim Q} \left[\left(\frac{P(x)}{Q(x)} \right)^\alpha \right] \tag{1}$$

where:

- $D_\alpha(P||Q)$ represents the Rényi divergence between distributions P and Q.
- α is the Rényi divergence order, which controls the trade-off between privacy and utility.
- $P(x)$ and $Q(x)$ are the probability distributions of outputs under two neighboring datasets.

One of the main advantages of RDP is its strong composition properties, enabling precise and efficient tracking of cumulative privacy loss over multiple computations. The RDP accountant accurately quantifies cumulative privacy loss across iterative computations, which is particularly beneficial

in scenarios involving repeated model updates, such as federated learning or deep learning applications [15]. Compared to traditional DP, RDP offers improved composition properties, providing a better balance between privacy protection and model utility. These strengths position RDP as an effective and powerful approach for privacy-preserving AI systems in practical, real-world scenarios.

B. Overview of RDP-FL

As shown in Figure 1, the RDP-FL framework integrates a parameter shuffling mechanism [4] to reduce privacy budget accumulation and moment accounted noise scaling to enhance privacy while preserving model utility. In this approach, multiple federated clients train their local models and generate model updates. Instead of sending these updates directly to the server, each client shuffles its model parameters to prevent any direct linkage between clients and their updates. After parameter splitting and shuffling, the updates are forwarded to the server-side, where moment accounted noise scaling is applied. This ensures that noise is injected dynamically based on gradient sensitivity before aggregation. By leveraging Rényi Differential Privacy (RDP) at the server level, privacy is preserved while maintaining good model accuracy. Once aggregated, the global model is updated and redistributed back to the clients for the next training round.

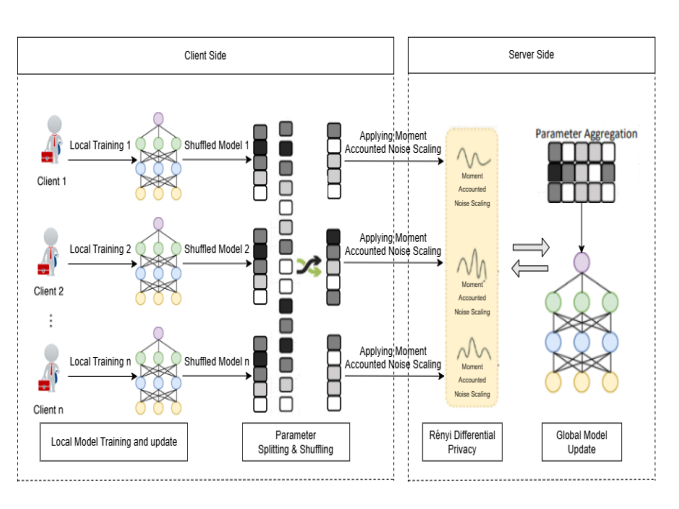

Fig. 1: Overview of our proposed model RDP-FL

1) Local Model Training and Parameter Shuffling: Each client performs local training on its private dataset without sharing raw data. Once training is complete, parameter updates are shuffled before transmission to the server. This step enhances privacy by breaking the direct correlation between clients and their updates, reducing the risk of inference attacks. The local model update follows:

$$w_i^{t+1} = \mathcal{F}(w_i^t, D_i, \theta) \tag{2}$$

where:

- w_i^t is the local model at client i at round t.
- $\mathcal{F}(\cdot)$ is the client specific update function.
- D_i is the local dataset at client i.
- θ represents the hyperparameters (e.g., learning rate, batch size, number of local iterations). The parameter updates are shuffled to obfuscate the origin of the updates, thereby reducing the risk of adversaries linking specific updates to individual clients. This shuffling mechanism is vital in maintaining privacy in a distributed setting, especially in the presence of inference attacks.

2) Moment Accounted Noise Scaling for Privacy Preservation: Federated Learning (FL) is a decentralized approach that reduces the exposure of sensitive client data by keeping data local. However, FL alone does not inherently ensure privacy. Traditional DP methods applied to FL typically use fixed noise injection, which may lead either to excessive perturbations that degrade model accuracy or insufficient noise that compromises privacy. To address this, we introduce a dynamic moment accounted noise scaling mechanism, which adapts noise levels based on model update sensitivity, thereby enhancing both privacy and model utility. This technique has been explored in other machine learning applications [16], but its direct application for privacy preservation in FL has not been extensively studied. Our approach integrates a dynamic privacy-aware noise scaling mechanism in RDP with FL before aggregation. Our method leverages RDP to compute an adaptive noise scale based on the sensitivity of model updates, allowing us to better regulate noise injection and improve privacy protection without sacrificing accuracy. The noise scale σ is computed as follows [14]:

$$\sigma = \frac{\Delta f}{\sqrt{2\alpha \log(1/\delta)}} \tag{3}$$

where Δf represents the sensitivity of the model update, defined as:

$$\Delta f = \max_{D,D'} \|\mathcal{F}(w, D, \theta) - \mathcal{F}(w, D', \theta)\| \tag{4}$$

with D and D' being neighboring datasets differing by a single data point, α the Rényi divergence order, and δ the privacy loss parameter. The cumulative privacy budget is tracked as:

$$\epsilon_T = \sum_{t=1}^{T} \frac{\Delta f_t}{\sqrt{2\alpha \log(1/\delta)}} \tag{5}$$

where Δf_t represents the gradient sensitivity at iteration t, dynamically tracking privacy expenditure to ensure more efficient noise injection while maintaining strict privacy guarantees. To further stabilize noise injection, we introduce a dynamic noise adjustment mechanism based on gradient variance [14], where the noise scale at iteration t is adapted as:

$$\sigma_t = \frac{\beta \sigma_{t-1}}{1 + \gamma \cdot \text{Var}(\nabla \mathcal{L}_t)} \tag{6}$$

where β is a decay factor, γ a learning rate-dependent parameter, and $\text{Var}(\nabla \mathcal{L}_t)$ the variance of gradients at iteration t.

Unlike static noise addition, which may either overprotect or underprotect model updates, this dynamic adaptation ensures that noise injection is proportional to gradient sensitivity changes, leading to a better balance between privacy and learning efficiency.

Our approach differs from previous works in three key ways. First, this is the first application of moment-accounted noise scaling for privacy protection in FL, whereas previous studies have applied this technique in other machine learning contexts without focusing explicitly on privacy in FL. Second, unlike existing work [4] that uses fixed noise addition, our method leverages RDP with dynamic noise scaling, achieving stronger privacy guarantees while maintaining model accuracy. Third, our method enables fine-grained noise adaptation by modulating the noise level based on gradient variance, effectively avoiding excessive perturbation (which harms accuracy) and insufficient noise (which weakens privacy). The decay factor β prevents abrupt fluctuations in noise levels, stabilizing training [14], while the learning rate-dependent parameter γ ensures optimal noise scaling without excessive perturbation [16]. This approach mitigates the risk of over-smoothing gradients, which can slow convergence, while also preventing privacy leakage from underestimated sensitivity variations. Overall, our proposed noise scaling mechanism provides a more robust and efficient privacy-preserving solution for FL, particularly in settings with heterogeneous client data and varying gradient sensitivities.

3) Privacy-Preserving Model Aggregation at the Server: After receiving the noisy values, the server aggregates these differentially private updates using moment accounted noise scaling to compute the new global model weights. These updated weights are then shared back with the clients, enabling them to update their local models while preserving privacy. The global model is computed as follows:

$$w^{t+1} = \frac{1}{N} \sum_{i=1}^{N} \tilde{w}_i^t \tag{7}$$

where:

- N is the number of participating clients.
- $\tilde{w}_i^t$ represents the noisy shuffled model update from client i.

IV. EXPERIMENTATION AND EVALUATION

Our RDP-FL framework was evaluated on the Medical-MNIST [17] and CIFAR-10 [18] datasets, covering both medical imaging and general object classification tasks, within a privacy-preserving Federated Learning (FL) environment. We employed an adapted VGG-like model optimized for CIFAR-10 object classification and Medical-MNIST medical imaging classification, ensuring a balance between computational efficiency and model accuracy.

As shown in Table I, our experimental setup consisted of 100 federated clients, each training on a partitioned subset of their respective datasets. For our primary evaluation, we adopted a privacy budget of $\epsilon = 5$ to balance privacy and model accuracy. Additionally, we further investigated the impact of varying the privacy budget by conducting complementary experiments with multiple values of ϵ, as detailed in Section IV-A. The training process involved 10 communication rounds for Medical-MNIST and 20 for CIFAR-10, during which clients locally updated their models and transmitted their updates to the central server for secure aggregation. Differential privacy was applied at the server side using moment-accounted noise scaling under the RDP framework. This setup allowed us to evaluate the effectiveness of RDP-FL, particularly in terms of model accuracy and the test privacy score. This experimental design facilitated a comprehensive assessment of our federated learning approach across datasets with varying complexities, providing insights into the adaptability and robustness of our privacy-preserving method.

TABLE I: Key parameters used for the Medical-MNIST and CIFAR-10 datasets

Dataset	**Clients (N)**	**Comm. Rounds (R)**	**Privacy Budget (ϵ)**
Medical-MNIST	100	10	5
CIFAR-10	100	20	5

A. Impact of Privacy Buget Parameter

In our set of experiments, we varied the privacy budget (ϵ) to observe its impact on model accuracy, following the evaluation protocol adopted in recent work on LDP for FL [4]. As shown in Figure 2, higher values of ϵ led to improved performance across datasets. For Medical MNIST, RDP-FL achieved consistently high accuracy, reaching 94% at ϵ = 3 and 96% at ϵ = 5, with only slight improvements as ϵ increased. In contrast, CIFAR-10 showed a more significant upward trend, reaching approximately 65% accuracy at ϵ = 3 and 75% at ϵ = 5, indicating its higher sensitivity to noise. This trend highlights the trade-off between privacy and model performance, where larger privacy budgets (weaker privacy guarantees) lead to better accuracy, particularly benefiting more complex datasets like CIFAR-10, which require greater data fidelity for optimal performance.

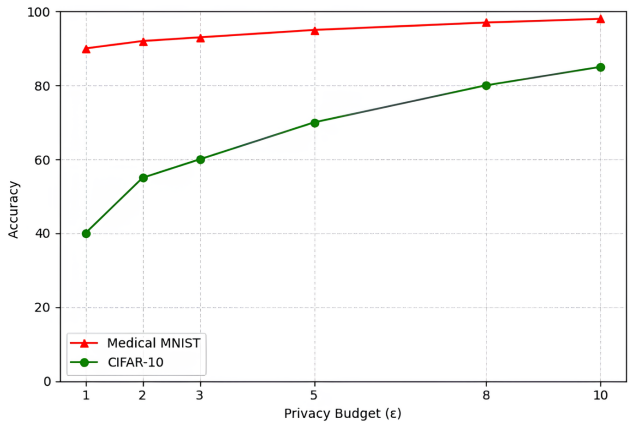

Fig. 2: Effect of privacy parameter ϵ

B. Comparative Evaluation of RDP-FL Model

We selected the Staircase Randomized Response Federated Learning (SRR-FL), which is based on an LDP mechanism, as the baseline model due to its recent effectiveness in mitigating privacy leakage in FL scenarios [4]. SRR-FL introduces discrete randomized noise to model updates, enhancing privacy protection without excessively compromising performance, making it a suitable baseline for comparison. We compare accuracy and privacy preservation between the SRR-FL and RDP-FL mechanisms on two datasets: Medical-MNIST and CIFAR-10, as detailed in Figures 3 and 4.

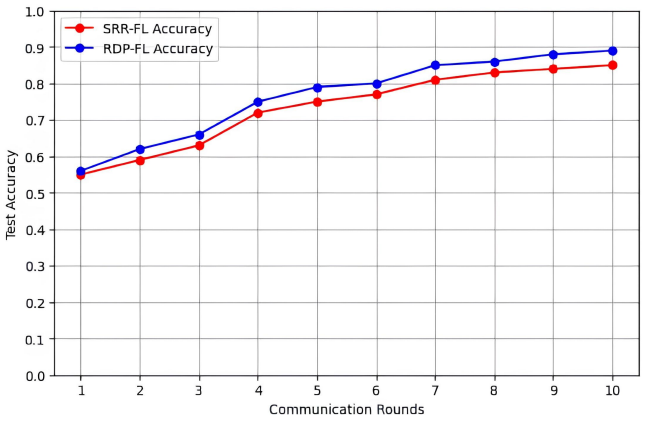

(a) Test Accuracy for the Medical-MNIST

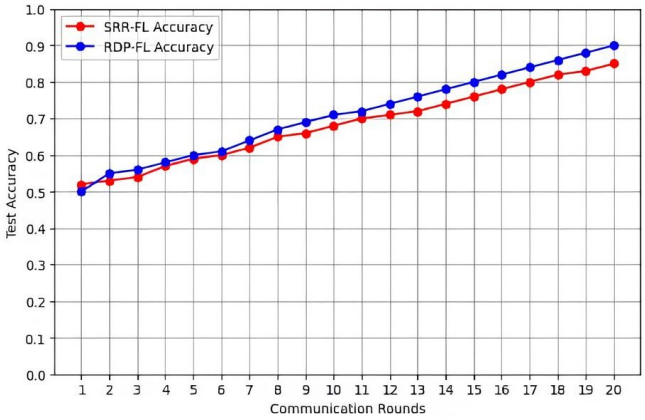

(b) Test Accuracy for the CIFAR-10

Fig. 3: Comparison of Accuracy between SRR-FL and RDP-FL.

Figure 3 illustrates the test accuracy progression over multiple communication rounds for both the Medical-MNIST and CIFAR-10 datasets, under a privacy budget of $\epsilon = 5$. For the Medical-MNIST dataset (Figure 3a), both SRR-FL and RDP-FL start with similar accuracy levels at around 60% in the first communication round. As training progresses, RDP-FL consistently outperforms SRR-FL, demonstrating superior learning efficiency. After the fifth communication round, RDP-FL exhibits a noticeable improvement, gaining a higher accuracy margin over SRR-FL. By round 10, RDP-FL reaches an accuracy of approximately 90%, whereas SRR-FL remains slightly lower at 88%. The accuracy gap between the two models increases in the later rounds, indicating that RDP-FL converges faster and generalizes better. The results confirm that RDP-FL achieves superior performance compared to SRR-FL while maintaining privacy guarantees, making it a more effective approach for federated learning in medical imaging applications. For the CIFAR-10 dataset (Figure 3b), both methods exhibit a steady increase in accuracy throughout the communication rounds. Initially, their performance is similar, but as training progresses, RDP-FL starts to demonstrate a marginal improvement. The highest accuracy is recorded at round 20, with RDP-FL achieving about 89%, compared to 87% for SRR-FL. The results clearly indicate that RDP-FL consistently outperforms SRR-FL in both datasets. RDP-FL achieves higher accuracy at every stage, particularly in the later communication rounds, demonstrating its effectiveness in federated learning with differential privacy. Figure 4 highlights privacy preservation using the test privacy score, which corresponds to the cumulative privacy loss ϵ_T computed as shown in equation 5. A lower score indicates stronger privacy. For the Medical-MNIST dataset (Figure 4a), both RDP-FL and SRR-FL start with high privacy scores, initially fluctuating in the first

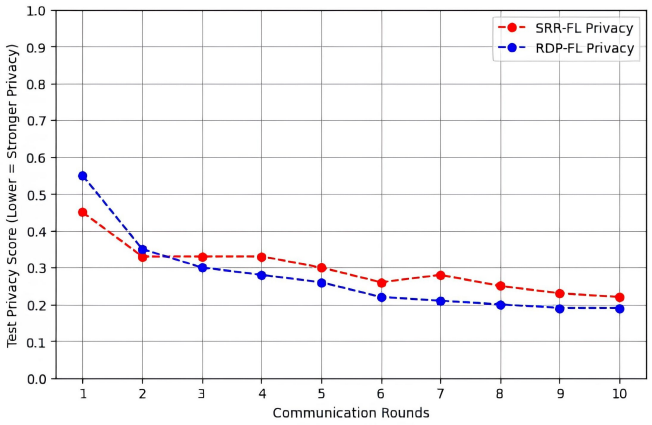

(a) Test Privacy for the Medical-MNIST

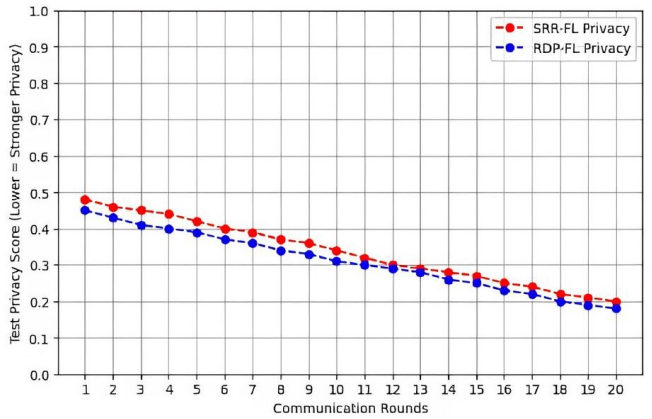

(b) Test Privacy for the CIFAR-10

Fig. 4: Comparison of Privacy between SRR-FL and RDP-FL.

few communication rounds. However, as training progresses, RDP-FL consistently achieves lower privacy scores than SRR-FL, indicating stronger privacy protection. By round 3, both models stabilize around 0.3, but from this point forward, RDP-FL maintains a clear advantage by further reducing the privacy score, outperforming SRR-FL in privacy preservation. At round 10, RDP-FL achieves a privacy score of approximately 0.2, whereas SRR-FL remains slightly higher, confirming that RDP-FL provides better privacy protection while also achieving superior accuracy. This demonstrates RDP-FL's ability to offer a better trade-off between privacy and utility, making it a more effective choice for federated learning in privacy-sensitive applications such as medical imaging. A similar trend is observed for CIFAR-10 (Figure 4b), where both methods start at 0.72, but RDP-FL maintains an advantage throughout training, reaching 0.50 at round 10 (vs. 0.52 for SRR-FL) and stabilizing at 0.42 by round 20, outperforming SRR-FL's 0.44, while also achieving superior accuracy (89% vs. 87%). These results highlight RDP-FL as a promising approach for federated learning where both strong privacy guarantees and high model utility are required. These detailed observations from Figures 3 and 4 highlight RDP-FL's capability to provide superior accuracy with efficient convergence, alongside robust privacy preservation, positioning it as a more practical solution for sensitive applications, particularly medical image classification. The adaptability of RDP-FL in dynamically adjusting noise levels enables an effective balance between accuracy and privacy. Additionally, because this adaptive approach avoids excessive noise, it helps maintain efficient model convergence, which is particularly beneficial for deployment in bandwidth and computationally constrained environments. Thus, RDP-FL adjusts noise based on gradient sensitivity, effectively balancing privacy and accuracy. By precisely tracking cumulative privacy loss through RDP, it overcomes common issues of fixed-noise approaches. Future research could explore its scalability on larger datasets and assess its robustness against advanced inference attacks in realistic settings.

V. CONCLUSIONS

This paper introduced a privacy-preserving Federated Learning (FL) approach that combines moment-accounted noise scaling with Rényi Differential Privacy (RDP) to enhance privacy while maintaining high model accuracy. The proposed method adjusts the privacy budget based on the sensitivity of model updates, reducing unnecessary noise injection while ensuring strong privacy guarantees against inference attacks. This adaptive mechanism effectively balances privacy protection and model utility. Our experimental evaluation on the Medical-MNIST and CIFAR-10 datasets demonstrates that RDP-FL achieves strong privacy protection while maintaining high predictive accuracy. The results indicate that RDP-FL provides a good balance between privacy and utility, making it a promising approach for differentially private federated learning. Future research could focus on improving the scalability of RDP-FL for large-scale federated learning deployments and evaluating its robustness against adversarial attacks in real-world applications. Additionally, further enhancements could integrate advanced optimization techniques to refine the privacy-utility trade-off, making the model even more efficient. Moreover, investigating its applicability to privacy-sensitive tasks such as medical diagnostics and financial fraud detection remains an important avenue for exploration. Future work will also explore the applicability of our approach to more sensitive medical data such as genomic datasets, and investigate its scalability to federated networks with thousands of clients.

REFERENCES

[1] OKDEM, Selcuk et OKDEM, Sema. Artificial Intelligence in Cybersecurity: A Review and a Case Study. Applied Sciences, 2024, vol. 14, no 22, p. 10487.

[2] FENG, Yunhao, GUO, Yanming, HOU, Yinjian, et al. A survey of security threats in federated learning. Complex & Intelligent Systems, 2025, vol. 11, no 2, p. 1-26.

[3] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated machine learning: Concept and applications. TIST 1, 2 (2019), 1–19.

[4] VARUN, Matta, FENG, Shuya, WANG, Han, et al. Towards Accurate and Stronger Local Differential Privacy for Federated Learning with Staircase Randomized Response. In : Proceedings of the Fourteenth ACM Conference on Data and Application Security and Privacy. 2024. p. 307-318.

[5] DWORK, Cynthia, ROTH, Aaron, et al. The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science, 2014, vol. 9, no 3–4, p. 211-407.

[6] Talaei, Mahtab, and Iman Izadi. "Adaptive Differential Privacy in Federated Learning: A Priority-Based Approach." arXiv preprint arXiv:2401.02453 (2024).

[7] ZHU, Gongxi, LI, Donghao, GU, Hanlin, et al. FedMIA: An Effective Membership Inference Attack Exploiting" All for One" Principle in Federated Learning. arXiv preprint arXiv:2402.06289, 2024.

[8] GUO, Pengxin, et al. A New Federated Learning Framework Against Gradient Inversion Attacks. arXiv preprint arXiv:2412.07187, 2024.

[9] TAYYEH, Huda Kadhim; AL-JUMAILI, Ahmed Sabah Ahmed. Balancing Privacy and Performance: A Differential Privacy Approach in Federated Learning. Computers, 2024, 13.11: 277.

[10] CHUANXIN, Zhou, YI, Sun, et DEGANG, Wang. Federated learning with Gaussian differential privacy. In : Proceedings of the 2020 2nd international conference on robotics, intelligent control and artificial intelligence. 2020. p. 296-301.

[11] Truex, S., Liu, L., Chow, K. H., Gursoy, M. E., & Wei, W. (2020, April). LDP-Fed: Federated learning with local differential privacy. In Proceedings of the third ACM international workshop on edge systems, analytics and networking (pp. 61-66).

[12] Chamikara, M. A. P., Liu, D., Camtepe, S., Nepal, S., Grobler, M., Bertok, P., & Khalil, I. (2022). Local differential privacy for federated learning. arXiv preprint arXiv:2202.06053.

[13] FENG, Shuya, MOHAMMADY, Meisam, HONG, Hanbin, et al. Universally harmonizing differential privacy mechanisms for federated learning: Boosting accuracy and convergence. arXiv preprint arXiv:2407.14710, 2024.

[14] MIRONOV, Ilya. Rényi differential privacy. In: 2017 IEEE 30th computer security foundations symposium (CSF). IEEE, 2017. p. 263-275.

[15] LÉCUYER, Mathias. Practical Privacy Filters and Odometers with Rényi Differential Privacy and Applications to Differentially Private Deep Learning. arXiv preprint arXiv:2103.01379, 2021.

[16] Heikkilä, M. A. (2024). On Joint Noise Scaling in Differentially Private Federated Learning with Multiple Local Steps. arXiv preprint arXiv:2407.19286.

[17] Jiancheng Yang, Rui Shi, and Bingbing Ni. 2021. MedMNIST Classification Decathlon: A Lightweight AutoML Benchmark for Medical Image Analysis. In IEEE ISBI. 191–195.

[18] KRIZHEVSKY, Alex, HINTON, Geoffrey, et al. Learning multiple layers of features from tiny images. 2009.