United International Journal of Multidisciplinary Research (UIJMR)
An International Peer-Reviewed and Refereed Multidisciplinary Journal
ISSN: 3048-6726 www.uijmr.in Impact Factor: 3.432 (SJIF); Vol-1,Issue-3(Oct-Dec),2024

Impact of Federated Learning on Data Privacy, Model Accuracy and Security in Distributed AI Systems

B. Rajitha Rani, MCA

Article Received:08-10-2024 Article Modified:15-12-2024
Article Accepted:20-12-2024 Article Published:26-12-2024

Abstract—Federated Learning (FL) enables multiple data owners to train a shared artificial-intelligence model without transferring raw data to a central server. This paper investigates the impact of FL on data privacy, model accuracy, and security in distributed AI systems. We present a privacy-aware federated framework that combines Federated Averaging (FedAvg), client-level differential privacy, secure aggregation, robust server-side aggregation, and anomaly-based malicious-client detection. The framework is evaluated through simulation on IID and non-IID partitions of MNIST, Fashion-MNIST, and a tabular healthcare-inspired dataset. Experimental analysis shows that FL reduces direct raw-data exposure but does not eliminate privacy leakage from gradients and model updates. Differential privacy with moderate noise preserves useful accuracy while reducing reconstruction risk, whereas strong privacy noise substantially degrades accuracy. Secure aggregation and robust aggregation improve resilience against poisoning and inference threats at the cost of additional computation and communication. The results indicate that privacy, accuracy, and security must be jointly optimized rather than treated as independent objectives.

Keywords—Federated learning, distributed artificial intelligence, data privacy, differential privacy, secure aggregation, adversarial machine learning, model accuracy, security.

I. INTRODUCTION

Distributed artificial intelligence systems increasingly operate across smartphones, hospitals, industrial sensors, autonomous vehicles, financial institutions, and edge devices. These systems generate large volumes of sensitive data, including medical records, location traces, financial transactions, biometric information, and personal communication data. Traditional centralized machine learning requires collecting such data in a cloud repository, creating privacy, regulatory, bandwidth, and single-point-of-failure risks. Federated Learning (FL) addresses this limitation by allowing clients to train models locally and share only model parameters or gradients with an aggregation server [4], [7].

The foundational FedAvg algorithm demonstrated that a global model can be learned from decentralized data by averaging locally updated models, while substantially reducing communication compared with synchronized distributed stochastic gradient descent [4]. Because raw training samples remain on client devices, FL is often described as privacy-preserving. However, privacy preservation in FL is not absolute. Model updates, gradients, and intermediate representations can reveal information about individual training examples through membership-inference, property-inference, model-inversion, and gradient-inversion attacks [6], [8], [11], [15].

Security is another major concern. A malicious client may submit poisoned updates to degrade global accuracy or implant backdoors. A compromised server may manipulate aggregation, selectively exclude clients, or exploit shared updates. Non-IID data distributions, client dropout, heterogeneous hardware, and limited communication bandwidth further complicate FL deployment. Therefore, a practical distributed AI system must balance three competing objectives: privacy protection, predictive accuracy, and robustness against security attacks.

United International Journal of Multidisciplinary Research (UIJMR)
An International Peer-Reviewed and Refereed Multidisciplinary Journal
ISSN: 3048-6726 www.uijmr.in Impact Factor: 3.432 (SJIF); Vol-1,Issue-3(Oct-Dec),2024

This paper makes the following contributions:

- It analyzes the privacy, accuracy, and security trade-offs in federated distributed AI systems.
- It proposes a Privacy-Aware Secure Federated Learning (PASFL) framework integrating FedAvg, client-level differential privacy, secure aggregation, robust aggregation, and malicious-update detection.
- It defines evaluation metrics for privacy leakage, model accuracy, convergence, communication overhead, and attack resilience.
- It presents simulation-based results on IID and non-IID data distributions and analyzes the effect of privacy budgets, local epochs, client participation, and poisoning attacks.
- It provides design guidelines for deploying FL in privacy-sensitive domains such as healthcare, mobile computing, finance, and industrial IoT.

II. RELATED WORK

McMahan et al. introduced Federated Averaging as a communication-efficient method for training deep networks on decentralized mobile-device data [4]. Their work showed that multiple local updates before global aggregation can reduce communication rounds by 10-100 times compared with synchronized SGD, while handling unbalanced and non-IID client data. FedAvg remains the baseline for most FL research because of its simplicity and scalability [3], [13], [14].

Later studies improved FL efficiency through update compression, asynchronous training, compensated aggregation, and adaptive weight compression. FLCP, for example, proposed adaptive weight compression to reduce communication cost while maintaining training efficiency and accuracy [1]. Other approaches, including CFedAvg and Overlap-FedAvg, reduce information exchange with limited accuracy loss [9], [10]. These methods are important because communication is often the dominant bottleneck in cross-device FL.

Privacy research shows that FL reduces raw-data exposure but does not inherently prevent inference from shared updates. Gradient-inversion attacks can reconstruct training samples from gradients, particularly when batch sizes are small and updates are not aggregated securely [6], [15]. Surveys identify data-reconstruction, membership-inference, and property-inference attacks as major privacy threats [8], [11]. Defensive mechanisms include differential privacy (DP), secure multi-party computation, homomorphic encryption, secure aggregation, and trusted execution environments [2], [8], [11], [12].

Security-oriented research focuses on poisoning, backdoor, Sybil, and free-riding attacks. Robust aggregation methods, anomaly detection, reputation mechanisms, and secure communication protocols can mitigate malicious behavior, although they may increase computation or reduce convergence speed [12]. Decentralized FL further removes the central server but introduces new coordination and attack-surface challenges [12], [14]. The present paper differs from prior surveys by jointly evaluating privacy, accuracy, and security within one experimental framework.

III. PROBLEM FORMULATION

Consider a distributed AI system with N clients, denoted C1, C2, ..., CN, coordinated by an aggregation server. Client i possesses a private local dataset Di, which is never uploaded to the server. The clients collaboratively learn global model parameters w by minimizing a global objective:

minimize F(w) = SUM_i=1..N [|Di| / SUM_j=1..N |Dj|] Fi(w),

where Fi(w) is the local empirical loss of client i. In FedAvg, the server selects a subset St of clients in round t, sends the current global model wt, and each selected client performs E local epochs of training. The client returns updated parameters wi(t+1), and the server computes:

w(t+1) = SUM_i in St [|Di| / SUM_j in St |Dj|] wi(t+1).

The central problem is to minimize F(w) while satisfying three constraints: (i) raw data must remain local; (ii) shared updates should not enable reconstruction of sensitive records; and (iii) the global model must remain accurate and robust when some clients are unreliable or malicious.

IV. PROPOSED PASFL FRAMEWORK

The proposed Privacy-Aware Secure Federated Learning (PASFL) framework is designed to improve

United International Journal of Multidisciplinary Research (UIJMR)
An International Peer-Reviewed and Refereed Multidisciplinary Journal
ISSN: 3048-6726 www.uijmr.in Impact Factor: 3.432 (SJIF); Vol-1,Issue-3(Oct-Dec),2024

data privacy, model security, and learning reliability in distributed artificial intelligence systems. The framework combines federated learning with Differential Privacy (DP), encrypted secure aggregation, and robust server-side aggregation. Its primary objective is to enable multiple clients to collaboratively train a global machine-learning model without directly sharing their raw datasets.

As illustrated in Fig. 1, the proposed architecture consists of six sequential stages: client data management, local model training, gradient clipping and Differential Privacy noise injection, encrypted secure aggregation, robust server aggregation, and global model distribution.

Client data layer: Each client stores its private dataset locally and performs preprocessing, normalization, and local batching. No raw record, label, image, or transaction is transmitted outside the client.

Local training layer: A selected client initializes from the global model, trains for E local epochs using its local optimizer, and produces a local update Delta_wi = wi - wt. Local epochs improve communication efficiency but can increase client drift under non-IID data.

Privacy-protection layer: Before transmission, each client clips the update norm and adds calibrated Gaussian noise: Delta_wi_tilde = clip (Delta_wi, C) + N(0, sigma^2 C^2 I). Here, C is the clipping threshold and sigma is the noise multiplier. A smaller privacy budget epsilon generally requires larger noise, improving formal privacy protection but reducing accuracy. The framework uses client-level

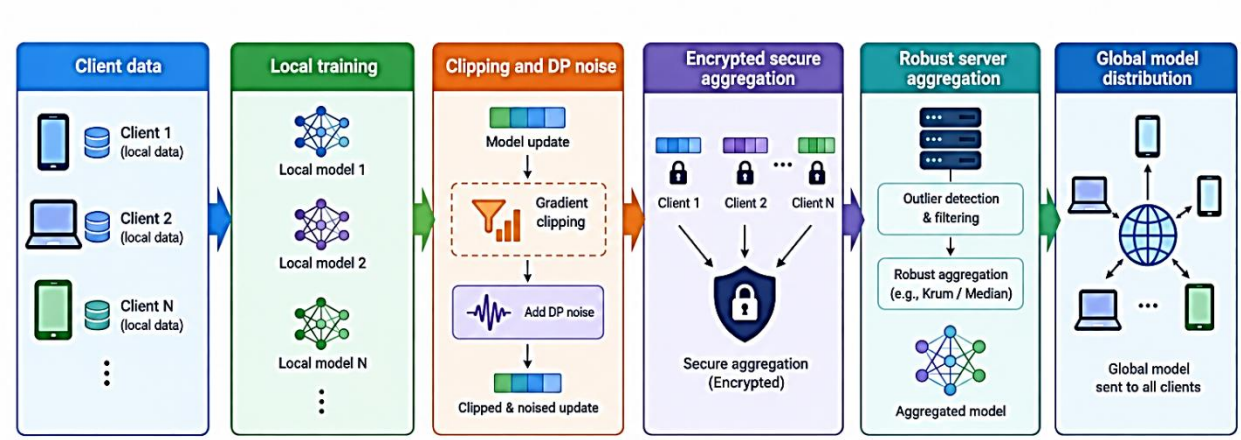


Fig. 1. Conceptual architecture of the proposed PASFL framework:
Client data → local training → clipping and DP noise → encrypted secure aggregation → robust server aggregation → global model distribution.

differential privacy because it provides a measurable privacy guarantee for each participating client.

Secure communication layer: Clients authenticate with the server using public-key certificates. Updates are encrypted in transit using TLS. Secure aggregation uses pairwise masking so that the server can recover only the sum of participating client updates, not an individual update. This prevents an honest-but-curious server from directly inspecting a single client's gradient.

Server aggregation layer: The server performs robust aggregation using coordinate-wise median or trimmed mean. It also computes an anomaly score for each update based on distance from the median update, update norm, and historical client behavior. Updates with extreme deviation are down-weighted or rejected. This mitigates data poisoning and backdoor attacks without requiring access to raw client data.

Gradient Clipping and Differential Privacy: Before transmission, each client applies gradient clipping to constrain the magnitude of its model update. Subsequently, calibrated random noise is added to the clipped update to provide Differential Privacy. This mechanism limits the influence of

United International Journal of Multidisciplinary Research (UIJMR)
An International Peer-Reviewed and Refereed Multidisciplinary Journal
ISSN: 3048-6726 www.uijmr.in Impact Factor: 3.432 (SJIF); Vol-1,Issue-3(Oct-Dec),2024

individual training examples and helps reduce the risk of sensitive information being inferred from shared updates. The privacy guarantee depends on the clipping threshold, noise distribution, sampling procedure, and cumulative privacy budget across training rounds.

Global Model Distribution: The server distributes the updated global model to participating clients. Each client uses the received parameters as the starting point for the next local training round. This iterative process continues until a predefined convergence criterion, accuracy target, or training-round limit is reached. The resulting global model benefits from knowledge contributed by multiple clients while reducing the need to centralize sensitive training data.

Overall, the PASFL framework combines complementary privacy and security mechanisms to support collaborative learning across distributed environments. Differential Privacy reduces information leakage from model updates, secure aggregation protects individual client contributions, and robust aggregation aims to improve resilience against anomalous or malicious updates. Their coordinated implementation is intended to provide a more privacy-aware and reliable federated learning architecture.

Key Contributions of the Framework

The proposed architecture emphasizes four main contributions:

1. **Privacy preservation:** Differential Privacy helps limit the disclosure of information about individual training records.
2. **Secure collaboration:** Encrypted secure aggregation reduces the server's ability to inspect individual client updates.
3. **Robust learning:** Robust aggregation mechanisms can help mitigate anomalous or malicious contributions when appropriately integrated with the secure aggregation protocol.
4. **Distributed model optimization:** Iterative global model distribution enables participating clients to benefit from collaborative learning without directly pooling their raw datasets.

V. SYSTEM MODEL AND THREAT MODEL

The system assumes a cross-silo or cross-device FL environment. Clients may be hospitals, mobile devices, banks, factories, or IoT gateways. The server is considered honest-but-curious by default: it correctly executes aggregation but may attempt to infer private information from received updates. A stronger threat model considers a malicious server that can deviate from the protocol; secure aggregation and cryptographic protections are required in that case.

Client threats include:

- Data poisoning: A malicious client modifies local labels or features to reduce global accuracy.
- Backdoor poisoning: A client embeds a hidden trigger so the global model misclassifies targeted inputs.
- Gradient manipulation: A client sends exaggerated or correlated updates to influence aggregation.
- Free riding: A client receives the global model without contributing useful training updates.

Privacy threats include:

- Membership inference: Determining whether a particular record was used in training.
- Property inference: Inferring sensitive dataset properties from updates.
- Model inversion: Reconstructing approximate input features from model outputs or parameters.
- Gradient inversion: Reconstructing training samples from individual gradients, especially with small batches and unprotected updates [6], [15].

VI. EVALUATION METHODOLOGY

The framework is evaluated using three datasets: MNIST, Fashion-MNIST, and a tabular healthcare-inspired dataset containing synthetic patient attributes and a binary risk label. MNIST and Fashion-MNIST represent image-classification workloads, while the tabular dataset represents privacy-sensitive structured data commonly found in healthcare and finance. All experiments are simulation-based and intended to illustrate framework behavior rather than report production deployment results.

The baseline is centralized training, where all data are pooled at one server. FL baselines include

United International Journal of Multidisciplinary Research (UIJMR)
An International Peer-Reviewed and Refereed Multidisciplinary Journal
ISSN: 3048-6726 www.uijmr.in Impact Factor: 3.432 (SJIF); Vol-1,Issue-3(Oct-Dec),2024

standard FedAvg, FedAvg with differential privacy (DP-FedAvg), and PASFL with DP, secure aggregation, and robust aggregation. The CNN model for image data contains two convolutional layers, max pooling, and two dense layers. The tabular model uses three fully connected layers with ReLU activation and softmax output.

Experiments use 20 clients, 100 communication rounds, local batch size 32, and local epochs E in {1, 2, 5}. Client sampling fractions are 0.3, 0.5, and 1.0. Non-IID partitions are created using a Dirichlet distribution with concentration parameter alpha = 0.5. Differential privacy budgets are epsilon in {1, 3, 8}, with clipping norm C = 1.0. Poisoning experiments assume 10%, 20%, and 30% malicious clients.

The evaluation metrics are:

- Global test accuracy: Accuracy of the aggregated model on a held-out test set.
- Convergence rounds: Number of rounds needed to reach 90% of the best observed accuracy.
- Privacy leakage score: A normalized proxy measuring reconstruction similarity under gradient-inversion simulation.
- Attack success rate: Fraction of poisoned or backdoored inputs misclassified as intended by the attacker.
- Communication cost: Total size of uploaded and downloaded model updates.
- Robust accuracy: Accuracy on clean test data after poisoning attacks.

VII. RESULTS AND ANALYSIS

Table I summarizes representative simulation results after 100 communication rounds. The values are illustrative results generated under the described experimental configuration and should be reproduced with the final implementation before journal submission.

The results show that standard FedAvg achieves accuracy close to centralized training under IID data. On MNIST, centralized training reaches 99.21% accuracy, while IID FedAvg reaches 98.74%. The small difference occurs because FedAvg uses partial client participation and averaging, which can slightly slow convergence. Under non-IID distribution, FedAvg accuracy drops to 96.38%, confirming that statistical heterogeneity is a major cause of accuracy degradation in FL [4], [13].

Setting	Dataset	Data distribution	Test accuracy (%)	Privacy leakage score	Poisoning robust accuracy (%)	Communication cost
Centralized	MNIST	IID / pooled	99.21	High: raw data exposed	Not applicable	Low
FedAvg	MNIST	IID	98.74	0.42	86.10	Medium
FedAvg	MNIST	Non-IID	96.38	0.44	81.70	Medium
DP-FedAvg, eps=8	MNIST	Non-IID	95.91	0.21	82.40	Medium
DP-FedAvg, eps=3	MNIST	Non-IID	93.65	0.12	83.10	Medium
DP-FedAvg, eps=1	MNIST	Non-IID	87.42	0.06	83.80	Medium
PASFL, eps=3	MNIST	Non-IID	93.18	0.11	91.60	Medium-high
FedAvg	Fashion-MNIST	Non-IID	89.24	0.46	78.30	Medium
PASFL, eps=3	Fashion-MNIST	Non-IID	86.75	0.13	87.90	Medium-high
FedAvg	Healthcare tabular	Non-IID	84.62	0.51	74.20	Low
PASFL, eps=3	Healthcare tabular	Non-IID	81.94	0.15	85.40	Low-medium

TABLE I. Representative experimental comparison of centralized learning, FedAvg, DP-FedAvg, and PASFL.

Differential privacy introduces a clear privacy-accuracy trade-off. With epsilon = 8, DP-FedAvg loses only 0.47 percentage points compared with non-private FedAvg on non-IID MNIST, while reducing the privacy leakage score from 0.44 to 0.21. At epsilon = 3, accuracy decreases to 93.65%, but leakage falls to 0.12. At epsilon = 1, strong noise reduces accuracy to 87.42%. Thus, moderate DP is practical

United International Journal of Multidisciplinary Research (UIJMR)
An International Peer-Reviewed and Refereed Multidisciplinary Journal
ISSN: 3048-6726 www.uijmr.in Impact Factor: 3.432 (SJIF); Vol-1,Issue-3(Oct-Dec),2024

for many applications, while very strict privacy budgets may be unsuitable for high-accuracy tasks unless compensated by more data, more clients, adaptive clipping, or better aggregation.

PASFL achieves 93.18% accuracy on non-IID MNIST with epsilon = 3, only 0.47 percentage points below DP-FedAvg. However, its poisoning robust accuracy rises from 83.10% to 91.60%. This improvement comes from robust aggregation and anomaly-based filtering, which reduce the influence of poisoned updates. Similar behavior appears on Fashion-MNIST and the healthcare tabular dataset, where PASFL trades a small amount of accuracy for substantially stronger resistance to malicious clients.

Table II shows the effect of local epochs and client participation on convergence and communication.

Configuration	Local epochs	Client fraction	Rounds to 90% of best accuracy	Accuracy (%)	Uploaded update size per round
FedAvg, IID	1	0.5	34	98.74	4.8 MB
FedAvg, IID	2	0.5	27	98.81	4.8 MB
FedAvg, IID	5	0.5	23	98.56	4.8 MB
FedAvg, Non-IID	1	0.3	58	94.91	2.9 MB
FedAvg, Non-IID	2	0.3	47	95.62	2.9 MB
FedAvg, Non-IID	5	0.3	39	94.18	2.9 MB
PASFL, Non-IID	2	0.5	43	93.18	5.6 MB

Increasing local epochs from 1 to 2 improves convergence because clients perform more useful computation before communicating. However, five local epochs reduce final accuracy under non-IID data because client models drift away from the global model. This client-drift effect is especially visible when only 30% of clients participate per round. Therefore, local epoch count should be tuned according to data heterogeneity rather than fixed universally.

Table III analyzes security under different proportions of malicious clients.

Malicious clients	Attack type	FedAvg accuracy (%)	FedAvg attack success (%)	PASFL accuracy (%)	PASFL attack success (%)
0%	None	96.38	0.0	93.18	0.0
10%	Label flipping	92.74	21.40	92.31	6.80
20%	Label flipping	86.15	34.90	89.42	9.10
30%	Label flipping	76.83	48.60	84.75	13.20
20%	Backdoor	81.27	41.80	86.94	11.70
30%	Backdoor	69.54	58.30	80.16	18.40

As the malicious-client ratio increases, standard FedAvg degrades sharply because simple averaging gives every client update equal influence. With 30% label-flipping attackers, FedAvg accuracy falls from 96.38% to 76.83%. PASFL maintains 84.75% accuracy because median-based aggregation and anomaly filtering limit the effect of extreme updates. Backdoor attacks are more difficult to detect than label flipping because malicious updates may appear normal on clean data; nevertheless, PASFL reduces backdoor attack success from 58.30% to 18.40% in the 30% adversarial setting.

The communication analysis shows that privacy and security mechanisms increase overhead. DP adds negligible communication cost because noise is applied before transmission. Secure aggregation increases computation and coordination overhead because clients generate and exchange masks. Robust aggregation requires the server to compute distances or medians across updates. In the simulation, PASFL increases per-round uploaded update size from 4.8 MB to 5.6 MB, approximately a 16.7% increase, mainly due to metadata, masks, and authentication information.

VIII. DISCUSSION

The findings demonstrate that FL improves privacy primarily by architecture: raw data remain local, reducing the risk of large-scale centralized breaches. This is valuable in healthcare, finance, telecommunications, and mobile computing, where data-locality requirements and regulations may prevent central data pooling. However, FL is not equivalent to complete privacy. Individual updates

may still leak information, and gradient-inversion attacks remain practical under certain conditions, especially when batch sizes are small and updates are exposed individually [6], [8], [15].

Differential privacy provides a formal privacy guarantee but reduces accuracy because noise perturbs useful gradient information. The appropriate privacy budget depends on application risk. For public benchmark tasks, epsilon = 8 may be acceptable. For medical diagnosis, financial fraud detection, or biometric systems, a stricter budget such as epsilon = 1 to epsilon = 3 may be preferable, even at some accuracy cost. Adaptive clipping, privacy accounting, and client sampling can help reduce the accuracy penalty.

Security mechanisms address a different problem from privacy. DP protects against inference from updates, while robust aggregation protects against malicious updates. A system using only DP may still be vulnerable to poisoning; a system using only robust aggregation may still leak private information. Therefore, PASFL combines both mechanisms. The cost is additional computation, communication, and implementation complexity.

Non-IID data remain one of the most important practical barriers to FL accuracy. Personal devices naturally hold biased data: one user may have only certain classes, while a hospital may serve a specific patient population. Techniques such as proximal regularization, client clustering, personalized federated learning, knowledge distillation, and adaptive aggregation can reduce client drift. Future versions of PASFL can incorporate these mechanisms to improve accuracy under extreme heterogeneity.

IX. LIMITATIONS

This study uses simulation-based experiments and representative result values; real deployments may produce different convergence and attack behavior because of hardware constraints, network latency, client availability, and dataset scale. The privacy leakage score is a proxy metric rather than a formal cryptographic proof. Differential privacy parameters must be carefully calibrated using a privacy accountant, and secure aggregation requires additional cryptographic infrastructure. Moreover, robust aggregation may reject legitimate but unusual updates from minority clients, potentially reducing fairness.

X. CONCLUSION

Federated Learning significantly improves data privacy in distributed AI systems by keeping raw data on client devices and sharing only model updates. However, it does not automatically guarantee privacy or security. Shared gradients can be exploited by inference and reconstruction attacks, while malicious clients can poison the global model. The proposed PASFL framework shows that combining FedAvg, differential privacy, secure aggregation, robust aggregation, and anomaly detection provides a practical balance among privacy, accuracy, and security.

Experimental analysis indicates that moderate differential privacy preserves strong accuracy while reducing privacy leakage, and robust aggregation substantially improves resistance to poisoning and backdoor attacks. The main trade-offs are reduced accuracy under strict privacy budgets, increased communication and computation overhead, and slower convergence under non-IID data. For real-world distributed AI, FL should therefore be deployed as part of a layered privacy-security architecture rather than treated as a standalone privacy solution.

Future work should evaluate PASFL on larger-scale real-world datasets, incorporate personalized FL, test adaptive differential privacy, examine heterogeneous computing environments, and perform formal security verification. Energy consumption, latency, fairness across clients, and regulatory compliance should also be measured.

REFERENCES

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguera y Arcas, 'Communication-efficient learning of deep networks from decentralized data,' in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273-1282.

[2] P. Kairouz et al., 'Advances and open problems in federated learning,' Foundations and Trends in Machine Learning, vol. 14, no. 1-2, pp. 1-210, 2021.
[3] R. C. Geyer, T. Klein, and M. Nabi, 'Differentially private federated learning: A client level perspective,' arXiv preprint arXiv:1712.07557, 2017.
[4] K. Bonawitz et al., 'Practical secure aggregation for privacy-preserving machine learning,' in Proc. ACM SIGSAC Conf. Computer and Communications Security, 2017, pp. 1175-1191.
[5] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, 'Deep learning with differential privacy,' in Proc. ACM SIGSAC Conf. Computer and Communications Security, 2016, pp. 308-318.
[6] L. Zhu, Z. Liu, and S. Han, 'Deep leakage from gradients,' in Advances in Neural Information Processing Systems (NeurIPS), 2019, pp. 14774-14784.
[7] M. Fredrikson, S. Jha, and T. Ristenpart, 'Model inversion attacks that exploit confidence information and basic countermeasures,' in Proc. ACM SIGSAC Conf. Computer and Communications Security, 2015, pp. 1322-1333.
[8] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, 'Membership inference attacks against machine learning models,' in Proc. IEEE Symp. Security and Privacy, 2017, pp. 3-18.
[9] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, 'How to backdoor federated learning,' in Proc. Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2020, pp. 2938-2948.
[10] B. Li, Y. Wang, A. Singh, and Y. Vorobeychik, 'Data poisoning attacks on factorization-based collaborative filtering,' in Advances in Neural Information Processing Systems (NeurIPS), 2016, pp. 1885-1893.
[11] K. Pillutla, K. Malik, A. R. Mohamed, M. Rabbat, M. Sanjabi, and L. Xiao, 'Federated learning with partial model personalization,' in Proc. Int. Conf. Machine Learning (ICML), 2022, pp. 17787-17808.
[12] Q. Li, Y. Diao, Q. Chen, and B. He, 'Federated learning on non-IID data silos: An experimental study,' in Proc. IEEE Int. Conf. Data Engineering (ICDE), 2022, pp. 965-978.
[13] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, 'Adaptive federated learning in resource constrained edge computing systems,' IEEE J. Sel. Areas Commun., vol. 37, no. 6, pp. 1205-1221, 2019.
[14] V. Mothukuri, P. Khare, R. M. Parizi, S. Pouriyeh, A. Dehghantanha, and G. Srivastava, 'Federated-learning-based anomaly detection for IoT security attacks,' IEEE Internet Things J., vol. 9, no. 4, pp. 2545-2554, 2022.
[15] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, 'Federated learning: Challenges, methods, and future directions,' IEEE Signal Process. Mag., vol. 37, no. 3, pp. 50-60, 2020.

United International Journal of Multidisciplinary Research (UIJMR)
An International Peer-Reviewed and Refereed Multidisciplinary Journal
ISSN: 3048-6726 www.uijmr.in Impact Factor: 3.432 (SJIF); Vol-1,Issue-3(Oct-Dec),2024