

SYSTEMATIC REVIEW OF PRIVACY-PRESERVING FEDERATED LEARNING: METHODS, ARCHITECTURES, PERFORMANCE, AND CHALLENGES

Rafli Arrasyid¹, **Nandini Ananda Syifa²**

^{1,2} Department of Information System Technology and Education, Indonesia University of Education, Bandung, Indonesia.

*Correspondence author: rafly111104@gmail.com

ARTICLE INFO

Article history:

Receive: March 19, 2026

Revised: March 31, 2026

Accepted: April 7, 2026

Available Online:

April 25, 2026

Published:

April 26, 2026

Keyword:

Differential Privacy,
Federated Learning,
Homomorphic Encryption,
Privacy-Preserving,
Systematic Literature Review

ABSTRACT

Federated learning enables collaborative model training without centralizing raw data, but model updates, gradients, communication metadata, and network topology can still expose sensitive information. This article presents a systematic literature review of 281 studies published between 2021 and 2026 on privacy-preserving federated learning. The review followed the PRISMA 2020 flow diagram for systematic reviews and applied title and abstract screening, full-text eligibility assessment, quality assessment, data extraction, and thematic synthesis. Extracted data covered privacy methods, federated learning architectures, performance outcomes, privacy analysis, limitations, open challenges, application contexts, and study design. The findings show that differential privacy is the most frequently reported approach, followed by homomorphic encryption, secure aggregation, secure multi-party computation, blockchain, trusted execution environments, and communication-efficient privacy techniques. The synthesis indicates that stronger privacy protection often introduces trade-offs in accuracy, utility, latency, computational overhead, communication cost, scalability, and fairness. Persistent challenges include privacy-utility trade-offs, non-IID data, topology-based leakage, attack resistance, heterogeneous clients, and inconsistent evaluation standards. The article contributes an integrated map of methods, architectures, empirical outcomes, and implementation challenges to guide future privacy-preserving federated learning research.

© 2026. This is an open access article under the [CC-BY 4.0](https://creativecommons.org/licenses/by-sa/4.0/) license.



INTRODUCTION

Federated learning (FL) has developed as a distributed machine learning approach that enables institutions, organizations, or devices to train a shared model without transferring raw data to a central server. This approach is widely discussed in healthcare, finance, the Internet of Things (IoT), edge computing, and cybersecurity because these domains manage sensitive data and require cross-organizational collaboration (Antunes et al., 2022).

Although FL reduces the need to centralize data, it does not automatically guarantee privacy. Model updates, gradients, parameter sizes, communication metadata, and network structures can still be exploited through model inversion, membership inference, reconstruction attacks, deep leakage attacks, or topology-based leakage (Baglin et al., 2024; Liu et al., 2022; Rangwala et al., 2025). The literature therefore develops privacy-preserving mechanisms such as differential privacy, homomorphic encryption, secure aggregation, secure multi-party computation, trusted execution environments, blockchain, and hybrid approaches.

The main problem in the literature is the fragmentation of evidence. Some studies emphasize model accuracy, others emphasize formal privacy guarantees, while other studies focus on communication efficiency, cryptographic overhead, robustness against Byzantine workers, or deployment readiness (Grivet Sébert et al., 2023; Xiang et al., 2023; Zheng et al., 2022). This fragmentation makes it difficult to understand how privacy methods, architectural forms, empirical outcomes, and implementation constraints are related.

Based on this problem, this article aims to synthesize the methods, architectures, performance outcomes, limitations, and open challenges of privacy-preserving federated learning. The review not only maps the methods used in the field but also examines their practical consequences for accuracy, utility, latency, computational cost, communication overhead, scalability, fairness, and system security.

Table 1. Research questions

Code	Research question
RQ1	What methods and mechanisms are used to preserve data privacy in federated learning?
RQ2	What federated learning architectures are used in privacy-preserving FL studies?
RQ3	How do privacy methods affect accuracy, utility, computational cost, communication cost, and scalability?
RQ4	What technical limitations and implementation challenges are reported most consistently?

METHOD

This study used a systematic literature review, with the main guidance drawn from SLR guidelines for software engineering and PRISMA 2020. The PRISMA type used in this article is the PRISMA 2020 flow diagram for systematic reviews, because the study applied staged identification, screening, eligibility assessment, and inclusion of studies (Kitchenham & Charters, 2007; Page et al., 2021).

The search was conducted in academic databases that index computer science and engineering literature, especially OpenAlex and Semantic Scholar. The search string combined three groups of terms: the FL concept, privacy methods, and performance evaluation. The publication period was limited to 2021-2026 to focus on recent developments following the increased adoption of FL in sensitive domains.

The core search string was: ("federated learning" OR "federated machine learning" OR "federated deep learning") AND ("privacy-preserving" OR "data privacy" OR "differential privacy" OR "homomorphic encryption" OR "secure aggregation" OR "secure multi-party computation") AND ("accuracy" OR "utility" OR "communication cost" OR "computational overhead" OR "empirical" OR "evaluation").

The inclusion criteria covered studies that discussed FL, presented privacy methods, explained an architecture or protocol, provided empirical evidence or performance evaluation, and had full text available for analysis. The exclusion criteria covered opinion

pieces without experiments, centralized machine learning without FL, articles without a privacy component, duplicates, and studies that did not provide sufficient methodological or evaluation information.

Table 2. Inclusion, exclusion, extraction, and quality strategy

Component	Details
Inclusion	Discusses FL and privacy, presents a technical method, explains an architecture or protocol, provides empirical results, and has full text available for analysis.
Exclusion	Opinion pieces without validation, centralized ML without FL, studies without a privacy component, duplicates, inadequate empirical evidence, or overly general method descriptions.
Extraction	Privacy methods, FL architecture, performance results, privacy analysis, limitations, open challenges, application context, study design, datasets, number of clients, and evaluation metrics.
Quality assessment	FL-privacy focus, method detail, empirical evidence, publication quality, limitation discussion, and originality. Each dimension was scored from 0 to 2 and used to determine synthesis eligibility.




Figure 1. Study selection flow diagram based on PRISMA 2020.

The initial search produced 1,000 records. After title and abstract screening, 699 studies were considered relevant for full-text screening, while 301 records were excluded. The full-text screening stage resulted in 281 studies that met the inclusion criteria, while 418 studies were excluded because they did not meet the relevance, methodological quality, empirical evidence, or full-text completeness requirements. The final corpus was then analyzed using thematic synthesis with non-exclusive coding categories.

The coding categories were non-exclusive. One study could be assigned to more than one privacy method, architecture, application domain, or challenge. A method was counted

only when the mechanism was used, evaluated, or presented as an explicit part of the system design, not when it was merely mentioned in the background section.

RESULTS AND DISCUSSION

Study characteristics

The final corpus consisted of 281 studies published between 2021 and 2026. The publication distribution showed a strong increase from 2023 and reached its highest point in 2025. The distribution was 17 studies in 2021, 42 in 2022, 48 in 2023, 67 in 2024, 96 in 2025, and 11 in 2026. The 2026 figure should be read as partial because the year had not been completed when the data were collected.

Application contexts were non-exclusive because one study could discuss more than one domain. Healthcare was the most dominant domain, with 208 studies (74.0%), followed by IoT, mobile, or edge systems with 147 studies (52.3%), finance with 114 studies (40.6%), cybersecurity or intrusion detection with 68 studies (24.2%), and autonomous systems with 26 studies (9.3%). This pattern is consistent with the need for data collaboration in sectors with high privacy risk (Al Tfaily et al., 2025; Antunes et al., 2022; Deshmukh et al., 2025; Shukla et al., 2025).

Privacy-preserving methods

Privacy-preserving methods in FL can be grouped into differential privacy, homomorphic encryption, secure aggregation, secure multi-party computation, blockchain, trusted execution environments, and communication-efficient privacy techniques. Table 3 shows the non-exclusive frequencies based on the 281 final studies.

Table 3. Non-exclusive frequency of privacy-preserving methods

Method	Number	Percentage	Main implication
Differential privacy	213	75.8%	Provides formal guarantees through noise and privacy budgets, but is sensitive to epsilon, clipping, number of clients, and non-IID data (Pustozero et al., 2023; Wang et al., 2022; Yin et al., 2025).
Homomorphic encryption	119	42.3%	Protects computation on encrypted data or updates, but often increases message size, latency, and computational cost (Ma et al., 2021; Park & Lim, 2022; Rahulamathavan et al., 2023).
Secure aggregation	79	28.1%	Hides individual updates during aggregation, especially when the server must not observe separate client contributions (Liu et al., 2022; Zheng et al., 2022).
Secure multi-party computation	74	26.3%	Distributes computation across several parties and is strong for collaborative scenarios, but increases protocol coordination (Singh et al., 2025; Xiang et al., 2023).
Blockchain	31	11.0%	Supports auditability and transaction verification, but may increase latency

Method	Number	Percentage	Main implication
			and integration complexity (Qammar et al., 2023).
Communication-efficient privacy	29	10.3%	Reduces bandwidth through top-k selection, quantization, sparsification, or selective LoRA, while still needing stable convergence (Miyata et al., 2025; Xiao et al., 2025).
Trusted execution environment	18	6.4%	Provides execution isolation, but depends on vendor trust assumptions and enclave capacity (Mo et al., 2021).

Differential privacy was the most dominant approach because it is relatively lightweight, easily combined with FedAvg, and supported by formal guarantees through epsilon and delta parameters. However, the review indicates that DP cannot be treated as a single universal solution. Output perturbation, DP-SGD, adaptive clipping, and dual-mode DP produce different outcomes depending on the number of clients, model size, non-IID data, and convergence phase (Pustozero et al., 2023; Wang et al., 2022; Yin et al., 2025; Zhang et al., 2025).

Homomorphic encryption and secure aggregation offer stronger protection against the disclosure of individual updates, but their consequences include larger ciphertexts, higher latency, and more complex key management. Studies on HE, FHE, and SmartNIC offloading show that hardware optimization can reduce host CPU load, but it does not always reduce total training time or communication time (Choi et al., 2024; Ma et al., 2021; Park & Lim, 2022; Rahulamathavan et al., 2023).

Hybrid approaches are often the most realistic option for production systems because a single mechanism is rarely sufficient to address every attack surface. Combining DP with HE, secure aggregation, SMPC, robust aggregation, or audit mechanisms can strengthen privacy while preserving utility, although it also increases system design complexity (Grivet Sébert et al., 2023; Singh et al., 2025; Xiang et al., 2023; Zheng et al., 2022).

Architecture and empirical performance

Centralized client-server architecture remained the most common pattern because it is easier to control and compatible with global aggregation. However, decentralized, peer-to-peer, and hierarchical FL are increasingly used when the system does not have a fully trusted single aggregator or when the number of clients is large. Table 4 summarizes the architectural patterns found in the corpus.

Table 4. Federated learning architectural patterns in the corpus

Architecture	Number	Percentage	Implication
Centralized client-server	270	96.1%	Easy to control and common in FedAvg, but the server becomes the main trust and aggregation point.
Decentralized or peer-to-peer	134	47.7%	Reduces dependence on a central server, but coordination, consensus, and topology-based leakage must be controlled (Piran et al., 2025; Rangwala et al., 2025).
Hierarchical or edge-assisted	81	28.8%	Uses regional aggregators or edge servers to reduce communication burden and support larger client populations (Singh et al., 2025).
Hybrid or split	13	4.6%	Combines architectural patterns, such as split learning, edge aggregation, or centralized and decentralized components.

Empirical performance was not uniform. DP with a moderate privacy budget often maintained accuracy, but an overly strict budget could sharply reduce utility. Healthcare studies show that DP can still produce high diagnostic accuracy when noise is calibrated, while epidemiological studies show that R2 values can remain close to the non-DP baseline under moderate epsilon values (Kerkouche et al., 2025; Shukla et al., 2025). In contrast, knowledge distillation with DP and output perturbation shows that increasing noise can reduce accuracy on more complex image datasets (Jiang et al., 2024; Pustozero et al., 2023).

In communication-constrained scenarios, top-k shuffled DP, sparse gradients, momentum gradient descent, and selective LoRA can substantially reduce communication cost. Weng et al. (2022) reported communication savings in DP-based FL with momentum, while Xiao et al. (2025) and Miyata et al. (2025) showed that parameter selection or low-rank adaptation can reduce communication overhead without necessarily causing large utility loss.

Table 5. Summary of representative empirical evidence

Context	Method	Main finding
Healthcare	DP for breast cancer diagnosis and epidemiology	Accuracy and utility can remain high under a moderate privacy budget; representative studies report 96.1% accuracy and R2 values close to non-DP baselines (Kerkouche et al., 2025; Shukla et al., 2025).
Finance	DP, HE, and SMC	An overly strict DP setting can reduce recall in smaller institutions, whereas HE and SMC tend to be more stable but more expensive (Wasif et al., 2025).
IoT and cybersecurity	Graph-based FL, secure aggregation, and DP	Communication reduction and high accuracy are possible, but device limitations, unstable connectivity, and heterogeneity remain dominant issues (Al Tfaily et al., 2025; Deshmukh et al., 2025).
Large models and communication	Top-k shuffled DP and selective LoRA	Communication reduction above 90% and up to 99% can be achieved in several configurations, although cross-dataset testing is still needed (Miyata et al., 2025; Xiao et al., 2025).
Attack resistance	DP with robust aggregation, FHE, and dual defense	Privacy protection must be combined with defenses against poisoning attacks, Byzantine workers, and deep leakage attacks (Baglin et al., 2024; Xu et al., 2024; Zhu & Ling, 2022).

Limitations and open challenges

The most consistent challenge is the privacy-utility trade-off. DP requires noise to limit individual contribution, but the same noise can slow convergence and reduce accuracy. HE and SMPC provide strong cryptographic privacy, but their computational and communication overhead can become a barrier when models are large or the number of clients increases.

Data heterogeneity further complicates the problem. In non-IID settings, smaller clients or underrepresented groups may experience greater performance degradation when privacy mechanisms are applied uniformly. Wasif et al. (2025) showed that privacy, fairness, and accuracy trade-offs must be treated as a system design problem rather than merely a parameter tuning issue. This finding is important for healthcare, finance, and IoT because each domain has a different risk tolerance.

Attacks against FL do not only target gradient content. Deep leakage attacks, poisoning attacks, Byzantine workers, and topology-based leakage show that privacy must be defended across multiple layers. Baglin et al. (2024), Rangwala et al. (2025), Xu et al. (2024), and Zhu & Ling (2022) demonstrate that secure systems need to combine content protection, robust aggregation, structural defense, and explicit threat-model auditing.

Table 6. Open challenges and research directions

Challenge	Main problem	Research direction
Privacy-utility trade-off	DP noise, obfuscation, and strong encryption can reduce accuracy or slow convergence.	Adaptive DP, personalized FL, dynamic privacy budgets, and consistent epsilon-delta reporting.
Cryptographic overhead	HE/FHE and SMPC increase latency, message size, and key management complexity.	Selective encryption, approximate HE, hardware acceleration, and measurable offloading strategies.
Non-IID data and fairness	Different data distributions make smaller or minority clients more vulnerable to performance loss.	Quality-aware aggregation, fairness-aware FL, drift-aware DP, and subgroup evaluation.
Emerging attacks	Deep leakage, poisoning, Byzantine attacks, membership inference, and topology-based leakage continue to evolve.	Layered defense, robust aggregation, structural noise, security audits, and explicit threat models.
Real deployment	Unstable connectivity, limited devices, cross-border regulation, and interoperability remain unresolved.	Lightweight FL, MLOps for FL, benchmark standards, and cross-domain compliance audits.

CONCLUSION

This review synthesized 281 studies on privacy-preserving federated learning published between 2021 and 2026. The findings show that differential privacy is the most dominant method, while homomorphic encryption, secure aggregation, secure multi-party computation, blockchain, trusted execution environments, and communication-efficient privacy techniques are used as complementary mechanisms depending on the threat model, device capacity, and domain requirements.

In terms of performance, privacy-preserving FL can maintain model utility when privacy parameters are calibrated moderately and adaptively. However, strict privacy budgets, non-IID data, large models, full encryption, and complex architectures can increase the risk of accuracy degradation, latency, communication overhead, computational overhead, and unequal performance across clients. Therefore, the selection of privacy methods cannot be separated from data characteristics, client scale, attack type, utility targets, and system resource constraints.

The main contribution of this article is an integrated mapping of privacy methods, FL architectures, performance impacts, and implementation challenges. The findings can support future research on hybrid approaches that are lighter, adaptive to heterogeneity, fairer for smaller groups or institutions, resistant to emerging attacks, and more ready for deployment in regulated production environments.

RECOMMENDATION

Future research should use cross-domain benchmarks that report accuracy, privacy budget, communication overhead, computational overhead, fairness, and robustness in a consistent way. Experimental studies also need to expand evaluation on non-IID data, large

client populations, resource-constrained devices, and cross-institutional deployment scenarios.

For practical implementation, system developers should define a threat model before selecting a privacy method. DP is suitable for lightweight formal privacy requirements, HE and SMPC are suitable for stronger cryptographic requirements, while hybrid approaches are more appropriate when a system must balance privacy, utility, fairness, and efficiency.

AI Use Statement

The authors acknowledge the use of ChatGPT 5.4 to explore preliminary research trends and support the initial identification of possible manuscript ideas. The tool was not used to make final methodological decisions, conduct study selection, perform data extraction, determine the findings, or replace the authors' scholarly judgment. All content, interpretations, references, and final claims were reviewed, verified, and approved by the authors.

Acknowledgements

The authors would like to thank the Department of Information System Technology and Education, Indonesia University of Education, for supporting the academic environment that facilitated this research.

Funding Information

This research received no external funding.

Author Contributions Statement

Author 1 contributed to conceptualization, methodology, formal analysis, investigation, resources, data curation, writing the original draft, visualization, and project administration. Author 2 contributed to conceptualization, methodology, validation, formal analysis, investigation, resources, writing review and editing, and supervision. Both authors read and approved the final manuscript.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Rafli Arrasyid	✓	✓			✓	✓	✓	✓	✓	✓	✓		✓	
Nandini Ananda Syifa	✓	✓		✓	✓	✓	✓			✓		✓		

Abbreviations:

C : Conceptualization

M : Methodology

So : Software

Va : Validation

Fo : Formal Analysis

I : Investigation

R : Resources

D : Data Curation

O : Writing – Original Draft

E : Writing – Review & Editing

Vi : Visualization

Su : Supervision

P : Project Administration

Fu : Funding Acquisition

Conflict of Interest Statement

The authors declare that there is no conflict of interest regarding the publication of this paper.

Informed Consent

Not applicable. This study is a systematic literature review and does not involve human participants.

Ethical Approval

Not applicable. This research relies entirely on previously published literature and does not involve human or animal subjects, therefore ethical approval is not required.

REFERENCES

- Al Tfaily, F., Ghalmane, Z., Brahmia, M. el A., Hazimeh, H., Jaber, A., & Zghal, M. (2025). Graph-based federated learning approach for intrusion detection in IoT networks. *Scientific Reports*, 15(1), 41264. <https://doi.org/10.1038/s41598-025-25175-1>
- Antunes, R. S., da Costa, C. A., K"uderle, A., Yari, I. A., & Eskofier, B. (2022). Federated learning for healthcare: Systematic review and architecture proposal. *ACM Transactions on Intelligent Systems and Technology*, 13(4), 1–23. <https://doi.org/10.1145/3501813>
- Baglin, I., Zhu, X., & Hadfield, S. (2024). *FEDLAD: Federated evaluation of deep leakage attacks and defenses*. <https://doi.org/10.48550/arXiv.2411.03019>
- Choi, S., Patel, D., Zad Tootaghaj, D., Cao, L., Ahmed, F., & Sharma, P. (2024). FedNIC: enhancing privacy-preserving federated learning via homomorphic encryption offload on SmartNIC. *Frontiers in Computer Science*, 6. <https://doi.org/10.3389/fcomp.2024.1465352>
- Deshmukh, A., Esteve de la Rosa, P. L., Rodriguez, R. V., & Dasari, S. (2025). Enhancing privacy in IoT-enabled digital infrastructure: Evaluating federated learning for intrusion and fraud detection. *Sensors*, 25(10), 3043. <https://doi.org/10.3390/s25103043>
- Grivet Sébert, A., Checri, M., Stan, O., Sirdey, R., & Gouy-Pailler, C. (2023). Combining homomorphic encryption and differential privacy in federated learning. *Proceedings of the 20th Annual International Conference on Privacy, Security and Trust (PST)*, 1–7. <https://doi.org/10.1109/PST58708.2023.10320195>
- Jiang, Y., Zhao, X., Li, H., & Xue, Y. (2024). A personalized federated learning method based on knowledge distillation and differential privacy. *Electronics*, 13(17), 3538. <https://doi.org/10.3390/electronics13173538>
- Kerkouche, R., Zunker, H., Fritz, M., & K"uhn, M. J. (2025). *Differentially private federated learning for localized control of infectious disease dynamics*. <https://doi.org/10.48550/arXiv.2509.14024>
- Kitchenham, B., & Charters, S. (2007). *Guidelines for performing systematic literature reviews in software engineering* (Technical Report EBSE-2007-01). Keele University and Durham University.
- Liu, Z., Guo, J., Yang, W., Fan, J., Lam, K.-Y., & Zhao, J. (2022). Privacy-preserving aggregation in federated learning: A survey. *IEEE Transactions on Big Data*, 1–20. <https://doi.org/10.1109/TBDATA.2022.3190835>
- Ma, J., Naas, S.-A., Sigg, S., & Lyu, X. (2021). Privacy-preserving federated learning based on multi-key homomorphic encryption. *International Journal of Intelligent Systems*, 37(9), 5880–5901. <https://doi.org/10.1002/int.22818>
- Miyata, T., Yang, L., Zhu, Z., Finnerty, P., & Ohta, C. (2025). Enhancing privacy and communication efficiency in federated learning through selective low-rank adaptation and differential privacy. *Applied Sciences*, 15(24), 13102. <https://doi.org/10.3390/app152413102>
- Mo, F., Haddadi, H., Katevas, K., Marin, E., Perino, D., & Kourtellis, N. (2021). PPFL: Privacy-preserving federated learning with trusted execution environments. *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 94–108. <https://doi.org/10.1145/3458864.3466628>

- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Park, J., & Lim, H.-K. (2022). Privacy-preserving federated learning using homomorphic encryption. *Applied Sciences*, 12(2), 734. <https://doi.org/10.3390/app12020734>
- Piran, F. J., Chen, Z., Zhang, Y., Qianyu, Z., Jiong, T., & Imani, F. (2025). *Privacy-preserving decentralized federated learning via explainable adaptive differential privacy*. <https://doi.org/10.48550/arXiv.2509.10691>
- Pustozero, A., Baumbach, J., & Mayer, R. (2023). Differentially private federated learning: Privacy and utility analysis of output perturbation and DP-SGD. *Proceedings of the 2023 IEEE International Conference on Big Data (BigData)*, 5549–5558. <https://doi.org/10.1109/BigData59044.2023.10386466>
- Qammar, A., Karim, A., Ning, H., & Ding, J. (2023). Securing federated learning with blockchain: A systematic literature review. *Artificial Intelligence Review*, 56(5), 3951–3985. <https://doi.org/10.1007/s10462-022-10271-9>
- Rahulamathavan, Y., Herath, C., Liu, X., Lambbotharan, S., & Maple, C. (2023). *FheFL: Fully homomorphic encryption friendly privacy-preserving federated learning with Byzantine users*. <https://doi.org/10.48550/arXiv.2306.05112>
- Rangwala, M., Sinnott, R. O., & Buyya, R. (2025). *Network structures as an attack surface: Topology-based privacy leakage in federated learning*. <https://doi.org/10.48550/arXiv.2506.19260>
- Shukla, S., Rajkumar, S., Sinha, A., Esha, M., Elango, K., & Sampath, V. (2025). Federated learning with differential privacy for breast cancer diagnosis enabling secure data sharing and model integrity. *Scientific Reports*, 15(1), 13061. <https://doi.org/10.1038/s41598-025-95858-2>
- Singh, J. P., Aqsa, A., Ghani, I., Sonani, R., & Govindarajan, V. (2025). Privacy-aware hierarchical federated learning in healthcare: Integrating differential privacy and secure multi-party computation. *Future Internet*, 17(8), 345. <https://doi.org/10.3390/fi17080345>
- Wang, N., Xiao, Y., Chen, Y., Zhang, N., Lou, W., & Hou, Y. T. (2022). Squeezing more utility via adaptive clipping on differentially private gradients in federated meta-learning. *Proceedings of the 38th Annual Computer Security Applications Conference, ACSAC '22*, 647–657. <https://doi.org/10.1145/3564625.3564652>
- Wasif, D., Chen, D., Madabushi, S., Alluru, N., Moore, T. J., & Cho, J.-H. (2025). Empirical Analysis of Privacy-Fairness-Accuracy Trade-offs in Federated Learning: A Step Towards Responsible AI. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(3), 2666–2677. <https://doi.org/10.1609/aies.v8i3.36746>
- Weng, S., Zhang, L., Feng, D., Feng, C., Wang, R., Klaine, P. V., & Imran, M. A. (2022). Privacy-preserving federated learning based on differential privacy and momentum gradient descent. *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 1–8. <https://doi.org/10.1109/IJCNN55064.2022.9889795>
- Xiang, Z., Wang, T., Lin, W., & Wang, D. (2023). Practical differentially private and Byzantine-resilient federated learning. *Proceedings of the ACM on Management of Data*, 1(2), 1–26. <https://doi.org/10.1145/3589264>

- Xiao, D., Fan, X., & Chen, L. (2025). Top-k shuffled differential privacy federated learning for heterogeneous data. *Sensors*, 25(5), 1441. <https://doi.org/10.3390/s25051441>
- Xu, R., Gao, S., Li, C., Joshi, J., & Li, J. (2024). Dual defense: Enhancing privacy and mitigating poisoning attacks in federated learning. *Advances in Neural Information Processing Systems*, 37, 70476–70498. <https://doi.org/10.52202/079017-2253>
- Yin, C., He, K., Li, Z., Kang, N., & Shi, J. (2025). Leveraging data drift to optimize differential privacy in federated learning via dual-mode switching. *Journal of King Saud University Computer and Information Sciences*, 37(8), 225. <https://doi.org/10.1007/s44443-025-00239-4>
- Zhang, H., Huang, H., & Peng, C. (2025). A novel user behavior modeling scheme for edge devices with dynamic privacy budget allocation. *Electronics*, 14(5), 954. <https://doi.org/10.3390/electronics14050954>
- Zheng, Y., Lai, S., Liu, Y., Yuan, X., Yi, X., & Wang, C. (2022). Aggregation service for federated learning: An efficient, secure, and more resilient realization. *IEEE Transactions on Dependable and Secure Computing*, 20(2), 988–1001. <https://doi.org/10.1109/TDSC.2022.3146448>
- Zhu, H., & Ling, Q. (2022). Bridging Differential Privacy and Byzantine-Robustness via Model Aggregation. In L. D. Raedt (Ed.), *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22* (pp. 2427–2433). International Joint Conferences on Artificial Intelligence Organization. <https://doi.org/10.24963/ijcai.2022/337>