Open Access

Journal of Artificial Intelligence & Robotics

Article type: Research Article

FLIP-science: A federated learning framework for privacy-preserving collaborative AI research

Faisal Majee*

Ulster University, UK.

Corresponding Author: Faisal Majee

Ulster University, UK.

Email: herts991@gmail.com

Received: May 23, 2026

Accepted: Jun 18, 2026

Published Online: Jun 25, 2026

Website: www.joaiar.org

Volume 3 [2026] Issue 1

Abstract

The growing reliance on artificial intelligence in scientific research has created a demand for large-scale datasets that institutions hold under strict privacy obligations. Centralized machine learning approaches require data aggregation into a single repository, conflicting with data governance and regulatory requirements. While federated learning supports decentralized model training, existing frameworks focus primarily on algorithm development rather than providing infrastructure for collaborative research workflows. This paper presents FLIP-Science, a federated learning framework that enables privacy-preserving collaborative AI research across multiple institutional nodes. The system introduces an adaptive trust-weighted aggregation mechanism that incorporates client reliability, dataset quality, and communication efficiency to improve model convergence under heterogeneous conditions. Experimental evaluation on standard benchmark datasets demonstrates that FLIP-Science achieves higher accuracy and faster convergence than conventional federated baselines, while maintaining scalable and responsive infrastructure performance suitable for real-world multi-institution deployments.

Keywords: Federated learning; Privacy-preserving AI; Collaborative machine learning; Distributed AI infrastructure; Secure model training; Trust-weighted aggregation.

Introduction

Artificial Intelligence (AI) and Machine Learning (ML) serve as essential technologies supporting scientific research, industrial development, and data-driven decision making across diverse disciplines. Advances in computational power, large datasets, and deep learning have enabled researchers to tackle problems ranging from healthcare diagnostics and financial forecasting to natural language processing and computer vision. Research collaborations across multiple institutions, however, require training datasets distributed across partner organizations. These datasets frequently contain sensitive, proprietary, or regulated content that institutions must protect under legal, ethical, and internal governance policies, making cross-organizational data sharing a significant challenge.

Data privacy regulations such as the General Data Protection Regulation (GDPR) impose stringent requirements on how personal and confidential information may be collected, processed, and transferred. Organizations face a fundamental tension between regulatory compliance and the need for large, diverse training sets. The traditional centralized machine learning paradigm, which requires aggregating data from multiple sources into a single repository, is increasingly incompatible with contemporary multi-institution research frameworks.

Federated learning has emerged as an effective solution by enabling decentralized model training across multiple data locations. Introduced by McMahan et al. [1], the approach allows participants to collaboratively train a shared model without exchanging local datasets. Each node trains a local model on its own data and shares only model updates with a central coordination server, which then redistributes an aggregated global model. This decentralized design enables model training over distributed data sources while preserving local data privacy. Federated learning has attracted substantial research attention for its benefits in healthcare analytics, financial systems, and Internet of Things (IoT) environments [2].

Citation: Majee F. FLIP-science: A federated learning framework for privacy-preserving collaborative AI research. J Artif Intell Robot. 2026; 3(1): 1042.

Despite these advantages, practical adoption of federated learning in collaborative research environments remains limited. Existing frameworks such as TensorFlow Federated, PySyft, and NVFlare are designed primarily for algorithm development rather than providing complete systems for collaborative AI research workflows. These platforms support decentralized model training but lack user-friendly interfaces, monitoring tools, role-based access control, and project management capabilities [3]. Scalability, heterogeneous system integration, and compatibility with existing infrastructure remain persistent obstacles [4-6]. The absence of unified platforms combining distributed training with collaborative research tools represents a major barrier to adoption outside controlled test environments [7,8].

This paper introduces FLIP-Science, a federated learning framework designed to enable privacy-preserving collaborative AI research across multiple academic institutions. The framework provides a web-based platform through which researchers can develop and evaluate machine learning models using federated learning. The main contributions of this work are: (1) a federated research framework supporting collaborative decentralized AI model development; (2) a unified system integrating diverse research workflow functions including access control and real-time monitoring; and (3) experimental results demonstrating performance improvements of the proposed adaptive aggregation strategy over standard baselines.

The remainder of this paper is structured as follows. Section II reviews related work on federated learning foundations, infrastructure, privacy-preserving techniques, and collaborative platforms. Section III describes the FLIP-Science framework and its novel adaptive trust-weighted aggregation mechanism. Section IV presents the system infrastructure. Section V describes the experimental setup. Section VI reports results and evaluation. Section VII discusses findings and limitations. Section VIII concludes the paper. Section IX outlines future work directions.

Related work

Federated learning foundations

Federated Learning (FL) is a machine learning paradigm enabling collaborative model development without centralizing data. The foundational work by McMahan et al. [1] introduced the Federated Averaging (FedAvg) algorithm, in which each client trains a local model and transmits updates to a central server for aggregation. FedAvg permits multiple local training steps before transmitting updates, reducing communication overhead. Subsequent research extended federated learning to address data heterogeneity, communication efficiency, and scalability. The FedProx algorithm [11] improves convergence when data distributions vary across clients. Recent work has explored adaptive aggregation and efficient communication strategies [4,12]. The decentralized nature of healthcare, mobile, and IoT data has established federated learning as a critical technology in these domains [2]. Most federated learning research has focused on algorithmic improvements while underserving system-level challenges, with recognized need for systems integrating federated algorithms with infrastructure supporting cooperative multi-user research workflows [4,5].

Federated learning infrastructure

Several frameworks have been developed to support federated learning deployment. Google's TensorFlow Federated (TFF) offers flexible algorithm development but lacks essential production features such as monitoring and collaborative tooling. OpenMined's PySyft enables privacy-focused federated learning through encrypted computation, but cryptographic overhead limits accessibility for non-technical users. NVIDIA's NVFlare provides strong operational capabilities and GPU support, but hardware-specific requirements reduce portability across heterogeneous environments. Recent studies highlight that existing federated learning infrastructures lack integrated support for collaborative workflows, including experiment tracking, user management, and real-time monitoring [9, 10]. Current frameworks remain incomplete as practical research platforms, requiring significant engineering investment before enabling real-world collaborative machine learning [9].

Privacy-preserving machine learning

Protecting user privacy is a central objective of federated learning. Although federated learning avoids sharing raw data, model updates shared during training can expose private information through inference attacks. Multiple techniques have been proposed: differential privacy introduces calibrated noise into model updates; secure aggregation enables encrypted model update aggregation; homomorphic encryption and secure multiparty computation allow machine learning operations on encrypted data. Recent work demonstrates that advanced privacy-preserving techniques are necessary for trustworthy federated AI systems [13,7], though they impose additional computational and communication overhead. Integrating privacy mechanisms into production federated learning systems while maintaining operational efficiency remains challenging in multi-institution environments [7].

Collaborative AI research platforms

Collaborative AI research platforms support distributed machine learning workflows across institutional networks, allowing researchers to share model architectures and computational workflows while preserving data privacy. Most currently available platforms restrict collaboration to single organizations and do not support multi-institutional workflows. Key challenges include designing access control systems, tracking distributed training activities, and maintaining transparency in joint research processes. Recent work identifies the need for integrated platforms unifying distributed training algorithms with collaborative workflow management, administrative control, and system visibility [9,4]. FLIP-Science is designed to address this gap by connecting federated learning theory with practical multi-institution research environments.

Flip-science research framework

Framework overview

The FLIP-Science framework enables secure joint development of AI models through federated learning across multiple organizations requiring data protection. Each participating institution conducts local training on its confidential data while sharing only model updates with the coordination server, fulfilling data protection and governance obligations. The workflow comprises five main stages: global

model initialization, distribution to client nodes, local training, parameter sharing, and federated aggregation. Clients update the model during each communication round, while the coordination server combines these updates to create an enhanced global model. The framework allows institutions to benefit from diverse distributed datasets while maintaining data confidentiality and minimizing raw data transfer requirements.

Federated learning process

FLIP-Science operates through synchronized interactions between client nodes and the central coordination server. Each client node functions as an independent institution, conducting local model training on its own dataset. The coordination server manages model distribution, collects updates, and performs aggregation. Training proceeds across multiple communication rounds: at the start of each round, the global model is distributed to all clients, which produce local updates that the server uses to build an improved global model. This decentralized design keeps raw data within institutional boundaries while enabling shared model development across heterogeneous learning environments.

Model aggregation strategy

The aggregation of local model updates is a fundamental component of federated learning. FLIP-Science initially adopts the Federated Averaging (FedAvg) algorithm [1], in which model updates are weighted according to local dataset size. Let K denote the number of clients, w_k the local model parameters at client k, and n_k the number of local training samples. The global model at communication round t+1 is computed as:

$$w_{+1}^{t} = \Sigma^{k}(n_k/n)\ w_k$$

where $n = \Sigma\ n_k$ is the total number of samples across all clients. Although FedAvg is effective in homogeneous settings, it assumes uniform reliability across clients, which is unrealistic in heterogeneous environments where data quality, computational capability, and communication reliability vary. This limitation motivates the enhanced aggregation strategy described in Section III-D.

Proposed enhancement: Adaptive trust-weighted aggregation

FLIP-Science introduces an adaptive trust-weighted aggregation mechanism that incorporates client reliability, dataset quality, and communication efficiency to improve aggregation robustness under diverse environmental conditions. Each client k is assigned a trust score T_k, a dataset quality factor Q_k, and a communication efficiency factor C_k. The adaptive aggregation weight for each client is computed as:

$$\alpha_k = (n_k \cdot T_k \cdot Q_k \cdot C_k) / \Sigma_l\ (n_l \cdot T_l \cdot Q_l \cdot C_l)$$

The global model is then updated as:

$$w^{t_{+1}} = \Sigma\ \alpha_k \cdot w_k$$

To further improve efficiency, a client selection mechanism excludes unreliable or delayed nodes from participating in a given round:

$$S^t = \{\ k \mid T_k \geq \tau\_T,\ C_k \geq \tau\_C\ \}$$

This approach reduces the influence of low-quality updates, improves convergence stability, and enhances communication efficiency. The proposed enhancement introduces an evaluation layer prior to aggregation, where client updates are filtered and weighted based on trust, data quality, and communication efficiency, ensuring that only reliable contributions influence the global model and improving robustness in heterogeneous settings.

Flip-science infrastructure

System architecture

The FLIP-Science platform provides a web-based, scalable system enabling federated learning operations across multiple research environments. The implementation follows a client-server architecture composed of five core components: the frontend interface, API layer, federated coordination engine, database layer, and monitoring components. This multi-layer design delivers effective distributed training while maintaining scalability, usability, and security. The frontend interface provides user access to collaborative machine learning capabilities including project creation, model submission, and training monitoring. RESTful APIs manage communication between frontend and backend components, handling authentication, model coordination, and system interactions. The federated coordination engine controls the entire training process, managing global model initialization, client node distribution, local update collection, and aggregation.

Frontend system

The frontend is built using the React framework, providing a responsive user interface organized into modular components. The system includes three principal elements: a user dashboard, a project management interface, and training monitoring panels. The dashboard displays current project status and system performance data, giving users visibility into active research projects. The project interface enables model submission, training parameter configuration, and participating node selection. The monitoring panels display key performance indicators—including accuracy and loss metrics—throughout each stage of the communication process. This design gives researchers an accessible entry point to federated learning workflows without requiring advanced knowledge of distributed systems.

Backend infrastructure

The backend is implemented using FastAPI, a high-performance Python framework supporting asynchronous request handling for concurrent workloads. This enables the system to manage multiple simultaneous client connections, essential for federated learning systems coordinating numerous remote nodes. The backend exposes RESTful APIs supporting authentication, project management, model coordination, and result retrieval. Federated learning operations are managed through global model distribution, local update collection, and multi-round aggregation. A relational database stores all system data including user credentials, project configurations, and training results, ensuring data consistency and enabling experimental reproducibility.

Security architecture

The FLIP-Science platform incorporates fundamental security mechanisms to protect user data and system operations. Role-Based Access Control (RBAC) governs user permissions, restricting system access according to each user's organizational role. JSON Web Tokens (JWT) provide secure, stateless user session authentication. All inter-component communication is encrypted using HTTPS protocols, protecting against unauthorized access and data interception during transmission.

These mechanisms collectively support secure collaboration between distributed organizations while ensuring compliance with data protection regulations.

Experimental setup

Datasets

The FLIP-Science framework was evaluated using standard benchmark datasets: MNIST, Fashion-MNIST, and CIFAR-10. These datasets are widely used in federated learning research due to the range of complexity they represent, enabling assessment of model accuracy and convergence characteristics. The MNIST dataset contains 70,000 grayscale images of handwritten digits. Fashion-MNIST provides a more complex classification task using clothing item images. CIFAR-10 consists of 60,000 colour images across ten object categories, presenting a more challenging scenario owing to higher dimensionality and visual diversity. Datasets were distributed across multiple client nodes to simulate a multi-institution research environment using two partition strategies: equal distribution and a Non-IID partition, reflecting data imbalance encountered across real institutional deployments. Each dataset was divided into five partitions.

Training environment

The experimental environment was configured to simulate a distributed federated learning setup comprising a central coordination server and five independent client nodes, each representing an independent research institution. Before transmitting updated model parameters to the coordination server, each client node performed local training using stochastic gradient descent across multiple local epochs. The coordination server performed model aggregation using both the standard FedAvg algorithm and the proposed adaptive trust-weighted aggregation method. Each experiment was repeated five times and the reported results represent averages across runs, reducing sensitivity to random initialization. The system was deployed on a server equipped with multi-core CPU and GPU support. Client nodes used containerized environments for simulation, enabling parallel execution of local training tasks. The framework was further tested under heterogeneous client conditions, including varied data distributions and simulated communication delays.

Evaluation metrics

FLIP-Science is evaluated across three categories of metrics. Model performance is assessed using standard classification measures: accuracy, precision, recall, and F1-score, characterizing predictive capability across diverse data conditions. Federated learning efficiency is measured by the number of communication rounds required for convergence, total training duration, and training stability at convergence. Comparative evaluation is performed between standard FedAvg, FedProx [11], and the proposed adaptive trust-weighted aggregation. System performance is assessed via infrastructure metrics including API response time, system throughput, and concurrent user capacity, evaluating the platform's operational scalability and responsiveness.

Results and evaluation

Model performance

Experiments were conducted under both centralized and federated learning settings across multiple benchmark datasets. Three approaches were compared: the standard FedAvg baseline, the FedProx algorithm [11], and the proposed adaptive trust-weighted aggregation. Each experiment was repeated five times; the reported results represent mean performance with standard deviation. Table I presents the comparison results.

Table 1: Model performance comparison across methods.

Method	Accuracy (%)	Precision	Recall	F1 Score
Centralized	92.4 ± 0.3	0.92	0.91	0.91
FedAvg [1]	89.1 ± 0.8	0.88	0.87	0.87
FedProx [11]	90.0 ± 0.6	0.89	0.88	0.88
FLIP-Science (Proposed)	91.2 ± 0.5	0.90	0.89	0.89

The results demonstrate a clear advantage for the proposed FLIP-Science framework. While centralized training achieves the highest accuracy, it requires unrestricted data aggregation unavailable in federated settings. The proposed method substantially reduces the performance gap while maintaining strict data decentralization. Compared with FedAvg [1], the proposed approach improves accuracy and reduces variance, indicating more stable training. The improvement over FedProx [11] reflects the additional benefit of incorporating trust scores, dataset quality, and communication efficiency directly into the aggregation process.

Convergence analysis

Convergence characteristics were examined across multiple communication rounds for each evaluated method. The FLIP-Science framework achieves faster convergence than both baselines. The adaptive aggregation mechanism reduces the influence of low-quality updates while increasing the weight assigned to reliable clients, accelerating learning. FLIP-Science reaches near-peak accuracy within fewer communication rounds, demonstrating improved communication efficiency. The convergence curve for the proposed method also exhibits lower oscillation, confirming its ability to sustain stable performance under varying data conditions. This behavior validates the effectiveness of the trust-weighted aggregation strategy in stabilizing model updates across distributed clients.

System performance

In addition to model accuracy, the operational performance of the FLIP-Science infrastructure was evaluated using system-level metrics including API response time, concurrent user capacity, and training latency. Table II summarizes the infrastructure performance results.

Table 2: Infrastructure performance metrics.

Metric	Value
API Response Time	120 ms
Maximum Concurrent Users	500
Maximum Dataset Size	10 GB
Average Training Time (5 GB)	2.5 hours

The results show that FLIP-Science maintains consistent performance during distributed training operations. The low API response time supports responsive user interaction, while the ability to handle up to 500 concurrent users demonstrates platform scalability. Training latency remains within acceptable bounds, confirming the system's capacity to manage federated learning operations across multiple distributed nodes.

Comparative analysis

A comparative evaluation assessed FLIP-Science against standard federated baselines in terms of accuracy, training time, and communication rounds. Table III presents the comparative results.

Table 3: Comparative performance of federated learning approaches.

Method	Acc. (%)	Train Time	Rounds	Monitoring/Workflow
FedAvg [1]	89.1	3.5 h	100	Limited
FedProx [11]	90.0	3.2 h	90	Limited
FLIP-Science	91.2	3.0 h	80	Integrated
Centralized	92.4	2.8 h	N/A	Not applicable

FLIP-Science outperforms standard federated baselines in both predictive performance and training efficiency. The framework achieves higher accuracy than FedAvg [1] while requiring fewer communication rounds, indicating faster and more stable convergence. The improvement over FedProx [11] demonstrates that the adaptive trust-weighted aggregation mechanism handles heterogeneous conditions more effectively. Beyond algorithmic performance, FLIP-Science delivers integrated monitoring, access control, and workflow management capabilities not available in standard federated learning frameworks. The comparative results confirm that FLIP-Science delivers advantages across both methodological and system dimensions, making it better suited for multi-institution real-world deployment.

Discussion

The experimental results confirm that FLIP-Science enables machine learning collaboration in distributed environments while maintaining data security and delivering scalable system performance. The framework's central contribution is its ability to support decentralized model training without requiring organizations to share protected data assets, directly addressing one of the most persistent obstacles in collaborative AI research. The adaptive trust-weighted aggregation mechanism not only preserves privacy but also enhances federated learning performance under heterogeneous conditions. The improvements in accuracy, variance reduction, and convergence speed over both FedAvg [1] and FedProx [11] demonstrate that weighting client contributions according to trust, dataset quality, and communication efficiency leads to better training outcomes and greater system reliability.

A key finding is that infrastructure-level optimization delivers substantial benefits for federated learning, improving both usability and performance without necessitating complex algorithmic changes. The federated learning literature has concentrated heavily on aggregation algorithms studied in isolation, yet this research demonstrates that unified systems combining coordination, monitoring, access control, and adaptive client management produce measurable performance gains. FLIP-Science thus represents both a system implementation and an infrastructure-based research platform for collaborative AI.

The framework's adaptive selection and weighting mechanisms address the heterogeneous conditions typical of real collaborative research environments, where participating organizations bring different computational resources, communication systems, and data quality standards. In this respect, FLIP-Science more accurately reflects actual research conditions than traditional federated learning systems operating under assumptions of uniform client reliability.

Several limitations should be acknowledged. Communication overhead in federated learning scales with the number of participating nodes; while the proposed method reduces unnecessary participation, communication costs remain. The current evaluation relies on benchmark datasets and simulated client heterogeneity, which supports reproducibility but may not fully represent actual institutional deployments. The trust, dataset quality, and communication factors used are intentionally straightforward; more sophisticated estimation methods could improve performance but would introduce additional complexity. Network latency and client imbalance will continue to affect synchronization and aggregation stability at large scales.

Conclusion

This paper introduces FLIP-Science, a federated learning system enabling privacy-preserving collaborative AI research in multi-institution settings. The framework addresses a core challenge in contemporary machine learning: building a shared global model from distributed, independently held datasets that cannot be centrally aggregated due to privacy policies, institutional governance, and regulatory requirements. Unlike traditional centralized approaches, FLIP-Science enables organizations to collaboratively train models without sharing any raw training data, combining privacy protection with collaborative knowledge discovery.

FLIP-Science integrates distributed learning with an adaptive trust-weighted aggregation method that accounts for client model quality, dataset characteristics, and network conditions, making the system well-suited for heterogeneous multi-institution environments. Experiments confirm that FLIP-Science achieves higher accuracy and faster convergence than standard federated baselines, with consistent and scalable platform performance. The framework's web-based infrastructure—including a React frontend and FastAPI backend—alongside real-time monitoring and role-based access control, makes federated learning more accessible and manageable for cooperative research environments. The adaptive aggregation approach demonstrates that infrastructure-aware optimization can improve federated learning efficiency without complex algorithmic modifications. FLIP-Science contributes to the broader goal of making trustworthy, scalable, and accessible federated learning a practical reality beyond simulated settings.

Future work

FLIP-Science shows considerable promise for collaborative AI research, and several directions warrant further investigation. First, integrating blockchain-based model validation could improve transparency, auditability, and trust by recording model updates and aggregation operations on an immutable ledger. Second, hierarchical federated learning architectures could reduce communication costs and improve scalability at large scale, by introducing intermediate aggregation layers that cluster clients into local communities before global aggregation. This approach is particularly relevant for geographically distributed or highly heterogeneous research networks.

A third avenue involves automated model governance, encompassing model versioning, policy-aware access control, and lifecycle auditing for regulatory compliance. Enhanced

privacy-preserving techniques, including differential privacy, secure multiparty computation, and encrypted aggregation, represent additional research priorities. Future work will also include evaluating FLIP-Science on real-world multi-institutional datasets and conducting ablation studies on the trust, quality, and communication factors in the aggregation mechanism.

Declarations

Use of artificial intelligence: In accordance with IJMAI publication policies, the authors acknowledge that OpenAI ChatGPT was used to assist in generating portions of the text in this paper.

References

1. McMahan HB, Moore E, Ramage D, Hampson S, Arcas BA. Communication-efficient learning of deep networks from decentralized data. In: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS). 2017: 1273-1282.
2. Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. npj Digit Med. 2020; 3: 1-7.
3. Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: concept and applications. ACM Trans Intell Syst Technol. 2019; 10: 1-19.
4. Liu B, Lv N, Guo Y, Li Y. Recent advances on federated learning: a systematic survey. Neurocomputing. 2024; 574: 127-145.
5. Chai D, Wang L, Yang L, et al. A survey for federated learning evaluations: goals and measures. IEEE Trans Knowl Data Eng. 2024; 36: 1-18.
6. Zhang Q, Chen M, Yang L. A survey on federated learning systems: vision, hype and reality. IEEE Trans Knowl Data Eng. 2023; 35: 1-17.
7. Bhanbhro J, et al. Issues in federated learning: empirical challenges and limitations. Sci Rep. 2024; 14: 1-14.
8. Shenoy S, et al. Federated learning for IoT: techniques, challenges, and applications. Future Internet. 2025; 17: 1-25.
9. Xu J, Glicksberg BS, Su C, et al. Federated learning for healthcare informatics: recent advances and future directions. Nat Mach Intell. 2024; 6: 12-24.
10. Li Q, Wen Z, He B, et al. A survey on federated learning: challenges, methods, and future directions. ACM Comput Surv. 2023; 56: 1-36.
11. Li T, Sahu AK, Talwalkar A, Smith V. Federated optimization in heterogeneous networks. Proc Mach Learn Syst. 2020; 2: 429-450.
12. Liu F, et al. A survey on federated learning from a multi-party computation perspective. Front Comput Sci. 2024; 18: 1-22.
13. Kairouz P, McMahan HB, Avent B, et al. Advances and open problems in federated learning. Found Trends Mach Learn. 2021; 14: 1-210.
14. Nguyen DC, Ding M, Pathirana PN, et al. Federated learning for smart healthcare: a survey. IEEE Internet Things J. 2022; 9: 370-388.
15. Bonawitz K, Ivanov V, Kreuter B, et al. Practical secure aggregation for privacy-preserving machine learning. In: Proceedings of the ACM SIGSAC Conference on Computer and Communications Security. 2017: 1175-1191.
16. Zhang C, Li S, Xia J, et al. BatchCrypt: efficient homomorphic encryption for cross-silo federated learning. In: Proceedings of the USENIX Annual Technical Conference. 2022: 493-506.
17. Hard A, Rao K, Mathews R, et al. Federated learning for mobile keyboard prediction. arXiv preprint arXiv:1811.03604. 2018.
18. Chen J, Pan X, Monga R, Bengio S, Jozefowicz R. Revisiting distributed synchronous SGD. In: Proceedings of the International Conference on Machine Learning. 2020: 1266-1275.
19. Chen C, et al. Advances in robust federated learning: a survey. IEEE Trans Big Data. 2025; 11: 1-20.
20. Dritsas E, et al. Federated learning for IoT: techniques, challenges, and applications. Future Internet. 2025; 17: 1-25.
21. Gadekallu TR, et al. Federated learning for big data: opportunities and challenges. Eng Appl Artif Intell. 2026; 128: 105-120.
22. Baduwal M, et al. Federated learning: core challenges, trends, and future directions. Computers. 2026; 15: 155-172.