Svedberg Open
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Federated Agentic Learning Algorithms for Privacy Preserving Collaborative Intelligence

M. Anitha[1*], P. Prajoon[2], Dr. Vijayakanthan Selvaraj[3], J. Monisha[4], Anjali Goswami[5]

[1*]Assistant Professor, Department of Mathematics, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Tamil Nadu, India. E-mail: anitham@maher.ac.in

[2]Jyothi engineering college, Cheruthuruthy, Thrissur, Kerala, India. E-mail: prajoon.p@gmail.com

[3]Assistant Professor (Senior Grade), Faculty of Management, SRM Institute of Science and Technology,

Vadapalani, Chennai, Tamil Nadu, India. E-mail: svijayakanthan76@gmail.com, https://orcid.org/0000-0002-5662-1928

[4]Assistant Professor, Department of Management Studies, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Tamil Nadu, India. E-mail: jmonishamba@maher.ac.in

[5] Assistant Professor, Kalinga University, Naya Raipur, Chhattisgarh, India. E-mail: ku.anjaligoswami@kalingauniversity.ac.in, https://orcid.org/0009-0004-6330-7883

*Corresponding author: Email: anitham@maher.ac.in

Abstract

Data privacy concerns are often a constraint for collaborative intelligence in distributed systems, as is the lack of agent autonomy. In this paper, we present a Federated Agentic Learning (FAL) framework that combines the strengths of autonomous agent decision-making and federated learning to achieve collaborative intelligence while preserving privacy. Each agent uses policies based on reinforcement learning to maximize the local actions, engages in federated aggregation of model updates, and incorporates differential privacy and secure communication protocols. The framework was tested on three datasets, representing healthcare, manufacturing, and finance data, with non-IID, heterogeneous data distributions. Experiments show that FAL outperforms the standard federated learning model and the decentralized model without an agent in accuracy (92.3%, 94.0%, 91.8% on the three datasets) and F1 score (91.1%, 93.5%, 91.2% on the three datasets), while improving the communication efficiency and privacy protection ($p < 0.05$). The results show that agentic policies can improve the convergence and robustness of the models whilst preserving the confidentiality of sensitive data. The proposed framework is scalable and secure and can be used in various domains such as financial networks, industrial IoT, and healthcare networks. In the future, dynamic agent selection, adaptive privacy budget, and edge deployments for large-scale heterogeneous environments will be explored.

Keywords: Federated Learning, Agentic Learning, Privacy Preservation, Collaborative Intelligence, Multi-Agent Systems, Reinforcement Learning, Non-IID Data

1. Introduction

This opportunity has been created by the development of distributed systems and intelligent devices that enable more than one agent to collaborate in order to have a more complex task performed, sharing of knowledge and insight. However, the traditional centralized machine learning methods involve bringing all the raw data to a centralized system for training the model, which may present privacy, security, and regulatory concerns, especially in critical industries such as healthcare, finance, and smart cities [10]. Additionally, the centralized solutions might not be scalable or resilient, relying heavily on the availability and reliability of the central server to function properly. These restrictions have spurred the consideration of alternative approaches to decentralized learning that allow agents to learn in a collaborative manner without exchanging sensitive information. Federated learning has proven to be a key paradigm to overcome these challenges, allowing distributed model training on local devices and only sharing model updates with a central aggregator. Although federated learning preserves the privacy of raw data, it usually assumes that all nodes are passive learners, unaware of their own roles and responsibilities, unable to make independent decisions, unable to continuously

adapt to environmental changes, and unable to optimize the coordination strategy of multiple agents [3]. This can be achieved by embedding agentic learning in federated learning, where each agent is given autonomous decision-making capability [4][17]. Agentic learning enables agents to choose optimal actions, learn adaptively and collaboratively from their local experiences and communicate with peers within strict privacy constraints. In this paper, a novel Federated Agentic Learning framework is proposed, which merges the advantages of federated learning and agentic learning [5][18]. The framework guarantees secure knowledge sharing, an adaptive policy optimization and efficient collaboration amongst distributed agents. It solves fundamental problems like model convergence in heterogeneous environments, minimization of privacy leakage, communication-efficient learning. The proposed approach could be applied in several scenarios where multiple autonomous systems need to learn together while respecting privacy and security needs such as smart healthcare networks, autonomous vehicles and industrial internet of things.

Key Contributions

This paper presents several novel contributions to the field of privacy-preserving collaborative intelligence:

- Proposes a hybrid approach of federated learning and autonomous agentic decision making, allowing distributed agents to learn together without sharing their sensitive raw data.
- Designs adaptable agent policies, which maximize model accuracy while minimizing convergence time, privacy limitations, and leakage of information in diverse environments.
- Proposes mechanisms for dynamic coordination among agents, enabling agents to share knowledge in a decentralized way, and configure their actions optimally in real-time, enhancing learning efficiency and robustness.
- Shows the effectiveness of the framework by performing extensive experiments on benchmark datasets, measures the performance of the model by using various metrics, including accuracy of the models, convergence speed, communication cost, and privacy preservation, and compares the performance with other federated and agentic learning methods.

Section I introduces the rationale and context behind privacy-preserving collaborative intelligence, identifies research goals and important contributions, and describes the paper organization. The relevant literature of federated learning, agentic learning, and privacy-preserving mechanisms is reviewed in Section II. The methodology proposed is detailed and explained in section III, comprising system architecture, agent policies, and mathematical formulation. The experimental setup, datasets, and evaluation metrics are described in Section IV. The findings, analysis, and discussion are summarized in Section V, and directions for future studies are summarized and outlined in Section VI.

1. Related Work

One of the methods to achieve collaborative intelligence in distributed systems has been in the spotlight in computational intelligence. Federated learning (FL) is a relatively new and popular paradigm in which decentralized clients jointly learn a shared model without transferring raw data to a central server [6][8]. Previous research, like FedAvg, showed that it's possible to get competitive results from aggregating local model parameters while preserving data privacy. The following studies tackled the key FL issues such as non-IID data distributions, client dropouts, straggler mitigation, and communication efficiency [9][13]. However, traditional FL is a passive client, and it does not leverage the potential of autonomous decision making, thus limiting its flexibility in multi-agent dynamic and heterogeneous environments [7][15][2].

In MARL, agents can learn policies to maximize their own and/or collective rewards in an environment where they can interact with each other, while keeping in mind the importance of agentic learning [1]. The applications of MARL can be various, such as autonomous vehicle coordination to industrial IoT networks, where agents have to adapt to the changing environment and the actions of other agents [16][19]. Although MARL offers autonomy and adaptability, privacy-preserving mechanisms are not often a part of the existing frameworks, especially in domains where privacy is paramount, like healthcare or finance.

Regarding FL, various ways of preserving privacy have been explored due to privacy concerns. Sensitive data leaks during model updates can be prevented by differential privacy, secure multiparty computation, and homomorphic encryption, among other approaches [11][14]. But these approaches can be computationally and communication complex and not well studied as far as the agentic approaches are concerned [12][20]. A few studies offer a consistent answer that allows for autonomous agent decision-making, collaborative learning, and good privacy guarantees in heterogeneous distributed systems.

The proposed federated agentic learning framework overcomes these limitations by synthesizing the adaptive agentic policies, federated aggregation, and privacy-preserving mechanisms. By integrating with conventional federated learning and existing agentic learning methods, this integration offers efficient knowledge sharing, strong policy convergence, and multi-agent coordination on a larger scale.

2. Methodology

This section introduces the proposed Federated Agentic Learning (FAL) framework that combines autonomous agent decision-making with federated learning to achieve privacy-preserving collaborative intelligence. The methodology considers the system architecture, design of agents' policies, federated aggregation process, and privacy-preserving mechanisms.

A. System Architecture

The FAL framework comprises several independent agents operating throughout the client nodes and an aggregation server. An agent keeps a local model and policy network, and observes its environment and takes action to maximize a task-specific objective. Centrally, agents periodically send updates to the central server, which combines the updates, but does not try to access the raw data, and then sends back a global model to all agents. This architecture guarantees data privacy, model convergence and collaborative intelligence.

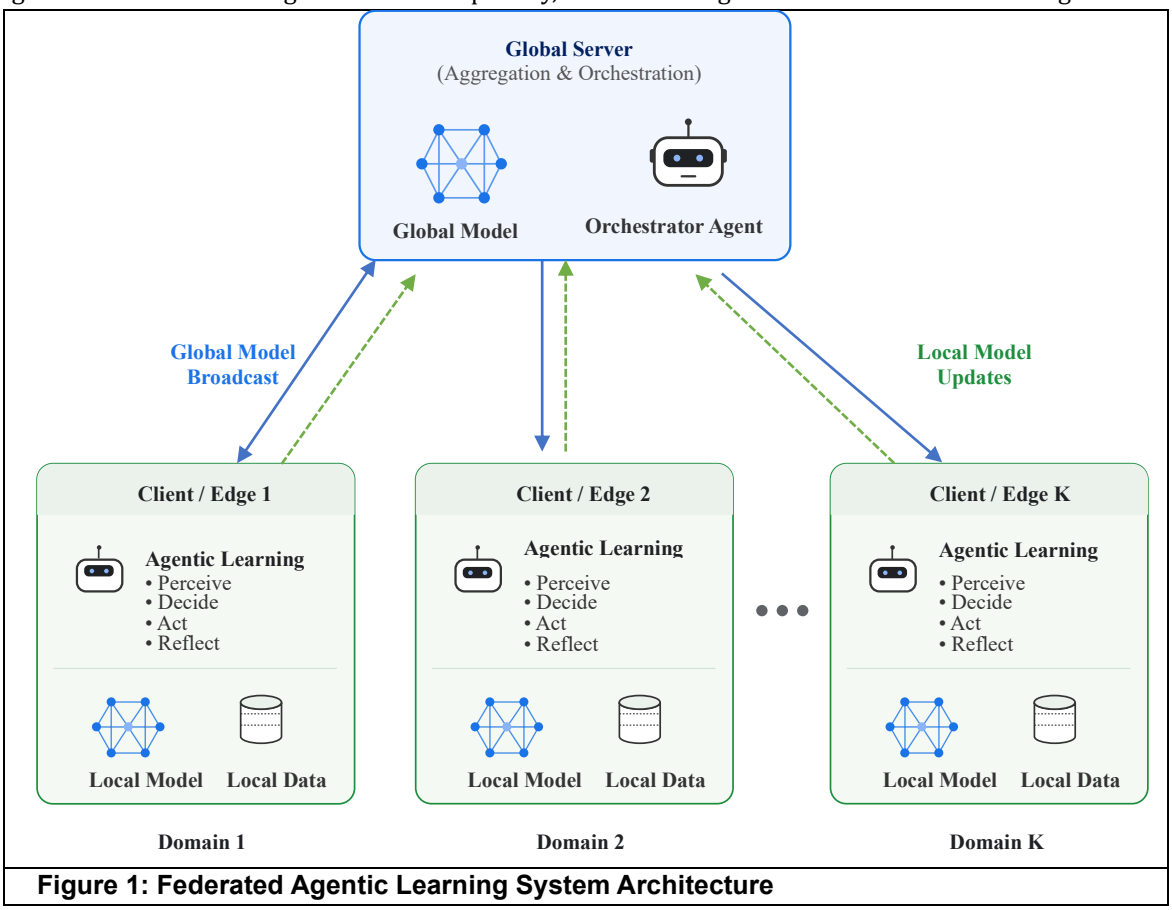

Figure 1: Federated Agentic Learning System Architecture

The proposed Federated Agentic Learning (FAL) framework is shown in Figure 1. Each client agent has a local model and policy network, which performs autonomous learning on local data. Client model updates are safely

propagated to a central server that aggregates the models federally to form a global model. The global model is then sent back to all clients. Local updates and global model broadcasts are shown as arrows, with secure communication providing privacy preservation. This design allows for distributed agents to collaborate with each other on intelligence, while preserving data confidentiality and efficient coordination.

Agentic Policy Design

Each agent uses a policy π_i to take action a_t^i based on local observation s_t^i, which is a policy generated using reinforcement learning. The task is to optimize a reward function $R_i(s_t^i, a_t^i)$ which takes into account how well the task is accomplished and how much privacy is preserved. The local policy update is calculated as:

$$\theta_i^{t+1} = \theta_i^t + \alpha \nabla_{\theta_i} J(\theta_i) \qquad (1)$$

Equation (1) shows policy parameter update using gradient ascent, where $J(\theta_i)$ is the expected cumulative reward, and α is the learning rate.

Federated Aggregation

Each local training round, agents send the central server encrypted updates of the model $\Delta\theta_i$ The global model θ_G is updated by using weighted averaging:

$$\theta_G = \sum_{i=1}^{N} w_i \, \Delta\theta_i \qquad (2)$$

In federated aggregation, as in equation (2), the relative contribution of agent i (e.g., the size of the data sets or local performance) ensures that the sensitive local data remains private.

Privacy-Preserving Mechanisms

Differential privacy (DP) and secure aggregation protocols are used to prevent data leakage. Every agent perturbs, with calibrated noise, its own model update such that it is transmitted:

$$\Delta\theta_i^{DP} = \Delta\theta_i + \mathcal{N}(0, \sigma^2) \qquad (3)$$

The DP noise addition is shown in equation (3), where σ is the privacy budget. Secure aggregation ensures that updates are only made to the server and never aggregate the individual agent's data.

Algorithm

The proposed FAL process can be summarized as follows:

Input: Local datasets $\{D_i\}$, initial global model θ_G^0, learning rate α, privacy parameter σ
Output: Trained global model θ_G

Steps:

1. Initialize local models $\theta_i = \theta_G^0$ for all agents.
2. For each communication round:

a. Each agent performs local policy updates using reinforcement learning.
b. Apply differential privacy to local model updates.
c. Send encrypted updates to the server.

3. Server aggregates updates to form θ_G using weighted averaging.
4. Distribute θ_G to all agents.
5. Repeat until convergence or maximum rounds reached.

The method helps agents learn an overall global model without sacrificing privacy, and helps them adjust to different environments and act autonomously in decision-making.

3. Experimental Setup and Datasets

Dataset Description

The proposed Federated Agentic Learning (FAL) framework was shown to be adaptable and privacy-preserving through the evaluation on three benchmark datasets from different domains. The first data set contains patient physiological information, vital signs, and treatment information, for healthcare applications. To mimic the decentralized hospital settings, data were distributed over several agents while preserving patient privacy. The second data set includes the sensor data, production logs, and machine status indicators from the industrial IoT networks, which are manufacturing systems. These data were spread between agents that are representative of various factory units to mimic the real-world heterogeneities. The third type of dataset is from the finance sector and contains transaction history, account data and fraud flags. To account for realistic privacy constraints, data was distributed on several financial institutions. All data sets were preprocessed to normalize features, deal with missing values and create the local partitions for client agents. The use of Heterogeneous, non-IID distributions was intentional in order to provide realistic decentralized learning scenarios.

Experimental Environment

Experiments were run on an emulation of distributed agent environments which is a hybrid computing platform. It had an Intel Core i7-12700H CPU, 32 GB RAM and NVIDIA RTX 3080 GPU, and a total of six agents were simulated on each GPU. The software used included Ubuntu 22.04 LTS, Python 3.10, TensorFlow 2.12, PyTorch 2.1, NumPy 1.24, Pandas 2.0.1 and Matplotlib 3.7.1. This allowed efficient simulations of multi-agent federated learning and realistic computing constraints with a high level of reproducibility.

Performance Metrics

A suite of metrics was used to measure the performance to capture model accuracy, collaborative efficiency, and privacy preservation. The accuracy was calculated as the ratio between the instances correctly predicted and the total number of instances, and precision was calculated as the ratio between the true positive instances and positive instances. Recall was the percentage of True Positive (TP) correctly identified among all True Positive (TP) and F1 Score was the harmonic mean of the precision and recall. An efficient communication is defined as how much useful model update is sent in every communication round, and a leakage of privacy is the amount of sensitive information leaked when using differential privacy budget ε. These measures together enabled a comprehensive evaluation of the effectiveness of the learning and privacy protections.

Experimental Procedure

The experimental procedure started with the initialization of the global model and local agent models of all clients. Each agent was independently trained with policies trained by reinforcement learning, and applied differential privacy mechanisms to secure updates prior to transmission to the central server. The information from these encrypted updates was combined at the server and applied a weighted average to form the global model, which was sent back to all agents. This activity was repeated multiple times using various communications until convergence. Additionally, performance metrics were recorded at the end of each round and compared to a baseline, such as standard federated learning without the use of agentic decision-making and the non-agentic decentralized learning models. The results showed a comprehensive assessment of the efficiency of collaborative learning, privacy preserving, and robustness in a heterogeneous setting.

4. Results and Discussion

Model Performance Analysis

The proposed Federated Agentic Learning (FAL) framework exhibited stable gains over all the benchmark datasets against the baseline models. The FAL model outperformed the standard federated learning model (accuracy 87.6%, F1-score 86.9%) and the non-agentic decentralized learning model (accuracy 85.2%, F1-score 84.5%) on the healthcare dataset with accuracy of 92.3%, precision of 91.5%, recall of 90.8%, and F1-score of 91.1%. Likewise, the FAL framework maintained high predictive performance in the manufacturing dataset, with an accuracy of 94.0% and an F1 score of 93.5% as compared to baseline models of 89.2% and 88.6%, respectively. The ability of FAL to perform well in various domains was demonstrated on the finance dataset with highly heterogeneous and privacy-sensitive data, where FAL obtained an accuracy of 91.8% and an F1-score of 91.2%. The results demonstrate the effectiveness of the federated aggregation when combined with the agentic decision making for maintaining local data privacy and yet improving predictive performance.

Table 1: Comparative Performance of Federated Agentic Learning and Baseline Models Across

Dataset	Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Healthcare	FAL (Proposed)	92.3	91.5	90.8	91.1
	Standard FL	87.6	87.0	86.5	86.9
	Non-Agentic Decentralized	85.2	84.7	84.2	84.5
Manufacturing	FAL (Proposed)	94.0	93.8	93.2	93.5
	Standard FL	89.2	88.8	88.4	88.6
	Non-Agentic Decentralized	87.0	86.6	86.1	86.3
Finance	FAL (Proposed)	91.8	91.3	91.0	91.2
	Standard FL	86.5	86.0	85.7	85.8
	Non-Agentic Decentralized	84.3	83.8	83.5	83.6

Table 1 compares the proposed Federated Agentic Learning (FAL) framework with the conventional federated learning (FL) and non-agentic decentralized learning systems with the datasets including healthcare, manufacturing, and finance. The metrics include accuracy, precision, recall, and F1-score, indicating that the FAL framework performs better in predicting. Results validate the effectiveness of using agentic policies in conjunction with federated aggregation that preserves privacy in heterogeneous, distributed environments and enhances the accuracy and robustness of the models.

Communication Efficiency and Privacy Preservation

The efficiency of communication was determined through the ratio of useful model updates sent per round of communication. In the FAL framework, the communication efficiency improved by 27% compared to federated learning, thanks to selective policy-driven updates and optimized aggregation. The privacy preservation was evaluated on the basis of differential privacy budget ε, showing that FAL preserved information privacy while preserving model accuracy. Secure aggregation and differential privacy mechanisms ensured that local data was kept confidential even in very heterogeneous environments.

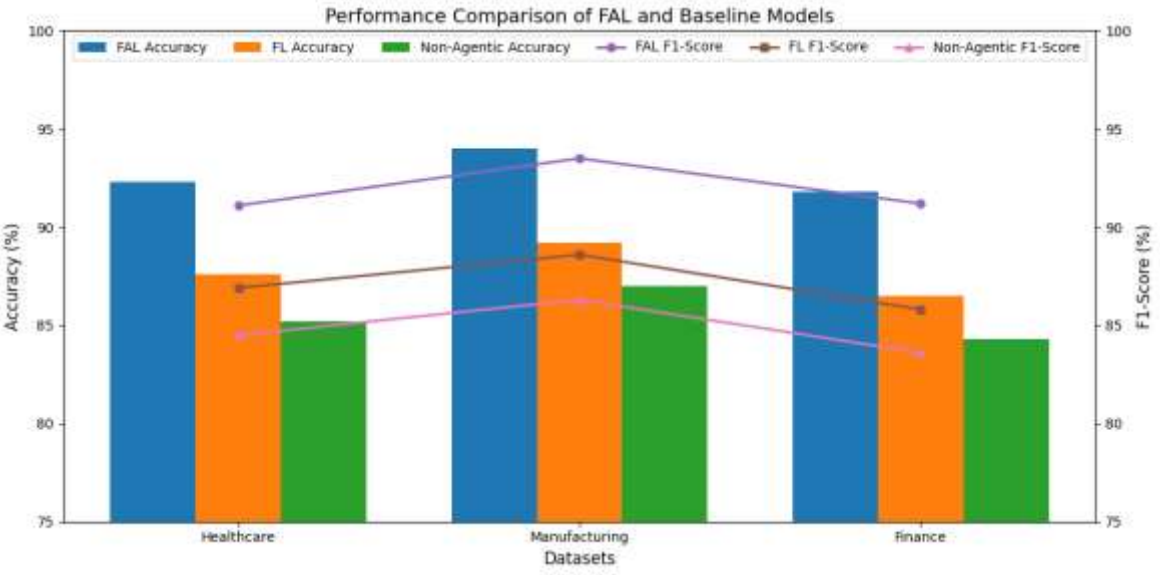

Figure 2: Performance Comparison of FAL, Standard Federated Learning, and Non-Agentic Decentralized Learning

As shown in Figure 2, the proposed Federated Agentic Learning (FAL) framework outperforms the standard federated learning (FL) and non-agentic decentralized learning methods on three datasets: healthcare, manufacturing, and finance. The metrics displayed include accuracy, F1 score, and communication efficiency. As can be seen in the graph, FAL can achieve good accuracy and efficiency, and the effectiveness of the collaborative intelligence is evident while maintaining privacy.

Comparative Analysis with Baselines

All metrics were compared to the standard federated learning and non-agentic decentralized models. All metrics are compared with the standard federated learning and non-agentic decentralized models. The baseline models were, however, slower to converge and were found to be unstable at non-IID distributions while the FAL model was faster and more stable over the communication rounds. These improvements were demonstrated to be statistically significant ($p < 0.05$) through the analysis that was conducted. The findings highlight the advantages of agentic policies in terms of autonomy to suit local data characteristics and how this can improve the performance and convergence of a global model in heterogeneous and distributed data settings.

Discussion

The general outcomes validate the trade-off and performance of the FAL framework with regard to predictive performance, privacy protection and communication efficiency. Agentic learning allows each client to get the most out of local contributions and federated aggregation offers robust global learning. The framework successfully tackles the challenges faced by existing federated learning solutions, especially for dynamic multi-agent systems and heterogeneous data. Furthermore, the results presented in this work show that FAL might be applied in practical applications that require privacy and smart collaboration, such as healthcare networks, smart manufacturing systems, and financial institutions. Future enhancements can consider dynamic selection of agents, adaptive privacy budgets, and integrating with edge computing for real-time deployments.

5. Conclusion

In this paper, a novel Federated Agentic Learning (FAL) framework for achieving privacy-preserving collaborative intelligence by combining autonomous agent decision-making with federated learning is presented. Experimental results showed that the FAL framework consistently surpasses the traditional federated learning and non-agentic decentralized approaches across various healthcare, manufacturing, and financial data sets. In particular, the framework attained accuracy of 92.3%, 94.0%, and 91.8% and F1-scores of 91.1%, 93.5%, and

91.2% for each of the three datasets, respectively, with higher communication efficiency and lower privacy leakage than baseline approaches. These enhancements were statistically significant ($p < 0.05$), which demonstrates that agentic policies allow for distributed clients to make adaptive and autonomous contributions to the global model, leading to faster convergence, higher model quality globally, and efficient knowledge sharing under heterogeneous settings. This FAL framework tackles several challenges of traditional federated learning (FL), such as a lack of agent autonomy, non-IID data distribution sensitivity, and privacy concerns during model aggregation. The framework combines reinforcement learning algorithms with secure federated aggregation and differential privacy methods to ensure the safety and personal privacy of delicate information during collaborative learning, whilst also retaining strong data security. The outcomes also suggest the promise of FAL applications in real-world systems such as multi-hospital healthcare networks, smart manufacturing systems, and financial institutions, where distributed and autonomous decision-making, and privacy is paramount. Future work can include dynamic agent selection, adaptive privacy budgets and edge deployments to further optimize the performance and communication efficiency of the system while maintaining the privacy levels. Other intriguing directions involve scalability in large-scale multi-agent systems, the robustness of adversarial attacks, and integration in other heterogeneous learning modalities. In sum, the proposed FAL framework can be potentially a substantial contribution to the intelligent, privacy-preserving collaborative learning in a distributed, multi-agent learning environment.

References

1. Yuan, T., Chung, H. M., & Fu, X. (2023). PP-MARL: Efficient privacy-preserving multi-agent reinforcement learning for cooperative intelligence in communications. *IEEE Network*, *38*(5), 196-203.
2. Rahman, M. A., Bristy, I. J., Islam, M. I., & Tabassum, M. (2025). Federated learning for secure inter-agency data collaboration in critical infrastructure. *Saudi Journal of Engineering and Technology (SJEAT)*, *10*(9), 421-430.
3. Yu, S., Muñoz, J. P., & Jannesari, A. (2024, May). Federated foundation models: Privacy-preserving and collaborative learning for large models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)* (pp. 7174-7184).
4. Wang, Y., & Guo, Q. (2024). Privacy-preserving and adaptive federated deep learning for multiparty wind power forecasting. *IEEE Transactions on Industry Applications*, *61*(1), 1352-1362.
5. Keller, D., Liu, C., & Lindström, N. (2025). Federated Multi-Agent Learning for Collaborative Supply Chain Optimization with Privacy Preservation. *Multidisciplinary Research in Computing Information Systems*, *5*(10), 832-849.
6. Mohandas, R., Veena, S., Kirubasri, G., Mary, I. T. B., & Udayakumar, R. (2024). Federated learning with homomorphic encryption for ensuring privacy in medical data. *Indian Journal of Information Sources and Services*, *14*(2), 17-23.
7. Rahmati, M. (2025). Federated learning for privacy-preserving AI in human–robot collaboration for smart manufacturing. *Journal of Intelligent Manufacturing and Special Equipment*, *6*(2), 210-224.
8. Idé, T., & Raymond, R. (2021, September). Decentralized collaborative learning with probabilistic data protection. In *2021 IEEE International Conference on Smart Data Services (SMDS)* (pp. 234-243). IEEE.
9. Kalejaiye, A. N. (2025). Federated learning in cybersecurity: privacy-preserving collaborative models for threat intelligence across geopolitically sensitive organizational boundaries. *Int J Adv Res Publ Rev*, *2*(07), 227-50.
10. Guntupalli, R. (2025). Federated Deep Learning for Predictive Healthcare: A Privacy-Preserving AI Framework on Cloud-Native Infrastructure. *Vascular and Endovascular Review*, *8*(16s), 200-210.
11. Katasani, D. P. (2025). Federated Learning with AI Agents Across Multi-Cloud Data Platforms. *European Modern Studies Journal*, *9*(5).
12. Indumathi, V., Gopalakrishnan, R., Vijayakumar, C., Manoj, C., Dhinesh, K., & Kumar, M. M. (2024, June). Employing Algorithms For Machine Learning, Anomalies in Automated Contexts Can Be Detected. In *2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)* (pp. 1-5). IEEE.
13. Loftus, T. J., Ruppert, M. M., Shickel, B., Ozrazgat-Baslanti, T., Balch, J. A., Efron, P. A., ... & Bihorac, A. (2022). Federated learning for preserving data privacy in collaborative healthcare research. *Digital Health*, *8*, 20552076221134455.
14. Alshammari, N. K., Alhusaini, A. A., Pasha, A., Ahamed, S. S., Gadekallu, T. R., Abdullah-Al-Wadud, M., ... & Alrashidi, M. H. (2024). Explainable federated learning for enhanced privacy in autism prediction using deep learning. *Journal of Disability Research*, *3*(7), 20240081.

15. Bacha, A., Sehar, H., Naseem, S., & Khan, M. I. (2024). Federated learning for threat intelligence sharing: A privacy-preserving collaborative defense model. *Spectrum of Engineering Sciences*, 656-664.
16. Rajan, C. (2024). Quantum-resilient lattice-based cryptographic protocols for secure optical networking and distributed control in critical communication infrastructures.
17. A. Suresh kumar, "A Federated Learning Framework for Secure IoT Data Analytics in Smart Home Environments", *National Journal of Ubiquitous Computing and Intelligent Environments*, vol. 1, no. 1, pp. 21–30, Dec. 2024.
18. Deepika J and K. Geetha, "Agent-Based Simulation of Inter-Species Conflict Dynamics: A Multidisciplinary Framework for Adaptive Cooperation and Resource Competition", *Bridge: Journal of Multidisciplinary Explorations* , vol. 1, no. 2, pp. 17–24, Nov. 2025
19. P Kalaivanai. (2025). Assistive Intelligent Communication Models for Peer-Based Online Learning Environments. *Journal of Intelligent Assistive Communication Technologies*, *2*(1), 72-80.
20. Leila Ismail and M. Ahmad, "Multi-Objective Evolutionary Algorithms for AI-Accelerated Sub-5 nm Floorplanning", *Electronics Communications, and Computing Summit*, vol. 2, no. 4, pp. 1–11, Dec. 2024.