WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

A Context-Aware LLM-Assisted Multi-Agent Framework for Intelligent Federated Learning Management and Decision Support

S. Priyanka[1], N. Yashema[2], S. Balaji[3]

[1] Department of Computer Science and Engineering Kingston Engineering College, Vellore, Tamil Nadu, India
[2]AP/CSE, [3]HOD/CSE
priyankasubash19@gmail.com | yash.engineering@kingston.ac.in | balajicse.engineering@kingston.ac.in

ABSTRACT :

Federated Learning is a distributed machine learning approach that enables multiple clients to collectively train models without sharing user's raw data, hence preserving user privacy. However, existing Federated Learning systems rely on static coordination mechanisms and lack intelligent decision-making capabilities for client selection, privacy management, and model monitoring. This project introduces a Context-Aware LLM-Assisted Multi-Agent Framework that combines Large Language Models with specialized agents to improvise Federated Learning management. The proposed framework consists of multiple autonomous agents, including Client Selection Agent, Privacy Agent, Trust agent, Performance Monitoring Agent, and LLM Decision Support Agent. These agents work together to analyze system context, suggest suitable actions, explain decision-making processes, and improve the Federated Learning workflow. The framework aims to improve scalability, transparency, and operational productivity while maintaining data privacy. Future enhancements include adaptive client selection, dynamic agent organization, multilingual interaction, and explainable Ai for Federated and Learning environments.

Keywords: Federated Learning, Multi-Agent Systems, Large Language Models (LLMs), Context-Aware Computing, Decision Support Systems, Distributed AI.

Introduction

Artificial Intelligence (AI) and Machine Learning (ML) are widely used to solve complex problems and support intelligent decision-making. Federated Learning (FL) is an emerging approach that enables multiple clients to collaboratively train machine learning models without sharing their raw data. However, managing federated learning environments involves challenges such as client selection, resource allocation, communication overhead, and model performance optimization.

Large Language Models (LLMs) provide advanced natural-language understanding and reasoning capabilities that can simplify complex system management tasks. By integrating LLMs with multi-agent systems, different specialized agents can collaborate to monitor the environment, analyze system conditions, and recommend suitable actions. A context-aware approach enables the system to generate recommendations based on the current state of the federated learning environment.

The proposed project, "A Context-Aware LLM-Assisted Multi-Agent Framework for Intelligent Federated Learning Management and Decision Support," aims to develop an intelligent framework for simplifying federated learning operations. The system uses an LLM to interpret user queries and coordinate specialized agents for monitoring, analysis, and decision support. A user-friendly dashboard presents important information, including client availability, training progress, model performance, and resource utilization.

The framework aims to improve management efficiency, support informed decision-making, and reduce the complexity of federated learning operations. It also considers privacy, security, transparency, and responsible AI practices. The proposed system can initially be evaluated using simulated federated clients and publicly available datasets, providing a practical foundation for further development and testing.

Literature Survey

McMahan et al. (2017): Introduced Federated Averaging (FedAvg), enabling multiple clients to train a shared model while keeping their raw data locally. This provides the foundation for the Federated Learning process in the proposed framework.

Konečný et al. (2016): Explored communication-efficient techniques for Federated Learning, highlighting the importance of reducing communication overhead during model updates.

Kairouz et al. (2021): Discussed privacy challenges and protection mechanisms in Federated Learning, supporting the need for a dedicated Privacy Agent.

Problem Statement

Federated Learning enables multiple clients to collaboratively train machine learning models without sharing their raw data. However, managing federated learning systems is challenging due to heterogeneous data, limited resources, communication overhead, and varying client performance. Traditional approaches often lack intelligent coordination and adaptive decision-making capabilities. There is a need for an automated framework that can understand the context of the learning environment and manage these challenges efficiently. The proposed system uses Large Language Models (LLMs) to analyze system conditions and support intelligent decision-making. A multi-agent architecture is employed to coordinate tasks such as client selection, resource allocation, and training optimization. The framework aims to improve learning efficiency, scalability, and resource utilization while preserving data privacy. It also provides context-aware recommendations to support effective federated learning management.

System Study

4A. Technical Feasibility

The proposed system is technically feasible as it can be developed using existing technologies such as Large Language Models (LLMs), multi-agent systems, web applications, databases, and Federated Learning (FL) frameworks. The frontend provides an interactive dashboard through which users can submit natural-language queries and monitor the federated learning environment. The backend incorporates an LLM orchestration layer that interprets user requests and coordinates specialized agents for intelligent analysis and decision support. A federated learning framework manages distributed clients and supports collaborative model training without sharing raw data. Initially, simulated clients and publicly available datasets can be used to evaluate the proposed architecture. This approach reduces infrastructure requirements and enables testing of the system's functionality, performance, and adaptability.

4B. Economic Feasibility

The proposed system is economically feasible because it can be developed using open-source software, publicly available datasets, and locally available computing resources. Simulated federated clients eliminate the need for multiple physical devices during initial development and testing. Open-source programming languages, machine learning libraries, web development tools, and database systems help minimize development expenses. An available LLM can be integrated based on the project's requirements and resource availability. Since the system primarily focuses on intelligent management, agent coordination, and decision support, specialized hardware may not be necessary for the initial prototype. Therefore, overall development and testing costs can be maintained at a reasonable level.

4C. Operational Feasibility

The proposed system is operationally feasible because it aims to simplify federated learning management through context-aware, natural-language interaction. Users can communicate their requirements without needing extensive knowledge of every technical parameter in the FL environment. The LLM interprets user queries and coordinates appropriate agents to analyze system conditions and generate relevant recommendations. The dashboard displays essential information, including client availability, training progress, model performance, resource utilization, detected risks, and suggested actions. This enables researchers, administrators, and FL managers to monitor system activities and make informed decisions. Context-aware recommendations allow the system to respond to changing environmental conditions instead of relying solely on fixed, predefined responses.

4D. Legal and Ethical Feasibility

The proposed system considers privacy, security, transparency, and responsible artificial intelligence principles. Federated learning enables raw data to remain at participating clients, reducing the need to transfer sensitive information to a central server. However, it does not guarantee complete privacy, as model updates and other system information may expose sensitive details. Therefore, appropriate security measures, access controls, secure logging, and minimal data collection should be implemented. The system should also ensure that user information is handled responsibly, and that LLM-generated recommendations are presented transparently. Human supervision should be maintained for important decisions, as AI-generated outputs may contain errors or biases. These measures help promote responsible, secure, and trustworthy management of the federated learning environment.

System Architecture

The proposed framework monitors the federated learning process using multiple specialized AI agents. The agents check client performance, security, resource usage, and training status. A decision agent combines these observations and selects a suitable action for each client. The LLM helps understand the collected information and provides simple explanations and recommendations. Only suitable client updates are forwarded to the

federated aggregation server. The server combines these updates to build an improved global model. The updated model is shared with the clients for the next training round.

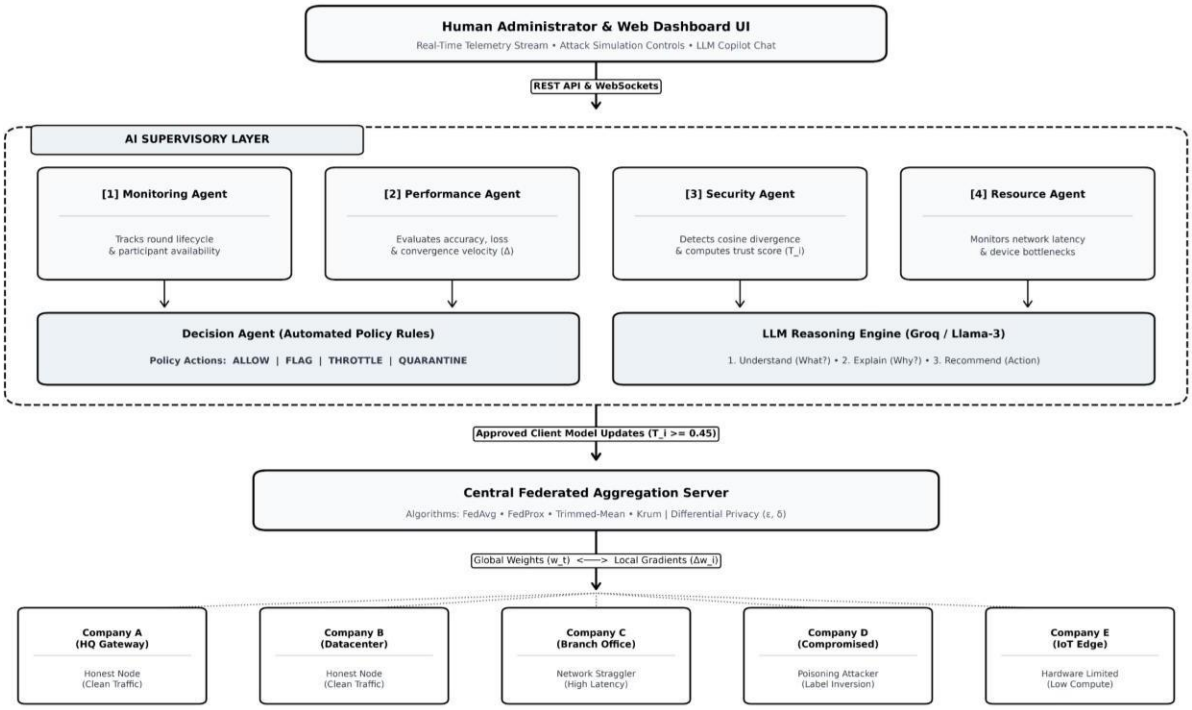


Fig. 1. Context-Aware LLM-Assisted Multi-Agent FL Supervisory Framework

Modules

6A. Client Selection Agent

6B. Privacy Agent

6C. Trust Agent

6D. Performance Monitoring Agent

6E. LLM Decision Support Agent

6A. Client Selection Agent

The Client Selection Agent is responsible for identifying the most suitable clients for each federated learning round. It evaluates various factors, including client availability, computational capability, network connectivity, resource utilization, and previous participation. Based on these factors, the agent selects clients that can contribute effectively to the global model training process. It avoids selecting clients with limited resources or unstable network connections that may cause delays or communication failures.

The agent can also consider historical performance and participation records to improve future selection decisions. Once suitable clients are identified, they receive the latest global model from the central server. Each selected client then performs local model training using own dataset without sharing raw data with other participants. This approach helps preserve data privacy while enabling collaborative learning. The Client Selection Agent aims to reduce unnecessary communication overhead, improve client participation, and support efficient utilization of available resources. Its context-aware selection strategy helps the federated learning system adapt to changing client conditions and maintain reliable training performance.

6B. Privacy Agent

The Privacy Agent is responsible for protecting sensitive information throughout the federated learning process. Although raw client data remains on local devices, model updates exchanged with the central server may still contain information that could reveal sensitive details. Therefore, the agent focuses on strengthening privacy protection during communication and model training. It can integrate privacy-preserving techniques such as differential privacy and secure aggregation to minimize the risk of information leakage. Differential privacy adds carefully controlled noise to model updates, while secure aggregation helps prevent the server from accessing individual clients' updates in their unprotected form.

The agent also monitors privacy-related requirements and checks whether appropriate protection mechanisms are applied before updates are shared. It works alongside other agents to ensure that privacy considerations are maintained during client selection, local training, and global model aggregation. By identifying potential privacy risks and recommending suitable protective measures, the Privacy Agent helps improve the security and trustworthiness of the federated learning environment. This approach supports collaborative model training while reducing the possibility of exposing confidential client information.

6C. Trust Agent

The Trust Agent is responsible for evaluating the reliability and behaviour of participating clients in the federated learning environment. It assesses client performance before their model updates are accepted for global model aggregation. The agent considers several factors, including previous participation, update consistency, contribution quality, training behaviour, and unusual activity. Based on these factors, it assigns a trust score to each participating client. Clients that demonstrate consistent and reliable behaviour may receive higher trust scores, while those exhibiting suspicious or abnormal patterns may receive lower scores.

The trust score can be used by the Client Selection Agent to identify reliable participants and by the aggregation process to determine the influence of individual model updates. The agent also helps detect potentially malicious clients that may attempt to manipulate the global model through misleading or harmful updates. By continuously monitoring client behaviour, the Trust Agent supports more reliable and secure federated learning operations. Its integration with other specialized agents helps reduce the impact of untrustworthy participants, improve model integrity, and strengthen the overall reliability of the proposed multi-agent framework.

6D. Performance Agent

The Performance Monitoring Agent is responsible for continuously observing and evaluating the overall performance of the federated learning process. It monitors important parameters such as model accuracy, training loss, convergence behaviour, communication latency, client response time, and training duration. The agent collects performance-related information from participating clients and the central server during each federated learning round. It analyzes changes in these parameters to identify performance degradation, training delays, communication bottlenecks, and unexpected behaviour. When performance issues are detected, the agent provides relevant information to the decision-support component for further analysis.

Based on the observed conditions, it can recommend suitable actions, such as adjusting training parameters, improving client selection, or optimizing resource allocation. The agent also tracks performance trends across multiple training rounds to evaluate the effectiveness of previous decisions. Its monitoring capabilities help identify potential problems at an early stage and support timely corrective actions. By coordinating with other specialized agents, the Performance Monitoring Agent contributes to maintaining stable training operations, improving resource utilization, and enhancing the overall efficiency of the proposed federated learning framework.

6E. LLM Decision Support Agent

The LLM Decision Support Agent acts as the intelligent reasoning and decision-making component of the proposed federated learning framework. It receives information from the Client Selection Agent, Privacy Agent, Trust Agent, and Performance Monitoring Agent to understand the overall state of the federated learning environment. Using Large Language Model capabilities, it analyzes the collected information, identifies potential issues, and evaluates possible actions based on the current system context. The agent provides recommendations such as accepting model updates, flagging suspicious clients, restricting participation, or excluding unreliable clients when necessary.

It can also suggest improvements to client selection, resource utilization, privacy protection, and training performance. Unlike systems that rely only on predefined rules, the LLM Decision Support Agent generates human-readable explanations describing the reasons behind its recommendations. These explanations help administrators understand system behaviour and evaluate the possible consequences of different actions. The agent can also consider feedback from previous decisions to support more consistent management. By coordinating information from multiple specialized agents, it provides context-aware recommendations that adapt to changing system conditions.

Scope of Future Enhancement

The proposed framework can be extended to support large-scale federated learning environments involving heterogeneous devices and multiple organizations. Advanced client-selection and resource-allocation algorithms can be incorporated to improve efficiency and adaptability.

Future enhancements can include stronger privacy mechanisms, improved trust management, and anomaly detection techniques to protect against malicious clients. Integrating locally deployed LLMs can further improve privacy and reduce dependency on external services.

Additional specialized agents can be introduced for fairness evaluation, energy optimization, and automated model performance monitoring. The framework can also be evaluated using standard datasets and benchmarks to measure accuracy, communication cost, training time, and decision-making efficiency.

Conclusion

The proposed Context-Aware LLM-Assisted Multi-Agent Framework for Intelligent Federated Learning Management and Decision Support combines federated learning, Large Language Models, multiagent systems, contextual reasoning, and intelligent decision support into a unified framework. Traditional federated learning focuses primarily on collaborative model training, while the proposed approach introduces an intelligent management layer that can understand user requirements, analyze the current FL environment, coordinate specialized agents, and recommendations. The use of multiple specialized agents allows different aspects of the FL environment, including monitoring, client selection, privacy, security, resource optimization, and performance analysis, to be handled independently while being coordinated through the LLM. The proposed system therefore aims to make federated learning management more adaptive, interactive, explainable, and efficient. With further development and experimental validation, the framework can be extended toward real-world privacy-preserving distributed AI applications and can provide a foundation for future research in LLM-based intelligent orchestration of federated learning systems.

References

[1] S. Saha, A. Hota, A. K. Chattopadhyay, A. Nag, and S. Nandi, "A multifaceted survey on privacy preservation of federated learning: Progress, challenges, and opportunities," Artificial Intelligence Review, vol. 57, Art. no. 184, 2024. https://doi.org/10.1007/s10462-024-10766-7

[2] A. Akhtarshenas, M. A. Vahedifar, N. Ayoobi, B. Maham, T. Alizadeh, S. Ebrahimi, and D. López-Pérez, "Federated learning: A cutting-edge survey of the latest advancements and applications," Computer Communications, vol. 228, Art. no. 107964, 2024. https://doi.org/10.1016/j.comcom.2024.107964

[3] B. Alotaibi, F. A. Khan, and S. Mahmood, "Communication efficiency and non-independent and identically distributed data challenge in federated learning: A systematic mapping study," Applied Sciences, vol. 14, no. 7, Art. no. 2720, 2024. https://doi.org/10.3390/app14072720

[4] X. Li, S. Wang, S. Zeng, Y. Wu, et al., "A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges," Vicinagearth, vol. 1, Art. no. 9, 2024. https://doi.org/10.1007/s44336-024-00009-2

[5] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, vol. 54, pp. 1273–1282, 2017. https://proceedings.mlr.press/v54/mcmahan17a.html

[6] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks," in Proceedings of Machine Learning and Systems, vol. 2, 2020. https://proceedings.mlsys.org/paper_files/paper/2020/hash/1f5fe83998a09396ebe6477d9475ba0c-Abstract.html

[7] G. Li, H. Hammoud, H. Itani, D. Khizbullin, and B. Ghanem, "CAMEL: Communicative Agents for 'Mind' Exploration of Large Language Model Society," in Advances in Neural Information Processing Systems, vol. 36, 2023. https://proceedings.neurips.cc/paper/2023/hash/a3621ee907def47c1b952ade25c67698-Abstract-Conference.html

[8] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. R. Narasimhan, and Y. Cao, "ReAct: Synergizing Reasoning and Acting in Language Models," in International Conference on Learning Representations, 2023. https://openreview.net/forum?id=WE_vluYUL-X